ETH zürich

Essays on Managing Uncertainty and Diversity

Jose Pablo Arrieta Navarro
19. Januar 2021

ETH Zürich
Department of Management, Technology, and Economics
Technology and Innovation Management chair
Weinbergstrasse 56/58
8092 Zürich

timgroup.ethz.ch

DISS. ETH NO: 27269

ESSAYS ON MANAGING UNCERTAINTY AND DIVERSITY

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

JOSE PABLO ARRIETA NAVARRO

MSc in Physics, Universidad de Costa Rica

born on 13.01.1987

citizen of Costa Rica

accepted on the recommendation of

Examiner: Prof. Dr. Stefano Brusoni

Co-examiners: Dr. Daniella Laureiro-Martínez, Prof. Dr. Chengwei Liu

2021

SUMMARY

Uncertainty is the main factor that drives organizations to constantly reinvent themselves. In an environment without uncertainty, organizations would—*eventually*—find optimal solutions to their problems and stop changing. However, in uncertain environments, an element of doubt lies behind every decision. At every step, organizations need to decide when to explore new possibilities or exploit old certainties. This dilemma is exacerbated since all the members of the organization see the world from different vantage points. These diverse points of view affect how they value options when making decisions, and how they interpret the repercussions of these decisions: The diversity of their viewpoints affects how they learn.

The concept of *diversity* is central to the literature on group decision-making, but much less so in that on organizational learning. This disparity is the focus of my dissertation, where I explore how the fact that people see the world in different ways affects how the organizations that they compose learn. I call this *viewpoint diversity*, to distinguish it from other, broader notions of diversity. In the four essays in my dissertation, I explore the *emergence*, *interdependencies*, and *strategic use* of *viewpoint diversity* through different conceptualizations, contexts, and methods.

In the first essay, co-authored with Daniella Laureiro-Martínez and Stefano Brusoni, we study how *viewpoint diversity emerges* from the cluster analysis of experienced managers' verbal protocols. We find two distinct strategies that the managers follow while solving strategic problems. In this essay, we conduct and replicate a behavioral experiment to explore further how people's attention can be manipulated towards one strategy or the other.

In the second essay, I focus my attention on the *interdependencies* resulting from having people learn in diverse ways. Even if people all want the best for their firm, but they see the world differently, and this *viewpoint diversity* will affect how their organization explores its environment. I use computational models to outline how different contingencies of *viewpoint diversity* affect the way organizations learn. *Viewpoint diversity* has a strong and nontrivial effect on exploration, it can lead to organizations that explore much more, but also much less, than organizations whose members all see the world in the same way, i.e., homogeneous organizations. The key to unlocking this apparent contradiction lies in unbundling the *interdependencies* of the organizations' *viewpoint diversity.*

In the third essay, co-authored with Chengwei Liu, we go a step further. We assume that people within organizations know that they have diverse views of the world and *use* their *viewpoint diversity strategically*. We follow a case from the venture-capital industry, where a firm followed a seemingly irrational decision rule. This rule on its own should not have been profitable—yet it actually enabled the firm to grow. We complement the case study with decision analysis and evolutionary models to show how, due to the homogeneities of the industry, this strategy was a rational way of managing that industry's uncertainty. We show that this "contrarian strategy" was valuable precisely because the decision-makers had diverse views of the world; homogeneous organizations could not have implemented it profitably.

In the fourth essay, co-authored with Roberto Fontana and Stefano Brusoni, we explore how organizations can use their product architecture to adapt to the market's uncertainty. This study follows an industry composed of 85 firms that enter the market at different times, and attain different levels of success. We build and empirically test a real-options model to determine the market-entry strategies firms should follow. We find support for our model's predictions: there is a dominant strategy for market entry. Yet, as in the third essay, it leaves blind spots that firms that *use viewpoint diversity strategically* can capture.

Overall, the four essays in the dissertation explore—at different levels of analysis and through different methods—the *emergence*, *interdependencies*, and *strategic use* of *viewpoint diversity.* For years, scholars have called for this topic to be studied, and this dissertation plays its part by highlighting the value of *viewpoint diversity* in organizational learning.

ZUSAMMENFASSUNG

Unsicherheit ist der Hauptfaktor, der Organisationen dazu treibt, sich ständig neu zu erfinden. In einem Umfeld ohne Unsicherheit würden Organisationen - *schließlich* - optimale Lösungen für ihre Probleme finden und dann aufhören sich zu verändern. In einem unsicheren Umfeld steht jedoch hinter jeder Entscheidung ein Element des Zweifels. Bei jedem Schritt müssen Organisationen entscheiden, wann sie neue Möglichkeiten erkunden oder alte Gewissheiten ausnutzen wollen. Dieses Dilemma wird noch verschärft, da alle Mitglieder der Organisation die Welt aua unterschiedlichen Perspektiven sehen. Diese diversen Perspektiven wirken sich darauf aus, wie Optionen bei der Entscheidungsfindung bewerten und wie Auswirkungen dieser Entscheidungen interpretiert werden: Die Diversität ihrer Perspektiven wirkt sich darauf aus, wie sie lernen.

Diversität ist ein zentraler Begriff in der Literatur zu Entscheidungsfindung in Gruppen, aber viel weniger beim Organisationales Lernens. Diese Disparität steht im Mittelpunkt meiner Dissertation, in der ich untersuche, wie die Tatsache, dass Menschen die Welt auf unterschiedliche Perspektiven sehen, das Lernen der Organisationen, die sie bilden, beeinflusst. Ich nenne es *Perspektivendiversität*, um es von anderen, weiter gefassten Begriffen von Diversität zu unterscheiden. In den vier Aufsätzen meiner Dissertation untersuche ich die *Entstehung*, die *Interdependenzen* und die *strategische Nutzung* der *Perspektivendiversität* durch verschiedene Konzeptualisierungen, Kontexte und Methoden.

Im ersten Aufsatz, den ich zusammen mit Daniella Laureiro-Martínez und Stefano Brusoni verfasst habe, untersuchen wir, wie sich *Perspektivendiversität* aus der Clusteranalyse der verbalen Protokolle erfahrener Manager *entsteht*. Wir fanden zwei unterschiedliche Strategien, die die Manager bei der Lösung strategischer Probleme verfolgen. In diesem Aufsatz führten wir ein Verhaltensexperiment durch und wiederholten es, um weiter zu untersuchen, wie die Aufmerksamkeit der Menschen in Richtung der einen oder anderen Strategie manipuliert werden kann.

Im zweiten Aufsatz richte ich meine Aufmerksamkeit auf die Interdependenzen, die sich daraus ergeben, dass Menschen auf unterschiedliche Weise lernen. Angenommen alle Menschen wollen das Beste für ihr Unternehmen, sehen aber die Welt auf unterschiedliche Perspektiven, wirkt sich das auf die Art und Weise aus, in der die Organisation ihr Umfeld erkundet. Ich benutze Berechnungsmodelle, um zu skizzieren, wie sich verschiedene Kontingenzen der *Perspektivendiversität* auf die Art und Weise auswirken, wie Organisationen lernen. Sie kann zu Organisationen führen, die viel mehr, aber auch viel weniger erforschen als Organisationen, deren Mitglieder alle die Welt auf die gleiche Weise sehen, d.h. homogene Organisationen. Der Schlüssel zur Auflösung dieses scheinbaren Widerspruchs liegt in der Entflechtung der *Interdependenzen* der *Perspektivendiversität* der Organisationen.

Im dritten Aufsatz, der gemeinsam mit Chengwei Liu verfasst wurde, gehen wir einen Schritt weiter. Wir gehen davon aus, dass die Menschen in Organisationen wissen, dass sie unterschiedliche Sichtweisen auf die Welt haben und dass sie die *Perspektivendiversität strategisch nützen*. Wir verfolgen einem Fall aus der Risikokapitalbranche, wo ein Unternehmen einer scheinbar irrationalen Entscheidungsregel folgte. Diese Regel allen hätte nicht gewinnbringend sein dürfen - und doch ermöglichte sie es dem Unternehmen, zu wachsen. Wir ergänzen die Fallstudie mit Entscheidungsanalysen und Evolutionsmodellen, um zu zeigen, wie diese Strategie aufgrund der Homogenität der Branche ein rationaler Weg um mit der Unsicherheit dieser Branche umzugehen. Wir zeigen, dass diese "konträre Strategie" wertvoll war, weil die Entscheidungsträger unterschiedliche Perspektiven hatten. homogene Organisationen hätten sie nicht gewinnbringend umsetzen können.

Im vierten Aufsatz, der gemeinsam mit Roberto Fontana und Stefano Brusoni verfasst wurde, untersuchen wir, wie Organisationen ihre Produktarchitektur nutzen können, um sich an die Unsicherheit des Marktes anzupassen. Diese Studie folgt einer Branche, die aus 85 Firmen besteht, die zu unterschiedlichen Zeiten auf den Markt kommen und unterschiedliche Erfolgsquoten erreichen. Wir bauen und testen empirisch ein Realoptionsmodell, um die Markteintrittsstrategien zu bestimmen, welche Unternehmen verfolgen sollten. und untersuchen, wie sie ihre Produktarchitektur nutzen können, um sich an die Unsicherheit des Marktes anzupassen. Wir finden Unterstützung für die Vorhersagen unseres Modells: Es gibt eine dominante Strategie für den Markteintritt. Doch wie im dritten Aufsatz hinterlässt sie blinde Flecken, die Unternehmen, die ihre *Perspektivendiversität strategisch nutzen.*

Insgesamt untersuchen die vier Aufsätze der Dissertation - auf verschiedenen Analyseebenen und mit unterschiedlichen Methoden - die *Entstehung*, die *Interdependenzen* und die *strategische Nutzung* der *Perspektivendiversität*. Seit Jahren fordern Wissenschaftlerinnen und Wissenschaftler die Untersuchung dieses Themas, und die vorliegende Dissertation trägt ihren Teil dazu bei, indem sie den Wert der *Perspektivendiversität* beim organisationalen Lernen hervorhebt.

CONTENTS

SUMMARY i

ZUSAMMENFASSUNG ii

INTRODUCTION 1

OVERVIEW OF PAPERS 5

SUMMARY OF PAPERS 6

THESIS CONTRIBUTIONS 13

REFERENCES 22

FULL PAPERS 27

PAPER 1 28

PAPER 2 159

PAPER 3 220

PAPER 4 263

ACKNOWLEDGEMENTS 298

INTRODUCTION

The assumption that people behave “as if” they single-mindedly maximize profits has been around for many years, and has remained a bone of contention throughout that time (Machlup, 1946; Oliver, 1947; Friedman, 1953:15). Scholars such as Jim March and Herbert Simon set forth a research program grounded in analyzing what *actually* happens in organizations (March & Simon, 1958). One key idea was to look at organizations as political systems composed of conflicting coalitions (March, 1962). This idea guided Cyert and March’s (1963:26) view that within an organization, while “people, i.e., individuals, have goals, collectivities of people do not”. In contrast to mainstream economics, this stream of work accepted that people in organizations have diverse points of view. However, current organizational learning theories scarcely consider the role of this “diversity” (Levinthal, 1997; Posen & Levinthal, 2012). In my dissertation, I aim to extend these theories by presenting how the *emergence*, *interdependencies*, and *strategic use* of *viewpoint diversity* affects organizational learning.

The business case for diversity is a strong one (Hunt, Layton, & Prince, 2015; Ely & Thomas, 2020; Pedulla, 2020). Among many possible metrics, people in organizations can differ in terms of their knowledge, demographics, and experience (DiTomaso, Post, & Parks-Yancy, 2007; Shore et al., 2009). The more we study diversity in organizations, the more paths to diversity we discover. Indeed, our studies follow the so-called *Anna Karenina* principle, whereby “all [homogenous organizations] are alike, but each [diverse organization] is [diverse] in its own way” (after Tolstoy, 1877:1). Initial studies of diversity in organizations focused on the diversity of knowledge, job types, and functional backgrounds (March, 1991; Simons, Pelled, & Smith, 1999; Kilduff, Angelmar, & Mehra, 2000). Later studies explored how demographic diversity helps build better organizations (Cox, 1994: Knight et al., 1999). Hence, a discussion emerged concerning “the ‘browning’ of the top management teams” and “shattering the glass ceiling” of organizations (Kumar & Puranam, 2012:19; Meyerson &

Fletcher, 2000:126; Bear & Woolley, 2011; Nielsen & Nielsen, 2013; Kim & Starks, 2016). To narrow down, in my dissertation, I focus on what I call *viewpoint diversity:* the fact that people see the world in different ways.

Viewpoint diversity is known to be beneficial for group decision-making and information aggregation (Stasser & Titus, 2003; Page, 2010; Woolley et al., 2010). It is particularly important for accurate decision-making under uncertainty (Canella et al., 2008; Tetlock & Gardner, 2016). In these situations, organizations aggregate their different members' points of view to find a way forward (Csaszar & Eggers, 2013). Argote (2013:134) explained how when organizations face situations "without a demonstrably correct answer, a majority… decision scheme characterizes how groups make decisions." Similarly, Kaplan (2008) showed how firms employ framing contests to choose their future strategy in uncertain environments, and build a coalition to implement the resulting vision.

The value of *viewpoint diversity* for organizational learning was first outlined by March (1962:669). March explained that the idea that organizations are composed of people who see the world in the same way is "apparently convenient for the construction of theories," but "is almost certainly wrong as a micro-description of a business." Evidence for the importance of *viewpoint diversity* in organizational learning comes from qualitative studies. For example, Ramus, Vaccaro, and Brusoni (2017) explained how hybrid logics led organizations to follow cycles of formalization and collaboration, where different points of view ultimately blend together into a cohesive framework. Similarly, Rerup and colleagues' work shows how diverse points of view help organizations triangulate their attention in reaction to crises, adapt the organization'points of view toward a common vision, balance conflicting goals while developing new products, and manage the conflicts inherent in learning from experience (Rerup, 2009; Rerup & Feldman, 2011; Salvato & Rerup, 2018; Rerup & Zbaracki, 2020). More generally, Gaba and Greve (2019) employed quantitative methods to show how firms manage safety and profit goals through political processes.

These studies provide concrete evidence of the importance of *viewpoint diversity* for organizational learning. However, the bulk of theoretical organizational learning models operationalize firms either as unitary agents, or as groups of agents who all see the world in the same way—i.e., who follow a superordinate goal (March, 1962: Levinthal, 1997; Greve, 2003; Denrell, Fang & Levinthal, 2004; Posen & Levinthal, 2012; Csaszar, 2013; Puranam & Swamy, 2016; Piezunka, Aggarwal, & Posen, 2020). In my dissertation, I build upon and extend these models to tackle the question: *How does viewpoint diversity affect how organizations learn—specifically, under uncertainty?*

Paper 1 studies the *emergence* of *viewpoint diversity*. After controlling for experience, incentives, time pressure, and demographic characteristics, we find that *viewpoint diversity emerges* from the data analysis of experienced managers' think-aloud protocols. Holding an idiosyncratic viewpoint should help an individual achieve precisely what they want. However, in groups, *viewpoint diversity* could lead individuals behave as if they follow conflicting goals, which in turn could affect how their organization learns. In Paper 2, I find this to be the case. I use agent-based models to study the *interdependencies* of *viewpoint diversity.* From the simulations, I outline two sets of contingencies that lead organizations that have *viewpoint diversity* to explore more than homogeneous organizations (i.e., the ones whose members share the same points of view) and the varied conditions that lead to decreased exploration rates.

Viewpoint diversity allows us to arm managers with a more varied set of strategies to face their environment than if we assume that all organization members behave "as if" they follow a single goal. I explore this idea in Paper 3, where I study how *viewpoint diversity* can help firms develop contrarian strategies that capture value hidden in the blind spots of the dominant firms' strategies, and equip managers with more approaches for adapting their organizations to change. Ashby (1956) explained that “only variety destroys variety.” As environments change, managers who acknowledge the strategic value of *viewpoint diversity* will have a wider variety of strategies at their disposal.

However, a common finding in organization theory is that, as time goes by, dominant firms in an industry tend to become more similar to each other (DiMaggio & Powell, 1983). While *viewpoint diversity* might enable firms to adapt more readily, one strategy will still tend to dominate. In Paper 4, we employ real-options models to estimate the optimal market-entry strategy in an uncertain market and test the model’s predictions empirically. The firms that employ the optimal strategies outperform their competition. Yet, as Paper 3 predicts, we find that a small minority of firms manage to survive and perform well by following a contrarian strategy that captures the value hidden in the blind spots of the dominant firms’ strategies.

Over 50 years have passed since March exhorted scholars to take a more diverse view of organizations (March, 1962). Advances in the field since then have created the three building blocks that I use to answer March’s call. They are: learning under uncertainty (Levinthal & March, 1993; Denrell & March, 2001; Posen & Levinthal, 2012; Piezunka, Aggarwal & Posen, 2020), information aggregation (Sah & Stiglitz, 1986; Christensen & Knudsen, 2010; Csaszar, 2012; Csaszar & Eggers, 2013), and their empirical study in controlled conditions (Laureiro-Martínez et al. 2015; Laureiro-Martinez & Brusoni, 2018; Csaszar, 2013; Csaszar & Laureiro-Martínez, 2018; Christensen & Knudsen, 2020). The four studies of my dissertation, which explore the *emergence*, *interdependencies*, and *strategic use* of *viewpoint diversity*, show that we can now extend our organizational learning models to include *viewpoint diversity* as a core construct.

OVERVIEW OF PAPERS

Paper	Title	Authors	*Viewpoint diversity*	Method	Contribution by Jose Arrieta	Status
1	Attention Processes Predict Problem-Solving Strategies: Evidence from Think-Aloud Protocols and Behavioral Experiments	Daniella Laureiro Martínez **Jose P. Arrieta** Stefano Brusoni	*Emergence*	Verbal protocol analysis + Behavioral experiment	Verbal protocol lab study: data analysis and interpretation of results Behavioral experiment: design of the study, data acquisition and analysis, and interpretation of results Drafting and revision of the paper	Preparing resubmission after 2nd R&R at *Organization Science*
2	More is Different: The Effect of Diversity of Preferences on Exploration	**Jose P. Arrieta**	*Inter-dependencies*	Agent-based model	Single author	Preparing first submission
3	In Search of Contrarian Opportunities from the Blind Spot of Majority Rule	**Jose P. Arrieta** Chengwei Liu	*Strategic use*	Case study + Evolutionary model	Conception and design of the study, design of the model, interpretation of the results, and drafting of the paper	Preparing new submission Rejected from *Academy of Management Journal*
4	On the Strategic Use of Product Modularity for Market Entry: Theory and Empirical Evidence	**Jose P. Arrieta** Roberto Fontana Stefano Brusoni	*Strategic use*	Real-options model + Quantitative	Design of the real-options model, data analysis, and drafting of the paper	Preparing new submission Rejected from *Management Science*

SUMMARY OF PAPERS

PAPER 1

In the first paper, co-authored with Daniella Laureiro-Martínez and Stefano Brusoni, we focus on finding and manipulating *viewpoint diversity.* The paper is composed of three studies. In the first study, we use verbal protocol analysis and controlled conditions to explore how a homogeneous sample of experienced managers solve strategic problems on which they do not receive performance feedback, i.e., an offline learning process (Ericsson & Simon, 1980; Fernandes & Simon, 1999). We use controlled conditions (e.g., equal incentives, time pressure, and problem description) to minimize the sources of variance (Hertwig & Ortmann, 2001; Fox, Best, & Ericsson, 2011). We then employ content-analysis methods to code the managers' thinking processes into an exhaustive set of problem-solving phases (Mintzberg, Raisinghani, & Théorêt, 1976; Neuendorf, 2002; Rangel, Camerer, & Montague, 2006). Following a quantitative approach, we go on to use sequence and cluster analysis to test for the *emergence* of different problem-solving strategies (Kaufman & Rosseeuw, 1990; Pentland, 2003; Salvato, 2009; Hennig, 2015). Out of a sample of 48 participants, two robust clusters *emerge,* i.e., two distinct problem-solving strategies. The clusters' robustness tells us that, even under controlled conditions, managers solve the problem in significantly different ways.

In the second study, we conduct a behavioral experiment to manipulate participants into using one problem-solving strategy or the other. The experiment is run online and employs tried and tested tasks for solving complex problems (Hall & Watson, 1970; Johnson & Johnson, 1982). We find evidence that we can manipulate participants' behavior. We find that participants in the treatment conditions increase their time spent on the task, and the additional cognitive effort is spent in the way requested by each treatment. In the third study, we run a preregistered replication of the second study and find broad support for the prior findings.

Overall, in Paper 1, we find that *viewpoint diversity emerges* from the cluster analysis of experienced managers' verbal protocols. We find two distinct strategies that the managers follow while solving strategic problems, and conduct and replicate a behavioral experiment to explore further how to manipulate people's attention towards either strategy.

PAPER 2

In the second paper, I study an open question from the first: namely, how does an organization composed of people who see the world in different ways gather experience and learn about its environment? This enquiry is relevant because, as March detailed in 1962, our theoretical models of organizational learning tend to assume that organizations are composed of either single agents, or agents who follow the same superordinate goal—i.e., the models lack *viewpoint diversity*.

This study tackles this question head-on and, in doing so, extends the canonical models of learning under uncertainty and decision aggregation (Posen & Levinthal, 2012; Christensen & Knudsen, 2010). I employ a multiattribute utility function from the literature on strategic positioning (Adner, Csaszar, & Zemsky, 2014). This utility function allows for one parameter, a preference, to determine how much utility an agent will perceive from an option described by two attributes. For example, if an option gives a high value of Attribute 1 and a low value of Attribute 2, an agent who prefers Attribute 1 will perceive a high utility from this option, while one who prefers Attribute 2 will not.

I build upon this utility function and borrow a standard decision model for organizational decision-making under uncertainty: the majority voting model (Christensen & Knudsen, 2010). I then formalize this model and create microstructural organizations (i.e., organizations composed of three, five, or seven agents) and explore how differences in these agents' preferences affect how the organizations explore their environment—i.e., how they learn under uncertainty (March, 1991; Puranam, 2018). The different preferences determine the measure of *viewpoint diversity*, which in this paper is termed *preference diversity.*

This paper builds on a strong foundation. Organizational learning scholars have developed an extensive understanding of how homogeneous organizations learn under uncertainty (Denrell & March, 2001; Posen & Levinthal, 2012; Puranam & Swamy, 2016; Csaszar, 2013). This foundation allows me to compare how organizations with *diverse preferences* deviate from the behavior of homogeneous organizations.

What I find is that *preference diversity* has a strong and nontrivial nuanced effect on organizational learning under uncertainty. *Preference diversity* can lead to either significantly increase or significantly decrease the amount of exploration a given organization performs. *Preference diversity* is also the key to understanding the contingencies that give rise to higher or lower exploration rates. If an organization has polarized preferences, it will achieve high exploration rates only if the preferences are unbiased. However, even small levels of bias reduce the exploration rate to levels much lower than those of homogeneous organizations. In contrast, if the organization has nonpolarized preferences, its exploration rate will be high, independent of preference bias. However, nonpolarized organizations achieve lower exploration rates than polarized unbiased ones.

The pursuit of organizational ambidexterity is a relevant research stream (O'Reilly & Tushman, 2013). Since March's original paper on the exploration-exploitation dilemma, organization scholars have searched for ways to increase exploration in organizations and avoid the "vulnerability of exploration" (March, 1991:73). Several mechanisms have been proposed for making organizations explore more and become more ambidextrous. All these mechanisms rely on separating activities by context, structures, or time (Turner, Swart, & Maylor, 2013). This paper puts forth a new way of managing an organization's exploration rate, i.e., *preference diversity*—a mechanism that does not require the separation of activities.

In Paper 2, I study the *interdependencies* that arise from having people learn in diverse ways. A firm's employees may all want what is best for the company, but if they see the world differently, organizational exploration will still be affected. I find that diverse organizations can explore much more or much less than homogenous organizations. The key to unlocking this ambiguity lies in unbundling the *interdependencies* of the organizations' *viewpoint diversity*.

PAPER 3

The third paper, co-authored with Chengwei Liu, explores an idea that emerged from the second paper: *viewpoint diversity* can be *used strategically*. The paper presents this idea by building on an empirical puzzle from the venture-capital industry (Liu et al., 2017) and explores how *viewpoint diversity* provides a solution.

Draper Fisher Jurvetson (DFJ) is a venture-capital firm founded in 1985 and currently valued at 5 billion US$. Early in its history, DFJ's founders decided to follow an unconventional rule when deciding which startups to invest in. The rule stated that DFJ would invest in a startup if just one of the three founders *really* liked it. However, they would refrain from investing if two or even all three of the founders liked it—i.e., an *antimajority voting* rule. The decision to use this rule is irrational under any performance measure.

Literature on information aggregation and social-choice theory explains that majority voting or consensus voting is the most efficient and effective way of aggregating individual decisions (Arrow, 1951; Dasgupta & Maskin, 2008; Christensen & Knudsen, 2010; Csaszar & Eggers, 2013). In contrast to theoretical expectations, DFJ grew significantly during the period that they employed *antimajority voting.* But, why did a seemingly irrational decision rule lead to consistent growth?

To solve this puzzle, we use *viewpoint diversity* to leverage the idea that startups are complex investment options that can be described with many different attributes, including intellectual property, the founding team, and the competitiveness of the market, among many others (Tata & Niedworok, 2020). However, a startup that is exceptional in one attribute but mediocre in many others might still be a highly profitable investment (Hampel, Tracey, & Weber, 2020). These startups will be left behind by firms that use majority voting or averaging decision rules (Csaszar & Eggers, 2013). The use of *antimajority voting* in organizations with *viewpoint diversity* allows for investing in startups with a few exceptional attributes even though the startup is not exceptional on average. In these firms, each agent will use different attributes to estimate the startup's value, and thus, the exceptional attributes will trigger an agent's evaluation and gain their vote. The other agents will not find the exceptional value, thus enabling the VC firm to capture a valuable option that other firms left behind by employing majority voting or averaging investment rules.

We find that *antimajority voting* can help a VC firm carve out a stable niche in a population of VC firms that all use *majority voting.* The resilience of *antimajority voting* firms is similar to the case of DFJ, which carved its position in a market composed of firms that consistently invest via majority voting or consensus. Additionally, the model shows that the firms that benefit the most from *antimajority voting* share more profound similarities with DFJ. These firms had agents who invested only if they *really* like an option, and their agents were highly dissimilar, and invested in lower-cost yet profitable startups disregarded by majority-voting firms. The fact that *antimajority voting* requires all these characteristics led to the idea that it is a complex strategy and explains why it remains uncommon, despite being profitable.

In contrast to the first two papers in the dissertation, Paper 3 finds that *viewpoint diversity* enables firms to create a defendable niche in the VC market. *Viewpoint diversity* explains why, even though *antimajority voting* can be seen as irrational, it can still be a good choice. We show that this "contrarian strategy" was successful precisely because the decision-makers had diverse views of the world. Homogeneous organizations could not implement it.

PAPER 4

The fourth paper is co-authored with Roberto Fontana and Stefano Brusoni. Paper 3 explored firms that employed their conscious understanding of *viewpoint diversity* to capture value. For this niche to be a viable option, “contrarian” firms need most of their peers to follow the industry’s dominant strategy. In Paper 4, we explore why such a strategy can emerge. We build a real-options model that outlines the appropriate entry strategies for firms in a market with different uncertainty levels. Specifically, it shows how product modularity enables firms to enter the market with a broader set of products when market uncertainty is high.

In a market defined by multiple standards battling for dominance, the best possible entry strategy for a firm is to enter the market early and with the standard that will ultimately emerge as dominant (Suarez, 2004). However, early in the market's history, the future dominant standard is uncertain (Wiegmann, de Vries & Blind, 2017). By entering early, the firm risks incurring switching costs later on—a risk that is offset by the chance of accruing early-mover advantages (Folta & O’Brien, 2004). One way for a firm to improve its chances of entering early and with the future dominant standard is by entering the market with a broader portfolio of products (Klingebiel & Rammer, 2014; Klingebiel & Joseph, 2016). This market-entry strategy has a higher cost, but leverages this cost with the lower expected switching costs and higher expected early-mover advantages. In the extreme, a firm that introduces *all* possible standards to the market early on will be certain to avoid switching costs and accrue early-mover advantages. If this firm manages to lower its development costs, this entry strategy can be a reliable way of managing the uncertainty inherent in its environment.

In this paper, we build a real-options model to outline the conditions under which an architectural innovation—product modularity—enables firms to lower the development costs of introducing multiple products to the market (Ulrich, 1994). We predict and test empirically how product modularity, by lowering development costs, can enable firms to broaden their product portfolio when market uncertainty is high, and narrow their product portfolio when such uncertainty is low.

Out of the 85 firms that form the industry we study, we find that more than half enter the market with the strategies predicted by our model (multiple modular products early on, and single products late in the market's history). The firms that replicate our model's predictions outperform their competitors. However, we also find that every possible market-entry strategy is employed. The minority of firms that employ product modularity early in the market's history, but only introduce a single standard, manage to survive longer than those that follow the model's dominant market-entry strategies. The resilience of these contrarian firms can be explained with the logic of Paper 3, where a minority of firms managed to carve out a niche in a market that followed a dominant strategy, even though their market-entry strategy was not the top-performing one.

The literature on modularity has focused on the benefits that product modularity brings to organizational problem-solving (Schilling, 2000; Brusoni, Prencipe, & Pavitt, 2001; Fang, Lee, & Schilling, 2010). This paper contributes to broadening the strategic uses of product modularity by showing how modularity can be used to diversify risk. Product modularity enables firms to introduce a broader product portfolio to the market by decreasing development costs. Thus, organizations can trade off higher development costs against the increased chance of entering the market with the right standard.

Overall, Paper 4 studies how organizations can use their product architecture to adapt to uncertainty in their market. We build and empirically test a real-options model to determine the market-entry strategies that firms should follow, and find support for our model predictions. We find that in this industry, there was an optimal way to enter the market. However, just as in Paper 3, the optimal way left behind a niche for firms that leverage their *viewpoint diversity* and act strategically.

THESIS CONTRIBUTIONS

Our society comprises people who grow up in different conditions and see the world from different vantage points. During the past few decades, the literature on decision-making has shown the great value that *viewpoint diversity* brings to group problem-solving (Hong & Page, 2004; Woolley et al., 2010; Tetlock & Gardner, 2016). Similar research has developed within the management community, and research on upper organizational echelons, particularly, has provided evidence on the importance of functional, demographic, and cultural diversity (Cox, 1994; Simons, Pelled, & Smith, 1999; Kim & Starks, 2016; Liu, 2020).

The research of my dissertation contributes to the pursuit of diversity in organizations. Specifically, the four papers study the *emergence, interdependencies,* and *strategic use* of *viewpoint diversity* in organizations. These papers contribute to our understanding of organizational learning by combining new *methods,* extending prior *theories,* and guiding *managers and practitioners* who aim to foster diversity in their organizations. In this section, I outline the contributions of my dissertation.

METHODOLOGICAL CONTRIBUTIONS

Over the past few years, methods such as natural language processing and choice analysis have been employed to quantify the differences in individuals' points of view and their effect on their organizations' decisions (Toubia et al., 2013; Hansen, McMahon, Velasco, 2014; Bailey, Strezhnev, & Voeten, 2017; Ganz, 2020). However, these methods are descriptive. They can estimate the levels of *viewpoint diversity* in organizations, but they cannot manipulate it or test its interdependencies. To study the *interdependencies* and *strategic use* of *viewpoint diversity*, we need a higher level of control. Behavioral experiments and simulations allow for this level of control, and are the main methods employed in my dissertation. I pair these methods with more classical ones, such as case studies and regression analysis, to explore the value of *viewpoint diversity* for organizational learning.

In the first paper of my dissertation, we employ a new way of uncovering *viewpoint diversity.* We couple content analysis with sequence and cluster analysis to extract two different problem-solving strategies from the think-aloud protocols of a group of experienced managers (Neuendorf, 2002; Ericsson & Simon, 1980; Pentland, 2003; Hennig, 2015). These think-aloud protocols were collected under controlled conditions (Fox, Best, & Ericsson, 2011). All managers received the same incentives, were under the same time pressure, and shared the same incentive schemes (Hertwig & Ortmann, 2001). Moreover, all the managers had similar work requirements and experience.

Despite this high level of control and sample homogeneity, two distinct problem-solving strategies emerged from our analysis. These two problem-solving strategies are uncorrelated to other demographic variables and represent a new variable to study their behavior. The emergence of this form of cognitive diversity—i.e., *viewpoint diversity*—provides an alternate route for exploring the effects of *viewpoint diversity* in organizations. Future studies could employ the two problem-solving strategies as an independent variable, and explore how groups of cognitively diverse individuals make decisions together in uncertain environments. This study would provide a real-world test of the results of the second paper of my dissertation.

The think-aloud protocols and content-analysis techniques used in Paper 1 are able to account for high levels of complexity and detail in managers' problem-solving. From each manager's transition matrices, we obtained 49 variables that described their verbal protocols in detail. This data was then fed into the cluster analysis, which was capable of reducing this complexity and pinpointing the one bit of information that distinguishes the managers' two main problem-solving strategies (Kaufman & Rosseeuw, 1990; Hennig, 2015). The use of highly granular data in the form of sequence data, coupled with unsupervised machine-learning methods, allowed for a much more nuanced evaluation of the participants' problem-solving than other content-analysis methods (e.g., word-counting, bag of words) could provide.

The combination of methods used in Paper 1 is uncommon within the organizational learning community. To help future scholars employ these methods, we uploaded a set of manuals as well as all the material we use to collect and analyze our data to an Open Science Framework repository. This paper includes two preregistered experiments: the third study of Paper 1 and a preliminary version of that study. The preregistrations are openly available at: https://osf.io/nvfdc and https://osf.io/a7sm5, respectively. The full repository is not yet public; the aim is to make it available after publication. However, the following is an anonymous view-only link to it: https://osf.io/eh5m2/?view_only=2bd6e1e7320548858fd872db4c658932

The second paper contributes to our development of theoretical organizational models. This study combines three well-known models of organizational learning: a) the canonical model under uncertainty (Posen & Levinthal, 2012); b) information aggregation (Christensen & Knudsen, 2010); and c) multi-attribute utility estimation (Adner, Csaszar, & Zemsky, 2014). The paper integrates these three models into a coherent framework. Its main methodological contribution is to show how each building block is important, but not essential for the results to be replicated.

Given that this contribution is purely methodological, it is presented only in the Appendices of Paper 2, where I replace and adapt every building block and describe how the results are replicated. The results are replicated because they depend on the interdependencies of the agents' preferences, not the specific implementation of each building block. The controlled manipulation of these interdependencies is the mechanism studied in this paper. The increase or decrease in the organization's exploration rates depends on the specific preferences of the organization. I can specify agents that learn faster, that use different utility functions, or follow different majority voting rules. But as I show in the Appendices, the exploration rate remains qualitatively similar to the ones presented in the paper as long as the organization's preferences remain unchanged.

The later papers in the dissertation are multi-method papers. Paper 3 is based on a puzzle from a case study that we solve via decision analysis and evolutionary models. Paper 4 tests the propositions of a real-options model with a quantitative exercise. Additionally, Paper 4 serves as a check of the assumptions of Paper 3. In Paper 3, we assume that organizations follow a dominant strategy, but do not explain why. Paper 4, in turn, builds a real-options model to pinpoint the dominant strategy that organizations follow, and how this strategy is dynamic and dependent on the market's uncertainty. These two papers allow for a cycle of theorizing and testing that helps explore the *strategic uses* of *viewpoint diversity* in organizations.

THEORETICAL CONTRIBUTIONS

The main contributions of this dissertation are theoretical. The papers explore a single research question: *How does viewpoint diversity affect how organizations learn—specifically, under uncertainty?* Each paper provides a different answer to this overarching question. Jointly, they help explain why, although the question has been open for over 50 years, our basic theoretical models have not accommodated *viewpoint diversity* as a foundation of how organizations learn in uncertain environments.

Viewpoint diversity is an under-explored construct in organizational learning (March, 1962; Gaba & Greve, 2019). Even though *viewpoint diversity* exists (Paper 1), and has significant and nontrivial effects (Paper 2) that allow firms to create a richer set of strategies than if everyone shared the same point of view (Papers 3 and 4), it is seldom studied. I claim that *viewpoint diversity* is an under-explored construct in organizational learning because it complicates our theorizing. The study of *viewpoint diversity* requires a combination of trusted building blocks (Paper 2 and 3), and the creation and testing of these building blocks is a time-consuming process. While scholars have undertaken this task in the past few decades (Sah & Stiglitz, 1986; Csaszar, 2012; Posen & Levinthal, 2012; Laureiro-Martinez et al., 2015), it is only in recent years that we have had robust building blocks to explore the effects of *viewpoint* *diversity* in organizational learning. My dissertation is an initial step that shows we now have the tools to respond to March's (1962) call.

The main theoretical contribution of Paper 1 is the emergence of diverse strategies for solving strategic problems. Prior studies assumed that managers solved strategic problems in a highly linear manner (Lipshitz & Bar-Ilan, 1996). We provide granular evidence to show that this is not the case; i.e., we find that managers solve problems in a highly cyclical manner. Our data also support the idea that there are significant differences in the way they solve these problems. Taking these results into account would require us, as scholars, to adapt our strategic decision-making theories to account for the concept of *viewpoint diversity*. Future studies could explore the conditions that are most appropriate to one problem-solving strategy or the other. From our study, we know that we can manipulate people into spending their time pursuing the requested problem-solving strategy. However, we were not able to improve the quality of their solutions. Therefore, future studies could explore ways of allocating people to different tasks that benefit from their preferred problem-solving strategies.

During the last decade, scholars have dramatically improved our understanding of how to design more accurate decision-making organizations, to the point of "approaching perfection" (Christensen & Knudsen, 2010:77). Paper 2 sounds a note of warning over these studies, as I find that organizations behave very differently when their agents have diverse preferences. Importantly, this form of *viewpoint diversity* affects organizational learning in a nuanced manner. A manager cannot blindly promote *viewpoint diversity;* instead, she needs to account for the interdependencies of her employees' preferences. Blind pursuit of *viewpoint diversity* could lead to a less adaptable organization.

In Paper 3, we theorize strategies for managers to use their company's *viewpoint diversity* and capture value from their market, even if it is contested. Prior literature in information aggregation had focused on the idea that organizations need to minimize errors.

For this reason, scholars have widely recommended the use of majority voting rules and averaging when making decisions in uncertain environments (Christensen & Knudsen, 2010; Csaszar & Eggers, 2013; Tata & Niedworok, 2020). However, actually *defining* an error requires a common reference frame within the organization, so that the different agents can agree on what constitutes a “hit” and a “miss.” Such a common frame of reference is standard within the literature on decision structures, but in Paper 3, we relax this assumption (Csaszar, 2013). In our paper, agents still make decisions that can be right or wrong, but since each one sees the world differently, they make different decisions and thus different errors. A manager who *uses* this diversity *strategicall*y can create organizations that capture value from the blind spot of the market’s dominant strategy. The *antimajority voting rule* that DFJ followed worked in this way. It allowed DFJ to create a market niche by investing in the valuable options left behind by the dominant firms in the industry.

This result hinges on the idea that organizations face a force that tends to make them similar to each other: a dominant strategy (DiMaggio & Powell, 1983). In Paper 4, we build a real-options model to outline the dominant strategies firms should employ when entering uncertain markets. Crucially, the model provides different predictions for different levels of market uncertainty. We corroborate these predictions and show that the firms who employ the market-entry strategies predicted by our model outperform their competition. However, as we also find, some contrarian firms remain behind—and even though they behave differently, they survive our study’s time window. These findings extend the literature on market entry by showing how dominant strategies can be dynamic and dependent on the market’s uncertainty.

With *viewpoint diversity,* we can create a wider variety of strategies for firms to adapt to their environment than when we assume that organizations are composed of people who all share the same view of the world. As Ashby (1956:206) explained, “only variety can destroy variety.” By acknowledging *viewpoint diversity,* we can create a more varied arsenal of strategies to help organizations adapt to the variety brought by changes in their environment. This wider strategic variety is central to Papers 3 and 4. In these papers, we show how *viewpoint diversity* can help organizations find the dominant niche, and also take contrarian opportunities to capture value when the dominant niche is crowded. It is too early to understand how a full inclusion of *viewpoint diversity* will adapt our theories of organizational learning. However, as diversity follows the *Anna Karenina* principle, the changes should be varied and nontrivial—two of the conditions that make *viewpoint diversity* a potentially fruitful research stream.

MANAGERIAL IMPLICATIONS

People see the world in different ways. A manager who acknowledges this fact will be able to create a wider variety of strategies than one who assumes that everyone in her firm has the same point of view on every issue. However, accepting the existence of *viewpoint diversity* is just the first step (Paper 1). A manager should understand how the different *interdependencies* of *viewpoint diversity* work together to create measures that help organizations adapt to change (Paper 2). Without this knowledge, the manager could be promoting measures that directly contradict her intentions. A manager who understands the complexity behind *viewpoint diversity* is able to understand what a market's dominant strategy can be (Paper 4), and create contrarian strategies that capture value from the blind spots of the dominant players in the market (Paper 3). The four papers in my dissertation follow this line of work, and help managers: a) acknowledge that *viewpoint diversity exists*, b) manage its *interdependencies*, and c) make *strategic use* of their organization's *viewpoint diversity*.

Viewpoint diversity emerges from studying the thinking processes of managers solving strategic problems (Paper 1). Their problem-solving processes are highly varied, but can be classified under two distinct strategies. A manager who acknowledges this diversity will have a more nuanced set of tools for allocating her employees to specific tasks. From the supplemental material of Paper 1, we know that each strategy is best apt at solving specific types of problems, one being best for ill-structured problems, and the other for ill-structured problems. However, in the different behavioral experiments we ran, we could not manipulate people into achieving these performance differences. Therefore, it is the manager's role to match her employees to the problem at hand, or even convene teams of diverse individuals when a more complex problem needs to be solved. In contrast, a manager who assumes that people solve problems homogeneously will not have access to this variety of actions and thus achieve lower performance.

The past decade has seen many calls to make organizations more diverse (Hunt, Layton, & Prince, 2015; Ely & Thomas, 2020). Organizations should foster diversity; however, diversity can lead to paradoxical effects (Kwak, 2003; Pedulla, 2020). As Paper 2 shows, *viewpoint diversity* is indeed a double-edged sword: Depending of the specific arrangement of preferences, it can lead to both increased or decreased exploration rates. In addition to asking managers to foster diversity in their organizations, we must also provide them with guidance for handling the intrinsic *interdependencies* of *viewpoint diversity*.

Paper 2 takes a first step in this direction. If a manager aims to create a more adaptable organization, fostering *viewpoint diversity* can be beneficial—but only if she implements adequate measures. The adequate measures are the two routes that lead to higher exploration. In the first route, the manager needs to promote differences in preferences, but create structures to minimize their polarization. The second route requires that the manager promotes polarization, but minimizes the bias of preferences. There are many more ways of creating organizations with *viewpoint diversity* that explore less than organizations whose agents all share the same point of view. However, by following the two routes outlined in the paper, managers can instead foster exploration and create more ambidextrous organizations.

A manager who understands the interdependencies of *viewpoint diversity* can select from a wider variety of strategic actions and adapt more effectively to the market. As Paper 3 shows, she can study her major competitors' behavior and create an appropriate set of rules to capture value from their blind spots. In doing so, the manager would create a contrarian strategy. Firms that employ these contrarian strategies can survive longer than firms that faithfully imitate the dominant players' strategy. Importantly, managers who disregard the value of *viewpoint diversity* will not discover these blind spots.

Contrarian strategies are important, but they are not the ones that capture the most value from a specific market. These dominant strategies are the focus of Paper 4, which shows how managers need to adapt and create strategies that parallel the uncertainty of their environment. In highly uncertain environments, managers should use every possible tool at their disposal to broaden their product portfolios and diversify their risk. In environments with lower uncertainty, managers should narrow down their product lines and only invest in the best options available. These high and low uncertainty cycles repeat themselves at the industry level, and thus managers should be careful to keep *viewpoint diversity* alive in their organizations. If they fail to do so, they will take longer to see change coming, and even lose key resources that could help them adapt to the next wave of change (Gavetti, 2005; Brusoni et al., 2001). Tools such as product modularity (Paper 4) can enable managers to invest in a broad portfolio of technologies, thus maintaining a diverse knowledge base in their organization.

The strategic and nuanced use of *viewpoint diversity* in organizations can help managers keep pushing measures to further inclusion and diversity at the workplace. Fostering diversity at work is an effective business strategy but, more importantly, a vital societal goal (Hunt et al., 2015; Ely & Thomas, 2020). After 50 years of calls for the study of diversity in organizations, we now have the tools to study it coherently. We can now help managers foster *viewpoint diversity* in their organizations and create a more inclusive society.

REFERENCES

Adner, R., Csaszar, F. A., & Zemsky, P. B. (2014). Positioning on a multi-attribute landscape. *Management Science*, 60(11), 2794-2815.

Argote, L. (2013). *Organizational learning*. Springer, Boston

Arrow, K. J. (1951). *Social Choice and Individual Values*, vol. 12. Yale University Press.

Ashby, W. R. (1961). *An introduction to cybernetics*. Chapman & Hall Ltd.

Bailey, M. A., Strezhnev, A., & Voeten, E. (2017). Estimating dynamic state preferences from United Nations voting data. *Journal of Conflict Resolution*, 61(2), 430-456.

Bear, J. B., & Woolley, A. W. (2011). The role of gender in team collaboration and performance. *Interdisciplinary science reviews,* 36(2), 146-153.

Brusoni, S., Prencipe, A., & Pavitt, K. (2001). Knowledge specialization, organizational coupling, and the boundaries of the firm: why do firms know more than they make?. *Administrative science quarterly*, 46(4), 597-621.

Cannella Jr, A. A., Park, J. H., & Lee, H. U. (2008). Top management team functional background diversity and firm performance: Examining the roles of team member colocation and environmental uncertainty. *Academy of management Journal*, 51(4), 768-784.

Christensen, M., & Knudsen, T. (2010). Design of decision-making org. *Management Science*, 56(1), 71-89.

Christensen, M., & Knudsen, T. (2020). Division of roles and endogenous specialization. *Industrial and Corporate Change*, 29(1), 105-124.

Cox, T. (1994). *Cultural diversity in organizations: Theory, research and practice.* Berrett-Koehler Publishers.

Csaszar, F. A. (2012). Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management Journal*, 33(6), 611-632.

Csaszar, F. A. (2013). An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organization Science*, 24(4): 1083–1101.

Csaszar, F. A., & Eggers, J. (2013). Organizational decision making: An information aggregation view. *Management Science*, 59(10): 2257–2277.

Csaszar, F. A., & Laureiro-Martínez, D. (2018). Individual and Organizational Antecedents of Strategic Foresight. *Strategy Science,* 3(3): 481–553.

Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. Englewood Cliffs, NJ, 2(4), 169-187.

Dasgupta, P., & Maskin, E. (2008). On the robustness of majority rule. *Journal of the European Economic Association,* 6(5), 949-973.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science,* 12(5), 523-538.

Denrell, J., Fang, C., & Levinthal, D. A. (2004). From T-mazes to labyrinths: Learning from model-based feedback. *Management Science*, 50(10), 1366-1378.

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American sociological review*, 147-160.

DiTomaso, N., Post, C., & Parks-Yancy, R. (2007). Workforce diversity and inequality: Power, status, and numbers. *Annual Reviews of Sociology,* 33, 473-501.

Ely, R. J., Thomas, D. A. (2020). Getting Serious About Diversity: Enough Already with the Business Case. *Harvard Business Review*, November-December Issue

Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, 87(3): 215.

Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science,* 21(3), 625-642.

Fernandes, R., & Simon, H. A. (1999). A study of how individuals solve complex and ill-structured problems. *Policy Sciences,* 32(3): 225-245.

Folta, T. B., & O'Brien, J. P. (2004). Entry in the presence of dueling options. *Strategic Management Journal,* 25(2), 121-138.

Fox, M. C., Ericsson, K. A., & Best, R. (2011). Do procedures for verbal reporting of thinking have to be reactive? A meta-analysis and recommendations for best reporting methods. *Psychological bulletin*, 137(2), 316.

Friedman, M. (1953). The methodology of positive economics. In *Essays in positive economics*, 3(3), 145-178.

Gaba, V., & Greve, H. R. (2019). Safe or profitable? The pursuit of conflicting goals. *Organization Science,* 30(4), 647-667.

Ganz, S. C. (2020). Conflict, Chaos, and the Art of Institutional Design. *Open Science Framework.* Retrieved from: https://osf.io/qjn5y/

Gavetti, G. (2005). Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. *Organization Science,* 16(6), 599-617.

Greve, H. R. (2003). *Organizational learning from performance feedback: A behavioral perspective on innovation and change.* Cambridge University Press.

Hall, J., & Watson, W. H. (1970). The effects of a normative intervention on group decision-making performance. *Human relations*, 23(4), 299-317.

Hampel, C. E., Tracey, P., & Weber, K. (2020). The art of the pivot: How new ventures manage identification relationships with stakeholders as they change direction. *Academy of Management Journal*, 63(2), 440-471.

Hansen, S., McMahon, M., & Rivera, C. V. (2014). Preferences or private assessments on a monetary policy committee?. *Journal of Monetary Economics,* 67, 16-32.

Hennig, C. (2015*). fpc: Flexible Procedures For Clustering*. R Package Version 2.1-6.

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists?. *Behavioral* and Brain Sciences, 24(3), 383-403.

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences,* 101(46), 16385-16389.

Hunt, V., Layton, D., & Prince, S. (2015). Diversity matters. *McKinsey & Company*, 1(1), 15-29.

Johnson, D. W., & Johnson, F. P. (1982). *Joining together* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall. 111-116.

Kaplan, S. (2008). Framing contests: Strategy making under uncertainty. *Organization Science*, 19(5), 729-752.

Kaufman L, Rousseeuw P. J. (0). Partitioning around medoids (program pam). In *Finding groups in data: an introduction to cluster analysis:* John Wiley & Sons. 68-125.

Kilduff, M., Angelmar, R., & Mehra, A. (2000). Top management-team diversity and firm performance: Examining the role of cognitions. *Organization science*, 11(1), 21-34.

Kim, D., & Starks, L. T. (2016). Gender diversity on corporate boards: Do women contribute unique skills?. *American Economic Review,* 106(5), 267-71.

Klingebiel, R., & Joseph, J. (2016). Entry timing and innovation strategy in feature phones. *Strategic Management Journal,* 37(6), 1002-1020.

Klingebiel, R., & Rammer, C. (2014). Resource allocation strategy for innovation portfolio management. *Strategic Management Journal,* 35(2), 246-268.

Knight, D., Pearce, C. L., Smith, K. G., Olian, J. D., Sims, H. P., Smith, K. A., & Flood, P. (1999). Top management team diversity, group process, and strategic consensus. S*trategic Management Journal*, 20(5), 445-465.
Kumar, N., & Puranam, P. (2012). *India inside: The emerging innovation challenge to the West*. Harvard Business Press.
Kwak, M. (2003). The paradoxical effects of diversity. *MIT Sloan Management Review*, 44, 7.
Laureiro-Martínez, D., & Brusoni, S. (2018). Cognitive flexibility and adaptive decision-making: Evidence from a laboratory study of expert decision makers. *Strategic Management Journal,* 39(4), 1031-1058.
Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration-exploitation dilemma: An fMRI study of attention control and decision-making performance. *Strategic Management Journal*, 36(3), 319-338.
Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14(8): 95–112.
Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science,* 43, 934-950.
Lipshitz, R., & Bar-Ilan, O. (1996). How problems are solved: Reconsidering the phase theorem. *Organizational Behavior and Human Decision Processes,* 65(1): 48-60.
Liu, C., Vlaev, I., Fang, C., Denrell, J., & Chater, N. (2017). Strategizing with biases: Engineering choice contexts for better decisions using the Mindspace approach. *California Management Review,* 59(3): 135–161.
Liu, C. (2020). Why do firms fail to engage diversity? A behavioral strategy perspective. Forthcoming at *Organization Science*.
Machlup, F. (1946). Marginal analysis and empirical research. *The American Economic Review,* 36(4), 519-554.
March, J. G., & Olsen, J. P. (2011). The logic of appropriateness. In *The Oxford handbook of political science.*
March, J. G., & Simon, H. A. (1958). *Organizations*. Wiley, New York.
March, J. G. (1962). The Business Firm as a Political Coalition. *The Journal of Politics*, 24(4): 662-678.
March, J. G., Sproull, L. S., & Tamuz, M. (1991). Learning from samples of one or fewer. *Organization science,* 2(1), 1-13.
March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1), 71-87.
Meyerson, D. E., & Fletcher, J. K. (2000). A modest manifesto for shattering the glass ceiling. *Harvard business review*, 78(1), 126-136.
Mintzberg, H., Raisinghani, D., & Théorêt, A. (1976). The structure of "unstructured" decision processes. *Administrative Science Quarterly*, 21(2): 246-275.
Neuendorf, K. A. (2002). *The content analysis guidebook.* Sage.
Nielsen, B. B., & Nielsen, S. (2013). Top management team nationality diversity and firm performance: A multilevel study. *Strategic Management Journal*, 34(3), 373-382.
Oliver, H. M. (1947). Marginal theory and business behavior. *The American Economic Review*, 37(3), 375-383.
O'Reilly III, C. A., & Tushman, M. L. (2013). Organizational ambidexterity: Past, present, and future. *Academy of management Perspectives*, 27(4), 324-338.
Page, S. E. (2010). *Diversity and complexity* (Vol. 2). Princeton University Press.
Pedulla, D. (2020, May 12), Diversity and Inclusion Efforts That Really Work. Harvard Business Review. Retrieved online from: https://hbr.org/2020/05/diversity-and-inclusion-efforts-that-really-work
Pentland, B. T. (2003). Sequential variety in work processes. *Organization Science*, 14(5), 528.

Piezunka, H., Aggarwal, V. A., & Posen, H. E. (2020). *Learning-by-Participating: The Dual Role of Structure in Aggregating Information and Shaping Learning.* SSRN, Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3425696

Posen, H. E., & Levinthal, D. A. (2012). Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science*, 58(3), 587-601.

Puranam, P., & Swamy, M. (2016). How initial representations shape coupled learning processes. *Organization Science*, 27(2), 323-335.

Puranam, P. (2018). *The microstructure of organizations*. Oxford University Press.

Ramus, T., Vaccaro, A., & Brusoni, S. (2017). Institutional complexity in turbulent times: Formalization, collaboration, and the emergence of blended logics. *Academy of Management Journal,* 60(4), 1253-1284.

Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making*. Nature Reviews Neuroscience*, 9(7): 545-556.

Rerup, C. (2009). Attentional triangulation: Learning from unexpected rare crises. *Organization Science*, 20(5), 876-893.

Rerup, C., & Feldman, M. S. (2011). Routines as a source of change in organizational schemata: The role of trial-and-error learning. *Academy of Management Journal,* 54(3), 577-610.

Rerup, C., & Zbaracki, M. J. (2020). The politics of learning from rare events. *Organization Science.* Forthcoming

Sah, R. K., & Stiglitz, J. E. (1986). The architecture of economic systems: Hierarchies and polyarchies. *The American Economic Review*, 76(4): 716–727.

Salvato, C. (2009). Capabilities unveiled: The role of ordinary activities in the evolution of product development processes. *Organization Science,* 20(2), 384-409.

Salvato, C., & Rerup, C. (2018). Routine regulation: Balancing conflicting goals in organizational routines. *Administrative Science Quarterly*, 63(1), 170-209.

Schilling, M. A. (2000). Toward a general modular systems theory and its application to interfirm product modularity. *Academy of management review*, 25(2), 312-334.

Shore, L. M., Chung-Herrera, B. G., Dean, M. A., Ehrhart, K. H., Jung, D. I., Randel, A. E., & Singh, G. (2009). Diversity in organizations: Where are we now and where are we going?. *Human resource management review*, 19(2), 117-133.

Simons, T., Pelled, L. H., & Smith, K. A. (1999). Making use of difference: Diversity, debate, and decision comprehensiveness in top management teams. Academy of management journal, 42(6), 662-673.

Stasser, G., & Titus, W. (2003). Hidden profiles: A brief history. *Psychological Inquiry*, 14(3-4), 304-313.

Suarez, F. F. (2004). Battles for technological dominance: an integrative framework. *Research Policy,* 33(2), 271-286.

Tata, A., & Niedworok, A. (2020). Is beauty in the eye of the beholder? An empirical study of how entrepreneurs, managers, and investors evaluate business opportunities at the earliest stages. Venture Capital, 22(1), 71-104.

Tetlock, P. E., & Gardner, D. (2016). *Superforecasting: The art and science of prediction*. Random House.

Tolstoy, L. N. (1877). *Anna Karenina*

Toubia, O., Johnson, E., Evgeniou, T., & Delquié, P. (2013). Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science,* 59(3), 613-640.

Turner, N., Swart, J., & Maylor, H. (2013). Mechanisms for managing ambidexterity: A review and research agenda. International Journal of Management Reviews, 15(3), 317-332.

Ulrich, K. (1994). Fundamentals of product modularity. In *Management of Design* (pp. 219-231). Springer, Dordrecht.

Weick, K. E. (1995). Sensemaking in organizations (Vol. 3). Sage.
Wiegmann, P. M., de Vries, H. J., & Blind, K. (2017). Multi-mode standardisation: A critical review and a research agenda. *Research Policy*, 46(8), 1370-1386.
Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686-688.

FULL PAPERS

"Happy families are all alike,
but every unhappy family is unhappy in its own way."
– Tolstoy (1877:1)

PAPER 1

"Organizations learn from experience.
Sometimes, however, history is not generous with experience."
– March et al. (1991:1)

Attention Processes Predict Problem-Solving Strategies: Evidence from Think-Aloud Protocols and Behavioral Experiments

Daniella Laureiro-Martínez[1], Jose P. Arrieta[1], and Stefano Brusoni[1]
[1]ETH Zürich, Switzerland

ABSTRACT

The cognitive micro-processes that managers use to solve strategic problems remain unknown. This paper relies on the less-studied cognitive approach to the attention-based view (ABV), to propose a theoretical lens that complements the study of problem-solving. Few studies have used primary, non-retrospective data to explore the process of problem-solving in the absence of feedback—perhaps due to methodological difficulties. To bridge this gap, this paper combines different methods in three studies. First, an exploratory lab study aims at understanding, with very fine-grained data, how strategic problems are solved in the absence of feedback. We employ think-aloud methods combined with content, sequence, and cluster analyses. We find that two problem-solving strategies emerge. One allocates more attention to the framing of the problem, and the other to the implementation of the solution. This result leads us to the second study, where we use a mixed factorial design experiment to pinpoint the causal mechanism that explains the emergence of the two strategies for solving strategic problems. We retest the hypotheses in a third, preregistered replication study. We find that manipulating attention towards cognitive processes related to framing increases deliberation aimed at restructuring the problem elements (i.e., a problem-focused strategy). In contrast, manipulating attention towards cognitive processes related to solution implementation increases reflection on the potential contingencies and consequences of the solution (i.e., a solution-focused strategy). We discuss how our findings can serve to extend research on problem-solving, the microstructure of organizations, and learning. We conclude by deriving managerial implications.

Keywords: attention; cognition; ill-structured; process; strategic problems;

1. INTRODUCTION

Since Cyert and March (1963), a foundational idea of both the behavioral theory of the firm and the Carnegie School has been that managers solve problems by adapting to feedback from their prior actions. Managers learn from experience—i.e., online (Levitt & March, 1988, Nelson & Winter, 1982). However, while this is true for some problems, it is by no means true for all. In fact, managers must very often act in environments where they not only lack relevant past experience, but must also solve the problem at hand without any possibility of experiential learning or receiving feedback on it (March, Sproull, & Tamuz, 1991). The need to solve problems in the absence of experience or feedback is perhaps most evident in strategic problems, which are characterized by their complexity (Simon, 1962), novelty, uncertainty (Mintzberg et al., 1976), and ambiguity (Nickerson & Zenger, 2004). Strategic problems involve high-stakes decisions (Eisenhardt & Bourgeois, 1988) that are often irreversible, and lead to outcomes that are very hard to predict (Ghemawat, 1991). Given these characteristics, it can be very costly, if not impossible, to engage in online learning when solving strategic problems. Hence, decision-makers must solve such problems in the absence of feedback or direct or relevant experience, through a process of offline learning.

Even though solving strategic problems is paramount for managers (Baer et al., 2013), and despite offline learning being at the core of strategic problem-solving, we know very little about *how* strategic problems are solved. Within the Carnegie school, Posen et al. (2018) call for us to "take a process approach" (2018: 240), claiming that "the literature has often been black-boxing the search process in the discussion of problemistic search[1], studying its antecedents and consequences without a rich connection to search itself" (Posen et al., 2018: 219). In a different stream, Langley et al. (1995) call for a departure from a middle-distance

[1] Problemistic search is the main problem-solving process in the behavioral theory of the firm (Cyert & March, 1963).

approach to organizational decision-making research, and exhort us to "zoom in closer to the people and processes under study" (Langley et al., 1995: 276). They argue that we should adopt a micro-perspective on attention, acknowledge inter-individual differences, and trace strategies in real time, in order to ensure that "perceptions are not biased by a knowledge of a final outcome, as has been the case in most decision making research (Schwenk, 1985)" (Langley et al., 1995: 276).

The attention-based view (ABV) has made less progress in explaining *how* attention can impact more radical forms of strategic change, as scholars have recognized (Ocasio, Laamanen, & Vaara, 2018). This is surprising when we consider that change can be regarded as a form of novel problem where the stakes are high, but the scope for experimentation is low. Moreover, most studies within the ABV have contributed to the organization-design perspective, rather than focusing on decision-makers themselves. When scholars have considered the human angle, they have dwelt on *what* decision-makers pay attention to, rather than *how* they attend to stimuli (i.e., their cognitive processes). In other words, the ABV has tended towards a structural rather than an individual view of attention, and emphasized the subject of attention over the manner in which attention is directed. In this paper, we respond to calls from the problem-solving and attention literatures by designing and implementing an empirical strategy to open the black box of strategic problem-solving in the absence of feedback.

This paper comprises three studies. Study 1 is aimed at building theory on *how strategic problems are solved*. We trace individuals' thoughts as they solve a strategic problem that is complex, ill-structured, and novel, involving high-stakes, irreversible decisions. We use a novel combination of methods that allows us to examine, precisely and in detail, the processes that managers follow while solving a strategic problem in a controlled environment. We combine think-aloud protocols with content analysis to give a clear view of how managers allocate their attention while solving a problem. Additionally, we use sequence and cluster analyses to extract and analyze the common strategies that managers follow while solving strategic problems.

When we analyze the process of problem-solving and the emergent patterns of attention, two clusters emerge. We find that in the absence of experience or feedback, decision-makers focus their attention on different problem-solving phases. One cluster focuses their attention on phases aimed at understanding various aspects of the problem, in order to obtain a rich framing of the situation. Another cluster, meanwhile, focuses their attention on engaging more deeply in simulating outcomes of potential solutions. Following Ocasio and Joseph (2018), we call these two emergent patterns of attention *strategies*. The first strategy, which we call *problem-focused*, allocates more attention to the phases related to the framing of the problem. The second strategy, which we call *solution-focused*, allocates more attention to the phases related to the implementation of the solution.

In our second study, we uncover the mechanisms that explain *why there are two different types of strategies*. To answer this question, we rely on an experimental study. This study uses a manipulation aimed at directing attention in different ways, to see whether we could observe the emergence of the processes found in the first study, and thus explain their causality. As participants solve the problems, they move a set of items into the ranking order they consider best for the goal they have been assigned. We use move analyses to uncover the underlying cognitive processes that precede the solution of the problem (Yu et al., 2012; Öllinger et al., 2013; Fedor et al., 2015). By studying these processes, we obtain a precise account of the differences induced in attention allocation according to the type of manipulation and the resulting behavioral changes. We thus achieve an understanding of the causal mechanism (i.e., the allocation of attention) that explains the emergence of the two strategies found in our first study.

We find that manipulating participants' attention towards different phases of problem-solving has the general effect of increasing the total deliberative effort they devote to solving the problem, with all manipulated participants devoting more time to solving the problem than those in the control condition. Interestingly, when asked to focus their attention on the framing of the problem, participants spend the additional deliberation effort on restructuring the problem elements, which could be observed as taking more moves to solve the problem. In contrast, participants asked to focus their attention on the solution use the additional deliberative effort to pause longer and reflect more deeply before each move. Both manipulated groups expend an equivalent deliberative effort on the task, but the way they allocate their attention is reflected differently in each group, according to the type of strategy each group develops. We then run a third study in which we preregister the hypotheses of the second study and replicate it.

On the basis of the findings from our three studies, our main contribution is to explain how attention allocation leads to the emergence of different strategies. Past studies have mainly focused on how experience is gained from solving problems, and the behavioral consequences of the experience thus gained. In contrast, we study strategic problems where direct or relevant experience is not available, and focus on the less-studied cognitive approach to the ABV in order to propose a theoretical lens that opens the black box of problem-solving micro-processes.

This paper makes at least four significant contributions. First, by relying on fine-grained data, it provides a highly detailed account of the processes that emerge when individuals solve problems in the absence of experience or feedback. In many strategic situations, it would not be feasible to gain experience through experimentation. For this reason, it is useful to understand precisely how managers solve problems in the absence of feedback. Understanding these processes complements existing models of attention, search, and problem-solving. It answers Posen and colleagues' (2018) call to open "black-boxed elements of the search process," describing—with very fine-grained and concurrent data—the emergence of different strategies.

Second, this paper's findings build theory that allows us to predict which strategies managers will use when they engage in solving strategic problems, and the mechanism that differentiates those strategies. Our theoretical contribution builds on the Carnegie School (Gavetti, Levinthal, & Ocasio, 2007) and tackles a gap in our knowledge of the processes of problem-solving (Langley et al., 1995; Posen et al., 2018) by combining theories of managerial problem-solving and decision-making (Langley et al., 1995; Nickerson & Zenger, 2004; Klingebiel & De Meyer, 2013; Felin & Zenger, 2016) with the ABV as a lens to study strategic problem-solving (Ocasio, 1997; 2011; Ocasio & Joseph, 2018). So far, most ABV studies have focused on the structural drivers of attention, rather than decision-makers themselves. The few studies that do consider decision-makers have focused on *what* attention is paid to (i.e. the targets or the stimuli attended to), rather than *how* decision-makers pay attention. We complement the ABV by focusing on the cognitive micro-processes that explain how attention allocation is the primary origin of problem-solving strategies. We test whether what causes the emergence of the different strategies is indeed the allocation of attention to different cognitive processes. With this focus, we expect to complement the ABV by opening up the black box of the problem-solving attentional processes in which decision-makers engage.

Third, this paper makes a methodological contribution by combining new and established techniques to establish a novel method for building and testing theory in two connected studies, in the manner of Reypens and Levine (2018). The first study is exploratory, and builds theory by exploring the processes that emerge under a controlled environment. It combines time-honored techniques (i.e., think-aloud protocols) with more recent ones (e.g., sequence analysis). The second and third studies, meanwhile, are confirmatory in nature. They test the findings of the first study using a mixed factorial design experiment with three conditions and each participant solving two problems. The combination of methods is a good example of the cycle of theory-building and theory-testing: methodologically complex, but foundational to the growth of scientific knowledge (Popper, 1963).

Fourth, this paper offers potentially valuable ideas to practitioners, as many management methods can be interpreted as methods that affect the attention focus of decision makers. For example, both Lean manufacturing and Six Sigma emphasize the analytical, solution-oriented phase of problem solving. Scenario analyses prompt people to think of problems and opportunities in the future, while Design Thinking tools push individuals to shift their attention cyclically. Our findings clearly show that methods that "manipulate" attention in different directions are useful—but different people might use them and learn from them with varying levels of effectiveness depending on their predispositions (Nickerson, Silverman, & Zenger, 2007: 215–216). This does not mean that people can or should not switch between different strategies, but rather that since they tend to follow different processes, certain management methods might work better with some people than others, in a predictable way. Since tools like Lean and Six Sigma are widely used in organizations, understanding how people actually solve problems and benefit from such approaches can help organizations shape them for a better fit with different individuals' developmental needs.

This paper is divided into nine sections. Following this introduction, section 2 presents our theoretical framework. Third, we introduce our methods for Study 1, and summarize its results in section 4. Sections 5 and 6 present the methods and results of Study 2, respectively. Sections 7 and 8 present the methods and results of Study 3. Finally, section 9 concludes with a discussion of managerial and theoretical implications.

2. THEORY

In this paper, we take a micro-level view of problem-solving. We base our work on the problem-solving perspective (Nickerson & Zenger, 2004), and use the ABV (Ocasio, 1997, 2011) as a lens that allows us to design a fine-grained study of the micro-processes involved in problem-solving as they unfold. In this section, we start by presenting the phenomenon of interest: strategic problem-solving. We then discuss what prior studies tell us about solving strategic problems: first, studies of the different phases that are involved in solving a strategic problem, and second, studies exploring the sequence in which such phases unfold.

2.1 Defining strategic problems

Strategic problems are different from other types of problem in several ways. They involve an *irreversible* decision; they involve *high stakes* with significant upsides or downsides; and they are *complex, novel*, and *ill structured*[2]. There is a rich literature on each of these five characteristics. For example, a strategic problem needs to involve high stakes because if there is no risk involved, making the decision incurs no potential cost or gain for the organization (Eisenhardt & Bourgeois, 1988). The same holds for irreversibility (Ghemawat, 1991): if the decision can be reversed without significant costs, then the problem is operational rather than strategic. Levinthal (1997) started an important discussion on how complexity provides a competitive advantage, and Gavetti, Levinthal, and Rivkin (2005) added to it by explaining how novelty opens up strategic opportunities. In addition, the structure of a strategic problem is a defining factor. In an ill-structured problem, the means-end relationships, initial states, end states, and actions tend to be ill defined, so the decision-maker can never be sure about the possible solutions they might reach. An ill-structured problem, therefore, is inherently uncertain.

Fernandes and Simon (1999: 226) defined six different dimensions that a problem must exhibit in order to be considered well structured. In sum, they explain that the goals, beginning and end state, actions, constraints and knowledge the problem solver can acquire need to be well defined. Each of the six dimensions can be more or less structured, and thus problems will

[2] Uncertainty is another key characteristic, but if the problem is ill-structured, it is necessarily also uncertain, as we explain below. For parsimony, we do not include uncertainty as a key characteristic of strategic problems.

differ greatly depending on the level of structure on each dimension, creating multiple types of ill-structured problems; in contrast, in any well-structured problem, all six characteristics are equally well defined. Perhaps for this reason, the literature on solving ill-structured problems appears to be rather dispersed. In the next section, we show how the literature has tried to understand how problems are solved by proposing different phases, and sequences of such phases. We then present a model that combines multiple research streams that study problem-solving.

2.2 How are strategic problems solved?

2.2.1 The phases of problem-solving

Problem-solving has been studied since at least the early 20th century (e.g., Dewey, 1910). Research has studied how organizations develop solutions (Mintzberg et al., 1976) and how individuals solve their everyday problems (Klein, 1997) or make judgments (Tversky & Kahneman, 1975). The literature has taken some giant leaps forward over the past 50 years, with significant attention being paid to the study of well-structured problems. This has given us a deep theoretical understanding of the process through which such problems are solved. For example, studies have shown how solutions are affected by the speed with which they are arrived at (Ratcliff & McKoon, 2008), the biases and heuristics of the decision-maker (Tversky & Kahneman, 1975; Gigerenzer & Goldstein, 1996), or the general effect of problem framing (Baer et al., 2013).

Studies have proposed models[3], usually following some sequential steps, that summarize the process by which well-structured problems are solved. However, in order to study strategic problems, we must depart from the safe harbor of well-structured problems and

[3] Depending on the study, models include more or fewer phases, and more or less rigid sequences. Simon's (1965) model contained three phases (*intelligence, design,* and *choice*). Later on, models tended to break these phases into sub-phases where specific actions took place. A recent and very well-established model by Rangel et al. (2008) starts with the phase of *representation*, where the decision-maker recognizes the different actions possible in this setting. There follows the *valuation* phase, where the value of each alternative is assessed, according to individual wishes. This is followed by the *action selection* phase, where a choice is made. The final phase is *outcome evaluation*, which evaluates the desirability of the choices; this assessment is then internalized through *learning* in case the same problem has to be solved again.

set sail for the choppier waters of ill-structured ones. In contrast with well-structured problems, ill-structured problems have been more sparsely studied, and fewer models are available. Three foundational models are presented in Table 1. For example, models proposed by Simon (1965) and Mintzberg et al. (1976) acknowledge that for ill-structured problems, the set of actions is not given, and knowledge is not complete. Thus the process of problem-solving involves an initial phase that is absent from models of well-structured problems. In this phase, called *design* in Simon's model, the set of possible directions and actions is delineated and studied. In a similar vein, Schwenk (1985) recognized that many strategic problems involve high stakes and irreversibility, and added a phase corresponding to the *implementation* of the solution.

In order to understand the process of problem-solving more precisely, we believe it is useful for a model to include more phases rather than fewer—even if some of them might not take place in every scenario. Therefore, in this study, we present an integration of the different phases that previous models have proposed. This results in 7 different phases that can take place while solving a strategic problem.

Table 1 shows our combined model (in the shaded column). It separates *goal formulation and problem identification* from Schwenk's (1985) model into two phases, while retaining *valuation* and *action selection* from the more recent and well-established model devised by Rangel, Camerer, and Montague (2008)—in our model, denoted as *evaluation* and *decision*. In addition, our proposed model preserves the phases of *implementation, implementation evaluation*, and *direction setting* common to other models.

Having reviewed the phases involved in solving strategic problems, we now turn to the sequencing of such phases as the decision-making process unfolds.

Table 1: Comparison of different problem-solving models and their phases

Prior Management Models			Combined problem-solving model	Neuroscience model
Simon (1965)	**Mintzberg et al. (1976)**	**Schwenk (1985)**		**Rangel et al. (2008)**
Intelligence gathering	Recognition	Goal formulation and problem identification	*Frame stating (FS)*	Representation
	Diagnosis		*Frame assuming (FA)*	
Design	Search	Strategic alternative generation	*Direction setting (DS)*	
	Design			
Choice	Screen	Evaluation and selection	*Evaluation (EV)*	Valuation
	Evaluation		*Decision (DE)*	Action selection
	Authorization			
—	—	Implementation	*Implementation (IM)*	—
—	—	—	*Implementation evaluation (IE)*	Outcome evaluation

2.2.2 The sequencing of problem-solving phases

In a landmark paper, Mintzberg et al. (1976) argued that the process through which managers solve problems and reach decisions involves multiple transitions, some between phases that do not follow the order expected from a linear model. They showed that in many decisions, managers allocate very little of their deliberation time to the initial phases of problem-solving, and cycle back and forth repeatedly between some others. Brightman (1978) explained how for complex problems, each phase of problem-solving is a micro-problem-solving process in itself, necessitating smaller cycles within the problem-solving process. Fernandes and Simon (1999) showed that while solving complex and ill-structured problems in a think-aloud protocol, individuals cycle through phases of analysis in different manners depending on their professional background (lawyer, physician, architect, or engineer). Their study focused on cognitive processes, not problem-solving phases, and presented only two participants per condition, limiting its generalizability. However, the manner in which Fernandes and Simon (1999) studied the process of problem-solving makes their paper a remarkable example, given the study of the process in real time (in contrast with most other studies, which use retrospective techniques) and the level of granularity, which supports a deeper understanding of the problem-solving process.

Langley et al. (1995) put forward a complex view of decision-making and problem-solving. In all these studies, the problem complexity resulted in "messy" sequences with dynamic linkages between phases that are attended to differently. To understand how attention is allocated to different phases, we build upon Simon's (1947) idea that decision-makers are influenced by both cognitive and structural aspects that drive their attention. We focus on the less-studied cognitive approach to the ABV to propose a theoretical lens that complements the study of problem-solving.

2.3 Strategic problem-solving and attention

Strategic problem-solving can be studied through different lenses and at different levels of analysis. The problem-solving perspective started at the organizational level (Nickerson & Zenger, 2004) before progressing to the meso level. This perspective limits its theorizing to teams, and excludes the micro-processes followed by individuals. However, this perspective does recognize that "numerous individual-level decision biases exist" (Baer et al., 2013: 200). In this paper, we study individual-level micro-processes and use the ABV as a lens to examine the antecedents of behavior. Crucially, the ABV allows us to infer the strategy of an individual from how they direct their attention (Ocasio, 2011; Ocasio & Joseph, 2018).

In the ABV, strategy is defined by attention rather than action—a departure from prior theories. For example, Andrews (1971) defined corporate strategy as "the pattern of decisions in a company that determines and reveals its objectives, purposes, or goals" (Andrews, 1971: 13). The ABV adopts a more processual view, defining strategy as "a pattern of attention" rather than a set of actions (Ocasio & Joseph, 2018: 289). Within this view, "what decision-makers do depends on what issues and answers they focus their attention on."

The ABV provides a meta-theoretical structure to explain how attention is a key resource to be managed by organizations (Ocasio, 1997, 2011; Ocasio & Joseph, 2005). The ABV can be seen as having two different pillars, set in an environment where a decision is made (Ocasio, 1997: 192). The first pillar is normative, and relates to the design of the organization and the internal processes that guide attention within a firm (Joseph & Wilson, 2018). The second pillar is descriptive: it focuses on the targets of decision-makers' attention. Ocasio & Joseph (2018) proposed an organization-level explanation of how organizational structure leads to the allocation of attention on a focused set of problems and creates value. In a stance that complements Joseph and Wilson (2018), we propose to study the second pillar of the ABV by focusing on how the allocation of attention results in the strategies used to solve a strategic problem (Ocasio & Joseph, 2018). Using novel methods, we focus on the cognitive micro-processes that explain how attention allocation causes problem-solving strategies to emerge. This is in contrast to most empirical ABV studies of how attention affects strategy-making, which have focused on the *content* of attention (i.e. the stimuli) rather than *how* attention is allocated.

A review of ABV papers published between 1997 and 2020 that cite Ocasio's foundational article of 1997 shows that, for the most part, the literature has addressed questions such as "*Where* is attention directed?" and "*What* is the content of attention?" For example, studies have focused on studying where environmental influences direct organizational attention (e.g. organizational architecture influencing the distribution of attention in Crilly and Sloan (2014), the structural elements and diverse composition of top management teams in Cho and Hambrick (2006), or organizational structures and processes and attention triangulation for learning in Rerup (2009)). Other studies, meanwhile, have focused on the content of the organization's attention, and the features to which organizational members attend (e.g. senior managers eyeing competitors in Levy (2005), CEOs monitoring cognitive group membership in Surroca, Prior, and Tribo-Gine (2016), top management's consideration of technological discontinuities in Maula, Keil, and Zahra. (2013), or CEOs pondering future events in Yadav, Prabhu, and Chandi,(2007)).

Within the ABV, studies that deal with problem-solving have focused primarily on *what* decision-makers focus their attention on. For example, Sullivan (2010) explores the characteristics of new problems that affect the attention paid to solving old problems, while Maslach, Branzei, Rerup, and Zbaracki (2018) explore how decisions about which features to attend to affect how organizations learn from rare events.

Cognitive science has a similar focus. First, studies that directly investigate the impact of attention on problem-solving are quite scarce, since research on attention focuses on lower-level processes such as perception of impulses, whereas problem-solving involves higher-level cognitive processes such as attention control (Rouinfar et al., 2014). The few studies that explicitly explore the role of attention in problem-solving mainly focus on visual attention—in particular, what decision-makers look at. For example, Thomas and Lleras (2007) and Grant and Spivey (2003) found that redirecting visual attention to problem-relevant information through cueing led to more accurate problem solutions.

To the best of our knowledge, just two studies have looked at *how* attention is allocated. The first, by Li, Maggitti, Smith, Tesluk, and Katila (2013), takes a pioneering look at how the decision-makers of the ABV pay attention. However, the data for the attention-intensity variables comes from a survey in which top managers had to remember how much time and energy they had devoted to certain issues and stimuli, which is likely to suffer from retrospective biases. In a second study, also very novel, Frankenberger and Sauer (2019) analyzed "attention intensity" using interview questions, where the central aim was to understand whether the focus and the intensity of attention could have an impact on business models. Overall, however, the direct link between the processes involved in paying attention and resultant problem-solving strategies remains unexplored. Moreover, to effectively understand how attention is paid, we need a novel method that does not rely on introspection.

Summing up, there is a gap in our understanding of how strategic problems are solved. Past studies have proposed that solving strategic problems must involve multiple iterations between different phases, deviating from linear models. Given the characteristics of strategic problems (i.e., novelty, lack of structure, complexity, high stakes, and irreversibility), we might expect decision-makers to devote much of their attention to structuring and simplifying the problem. Indeed, as Einstein famously observed, "The formulation of a problem is often more essential than its solution…" (Einstein & Infeld, 1938: 92). We do not expect the process of solving a strategic problem to follow a linear sequence. Instead, we follow March's idea that managers must solve the most difficult problems, and that "unfortunately, God gave all the easy problems to the physicists. It is difficult. It's a world that is complex, that is shifting all the time" (Dong, March, & Workiewicz, 2017: 12). Due to this complexity and lack of structure, we expect the problem-solving process to be characterized by multiple iterations across different phases. In addition, we expect decision-makers to devote more attention to phases that aim at framing the problem and structuring its key elements, to familiarize themselves with the novel and complex situation and give it some structure. In Study 1, we carry out an exploratory study to investigate this initial expectation and get a clearer picture of the attentional processes that affect how strategic problems are solved.

3. STUDY 1

"To grasp cognition in action," Reypens and Levine (2018) recommend that we "combine experiments with protocol analysis." Following this methodology, our first study focuses on protocol analysis to understand cognition in action (Andrews, 1971).

In this study, we examine the problem-solving process of experienced managers by employing a combination of think-aloud protocols (Ericsson & Simon, 1980) and content, sequence, and cluster analyses. We use these techniques firstly to collect fine-grained data, and then to reduce its dimensionality in a structured way, in order to avoid discarding meaningful insights.

We present our methodology in three parts: the development of the problem, the data collection, and finally, the data analyses.

3.1 Strategic problem: The “Karabayos” problem

In this study, we employ a problem that has been tested and validated in a prior management study (Laureiro-Martínez & Brusoni, 2018): the “Karabayos” problem. The problem requires participants to imagine themselves as the leader of a small aboriginal tribe, managing limited resources, under threat from external invaders. The objective of the tribal leader is to keep the tribe safe. Setting the problem in a distant geographical location and non-business context does not prevent it from fulfilling all the essential characteristics of a strategic problem. In fact, the task shares many commonalities with difficult situations that managers, team leaders, and entrepreneurs might face when leading their groups to a common goal. First, the problem is both complex and ill structured: a starting point is provided, but it entails contradictions. The problem involves several major uncertainties: the time available to achieve the goal (i.e., how long the leader has to save the tribe, when enemies will attack, etc.), the reactions from relevant stakeholders (e.g., the level of resistance to their actions decision-makers encounter), and even how the primary goal is defined (e.g., it could be to save the current generation alone, or to ensure that future generations can survive)—among others. Moreover, neither the possible actions nor their outcomes are well defined (e.g., can I communicate with the “enemy?”), and there is a potentially infinite range of alternatives to explore. The “Karabayos” problem also presents participants with a high-stakes, irreversible decision. The tribe might survive, or it might perish, and there is no possibility of receiving any process or potential performance feedback as events unfold. An additional advantage of using this task is that we wanted to avoid past experience with the specific problem setting, and it was easy to find participants without experience in the context of the tribal leader problem.

3.2 Data collection: Think-aloud protocols

The think-aloud protocols used were initially collected for the study by Laureiro-Martinez & Brusoni (2018), but the coding and analyses carried out for this study were completely different. Think-aloud protocols follow a similar temporal flow as silent thinking (Ericsson & Simon, 1998), and provide the researcher with an unobtrusive and more accurate reflection of the thinking process than retrospective verbal protocol analysis, or descriptions and explanations of the thinking process (Kuusela & Paul, 2000; Ericsson & Simon, 1998).

Participants were assessed individually in a quiet and secluded location and given written and verbal instructions, which in turn were preceded by training sessions. We followed Ericsson and Simon's (1998) method to instruct participants about how to produce consistent, non-reactive verbalized thoughts during problem completion. More specifically, we emphasized how "thinking aloud" differs from "verbal reporting." All participants completed a minimum of three exercises to make them familiar and comfortable with the think-aloud method. After each exercise, participants received verbal feedback. The study's problem was presented only when the participant felt comfortable with the method, and the researcher was satisfied with the technical aspects of the verbalizations (i.e., the speed, vocalization, and type of language did indeed reflect thinking, and not a retrospective verbalization). Participants required as many as six familiarization exercises before sufficient reliability was achieved.

3.2.1 Potential issues with think-aloud protocols

There are three main issues related to the use of think-aloud protocols that can affect the reliability of the data. Below, we summarize each of them and present our solutions, consistent with the state of the art as described in Ericsson (2003) and Fox, Ericsson, and Best (2011).

First, the method might put pressure on participants, leading to biased responses. To avoid this, we took three main measures. First, during the training phase, we clearly distinguished between thinking aloud (what was requested: that is, a mere vocalization of inner speech) and verbal reporting (what was not requested: the product of additional reflective and analytic introspections) (Fox, Ericsson, & Best, 2011). Second, we informed participants that, during their search for a solution, no interaction with the researcher would be allowed. The researcher would intervene only if the participant failed to think aloud, and then merely remind them to verbalize their thoughts. During the training phase, we told the participants about our interest in their thinking process. Third, in order to avoid differences among participants due to time pressure, we told participants that they could work on this task for as long as they wanted. We had reserved two hours of the participants' time, and all participants took less than half of this, the longest taking 51 minutes. Hence, participants had no reason to rush their answers. They were directed to signal when they had arrived at a solution, which indicated the end of the task.

A second issue is that the verbalization might not reflect the actual thinking process, but a narrative created by the participant to paraphrase their thinking process to the researcher. To address this issue, we used concurrent think-aloud protocols. The participants had to verbalize their thoughts during the process of solving the problem "in real time," without seeing the problem beforehand, which avoids retrospective and introspective biases.

A third issue is that the verbalization might reflect the talkativeness of the participant, rather than their thinking process. To prevent this, we carefully instructed participants not to discuss, describe, or explain how they solved the problem. Instead, we told them to remain focused on solving it, and to verbalize those thoughts that emerged in their attention while generating the solution under normal (silent) conditions. Having completed the task, we asked the participants to restate their final answer, and then the debriefing started. The aim of the debriefing was for us to understand the general experience while solving the problem and to check whether the participants had experience with this kind of problem; none had any experience in a similar context. All participants gave serious thought to how to solve the difficult problem they were confronted with, and many reported that they had empathized with the role of the tribal leader.

3.2.2 Sample

The sample, as in Laureiro-Martinez and Brusoni (2018), comprised 49 participants. All were managers, executives in multinational firms, founders of small companies, or unit managers in medium-sized organizations. Participants had at least four years' management experience, were responsible for budget allocation, and played a leadership role in a group with at least two other members. Fifteen of the participants were entrepreneurs, while 34 were experienced managers working within firms. The sample consisted of 40 men and nine women, with an average age of 35 years (s.d. = 6.7 years, with a range between 24 and 47 years). Participants were offered financial and non-financial incentives for taking part (detailed in the Supplemental Materials of Laureiro-Martinez & Brusoni, 2018).

The processing of think-aloud protocols is complex and time-intensive. For this reason, previous studies based on think-aloud protocols have worked with 15 or fewer participants (Grégoire, Barr, & Shepherd, 2010; Sarasvathy, Simon, & Lave, 1998; Isenberg, 1986; Fernandes & Simon, 1999). Our sample, while still small for quantitative analysis, is larger than that of similar studies.

3.3 Data analysis

3.3.1 Content analysis

After we collected the think-aloud protocols, they were transcribed verbatim by research assistants involved in the project. To ensure the quality of the data, we followed transcription protocols that helped us to systematically prepare the transcripts and minimize any threats to their quality (McLellan et al., 2003; Poland, 1995). We analyzed participants' verbalization using content-analysis techniques (Krippendorff, 2012; Neuendorf, 2002). For each protocol, we analyzed the content and, with the help of three independent raters, selected the specific passages that represented “chunks of thought” corresponding to specific problem-solving phases. The protocols were coded according to the seven phases of the combined model presented in the theory section and the shaded column of Table 1.

Table 2 presents a more detailed view of how the coding was operationalized. The table presents the seven phases of problem-solving, a short description of the construct and the type of processes involved, and a quote representative of each specific code. The initial phase is *frame stating* (FS), in which the participant analyzes the problem by repeating or paraphrasing the data mentioned in the problem description (a text that was provided to each participant). *Frame assuming* (FA) follows when the participant develops their own hypotheses and assumptions about the problem at hand and begins taking them for granted, even when they were not mentioned in the problem description. *Direction setting* (DS) consists of defining general paths one intends to follow without stating a specific proposal, or generating alternative proposals for what to implement later on. *Evaluation (EV)* is when the participant judges the merits of a proposed path, and considers the solution without evaluating specific details of it. The *decision (DE)* phase is when the participant manifests a clear choice regarding what they intend to do. In *implementation* (IM), the participant designs the sequence of actions to carry out their proposals. The seventh stage is *implementation evaluation (IE),* where the participant evaluates the feasibility of their implementation. We codified any unintelligible sounds as *babble*.

The raters were tasked with coding every word of the think-aloud protocol into one of the seven phases of problem-solving or babble. We should highlight that in order to achieve a more objective interpretation of the think-aloud protocols, the researchers were involved in refining and piloting the code, but not in the actual coding process.

Table 2: Problem-solving phase coding definitions

Problem-solving phases	**Description**	**Examples of verbalized thoughts (transcribed verbatim)**
Frame stating (FS)	Repeating the data mentioned in the text of the problem	"…so our area want to be left alone we are vulnerable that we have understood for a good reason … I mean here I do not have other information problems diseases a very small zone lack of food…"
Frame assuming (FA)	Development of hypotheses not mentioned in the problem	"… for millennia and before me, my father, my grandfather, and all the others one after the other without having to face things that were more difficult go hunting sometimes or collect fruit…"
Direction setting (DS)	Defining a general path of actions to be followed and generating proposals about what should be done	"… we can also be a means for, a means to attract, for your region, we can, we can make people, we can, we can help you make I do not know a museum something we can make lessons to teach city kids how to love the forest…"
Evaluation (EV)	Evaluating and judging the proposal and considered their strategy without evaluating specific details	"… sending two or three people can be interesting… even though most likely those two or three won't return…"
Decision (DE)	Making an explicit choice about what intended actions	"…however I will try to dialogue this for sure I will try three key points dialogue with another civilization support from my group and away and an alternative in case of failure of dialogue…"
Implementation (IM)	Designing a sequence of actions required to carry on the proposed actions	"…slow calm we arrive in front of a representative we try with presents with kids with women and with men with those most intelligent to craft a speech even with gestures drawing we ask for help and we see if they help if not we try alone we do not explain where we are because if we explain because if we have to try at least they don't know where we are… we return…"
Implementation evaluation (IE)	Evaluating the possible actions' outcomes	"…is clear that it is not easy because probably out the jungle a someone some member of my tribe will hardly survive but is an endeavor to try…" "…if the two people [that were sent away before] should not return however 46 people will still be alive if instead return with a positive answer we have solved at least for some time long enough the problem…"

In order to ensure the robustness of our results, we calculated two measures of inter-coder reliability. The first was the average percentage of agreement, which was 93.4%. Average agreement is useful in the case of simple codes, but when the data is complex, prior studies recommend using Cohen's Kappa. We found a value of 0.51 for this metric. Both values are satisfactory for the type of text we studied (Neuendorf, 2002; Lombard, Snyder, & Bracken, 2002).

3.3.2 Code merging

Each rater provided a fully coded transcript for each participant's protocol. Although we achieved high reliability, a perfect match for every word in every protocol is almost impossible. However, a prerequisite for sequence analysis is that each passage must be assigned a single code. We therefore followed a second content-analysis process where we compiled the coded transcripts of each rater and followed a simple process of code merging.

By code merging, we mean taking multiple codes for a single passage and converting them to a single code. Our code-merging process had three steps. First, in cases of consensus between the three raters, we kept the agreed-upon code. Second, in cases of partial agreement (i.e., two raters select the same code, and one disagrees), we saved the value chosen by the majority. Third, in cases of complete disagreement between the raters (i.e., all three assign different codes), two authors conferred and selected the appropriate code for the passage in question from the three codes proposed by raters.

The output of these three steps was a fully coded transcript in which every passage was coded into a single problem-solving phase. This resulted in a sequence of phases for each participant that represented their entire problem-solving process. At this stage, we removed the *babble* codes (i.e., unintelligible sounds, which accounted for 2.8% of the protocols in total) from the sequence, since they do not represent the problem-solving process.

3.3.3 Sequence analysis

Next, we shifted our attention from the content of the phases to the transitions between them. Although the duration of phases can vary widely, we assigned them all the same unitary length for the purposes of this analysis, in order to focus solely on transitions.

Figure 1 illustrates the problem-solving sequences of two participants: Person A and Person B. Problem-solving phases are shown as color-coded rectangles, defined in the key shown above Figure 1.

The problem-solving sequences of Person A and Person B differ considerably, although both employ all the phases of problem-solving. Person A focuses more on *frame stating* and *assuming*, and only spends time on *implementation* and *implementation evaluation* towards the end. Person B, in contrast, performs *frame assuming* and *frame stating* on far fewer occasions, and performs *implementation* and *implementation evaluation* earlier and more often.

Person A follows a more standard way of solving a problem, devoting considerable attention to understanding the situation first, and only then making decisions and implementing the solution. In contrast, Person B performs many *decision* and *implementation* rounds throughout the protocol, with considerably less *framing.*

3.3.4 Transition matrices

We reduced the variance of the information comprising the problem-solving sequences by creating transition matrices. Transition matrices provide comparable summaries of the participants' problem-solving processes. Our work on transition matrices is based on the research by Lipshitz and Bar-Ilan (1996), who developed a lag-analysis of transition probabilities between problem-solving phases in recollections of successful and unsuccessful problem-solving cases in military organizations.

Figure 1: Example of problem-solving sequences of two individuals

Lipshitz and Bar-Ilan (1996) focus their analysis on a matrix in which each cell represents a transition between phases. The starting phase of the transition is given by the row of the cell, while the destination phase is denoted by its column. These structures were originally referred to as "transition matrices" by Gibbs et al. (1971). Transition matrices are used to study similarities between sequences, focusing on the number of transitions between elements in the sequence (i.e., problem-solving phases). For instance, in the case of Lipshitz and Bar-Ilan (1996), their focus was on studying the order of events rather than their duration, as is commonly the case in other analyses.

In this study, we have seven phases, which give rise to 42 transitions between phases. The 42 values are entered in the off-diagonal cells of the transition matrix and represent all the transitions made by the participant during their problem-solving process. We normalized these values to obtain transition numbers that were comparable between participants, i.e., for each protocol the sum of all transitions (i.e., off-diagonal cells) sum up to 1. In addition to the transition between phases, we included the percentage of time spent in each of the seven problem-solving phases. In sequence analysis, it is common to have transitions within the same phase. However, think-aloud protocols do not have clear transitions within thoughts for coding within-phase transitions. As a proxy for within-phase transitions, we take the percentage of time spent on each phase.

The 42 normalized transitions and seven time allocations comprised the data we used to compare the problem-solving processes of the participants in our sample. Although we created 49 variables to characterize a problem-solving process, we performed cluster analysis on this data to reduce the dimensionality of this data to one categorical variable.

3.3.5 Example of linear problem-solving

Table 3 and Figure 2 illustrate how to read the transition matrices we use in this study. For illustrative purposes, we start by using a simple linear model. As stated above, in a transition matrix such as Table 3, the row denotes the starting phase and the column denotes the destination phase, with each value denoting the frequency of that transition. For example, the transition between *frame stating* and *frame assuming (FS → FA)* was made 16.7% of the time.

The transition matrix of Table 3 depicts a linear model because there are only transitions in the cells next to the diagonal (every other cell has a 0 value). The key (uppermost bar) of Figure 1 depicts this transition matrix in sequence form: seven phases, one after the other, from *frame stating* to *implementation evaluation*. There is no circling back to *frame stating*, as the value of that transition is zero. Since there are only six transitions, each represents 16.6% of the total. Additionally, each phase of problem-solving has the same duration: 14.3% of the total.

Figure 2 provides a visualization of the transition matrix depicted in Table 3. The sizes of the circles denote the duration of each phase, while the widths of the connecting lines denote the frequency of the transitions between them. In this simple linear model, all the circles are of equal size, since each phase lasts for the same amount of time. Similarly, the linking lines are all equally wide, as each transition is made an equal number of times, i.e. once. In the results section below, we present more sophisticated visualizations for non-linear cases where the transition frequencies and percentage of time spent vary.

Table 3: Transition matrix for a flat and linear sequence

Flat & Linear	**→ FS**	**→ FA**	**→ DS**	**→ EV**	**→ DE**	**→ IM**	**→ IE**
FS →		16.67	0.00	0.00	0.00	0.00	0.00
FA →	0.00		16.67	0.00	0.00	0.00	0.00
DS →	0.00	0.00		16.67	0.00	0.00	0.00
EV →	0.00	0.00	0.00		16.67	0.00	0.00
DE →	0.00	0.00	0.00	0.00		16.67	0.00
IM →	0.00	0.00	0.00	0.00	0.00		16.67
IE →	0.00	0.00	0.00	0.00	0.00	0.00	
% of thinking time	14.3	14.3	14.3	14.3	14.3	14.3	14.3

Figure 2: Visualization of a linear model transition matrix

3.3.6 Clustering

The transition matrices for each participant provide information about how they allocated their attention across the seven phases of problem-solving. We use clustering algorithms to extract the commonalities between the transition matrices. By using clustering, we can classify the common patterns of attention our participants used when solving a strategic problem—that is, the strategies they followed (Ocasio & Joseph, 2018: 289). We did not cluster via the optimal matching of sequences, as in Salvato (2009), because sequence length varied significantly between participants and led the optimal matching algorithm to give spurious results. Namely, it matched protocols by sequence length; this occurred regardless of algorithm setting. We chose to use transition matrices instead, as they are indifferent to the length of the sequence.

We employed a clustering method called *partitioning around medoids* (Kaufman & Rosseeuw, 1990). This method selects the best number—k—of clusters for a data set, and groups the rest of the participants around a set of the k most representative participants, called "medoids." The benefit of this method compared to others is that its clustering output is consistent and deterministic. The categorical variable assigns the same observations to the same cluster every time—something that k-means and other non-medoid clustering methods cannot do, except for clearly separate data sets.

We clustered our data using the *pamk* method from R's library *fpc* (Hennig, 2015). The *pamk* method starts by developing a variance ratio criterion (Calinski-Harabsz index) to determine the number of clusters, and then goes on to estimate whether there is a real benefit from splitting the dataset into two clusters (Duda-Hart test). The procedure is followed by an iterative process known as the building phase. In this phase, a total of k participants are selected and designated as medoids. Subsequently, a matrix of each of the remaining participants' dissimilarity from every medoid is calculated. Finally, the algorithm places each of the remaining participants into a cluster, minimizing dissimilarity among clusters. The following phase is swapping: one participant is exchanged for a medoid, the average dissimilarity of the new configuration is calculated, and if it is lower than the original, then the change is saved. The process continues until it results in the set of k medoids that provides the lowest average dissimilarity from its cluster members.

After the *pamk* function was completed, we were left with a categorical variable that assigned each think-aloud protocol to one cluster. In our case, it is a dichotomous variable. This dichotomous variable is the outcome of the structured dimensionality reduction procedure of this study. We started with 49 think-aloud protocols, all completely different. We carried out content analysis and merged the coding differences we found. We conducted a sequence analysis and transformed these sequences into transition matrices to create comparable data structures that captured the problem-solving processes of every participant, independent of their length. We then used a robust clustering procedure, *partitioning around medoids*, to create a single variable that summarizes the similarities between the think-aloud protocols.

In the results section, we use this clustering variable to characterize how the participants assigned to each value solved the "Karabayos" problem. From this, we can reach an understanding of how these problems are solved, and how problem-solving approaches differ. Instead of seeing strategic problem-solving as a homogenous process, we can study the commonalities and differences within it. Having reduced our data to a single key variable, we can use it to attain insights into how managers solve problems in the absence of feedback.

3.3.7 Performance scoring

We also coded the performance of the solutions given to the "Karabayos" problem. Two coders (different from those who coded the problem-solving phases) were assigned all 49 protocols, and each coder independently assigned a score based upon how well the participant's solutions fulfilled the problem's objective—namely, to "save your tribe." To do this, the coders first read the participant's entire protocol to familiarize themselves with the participant's problem-solving. They then classified the solution into one of three categories: "solved the problem and reached the objective," "somewhat likely to achieve the objective," and "unlikely to achieve the objective." After this, the coders reread the protocol and scored the solution on a scale of 1 to 10. The scores exhibited acceptable interrater reliability of 92.2%. After all the scores were assigned, the two coders conferred in order to reach agreement on those cases where their respective scores differed. We used the agreed-upon score as our *performance* value.

3.3.8 Control measures

We collected a further set of variables to explore alternative explanations for the clustering results. From the task, we recorded the total time spent solving the problem (*protocol duration*). We also asked participants to perform two further tasks to gauge their cognitive skills. Participants answered a 10-question Raven's Progressive Matrices test, which is correlated with abstract thinking (Laureiro-Martínez, 2014). Participants also solved a "Tower of Hanoi" task, which is known to measure planning and generativity skills (Laureiro-Martínez, 2014). We recorded the time it took them to finish the task (with longer times indicating worse planning). Finally, we added controls for individual characteristics: *age*, *gender* (female = 1, male = 0), and *profession* (entrepreneur = 1, manager = 0).

4. STUDY 1: RESULTS

In this section, we present the results obtained from the analysis of the participants' think-aloud protocols. We start with an introduction to the transition matrix of the full sample, and then explain the clustering procedure employed and characterize the different strategies that emerge from the clustering algorithm. We end by presenting an assessment of other possible explanations for the strategies and differences found.

4.1 Full-sample transition matrix

Table 4 presents the average transition matrix for all 48 participants[4]. The participants' transition matrices allowed us to study their patterns of attention as they solved the "Karabayos"

[4] As the raters coded the protocols, they informed us that one protocol was quite different to the others in that the thoughts of the participant were mainly devoted to numerical calculations, based on assumptions and not information provided in the problem. After coding, we compared the protocol to the others and decided to remove it from the sample. On average, participants spent 68% of their time on the frame stating, direction setting, and

problem. In this study, we assess a participant's allocation of attention by examining the percentage of time they spend on each problem-solving phase, and the number of transitions they make to and from that phase. The participants differed greatly in how they allocated their attention to the different phases of problem-solving, and we use these differences to understand how they solved the "Karabayos" problem.

The bottom row of Table 4 presents the thinking time spent on each specific phase. These values show that participants spend more time in some phases than others. For instance, on average, the participants spent 25.4% of thinking time in the *direction setting* phase, and just 3.3% on making a *decision*.

The off-diagonal values present the transitions between the different phases of problem-solving. For instance, the most common transition is from *direction setting* to *evaluation*, which represented 16.5% of all transitions. Sixty-six percent of all transitions are generated between directly neighboring phases (e.g. *FS* → *FA or DE* → *EV*), whereas longer jumps are less common. Second-order transitions such as *FA* → *EV* represent 17.4% of the total, and third- or higher-order transitions just 16.6%. These results help us replicate what one would expect from prior studies such as that of Mintzberg et al. (1976), who proposed a problem-solving model where transitions were complex and took place between *all* phases, not just adjacent ones.

4.2 Clustering

We input the participants' transition matrices into the partitioning around medoids (*pamk*) method. Each matrix has 49 variables: seven representing the thinking time spent on each phase, and 42 from the transitions between phases. Two clusters emerged from the *pamk* method: one comprising 20 participants and the other 28. In the Supplemental Materials we show robustness

---

evaluation phases, and 17.6% on the implementation and implementation evaluation phases. This participant, however, was a clear outlier: they spent 17.3% of the time on the framing, direction setting, and evaluation phases, and 63.8% on the implementation. The sample's median Mahalanobis distance to the average time spent on the problem-solving phases was 4.80 and the 75% percentile 6.96 (Mahalanobis, 1936). The removed protocol had a Mahalanobis distance of 21.94. On any measure of normality, the protocol was the least normal by a large margin (Rasmussen, 1988). The observation by the raters and the quantitative measures led us to remove the protocol from our sample.

checks on the clustering that indicate robust cluster assignment even upon removal of participants.

Table 4: Full-sample transition matrix

Full sample	**→ FS**	**→ FA**	**→ DS**	**→ EV**	**→ DE**	**→ IM**	**→ IE**
FS →		8.99	3.98	1.48	0.39	0.25	0.04
FA →	7.23		6.64	2.10	0.53	0.61	0.09
DS →	1.94	2.74		16.76	1.18	3.01	0.20
EV →	2.75	2.68	11.87		3.10	2.33	0.44
DE →	0.73	0.39	0.70	1.87		1.02	1.03
IM →	0.39	0.75	1.67	1.67	0.99		2.89
IE →	0.21	0.39	0.63	0.97	0.73	1.63	
% of thinking time	10.6	19.1	25.8	24.5	3.3	13.0	3.7

Each cluster represents an emerging pairing of the patterns of attention used by the participants in the study. Following Ocasio and Joseph (2018), we refer to these emergent patterns of attention as the participants' "problem-solving strategies." From now on, therefore, we no longer refer to clusters, but to problem-solving strategies.

4.3 Transition matrices per strategy

The first step to characterize a strategy is to understand its transition matrix. To do so, we generated the transition matrix for each strategy by averaging the transition matrices of the participants who followed it. We called the strategy followed by the first 28 participants the *problem-focused* strategy, and that followed by the other 20 the *solution-focused* strategy, for reasons outlined below.

Tables 5 and 6 present the average of the transition matrices of the participants who followed the *problem-focused* and *solution-focused* strategies, respectively, showing how they differ in terms of the attention they allocate to four of the seven phases of problem-solving. Adherents of the *problem-focused* strategy attend more to the *frame stating* and *frame assuming* phases; they spend longer on them, and transition to and from them more often too. Those adopting a *solution-focused* strategy, meanwhile, tend to favor the *implementation* and *implementation evaluation* phases.

Table 5: Problem-focused strategy transition matrix

Problem-focused	**→ FS**	**→ FA**	**→ DS**	**→ EV**	**→ DE**	**→ IM**	**→ IE**
FS →		11.61	4.70	1.81	0.47	0.34	0.07
FA →	10.21		7.77	2.21	0.36	0.16	0.07
DS →	2.12	3.04		18.09	1.24	1.60	0.08
EV →	3.93	3.42	11.71		3.63	0.99	0.16
DE →	0.57	0.58	0.83	1.82		0.59	0.16
IM →	0.23	0.43	0.96	1.21	0.08		1.11
IE →	0.20	0.13	0.20	0.19	0.39	0.54	
% of thinking time	14.2	24.0	25.3	25.5	3.5	6.3	1.3

The differences in the four *frame* and *implementation* phases provide the strongest differences between the two groups (p-value < 0.001 and |t-statistic| > 3.5, for the four comparisons, shown in the bottom rows of Table 6). Interestingly, both strategies allocate almost equal attention to *direction setting, evaluation,* and *decision*.

We call the first strategy *problem-focused* because it allocates more attention to the *framing stating* and *framing assuming* phases. Examples of these phases can be seen in Table 2. For these phases, the coding scheme asked raters to identify verbalized thoughts that focused on empathizing to assess the situation; developing hypotheses or assumptions to gain an understanding of the problem; or analyzing the problem by recalling the available information. The thoughts coded in these phases relate strongly to *problem-focused* behavior.

In contrast, the *solution-focused* strategy allocated more attention to the *implementation* and *implementation evaluation* phases. For these two phases, the coding scheme asked raters to select passages where the participants designed sequences of actions that could unfold during the solution of the problem; anticipated how events would play out; or evaluated the feasibility of their solutions. These were situations where the participant was strongly *solution-focused*.

Figure 3 shows visualizations of the transition matrices presented in Tables 5 and 6. Here, in contrast to the simple, linear problem-solving process depicted in Figure 2, the transition processes are much more complex. We see that the strategies use their time unevenly: the circle's diameter is proportional to the time each strategy spends on each phase. Additionally, since participants can transition between any two phases, *any* two circles may be linked (not just those that are adjacent in the linear model). Figure 3 shows that some transitions are more common than others, as the thickness of the lines are proportional to their use by the strategies. Figure 3 is simplified in one respect: the lines represent the sum of the transitions between two phases in *both* directions—for example, *FS* → *DE* + *DE* → *FS*. To reflect this, we here replace the directional arrows of Figure 2 with simple lines.

Figure 3 shows that the strategies are most strongly differentiated by the phases they attend to the most—the *focus of attention*—and the number of transitions to and from this focus of attention. These differences result in the two contrasting attention patterns that we see emerge from the data. The *problem-focused* strategy took the *frame stating* and *frame assuming* phases as its *focus of attention*, whereas the *solution-focused* strategy focused its attention on the *implementation* and *implementation evaluation* phases.

Table 6: Solution-focused strategy transition matrix

Solution-focused		**→ FS**	**→ FA**	**→ DS**	**→ EV**	**→ DE**	**→ IM**	**→ IE**
FS	→		5.32	2.97	1.02	0.26	0.13	0.00
FA	→	3.06		5.06	1.95	0.78	1.25	0.13
DS	→	1.69	2.31		14.88	1.11	4.98	0.35
EV	→	1.10	1.64	12.10		2.37	4.21	0.84
DE	→	0.94	0.13	0.51	1.94		1.63	2.23
IM	→	0.62	1.19	2.66	2.32	2.27		5.37
IE	→	0.23	0.75	1.23	2.08	1.22	3.15	
% of thinking time		5.7	12.2	26.4	23.2	3.1	22.4	6.9
t-test	**p-value**	0.000	0.001	0.786	0.553	0.681	0.000	0.001
	t-statistic	-4.183	-3.518	0.273	-0.598	-0.414	6.509	3.881

Note: A positive t-value implies that the *solution-focused* strategy had a higher value than the *problem-focused* strategy.

Figure 3: Visualizations of the problem- and solution-focused strategy transition matrices

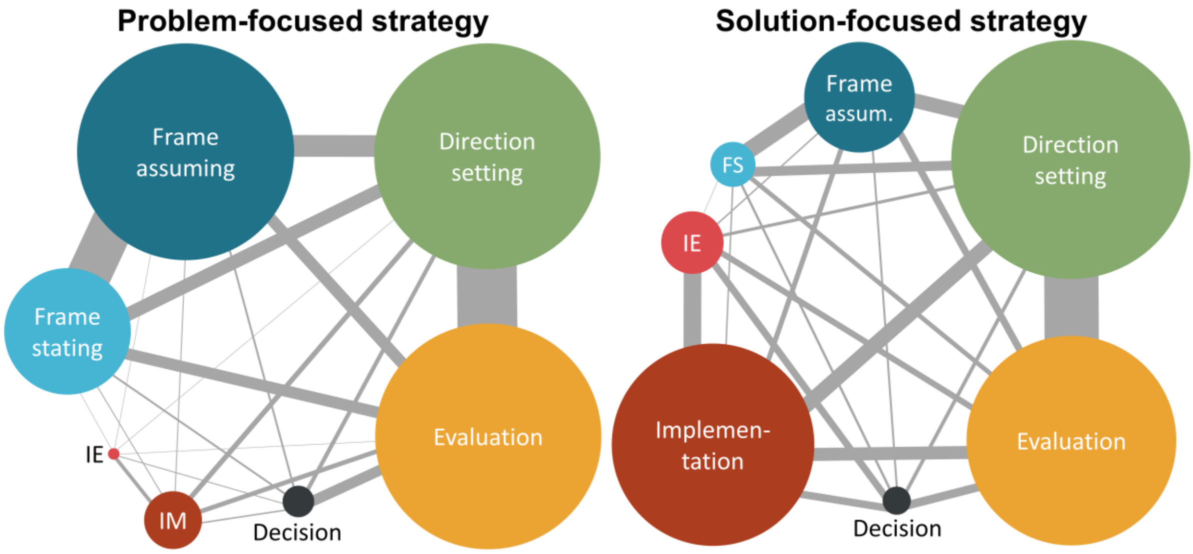

In Figure 1, we showed the problem-solving sequence of two sample participants who followed each strategy. Person A followed the *problem-focused* strategy, while Person B followed the *solution-focused* strategy. From Figure 1, one can observe how the differences in the transition matrices emerge, as people similar to Person A direct their attention towards *frame stating* and *frame assuming*. In contrast, people similar to Person B attend more often to *implementation* and *implementation evaluation*.

The ninth rows in Tables 5 and 6 show the proportion of total thinking time spent in the different phases of problem-solving by each group. We calculated that the *solution-focused* strategy group devoted 3.5 times more attention to *implementation* and 5.3 times more attention to *implementation evaluation* than the *problem-focused* strategy group.

There are two possible reasons for the difference between the strategies in these two phases. Either the *solution-focused* strategy transitions into these phases just as often as the *problem-focused* strategy and then spends more time in them, or it transitions more often into these phases but spends a similar period there each time. By counting the number of instances of participants transitioning into the *implementation* phases, we corroborated the latter explanation. We found that the *solution-focused* strategy transitions 3.7 times more often into *implementation,* and 6.2 times more frequently into *implementation evaluation,* than the *problem-focused* strategy does. Thus, what differentiates the two strategies is not that the periods of attention last longer, but that the relevant phases are attended to more frequently—that is, the two strategies pattern attention differently by transitioning more or less between phases.

A similar analysis shows that the *problem-focused* strategy attends 2.5 times longer to *frame stating* and twice as long to *frame assuming* as the *solution-focused strategy*. Conducting a deeper analysis, we observe the same reason as before, only reversed: the *problem-focused* strategy transitions twice as often into *frame stating* and 1.7 times more often into *frame assuming* than the *solution-focused strategy*. In this case, per occasion, those following the *problem-focused* strategy spent around 25% less time every time they attended to the *framing* phases than *solution-focused* participants did, but as they made the transition much more often, the cumulative attention they spent was greater.

In the remainder of this paper, we will refer to allocating more attention to the *framing* phases as the *problem-focused* strategy, and to allocating more attention to the *implementation* phases the *solution-focused* strategy. We do this because of the finding that the amount of time spent on each phase is proportional to the number of transitions to and from the phases.

4.4 Alternative explanations

Table 7 contains the descriptive statistics and zero-order correlations between the *strategy* categorical variable (0 for the *problem-focused* strategy and 1 for the *solution-focused* strategy), *protocol duration*, cognitive skills, and demographic characteristics of our participants. We find that the strategy followed by our participants is not significantly correlated to most variables. Interestingly, the *solution-focused* strategy is positively correlated to *performance*: participants who used this strategy performed about 14% (t-test p-value = 0.003, t-statistic = 3.095) better than those who followed the *problem-focused* strategy. Similarly, *protocol duration* was correlated to *performance*. However, as *protocol duration* and *strategy* are uncorrelated, each might provide a separate avenue for higher *performance*.

Table 7: Descriptive statistics for the "Karabayos" problem

	1.	**2.**	**3.**	**4.**	**5.**	**6.**	**7.**	**8.**
1. Performance	1							
2. Solution-focused strategy	0.404 (0.004)	1						
3. Protocol duration (min.)	0.464 (0.001)	0.120 (0.416)	1					
4. Planning & generativity (s)	0.223 (0.146)	0.268 (0.078)	0.091 (0.559)	1				
5. Abstract thinking	0.245 (0.104)	0.016 (0.919)	0.174 (0.253)	-0.126 (0.417)	1			
6. Age (years)	0.059 (0.692)	0.093 (0.529)	-0.192 (0.19)	-0.167 (0.278)	-0.018 (0.905)	1		
7. Gender: Female	-0.102 (0.491)	0.027 (0.855)	-0.174 (0.236)	-0.213 (0.164)	-0.174 (0.254)	-0.255 (0.081)	1	
8. Profession: Entrepreneur	0.002 (0.989)	0.068 (0.644)	0.256 (0.079)	0.106 (0.493)	0.069 (0.654)	-0.126 (0.394)	-0.094 (0.527)	1
M	6.13	0.417	12.57	285.6	7.444	35.46	9 of 48	15 of 48
SD	1.02	0.498	9.65	171.5	1.617	6.81		

Note: p-value of pairwise correlations shown in parenthesis

Additionally, we found a marginal mean difference (t-test p-value = 0.088, t-statistic = 1.756) between the *planning and generativity* skills of the participants of the two groups. Specifically, participants who followed the *problem-focused* strategy tended to be marginally better at *planning and generativity*. This difference could be due to the fact that the *problem-focused* strategy spends longer attending to the *framing* of the problem—a key skill within the task we used to measure *planning and generativity,* the "Tower of Hanoi" (Laureiro-Martínez, 2014).

4.5 Transition to Study 2

In Study 1, we find that managers solve strategic problems by following one of two alternative strategies. Note that this result is not obvious. We could have found that there were too few similarities among the strategies to cluster them together—or, alternatively, that there were as many different strategies as participants in the sample. We could have also found that there was a single, dominant process that described the strategies developed by most participants. Instead, we found two patterns that describe strategies that have enough commonalities within a cluster, but enough differences to fall into two clearly differentiated clusters.

These strategies appear to differ in terms of the amount of attention spent on cognitive processes related to the *framing* or *implementation* phases of problem-solving. However, beyond the descriptive finding, we cannot make a causal claim, as we do not know whether the allocation of attention to different cognitive processes is what *causes* such strategies to emerge. In other words, with Study 1, we are able to describe the emergence of two different strategies under a controlled environment. With Study 2, we aim to explain the cause of the different strategies by manipulating the cognitive processes related to different problem-solving phases.

Specifically, in Study 2, we manipulate the attention participants pay to the *framing* or *implementation* phases of problem-solving, and compare their behavioral changes to a control condition. To estimate the behavioral changes, we conduct a study in which each participant solves two problems, to compare how participants behave before and after the corresponding manipulation.

Three outcomes are possible from this experimental study. First, we might find that we cannot manipulate the allocation of attention, and thus there is no behavioral change between the two treatment conditions and the control condition. Second, we might find that the two manipulations do change the participants' behavior, but that the behavioral change is the same or indistinguishable in both conditions, thus failing to illuminate the cause of the two different strategies. Finally, we could find that each manipulation of attention affects each condition differently. This outcome, in turn, can be manifested with clearly differentiated strategies that correspond to what we can expect from the theorizing derived from the findings obtained on Study 1, or refute those expectations. Below, we develop our expectations from the results of Study 1.

We study strategic problems: those that are complex, ill structured, and novel, involving high-stakes, irreversible outcomes. Study 1 suggests that the process that participants adopt corresponds to the way their attention is allocated. If so, requiring participants to focus their attention on either the framing of the problem or the implementation of the solution should translate into changes in their behavior.

Participants who pay more attention to the framing of the problem will notice that the problem is new to them and hard to comprehend due to its lack of structure. They will put more effort into understanding the various elements of the problem and their relations. While this deliberative effort is devoted to better framing the problem and its structure, multiple thoughts will appear, aimed at connecting the elements of the problem; creating and revising a hierarchy of goals and priorities; and, according to that evolving framework, thinking about the structure of the problem (Baer et al., 2013). The problem-solving process might thus require a higher number of thoughts than the process of a participant in a control condition. Therefore, we can propose that:

Proposition 1: Increased attention to the cognitive processes related to the framing of the problem will lead to a *problem-focused* strategy.

In contrast, participants who are asked to pay more attention to the implementation of the solution will focus on thinking about what is at stake; reflecting on how to minimize potential negative outcomes; and developing detailed thinking about possible contingencies and the future consequences of potential solutions (Schacter, Benoit, & Szpunar, 2017). By focusing more on the implementation of the solution to the problem, the participant might end up devoting more time and deliberation to each thought than a participant in a control condition. A solution-focused strategy will be associated with every thought requiring more time to be performed. Therefore, we can propose that:

Proposition 2: Increased attention to the cognitive processes related to the implementation of the solution will lead to a *solution-focused* strategy.

Testing these two propositions requires a method where we can track the thinking process of problem-solving and decompose it into multiple thoughts. In addition, we need to do this twice: once before a manipulation takes place, and once afterwards. While think-aloud protocols would still be a very useful technique, the data analyses would prove very expensive—not just because the analyses would have to be done one by one, but also because they would have to be performed twice per participant. For a sample size that would support a three-by-two mixed factorial design, such as the one we will use in Study 2, we would have to collect and analyze over 1000 protocols. Rank-based tasks combined with mouse-move analyses provide an excellent option for our needs, as they allow us to present a strategic problem and track the thinking of the participants as it unfolds in real time, using the computer to measure the movements of the mouse (Freeman, 2018). While the thoughts are not verbalized in this case, mouse moves are used as a proxy for thoughts (see Data Analysis section). In the next section, we present the two problems we employ to investigate the effect of manipulating the allocation of attention on strategic problem-solving.

5. STUDY 2

In the previous study, we uncovered two types of strategies that managers employ when solving strategic problems. In this section, we present the methodology we used to investigate the causal antecedent of these strategies. We first introduce the tasks used, continue with data collection, and finally present the data analysis and results. More details on all tasks can be found in the Supplemental Materials.

5.1 Strategic problems: "NASA survival" and "winter survival"

The "winter survival" problem by Johnson and Johnson (1982: 111) and the "NASA survival" problem by Hall and Watson (1970) are two tasks that allow us to observe the problem-solving processes of the participants who perform them. Additionally, the two tasks are commonly used in research and management education (Baker & Paulson, 1995; Joshi et al., 2005; Lane et al., 1982; Yetton & Bottger, 1983). Both tasks require participants to think as a leader who must make decisions for a group they were responsible for. The "winter survival" problem is set on a winter's day in the mountains of Manitoba in Canada. Minutes ago, a plane carrying the leader and their group has crashed into a lake. The survivors have collected a list of 12 items that the leader must rank based on their importance to the group's survival. The "NASA survival" problem has a very similar structure, but it is set on the moon. The participant is asked to imagine that the lunar module carrying the crew was forced to land 300 kilometers away from its destination, a meeting point. Now, in order to survive, the participant must rank 15 items in terms of their importance in allowing the crew to reach the destination.

Although both problems initially appear to be far removed from a managerial setting, they actually fulfill the requirements for being considered strategic problems. They are *complex*, given the number of interrelated items that the participant needs to rank in order. They are *ill structured*, as there is no clear information about the exact means for achieving the goals. The list of items is not well related to survival; some items are of very little use, and it might even be better to leave them behind. The role of external agents is uncertain; it is not clear whether or not help is coming, or if the group is alone. Both problems are *high-stakes* and also *irreversible*. After the leader has finalized a solution and moves on to implement it (e.g. the group sets off on foot, or starts to build a fire), every choice will have a cost that cannot be recovered. An item ranked too low might be left behind and create problems along the way. At their core, both problems share many commonalities with very difficult situations a manager or an entrepreneur might face when trying to ensure the survival of their business unit or small firm in the face of multiple constraints and limited resources. However, in these problems, the contexts are *novel* to the participants, which prevents them from directly taking past experience into account.

In the computer interface that participants use to tackle the problems, they are asked to rank a list of items by dragging them from a column on the left to a column on the right. Participants can and do reposition items on the right while they think through their solution. This interface allows us to explore the problem-solving process participants follow, in the form of drag-and-drop events and the time it takes participants to carry out the movements (Fedor et al., 2015)—not just their overall reaction times and solutions (Yu et al., 2012). We can study the time they take to make each move, and how that move comes about. While we cannot record the thinking processes directly, the events we can observe provide a proxy for the problem-solving process of the participants (Öllinger et al., 2013).

By using the “NASA survival” and “winter survival” problems, we can operationalize our expectations on the increase of number of thoughts or their duration in specific ways. “Today it is relatively uncontroversial that thinking can be represented as a sequence of thoughts (relatively stable cognitive states) interspersed by periods of processing activity” (Ericsson & Simon, 1998: 180). We capture chunks of thought as any drag-and-drop move that a participant performs before arriving at their finished solution. We count the thoughts and measure their duration as the time between the preceding move (or the start of the test, for the first move) and the current one.

5.2 Hypotheses

The two tasks in the experiment require participants to rank-order items. The rank ordering process gives us a proxy for the concurrent problem-solving process of the participant. We operationalize this problem-solving process with three variables: the *total time* allocated to each task, the *number of moves* performed, and the *time per move* that the participant took after the first move. Note that the first and third variables are not directly related, as participants need to read the task before the first move is done.

A *problem-focused* strategy will involve an effort to connect a multiplicity of problem elements as the participant tries to give the problem a structure and define goals and priorities. A *problem-focused* strategy will be associated with a higher number of thoughts, measured as the number of drag-and-drop moves. With this, we operationalize Proposition 1 as two sub-hypotheses:

Hypothesis 1a: Increased attention to the *framing of the problem* will lead to an increased deliberative effort, operationalized as a higher *total time* allocated to the task when compared to the control condition.

Hypothesis 1b: Increased attention to the *framing of the problem* will lead to an increased effort in structuring problem elements, operationalized as a higher *number of moves* when compared to the control condition.

A *solution-focused* strategy will involve an effort to reflect on developing and maturing the possible solution. A *solution-focused* strategy will be associated with a greater depth to each thought, measured as the time between moves. With this, we operationalize Proposition 2 as two sub-hypotheses:

Hypothesis 2a: Increased attention to the implementation of the solution will lead to an increased deliberative effort, operationalized as a higher total time allocated to the task when compared to the control condition.

Hypothesis 2b: Increased attention to the implementation of the solution will lead to an increased effort in developing possible outcomes, operationalized as a higher time per move when compared to the control condition.

5.3 Data collection: Online experiment

We performed an online experiment that studied the behavioral changes that develop as a consequence of manipulating the allocation of attention towards either the framing of the problem (*framing-focused*) or the implementation of the solution (*implementation-focused*), or allowing the task to unfold without intervention (*control* condition). By asking participants in different treatment conditions to focus on the problem or the solution, we can compare how their behavior changes in comparison to a control condition and infer why the two strategies exist in the first place.

We ran four pretests and two pilot studies before the online experiment took place, and one small scale follow-up study after the online experiment took place. The main benefit of these, besides refining the problems and the computer interface, was the debriefing interviews, which provided qualitative evidence about the problem-solving processes, complementing the quantitative measures obtained from the experiment. Each of these studies is described in the Supplemental Materials.

5.3.1 Research design

We perform a three-conditions-by-two-tasks mixed factorial design experiment (Oehlert, 2000). The experimental procedure began with all participants performing the first task (the "winter survival" problem) without being manipulated. After this, the participants were split into three groups: one control group and two treatment groups (*control*, *framing-focused*, and *implementation-focused*). The two treatment groups were presented with manipulations that aimed at increasing participants' allocation of attention towards either the framing of the problem or the implementation of the solution.

After the manipulation, all participants performed the second task (the "NASA survival" problem). The mixed factorial design allows us to study the behavioral change of the participants as an effect of the treatment. In comparison to a between-subject design, we can use the participants' measures before the manipulation as a baseline for the treatment effect, thus reducing variation in the analyses.

In contrast to a within-subject design, not all participants are exposed to every treatment, allowing us to estimate the average treatment effects for each condition without having task-order bias. The order effect does not bias our estimation as the task order affects all participants in the same way; when we compare between conditions, the order effect cancels outs and allows us to have an unbiased estimator. The lack of order-effect bias does not mean that the task order has no effect on the participant's responses, but rather that the effect is canceled out by the comparison between conditions. By employing a mixed-factorial design, we are able both to use within-subject design (to use participants' responses as their own baselines) and also to compare between participants (to estimate treatment effects without having to counterbalance the order of the tasks).

5.3.2 Manipulations: Framing-focused and implementation-focused

The manipulation was shown between the "NASA survival" and "winter survival" problems, so we could compare the groups before any change and study the behavioral change of every participant after the manipulation—for example, whether they spent longer on the task, or performed more moves.

We asked participants to direct their attention to engage in more of the cognitive processes related to the different phases of problem-solving that were characteristic of each type of strategy. The manipulations recommended that participants "spend more of their time thinking about" either "the framing of the problem" or "the implementation of the solution." In the framing-focused manipulation, we explained to participants that the framing of the problem involves the following mental activities:

- Analyzing the problem by recalling the available information
- Empathizing to identify with the situation
- Developing hypotheses/assumptions to gain an understanding of the problem.

In the implementation-focused manipulation, we explained to participants that implementing a solution involves the following mental activities:

- Designing the sequence of actions that could unfold during the solution of the problem
- Anticipating how events will play out
- Evaluating the feasibility of the solutions.

Finally, the control condition was asked to solve the problem in whatever way felt natural to them. We took the mental activities that we asked participants to follow directly from the coding scheme of the "Karabayos" task, in order to minimize the over-interpretation of our findings. We include more detail on the manipulation and research procedure in the Supplemental Materials.

5.3.3 Sample

We conducted the behavioral experiment through the platform Prolific, as a way of recruiting participants to our study. Prolific is a dedicated research-subject pool and recruiting platform, employed in multiple studies in recent years. For comparisons between Prolific and other online participant recruitment platforms, see Palan and Schitter (2018) and Peer et al. (2017).

In our study, we allowed participants with a broader background than just managers, but we did filter for three participant attributes in order to obtain homogenous behavior and a comparable sample. First, we required participants to have a common minimum education level (i.e. at least a bachelor's degree). Secondly, we required English to be their first language, to ensure they could comprehend the instructions. Third, we selected participants aged 55 and under, as the task required interaction in a drag-and-drop setting and a younger age could be correlated with more familiarity with computer interfaces.

Our initial selection comprised 523 participants. We excluded 51 participants who had experience in survival training because we wanted to replicate the conditions of the "Karabayos" problem, where participants had neither experience of leading Amazonian tribes, nor access to feedback. The experiment ultimately included 472 participants: 276 in the control condition, 97 in the framing-focused condition, and 99 in the implementation-focused condition. We employed the G*Power 3.1 application to specify the experiment's sample size (Faul, Erdfelder, Buchner, & Lang, 2009). The experiment has three conditions and two measures. A between-factors repeated measures ANOVA specifies 294 participants (98 per condition) to achieve both alpha (probability of rejecting a true null hypothesis) and beta (probability of rejecting a false null hypothesis) of 0.95, and a small effect size of 0.2. We did not have a strong stopping rule, as we needed to determine whether the participants lacked survival training, which prevented us from having exactly 98 participants in the two manipulated conditions. The control condition was larger than the other two because we wanted to validate findings from prior piloting sessions. Overall, the experiment has enough statistical power even if one of the samples is larger than the other.

5.3.4 Incentives

In order to elicit strong commitment to the tasks, we created a three-level incentive scheme. The base payment rate in Prolific is £5 (British pounds sterling; GBP) per hour. Our study took on average 30 minutes in total, for a total base payment of £2.50. The top 25% of performers on both tasks received twice the hourly payoff of the platform (£5), the middle 50% received 1.5 times the hourly rate (£3.75), and the bottom 25% received the hourly rate (£2.50). As participants self-selected to take part in the online platform, doubling the payoff for high performance is deemed an attractive way to increase the saliency of the task (Hertwig & Ortmann, 2001).

5.4 Data analysis

As in the “Karabayos” task, we focused on the processes that participants employed to solve the problem and find a solution. To uncover participants’ thought sequences, we employed move analyses—a process-tracing method. In the past decade, tracking mouse movements has become a popular method in psychological science (Freeman, 2018). Move analyses focus on isolating the moves as a way to proxy for the thoughts that unfold while solving a problem (Yu et al., 2012; Öllinger et al., 2013; Fedor et al., 2015). Although reaction-time measures and other temporally sensitive methodologies (for example, EEG) shed light on the cognitive and neural processing of a response, mouse tracking offers a "more direct measure of the evolution of a particular response" (Freeman, 2018). For example, to investigate the stages of the insight problem-solving process in a five-square problem, Fedor et al. (2015) recorded the drag-and-drop movements of the mouse, along with keyboard strokes. They found support for the notion of mouse movements as a behavioral proxy for ongoing cognitive processing. Further, and vital for the purposes of our study, analyzing the sequence of mouse movements has been proven as a temporally fine-grained measure that even reveals participants' tentative commitment to alternative choice responses (Freeman & Ambady, 2009). For instance, Yu et al. (2012) conducted an analysis of mouse trajectories to explore the underlying processes of the implicit association test, and the results demonstrated continual movements towards alternate responses before choosing the correct one. Analyzing movement trajectories of the mouse, as in a simple drag-and-drop movement, allows us to reveal the microstructure of real-time decisions by looking more closely into ongoing cognitive processing. In addition, in this study's analyses, we compared the changes in behavior that resulted in the three different conditions.

5.4.1 Dependent variables

We measured three main variables that allowed us to infer differences in the way each participant thought while solving the problem. We based our measurements on studies of move analyses that aim at understanding the deliberation that takes place during real-time problem-solving by measuring mouse clicks, and drag-and-drop moves of elements while solving a problem. The assumption is that the moves represent steps involved in the deliberation process (Yu et al., 2012; Öllinger et al., 2013; Fedor et al., 2015; Ormerod et al., 2002).

The first measure is the *total time* spent. This variable includes both the time spent reading the task and the time taken to move the items to create the final ranking. Therefore, this measure reflects the total effort and attention the participants put into the task.

Second, we measure the *number of moves* each participant performs. There is a minimum number of moves the participant can make, imposed by the number of items that must be ranked. For the "NASA survival" problem, this lower bound is 15 moves, and for "winter survival" it is 12 moves. Any additional move above this threshold represents the refinement or correction of a previous idea. Since the lower bound is the same for everyone, the total number of moves can be used as a proxy for the number of thoughts a participant engaged in during the problem-solving process.

Third, we calculate the time between the first and last moves and divide it by the number of moves. The *time per move* is a measurement of the amount of deliberation involved in each thought. Some moves will involve more thought than others, and some processes will involve more or fewer moves. Putting the three variables together, we can explore the behavioral changes that arise from manipulating the *focus of attention* in strategic problem-solving.

We study these three variables because they provide us with a proxy for the participants' problem-solving process. By comparing their values before and after the manipulation, and how the changes compare to the control condition, we can find an answer to the research question of Study 2, namely: *Why are there two strategies for solving strategic problems?*

Finally, both tasks include an optimal ranking created by an expert (Hall & Watson, 1970; Johnson & Johnson, 1982). We estimate the performance of the solutions by calculating the distance of the participant's ranking from the optimum. For example, if the participant placed item A in the first position, and it is supposed to be in position 7, we add a distance of 6 to the first item. We sum the distance of all the other items in the task to calculate the participant's *performance.*

5.4.2 Control measures

For each participant, we use demographic variables as control variables—specifically, their *age*, *gender* (1 for female, 0 for male), *postgraduate* education (1 if they have a master's degree or above, 0 if not), and whether they read more than twice a week, and are thus classified as a *reader* (1 if they do, 0 otherwise). We use these variables as control measures for the behavioral and performance metrics. Overall, the average age was 34.9 years (s.d. = 8.5 years with a range from 21 to 56 years); 269 of the 478 participants were female; 161 had a postgraduate degree; and 219 read more than twice a week.

5.4.3 Behavioral change

We estimated the behavioral and performance change of every participant due to the manipulation. As the two tasks have different numbers of items to rank, we could not directly compare the performance of behavioral variables, so we standardized our variables to study the changes in behavior between the two tasks. Specifically, we calculated the distance in units of standard deviation from the mean of the control condition after the manipulation, and subtracted the distance before. This is calculated using the following formula[5]:

$$Change\ Value(i) = \frac{Value_{NASA}(i) - \ Mean_{NASA}(control)}{Std.Dev._{NASA}\ (control)} - \frac{Value_{Winter}(i) - \ Mean_{Winter}(All)}{Std.Dev._{Winter}\ (All)}$$

For example, if a participant (i) took the average time to finish the "winter survival" problem, and, after the manipulation, their time in the "NASA survival" problem was 0.33 standard deviations higher than the average, we stored a change of +0.33 standard deviations. This analysis allows us to study in greater detail the behavioral changes that happen to every individual, and not just the entire group.

In the results section below, we present only changed variables, namely *total time*, *time per move*, *number of moves*, and *performance.* The control variables do not change and are presented in simple form.

[5] We use the values from the control condition only to calculate the mean and standard deviation in the "NASA survival" exercise. We do this filtering to avoid diluting the effect of the manipulation through increased standard deviations or changed means. Therefore the mean and standard deviation used for standardization come from untreated participants.

6. STUDY 2: RESULTS

Table 8 presents the descriptive statistics of the four main variables of Study 2; a descriptive and first-order correlation table is included as Table A.4 in the Supplemental Materials. From Table 8, we can observe that, on average, participants performed 66% more moves than the minimum number required for each task. Increases in the number of moves used allow us to explore the differences in the behavior of the participants. First, however, we present a short example of how the measure of behavioral change is calculated. From the results in Table 8, a participant who performed 20 moves in the "winter survival" problem and 33 in the "NASA survival" problem is stored as a behavioral change of 0.33 standard deviations. This is the equivalent of the example from the methods section.

Table 9 presents four ordinary least square regressions; we obtained the same results with robust regressions. The robust regression results are shown in Table A.5 of the Supplemental Materials. Each regression uses the same covariates, but focuses on a different dependent variable. Model 1 presents the change in *total time*. Both treatment conditions have beta values whose 95% confidence intervals are positive. That is, in both cases, the treatment led to participants spending more time on the task. This provides support to Hypotheses 1a and 2a.

Table 8: Descriptive statistics for Study 2

Task		Total time	Time per move	# of moves	Performance
"Winter survival" problem	M	324.5	10.17	20.05	45.13
	SD	287.2	11.37	7.37	9.54
"NASA survival" problem	M	363.0	8.85	28.92	49.34
	SD	208.8	5.92	12.31	15.48

Model 1 showed that the participants of both conditions spent more time overall on the problem after the manipulation. However, they used the time in different ways. Interestingly, when analyzing the time spent per move, Model 2 shows that the participants who were asked to focus on the solution spent longer *time per move*, whereas the framing-focused condition spent a similar amount of *time per move* as the control condition, thus supporting Hypothesis

1b. Interestingly too, we found the opposite when analyzing the number of moves in Model 3. Framing-focused participants increased the *number of moves* when compared to the control condition. In contrast, the implementation-focused participants performed a similar number of moves as the control condition. Jointly, these results give support to Hypothesis 2b. Finally, Model 4 focuses on the change in *performance* after the manipulation; focusing on the problem or the solution did not affect *performance*. However, behavioral changes were present.

Table 9: Study 2 OLS Regressions of behavioral change and performance

	Dependent variable:			
	Total time (1)	**Time per move** (2)	**# of moves** (3)	**Performance** (4)
Framing condition	0.274 (0.042, 0.506)	0.080 (-0.172, 0.332)	0.310 (0.061, 0.558)	-0.084 (-0.379, 0.210)
Implementation condition	0.276 (0.046, 0.505)	0.343 (0.094, 0.592)	0.031 (-0.215, 0.277)	-0.074 (-0.365, 0.218)
Gender	-0.221 (-0.405, -0.036)	-0.146 (-0.346, 0.054)	-0.046 (-0.243, 0.152)	-0.030 (-0.264, 0.204)
Age	0.003 (-0.008, 0.014)	0.002 (-0.009, 0.014)	-0.008 (-0.020, 0.003)	0.013 (-0.0005, 0.027)
Postgraduate	-0.251 (-0.447, -0.055)	-0.123 (-0.335, 0.089)	-0.069 (-0.279, 0.140)	0.096 (-0.152, 0.344)
Reader	-0.037 (-0.227, 0.153)	-0.164 (-0.370, 0.043)	0.223 (0.019, 0.427)	-0.107 (-0.348, 0.134)
Constant	0.049 (-0.358, 0.457)	0.059 (-0.383, 0.502)	0.199 (-0.238, 0.636)	-0.433 (-0.950, 0.084)
Observations	472	472	472	472
R2	0.043	0.029	0.023	0.010
Adjusted R2	0.030	0.017	0.010	-0.003
Residual Std. Error	0.999 (df = 465)	1.084 (df = 465)	1.071 (df = 465)	1.267 (df = 465)
F Statistic (df = 6; 465)	3.461 (p = 0.002)	2.317 (p = 0.032)	1.832 (p = 0.091)	0.774 (p = 0.591)
Note:		95% confidence intervals shown in parenthesis		

We employed G*Power 3.1 to estimate the ex-post effect sizes for the findings of this study (Faul, Erdfelder, Buchner, & Lang, 2009). The effect size for the *total time* is 0.136, for the *time per move* 0.134, and for the number of moves 0.113. For the three cases, the effect size is small, with an average value of 0.127.

Interestingly, while both manipulations led to an overall increase in deliberative effort (total time employed by the participants), each type of focus led to this additional deliberation being employed in two very different ways. Participants asked to focus on the framing of the problem spent their time engaging in more thoughts, represented by 3.81 more moves than the *control* and *implementation-focused* conditions. Each move was preceded with the same amount of deliberation as the moves of the control condition. In contrast, the participants who were asked to focus on the implementation of the solution conducted about 20% more deliberation before every move, but their total number of moves was similar to those in the control condition. Each thought took longer, but no additional thoughts were needed to solve the problem.

Combining the moves measures with debriefing interviews, we can infer that participants in the *framing-focused* condition restructured the way they defined the problem and adjusted their priorities more often than participants in the other conditions, probably due to a constant updating of their definition of the problem. An increase in the focus on the solution, meanwhile, led participants to perform the same number of thoughts as the control condition, but each thought involved more deliberation than in the other conditions. This might be because they delved deeper into their thoughts about how their solutions might unfold into the future.

7. STUDY 3: REPLICATION

To test the reliability of our claims, we replicated Study 2 using the same tasks, incentive scheme, and measures. Furthermore, we ran the study on the same platform as Study 2, and participants were sampled with the same criteria. The study was pre-registered and is publicly accessible at: https://osf.io/nvfdc

7.1 Changes implemented in the replication

Study 3 deviates from Study 2 in two respects. First, it has a larger sample size, and second, it includes control scales related to coping mechanisms due to the ongoing pandemic at the time of data collection (Taylor, 2019).

7.1.1 Sample size estimation

We employ a larger sample size as the average ex-post effect size of the three main results of Study 2 was 0.127. Study 2 was designed for effect sizes of 0.2 and, given the smaller effects, a larger sample was needed (Baguley, 2004). To calculate the new sample size, we maintain the specification of 5% commission and omission errors for the three-condition-by-two-task experimental design. We employ the G*Power 3.1 application and use a 0.125 effect size to determine that the sample size required is 747 participants: 249 per condition (Faul, Erdfelder, Buchner, & Lang, 2009)[6].

7.1.2 Pandemic control variables

This study investigates how participants solve novel, complex, and ill-structured tasks in the absence of feedback. The tasks in Study 3 comply with the last three of these characteristics. The tasks are related to survival, either on the moon or in freezing, mountainous terrain (Hall & Watson, 1970; Johnson & Johnson, 1982). To guarantee novelty, in Study 2 we excluded people with survival training from the dataset. We provided these control questions as a follow-up survey. The participants were not informed of the second survey when they performed the main experimental task[7]. We did this to maintain the same expectations from Study 2 and avoid introducing variance to the replication. Almost all participants answered the controls (92.8%).

[6] The ex-post power estimations are not expected to give a completely accurate effect size, but are to be used as guides for updating the future research design. For Study 3, we use an effect size of 0.125 because it is close to the average effect of Study 2 and allows us to design an experiment that is strong enough to test the hypotheses.

[7] We did not include the pandemic control variables in the same survey as the main study because we learned through the pandemic pilot study (detailed in the Supplemental Materials) that participants allocate their time and

We retained these exclusion criteria in Study 3. However, we collected the data in May 2020, at a time when media coverage was dominated by the global coronavirus crisis. The increased coverage of disease in the media could affect the preparedness of the participants and confound our results. To account for this, we include six separate control scales. We compiled these scales from work in psychology to account for the crisis-coping styles of the participants, as well as how fearful they were of COVID-19, and how they had been financially affected by the crisis (Taylor, 2019). We employed four validated scales to estimate the participant's *fear of COVID-19* (Ahorsu et al., 2020), *monitoring/blunting coping styles* (Steptoe, 1989), *perceived vulnerability to disease* (Duncan et al., 2009), and *intolerance to uncertainty* (Carleton et al., 2007).

Additionally, we asked participants if either they or someone in their household had lost their job or a significant part of their income due to the COVID-19 crisis (*job loss*). More than a third of participants, 33.9%, reported that they either they (20.3%) or someone in their household (23.0%) lost their job or a significant part of their income during the COVID-19 crisis. Finally, we asked the participants to estimate the importance that their work doing online experiments at Prolific.ac has for them now compared to before the crisis, both in general and for their personal income (*Prolific importance*). The median participant response stated that they found Prolific equally as important now as before the COVID-19 crisis. However, 11% of participants stated that they found Prolific more or much more important than before. In contrast, less than 0.3% of the participants said they found Prolific less important than before the COVID-19 crisis.

effort based upon the time advertised to complete the survey. When we put the main study and the pandemic controls together, participants adapted their time use and spent longer in the performance-related tasks ("NASA survival" and "winter survival" problems). The increased time changes the saliency and behavior of the participants, making for a bad replication of the study (Hertwig & Ortmann, 2001). By separating the main study and the pandemic controls, we avoid this problem.

8. RESULTS

The results of Study 3 mostly support the findings of Study 2. The increase in *time per move* for the participants in the *implementation-focused* condition is only marginally significant when we employ the control variables of Study 2, but the use of the pandemic controls provides support to the increase with 95% confidence. Below, we present the results in more detail. We do this in three steps. First we compare the two studies in broad terms to confirm that Study 3 does not differ from Study 2 in any major way. Second, we show the differences we find in terms of hypothesis testing. Third, we summarize the hypothesis testing.

8.1 Comparison of Studies 2 and 3

Table 10 presents the descriptive statistics of the four main variables of Study 3; a descriptive and first-order correlation table is included as Table A.10 in the Supplemental Materials. The descriptive statistics of Study 3 closely match those of Study 2, which it replicates. Table 11 shows the t-tests of comparing each of the variables of interest of Study 2 and Study 3—that is, the *total time, time per move, number of moves,* and *performance* in both tasks. For these t-tests, we put together three conditions of each study and compare them to the other study. In doing this we compare the 472 participants of Study 2 with the 747 participants of Study 3.

Table 10: Descriptive statistics of main variables for Study 3

Task		Total time	Time per move	# of moves	Performance
"Winter survival" problem	M	337.2	10.20	20.90	44.2
	SD	233.7	8.25	7.04	11.98
"NASA survival" problem	M	383.8	9.55	28.19	48.82
	SD	260.7	6.98	10.38	15.76

Table 11: Comparison of main variables of Study 2 and Study 3

Task		Total time	Time per move	# of moves	Performance
"Winter survival" problem	p-value	0.419	0.949	0.045	0.144
	t-statistic	0.809	0.065	2.010	-1.461
"NASA survival" problem	p-value	0.125	0.062	0.283	0.573
	t-statistic	1.537	1.868	-1.073	-0.564

Note: A positive t-value implies that Study 3 had a higher value than Study 2.

Only the t-test of comparing the *number of moves* in the “winter survival” problem has a significant deviation between Studies 2 and 3 (p-value = 0.045, t-statistic = -2.010; i.e., Study 3 used more moves than Study 2). This is due to the fact that in Study 3, participants made about 0.85 more moves in this task. This is not a large deviation if we consider that participants in both studies made 66% more moves than the minimum required. Additionally, there is one marginally significant deviation in the case of the *time per move* in the “NASA survival” problem. The participants of Study 3 spent on average 0.7 seconds longer on every move—about 8% longer than in Study 2.

These two deviations are smaller in effect size than what Study 3 is designed to measure. The ex-post effect size is 0.052 in the case of the *number of moves* in the “winter survival” problem, and 0.058 in the case of the *time per move* in the “NASA survival” problem. These effects are much smaller than the effects that the experiment is designed to detect, i.e., 0.125. We find them in the comparisons of Table 11 because when we combine both studies the joint sample is composed of 1219 participants. The lack of deviations of a size similar to or larger than that which Study 3 is designed to test gives us support to consider the study an adequate replication of Study 2.

8.2 Hypotheses testing

The results of Study 2 are mostly replicated in Study 3. In Study 3 we find broad support for three of the four hypotheses. These results are shown in Tables 12. In Model 1, we find that the increase in *total time* of participants in both the *framing-focused* and *implementation-focused* conditions is higher than those in the *control* condition, providing support for Hypotheses 1a and 2a. In Model 3, we find support for Hypothesis 1b, as the increase in the *number of moves* of the participants in the *framing-focused* condition is significantly different from the coefficient of the *control* condition. The support for these three hypotheses remains if we change the demographic control variables for pandemic controls (Table 13), when we add the demographic and pandemic sets of control variables together (Table A.11), or when we perform robust regressions instead of ordinary least squares regressions (Table A.12, A.13, and A.14). Although the results Study 3 broadly align with Study 2, we do find deviations.

Table 12: Study 3 OLS Regressions of behavioral change and performance

	Dependent variable:			
	Total time (1)	**Time per move** (2)	**# of moves** (3)	**Performance** (4)
Framing condition	0.546 (0.302, 0.790)	0.225 (-0.037, 0.487)	0.324 (0.116, 0.532)	0.047 (-0.158, 0.252)
Implementation condition	0.296 (0.051, 0.540)	0.208 (-0.055, 0.470)	0.110 (-0.098, 0.319)	0.047 (-0.159, 0.252)
Gender	-0.025 (-0.227, 0.177)	0.015 (-0.202, 0.231)	-0.057 (-0.230, 0.115)	-0.056 (-0.225, 0.114)
Age	0.015 (0.002, 0.028)	0.007 (-0.007, 0.022)	-0.002 (-0.014, 0.009)	0.008 (-0.003, 0.019)
Post Graduate	0.011 (-0.200, 0.222)	-0.026 (-0.252, 0.201)	-0.187 (-0.367, -0.006)	0.123 (-0.055, 0.300)
Reader	0.023 (-0.183, 0.229)	0.102 (-0.119, 0.323)	0.148 (-0.028, 0.324)	-0.080 (-0.253, 0.093)
Constant	-0.476 (-0.975, 0.022)	-0.235 (-0.771, 0.301)	0.022 (-0.404, 0.448)	-0.274 (-0.694, 0.145)
Observations	747	747	747	747
R2	0.032	0.008	0.023	0.006
Adjusted R2	0.024	0.0003	0.015	-0.002
Residual Std. Error	1.384 (df = 740)	1.487 (df = 740)	1.182 (df = 740)	1.165 (df = 740)
F Statistic (df = 6; 740)	4.103 (p = 0.000)	1.033 (p = n.s.)	2.860 (p = 0.009)	0.732 (p = n.s.)
Note:		95% confidence intervals shown in parenthesis		

8.3 Deviations of Study 3

The first deviation between Study 2 and 3 is shown in Model 2 of Table 12. The increase in *time per move* we found for the *implementation-focused* condition in Study 2 is not significant when we include the demographic control variables. A t-test shows that the effect is marginally significant (p-value = 0.079, t-statistic = 1.762), but when the control variables are included, the significance goes away. However, if we use the pandemic control variables, the significance of the results returns to the level we found in Study 2.

We show the regressions with the pandemic controls, rather than demographic control variables, in Table 13. In Model 2 of Table 13, we find that the regression coefficient of the *implementation-focused* condition is significantly different from that of the *control* condition. Therefore, the pandemic control variables account for some of the changes in the participants' behavior. These results are maintained if we use both sets of controls together; these are shown in Table A.11 of the Supplemental Materials[8]. The fact that the increase in *time per move* is present only in a subset of the regressions indicates weaker support for Hypothesis 2b, in comparison to the other three sub-hypotheses. The second deviation is shown in Model 1 of Table 12, where the coefficient for the *framing-focused* condition in Study 3 (Table 12) is almost double the value of Study 2 (Table 9). In Study 2, both the *framing-focused* and the *implementation-focused* conditions increased their *total time* by about 0.274 standard deviations when compared to the *control* condition. However, in Study 3, the increase in time for the *framing-focused* condition was 0.546. This stronger increase in *total time* in the case of the participants in the *framing-focused* condition leads to a marginally significant difference when compared to the increase of the *implementation-focused* condition, whose coefficient remained similar to that of Study 2 (p-value = 0.085, t-statistic = 1.724).

We employed G*Power 3.1 to estimate the ex-post effect sizes for the findings of this study (Faul, Erdfelder, Buchner, & Lang, 2009). The effect size for the *total time* is 0.222, for the *time per move* 0.141, and for the number of moves 0.106. Two of the three effect sizes increased in comparison to Study 2. The increase in effect size for the *total time* appeared as the participants in the *framing-focused* condition were led to spend more time on the "NASA survival" problem.

[8] The results of Models 1, 3, and 4 are aligned between Tables 12 and 13. It is important to note that almost a third of the participants lost their job during the COVID-19 crisis (*job loss* variable). The only regression that is affected by the added pandemic control variables is the one for the *time per move* in Model 2.

Table 13: Study 3 including pandemic controls

	Dependent variable:			
	Total time (1)	**Time per move** (2)	**# of moves** (3)	**Performance** (4)
Framing condition	0.563 (0.315, 0.812)	0.245 (-0.017, 0.506)	0.339 (0.125, 0.553)	0.032 (-0.177, 0.240)
Implementation condition	0.333 (0.081, 0.585)	0.297 (0.033, 0.562)	0.125 (-0.091, 0.342)	0.014 (-0.197, 0.225)
Fear of COVID-19	0.0001 (-0.018, 0.018)	-0.011 (-0.030, 0.007)	0.004 (-0.011, 0.020)	-0.002 (-0.017, 0.013)
Monitor/Blunt Scale	-0.033 (-0.081, 0.015)	0.002 (-0.048, 0.053)	-0.015 (-0.056, 0.027)	0.017 (-0.023, 0.057)
Perceived Vulnerability	-0.003 (-0.020, 0.014)	-0.003 (-0.020, 0.015)	0.005 (-0.009, 0.020)	-0.004 (-0.018, 0.010)
Intolerance to Uncertainty	-0.001 (-0.014, 0.012)	0.013 (-0.001, 0.027)	-0.014 (-0.025, -0.002)	-0.002 (-0.013, 0.009)
Job Loss	0.107 (-0.124, 0.337)	0.229 (-0.013, 0.472)	0.093 (-0.105, 0.291)	-0.108 (-0.301, 0.084)
Prolific Importance	0.030 (-0.056, 0.115)	-0.081 (-0.171, 0.009)	0.056 (-0.017, 0.129)	0.017 (-0.054, 0.088)
Constant	0.306 (-0.881, 1.493)	0.388 (-0.859, 1.634)	-0.276 (-1.296, 0.745)	0.061 (-0.932, 1.055)
Observations	710	710	710	710
R2	0.031	0.020	0.026	0.003
Adjusted R2	0.020 (df = 701)	0.009 (df = 701)	0.015 (df = 701)	-0.008 (df = 701)
Residual Std. Error	1.385	1.455	1.191	1.159
F Statistic (df = 8; 701)	2.821 (p = 0.004)	1.807 (p = 0.073)	2.339 (p = 0.017)	0.307 (p = n.s.)
Note:		95% confidence intervals shown in parenthesis		

Overall, from the replication, we find general support for the findings of Study 2: by manipulating the attention of a participant to focus on a specific phase of problem-solving, we are able to affect their behavior in a consistent way. Although one result was less significant than before, we find support for the idea that the attention treatments lead to increases in the *total time* spent on the tasks, and that focusing the participants' attention on the *framing of the problem* leads them to increase the *number of moves* needed to finalize the task.

9. DISCUSSION

In this paper, we study the micro-processes of strategic problem-solving. These micro-processes have frequently been treated as black boxes (as critiqued in Langley et al., 1995; Posen et al., 2018) and omitted from the theorizing of the problem-solving perspective (as highlighted in Baer et al., 2013). Using precise exploratory methods, we open the black box. Inside it, we discover two alternative strategies that reflect the way managers go about framing, analyzing, and ultimately solving strategic problems when they have neither experience nor feedback. Despite their unfamiliar settings, the problems we chose for our participants to solve share the fundamental characteristics of strategic problems, and have many parallels with the type of problems managers face in real-world organizational settings.

This paper contributes to literature and managerial practice in several ways. Our first contribution is to describe the emergence of these two strategies using exploratory methods. We used think-aloud protocols and a structured data analysis process to allow the two strategies to emerge from the data. Our methods did not pre-specify the number of strategies; we could have found any number of them—or none at all—yet only two emerged. We found that the two strategies seem to differ in how they allocate their attention to different problem-solving phases: either framing or implementation, each of which requires different types of cognitive processes. Building on this finding, we developed a behavioral experiment to test whether and how shifts in attention focus affected the choice of problem-solving strategy. We found that by manipulating participants' attention, we could indeed influence which strategy they adopt.

Our second contribution is to develop a theory that allows us to make predictions about how the allocation of attention drives the way strategic problems are solved. This contribution lies at the intersection of multiple theories of organizations—in particular, the Carnegie School (Gavetti, Levinthal, & Ocasio, 2007), the ABV (Ocasio, 1997; 2011; Ocasio & Joseph, 2018), and theories of managerial problem-solving and decision-making (Langley et al., 1995; Nickerson & Zenger, 2004; Klingebiel & De Meyer, 2013; Felin & Zenger, 2016). We use the

ABV to conceptualize the processes that precede the formation of strategies. The ABV adopts a processual view, defining strategy as “a pattern of attention” rather than a set of actions (Ocasio & Joseph, 2018: 289). Studying attention allows us to answer the call to study how the process of solving managerially relevant problems unfolds beyond the mere linear sequences of decomposed phases, and capture the unfolding micro-processes through which different people solve problems in different ways (Langley et al., 1995; Posen et al., 2018). Uncovering such processes is helpful to understand how managers solve problems and learn, even when there is no possibility of receiving feedback (March, Sproull, & Tamuz, 1991).

Redefining strategies from *what* attention is focused toward *how* is attention focused (Ocasio, 2011; Ocasio & Joseph, 2018) allows us to illuminate the processes that precede actions, and thus understand where differences come from. We conceive attention as a dynamic resource (Bansal, Kim, & Wood, 2018), and like Bansal et al. (2018), we find that attention can be focused, but not spread too thinly: a problem-solving strategy emphasizes attending either to the problem itself or to its solutions, but not both equally. confirming the idea that “The accurate planning and performance of strategic actions and the speed of their execution require that individual and group decision-makers concentrate their energy, effort, and mindfulness on a limited number of issues and tasks” (Ocasio, 1997: 203). Such a need for focus is once again evident when we manipulate attention in Studies 2 and 3. We observe that compared to the control condition, the total thinking time is higher in the two manipulated conditions. This might reflect the fact that, since attention is limited, people tend to conserve this scarce resource, thus deliberating less in the control condition than in the two manipulated conditions when they are instructed to perform certain mental processes. Future studies should investigate how attention is focused under conditions of higher activity load (Castellaneta & Zollo, 2014) when, for example, multiple demands on attention might affect the strategies that emerge, or what abilities help some individuals to more flexibly switch between problem-solving strategies (Laureiro-Martínez & Brusoni, 2018).

Complexity and uncertainty have canonical representations in the behavioral theory of the firm: the "NK landscape" for complexity (Levinthal, 1997; Billinger, Stieglitz, & Schumacher, 2013) and the "N-arm bandit" for uncertainty (Posen & Levinthal, 2012; Laureiro-Martínez et al., 2015). Future studies could build on these representations and use think-aloud protocols to trace search processes. Process studies based on think-aloud protocols hold the potential to directly observe, validate, and refine the model proposed by Cyert and March (1963) when appropriate. Some steps in this direction have been taken by Reypens and Levine (2018), but should be extended to the environments of Billinger et al. (2013) and Laureiro-Martínez et al. (2015) as they map on to the canonical representations. Within the studies of microfoundations of strategy, recent studies show how individuals' specific traits influence their problem-solving. More specifically, this stream of research has shown that cognitive flexibility (Laureiro-Martínez & Brusoni, 2018) and strategic intelligence (Levine, Bernard, & Nagel, 2017) can be seen as antecedents of adaptive decision-making. Our study adds the concept of problem-solving strategies to this repertoire—but, in contrast to prior studies, we show that strategies can be changed by shifting the focus of attention. Future studies could investigate how managers change their problem-solving strategies in response to shifts in attention, whether caused by the manager's own attention focus, and/or the way their attention is directed by organizational structures (Ocasio, 1997).

The knowledge of how strategic problems are solved can serve as the foundation of research on the microstructure of organizations (Puranam, 2018). The microstructural approach argues that by accumulating knowledge on the smallest organizational forms—dyads and triads—we can build organization science from the microstructures up. We agree with this view—yet, as Felin and Foss (2005: 441) point out, "there is no organization without individuals." Often, individual-level heterogeneity is acknowledged merely by controlling for variables such as gender, age, or level of education—if it is not simply disregarded altogether. In fact, as we show in this paper, individual heterogeneity in the way attention is allocated matters greatly. It leads to real differences in the way the most challenging problems are solved, and ultimately engenders the strategies that leaders induce their organizations to follow. Only by understanding how the individual members of a dyad or triad solve problems can we make a judgment on how best to organize them. Future studies could use our findings as a point of departure, and use the methods from Study 1 to investigate whether and how attention is allocated differently when a strategic problem is solved by two individuals working together. This might result in a robust tool to answer "fundamental and universal problems of organizing (that relate to how they aggregate their members' efforts)" (Puranam, 2018: 1). How does a strategy emerge from the interaction of two different problem-solving processes? What happens if two different strategies emerge within a dyad? Do contrasting decision-making strategies complement each other, or simply lead to conflict within the dyad? In a similar vein, the methods of Study 1 could be used in combination with, for example, the task from Cohen and Bacdayan (1994), who recorded the routine formation process of dyads who do not communicate. Future studies could employ think-aloud protocols and record the thinking processes that develop as two participants' strategies form, and how routinized patterns of actions unfold from those emergent strategies. For example, the routines in Cohen and Bacdayan (1994) require cooperation between the agents, but if two agents follow different strategies for doing their routine work, they will need to find a solution that is both fit for purpose and mutually acceptable. Such "good enough" routines are what Nelson and Winter (1982) call "routines as truces." Even though such interaction is a foundational concept in the evolutionary theory of economic change and the behavioral theory of the firm (Cyert & March, 1963), its micro-processes have barely been studied. The methodology of Study 1 could enable future studies to investigate these processes.

Our third contribution is methodological. This paper's combination of methods is a good example of the cycle of theory building and theory testing—methodologically complex, but foundational to the growth of scientific knowledge (Popper, 1963). This paper builds upon and extends prior work on sequence analysis. For example, Salvato (2009) used sequence analysis to study the role of routine activities in the evolution of new product development (NPD) processes, to reveal a firm's capabilities. We take this approach to the micro level, and use sequence analysis to study the role of problem-solving phases in the crafting of solutions to strategic problems. Such solutions are the building blocks of the NPD capabilities that Salvato (2009) studied.

This paper's substantial empirical efforts are now condensed and publicly shared for others' use. We started with a broad and important research question and employed micro-level methods to study it in detail in the controlled conditions of Study 1. We analyzed the findings and allowed them to aggregate into two propositions. We created a behavioral experiment to operationalize the propositions into testable hypotheses. We tested the hypotheses in Study 2 and replicated most of the results in Study 3. The replication was preregistered and followed strict guidelines to show that the results were in line with the initial study. We are sharing the data-collection and data-analysis protocols employed in all studies of this paper to allow other researchers to test our results and claims. We hope that by having access to all tools used to collect and analyze the data, future researchers can continue studying important organizational processes in detail, creating hypotheses and testing them as was done in this paper.

Our fourth and final contribution is to practice. Organizations put significant effort into managing their strategic processes, and adopt a variety of management systems and that support decision-making. Our study proposes a different way to look at such systems, i.e. as devices to influence attention. Existing management systems and tools can be seen as the field equivalents of the attention manipulations in our study. For example, tools such as scenario analysis emphasize the importance of thinking beyond the status quo to identify future problems that might impact an organization before they arise. Technology-scouting tools push managers to identify new technologies for which they might have no immediate use, but could potentially meet a customer need; in the language of this study, they push managers to adopt problem-focused strategies. Other management systems and tools, such as Six Sigma or Lean management (e.g. Liker & Morgan, 2006; Schroeder et al., 2008), are highly analytical, and tend to emphasize the importance of thinking about data and concrete applications. Thus, they push managers to adopt solution-focused strategies. Interestingly, different sets of methods are used by different people in the organization, with senior decision-makers probably focusing more on problem-oriented management systems and tools, while middle managers rely more on Lean-style methods. Our approach highlights the fact that people might have a preferred problem-solving strategy, irrespective of their rank or seniority, so it would be wise for an organization to consider whether to develop their future leaders by training them with a range of tools. One possible recommendation would be to alternate between or cycle through different types of tools, so people can sometimes work in a way that emphasizes the attention focus that is more natural to them, and at other times work with a tool that emphasizes a different focus. Another possible recommendation would be to use methods that allow people to cyclically consider both problem-definition and problem-solving. For example, in Design Thinking, participants first focus attention on the problem, then shift to considering solutions, and then iterate (Elsbach & Stigliani, 2018). A shifting focus of attention is also present in the scientific approach to entrepreneurial decision-making advocated by Camuffo et al. (2019). In their approach, startups first attend to the problem by building key performance indicators, and then test their indicators, shifting their attention to the solution. After testing, new indicators are built, and the cycle continues. While we are still some way from establishing the external validity of our findings, we believe that the analysis of management tools as attention-focusing devices could provide us with an excellent context to extend our ideas to field studies, whether qualitative or experimental in nature.

In conclusion, this study broadens our understanding of how strategic problems are solved. Prior studies relied on distant analogies that only reflected certain characteristics of strategic problems. For example, Newell and Simon (1972) studied chess, which is admittedly complex, but far more structured than the problems managers typically face. With this study, we bring process-level data to tasks that, though apparently far-fetched in the contexts they involve, nevertheless require participants to grapple with problems that are ill structured, complex, and novel, with high-stakes, irreversible outcomes. Thus, we show how people grapple with a type of problem that managers face under conditions that are representative of the complexities and rapid change of their actual organizational lives.

10. SUPPLEMENTAL MATERIALS

This paper encompasses several different methods, each with their respective data sets. In this appendix of supplementary material, we aim to make our processes for data collection and data analysis fully transparent. Study 1 is qualitative in nature and employs think-aloud protocols. Studies 2 and 3 are behavioral experiments, one of which is a preregistered replication. Below we present the tasks assigned to study participants in greater detail, as well as robustness checks for the clustering employed in Study 1.

Further descriptions of the tasks and data are available at the Open Science Framework repository for this project. This information should enough to replicate the studies and analyses we do in this study. The repository is available online, and the written documents are presented in section 11 (osf.io/eh5m2/?view_only=2bd6e1e7320548858fd872db4c658932).

In this repository, we included the data-collection and data-analysis protocols used in all the main studies of this project, as well as all preregistered data collections. For Study 1, we included a guide for collecting think-aloud protocols, along with two fully coded verbal protocols. These protocols are the verbatim transcripts of Person A and B we present in Study 1. We further explain, step by step, how we coded the data, created sequences from the coded transcripts, and stored these sequences in transitions matrices. We then provide the 49 transition matrices from Study 2, and the code we use to create each table and robustness check of this study, including the robustness check available in these supplementary materials.

In the Open Science Framework repository, we also include the surveys used to collect the data for Studies 2 and 3. These surveys include the tasks employed in the study and all the extra questions we asked the participants. For the purposes of this project, we report the main variables used in hypothesis testing. However, we collected much more detailed data. Therefore, in the repository, we include this data, as well as the scripts used to transform it from time-traced mouse-clicks to the aggregate variables used in hypothesis testing. This data is important to allow future researchers to trace the data analysis we did to the one we stated in the preregistration of Study 3—but, more generally, to allow future researchers to use the data more freely and not be bound to our variable definitions.

Below, we present the three tasks employed in this study. Furthermore, we explain when and how we collected data, and how the data collection was designed and handled.

10.1 Tasks used in this paper

The three problems used in this paper have been published and validated. We present each below:

10.1.1 "Karabayos" problem (Laureiro-Martínez & Brusoni, 2018)

In Study 1, we rely on a task that was used as an ill-structured problem by Laureiro-Martínez and Brusoni (2018). This task asks participants to imagine themselves as "the leader of the Karabayos," an Amazonian tribe. The Karabayos tribe faces a set of threats and challenges, and it is the task of its leader to imagine what to do to save the tribe. The full text and layout of the task as given to the participants are shown in Figure A.1

10.1.2 "Winter survival" problem (Johnson & Johnson, 1982: 111)

Participants in Study 2 were presented first with the "winter survival" problem of Johnson and Johnson (1982: 111). This task requires participants to imagine themselves at the site of a plane crash that they have just survived. It is now midday, the temperature is freezing, and they find themselves in a forest with a group of fellow survivors who will follow their commands. They have a list of 12 items they can use to survive the night, and they must rank these items based on their importance to their survival. The participants are told there is a town 30 kilometers away; this information complicates the problem, as participants have to determine whether they prefer to stay near the crash site and await rescue, or walk to the town in the freezing temperatures.

To finish the problem, participants need to move the 12 items into the order they deem appropriate. Figure A.2 presents the graphical user interface that participants used during the task, which includes the full text of the task. Figure A.2 shows the task before any moves have been made. In order to finish the task, the participant needs to drag and drop all items from the left column to the right column. After all the items have been placed, a button asking the participant to "confirm final ranking" appears on the screen. This button is shown in Figure A.3.

Survival experts gave Johnson and Johnson (1982) an optimal ranking of the items. They determined that, first of all, staying near the crash site is the best decision. To survive until help arrives, it is of utmost importance to prioritize heating and food supplies. Table A.1 shows the 12 items in the order they are shown to the participants, along with the correct positions in which they should be ranked, according to the experts.

Figure A.1: Text and layout of the problem used in Study 1

Imagine that you are the leader of the Karabayos, a small tribe (22 women and 26 men) in the Amazon rainforest. There are hundreds of tribes in the world that, like yours, have never had any contact with other peoples and are scattered in the vast jungles of South America, New Guinea and the Indian Ocean.

You all have one characteristic in common: you are the most vulnerable people in the world and you want to be left in peace. And for a good reason! The history of contact between indigenous tribes and the rest of the world has always been particularly unfortunate.

Contact with other people is almost always a disaster for these types of tribes, who have lived according to a lifestyle largely intact for more than 10,000 years. Your lifestyle does not include the use of television, microwaves, or cars. You never feel the need of any of these. Most of these tribes live in hidden places inside the forest. However, many of these hidden places are getting closer to the areas under the control of rubber producers, loggers, settlers and drug traffickers, which endanger the survival of the tribes.

Even when the loggers do not kill any tribe members directly (which often happens), after contact with other peoples the tribes are decimated within a year or two from many diseases (such as influenza, measles, chicken pox) against which these tribes have no immunity.

The peace and harmony of your tribe, and the abundance of your lands, which for centuries have allowed you to live in balance with nature is constantly endangered by the approach of "civilization".

Your group lives in an area that for years has been full of fruit trees and animals of all kinds. Traditionally, men hunt with bows and blowguns, while women stay at home to take care of children. You are aware that some parts of the area you live in are bordering areas of the "whites" and, for years, your people have avoided coming into contact with them. In recent years, you have realized that the trees no longer produce as many fruits as they used to, and that many of the animals you used to hunt have disappeared. You ask yourself what you should do to save your tribe.

We adapted the units to the International System of Units and changed the names of the items slightly after pilot sessions showed us that some items were hard for participants to understand (e.g., "can of shortening"). We kept the original formulation from Johnson and Johnson (1982) and added short explanations in parentheses. Additionally, on the tiles that participants needed to drag and drop, we added a small image of the item to make it easier to recognize. We took the same approach in the case of the "NASA survival" problem presented in the next section.

Figure A.2: User interface and layout of the "winter survival" problem

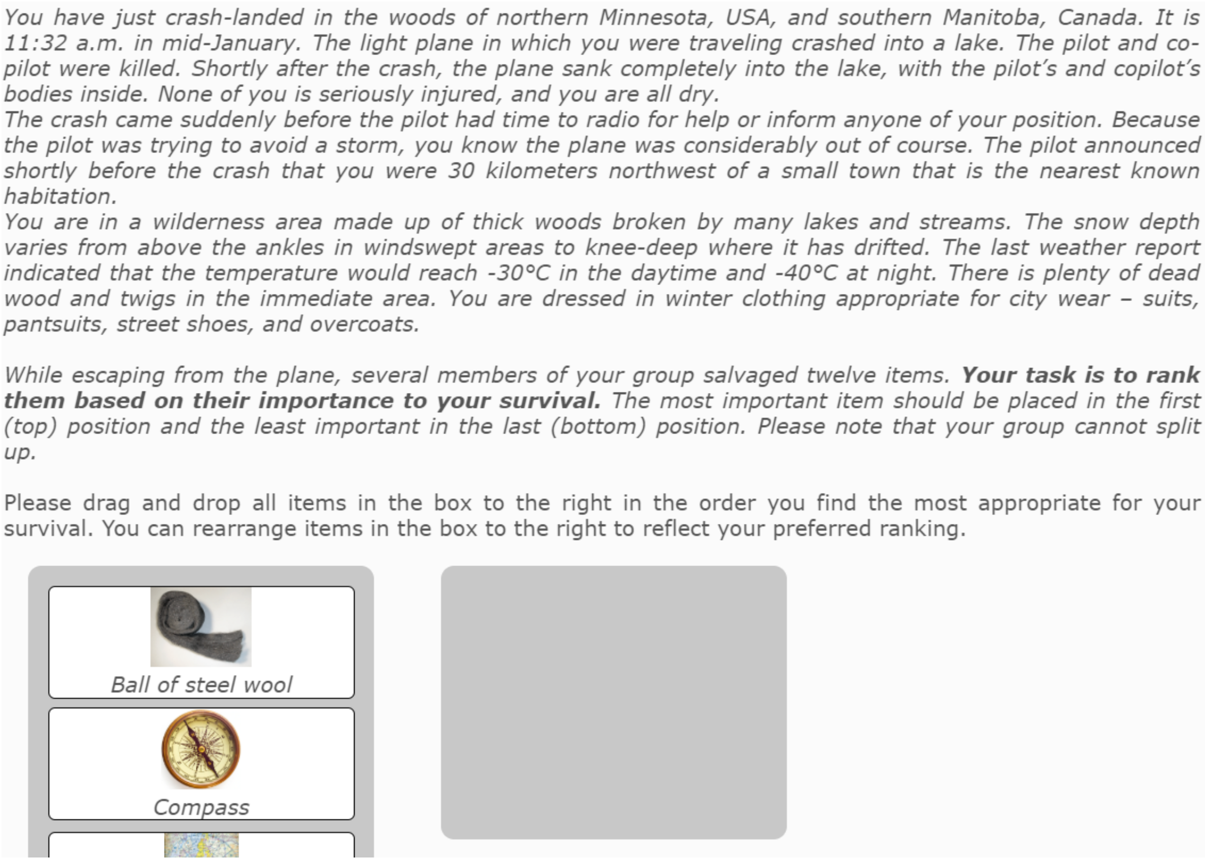

Table A.1: List of items used in the "winter survival" problem

Order Shown	Correct Ranking	Item
1	2	Ball of steel wool
2	12	Compass
3	1	Cigarette lighter (without fluid)
4	11	Sectional air map made of plastic
5	3	Extra shirt and pants for each survivor
6	10	Quart of 100-proof whiskey
7	8	Newspaper (one per person)
8	6	Hand ax
9	9	Loaded 0.45 caliber pistol
10	5	6x6 meter of heavy-duty canvas
11	4	Can of shortening (margarine)
12	7	Family size chocolate bar (one per person)

10.1.3 “NASA survival” problem (Hall & Watson, 1970)

After performing the “winter survival” problem, participants were shown the manipulations. The manipulations are explained in section 7.3. After the manipulation, participants solved the “NASA survival” task (Hall & Watson, 1970). This task, like the “winter survival” problem, requires participants to imagine themselves at a crash site. In this case, however, they have crashed on the surface of the moon along with the crew of their lunar module. There is only one possible course of action: they must try and reach the meeting point. To reach the meeting point, the participants are given 15 items. The task is to rank these 15 items in terms of their importance in allowing the crew to reach the meeting point. As in the “winter survival” problem, Hall and Watson (1970) obtained an expert ranking and used this ranking to compare the participants’ responses to an objective measure. Table A.2 shows the items in the order shown to the participants, and the correct position each item should be ranked.

Figure A.3 presents the implementation-focused condition’s graphic interface. It shows the interface after all items were placed in the right-hand column. It also shows the “confirm final ranking” button that, once clicked, takes participants to the next stage in the experiment. The sentence in red text changed in every experimental condition, as the next section explains in more detail.

Table A.2: List of items used in the “NASA survival” problem

Order Shown	Correct Ranking	Item
1	15	Box of matches
2	6	15 meters of nylon rope
3	13	Portable heating unit
4	12	One case of dehydrated pet milk
5	3	Stellar map (of how constellations look on the moon)
6	14	Magnetic compass
7	10	Signal flares
8	5	Solar-powered FM receiver transmitter
9	4	Food concentrate
10	8	Parachute silk
11	11	Two 0.45 caliber pistols
12	1	Two fifty-kilo tanks of oxygen
13	9	Life raft
14	2	20 liters of water
15	7	First aid kit containing injection needles

Figure A.3: User interface and layout of the "NASA survival" problem, implementation-focused condition

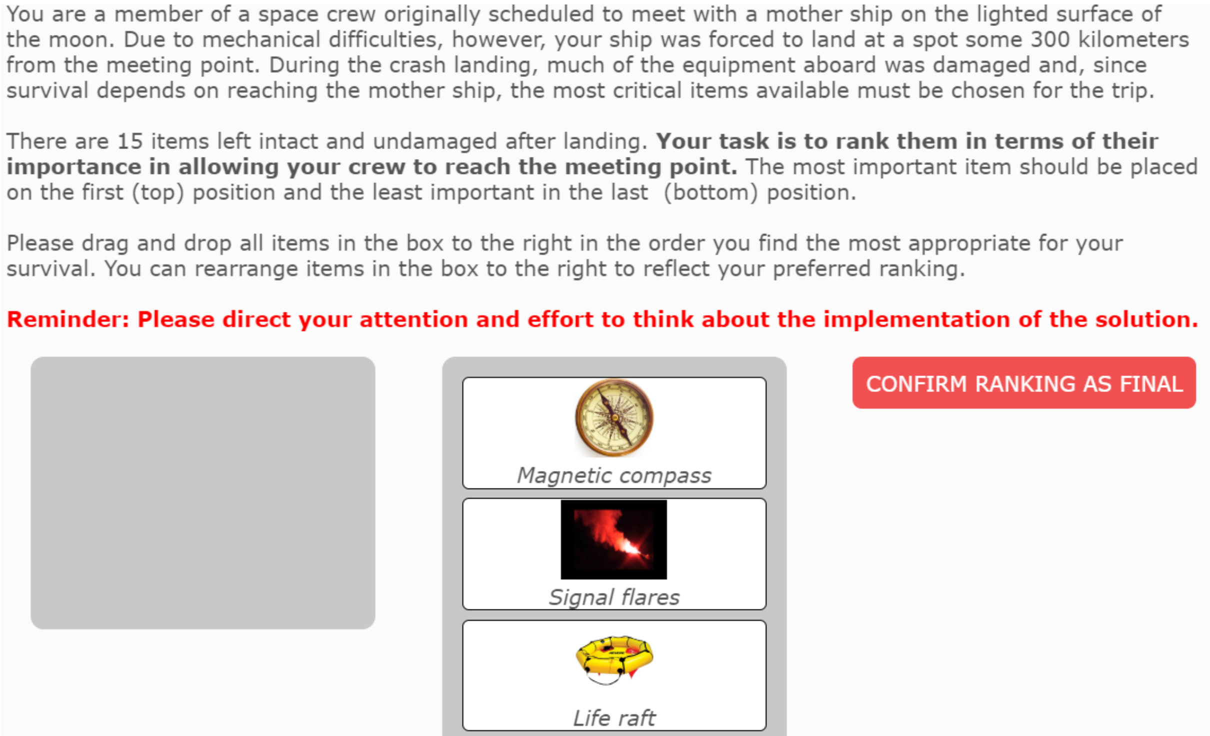

10.2 Research design of Study 2

Study two followed a mixed factorial experimental design, as it mixed a between-subject design and a within-subject design (Oehlert, 2010; Anderson & McLean, 2018). This is because the experiment had three experimental conditions (between-subject design), and each participant performed two tasks—one before the manipulation, and one afterwards (within-subject design). We follow this procedure to better pinpoint the causality of the attention-focus mechanism by reducing the amount of unexplained variance in our analyses, and studying only the behavioral change induced by the manipulations.

We hosted the online experiment using the Qualtrics platform. Qualtrics provides multiple tools for creating tasks, such as Likert scales, multiple-choice quizzes, and even drag-and-drop interfaces. However, in our experiment, we wrote our own interface in JavaScript. We did this in order to capture detailed timings of every movement the participant made while solving the task—every click, drag, and drop. We recorded and timed each of these events, as well as the time it took the participant to read the task, and to leave the task once their ranking was finished. The motivation for this level of detail was to understand the thinking processes of the participants beyond reaction time and solution performance (Yu et al., 2012). We followed the methods of studies that have used mouse actions to provide a "*more direct measure of the evolution of a particular response*" (Freeman, 2018). It has been found that mouse movements are more representative of cognitive processing than other subjective measures, such as self-reports (Fedor et al., 2015), and have proven a reliable means for examining continuous cognitive processes in real time (Yu et al., 2012), providing a good proxy for the analysis of thinking processes (Ollinger et al., 2013).

We recruited participants to our experiment through the online platform Prolific, and paid them based upon their performance. The top quartile of participants received twice the base rate of the platform for the 30 minutes of solving the experiment: a total of £5 (British pounds sterling; GBP). The middle half received 1.5 times the base rate of the platform: £3.75. The bottom quartile was paid £2.50: the base rate for the half an hour the experiment required of the participants.

Figure A.4 presents a depiction of the experimental procedure of Study 2. An experimental session unfolded according to the following sequence. First, Prolific referred us a participant. The participant arrived at the experiment and was welcomed. We requested their compliance with the conditions of the experiment and presented the incentives scheme, and the participant performed an attention check regarding the incentive scheme. After passing the attention check, the participant performed the first task: the "winter survival" problem. From this task, we saved the final ranking; the reading, processing, and total time spent on the task; the number of moves; and the time each move took.

Figure A.4. Visualization of the research procedure

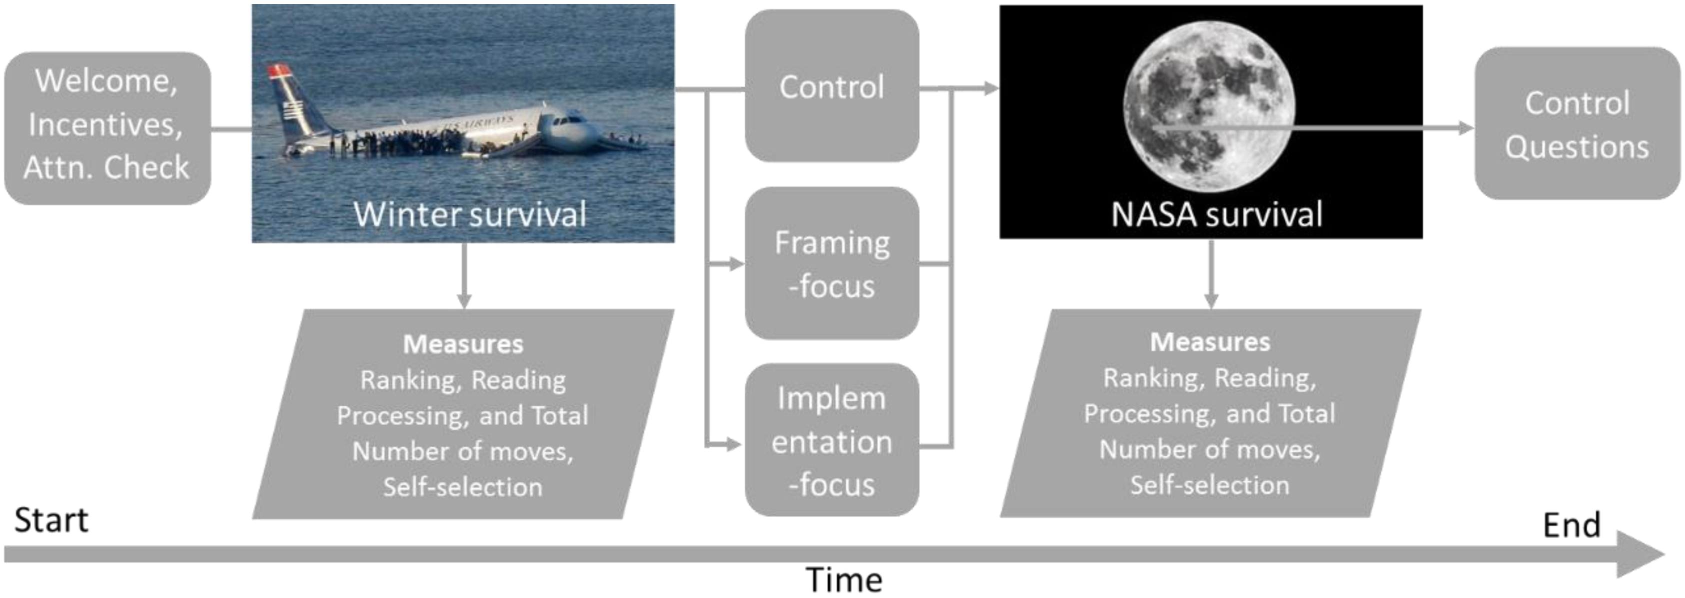

After the participant confirmed the solution, they were asked to write down the motivation for how they built the ranking, and afterward to explain to us the process of how they built the ranking. Specifically, we asked them to "imagine what a tape recorder would play if it had recorded what you thought about when ranking the items." After they finished writing, we asked the participants to select, from a list of prototypical processes, the one that most closely resembled their own thinking process. Given that they had just spent time writing about their thinking process, they had a deliberate comparison from which to self-select. Additionally, participants needed to answer an attention check at this point.

Up to this point, the experience of every participant in the study had been the same (i.e., the three conditions were hitherto indistinguishable). The next step was the manipulation, which was shown as a transition page between the two tasks. We created three pages, one per manipulation, in order to incentivize all participants in a similar way except for the focus of their attention.

In the manipulation, we wanted participants to focus their attention on the specific phases of problem-solving on which the strategies of Study 1 differed. To do this, we took the coding scheme the raters used to code the problem-solving phases in Study 1 and created recommendations for the participants directly out of the coding scheme. By using the same language as the coding scheme, we can be closer to the original difference in attention focus found in Study 1, even if this difference is not linked to a theoretical or cognitive mechanism.

The manipulations had three parts. Figures A.5, A.6, and A.7 show screen captures of the graphical interface the participants saw on each the three conditions. The first part explained that experts recommend individuals should spend most of their time thinking about either: a) "the problem in the way that feels most natural to them" in the case of the *control* condition; b) "the framing of the problem" in the case of the *framing-focused* condition; or c) "the implementation of the solution" in the case of the *implementation-focused* condition. The literature on naturalistic decision-making (Klein, 2017), planning (Steiner, 2010), and forecasting (Tetlock & Gardner, 2016) give recommendations that could be interpreted as suggestions to focus in ways that follow either of the three recommendations. So, given that there is no scientific consensus, and we can find research findings backing all three manipulations, we did not deceive the participants. The three recommendations could lead to either better performance or no performance difference; in this study, we give support to the latter.

The second part of the manipulation was present only in the treatment conditions (i.e., the *framing-focused* and *implementation-focused* conditions). This part explained what thinking about the "framing of the problem" or "implementation of the solution" means. The explanation was given in three bullet points. The bullet points were précis of the coding schemes for the phases used most differently by the two strategies identified in Study 1. Namely, the three bullet points of the *framing-focused* condition came from summarizing the *frame stating* and *frame assuming* phases. In the case of the *implementation-focused* condition, the bullet points were summaries of the coding schemes of the *implementation* and *implementation evaluation* phases.

The third and final part of the manipulation slide was a reminder for the participants to focus their attention, as shown in Figures A.5, A.6, and A.7. After reading the manipulations, the participants started the "NASA survival" problem. The interface for the three conditions was the same, except for the sentence shown in red (see Figure A.3). In this sentence, we reminded participants to "Please direct your attention and effort" in either: a) "ways that feel natural to you"; "to think about the framing of the problem"; or "to think about the implementation of the solution." The sentence is the last part of the manipulation we gave the participants.

Figure A.5: Manipulation for the control condition of Study 2

Welcome to the second task.

Your solution to this problem will affect your performance evaluation. There is no time limit.

Research and experts in problem-solving recommend that in order to achieve higher performance, individuals should spend most of their time thinking about the problem in the way that feels more natural to them.

It is crucial that while solving the next task you direct your attention and effort to think about the problem in the way that seems most natural to you.

On the next page, please carefully read the problem description and answer accordingly.

Figure A.6: Manipulation for the framing-focused condition of Study 2

Welcome to the second task.

Your solution to this problem will affect your performance evaluation. There is no time limit.

Research and experts in problem-solving recommend that in order to achieve higher performance, individuals should spend most of their time thinking about the framing of the problem.

The framing of the problem involves important mental activities linked to higher performance, such as:

- **Analyze the problem by recalling the available information**
- **Empathizing to identify with the situation**
- **Develop hypotheses/assumptions to gain an understanding of the problem**

It is crucial that while solving the task you direct your attention and effort to think about the framing of the problem.

On the next page, please carefully read the problem description and answer accordingly.

Figure A.7: Manipulation for the implementation-focused condition of Study 2

Welcome to the second task.

Your solution to this problem will affect your performance evaluation. There is no time limit.

Research and experts in problem-solving recommend that in order to achieve higher performance, individuals should spend most of their time thinking about the implementation of the solution.

The Implementation of the solution involves important mental activities linked to higher performance, such as:

- **Design the sequence of actions that could unfold during the solution of the problem**
- **Anticipate how events will carry out**
- **Evaluate the feasibility of the solutions**

It is crucial that while solving the task you direct your attention and effort to think about the implementation of the solution.

On the next page, please carefully read the problem description and answer accordingly.

We recorded the same variables in this task as in the "winter survival" problem (e.g., rankings, times, number of moves, process self-selection). After the participant finished the "NASA survival" problem, we asked them to answer a set of control questions. In between the control questions, we included an attention check. The control questions had two aims: first, to understand whether the participants had understood the task and paid attention, and second, to collect demographic information about the participant. The main questions on this second theme were *age*, *gender*, *education level*, *reading habits,* and *experience in survival training*. These questions were important during the piloting process and helped us to improve the experimental design.

We removed participants with experience in survival training because they might not see novelty in the problems—or, at least, less than other participants did. A participant who sees the problem as less novel will also see it as less strategic. Study 2 aims at replicating the conditions of Study 1, where the problem was novel to all participants. Thus, removing the participants with experience in the contexts presented was a way of maintaining the similarity between the studies.

After the control questions were answered, we thanked the participants and referred them back to Prolific, where they could claim remuneration for partaking in our experiment. The experiment was held at the end of May 2018. The experiment included 523 participants from which 51 were removed due to survival training. Out of the 472 remaining participants, 276 were in the control group, 97 in the *framing-focused* group, and 99 in the *implementation-focused* group. We estimated sample size based upon G*Power 3.1 for a three condition by two task experiment with omission and commission error rate of 5% (Faul et al., 2009). The control condition was larger in order to understand the baseline behavior better, and to have a better understanding of the correlational results we found while piloting this experiment.

10.3 Piloting of Study 2

Study 2 was preceded by four pretests and two pilot studies. The first four pretests were mainly aimed at selecting the tasks, refining the instructions, and polishing the computer interface. The two pilot studies were mainly aimed at refining the manipulation and incentive scheme. After each study, we held debriefing sessions.

For the initial four pretests we included three tasks” the “NASA survival” task, the “winter survival” task, and a case analysis regarding the LEGO Corporation circa 2003 obtained from Wellian (2010), which we later eliminated. We had initially attempted to capture participants’ thinking process by asking them to briefly describe it. The four pretests took place during March–August 2017 and were carried out with samples of 21, 19, 60, and 61 participants. The pilots were run on the Prolific platform, and participants incentivized with a flat base rate. From the pretests, in addition to our main goal of polishing the interface, we learned that participants found it very difficult to abstract their own thinking process in a clean manner. They mixed up the description of their thinking process with justifications about the motivation for their choices. In later studies, we asked participants first to explain their motivation, later explain their process, and finally match their thinking processes to the prototypical quotes we wrote to resemble *implementation-focused* or *framing-focused* problem-solving process. This seemed to help, but still, participants acknowledged that they were frequently unaware of their own thinking process, meaning that the self-classification measures were of limited use. As a result, we dropped the case analysis regarding the LEGO Corporation after the first pilot. This was mainly because we needed each participant to solve two problems for our experiment, and solving two such lengthy cases could easily lead to cognitive depletion—an interesting topic for future studies, but not the purpose of ours.

After the pretests, we explored different ways to manipulate the participants’ attention. Our first attempt at creating a manipulation that could be at the root cause of the different strategies observed in Study 1 involved using episodic future thinking. Episodic future thinking is the ability of the brain to imagine future events in detail (Schacter, Benoit, & Szpunar, 2017). We attempted to recreate the solutions observed in the solution-focused cluster by inducing episodic future thinking. In October 2017, we ran the first pilot study, the first with enough statistical power to refute hypotheses (Faul et al., 2009). The first pilot study was run in Prolific and involved 158 participants. Our assumption for this pilot study was that the *implementation-focused* strategy might use more episodic future thinking, and imagine how the solution could be implemented in the future at the expense of focusing on the problem in the present. In contrast, the *framing-focused* strategy would be more present-focused. To test this, we ran a pilot study using with two conditions (high episodic future thinking, control), but the manipulation was not effective. We did not find an effect on the manipulation checks regarding the use of episodic future thinking and temporal focus, and we did not find an effect on performance or behavior. We found some correlational results on the participants' self-selection, but nothing causal.

After the first pilot study, we decided to remove our interpretation from the design of the experiment and strip down the manipulation to the core of the actual differences we found in Study 1 (i.e., the attention paid to different phases of the problem). We then created manipulations to incentivize participants to focus their attention on either *implementation* and *implementation evaluation* or *frame stating* and *frame assuming.* We tested the manipulations in April 2018 with our second pilot study, which included 46 participants. The key learning from this pilot was that participants forgot the incentive scheme at the end of the experiment. We needed to incentivize participants' performance more directly and check whether they had understood the incentive scheme. We found out that an attention check at the end was too late. Better results were obtained by inserting an attention check on the incentive scheme at the start of the experimental procedure. With this amendment, we ran the experiment shown in Study 2 in May 2018.

Finally, we ran a small scale, follow-up study with a group of 51 student volunteers during June 2018. This study helped us understand individuals' actual thinking processes, and allowed for an in-depth qualitative debriefing of the final experimental design presented in the paper. We validated our expectations, but more importantly, we learned two points. First, we gained a qualitative understanding of the quantitative measures we obtained from the behavioral experiments we had run in the past. For example, participants elaborated on what was going on in their heads when they were taking a long time before a drag-and-drop move (i.e., they were thinking carefully about how the solution in their heads might unfold through time), or what was going on when they were making many moves (i.e., they were trying to connect different elements of the problem, and thinking about their use and possibilities for different sub-goals). Second, we checked for the ecological validity of the task by asking the students—most of whom had years of experience as managers—to draw parallels between the experiment tasks and some real work situations. While not all were able to perceive parallels straight away, some made immediate connections with personal challenges they had faced as managers, when they had ensured the survival of their startup, or their business unit, and the jobs of their employees under stressful conditions. Overall, while time-consuming, the various studies were a fundamental source for building the final research design for Study 2 and for understanding its robustness.

10.4 List of variables used in Study 2

In Study 2, we collected detailed measures of the participants' problem-solving processes. In the paper, we discuss four behavioral change variables, named *total time*, *time per move*, *number of moves*, and *performance*. To build each of these variables, we need one variable from the "winter survival" problem (before the manipulation) and one from the "NASA survival" problem (after the manipulation). In addition to these variables, we collected the age, gender, education level, and reading habits of the participants, referred to as *reader*. Finally, we collected two variables that relate to the choice participants made when they self-selected their thinking process into *framing-focused* or *implementation-focused*. The variable *implementation-focused* is a dichotomous variable that has a value of 1 if the participant said the text closest to their thinking process was a prototypical *implementation-focused* process, and 0 if they chose a *framing-focused* process. The variables and brief descriptions are shown in Table A.3. Table A.4 shows descriptive values and first-order correlations of the variables collected for Study 2.

Table A.3 List of variables in Study 2

	Variables	**Definition**
Experimental Condition	1. Control condition	Member of the *contro*l condition (1 if yes, 0 if no)
	2. *Framing focused* condition	Member of the *framing focused* (1 if yes, 0 if no)
	3. *Implementation focused* condition	Member of the *implementation focused* condition n (1 if yes, 0 if no)
Demographics	4. Gender	Gender of the participant (1 if female, 0 if male)
	5. Age	Age of the participant in years
	6. Postgraduate	The education level of the participant (1 if holds a masters degree or higher, 0 if lower)
	7. Reader	Reading habits of the participant (1 if more than twice a week, 0 if less)
Behavioral Change	8. Performance	Change in standardized performance between tasks; negative value implies improvement [s.d.]
	9. Total Time	Change in standardized total time between tasks [s.d.]
	10. Time per Move	Change in standardized time per move between tasks [sd.]
	11. # of Moves	Change in standardized total time between tasks [s.d.]
“NASA survival” problem	12. Performance	Performance in the “NASA survival” problem; lower is better, 0 is perfect
	13. Total Time	Total time spent in the “NASA survival” problem [s]
	14. Time per Move	Time per move in the “NASA survival” problem [s]
	15. # of Moves	Number of moves in the “NASA survival” problem
	16. Self-selection *Implementation focus*	Self-selected a prototypical *implementation-focused* problem-solving process text (1 if yes, 0 if no) in the “NASA survival” problem
“Winter survival” problem	17. Performance	Performance in the “winter survival” problem; lower is better, 0 is perfect
	18. Total Time	Total time spent in the “winter survival” problem [s]
	19. Time per Move	Time per move in the “winter survival” problem [s]
	20. # of Moves	Number of moves in the “winter survival” problem
	21. Self-selection *Implementation focus*	Self-selected a prototypical *implementation-focused* problem-solving process text (1 if yes, 0 if no) in the “winter survival” problem

It is important to note that our data collection method allowed us to record the step-by-step process by which participants built their solutions. This allowed us to create analyses similar to the move-and-time analysis from Fedor, Szathmáry, and Öllinger (2015). This data is not included in this document, as here we are focused on the effects of the manipulation on the behavioral change.

Table A.5 presents the equivalent robust regressions of the ordinary linear squares regressions shown in Table 9 of the main manuscript.

Table A.4: Descriptive and first-order correlation table of the variables of Study 2 (part 1 of 3)

	Variables	1.	2.	3.	4.	5.	6.	7.
Experimental Condition	1. Control condition	1						
	2. Framing focused cond.	-0.604 (0.000)	1					
	3. Implementation cond.	-0.611 (0.000)	-0.262 (0.000)	1				
Demographics	4. Gender	0.006 (0.895)	-0.024 (0.600)	0.017 (0.719)	1			
	5. Age	-0.001 (0.989)	-0.012 (0.795)	0.013 (0.783)	0.027 (0.556)	1		
	6. Postgraduate	-0.038 (0.415)	0.010 (0.827)	0.035 (0.442)	-0.070 (0.129)	-0.034 (0.460)	1	
	7. Reader	0.034 (0.462)	-0.095 (0.040)	0.053 (0.252)	0.122 (0.008)	0.173 (0.000)	0.200 (0.000)	1
Behavioral Change	8. Performance	0.028 (0.546)	-0.017 (0.71)	-0.017 (0.718)	-0.017 (0.717)	0.080 (0.082)	0.024 (0.601)	-0.020 (0.667)
	9. Total Time	-0.131 (0.005)	0.083 (0.071)	0.075 (0.102)	-0.102 (0.027)	0.021 (0.652)	-0.109 (0.018)	-0.055 (0.230)
	10. Time per Move	-0.097 (0.035)	0.004 (0.931)	0.114 (0.014)	-0.070 (0.130)	0.006 (0.898)	-0.059 (0.198)	-0.087 (0.060)
	11. # of Moves	-0.073 (0.114)	0.104 (0.023)	-0.015 (0.738)	-0.011 (0.817)	-0.048 (0.293)	-0.005 (0.921)	0.073 (0.112)
NASA Survival Exercise	12. Performance	0.046 (0.316)	-0.025 (0.587)	-0.031 (0.500)	-0.050 (0.274)	-0.016 (0.723)	-0.032 (0.492)	-0.014 (0.769)
	13. Total Time	-0.036 (0.430)	0.013 (0.778)	0.031 (0.499)	-0.005 (0.920)	0.144 (0.002)	-0.079 (0.087)	-0.007 (0.878)
	14. Time per Move	-0.030 (0.511)	-0.038 (0.406)	0.075 (0.105)	0.008 (0.861)	0.100 (0.030)	-0.048 (0.300)	-0.033 (0.469)
	15. # of Steps	-0.041 (0.372)	0.103 (0.025)	-0.053 (0.255)	0.013 (0.782)	-0.022 (0.628)	-0.051 (0.273)	0.044 (0.343)
	16. Self-sel. *Impl. focus*	-0.017 (0.712)	-0.067 (0.147)	0.087 (0.059)	0.042 (0.368)	-0.013 (0.779)	0.027 (0.552)	-0.006 (0.900)
Winter Survival Exercise	17. Performance	0.011 (0.816)	-0.003 (0.945)	-0.010 (0.831)	-0.029 (0.531)	-0.118 (0.011)	-0.062 (0.179)	0.012 (0.800)
	18. Total Time	0.097 (0.034)	-0.072 (0.119)	-0.047 (0.313)	0.099 (0.032)	0.118 (0.011)	0.035 (0.449)	0.049 (0.285)
	19. Time per Step	0.077 (0.096)	-0.042 (0.367)	-0.051 (0.264)	0.084 (0.068)	0.091 (0.049)	0.018 (0.689)	0.062 (0.178)
	20. # of Steps	0.039 (0.400)	-0.013 (0.773)	-0.034 (0.464)	0.024 (0.607)	0.031 (0.506)	-0.044 (0.345)	-0.037 (0.426)
	21. Self-sel. *Impl. focus*	-0.041 (0.377)	0.052 (0.263)	-0.002 (0.967)	-0.030 (0.517)	0.010 (0.820)	-0.001 (0.981)	0.006 (0.905)
	M **SD**	276 of 472	97 of 472	99 of 472	269 of 472	34.88 8.51	161 of 472	219 of 472

Note: p-value of the pairwise correlations shown in parenthesis

The mean value and standard deviation of the behavioral change variables (8-11) are not 0 and 1 respectively, because we use only the data of the participants in the control condition to standardize the “NASA survival” problem data. The behavior of the other condition changed; thus, the mean and the standard deviation do as well.

Table A.4: Descriptive and first-order correlation table of the variables of Study 2 (continuation 2 of 3)

	Variables	8.	9.	10.	11.	12.	13.	14.
Behavioral Change	8. Performance	1						
	9. Total Time	-0.065 (0.159)	1					
	10. Time per Move	-0.032 (0.489)	0.662 (0.000)	1				
	11. # of Moves	-0.073 (0.112)	0.263 (0.000)	-0.139 (0.002)	1			
NASA Survival Exercise	12. Performance	0.631 (0.000)	-0.114 (0.014)	-0.119 (0.010)	-0.095 (0.038)	1		
	13. Total Time	-0.043 (0.350)	0.489 (0.000)	0.353 (0.000)	0.072 (0.119)	-0.259 (0.000)	1	
	14. Time per Move	-0.039 (0.396)	0.283 (0.000)	0.536 (0.000)	-0.151 (0.001)	-0.208 (0.000)	0.695 (0.000)	1
	15. # of Steps	-0.045 (0.327)	0.227 (0.000)	-0.117 (0.011)	0.523 (0.000)	-0.112 (0.015)	0.359 (0.000)	-0.162 (0.000)
	16. Self-sel. *Impl. focus*	-0.083 (0.073)	0.052 (0.256)	0.002 (0.959)	0.047 (0.313)	-0.123 (0.007)	0.029 (0.533)	-0.001 (0.990)
Winter Survival Exercise	17. Performance	-0.638 (0.000)	-0.031 (0.507)	-0.078 (0.09)	-0.002 (0.964)	0.194 (0.000)	-0.202 (0.000)	-0.157 (0.001)
	18. Total Time	0.025 (0.595)	-0.544 (0.000)	-0.332 (0.000)	-0.197 (0.000)	-0.133 (0.004)	0.466 (0.000)	0.381 (0.000)
	19. Time per Step	-0.003 (0.946)	-0.449 (0.000)	-0.572 (0.000)	0.006 (0.899)	-0.072 (0.118)	0.289 (0.000)	0.386 (0.000)
	20. # of Steps	0.035 (0.442)	-0.064 (0.163)	0.037 (0.422)	-0.574 (0.000)	-0.005 (0.918)	0.267 (0.000)	0.007 (0.884)
	21. Self-sel. *Impl. focus*	-0.033 (0.477)	0.067 (0.149)	0.039 (0.392)	0.024 (0.609)	0.060 (0.191)	-0.013 (0.779)	-0.011 (0.814)
	M	-0.039	0.029	0.025	0.033	49.34	363.0	8.853
	SD	1.265	1.015	1.093	1.077	15.48	208.8	5.921

Table A.4: Descriptive and first-order correlation table of the variables of Study 2 (continuation 3 of 3)

	Variables	15.	16.	17.	18.	19.	20.	21.
NASA Survival Exercise	15. # of Steps	1						
	16. Self-sel. *Impl. focus*	-0.007 (0.878)	1					
Winter Survival Exercise	17. Performance	-0.054 (0.242)	-0.018 (0.697)	1				
	18. Total Time	0.114 (0.013)	-0.025 (0.581)	-0.163 (0.000)	1			
	19. Time per Step	-0.029 (0.533)	-0.003 (0.946)	-0.068 (0.142)	0.734 (0.000)	1		
	20. # of Steps	0.398 (0.000)	-0.057 (0.217)	-0.050 (0.282)	0.322 (0.000)	-0.034 (0.462)	1	
	21. Self-sel. *Impl. focus*	-0.003 (0.955)	0.031 (0.501)	0.101 (0.028)	-0.080 (0.083)	-0.054 (0.244)	-0.028 (0.546)	1
	M	28.92	176 of 472	45.13	324.5	10.17	20.05	53 of 472
	SD	12.31		9.54	287.2	11.37	7.37	

Note: p-value of the pairwise correlations shown in parenthesis

The mean value and standard deviation of the behavioral change variables (8-11) are not 0 and 1 respectively, because we use only the data of the participants in the control condition to standardize the “NASA survival” problem data. The behavior of the other condition changed; thus, the mean and the standard deviation do as well.

Table A.5: Robust linear regression of performance and behavioral change of Study 2

	Dependent variable:			
	Total time (1)	**Time per move** (2)	**# of moves** (3)	**Performance** (4)
Framing-focused	0.227 0.094, 0.361)	0.027 (-0.095, 0.149)	0.334 (0.154, 0.513)	-0.040 (-0.329, 0.248)
Implementation-focused	0.143 (0.011, 0.275)	0.181 (0.060, 0.301)	0.029 (-0.149, 0.206)	-0.023 (-0.309, 0.263)
Gender	-0.126 (-0.232, -0.020)	-0.025 (-0.122, 0.072)	0.015 (-0.128, 0.157)	-0.012 (-0.241, 0.217)
Age	0.005 (-0.002, 0.011)	0.005 (-0.0002, 0.011)	-0.007 (-0.016, 0.001)	0.013 (-0.001, 0.026)
Postgraduate	-0.048 (-0.161, 0.065)	-0.005 (-0.108, 0.098)	-0.100 (-0.252, 0.051)	0.128 (-0.115, 0.372)
Reader	-0.034 (-0.144, 0.075)	-0.138 (-0.238, -0.038)	0.087 (-0.060, 0.234)	-0.139 (-0.376, 0.098)
Constant	-0.146 (-0.381, 0.088)	-0.169 (-0.383, 0.046)	0.203 (-0.112, 0.519)	-0.450 (-0.957, 0.057)
Observations	472	472	472	472
Residual Std. Erro	0.491 (df=465)	0.465 (df=465)	0.679 (df=465)	1.191 (df=465)
Note:		95% confidence intervals shown in parenthesis		

10.5 Replication studies

As part of the first revision to the paper, we ran two new experiments. Both experiments were preregistered in the Open Science Framework and publicly available. The first data collection is what we refer to as the pandemic pilot study, and it is available at: https://osf.io/a7sm5. The pandemic pilot study is a preregistered experiment where we introduced the pandemic control variables for the first time. We included the pandemic control variables at the end of the study and increased participants' payments above the minimum levels required by the hosting platform, Prolific.ac. In this study, we found that the participants behaved differently to Study 2, spending much more time than before on each task.

We learned from this mistake and, in the second preregistered experiment, we collected the pandemic control variables in a follow-up study. The second preregistration is an exact replication of Study 2, but with a larger sample size and a set of pandemic control variables collected in a follow-up survey to keep the replication as close as possible to the original experiment. The data collected for this preregistration is what we refer to as Study 3, the preregistration is available at: https://osf.io/nvfdc.

10.5.1 Pandemic pilot study

As part of the first revision of this paper, we preregistered an experiment where we aimed at testing whether the task order could have biased the results of Study 2. In doing this, we followed the same experimental design as in Study 2—that is, a three-condition by two-task design. The sample size was chosen to test effects of size 0.2 or larger, thus requiring 98 participants per condition, or 294 in total. The experiment employed counterbalancing, and thus we also had two different additional task orders, giving a total of six experimental blocks, each with 49 participants.

The experiment was held in late April 2020, in the midst of the COVID-19 pandemic. During the pandemic, survival is more present in people's minds. This is problematic for our study, as we aim at following the problem-solving of participants as they solve a task that is novel to them. In Study 2, we had excluded people who had survival training from the study, in order to ensure the novelty of the tasks. We kept that filter in this study, but given that the pandemic brings survival to mind, we also include a set of pandemic control variables. These control measures are taken from psychology, which has studied the different ways that the threat of disease affects people's decision-making (Taylor, 2019).

We employed six different pandemic control scales in total. Three of them are taken from Taylor (2019): the Monitoring/Blunting coping style measure (Miller, 1987; Steptoe, 1989), Perceived Vulnerability to Disease (Duncan et al., 2009), and the Intolerance to Uncertainty scale (Carleton et al., 2007). These scales gave us a comprehensive view of how the participants would react to a disease or threat. However, they did not control whether the participants see the COVID-19 pandemic as a threat or not. For that, we employed three other scales: first, the Fear of COVID-19 scale from Ahorsu (2020) and two sets of financial questions of our own making. One of the scales asks participants if they or someone else in their household has lost their job or a significant part of their income in the crisis resulting from the COVID-19 pandemic. The second set of questions asks whether the participants find the work they do in Prolific to be more important now, during the COVID-19 pandemic, than before. We asked this both in general, and also financially. Table A.6 summarizes the pandemic control variables added to the pandemic pilot study.

The combination of the six scales helps us obtain a more nuanced view of how COVID-19 has affected the participants. In all cases, the higher the scales, the more a participant has been affected, or the more they would react to perceived diseases or threats. Therefore, in using the six scales as control variables, we can account for part of the latent change in experimental conditions that COVID-19 brings to our data collection.

Table A.6 List of variables added in the replication studies

	Variables	**Definition**
Pandemic Controls	22. Fear of COVID-19	Measured responses to the Fear of COVID-19 scale by Ahorsu et al. (2020)
	23. Monitor/Blunt Scale	Abbreviated version of the Monitoring/Blunting scale for psychological coping with threats (Steptoe, 1989)
	24. Perceived Vulnerability to Disease	Perceived Vulnerability to Disease 15-item scale (Duncan et al., 2009)
	25. Intolerance to Uncertainty	Twelve-item Intolerance to Uncertainty Scale by Carleton et al. (2007)
	26. Job Loss	Two yes/no questions asking whether the participant or someone in their household has lost their job due to the COVID-19 crisis
	27. Prolific Importance	Two five-level Likert scale questions relating to the relative importance of the work in Prolific in general and financially for the participant

10.5.1.1 Changes to research design

The use of counterbalancing and the inclusion of the pandemic control variables are the two changes made in this pilot study when compared to Study 2. The use of counterbalanced design leads to a three-condition by two-task-order factorial design with six blocks. We follow Study 2 and use 0.2 effect size and 5% commission and omission error rates to determine the sample size. In total we require 49 participants per block, giving a total of 98 participants per condition and 294 participants in the full study.

The addition of the pandemic control variables increases the time required to finish the experiment. In Prolific.ac, the platform that provided the participants, participants were shown the list of open studies, the base payment available from the study, and the expected time to finish it. The inclusion of the pandemic control variables led to an increase in the stated duration of the study. In Study 2, we told participants that the study would take 30 minutes. In the pandemic control pilot study, we told participants that they would require 45 minutes to finish. This increase in time led to a change in the participants' behavior: they not only spent longer on the entire study, but allocated more time to solve the "winter survival" and "NASA survival" problems. We explain these differences in the next section.

10.5.1.2 Comparison of Study 2 and Pandemic Pilot

In this study we give participants performance-based incentives in the form of a bonus. The incentive scheme is the same as in Study 2, and top performers can earn up to twice as much as the lower-performing participants (Hertwig & Ortmann, 2001). This performance-based payment is calculated from the solutions given to the "NASA survival" and "winter survival" problems. These tasks are shown at the beginning of the study, and participants know that only these two tasks affect their performance.

In the pandemic pilot study, we find that the baseline behavior of the participants changed when compared when compared to Study 2. Table A.7 shows the mean and standard deviation of the *total time, time per move, number of moves,* and *performance* of all participants in the pandemic pilot study—a total of eight variables[9].

[9] Note that the table mixes the order in which participants see the tasks. The large deviations are also present if we show instead the only the baseline of people who saw each task

Table A.8 shows the comparison of the behavior of the participants in Study 2 and the pandemic pilot study. We calculate the t-test on each of the eight variables in Table A.7 between all the participants in Study 2 (N=472) and all the participants of the pandemic control study (N=294). We find that the *total time* spent on each problem by participants in the pandemic pilot study is much higher than the time spent in Study 2. The same is true for the *time per move*, although only marginally in the “winter survival” problem. Finally, the *number of moves* in the “winter survival” problem is higher in the pandemic pilot study than in Study 2. In general, participants spent more effort in the pandemic pilot study than in the Study 2.

We argue that participants expended more effort because we told them that they would require 50% longer to finish the experiment, and because they know that their payment depends only on their performance in the “winter survival” and “NASA survival” problems. We changed this in Study 3, the second preregistered experiment we performed during the review process.

Table A.7: Descriptive statistics for the pandemic pilot study

Task		Total time	Time per move	# of moves	Performance
“Winter survival” problem	M	384.29	11.40	21.38	43.92
	SD	229.33	6.78	9.21	11.94
“NASA survival” problem	M	464.78	11.01	29.08	44.49
	SD	361.61	9.68	11.34	15.42

Table A.8: Comparison of main variables of Study 2 and the pandemic pilot study

Task		Total time	Time per move	# of moves	Performance
“Winter survival” problem	p-value	0.002	0.061	0.037	0.141
	t-statistic	3.181	1.874	2.091	-1.473
“NASA survival” problem	p-value	0.000	0.001	0.852	0.000
	t-statistic	3.959	3.448	0.186	-4.227

Note: A positive t-value implies that the pandemic pilot study had a higher value than Study 2.

10.5.1.3 Task-order bias

The mixed-factorial design of this experiment, as well as Studies 2 and 3 (presented below), should not be affected by standard task-order bias. In our studies, we use the participants as their own baselines, given that they perform two tasks. In a simple within-subject design experiment, this would lead to a task-order bias. This does not happen in the mixed factorial design, because the average treatment effect is calculated between experimental conditions (i.e. between participants) and thus the task-order effect is canceled out. The effect is canceled out because we compare between experimental conditions, and the task order can be assumed to affect all conditions the same.

In the pandemic pilot study we employed a counterbalance research design to test the assumption that our results were not biased by the order in which the participants solved the tasks. We find support for this expectation by failing to reject the null-hypotheses of zero difference between the behavior of participants if they solve the “NASA survival” problem first or the “winter survival” problem first. The results are shown in Table A.9. This table presents the four standardized change variables, the main variables of this study, and compares the mean of participants who solve the “winter survival” problem first with the mean of the participants who solve the “NASA survival” problem first. The four t-tests are non-significant, the lowest p-value is 0.206 with the highest absolute t-statistic being 1.27.

The addition of the pandemic control variables increased the effort participants made to solve the tasks during the pandemic pilot study. This deviation makes the study a failed replication of Study 2. However, the pandemic pilot study did allow us to test our expectation that the mixed factorial design should not be biased by the order of the tasks. Having found support for this expectation, we proceeded to remove the counterbalancing of task order in Study 3 (shown below).

Table A.9: Effect of task order

Variable	Mean value		p-value	t-statistic
	Winter then NASA	NASA then Winter		
Std. Change in Total Time	0.020	-0.073	0.377	0.885
Std. Change in Time per move	-0.053	-0.088	0.763	-0.302
Std. Change in Number of Moves	0.060	-0.089	0.189	1.317
Std. Change in Performance	0.083	-0.043	0.356	0.924

Note: A positive t-value implies that the participants who solved the "winter survival" problem first had a higher value than the participants who solved the NASA survival problem first.

10.5.2 Study 3

From the pandemic pilot study, we learned that in order to replicate Study 2, we needed to keep the instrument used for collecting the responses as close as possible to the one used in Study 2. To do this in Study 3, we used the exact survey employed in Study 2, with the same remuneration scheme, variable definition, and data analysis scripts. The direct replication allowed us to gain trust in the validity of the results, and in the way the variables were defined and hypotheses tested in Study 2.

10.5.2.1 Changes to the research design

To guarantee that the data collection in Study 3 was a similar as possible to Study 2, we used the same survey as in Study 2. However, several days after the experiment was finished and all participants had given their responses, we opened up a new study in which only respondents of Study 3 were allowed to participate. This second study included the pandemic control variables. There was a small dropout rate, as some participants did not respond to both surveys, but over 92.5% of participants responded to the second survey.

In addition to using a second survey, we employed a larger sample in Study 3 than in Study 2. We did this because ex-post effect-size analysis in Study 2 showed us that the average effect size of the results on *total time, time per move,* and *number of moves* in Study 2 (Table 9) was 0.127. To calculate the effect sizes we used the G*Power 3.1 application—the same application we used for estimating the sample size required for designing a statistically sound experiment (Faul, Erdfelder, Buchner, & Lang, 2009). To account for the smaller effect sizes, we estimated the sample size for Study 3 and found that an experiment with three conditions and two tasks that can detect an effect size of at least 0.125 with 5% omission and commission errors, would require 747 participants, or 249 per condition. We use this value in Study 3 to replicate Study 2 in a sound manner.

10.5.2.2 Comparison of Study 2 and Study 3

We discuss this comparison in more detail in the main paper. In sum, Study 3 differs slightly from Study 2, but the deviations are less than half the size of what the experiment is meant to measure, and thus we argue that Study 3 is a good replication of Study 2.

Table A.10 shows descriptive values and first-order correlations of the variables collected for Study 3.

10.5.2.3 Hypothesis Testing

In the manuscript, we employ two regression tables to motivate the hypothesis testing (Tables 12 and 13). We present a larger regression in Table A.11. Additionally, we make reference to the robust regressions of each of the main regression tables (Tables A.12, A.13, and A.14).

Table A.10 Descriptive and first-order correlation table of the variables of Study 3 (part 1 of 3)

	Variables	1.	2.	3.	4.	5.	6.	7.
Experimental condition	1. Control condition	1						
	2. Framing focused cond.	-0.500 (0.000)	1					
	3. Implementation cond.	-0.500 (0.000)	-0.500 (0.000)	1				
Demographics	4. Gender	-0.006 (0.876)	-0.080 (0.028)	0.086 (0.019)	1			
	5. Age	-0.018 (0.629)	-0.050 (0.170)	0.068 (0.064)	0.014 (0.693)	1		
	6. Postgraduate	0.034 (0.355)	-0.014 (0.703)	-0.020 (0.586)	0.034 (0.348)	-0.037 (0.307)	1	
	7. Reader	-0.059 (0.107)	0.010 (0.795)	0.049 (0.177)	0.050 (0.174)	0.201 (0.000)	0.111 (0.002)	1
Behavioral change	8. Performance	-0.016 (0.660)	0.008 (0.827)	0.008 (0.825)	-0.023 (0.531)	0.043 (0.235)	0.043 (0.242)	-0.018 (0.619)
	9. Total Time	-0.143 (0.000)	0.131 (0.000)	0.013 (0.729)	-0.013 (0.716)	0.080 (0.028)	-0.003 (0.926)	0.031 (0.395)
	10. Time per Move	-0.072 (0.051)	0.036 (0.320)	0.035 (0.338)	0.007 (0.854)	0.046 (0.207)	-0.008 (0.828)	0.045 (0.218)
	11. # of Moves	-0.092 (0.012)	0.111 (0.002)	-0.019 (0.605)	-0.030 (0.411)	-0.004 (0.921)	-0.071 (0.054)	0.053 (0.149)
“NASA survival” exercise	12. Performance	-0.003 (0.927)	-0.035 (0.343)	0.038 (0.298)	-0.011 (0.773)	-0.041 (0.268)	-0.020 (0.576)	-0.061 (0.094)
	13. Total Time	-0.152 (0.000)	0.064 (0.078)	0.087 (0.017)	0.006 (0.877)	0.170 (0.000)	-0.023 (0.531)	0.084 (0.022)
	14. Time per Move	-0.101 (0.006)	-0.008 (0.824)	0.109 (0.003)	0.026 (0.472)	0.133 (0.000)	-0.029 (0.426)	0.104 (0.004)
	15. # of Steps	-0.037 (0.311)	0.092 (0.012)	-0.055 (0.136)	-0.044 (0.229)	0.007 (0.848)	-0.015 (0.672)	0.036 (0.329)
	16. Self-sel. *Impl. focus*	-0.049 (0.18)	-0.037 (0.308)	0.086 (0.018)	-0.001 (0.981)	0.027 (0.459)	0.006 (0.861)	-0.015 (0.691)
“Winter survival” exercise	17. Performance	0.015 (0.679)	-0.046 (0.209)	0.031 (0.400)	0.016 (0.671)	-0.093 (0.011)	-0.072 (0.051)	-0.044 (0.234)
	18. Total Time	-0.024 (0.511)	-0.088 (0.017)	0.112 (0.002)	0.027 (0.460)	0.140 (0.000)	-0.029 (0.424)	0.081 (0.028)
	19. Time per Step	-0.042 (0.253)	-0.066 (0.071)	0.108 (0.003)	0.029 (0.432)	0.128 (0.000)	-0.031 (0.395)	0.087 (0.018)
	20. # of Steps	0.070 (0.057)	-0.034 (0.353)	-0.036 (0.330)	-0.011 (0.763)	0.012 (0.748)	0.068 (0.065)	-0.025 (0.498)
	21. Self-sel. *Impl. focus*	-0.044 (0.227)	0.000 (1.000)	0.044 (0.227)	0.015 (0.690)	0.096 (0.009)	0.037 (0.315)	-0.010 (0.775)
Pandemic controls	22. Fear of COVID-19	-0.014 (0.705)	-0.009 (0.819)	0.023 (0.540)	0.243 (0.000)	-0.076 (0.043)	0.012 (0.752)	0.001 (0.985)
	23. Monitor/Blunt Scale	-0.051 (0.174)	0.022 (0.567)	0.030 (0.430)	0.028 (0.454)	-0.073 (0.052)	0.050 (0.183)	-0.007 (0.860)
	24. Perceived Vulnerability	0.057 (0.126)	-0.018 (0.624)	-0.039 (0.296)	-0.081 (0.031)	-0.123 (0.001)	-0.085 (0.023)	-0.032 (0.400)
	25. Intolerance to Uncertainty	-0.005 (0.897)	-0.028 (0.455)	0.033 (0.376)	0.126 (0.001)	-0.095 (0.011)	-0.063 (0.093)	-0.068 (0.070)
	26. Job Loss	0.025 (0.506)	-0.019 (0.605)	-0.006 (0.883)	0.045 (0.232)	-0.024 (0.519)	-0.028 (0.464)	0.028 (0.449)
	27. Prolific Importance	0.050 (0.185)	-0.040 (0.289)	-0.010 (0.792)	0.079 (0.036)	-0.056 (0.135)	-0.013 (0.722)	-0.040 (0.289)
	M **SD**	249 of 747	249 of 747	249 of 747	424 of 747	34.21 7.60	256 of 747	337 of 747

Note: p-value of the pairwise correlations shown in parenthesis

The mean value and standard deviation of the behavioral change variables (8-11) are not zero and one respectively, because we use only the data of the participants in the control condition to standardize the “NASA survival” problem data. The behavior of the other condition changed; thus, the mean and the standard deviation do as well.

Table A.10 Descriptive and first-order correlation table of the variables of Study 3 (continuation 2 of 4)

	Variables	8.	9.	10.	11.	12.	13.	14.
Behavioral change	8. Performance	1						
	9. Total Time	-0.087 (0.017)	1					
	10. Time per Move	-0.037 (0.316)	0.640 (0.000)	1				
	11. # of Moves	-0.136 (0.000)	0.255 (0.000)	-0.146 (0.000)	1			
"NASA survival" exercise	12. Performance	0.598 (0.000)	-0.076 (0.039)	-0.062 (0.091)	-0.049 (0.185)	1		
	13. Total Time	-0.056 (0.129)	0.761 (0.000)	0.468 (0.000)	0.147 (0.000)	-0.142 (0.000)	1	
	14. Time per Move	0.003 (0.926)	0.451 (0.000)	0.772 (0.000)	-0.132 (0.000)	-0.114 (0.002)	0.671 (0.000)	1
	15. # of Steps	-0.101 (0.006)	0.229 (0.000)	-0.141 (0.000)	0.612 (0.000)	-0.111 (0.002)	0.272 (0.000)	-0.160 (0.000)
	16. Self-sel. *Impl. focus*	0.021 (0.576)	0.005 (0.881)	0.025 (0.497)	0.039 (0.286)	0.039 (0.288)	-0.010 (0.781)	-0.003 (0.943)
"Winter survival" exercise	17. Performance	-0.532 (0.000)	0.022 (0.550)	-0.023 (0.535)	0.107 (0.003)	0.361 (0.000)	-0.086 (0.019)	-0.124 (0.001)
	18. Total Time	0.040 (0.273)	-0.273 (0.000)	-0.203 (0.000)	-0.139 (0.000)	-0.105 (0.004)	0.417 (0.000)	0.364 (0.000)
	19. Time per Step	0.060 (0.103)	-0.288 (0.000)	-0.349 (0.000)	0.022 (0.556)	-0.076 (0.038)	0.293 (0.000)	0.326 (0.000)
	20. # of Steps	0.055 (0.131)	-0.061 (0.098)	0.023 (0.529)	-0.539 (0.000)	-0.061 (0.096)	0.114 (0.002)	-0.013 (0.722)
	21. Self-sel. *Impl. focus*	0.034 (0.347)	-0.045 (0.224)	-0.051 (0.164)	-0.002 (0.951)	0.026 (0.485)	-0.016 (0.659)	-0.035 (0.334)
Pandemic controls	22. Fear of COVID-19	-0.015 (0.688)	0.002 (0.966)	-0.022 (0.553)	-0.002 (0.958)	0.079 (0.034)	-0.015 (0.688)	-0.044 (0.238)
	23. Monitor/Blunt Scale	0.026 (0.494)	-0.039 (0.302)	0.008 (0.829)	-0.015 (0.693)	0.013 (0.738)	-0.026 (0.492)	0.032 (0.398)
	24. Perceived Vulnerability	-0.020 (0.600)	-0.025 (0.504)	-0.012 (0.752)	0.015 (0.690)	0.020 (0.598)	-0.051 (0.171)	-0.017 (0.657)
	25. Intolerance to Uncertainty	-0.019 (0.618)	-0.005 (0.892)	0.052 (0.170)	-0.075 (0.044)	0.032 (0.399)	0.017 (0.657)	0.069 (0.068)
	26. Job Loss	-0.038 (0.317)	0.034 (0.371)	0.050 (0.183)	0.045 (0.236)	0.028 (0.451)	0.051 (0.173)	0.061 (0.105)
	27. Prolific Importance	0.004 (0.917)	0.022 (0.567)	-0.046 (0.220)	0.048 (0.201)	0.054 (0.153)	0.072 (0.055)	-0.002 (0.947)
	M	0.005	0.318	0.210	0.056	48.82	383.8	9.551
	SD	1.163	1.401	1.487	1.191	15.76	260.7	6.976

Note: p-value of the pairwise correlations shown in parenthesis.

The mean value and standard deviation of the behavioral change variables (8-11) are not 0 and 1 respectively, because we use only the data of the participants in the control condition to standardize the "NASA survival" problem data. The behavior of the other condition changed; thus, the mean and the standard deviation do as well.

Table A.10 Descriptive and first-order correlation table of the variables of Study 3 (continuation 3 of 4)

	Variables	15.	16.	17.	18.	19.	20.	21.
"NASA survival" exercise	15. # of Steps	1						
	16. Self-sel. *Impl. focus*	0.014 (0.7)	1					
"Winter survival" exercise	17. Performance	-0.001 (0.987)	0.017 (0.638)	1				
	18. Total Time	0.083 (0.024)	-0.023 (0.533)	-0.158 (0.000)	1			
	19. Time per Step	-0.026 (0.476)	-0.041 (0.264)	-0.150 (0.000)	0.839 (0.000)	1		
	20. # of Steps	0.336 (0.000)	-0.032 (0.389)	-0.129 (0.000)	0.254 (0.000)	-0.054 (0.144)	1	
	21. Self-sel. *Impl. focus*	0.058 (0.114)	0.061 (0.094)	-0.013 (0.721)	0.038 (0.295)	0.024 (0.519)	0.064 (0.079)	1
Pandemic controls	22. Fear of COVID-19	-0.010 (0.796)	-0.009 (0.812)	0.100 (0.008)	-0.026 (0.494)	-0.036 (0.34)	-0.008 (0.830)	-0.007 (0.860)
	23. Monitor/Blunt Scale	-0.023 (0.535)	-0.038 (0.307)	-0.017 (0.659)	0.017 (0.661)	0.038 (0.31)	-0.007 (0.847)	0.006 (0.873)
	24. Perceived Vulnerability	-0.071 (0.058)	-0.008 (0.825)	0.043 (0.248)	-0.043 (0.251)	-0.008 (0.830)	-0.094 (0.012)	-0.047 (0.212)
	25. Intolerance to Uncertainty	-0.072 (0.056)	-0.020 (0.589)	0.055 (0.146)	0.033 (0.376)	0.029 (0.446)	0.013 (0.721)	-0.005 (0.897)
	26. Job Loss	0.013 (0.721)	0.003 (0.932)	0.073 (0.052)	0.030 (0.419)	0.019 (0.618)	-0.039 (0.302)	0.017 (0.654)
	27. Prolific Importance	0.047 (0.212)	-0.028 (0.450)	0.051 (0.172)	0.080 (0.032)	0.069 (0.067)	-0.007 (0.85)	0.008 (0.825)
	M	28.19	274 of 7474	44.23	337.2	10.21	20.91	87 of 7474
	SD	10.38		11.98	233.7	8.25	7.04	

Table A.10 Descriptive and first-order correlation table of the variables of Study 3 (continuation 4 of 4)

	Variables	22.	23.	24.	25.	26.	27.
Pandemic controls	22. Fear of COVID-19	1					
	23. Monitor/Blunt Scale	0.088 (0.020)	1				
	24. Perceived Vulnerability to Disease	0.061 (0.106)	0.089 (0.018)	1			
	25. Intolerance to Uncertainty	0.374 (0.000)	0.086 (0.022)	0.086 (0.022)	1		
	26. Job Loss	0.106 (0.005)	0.102 (0.006)	0.020 (0.602)	0.075 (0.045)	1	
	27. Prolific Importance	0.163 (0.000)	0.125 (0.001)	0.024 (0.528)	0.188 (0.000)	0.309 (0.000)	1
	M	15.17	8.769	58.32	35.43	0.320	6.897
	SD	6.14	2.167	6.13	8.62	0.467	1.287

Note: p-value of the pairwise correlations shown in parenthesis.

Table A.11: OLS regressions of Study 3 with Demographic and Pandemic Controls

	Dependent variable:			
	Total time	**Time per move**	**# of moves**	**Performance**
	(1)	(2)	(3)	(4)
Framing condition	0.567 (0.317, 0.816)	0.243 (-0.020, 0.505)	0.318 (0.104, 0.533)	0.042 (-0.167, 0.251)
Implementation condition	0.321 (0.068, 0.574)	0.279 (0.013, 0.546)	0.110 (-0.107, 0.327)	0.024 (-0.187, 0.236)
Gender	-0.055 (-0.269, 0.159)	0.007 (-0.218, 0.233)	-0.069 (-0.253, 0.115)	-0.076 (-0.255, 0.103)
Age	0.013 (-0.0005, 0.027)	0.007 (-0.008, 0.021)	-0.003 (-0.015, 0.009)	0.009 (-0.003, 0.021)
Postgraduate	-0.014 (-0.232, 0.203)	-0.059 (-0.288, 0.170)	-0.209 (-0.396, -0.022)	0.132 (-0.050, 0.314)
Reader	-0.001 (-0.212, 0.210)	0.091 (-0.131, 0.313)	0.131 (-0.051, 0.312)	-0.100 (-0.276, 0.077)
Fear of COVID-19	0.002 (-0.017, 0.020)	-0.011 (-0.031, 0.008)	0.006 (-0.010, 0.022)	-0.0001 (-0.016, 0.015)
Monitor/Blunt Scale	-0.030 (-0.079, 0.018)	0.005 (-0.046, 0.055)	-0.012 (-0.053, 0.030)	0.016 (-0.024, 0.057)
Perceived Vulnerability	-0.002 (-0.019, 0.015)	-0.002 (-0.020, 0.016)	0.003 (-0.012, 0.018)	-0.002 (-0.017, 0.012)
Intolerance to Uncertainty	-0.0002 (-0.013, 0.013)	0.014 (-0.0001, 0.027)	-0.014 (-0.025, -0.003)	-0.001 (-0.012, 0.010)
Job Loss	0.107 (-0.124, 0.337)	0.223 (-0.020, 0.466)	0.080 (-0.118, 0.279)	-0.099 (-0.292, 0.094)
Prolific Importance	0.033 (-0.053, 0.118)	-0.079 (-0.168, 0.011)	0.058 (-0.015, 0.131)	0.018 (-0.053, 0.090)
Constant	-0.289 (-1.658, 1.080)	0.039 (-1.401, 1.479)	-0.041 (-1.216, 1.134)	-0.345 (-1.490, 0.800)
Observations	710	710	710	710
R2	0.037	0.023	0.035	0.011
Adjusted R2	0.02	0.006	0.019	-0.006
Residual Std. Error	1.385 (df = 697)	1.457 (df = 697)	1.189 (df = 697)	1.158 (df = 697)
F Statistic (df=12, 697)	2.212 (p = 0.010)	1.371 (p = 0.175)	2.13 (p = 0.014)	0.637 (p = n.s.)
Note:	95% confidence intervals shown in parenthesis			

Table A.12: Robust linear regressions of performance and behavioral change of Study 3

	Dependent variable:			
	Total time (1)	**Time per move** (2)	**# of moves** (3)	**Performance** (4)
Framing-condition	0.228	-0.009	0.199	0.050
	(0.108, 0.349)	(-0.128, 0.110)	(0.029, 0.369)	(-0.151, 0.250)
Implementation-condition	0.154	0.044	0.060	0.074
	(0.034, 0.275)	(-0.075, 0.163)	(-0.111, 0.230)	(-0.127, 0.275)
Gender	-0.057	0.018	-0.071	-0.100
	(-0.156, 0.043)	(-0.080, 0.117)	(-0.211, 0.070)	(-0.266, 0.066)
Age	0.009	0.008	-0.007	0.008
	(0.002, 0.016)	(0.002, 0.015)	(-0.016, 0.003)	(-0.003, 0.019)
Postgraduate	-0.020	0.033	-0.178	0.169
	(-0.124, 0.084)	(-0.071, 0.136)	(-0.325, -0.031)	(-0.004, 0.343)
Reader	0.139	0.057	0.132	-0.067
	(0.038, 0.241)	(-0.044, 0.157)	(-0.012, 0.275)	(-0.237, 0.102)
Constant	-0.257	-0.272	0.236	-0.271
	(-0.503, -0.011)	(-0.516, -0.028)	(-0.111, 0.584)	(-0.681, 0.140)
Observations	747	747	747	747
Residual Std. Error	0.612 (df = 740)	0.601 (df = 740)	0.837 (df = 740)	1.126 (df = 740)
Note:		95% confidence intervals shown in parenthesis		

Table A.13: Robust linear regressions of Study 3 including pandemic controls

	Dependent variable:			
	Total time (1)	**Time per move** (2)	**# of moves** (3)	**Performance** (4)
Framing condition	0.233 (0.111, 0.355)	0.001 (-0.121, 0.122)	0.232 (0.058, 0.406)	0.033 (-0.175, 0.240)
Implementation condition	0.172 (0.049, 0.296)	0.078 (-0.045, 0.201)	0.073 (-0.103, 0.249)	0.058 (-0.152, 0.268)
Fear of COVID-19	-0.007 (-0.016, 0.002)	-0.008 (-0.017, 0.0003)	-0.001 (-0.013, 0.012)	-0.002 (-0.017, 0.013)
Monitor/Blunt Scale	-0.013 (-0.037, 0.010)	0.002 (-0.022, 0.025)	-0.018 (-0.051, 0.016)	0.017 (-0.023, 0.057)
Perceived Vulnerability	0.001 (-0.007, 0.009)	0.002 (-0.006, 0.010)	0.0001 (-0.012, 0.012)	-0.004 (-0.018, 0.010)
Intolerance to Uncertainty	-0.004 (-0.011, 0.002)	0.001 (-0.005, 0.008)	-0.012 (-0.021, -0.003)	-0.005 (-0.016, 0.006)
Job Loss	0.088 (-0.025, 0.201)	0.029 (-0.084, 0.141)	0.129 (-0.032, 0.291)	-0.113 (-0.305, 0.079)
Prolific Importance	-0.004 (-0.046, 0.038)	-0.018 (-0.060, 0.023)	0.018 (-0.041, 0.078)	0.017 (-0.054, 0.088)
Constant	0.372 (-0.210, 0.954)	0.132 (-0.447, 0.712)	0.353 (-0.477, 1.184)	0.151 (-0.839, 1.140)
Observations	710	710	710	710
Residual Std. Error	0.632 (df = 701)	0.588 (df = 701)	0.874 (df = 701)	1.122 (df = 701)
Note:	95% confidence intervals shown in parenthesis			

Table A.14: Robust regressions of Study 3 with Demographic and Pandemic Controls

	Dependent variable:			
	Total time	**Time per move**	**# of moves**	**Performance**
	(1)	(2)	(3)	(4)
Framing condition	0.226 (0.106, 0.346)	-0.002 (-0.123, 0.120)	0.197 (0.022, 0.372)	0.043 (-0.164, 0.249)
Implementation condition	0.155 (0.033, 0.277)	0.069 (-0.054, 0.192)	0.057 (-0.120, 0.234)	0.071 (-0.139, 0.280)
Gender	-0.042 (-0.145, 0.061)	0.041 (-0.063, 0.145)	-0.057 (-0.207, 0.092)	-0.111 (-0.288, 0.066)
Age	0.008 (0.001, 0.014)	0.008 (0.001, 0.014)	-0.008 (-0.018, 0.002)	0.008 (-0.004, 0.019)
Postgraduate	-0.040 (-0.145, 0.065)	0.021 (-0.085, 0.127)	-0.194 (-0.346, -0.041)	0.168 (-0.012, 0.348)
Reader	0.124 (0.022, 0.226)	0.053 (-0.050, 0.156)	0.120 (-0.028, 0.267)	-0.076 (-0.251, 0.099)
Fear of COVID-19	-0.006 (-0.015, 0.003)	-0.009 (-0.018, 0.0002)	0.0005 (-0.012, 0.013)	0.0004 (-0.015, 0.016)
Monitor/Blunt Scale	-0.010 (-0.034, 0.013)	0.003 (-0.021, 0.026)	-0.016 (-0.049, 0.018)	0.017 (-0.023, 0.056)
Perceived Vulnerability	0.002 (-0.006, 0.010)	0.003 (-0.005, 0.012)	-0.003 (-0.015, 0.009)	-0.002 (-0.016, 0.012)
Intolerance to Uncertainty	-0.003 (-0.009, 0.003)	0.002 (-0.004, 0.008)	-0.013 (-0.022, -0.004)	-0.004 (-0.015, 0.007)
Job Loss	0.086 (-0.025, 0.197)	0.025 (-0.088, 0.137)	0.121 (-0.041, 0.282)	-0.098 (-0.289, 0.093)
Prolific Importance	-0.001 (-0.042, 0.040)	-0.015 (-0.056, 0.027)	0.017 (-0.043, 0.076)	0.020 (-0.051, 0.091)
Constant	-0.041 (-0.700, 0.617)	-0.300 (-0.965, 0.365)	0.860 (-0.097, 1.817)	-0.256 (-1.388, 0.876)
Observations	710	710	710	710
Residual Std. Error	0.606 (df = 697)	0.588 (df = 697)	0.871 (df = 697)	1.121 (df = 697)
Note:	95% confidence intervals shown in parenthesis			

10.6 Robustness checks: Are two clusters the norm in Study 1?

As a robustness check for the clustering procedure in Study 1, we reran the *pamk* clustering analysis in a specific way (Hennig, 2015). We removed one participant from the sample and clustered the remaining 47. We repeated this 48 times, removing each participant once and clustering the remaining participants together. The motivation for this is that when participants are removed, the medoids can change, and the partitioning-around-medoids clustering method can potentially determine that a different number of clusters are needed, that new medoids are found, or that participants should be clustered together in very different ways. By removing one participant at a time, we could observe how robust the results of the full-sample clustering were.

We removed each participant from the clustering once. This left us with 48 categorical clustering variables. After doing this for each participant, we realized two things. First, in 46 of the 48 cases, the partitioning-around-medoids method selected two groups as the best number for the analysis. Thus, the results we show in Study 1 were not a fluke, but the common number of clusters for our dataset. Of the two cases where more than two clusters were selected, we found that in one case, the protocols were separated into three groups. Interestingly, only one person (originally classified as *implementation-focused*) was placed alone in the third group. In the other case, the protocols were clustered into four groups. We found that the *framing-focused* group remained unchanged, but the *implementation-focused* group was separated into three clusters. Of the 20 participants of the *implementation-focused* group, 12 were selected into one group, seven into another, and the participant who was assigned their own one-person cluster before was selected into their own cluster once again. From this, we obtain our first finding—namely, that two clusters are most often the best way to separate the think-aloud protocols. Additionally, in the rare case when more clusters are required, the separation into *framing-focused* is stable, and the only thing that changes is how the *implementation-focused* cluster is defined—mainly due to a single protocol.

Our second finding goes down one level and studies the misclassifications. We find that misclassifications are uncommon, and occur with protocols that are on the border between the two strategies. To estimate this, we performed a second set of analyses on the 46 cases where the clustering gave two groups. In these cases, cluster assignment varied: some participants might be assigned to the *framing-focused* group on one occasion, and to the *implementation-focused group* on another. We call the cases where a participant was classified in a different cluster than in the 48-group cluster a misclassification. We estimated how frequent the misclassifications were, and discovered that of the 2160 (46x48-48) classification events, only 141 were cases where a protocol was classified differently than in the full sample. The protocols were classified correctly 93.4% of the time. Additionally, six participants accounted for 90% of the misclassifications (127 of 141); these participants were classified around 50% of the time as *framing-focused* and 50% of the time as *implementation-focused*. We learned that misclassifications are uncommon, happening around 6.6% of the time, and that they are over-represented among a few participants (six out of 48) who are at the border of being classified as *framing-focused* or *implementation-focused.* This finding, together with the finding that two clusters are a robust separation of the think-aloud protocols, led us to see the results of the full-sample clustering as robust— in terms of both the number of clusters, and the classifications of each participant to each cluster.

11. OPEN SCIENCE FRAMEWORK MATERIAL

We uploaded description of the tasks, and empirical strategies and protocols used to collect and analyze the data of all studies as a Open Science Framework repository (OSF). The material we uploaded included four written documents, mostly focused on Study 1, as well as dataset and data analysis scripts. In this section, I present the four written documents as uploaded at OSF (osf.io/eh5m2/?view_only=2bd6e1e7320548858fd872db4c658932).

11.1 The Leader of the Karabayos Study

The Leader of the Karabayos is an ill-structured and complex problem (Fernandes & Simon, 1999). The study we conducted using this problem was qualitative in nature. We employed think-aloud protocols to listen to the concurrent thoughts of experienced managers as they found a solution to the problem (Ericson & Simon, 1980). We had a large sample of participants in this study and we performed content analysis on the think-aloud protocols to understand the sequences of problem-solving phases each participant went through while crafting a solution to "The Leader of the Karabayos problem". We then took these sequences of actions and counted the number of transitions between distinct problem solving phases. We built a transition matrix, and compared the transition matrices of the participants using cluster analysis to let the major differences between the problem-solving sequences emerge. We found two clusters. These two clusters represented ways in which the managers crafted a solution to the problem, these two clusters when observed more in detail showed how the participants' attention focused guided how the participants solved the problem. This attention focus was intriguing for us and we then decided to create a way of manipulating it and testing how the shift of the attention focus affects the way participants solve strategic problems. The manipulation of attention focus and testing of its effects are done with the set of studies on the "Winter and NASA survival problems" that are also available in this repository and they provide a follow up to "The Leader of the Karabayos" first study.

For this study, we collected the responses of 48 managers with at least 5 years of experience. The data was collected by Daniella Laureiro-Martinez and a separate coding of the data has been used in a prior work with Stefano Brusoni (Laureiro-Martinez & Brusoni, 2018). The data of this study uses the same protocols but a completely different coding protocol. Thus the two studies are significantly different in terms of contents but also more data analysis methods.

In this repository, we include an explanation of how to collect think-aloud protocols and how to do content analysis in a way that follows our process. We provide the task and experimental protocols we used when collecting data, this includes a note log, checklist, and the exercises we did to train the participants to think-aloud. This is included in the Data Collection folder. In the Qualitative Analysis folder, we provide two fully coded protocols that are used in the manuscript as examples and show how the example sequences are built from the text up. The explanation should be sufficient to understand how our data was made. We are open to sharing the other 46 protocols but we would prefer to talk to whom may be interested first. We say this as the data is a bit more personal than behavioral data.

In the Quantitative Analysis folder, we include the data we used for clustering. That is the transition matrices of each participants, as well as the time they spent one each problem solving phase and a series of demographic control variables and cognitive scales. We then perform clustering and create the tables and analysis, and robustness checks presented in the paper. We do this in order to provide full transparency of how we analyze that data. The combination of the data collection, qualitative, and quantitative data analysis process should give a much clearer view to how "The Leader of the Karabayos" study was developed.

11.2 The Leader of the Karabayos Task

11.2.1 Study Protocol

The study included 49 Italian speaking participants. Each participant was given two hours to find a solution to the Karabayos problem, shown below. The participants spent significantly less than this maximum performing the task. The task was held in quiet rooms and a experimenter was with the participant at every moment. Before the task started the partcipant went through a series of exercises to exercise their thinking aloud. We provide a document detailing a set of exercises we employed, as well as a checklist and notetaking log we find useful when having participants think aloud[10]. These procedures are based in part on Ericsson (2003) and Fox, Ericsson, & Best (2011) but also based on the experience gathered while collecing data on ”The Leader of the Karabayos” task.

Our participant completed at least three training tasks and then started the Karabayos task. The task started with the participant receiving a printed version of the task and reading the problem aloud. After reading each participant would start thinking aloud. The experimenter would sit with the back to the participant and would only interrupt in case the participant stopped verbalizing its thoughts. With time the participant would find a solution and finish the task. Having completed the task, we asked the participants to restate their final answer, and then the debriefing started. The aim of the debriefing was for us to understand the general experience while solving the problem and to check whether the participants had experience with this kind of problem; none had any experience in a similar context. All participants appeared motivated while solving the problem and many reported that they had empathized with the role of the tribal leader and had given serious thought to how to solve the difficult problem they were confronted with.

[10] To avoid repetition, you can find the task description above in Figure A.1

11.3 Helping participants Think Aloud

Thinking aloud is not something everyone can do at a first try. For this reason, it is important to be prepared. We compiled a a checklist, a notetaking log and a set of exercises to give structure to the process of collecting a participant's thought process while they think aloud. This document introduces the three.

11.3.1 Exercises for getting a participant to Think ALOUD

Although several books and articles have been written about think-aloud protocols, the process is still complicated. It requires both a knowledgeable researcher and a participant willing to receive guidance, and go through exercises to think aloud in an open and honest way. The aim of the exercises is to guide the participants and provide feedback on their thinking aloud quality.

The first step is to provide participants with an explanatory text on the methodology. The text goes along these lines:

> *"We are interested in knowing your thoughts as you come up with the answers to the problems in this study. In order to do this, I am going to ask you to think aloud as you work on the problem. What I mean by "think aloud" is that I want you to verbalize your thoughts out loud right as they appear in your head, without filtering them and as comfortably as you can. Please don't plan or explain what you say. Just act as if you are alone and speaking to yourself. I will just listen and only intervene by saying, "Please remember to think aloud" if you are silent for a long time."*

After providing these instructions, a good first exercise is to ask the participant to think aloud about a moment when they were alone and packing their bags for a trip. Many participants say aloud what they needed to put in the bag, or ask themselves whether they had already packed something. This is a broad task that takes the participant by surprise, and sometimes they find it difficult. Other assignments that work well are thinking aloud while going through a grocery shopping list, or mentally rehearsing a recipe.

Next, comes the second exercise. We ask the participant to think aloud while they perform a mathematical operation. For example: *'Think aloud while you multiply 237 x 84'*. This task is very effective, as participants immediately switch from a 'telling' to a 'thinking' mode. For the researcher, this is very illustrative – first, because it fulfils the purpose of the training, and second, because it lets you listen to the pace and tone that characterize the participant's thinking. Choose some numbers that are demanding to handle mentally, but not overly so.

Another set of instructions and examples that works very well are the following (Inspired by the supplementary materials of Fox, Ericsson, and Best (2011) found in http://supp.apa.org/psycarticles/supplemental/a0021663/BUL-Fox20090266-RR-F6.zip):

> *"Please share the thoughts that you have as they occur to you.* ***Do not explain them to me.*** *Just verbalize what you are thinking, it does not matter if it doesn't seem grammatically correct or if you think that it doesn't make sense. It is important and useful that you just express what is on your head. Let's practice with one more example. Are you ready?*
> *What is the sixth letter after B?*
> *Participant thinks-aloud*
> *Thank you"*

Depending on the answer that was obtained, say or not say the following:

> *"Chances are that the letter "H" didn't immediately occur to you after hearing the question. You probably had to go through several steps to find the answer. Had you summarized your thinking during the last question rather than reporting the sequence of actual thoughts aloud, you might have said that you found the letter H by counting through the alphabet. But, when people solve this problem out loud, they usually say a sequence of individual letters, such as B, then C, D, E, F, and G, before the answer H. Because we are interested in knowing the thoughts you had as you answered the question, we wish to have the most accurate, detailed report of thoughts as possible, instead of a summary of those thoughts."*

If required, another alternative is to show participants a short video of somebody else thinking aloud. Demonstrating by concrete example gets the point across better than most abstract explanations; however, the risk is that participants might imitate the person in the video.

This is recommended only when instructions and exercises fail, or when there is no time or possibility to use them.

For the data to be valid, it is important that participants can think alone. Minimize your intrusion. Rather than sitting opposite the participant, you could sit next to them, so they feel more comfortable and less intimidated (Nunan, 1992). Take informal notes on the participant’s behavior and tone of voice, as well as any environmental events that might affect their behavior. Make sure you relate these observations to the participants’ progress throughout the task.

11.3.2 Think aloud session checklist and log

It is important to plan ahead for the data collection. In think-aloud protocols, We have compiled a checklist of aspects that we find important to have and go through when collecting a verbal protocol.

Similarly, while being with the participant is important to log information about the session. This information is useful for understanding potential issues or questions that emerge later in the data analysis process. The Checklist and the Notetaking Log are available at the end of this document.

11.3.2.1 CHECKLIST

Before starting make sure to have

- 2 recorders
- Extra batteries
- Pages for taking notes
- Page with instructions to help participants think aloud.
- Task page
- If desired: white page and pencil

Protocol to follow:

1. Warm-up: make sure the situation is comfortable, the room is ok, and the seat is really opposed to us
2. Instructions: highlight in particular (aloud, though it is in the text instructions):
   i. Talk as much as possible and take as much time as needed
   ii. Provide a rich answer as if you really were in the problem: details are very welcome!
   iii. Not only say **what** to do but also **how** to do it
   iv. Really ignore us, imagine we are not here, we'll take notes and be as silent as possible: the cell phones and devices are now off!
   v. In the desk there is one page and one pencil, feel free to use them if doing so **helps you talk** and develop your strategy
3. (after the participant stops talking): Thank and ask to wrap-up / summarize the strategy

11.3.2.2 THINK ALOUD LOG # ____________

Date Participant's name Place	Time starts reading case:____________ Time starts thinking aloud:____________ Time stops thinking aloud:____________ Number of training exercises:__________

Timing	What was said	Ideas for coding	My observations
		Assumptions?	

<table>
<tr><td>Was there…
noise or other conditions?
something to note on the participant’s mood?
any other particular factors</td><td>Notes on my mood…
How well did I explain the method?
How was the rapport?</td></tr>
<tr><td colspan="2">Total time thinking aloud:_____
My perception on the quality of the verbalization:____/4.
My perception on the quality of the solution:____/4.
How involved in the situation:____/4.
Other factors:</td></tr>
<tr><td colspan="2">Notes on debriefing and further reflection</td></tr>
</table>

11.4 Content and Sequence Analysis of the Karabayos Task

11.4.1 Qualitative Data analysis

In this document, we provide a succinct description of the qualitative data analysis performed after the 49 think-aloud protocols of the Karabayos task were collected. In this document, we provide two full transcripts, and we are open to share others upon request. The first step after recording the protocols was to transcribe them verbatim. Following this, a group of three rates coded each word of the protocols into seven different phases of problem solving. Table 1 presents a description of each of the phases and an example of translated quotes of each phase[11]

Table 1: Problem-solving phase coding definitions

Problem-solving phases	Description	Examples of verbalized thoughts (transcribed verbatim)
Frame stating (FS)	Repeating the data mentioned in the text of the problem	"…so our area want to be left alone we are vulnerable that we have understood for a good reason … *I* mean here *I do* no*t* have other information problems diseases a very small zone lack of food…"
Frame assuming (FA)	Development of hypotheses not mentioned in the problem	"… for millennia and before me, my father, my grandfather, and all the others one after the other without having to face things that were more difficult go hunting sometimes or collect fruit…"
Direction setting (DS)	Defining a general path of actions to be followed and generating proposals about what should be done	"… we can also be a means for, a means to attract, for your region, we can, we can make people, we can, we can help you make I do not know a museum something we can make lessons to teach city kids how to love the forest…"
Evaluation (EV)	Evaluating and judging the proposal and considered their strategy without evaluating specific details	"… sending two or three people can be interesting… even though most likely those two or three won't return…"
Decision (DE)	Making an explicit choice about what intended actions	"…however I will try to dialogue this for sure I will try three key points dialogue with another civilization support from my group and away and an alternative in case of failure of dialogue…"
Implementation (IM)	Designing a sequence of actions required to carry on the proposed actions	"…slow calm we arrive in front of a representative we try with presents with kids with women and with men with those most intelligent to craft a speech even with gestures drawing we ask for help and we see if they help if not we try alone we do not explain where we are because if we explain because if we have to try at least they don't know where we are… we return…"
Implementation evaluation (IE)	Evaluating the possible actions' outcomes	"…is clear that it is not easy because probably out the jungle a someone some member of my tribe will hardly survive but is an endeavor to try…" "…if the two people [that were sent away before] should not return however 46 people will still be alive if instead return with a positive answer we have solved at least for some time long enough the problem…"

[11] The protocols were collected and coded in Italian. Table 1 has translated versions of the quotes for simplicity of explanation.

The output of the coding of the seven phases gave us three documents that had each word coded into problem solving phases. The raters achieved a high agreement of 93.4% with a Cohen's Kappa of 0.51. However, to perform sequence analysis we require that every word is coded into only one phase of problem solving. For this reason, two of the co-authors of the study reviewed the cases were there was not even a majority of the raters agreeing in a code and assign it to the appropriate phase of problem solving. After this process we had a fully coded transcript. In the next sections, we show two of these transcripts. We color each of the phases according to the colors shown in Figure A.8. By doing this we highlight how complex the sequences of phases are in a single think aloud protocols.

The next step in the process is to collapse the length of each phase. In our study we follow the transitions between phases and for this reason we care more about the number of changes between one phase and the other and not about the amount of time spent before a transition is made. After collapsing, the phases, the transcripts become equivalent to the colored sequences in Figure 1[12]. The two people shown have very different problem solving processes, one uses frame stating and frame assuming significantly at the start; whereas the other employs implementation and implementation evaluation much earlier and often during problem solving.

To compare the sequences of each participant we create transition matrices. These matrices account for the number of transitions each agent did between each pair of phases and the percentage of time spent on each phase of problem solving. Tables A.15 and A.16 show the transition matrices for Person A and Person B. To analyze the data, we further normalize the transition matrices so that instead of having an exact number of transitions we keep only the percentage of transitions between phases. By doing this, we are able to have an intensive measure that does not vary with the extent of the protocol. These matrices are then used in the quantitative part of our data analysis. This data analysis is shown in a separate document.

[12] The figure that we reference here is Figure 1 in Section 3.3.3 Sequence Analysis

Table A.15: Transition Matrix of Person A Before Normalization

Person A	→ FS	→ FA	→ DS	→ EV	→ DE	→ IM	→ IE
FS →		3	0	0	0	0	0
FA →	1		3	3	1	1	0
DS →	1	2		18	1	8	1
EV →	1	2	22		3	1	0
DE →	0	0	0	3		3	1
IM →	0	1	5	4	2		2
IE →	0	0	1	2	0	1	
% of thinking time	10.80	11.10	7.92	21.00	21.30.	1.37	23.90

Table A.16: Transition Matrix of Person B Before Normalization

Person B	→ FS	→ FA	→ DS	→ EV	→ DE	→ IM	→ IE
FS →		11	7	2	2	1	0
FA →	11		2	4	0	0	0
DS →	3	2		13	0	2	0
EV →	9	3	8		2	2	0
DE →	0	0	0	4		0	0
IM →	0	0	2	1	0		1
IE →	0	0	1	0	0	0	
% of thinking time	5.70	29.9	14.70	15.60	31.30	2.13	6.20

11.4.2 Transcripts of representative participants

11.4.2.1 Transcript Person A

Per forza perché non abbiamo contatti con il mondo. sicuramente causata dall'ignoranza. ma sicuramente anche perché non lo conoscono. Se lo conoscessero ci sarebbe un po' più di curiosità. Allora innanzitutto sono il leader e questo già mi piace quindi posso posso decidere io cioè nel senso che la cosa più logica sarebbe cercare di capire il punto di vista degli altri cercare di capire gli altri esattamente cosa vogliono e in qualche maniera cercare di fare pace con loro però più che fare pace con loro è una questione di sottometerci a loro forse la cosa la cosa più semplice potrebbe essere quella di perché qua comunque siamo 22 più 26 siamo troppo pochi siamo senza niente se vogliamo sopravvivere dobbiamo sottostare alle regole degli altri e questa è sicuramente una cosa il problema però qua è anche quello delle malattie e vediamo dov'è che dice? Allora anche quando i disboscatori non li uccidono direttamente la popolazione viene falcidiata da numerose malattie quindi sostanzialmente il nostro problema non sono le malattie il nostro problema è il fatto che comunque ci invadono la terra e in qualche maniera noi dobbiamo scappare e oltre a questo la zona dove in qualche maniera noi non siamo protetti ci è stretta ci è piccola perché non ci basta più per vivere quindi tra virgolette siamo noi a dover violare le zone loro ---. Ok allora gli elementi in sostanza sono questi qua. Allora abbiamo problema malattia. Abbiamo il problema del fatto che la terra che abbiamo è il tuo gruppo vive in un'area che per anni è stata piena di piante da frutto e animali di ogni genere. Per tradizione cacciano con archi e cerbottane e le donne restano a casa a prendersi cura dei bambini. Sei consapevole del fatto che alcune parti dell'area in cui vivi sono zone dei bianchi e per anni la tua gente ha evitato di entrare in contatto con loro. Negli untimi anni ti sei reso conto che gli alberi non producono più tanti frutti come un tempo che molti degli animali che cacciavi sono spariti quindi comunque il nostro problema è la zona troppo strettae quindi non ci permette di vivere e quindi manca il cibo. Manca il cibo. Mi hanno detto che c'era un altro problema che ovviamente però dove c'è l'altro non mi ricordo che vediamo un po' allora e il problema ovviamente sono i bianchi allora volete essere lasciati in pace e per una buona ragione la storia dei contatti tra tribù indigene e resto del mondo è sempre stata particolarmente infelice. E questo lo sappiamo. Il contatto è quasi sempre un evento disastroso vivono secondo uno stile di vita sostanzialmente intatto. ok Le loro vite non contemplano l'uso di televisione. La maggioranza di queste tribù vive in luoghi nascosti dentro la foresta. Molte di queste zone nascoste tuttavia sono diventate via via più vicine alle zone sotto il controllo di produttori di gomma ok mettono a repentaglio la sopravvivenza delle tribù di indigeni. Anche quando i disboscatori non li uccidono direttamente il problema è il contatto con i bianchi perché i bianchi ci uccidono. Cosa che spesso succede la popolazione viene falcidiata entro un anno o due da numerose malattie. A parte che se effettivamente entro un anno o due il problema è grave perché non abbiamo più tempo e quindi bisogna trovare una soluzione in fretta. Ok allora se io dovessi mettermi in contatto con questi qua prima che mi uccidono e quindi anche devo ragionare come farlo sicuramente non so come capirci. Che lingua parlare. Vediamo di capire a gesti però probabilmente i miei gesti non è detto che siano uguali ai loro diciamo? diversi però io come faccio a saperlo? Se ho vissuto per 10000 anni fra virgolette dentro questa dentro questa zona quindi non ho avuto i contatti con altri. Allora quindi possiamo morire per malattie perché ci manca cibo o perché i bianchi ci fanno fuori. Allora negli ultimi anni ti sei reso conto che gli alberi non producono ti chiedi cosa fare per mantenere la tua tribù salva. bisogna capire anche qui le altre zone quanto sono quanto sono grandi e se ci possono essere altre zone dove i bianchi non ci sono per riuscire a spostarci e a vivere. È ovvio che non andiamo avanti così all'eterno però almeno riusciamo a vivere un po' di più. Quindi non è una soluzione definitiva per tutte le generazioni ma almeno tiriamo tiriamo avanti un po'. Allora vediamo un po' quanto è grande è la zona. La pace e l'armonia l'abbondanza delle tue terre che per millenni ti hanno permesso di vivere in equilibrio con la natura è costantemente messa in pericolo dall'avvicinamento della civiltà. Il tuo gruppo vive in un'area che per anni è stata piena di piante eccetera. Bé sicuramente uno dei modi è cercare di capire cosa c'è dall'altra parte cosa porta questa civiltà ma io ho dei contatti con loro? So che si esiste la televisione il microonde la macchina oppure qualsiasi genere di abitudini? --- *rilegge veloce alcuni pezzi ma non si capisce.* Il tuo gruppo vive in un'area che per anni è stata piena di piante da frutto e animali di ogni genere. Per tradizione. E sì i problemi sono sicuramente questi qua però non riesco a trovare così al volo una soluzione sempre che esista una soluzione cioè nel senso una soluzione ipotetica. Allora malattia e qua bisogna cercare di capire come come curarci come però anche lì come si fa a curarsi io non so assolutamente nulla non mi curo mai la zona è troppo stretta manca il cibo. Allora sicuramente nell'immediato bisogna cercare il cibo cioè nel senso se non mi dan da mangiare --- dobbiamo riuscire ad andare avanti ma però è una cosa una scena un po' da film del tipo il più grande e grosso Rambo vado lì sequestro due trafficanti di droga gli prendo le armi e cerco cerco di capire come funziona però ovviamente e non è la soluzione. Bé sì sicuramente l'unica soluzione che vedo e che appunto una tribù da 48 persone non ha alcun genere di forza quindi in qualche maniera si deve si deve adeguare però qui bisogna capire chi chi ci può aiutare di questi questi qua che stanno occupando la zona e quindi cercare di capire da chi andare con chi in qualche maniera stabilire il primo contatto per l'integrazione. Allora vediamo un po' Allora produttori di gomma di noi non gliene frega niente quindi sicuramente allora poi vediamo un po' quindi i disboscatori non sono sicuramente amici perché i disboscatori arrivano e ci ammazzano quindi no disboscatori. Allora devo cercare altri alleati o comunque alla fine alleati non sono perché ci sottometteranno però almeno noi riusciamo a sopravvivere però bisogna capire come possiamo attaccarci alle nostre abitudini. Immagina di essere il leader che vivi in una foresta. Ci sono un centinaio di tribù nel mondo che come la tua non sono mai entrate in contatto con gli altri popoli disperse nelle immense giungle eccetera. Avete tutti una caratteristica siete i popoli più vulnerabili al mondo e volete essere lasciati in pace. E però volete essere lasciati in pace lasciati in pace lasciati in pace . E per una buona ragione la storia dei contatti tra tribù indigene e resto del mondo è sempre stata particolarmente infelice. Bè sicuramente una cosa di cui bisogna rendere conto è che non saremo felici alla stessa maniera in cui lo eravamo prima quindi cerchiamo la felicità più o meno in un altro modo saremo sempre infelici perché non abbiamo non abbiamo chance. Allora niente disboscatori allora poi abbiamo i colonizzatori produttori di gomma e trafficanti di droga lo so e però come faccio a capire io chi sono chi sono gli uni chi sono gli altri Allora abbiamo i no colonizzatori e questo non c'era allora disboscatori no disboscatori abbiamo detto di no colonizzatori poi chi abbiamo ancora? Trafficanti di droga e abbiamo anche i produttori di gomma. Allora trafficanti di droga sicuramente no ci fanno fuori però io come faccio a capire chi sono i trafficanti di droga? I colonizzatori allora i colonizzatori e i produttori di gomma i produttori di gomma vengono lì per produrre la gomma ma perché vengono qui a produrre la gomma? Le loro vite non contemplano l'uso di. La maggioranza di queste tribù vive in luoghi. Molte di queste zone nascoste tuttavia stanno diventando via via più vicine alle zone sotto il controllo di produttori di gomma e ma non dice niente su perché perché vengono lì i produttori di gomma? Bé questo bisogna capire i produttori di gomma perché vengono lì e se noi gli diamo fastidio o meno e i colonizzatori. Allora i colonizzatori vengono lì per restarci e quindi ci saranno ci saranno sempre e noi in qualche maniera potremmo e però sì ho detto così perché conosco il mondo ma io non conosco il mondo non so come funzionano le cose quindi non è che questi arrivano qua e gli faccio io la casa però almeno sopravvivo I produttori di gomma sicuramente l'unico modo non so se è l'unico modo però una una delle cose che potrei cercare di fare è in qualche maniera mettermi in contatto con i con i colonizzatori perché tra virgolette è per me una tribù nuova cioè nel senso sono vicini di casa e con i vicini di casa io devo riuscire a convivere. Allora mi proteggeranno dai disboscatori? Boh forse mi faranno capire come scappare dai disboscatori perché sono quelli che mi ammazzano poi magari mi danno una mano per le varie malattie mi danno delle medicine a loro cosa gliene frega di me? Niente però magari uno di questi un minimo di cuore ce l' avrà. Sono in circolazione da millenni. Ok. La pace e l'armonia l'abbondanza delle terre ok. Però questo qua per cercare una soluzione diciamo più a lungo termine è cercare di farci gli amici i colonizzatori. I produttori di gomma forse non lo so però sì forse è più logico i colonizzatori però ripeto queste cose qua come faccio a saperle se non so chi è chi? Per me sono tutti uguali per me sono dei bianchi e alcuni mi ammazzano altri altri invece altri no però non li distinguo credo che per me siano tutti esattamente esattamente uguali poi non distinguo neanche il loro modo di fare di vestirsi non distinguo il fatto che possano guardarmi o meno o forse lo sto imparando forse comincio a capirci qualche cosa. Allora --- per anni piena di piante va bé per tradizione gli uomini cacciano con archi mentre le donne restano a casa. Sei consapevole del fatto allora questi dettagli che le donne che gli uomini cacciano con archi e cerbottane mentre le donne restano a casa è un dettaglio che mi può servire? Sei consapevole del fatto che alcune parti dell'area in cui vivi confinano con zone dei bianchi e per anni la tua gente ha evitato di entrare in contatto con loro. Ok quindi sicuramente l'idea chiave è entrare in contatto con loro entrare in contatto con loro allora sì bisogna Entrare in contatto con i bianchi e questo è sicuramente la cosa che va fatta perché oggi come oggi non possiamo più evitare perché altrimenti si muore di fame e questi qua ci fanno fuori a parte che se ci uccidono siamo 48 siamo già così pochi che e poi bisogna vedere anche quanto siamo vecchi 22 donne 26 uomini. Il fatto che siamo 22 donne e 26 uomini questo cosa dice? Bambini. Come informazioni 22 donne e 26 uomini è un'informazione cosa sappiamo ancora? Sappiamo che sono questi qua che vengono sappiamo che abbiamo problemi con le malattie sappiamo che i disboscatori ci uccidono sappiamo che è stata piena di piante e animali di ogni genere. Per tradizione gli uomini cacciano con archi mentre le donne che per anni è stata piena di piante da frutto e animali di ogni genere allora gli animali sono scappati. Gli animali sono scappati quindi uno dovrebbe andare via --- Sei consapevole del fatto che alcune parti dell'area in cui vivi confinano ok. Negli ultimi anni ti sei reso conto che gli alberi non producono più tanti frutti come un tempo e che molti degli animali che cacciavi sono spariti. entrare in contatto con i bianchi questo sicuramente è il modo per affrontare per affrontare la cosa. Come fare a entrare i contatto con i bianchi? Questo lo dobbiamo capire e ma queste cose qua sembra un film. Allora io non sono tanto un grande appassionato di queste storie storielle. Lo decido io. --- Allora la nostra zona essere lasciati in pace siamo vulnerabili questo l'abbiamo capito e per una buona ragione. La storia dei contatti tra tribù indigene e resto del mondo è sempre stata particolarmente infelice. Nel senso qua non ho altre informazioni. problemi malattie zona troppo stretta manca il cibo contatti con i bianchi che ci fanno fuori. Stare lontano dai disboscatori quindi sicuramente se vediamo che stanno tagliando gli alberi stiamo alla larga perché questi qua sono pericolosi. Trafficanti di droga che ci fanno i trafficanti di droga qui? Come facciamo a evitarli? Colonizzatori e i colonizzatori quindi teori-camente li riconosco perché vedo che stanno mettendo su una casa ma io cosa ne so che quella lì è una casa? Temporanea ma è una casa perché per me una casa dovrebbe stare sulle piante. Eh sì non lo so questo perché io per me una casa è su un albero non riesco a distinguere sicuramente queste cosa. Sì perché le informazioni sono queste. Allora il tuo gruppo vive in un'area che per anni è stata piena di piante e animali. Negli untimi anni ti sei reso conto che ci siamo resi conto --- Ti chiedi cosa fare. A cosa mi serve sapere che sono 22 donne e 26 uomini? Boh facciamo che abbiamo 4 uomini in più ci servono o non ci servono? Boh non lo so sacrifichiamo 4 uomini? Non sappiamo quanti sono i bambini il contatto è quasi sempre un evento disastroso vivono secondo uno stile di vita sostanzialmente intatto da più di 10.000. Le loro vite ok. Vive in luoghi nascosti dentro la foresta. Molte di queste zone nascoste tuttavia stanno diventando via via più vicine alle zone sotto il controllo. Sì sicuramente entrare in contatto è spostarsi un po' più lontano la zona non è così piccola. L'altra l'altra soluzione potrebbe essere anche spostarsi da questa zona migrare perché così nell'immediato io come faccio a sapere come entrare in contatto con i bianchi? Questi bianchi cosa mi fanno? C'è un bel rischio e sicuramente di loro ho paura perché mi ammazzano. Ragionando ragionandola invece ragionandola in quel in quel modo lì cioè da chi sta da chi vive lì non conosce il resto il resto del mondo forse la cosa più immediata che mi verrebbe perché io sono più vicino forse agli animali che alla civiltà mi sposto migro quindi questa è un'altra soluzione soluzione due è migrare. Migrare devo stare attento a attento a altre tribù attento a altri bianchi quindi devo evitarli bisogna vedere come conosco tutta la foresta però teoricamente la conosco meglio degli altrifino a quando non entro in contatto con con i bianchi in qualche maniera pacifica ma questo deve andarmi devo devo essere fortunato anche se il pensiero era quello deve andarmi di culo. Allora sto attento a altre tribù sto attento ai bianchi devo cercare zone con cibo. Zone con cibo. Il fatto che lì non mi non mi basti perché effettivamente io comunque vado in giro per la foresta a piedi e quindi io ho un raggio di azione sicuramente più limitato quindi se io mi sposto sposto anche la zona dove andavo dove vado a cercare il cibo. Quindi il contatto con i bianchi è un rischio però è la soluzione definitiva migrare è una soluzione temporanea la devo decidere io forse la soluzione sarebbe migrare almeno un po' perché così riusciamo a trovare cibo perché lì non ce la facciamo più a vivere perché io non so quanto tempo ci mettiamo ad entrare in contatto con i bianchi e a trovare in qualche maniera una pace --- eh però ovviamente non posso pensare di spostarci sempre perché comunque vedo che il problema c'è e grave peraltro. È vero che se vado a trattare con i bianchi ci sottometteranno e quindi ho trovato la soluzione per cui loro non ci ammazzano più però siamo sottomessi e invece noi siamo un popolo a cui piace essere lasciati in pace. Se migriamo invece scappiamo da loro tra virgolette possiamo trovare una pace però forse loro arriveranno ancora si avvicineranno ancora anche perché questi arrivano con tanti mezzi che sono molto più potenti delle nostre gambe quindi migrare sicuramente come una soluzione temporanea per riuscire a sopravvivere. Però bisogna capire se continuare a migrare oppure cercare di entrare in contatto. Continuare a migrare manteniamo la pace ma non abbiamo mai la pace perché continuiamo a migrare e invece entrare in contatto con i bianchi arriviamo ad avere un certo tipo di pace nel senso che comunque non ci non ci ammazzeranno più però ci sfrutteranno e quindi noi abituati a vivere in libertà chissà come facciamo. Poi di fatto siamo così pochi che dove vuoi che andiamo? Quindi direi andrei a spostare Allora quindi le donne restano a casa a prendersi cura dei bambini allora innanzitutto andiamo a evacuare le donne e i bambini. Dopo di che noi cerchiamo di stare uomini in mezzo --- in mezzo perché stanno con le donne e con i bambini comunque gli portano il cibo eccetera e dall'altro lato qua donne e bambini mentre dall'altro lato cercano di procurare il cibo e poi oltre a procurare il cibo che comunque è quello che ci serve per vivere giorno per giorno cercare di capire come trattare con i bianchi. io sono il leader di questo gruppo quindi sicuramente il fatto che gli altri vanno a procurare il cibo invece il mio compito è trattare con i bianchi. E quindi vado a parlarci io cerco di capire chi sono i bianchi con cui parlare e poi vado ad affrontare però sicuramente dobbiamo emigrare fino a quando non si trova una soluzione però la soluzione è sicuramente trovare in qualche maniera la pace la pace con i bianchi sperando che quelli che troviamo siano i colonizzatori perché colonizzatori lo so io qui e non io lì però questa qua è una questione di fortuna perché questi sicuramente son quelli che non hanno interesse ad ammazzarci. Loro vengono lì a colonizzare la zona perché hanno altri perché hanno altri interessi e comunque vengono vengono a stare lì quindi magari potrebbero sfruttarci però meglio che ci sfruttino piuttosto che ci ammazzino. vediamo un po' se questo può essere la soluzione. Allora le donne e i bambini via procuriamo gli portiamo da mangiare e quindi gli uomini ci pensano a questo e io vado a trattare con i bianchi non so come però poi lo capirò. Allora ricapitolando problemi malattie quindi ci spostiamo e la tamponiamo un po' però prima o poi ci arriverà magari gli uomini sono un pochino più forti altre parti zona troppo stretta manca il cibo ok l'abbiamo risolto perché un po' ci siamo accostati e il contatto con i bianchi è un problema era quello che in qualche maniera l'integrazione è la soluzione definitiva. Con chi questo deve andarci bene? Disboscatori li ho individuati perché comunque mi buttano giù gli alberi e ci ammazzano mentre gli altri non buttano giù gli alberi e non ci ammazzano quindi vado dagli altri. Forse dai disboscatori ho visto le armi i trafficanti di droga teoricamente sono armati e --- produttori di gomma sono lì a lavorare non ho capito perché ci sono i produttori di gomma colonizzatori. Entrare in contatto quindi soluzione uno è entrare in contatto soluzione due è migrare però devo stare attento alle altre tribù devo stare attento ai bianchi e devo cercare di andare con cibo però alla fine a questo punto evacuiamo le donne e i bambini così almeno riusciamo a tenerli più tranquilli riusciamo almeno loro sopravvivono. Gli uomini fanno due cose da un lato comunque rimangono in contatto con le donne e coi bambini perché gli procurano il cibo dall'altro lato vanno comunque in giro per cercare il cibo quindi dall'altra parte dove ci siamo spostati e io vado a parlare con l'uomo bianco. E spero di trovare l'uomo bianco giusto che non mi ammazzi e che in qualche maniera mi permetta di avere un minimo di rapporto un minimo di pace.

11.4.2.2 Transcript Person B

Allora quindi prima di tutto se mi collocassi nella mente del leader dei Karabayos teoricamente non saprei tutte queste cose quindi io teoricamente vedrei semplicemente delle cose sconosciute. Quindi certamente il mio processo decisionale sarebbe diverso da quello di una persona invece che ha una cultura completamente diversa che è la mia e che so esattamente che le cose probabilmente sono dei grossi pericoli. Quindi senza sapere che queste in realtà sono minacce senza sapere del fatto che altre tribù come la mia sono state distrutte decimate la prima cosa è cercare di fare un contatto con queste persone. Quindi io inizialmente non avrei questo timore non saprei inizialmente che questi sono dei nemici. Anche se probabilmente queste tribù vedono tutti gli estranei come nemici. Quindi la prima cosa che farei è nel momento in cui ti arrivano queste persone iniziare anche in modo probabilmente ostile secondo la tradizione che ci ha mantenuti in vita così di insomma di attaccarli. Mi renderei subito conto che questi sono più potenti di me quindi questi invasori hanno delle armi che io non riesco a contrastare e non riesco nemmeno a. Cioè e non so neanche prevedere che un contatto con loro vorrebbe dire probabilmente anche beccarsi una malattia che ci decima subito. Però mi renderei subito conto del fatto che queste persone sono. Rimarrei sempre ostile fino a che non riesco a stabilire un contatto ma rimarrei mi renderei subito conto che non posso combatterli con le mie cerbottane e le mie armi. Quindi probabilmente capito che è una minaccia, capito che questi sono molto pericolosi probabilmente cercherei di insomma dichiarare una specie di stato di guerra per quanto siamo una cinquantina di persone. E quindi. Bè innanzi tutto cercare di non farsi trovare a casa e non farsi trovare nelle nelle cose più vulnerabili. Quindi sempre pensando di avere una tecnologia zero e e non sapendo nulla del nemico che ho davanti probabilmente intanto cercherei di fare bo delle capanne alternative. Cioè riuscire a nasconderci nella foresta dove siamo sicuramente più più forti di loro. Quindi non vivere nelle nostre capanne che sono ben visibili tutte raccolte ma stare per esempio in piccoli rifugi in zone della foresta disperse. Quindi in modo tale da non farci beccare di notte arrivano lì ci bruciano tutte le nostre capanne ci ammazzano tutti in una botta sola. Quindi intanto mi organizzerei in questo modo. Portare via le cose di valore e cercare di essere [frase interrotta] di valore inteso come che servono per la sopravvivenza e disperderci organizzando una specie di rete di contatti in modo tale da stare lontani dal villaggio dove siamo più vulnerabili e stare comunque nascosti. Sicuramente il nostro vantaggio rispetto a questi è che noi conosciamo bene il territorio mentre loro hanno una tecnologia più alta ma ci stanno arrivando. Quindi bo cercare di portarli in una zona dove non abbiamo noi il vantaggio cioè diciamo più nel dentro la foresta. Però chiaramente loro finché non ci vedono possono andare avanti tranquillamente a a disboscare a fare le loro le loro porche cose sul nostro territorio. Una volta comunque [frase interrotta] cioè la prima cosa non cercherei di fare subito [frase interrotta] cioè dopo il primo contatto dove capisco che che non non vinceremo mai con le nostre tecniche standard mi allontanerei dal villaggio per mantenere le persone e cercherei soprattutto inizialmente di studiarli cioè di osservarli mentre fanno qualcosa cercando di capire un po' di più del loro stile di vita. Prima di attaccarli a viso aperto una volta capito appunto che sono nemici me ne starei lontano osservandoli cercando di capire le loro le loro abitudini. Quindi probabilmente capirei il potere non so delle loro armi il potere del delle loro delle loro macchine. E ne capirei un po' di più ecco. Cioè prima di tutto vorrei capire chi è l'invasore. Conoscere qualcosa di più di loro prima di fare mosse azzardate. Quindi una volta che ho le mie persone tutto sommato al sicuro e qualcuno che riesce a tenerli d'occhio cercherei prima di tutto di capire queste persone che che cosa sono. E probabilmente riuscirei a capire che essenzialmente capirei appunto le loro armi mentre li vedrei forse usarle e vedrei come come si muovono. Allora a questo punto potrei provare due strade. Uno è cercare di stabilire un contatto diretto quindi la lingua non sarà sicuramente facile riuscire a comunicare. Certamente non andrei lì con le mani in alto a cercare di parlare perché è molto rischioso.

Quindi un contatto potrebbe essere. Una è un contatto pacifico l'altro è cercare di catturarne qualcuno. Quindi nel momento in cui noi ci non ci facciamo vedere le persone lì non si rendono conto che c'è una tribù effettivamente nascosta nella foresta. Tenerli d'occhio e quando c'è qualcuno che si allontana. Quindi diciamo strategia uno è contatto amichevole adesso vediamo poi come la sviluppo. Strategia due invece è ho capito che sono dei nemici e cercherò in qualche modo di trattarli come tale. Quindi non vuol dire sterminarli ma vuol dire cercare di di non usare mezze misure. Quindi bo forse la prima mossa che cercherei di fare è catturarne qualcuno Sicuramente so che catturare qualche nemico è già ti mette in una posizione di vantaggio rispetto all'altro quindi. Uccidere qualcuno forse non non vale la pena fai soltanto scatenare una reazione esagerata. Catturare qualcuno vuol dire da un lato avere in mano qualcosa che potrebbe essere usato per mercanteggiare e dall'altro cercare anche di entrare in contatto con loro capire di più della loro cultura magari cercare di non dico di parlargli assieme imparando la lingua che sarebbe molto difficile però bo in qualche modo cercare di comunicare con col prigioniero. La strategia se vuoi di più aggressiva è ne catturiamo qualcuno uno e poi quando è lì lo trattiamo anche bene ma cerchiamo di di comunicare con lui. La strategia due che la sto abbandonando è andare no no non ha senso. La strategia due di comunicare direttamente non la vedo fattibile dai. Ora mi focalizzo. Una volta che li ho studiati ho capito come si muovono non appena qualcuno si muove nella foresta da solo e lo vedo vulnerabile cerchiamo di gli saltiamo su in 26 contro uno e lo e lo pigliamo. Ce lo portiamo via e lo portiamo in un posto dove gli altri non difficilmente ci raggiungano ecco quindi. Che per arrivarci dovrebbero veramente entrare nella foresta dove in realtà siamo superiori noi. E tratterei bene il prigioniero non lo non gli farei gran [frase interrotta] anzi magari poi lo potrei anche rilasciare dopo che ho capito qualcosa. Cercherei comunque di stabilire un contatto con questi cercare perlomeno di capire che cosa che cosa questo da lui che intenzioni hanno. Ovviamente se hanno intenzioni ostili non mi diranno che sono venuti per disboscare la foresta e per e per non so distruggere il nostro villaggio. Però ecco questa sarebbe la prima mossa poi vediamo la loro reazione. Contemporaneamente mi renderei conto del fatto che c'è un mondo là fuori che noi non abbiamo mai non abbiamo mai avuto contatti e che potrebbe essere minaccioso. Non potrei prendere mio figlio e mandarlo a studiare in un'università americana che forse sarebbe la cosa migliore per riuscire a capire meglio di queste cose però potrebbe essere un tentativo cioè il fatto di rendersi conto in realtà che i tempi son cambiati purtroppo non volenti o nolenti cioè diciamo questo qui è un passaggio difficile da fare eh perché è molto difficile che soltanto parlando con una persona catturata se vuoi mi riesca a rendere conto di qualche centinaio di anni di storia in cui popoli più grandi e piccole tribù come la mia sono state decimate. Però supponiamo che in qualche modo perché sono un capo perspicace capisco che i tempi sono cambiati e che in un modo o nell'altro la nostra vita così com'è non può essere più fattibile. Perché se io anche metti che è un gruppo di 20 persone io riesco a ucciderli tutti ne arriveranno altri e quindi capisco che rimanderei soltanto la fine. Allora a questo punto dovrei sicuramente cercare di conoscere meglio cosa c'è là fuori dalla foresta. In realtà con 26 uomini e 22 donne l'unica cioè la cosa più sicura è cercare di scappare. Cercare di rifugiarsi sempre più nella foresta e e vivere di un lento declino. Ok. Ci sarebbe anche un'altra cosa che potrei fare. Adesso dipende dai contatti che ho coi miei coi miei se vuoi vicini. Se io ho sempre combattuto con altre tribù come la mia magari a distanza di qualche decina di chilometri forse sarebbe il momento di iniziare a parlare anche con loro. Cioè cercare di fare un un non so come li chiameranno nella foresta però un summit tra altri capi tribù vicini per sapere se loro hanno qualche informazione in più sul nostro invasore. Quindi chiedergli se loro hanno già avuto contatti se sanno eventualmente di altri di altri popoli che cioè piccole tribù che sono state eventualmente sterminate o no. Quindi cercare in qualche modo di imparare dagli eventuali errori successi degli altri. Quindi se io fino a qualche mese fa vivo tranquillo nella mia foresta sapendo che ho dei vicini di tribù coi quali non andavo d'accordo ma che non ci siamo dati troppo fastidio adesso comincerei perlomeno a dire quelli sono molto più diversi da me che non i miei tradizionali nemici della foresta. E quindi cercherei di andare da loro e chiedergli perlomeno se loro sanno qualcosa in più di me e se hanno in mente di cioè non so se hanno già avuto notizie da altri altre cose. E quindi forse capirei che tutti quelli che fino ad adesso sono entrati in contatto sono finiti male. E quindi dovremmo sicuramente unirci. Adesso io non so quante altre tribù. Se siamo veramente l'ultima tribù rimasta nel raggio di mille chilometri ok siamo direi finiti e quindi cercherei per lo meno di vivere il più possibile e quindi fare gli erranti della foresta fino a che pian piano declineremo. Però se riesco a trovare diciamo degli alleati in questa cosa nella nostra tribù cioè nelle tribù vicine almeno fare un minimo di così di massa. Almeno allearsi insieme contro il nemico. Poi che cosa potremmo fare contro questa civiltà che arriva difficilmente potremmo da soli far qualche cosa. Però per lo meno la prima cosa che farei con una massa del genere è cercare di avvicinarmi di più alla cultura. Quindi bo cercare se tra di noi c'è qualcuno che per qualche ragione è in grado di comunicare. Quindi supponiamo di avere di essere riusciti a mettere insieme una decina di tribù con 500 persone. Non son tante però insieme agli altri capi ovviamente di decidere di di organizzarci insieme di unire le forze. E cercherei prima di tutto di entrare più velocemente in contatto con questa nuova cultura. Cioè capire meglio il nostro nemico. Capire.

Perché noi in questo momento vediamo soltanto dei degli degli avventurieri delle persone che sono qui puramente per disboscare e per e per guadagnarci qualcosa. Ma se riusciamo ad andare oltre, io questo non lo so, non lo so ancora. Pero io non lo so come capo tribù però come come Marco Bianchessi che sta facendo questa intervista lo so diciamo dovremmo una volta conosciuti i nostri nemici sapere anche che oltre alla faccia che vediamo cioè dell'avventuriero che vieni qui a disboscare c'è comunque dietro un una civiltà con la quale si può anche cercare di dialogare di intavolare un discorso. Quindi se io sapessi in quel momento che c'è un faccio parte di un'entità politica che si chiama Brasile che ha un leader che in qualche modo anche per lo meno a livello internazionale è pressato dal fatto che non può distruggere una civiltà così potrei se lo sapessi cercare di far sentire la mia voce. Quindi se sapessi che questa civiltà che in realtà sta cercando di distruggermi nella sua nella sua faccia che vedo dietro ha anche altre cose che potrebbero proteggermi cercherei di saltare oltre questa faccia e cercare di appunto far sentire la mia voce. Che potrebbe voler dire qualunque cosa. Potrebbe voler dire trovare un alleato Non so immagino ci siano delle missioni. E quindi cercare di saltare oltre gli avventurieri e andare a parlare con qualcuno che invece potrebbe essermi alleato e non nemico. Ora questa cosa qui io per il momento non la so. Non so che potrei trovare degli alleati in una missione o in un anche in un qualche ente internazionale di salvaguardia di popolazioni come la mia. Però ecco l'unico modo che avrei per saperlo è cercare di studiare meglio il mio nemico. Quindi io prima di averlo studiato non non posso saperlo. Però l'unica speranza che ho è che --- abbia dal mio punto di vista una vulnerabilità dietro e quindi conoscendolo riuscirei a capirlo. Come faccio a conoscere meglio il mio nemico? Allora abbiamo detto prima posso cercare di mettermi intanto in comunicazione quindi vorrebbe dire o trovare qualcuno nelle nostre tribù che ha mai parlato con queste persone conosca il portoghese che o qualcuno o che qualcuno di loro abbia un interprete quindi riuscire a almeno a comunicare. Ovviamente nel loro interesse non c'è quello di dirmi guarda che se tu ti appelli a un organismo internazionale ti salvano. Il loro interesse è di mettermi a tacere il più possibile. Quindi devo riuscire a bypassare queste questo questo scopo. Quindi riassumendo cosa faccio. Prima di tutto allora abbiamo detto li approccio come sono abituato ad approcciare un nemico quindi probabilmente gli tiro un po' di frecce e loro mi sparano. Capisco che loro sono più forti di me quindi mi ritiro e anzi cerco di essere il più vulnerabile il meno vulnerabile possibile. Quindi cambio stile di vita velocemente e mi ritiro nella foresta lasciando il villaggio incustodito quindi lasciando meno vulnerabilità quindi salvo le persone che ho. Apro un un contatto con tutti miei nemici tradizionali di fronte a questo nuovo nemico comune e cerco di imparare il più possibile da osservandoli, eventualmente cercando di entrare in contatto con singoli. Potrei anche non prenderlo catturarlo ma semplicemente diciamo ne prendo uno lo porto via lo tratto bene cerco di parlargli e lo rilascio subito. Quindi in modo totalmente non ostile. In modo tale da avere un un minimo di contatto. Potenzialmente cercare di anche andare oltre quindi uscire dalla foresta. E' chiaro che non è facile perché probabilmente fuori dalla foresta un qualunque qualcuno della mia tribù difficilmente sopravvivrebbe però è un tentativo da fare. Quindi cercare di mandare così una spia fuori dalla foresta per cercare di capire meglio chi è il mio nemico. E prima o poi scoprirei qual è il punto debole che in questo caso probabilmente sarebbe il fatto di far sentire la mia voce al di fuori della del della foresta. Perché allora l'unica soluzione vera per poter continuare col nostro stile di vita è quella di di farsi sentire a livello internazionale. Poi in realtà insomma è una speranza abbastanza vana perché se ci sono degli interessi molto grossi dietro probabilmente è molto più facile che qualcuno ci ci si dimentichi di noi e ci faccia sparire. Ecco nel momento in cui dovessi veramente capire che la situazione è disperata --- l'unica cosa per riuscire veramente a salvare le mie persone è quella di di scappare. Insomma non non penso che riusciamo a. Cioè non ha senso pensare di essere 50 persone che che cambiano il il mondo. Si c'è anche il fatto che comunque sia si anche nel caso nel caso migliore in cui dovessi riuscire a in qualche modo ottenere un bando per lo sfruttamento della mia foresta c'è il fatto che comunque sia la il nostro stile di vita era già minacciato prima che arrivassero i bianchi. Allora dunque io sto rileggendo bene il problema sono già consapevole del fatto che queste zone controllate dai bianchi sono zone a rischio e noi abbiamo sempre evitato di entrare in contatto con loro proprio per questo. Quindi diciamo non sono arrivati dal nulla li abbiamo già probabilmente sotto controllo. E comunque la foresta non produce più tanto gli animali non non non ce ne sono più. Quindi già comunque sia anche se questi non arrivano proprio ad invadere la nostra la nostra terra sicuramente il nostro stile di vita è destinato a a scomparire. Allora non mi illudo che potrei prendere le nostre tribù e convertirla a un a uno stile di vita occidentale. Ecco un'altra cosa che forse una volta capito bene chi è il nostro nemico potrei cercare di sfruttarli. Cioè una volta capito che loro sono più forti di me che loro hanno comunque una civiltà una tecnologia superiore e hanno potrebbero addirittura aiutarci a sopravvivere nel momento in cui la nostra stile di vita attuale non è più sostenibile quindi non abbiamo più frutta e animali da caccia. Cercherei a questo punto di anche di capire che cosa potremmo noi fare per questi in modo tale che sia qualcosa che a loro serve veramente. Quindi diciamo una volta assicurata la sopravvivenza fisica una volta assicurato il fatto che non siamo minacciati proprio di fisicamente di morte nel senso che non ci arrivano lì di notte e ci ammazzano tutti una volta stabilito che non non rischiamo questo pericolo imminente c'è il secondo il secondo momento in cui noi dobbiamo decidiamo come vivere come riuscire a campare in un mondo che in realtà sta cambiando e anche se ci lasciano nella nostra isola nelle nostre quattro alberi attorno frutta e animali non ci saranno più quindi noi non potremmo più comunque vivere. Quindi cioè una volta superato il pericolo iniziale del dello scontro diretto bisognerà cercare di mettere in piedi una microeconomia che permetta alla nostra nazione alla nostra tribù di vivere in un mondo diverso. Quindi adesso non so che cosa potrebbe essere. Lì dovrò guardarmi in caso a vedere se bo manufatti artigianali piuttosto che direttamente essere noi a sfruttare la foresta in un modo diverso. Quindi il fatto di non so. Una volta c'era il caucciù nelle nelle foreste quindi magari c'è qualche altra qualche altra risorsa che serve ai bianchi quindi potremmo essere noi una volta padroni della nostra del nostro territorio non tanto a cacciare e a raccogliere frutti ma a fare un qualcosa che sia di valore per per gli altri. Qui non è facile però cercare di riconvertire in qualche modo il nostro stile di vita dando qualcosa di valore per avere ovviamente in cambio tutti i frutti della civiltà. Quindi non non penso che diventeremo mai ingegneri scienziati medici e architetti ma cercare di sfruttare in un modo diverso il nostro territorio non per la caccia e la raccolta ma per qualcosa che più si adatti allo stile di vita che in ogni caso se vogliamo sopravvivere dovremo adottare. Quindi io vedo sicuramente ineluttabile il fatto che noi dobbiamo cambiare stile di vita. Diciamo lo vedo da occidentale purtroppo da da da capo tribù spero di rendermene conto al più presto. Una volta che mi rendo del fatto che è inutile cercare di mantenere lo stile precedente cercare di non tanto copiare ma vedere che cosa potrebbe essere di valore per il per i nuovi arrivati la mia tribù. Quindi bo mettere in piedi un'economia sostenibile basata sulle mie risorse che sono chiaramente poche. Però possono possono essere quelle cioè delle persone delle mie competenze cioè la capacità probabilmente di conoscere la foresta cercare qualcosa che e qui si potrebbe andare per tentativi che abbia un certo interesse nei miei nemici. Quindi di fatto far parte senza diventare appunto quello che non sono cercare di far entrare un po' nel nel giro del mio nemico. Adesso non so che cosa cosa potrebbe essere perché io entrerei in contatto con le persone che sono venute a cercare legno piuttosto che gomma cose bè trafficanti di droga non lo so. bè perché no per quello che mi interessa potrei essere io un corriere della droga per loro. Nel momento in cui so che è un business redditizio che può permettere alla mia tribù di sopravvivere potrei essere quello che che conoscendo la foresta apre delle nuove vie della droga piuttosto che che invece sfruttare sfruttare la foresta. Quindi da un punto di vista della della sopravvivenza della mia della mia tribù cercherei qualcosa che possa essere usata come vera merce di scambio con col nuovo col nuovo arrivato ecco. Bè difese immunitarie lo sappiamo questo è un problema nel senso che quasi sempre le tribù vengono distrutte proprio dalla malattie nuove che non sono che non. Quindi il fatto di mantenere comunque dei link dei contatti chiamiamoli commerciali con coi nuovi arrivati mantenendo però una separazione quindi evitando di entrare troppo in contatto. Si comunque riassumendo cercherei dopo la prima fase di avere di di sventare la minaccia di un proprio di un genocidio verso la mia tribù adattarmi in qualche modo dando sfruttando quello che io posso dare che adesso non so ancora che cosa sia però sfruttare al meglio le nostre risorse per per entrare diciamo nell'economia del del nuovo arrivato. E per far questo si la prima cosa da da fare è cercare subito di capire di studiare i nostri nemici. Quindi nel momento in cui io mi rendo conto che a loro interessa un qualcosa di particolare che io ho vado io a offrirglielo spontaneamente ma cercando di ottenere in cambio il più possibile. Che prima cosa il fatto di non essere invaso seconda cosa il fatto di avere tutte quelle cose che a noi servono. Che non son tante alla fine. Siamo abituati a vivere di poco andremo avanti a vivere di poco. Cibo probabilmente e utensili che otterremo da da loro. Quindi cerchiamo di impostare un'economia più di di scambio che non di cacciatori e raccoglitori. Che è l'unica possibilità per la tribù di sopravvivere. E poi vediamo. Quindi i punti così fondamentali sono. Capire chi ho di fronte trovare altri amici cioè ex nemici che potrebbero però diventare alleati in questa nuova nuova sfida. Capire il più presto possibile che le cose sono cambiate e capire soprattutto come stanno cambiando e cercare di adattarci a questo nuovo nuovo stato. Quando va tutto male nel momento in cui non trovassi gli alleati non riuscissi assolutamente a capire che cosa --- che cosa qual è il nemico, qual è la cosa che potrebbe permettermi di vivere in questo nuovo nuovo mondo a quel punto niente cercherei di guadagnare tempo il più tempo possibile. [35'46''] Cioè non non cercherei subito lo scontro perché avrei già subito capito che vorrebbe dire morire gloriosamente ma ma morire e quindi cercherei di così essenzialmente ritirarmi il più possibile, avere il meno contatti possibile con loro. Magari che è solo un modo per guadagnare tempo, è un modo di sopravvivere con la mia tribù fino a che non scoprirei qual è la la via. Quindi o sono fortunato e capisco subito qual è il modo nuovo di di vivere quindi riesco subito a installare un'economia che mi permette di di scambiare beni con il nuovo arrivato e quindi di vivere in qualche modo se non proprio da amici, almeno tollerandosi. O in alternativa se non ho la soluzione non non vado allo scontro ma --- mi ritiro guadagnando tempo insomma. Fino a che se sono fortunato riesco a capire piano piano continuando sempre a a cercare informazioni su su sui nuovi arrivati cercando di capire meglio chi ho di fronte e intanto sopravvivendo magari meno bene ma spostandosi in zone sempre più inaccessibili della foresta però almeno continuiamo a vivere. E se siamo sfortunati piano piano ci estingueremo se siamo invece fortunati riusciremo a trovare qualche soluzione. Più che soluzione qualche nuovo nuovo sistema nuova economia per per sopravvivere. Che sarà sicuramente un'economia che comunque contempla gli scambi con i nuovi arrivati. Cioè non non penso che l'isolamento sia una soluzione e nemmeno il l'affrontare a viso aperto convinto di poterli rimandare indietro ecco. Quindi visto che delle tre possibilità. Ci sono tre possibilità in assoluto. Il instaurare un rapporto con questi nuovi evitare totalmente i contatti quindi ritirandosi oppure arrivare allo scontro diretto. Allora l'ultimo scontro diretto vuol dire farla finita alla svelta. il ritirarsi vuol dire probabilmente declino lento ma ma definitivo. Il terzo è più una scommessa. Cioè il fatto di cercare di instaurare un nuovo modo di di vivere. Non è assolutamente la certezza totale ma tra le tre è quella che mi da qualche speranza di poter portare alla fine di far sopravvivere la mia tribù. Quindi niente bisognerà mettersi sicuramente dover rinunciare a gran parte delle nostre abitudini. Quindi sicuramente sarà uno sforzo il fatto di dover reinventarsi. Per reinventarsi dovremo ripeto ancora capire bene il nuovo mondo. [si stende sulla sedia e si gira verso di me] Quindi prima di tutto capire chi abbiamo di fronte cercare di capire che il mondo degli altri è basato sugli scambi è basato su sul dare avere cosa che probabilmente nel nostro mondo invece non lo è e una volta capito questo cercare qualcosa del quale noi siamo ricchi e per poter mercanteggiare. E diventare non so per esempio potremmo scoprire che le nostre collanine gli piacciono tantissimo è vero vivremo di collanine però per lo meno entriamo nell'economia e cominciamo a mercanteggiare. Poi magari inventiamo qualche qualche cos'altro che può essere utile. Però ecco l'unica via di salvezza della mia tribù è quella di intanto appunto difenderci subito e non farci

sterminare subito. Quindi innanzi tutto proteggerci a riccio scappare difenderci e poi cercare di insinuarsi di di di convivere con questa nuova realtà. Ecco. E poi pregare il totem. Che è il nostro. Alla fine ci sarà sempre qualche qualcuno che dirà preghiamo bruciamo qualche animale al tótem. E va bè non so se serve

REFERENCES

Ahorsu, D. K., Lin, C. Y., Imani, V., Saffari, M., Griffiths, M. D., & Pakpour, A. H. (2020). The fear of COVID-19 scale: development and initial validation. *International journal of mental health and addiction*.

Anderson, V. L., & McLean, R. A. (2018). *Design of experiments: a realistic approach*. Routledge.

Andrews, K. R. (1971). *The concept of corporate strategy*. New York.

Baer, M., Dirks, K. T., & Nickerson, J. A. (2013). Microfoundations of strategic problem formulation. *Strategic Management Journal*, 34(2): 197-214.

Baguley, T. (2004). Understanding statistical power in the context of applied research. *Applied ergonomics*, 35(2), 73-80.

Baker, H. E., & Paulson, S. K. (1995). *Experiential exercises in organization theory*. Prentice Hall.

Bansal, P., Kim, A., & Wood, M. O. (2018). Hidden in plain sight: The importance of scale in organizations' attention to issues. *Academy of Management Review*, 43(2), 217-241.

Billinger, S., Stieglitz, N., & Schumacher, T. R. (2013). Search on rugged landscapes: An experimental study. *Organization Science*, 25(1), 93-108.

Brightman, H. J. (1978). Differences in ill-structured problem-solving along the organizational hierarchy. *Decision Sciences*, 9(1): 1-18.

Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2019). A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science.*

Carleton, R. N., Norton, M. P. J., & Asmundson, G. J. (2007). Fearing the unknown: A short version of the Intolerance of Uncertainty Scale. *Journal of anxiety disorders*, 21(1), 105-117.

Castellaneta, F., & Zollo, M. (2014). The dimensions of experiential learning in the management of activity load. *Organization Science*, 26(1), 140-157.

Cho, T. S., & Hambrick, D. C. (2006). Attention as the mediator between top management team characteristics and strategic change: The case of airline deregulation. *Organization Science*, 17(4), 453-469.

Cohen, M. D., & Bacdayan, P. (1994). Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organization science*, 5(4), 554-568.

Crilly, D., & Sloan, P. (2014). Autonomy or control? Organizational architecture and corporate attention to stakeholders. *Organization Science*, *25*(2), 339-355.

Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. *Englewood Cliffs, NJ*, *2*(4), 169-187.

Dewey J. (1910). *How we think*. Boston. MA: DC Heath.

Dong, J., March, J. G., & Workiewicz, M. (2017). On organizing: an interview with James G. March. *Journal of Organization Design*, 6(1), 14.

Duncan, L. A., Schaller, M., & Park, J. H. (2009). Perceived vulnerability to disease: Development and validation of a 15-item self-report instrument. *Personality and Individual differences*, 47(6), 541-546.

Einstein A, Infeld L. (1938). *The Evolution of Physics.* Simon and Schuster: New York.

Elsbach, K. D., & Stigliani, I. (2018). Design thinking and organizational culture: A review and framework for future research. *Journal of Management*, 44(6), 2274-2306

Eisenhardt, K. M., & Bourgeois III, L. J. (1988). Politics of strategic decision making in high-velocity environments: Toward a midrange theory. *Academy of Management Journal*, 31(4), 737-770.

Ericsson, K. A. (2003). Valid and non-reactive verbalization of thoughts during performance of tasks towards a solution to the central problems of introspection as a source of scientific data. *Journal of Consciousness Studies*, 10(9-10): 1-18.
Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, 87(3): 215.
Ericsson, K. A., & Simon, H. A. (1998). How to study thinking in everyday life: Contrasting think-aloud protocols with descriptions and explanations of thinking. *Mind, Culture, and Activity*, 5(3): 178-186.
Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G* Power 3.1: Tests for correlation and regression analyses. *Behavior research methods*, *41*(4), 1149-1160.
Felin, T., & Foss, N. J. (2005). Strategic organization: A field in search of micro-foundations. Strategic Organization. 3(4), 441-455
Felin, T., & Zenger, T. R. (2016). CROSSROADS—Strategy, Problems, and a Theory for the Firm. *Organization Science*, 27(1), 222-231.
Fernandes, R., & Simon, H. A. (1999). A study of how individuals solve complex and ill-structured problems. *Policy Sciences*, 32(3): 225-245.
Fox, M. C., Ericsson, K. A., & Best, R. (2011). Do procedures for verbal reporting of thinking have to be reactive? A meta-analysis and recommendations for best reporting methods. *Psychological Bulletin*, 137(2), 316.
Frankenberger, K., & Sauer, R. (2019). Cognitive antecedents of business models: Exploring the link between attention and business model design over time. *Long Range Planning*, *52*(3), 283-304.
Freeman, J. B. (2018). Doing psychological science by hand. *Current Directions in Psychological Science*, 1-6. Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological science*, *20*(10), 1183-1188.
Gavetti, G., Levinthal, D., & Ocasio, W. (2007). Perspective—Neo-Carnegie: The Carnegie school's past, present, and reconstructing for the future. *Organization Science*, 18(3), 523-536.
Gavetti, G., Levinthal, D., & Rivkin, J. W. (2005). Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal*, 26(8): 691-712.
Ghemawat, P. (1991). *Commitment*. Simon and Schuster.
Gibbs, A. J., Dale, M. B., Kinns, H. R., & MacKenzie, H. G. (1971). The transition matrix method for comparing sequences; its use in describing and classifying proteins by their amino acid sequences. *Systematic Biology*, 20(4): 417-425.
Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4): 650.
Grant, E. R., & Spivey, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, *14*(5), 462-466.
Grégoire D. A., Barr P. S., Shepherd D. A. (2010). Cognitive processes of opportunity recognition: The role of structural alignment. *Organization Science* 21(2): 413-431.
Hall, J., & Watson, W. H. (1970). The effects of a normative intervention on group decision-making performance. *Human relations*, *23*(4), 299-317.
Hennig, C. (2015). *fpc: Flexible Procedures For Clustering*. R Package Version 2.1-6.
Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists?. *Behavioral and Brain Sciences*, 24(3), 383-403.
Isenberg D. J. (1986). Thinking and managing: A verbal protocol analysis of managerial problem-solving. *Academy of Management Journal* 29(4): 775-788.

Johnson, D. W., & Johnson, F. P. (1982). *Joining together* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall. 111-116Joseph, J., & Wilson, A. J. (2018). The growth of the firm: An attention-based view. *Strategic Management Journal*, 39(6), 1779-1800.

Joshi, M. P., Davis, E. B., Kathuria, R., & Weidner, C. K. (2005). Experiential learning process: Exploring teaching and learning of strategic management framework through the winter survival exercise. *Journal of Management Education*, *29*(5), 672-695.

Kaufman L, Rousseeuw P. J. (1990). Partitioning around medoids (program pam). *Finding groups in data: an introduction to cluster analysis*: 68-125.

Klein, G. (1997). The recognition-primed decision (RPD) model: *Looking back, looking forward. Naturalistic decision making*: 285-292.

Klingebiel, R., & De Meyer, A. (2013). Becoming aware of the unknown: decision making during the implementation of a strategic initiative. *Organization Science*, 24(1): 133-153.

Krippendorff K. (2012). *Content analysis: An introduction to its methodology*. Sage.

Kuusela, H., & Paul, P. (2000). A comparison of concurrent and retrospective verbal protocol analysis. *The American Journal of Psychology*, 113: 3.

Lane, I. M., Mathews, R. C., Chaney, C. M., Effmeyer, R. C., Reber, R. A., & Teddlie, C. B. (1982). Making the goals of acceptance and quality explicit: Effects on group decisions. *Small Group Behavior*, 13(4), 542-554.

Langley, A., Mintzberg, H., Pitcher, P., Posada, E., & Saint-Macary, J. (1995). Opening up decision making: The view from the black stool. *Organization Science*, 6(3): 260-279.

Laureiro-Martínez , D. (2014). Cognitive control capabilities, routinization propensity, and decision-making performance. *Organization Science*, 25(4): 1111-1133.

Laureiro-Martínez, D., & Brusoni, S. (2018). Cognitive flexibility and adaptive decision-making: Evidence from a laboratory study of expert decision makers. *Strategic Management Journal*, 39(4), 1031-1058.

Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration–exploitation dilemma: An fMRI study of attention control and decision-making performance. *Strategic Management Journal*, 36(3), 319-338.

Levine, S. S., Bernard, M., & Nagel, R. (2017). Strategic intelligence: The cognitive capability to anticipate competitor behavior. Strategic Management Journal, 38(12), 2390-2423.

Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43(7), 934-950.

Levitt, B., & March, J. G. (1988). Organizational learning. *Annual review of sociology*, 14(1), 319-338.

Levy, O. (2005). The influence of top management team attention patterns on global strategic posture of firms. Journal of Organizational Behavior: *The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, 26(7), 797-819.

Li, Q., Maggitti, P. G., Smith, K. G., Tesluk, P. E., & Katila, R. (2013). Top management attention to innovation: The role of search selection and intensity in new product introductions. Academy of Management Journal, 56(3), 893-916.

Liker, J. K., & Morgan, J. M. (2006). The Toyota way in services: the case of lean product development. *Academy of management perspectives*, 20(2), 5-20.

Lipshitz, R., & Bar-Ilan, O. (1996). How problems are solved: Reconsidering the phase theorem. *Organizational Behavior and Human Decision Processes*, 65(1): 48-60.

Lombard, M., Snyder-Duch, J., & Bracken, C. C. (2002). Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human communica*

Mahalanobis, P. C. (1936). On the generalized distance in statistics. *Proceedings of National Institute of Science of India*. 12, 49-55

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1), 71-87.
March, J. G., Sproull, L. S., & Tamuz, M. (1991). Learning from samples of one or fewer. *Organization science*, 2(1), 1-13.
Maslach, D., Branzei, O., Rerup, C., & Zbaracki, M. J. (2018). Noise as signal in learning from rare events. *Organization science*, *29*(2), 225-246.
Maula, M. V., Keil, T., & Zahra, S. A. (2013). Top management's attention to discontinuous technological change: Corporate venture capital as an alert mechanism. *Organization Science*, 24(3), 926-947.
McLellan, E., MacQueen, K. M., & Neidig, J. L. (2003). Beyond the qualitative interview: Data preparation and transcription. *Field methods*, 15(1), 63-84.
Mintzberg, H., Raisinghani, D., & Theoret, A. (1976). The structure of "unstructured" decision processes. *Administrative Science Quarterly*, 21(2): 246-275.
Nelson, R., & Winter, S. (1982). *An Evolutionary Theory of Economic Change*, Harvard Univ. Pres, Cambridge.
Neuendorf, K. A. (2002). *The content analysis guidebook*. Sage.
Newell, A., & Simon, H. A. (1972). *Human problem solving* (Vol. 104, No. 9). Englewood Cliffs, NJ: Prentice-Hall.
Nickerson, J. A., & Zenger, T. R. (2004). A knowledge-based theory of the firm—The problem-solving perspective. *Organization Science*, 15(6): 617-632.
Nickerson, J. A., Silverman, B. S., & Zenger, T. R. (2007). The problem of creating and capturing value. *Strategic Organization*, 5(3), 211-225.
Ocasio, W. (1997). Towards an attention-based view of the firm. S*trategic Management Journal*, 18(S1), 187-206.
Ocasio, W. (2011). Attention to attention. *Organization Science*, 22(5), 1286-1296.
Ocasio, W., & Joseph, J. (2005). An attention-based theory of strategy formulation: Linking micro-and Macroperspectives in strategy processes. In *Strategy Process* (pp. 39-61). Emerald Group Publishing Limited.
Ocasio, W., & Joseph, J. (2018). The attention-based view of great strategies. *Strategy Science*, 3(1), 289-294.
Ocasio, W., Laamanen, T., & Vaara, E. (2018). Communication and attention dynamics: An attention-based view of strategic change. *Strategic Management Journal*, *39*(1), 155-167.
Oehlert, G. W. (2000). *A first course in design and analysis of experiments.* W. H. Freeman
Öllinger, M., Jones, G., Faber, A. H., & Knoblich, G. (2013). Cognitive mechanisms of insight: the role of heuristics and representational change in solving the eight-coin problem. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(3), 931.
Ormerod, T. C., MacGregor, J. N., & Chronicle, E. P. (2002). Dynamics and constraints in insight problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(4), 791.
Palan, S., & Schitter, C. (2018). Prolific. ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17, 22-27.
Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. J*ournal of Experimental Social Psychology*, 70, 153-163.
Poland, B. D. (1995). Transcription quality as an aspect of rigor in qualitative research. *Qualitative inquiry*, 1(3), 290-310.
Popper, K. R. (1963). Conjectures and Refutations. Routledge and Kegan Paul.

Posen, H. E., Keil, T., Kim, S., & Meissner, F. D. (2018). Renewing Research on Problemistic Search—A Review and Research Agenda. *Academy of Management Annals*, 12(1), 208-251.
Posen, H. E., & Levinthal, D. A. (2012). Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science*, 58(3), 587-601.
Puranam, P. (2018). *The microstructure of organizations*. Oxford University Press.
Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7): 545-556.
Rasmussen, J. L. (1988). Evaluating outlier identification tests: Mahalanobis D squared and Comrey Dk. *Multivariate Behavioral Research*, 23(2), 189-202.
Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4): 873-922.
Rerup, C. (2009). Attentional triangulation: Learning from unexpected rare crises. *Organization Science*, *20*(5), 876-893.
Reypens, C., & Levine, S. S. (2017). To grasp cognition in action, combine behavioral experiments with protocol analysis. In Methodological challenges and advances in managerial and organizational cognition (pp. 123-146). Emerald Publishing Limited.
Reypens, C., & Levine, S. S. (2018). Behavior in Behavioral Strategy: Capturing, Measuring, Analyzing', *Behavioral Strategy in Perspective* (Advances in Strategic Management, Volume 39).
Rouinfar, A., Agra, E., Larson, A. M., Rebello, N. S., & Loschky, L. C. (2014). Linking attentional processes and conceptual problem solving: visual cues facilitate the automaticity of extracting relevant information from diagrams. *Frontiers in psychology*, 5, 1094.
Salvato, C. (2009). Capabilities unveiled: The role of ordinary activities in the evolution of product development processes. *Organization Science,* 20(2): 384-409.
Sarasvathy, S. D., Simon H. A., & Lave L. (1998). Perceiving and managing business risks: Differences between entrepreneurs and bankers. *Journal of Economic Behavior & Organization* 33(2): 207-225.
Schacter, D. L., Benoit, R. G., & Szpunar, K. K. (2017). Episodic future thinking: Mechanisms and functions. *Current opinion in behavioral sciences*, 17, 41-50.
Schroeder, R. G., Linderman, K., Liedtke, C., & Choo, A. S. (2008). Six Sigma: Definition and underlying theory. *Journal of operations Management*, 26(4), 536-554.
Schwenk, C. R. (1985). The use of participant recollection in the modeling of organizational decision process. *Academy of Management Review*, 10(3): 496-503.
Simon, H. A. (1947). *Administrative behavior. A study of decision-making processes in administrative organization*. New York: Free Press
Simon, H. A. (1962). The architecture of complexity, *Proceedings of the American Philosophical Society*, 106(6), 467-482.
Simon, H. A. (1965). *Administrative behavior* (Vol. 4). New York: Free Press.
Steiner, G. A. (2010). *Strategic planning*. Simon and Schuster.
Steptoe, A. (1989). An abbreviated version of the Miller behavioral style scale. *British Journal of Clinical Psychology*, 28(2), 183-184.
Sullivan, B. N. (2010). Competition and beyond: Problems and attention allocation in the organizational rulemaking process. *Organization Science*, 21(2), 432-450.
Surroca, J., Prior, D., & Tribo Gine, J. A. (2016). Using panel data dea to measure CEOs' focus of attention: An application to the study of cognitive group membership and performance. *Strategic Management Journal*, 37(2), 370-388.
Taylor, S. (2019). *The Psychology of Pandemics: Preparing for the Next Global Outbreak of Infectious Disease*. Cambridge Scholars Publishing.

Tetlock, P. E., & Gardner, D. (2016). *Superforecasting: The art and science of prediction*. Random House.

Thomas, L. E., & Lleras, A. (2007). Moving eyes and moving thought: On the spatial compatibility between eye movements and cognition. *Psychonomic bulletin & review*, *14*(4), 663-668.

Tversky, A., & Kahneman, D. (1975). Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*. Springer Netherlands.

Wellian, E. (2010), *LEGO: Consolidating distribution*, Case Study, IMD International

Yadav, M. S., Prabhu, J. C., & Chandy, R. K. (2007). Managing the future: CEO attention and innovation outcomes. *Journal of marketing*, 71(4), 84-101.

Yetton, P., & Bottger, P. (1983). The relationships among group size, member ability, social decision schemes, and performance. *Organizational Behavior and Human Performance*, *32*(2), 145-159.

Yu, Z., Wang, F., Wang, D., & Bastin, M. (2012). Beyond reaction times: Incorporating mouse-tracking measures into the implicit association test to examine its underlying process. *Social Cognition*, 30(3), 289-306.

PAPER 2

"Sense may be in the eye of the beholder,
but beholders vote and the majority rules."
– Weick (1995:6)

More is Different: The Effect of Diversity of Preferences on Exploration

Jose P. Arrieta[1]
[1]ETH Zürich, Switzerland

ABSTRACT

Models of organizational learning under uncertainty tend to assume that organizations are composed of either one single agent or multiple agents whom all share their points of view – homogeneous preferences. In this paper, I relax this assumption and study how organizations composed of people with diverse views of the world – diverse preferences – learn from their uncertain environments. I find that *preference diversity* can lead to large and nontrivial changes in the exploration rate of organizations. *Preference diversity* is a double-edged sword. It can lead to both increased exploration rates but also decreased exploration. The key behind an increase or decrease in exploration is given by the specific *preference range, bias,* and *polarization* of the organization. A manager who wishes its organizations to explore more has two routes for achieving higher exploration through *preference diversity*. First, the manager can create measures that promote the agent's *preference polarization.* Conversely, the manager could incentivize polarization and create measures to control the *preference bias.* However, even small deviations on a single agent's preferences can erode the organization's exploration rate in polarized organizations. Polarized and biased organizations, explore much less than organizations whose agents share their preferences.

Keywords: Diversity, Preferences, Organizational Learning, Majority-voting

1. INTRODUCTION

In his landmark 1991 paper, Jim March explored the different ways in which "reason inhibits foolishness" and leads to the "vulnerability of exploration" (March, 1991:73). Since then, scholars have kept designing ways to help organizations explore more and become "ambidextrous" (Denrell & March, 2001; Greve 2007, March, 2010; O'Reilly & Tushman, 2011; Posen & Levinthal, 2012; Csaszar, 2013). With time, a robust knowledge base has developed around the exploration-exploitation dilemma. However, in contrast to March's work, the research in exploration-exploitation has built on the assumption that organizations are composed of either one agent who makes all decisions of many agents who behave "as if" all see the world in the same way (Friedman, 1953; March, 1962; Denrell, Fang & Levinthal, 2004; Laureiro-Martínez et al., 2015; Puranam & Swamy, 2016; Piezunka, Aggarwal, & Posen, 2020). An assumption that March himself claimed "is almost certainly wrong as a micro-description of [the] business firm" (1962:666). In this paper, I relax this assumption and study how organizations composed of people with diverse points of view learn under uncertainty and, specifically, how they explore their environments.

Diverse points of view can happen for three main reasons: differences in goals, beliefs, or mental representations (Cyert & March, 1963; Puranam & Swamy, 2016; Csaszar & Levinthal, 2016). If people disagree in at least one of these three reasons, their behavior will appear as if they follow conflicting goals. Given a shared problem, people with differences in their goals, beliefs, or mental representations will have different preferences over which actions to promote. They have what I call *diverse preferences*[1].

Evidence for the importance of *preference diversity* in organizational learning comes from qualitative studies. For example, Ramus, Vaccaro, and Brusoni (2017) explained how hybrid logics led organizations to formalization and collaboration cycles, where the different

[1] I use the term *preference diversity* instead of heterogeneity, to connect the literature of complex adaptive systems, upper echelon theory, and human resource management (Page, 2010; Boone & Hendriks, 2009; Shore et al., 2009).

points of view blend together into a cohesive framework. Similarly, Rerup and colleagues' work shows how diverse points of view help organizations triangulate their attention in reaction to crises, adapt the organization'points of view toward a common vision, balance conflicting goals while developing new products, and manage the conflicts inherent in learning from experience (Rerup, 2009; Rerup & Feldman, 2011; Salvato & Rerup, 2018; Rerup & Zbaracki, 2020). Similarly, Argote (2013:134) explained how when organizations face situations "without a demonstrably correct answer, a majority… decision scheme characterizes how groups make decisions". Kaplan (2008) showed how firms employ framing contests to choose their future strategy in uncertain environments and build a coalition to implement this vision. More generally, Gaba and Greve (2019) employed quantitative methods to show how firms manage safety and profit goals through political processes.

To study the effects of *preference diversity* in organizational learning under uncertainty, I employ agent-based models that simulate multi-agent organizational learning processes (Christensen & Knudsen, 2010; Posen & Levinthal, 2012; Puranam et al., 2015). The use of agent-based models allows me to ascribe each agent with static and distinct preferences that completely define its utility function (Adner et al. 2014)[2]. In turn, the organizations I study are defined by the preferences of the agents who compose it, one preference per agent.

The specific positioning of these preferences is the mechanism explored in this study. However, when studying organizations with multiple preferences, I find that "more is different", the behavior of the organizations cannot be described by one single measure of diversity as each organization is "different in its own way" (Anderson, 1972:393; Tolstoy, 1877:1). Yet, as long as the organizations are small, I can create summary measures that enable to compare organizations with one another (Puranam, 2018).

[2] The agents in an organization are all equal except for their preference over the attributes of the options. They follow the same updating process, and share their learning parameters. On every period all agents are given the same feedback. The only differences is that given that they have diverse preferences, each will value the feedback differently and thus their willingness to invest in the option in the future will vary.

I employ three measures to unbundle the preferences of an organization. *Preference range* accounts for how different the most distinct points of view are from each other. This measure distinguishes homogeneous organizations (zero *range*) from organizations with *diverse preferences* (nonzero *range*). *Preference bias* accounts for the preferences of the median voter. The median voter is pivotal for the decision-making of majority voting systems (Congleton, 2004; Holcombe, 2006). If the median voter is closer to one coalition in the organization, then this coalition's interests will gain a majority of votes. However, the median voter effect is limited by the polarization of points of view between the coalitions (Kamada & Kojima, 2014; Ganz, 2020). Organizations with high *preference polarization* act as "echo chambers," with the coalitions' views being almost orthogonal to one another and giving more power to the median agent (Baumann et al., 2020). I find that the specific *preference range, bias,* and *polarization* determine the exploration rate of organizations with diverse preferences.

Preference diversity can lead to large and nontrivial variations in the exploration rate of organizations. Managers can increase the exploration rate in two ways. First, the manager can create measures that promote the agent's *preference polarization.* Conversely, the manager could incentivize polarization and create measures to control the *preference bias.* In polarized organizations, controlling the *preference bias* is very important. Even small deviations on a single agent's preferences could erode the organization's exploration rate and lead to rates much lower than homogeneous organizations.

The paper is structured into four sections. Section 2 introduces the theory used to build the organizational learning model. Section 3 introduces the model and how it is put together. Section 4 shows the results. I finalize the paper with the discussion and limitations. I include extensive appendices, where I test every aspect of the model to outline its generality.

2. THEORY

The model I will present in the next section builds on three building blocks—first, a model of learning under uncertainty. Second, a decision structure to aggregate individual decisions. Third, a multi-attribute utility function defined by a single and static preference. Below, I describe the relevant literature for each of these building blocks[3].

2.1 Learning under uncertainty

March (1962:667) explains that the effects of having different points of view should be most salient when the environment is not stable. As shown by Csaszar & Levinthal (2016), in stable environments, people can use performance feedback to learn which points of view lead to higher performance. With this knowledge, organizations can develop rules to avoid the preferences that lead to lower performance. In general terms, an environment that is not stable has some level of uncertainty (Knight, 1921; March, 1991).

The canonical task for learning under uncertainty is the "N-arm bandit" task (Posen & Levinthal, 2012). In the "N-arm bandit", the agent needs to choose one option from the N options available in the environment (e.g., the N-arms of the bandit). Every time an agent chooses an option, it receives feedback[4]. The agents are asked to do this M times and try to maximize the accrued utility. The N-arm Bandit task has been studied in depth in organizational learning but only with single-agents or agents with shared points of view, i.e., homogeneous organizations (Denrell & March, 2001; Posen and Levinthal, 2012; Laureiro-Martínez et al., 2015: Puranam & Swamy, 2016). To step away from studying homogeneous organizations, I need to extend the "N-arm bandit" task to include a way to operationalize *preference diversity* and a joint decision-making process for the organization. I explain these in the next sections.

[3] This study is qualitatively different from Piezunka's et al. (2020) who studied how diversity in initial beliefs affects organizational learning. At the start of their simulations, agents with homogeneous preferences differ in their initial beliefs on the value of the N-options. Belief diversity leads to a lengthier learning process but is not related to diversity of preferences, all agents agree on how to achieve the goal, not so with *preference diversity.*

[4] This feedback may or may not be dynamic. Dynamism is the key difference between the N-arm bandit formulation of simulation studies (Denrell & March, 2001; Posen & Levinthal, 2012; Puranam & Swamy, 2016) and behavioral experiments (Daw et al., 2005; Laureiro-Martínez et al. 2015).

2.2. Agents with diverse preferences

To study organizations with *diverse preferences,* we need “a process by which decisions are reached without an explicit comparison of utilities” (March, 1962:666). In such a model, each agent “intends only [its] own gain, and [it] is in this, as in many other cases, led by an invisible hand to promote an end which was not part of [its] intention” (Smith, 1776, book IV chapter II). This description allows a more general formulation of organizational decision making, away from the “visible hand” of a manager who imputes a superordinate goal to the whole organization (Chandler 1977).

Options in an environment can be specified by a vector of multiple attributes ($\vec{x} = [x_1, x_2, \ldots x_j]$). However, our brains take all these inputs and create one utility measure, which is then used for comparison compare with the beliefs about the utility of other options (Rangel, Camerer, & Montague, 2008). Multi-attribute utility theories build functions that take the different attributes that define an option and, through weights and transformations, output one single and deterministic utility for the agent (Butler, Morrice, & Mullarkey, 2001). After this process is done, the agent can explore its environment as if this utility characterized the option.

Hybrid organizations provide a classic example of multi-attribute utility theories. In hybrid organizations, agents differ in the “weight[s] that different [agents] place on different attributes” that define the organization's problems (e.g., safety, profit) (Gaba & Greve, 2019). This preference diversity leads to conflict and the need for a more formalized decision-making process (Ocasio & Thornton, 1999; Ramus et al., 2017). In this study, I use Adner et al. (2014:2798) [5] multi-attribute utility function to study how majority voting organizations, defined by agents with *diverse preferences* (i.e., weights), explore their environment.

[5] Adner et al. (2014) used preferences to operationalize customers and explain why firms introduce different products to a market. I take a mirror approach and use their preferences to influence how a group of agents – a majority voting organization - learns about the utility of a set of N different options available in a new market. The utility function follow Protagoras maxim (Weick, 1995: 137), However, more complex utility functions can be built that do not fully comply with this maxim (Butler et al., 2001). Attributes that follow Protagoras maxim are: cost/quality, RAM/CPU, computing power/battery life, ease of use/feature richness, value to self/value to others,

2.3 Decision structures

Sah and Stiglitz (1986) formalized the stream of research on decision structures. This research continued and expanded within the management community with studies by Christensen and Knudsen (2010) and Csaszar (2013). The main finding of the decision structure literature is that the agents' structural arrangement directly affects the organizations’ commission and omission error rates[6]. With this knowledge, we can design optimal decision structures to the point of “approaching perfection” (Christensen & Knudsen, 2010:77). However, this “perfection” only applies if the organization’s members have the same preferences. In an organization with *diverse preferences,* “hits” and “misses” depend on the agents' preferences. Organizations with optimized structures might experience large deviations from their expected behavior when their agents' preferences are not used to design the organizations.

In an uncertain environment, where the organizations lack “demonstrably correct answer[s]” Argote (2013:134) explained how organizations employ majority voting structures to decide. Majority voting is known to be the “better” decision rule when looking for aggregating the preferences of one group of diverse agents (Kollman, Miller, & Page, 1997; Dasgupta & Maskin, 2008). Through majority voting, the organization can reach a decision, a process compatible with Cyert and March’s view that “people (i.e., individuals) have goals; collectivities of people do not” ([1963]2013:26).

The behavior of majority voting structures is highly important to the literature in decision structures. Specifically, triads are the main structure used to minimize error rates in Christensen and Knudsen (2010). It is the smallest structure that achieves lower commission and omission error rates than an individual. Although majority voting can happen in organizations with more than three agents, it is widely observed that “the further expansion to

speed/safety, social goals/profit (Murphy, Ackerman, & Handgraaf, 2011; Ramus et al., 2017; Gaba & Greve, 2019).

[6] For example, majority voting organizations with three members (i.e., triads) are more prone to commission errors than majority voting organizations of 3+ members but less prone than polyarchies of three members.

four or more by no means correspondingly modifies the group any further" (Simmel 1995:138). This notion supports my focus on microstructural organizations (Puranam, 2018).

In majority voting structures, the median voter has a significant power to rule the group's decisions, be it organizations or society (Congleton, 2004). The median voter theorem states that "a majority rule voting system will select the outcome most preferred by the median voter" (Holcombe, 2006:155). The median voter has this power because it has what Krehbiel (1998) calls a "pivotal role" that enables her to move decisions around her interests. For this reason, it is common within common strategy in spatial bargaining models of social choice to carefully include or exclude people from the decision process to try and control who the median voter will be (Baron, 1991; Ganz, 2020).

3. MODEL

This paper extends prior learning under uncertainty models by exchanging the single-agent decision process by a majority voting one (Posen & Levinthal, 2012; Christensen & Knudsen, 2010). Additionally, I include a multi-attribute utility function to assign each agent its own preference and create organizations with preference diversity (Adner et al., 2014). Below, I detail how I build the model, but before I present an example to ground the model's logic.

3.1 Example

A venture capital (VC) firm decides to enter a new investment market. It knows that the startups in the new market can be split into N different sectors and that their value can be estimated by two attributes, e.g., the quality of the business model and the intellectual property (IP) of the startup (Tata & Niedworok, 2020). The VC firm is unclear about which sector is best and can only improve its knowledge from experiential feedback. The VC needs to first invest in a startup and only after it receives a measure of the quality of the business model and IP of the startup. The VC firm has a standardized way of entering a new market. Three people are in charge of making the investments. The three people differ only in their preferences over business models and IP. Two people are carefully chosen, Agent A prefers business models, Agent B prefers IP.

The VC adds a third person to avoid split decisions and impasses, but this person's preferences are not carefully chosen (Agent C). I will revisit this example in the next sections and check how different *preference diversity* measures (i.e., *range, bias,* and *polarization*) affect how this VC firm explores the new environment.

3.2 Reinforcement learning

Figure 1 depicts reinforcement learning as four main steps: *Evaluation*, *Update*, *Proposal*, and *Selection* (Sutton & Barto, 1998). The process starts when the agents receive feedback from the environment. The feedback comes in the form of two attributes that define the performance of the prior investment decision ($\vec{x} = [x_1, x_2]$)[7]. All agents receive exactly the same feedback. They *evaluate* the feedback, each in their own and specific way, due to their individual preferences. After evaluation, each agent *updates* its own beliefs about the utility of the option they just invested in. The agents then *propose* an option for the organization to choose from.

In the next step, *selection*, the proposals are collected, and a choice is made by majority voting. All agents have a say and a vote in the decision. In contrast to a single-agent-based model, here, the individual agent cannot single-handedly decide. An option is selected only after a majority of agent chooses it. To choose an option, an individual agent needs others to invest in the option they prefer. Below I present the four stages of the model. I start with the *selection* process because this is the stage in which I introduce the most assumptions,

Figure 1: Depiction of the reinforcement learning model used in this study

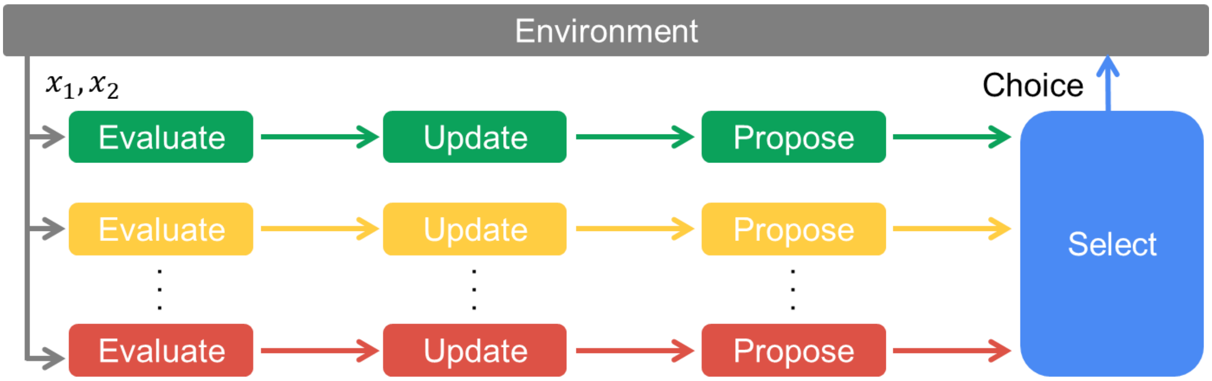

3.2.1 Selection process

In the selection process, the organization aggregates each agent's proposals and comes to one single choice. I have implemented this in several ways. The results are qualitatively similar as long as this selection process maintains an unbiased majority voting logic[7]. The first step in the selection process is to check if a majority of agents proposed the same option. If most agents proposed the same option, then the organization choses this option. If a majority is not reached at this point, the second step is to have the agents vote for the proposals. If an option receives the majority of the votes, then this option is chosen. However, given that there are more options than agents, it is not always the case that an option has a majority of the votes[8]. I call this situation a *split minority outcome.* In this situation, a decision rule considers the options the agents voted for and chooses one of them at random. This decision rule leads to exploration when the agents are indecisive[9]. For example, in the case of a triad, each option has a one-third chance of being selected, leading to at least a two-thirds chance that a new option will be chosen (i.e., one different from the previous choice).

3.2.1.1 Example revisited

The VC firm had selected three agents to make investment decisions in a new environment. Additionally, the VC firm requires each agent to bring every meeting a proposal of a sector to invest in the next period. They choose the sector they will invest via majority voting. The voting process goes as follows: if two of them proposed one sector, they would invest in a startup of that sector. If, however, there is no majority at the start of the meeting (e.g., all three proposed a different sector), they will vote and decide which sector to invest. However, if after voting gets a majority of the votes, they are told to choose randomly one of the three proposals.

---

[7] Appendix 6.4, present the results of other selection processes in all cases the results are qualitatively unchanged.
[8] I limit myself to cases where the number of options is larger than the number of agents.
[9] Appendix 6.3, shows what happens if I change the split minority outcome rule to one that maintains exploitation of previous options. The results are qualitatively the same and for this reason I do not explain them here.

The agents come prepared for the first meeting. However, given that neither has experience in this new and highly uncertain market, all choose a sector at random. Agent F proposes sector G, Agent B proposes sector A, and Agent C proposes sector F. Given that sector F is proposed by two of three agents, the organization selects to invest in a startup of sector F.

To summarize, there are three ways of selecting an option for investment: a) a majority proposes one option, b) a majority votes for one option or c) an option selected by a *split minority outcome* rule. After an option is selected, it is sent to the environment. The environment reacts and outputs a performance feedback for this choice in the form of two attributes. The feedback starts the next period of the reinforcement learning model, evaluation.

3.2.2 Evaluation

I employ the utility function from Adner et al. (2014:2978) to operationalize the agents' preferences. This utility function has three inputs, the two attributes of an option ($\vec{x} = [x_1, x_2]$), and the agent's preference, α. This preference weights how much utility the agent accrues from each attribute[10]. The utility function is defined by:

$$U(x_1, x_2, \alpha) = \alpha \cdot \log(1 + x_1) + (1 - \alpha) \cdot \log(1 + x_2) \quad (1)$$

Before each simulation starts, I assign each agent with a specific preference. These preferences are static. The agents' preferences are unchanged throughout the whole time the agents learn from the environment. They allow us to compare organizations with one another.

Prior models of learning under uncertainty have not included an evaluation stage because their organizations were homogeneous, i.e., composed of a single agent or agents whom all shared their preferences (Posen & Levinthal, 2012; Puranam & Swamy, 2016). Agents that share their preferences accrue the same utility from one option, and thus, the models can be implemented without the evaluation stage. If agents have diverse preferences, then we need an agent-specific evaluation stage.

[10] Appendix 6.5 presents results on the use of other utility functions outside the one from Adner et al. (2014). These results show little deviations due to exchanging utility function.

3.2.2.1 Example revisited

After investing a startup of sector F, the VC firm receives feedback on its performance. The startup had a bad business model ($x_1 = 0.2$) and a good IP ($x_2 = 0.8$). All agents are given the same information. They have no ambiguity about the feedback they receive. However, each agent will have a different view of how valuable they find the startup. Using Equation 1, Agent A, the one with a strong preference for business models, can be given $\alpha_A = 0.9$, and it will receive a value of $U(x_1 = 0.2, x_2 = 0.8\,, \alpha = 0.9) = 0.097$. Agent B, the one with a strong preference for IP, $\alpha_B = 0.1$ will receive a value of $U(x_1 = 0.2, x_2 = 0.8\,, \alpha = 0.1) = 0.238$. Agent C goes through a similar evaluation process. If we set its preference at $\alpha_C = 0.65$ it will receive a value of $U(x_1 = 0.2, x_2 = 0.8\,, \alpha = 0.65) = 0.194$. Agents A and B differ quite significantly in regards to the value they ascribe to the startup. In the future, Agent A will likely propose Sector F again, whereas Agent B is less likely.

3.2.3 Update and Propose

After evaluating the performance feedback, each agent *updates* its beliefs about the option's value and proposes a new option. I replicate exactly the *update* and *propose processes* from Puranam and Swamy (2016). The agent's beliefs are updated through an exponential recency weighted average process (Sutton & Barto, 1998). The values are updated by a constant ϕ parameter common to all agents. After the *update process*, the agent proposes a new option through a softmax function that uses the beliefs of the value of each option as input and a τ parameter that controls the amount of exploration between options.

All agents share both learning parameters, ϕ, and τ, and these parameters do not change during the simulations. I optimize the ϕ and τ parameters so that the agents in homogeneous organizations accrue the maximum amount of utility. In Appendix 6.2[11], I present the

[11] The optimization process of the ϕ and τ values achieves a goal set by Csaszar (2013) for the case of majority voting organizations, i.e., to understand how exploration and exploitation of multi-agent organizations changed with the number of agents and the decision structure. In Appendix 6.2, I show how larger majority voting

optimization process in detail and how it changes according to the number of learning periods and the organization's number of agents. The results replicated for a broad range of ϕ and τ values, different update (1/k+1), and proposal algorithms (ε-greedy, aspiration level search).

3.3 Mechanism: Diverse preferences

Agents in the model are indistinguishable from each other except by their preference (α) over the two attributes (x_1, x_2). Everything else that defines an agent (e.g., processes, parameters) is shared with the other agents in the organization. An organization is defined by the number of agents that compose it and the preferences of its agents. To describe the preferences of small organizations, we can list their preferences. A triad can be uniquely described by $\vec{\alpha} = [\alpha_A, \alpha_B, \alpha_C]$. A homogeneous organization is one whose agents all agents share the same preference. Organizations with *preference diversity* have at least one agent with different preferences from the other members.

One simplification I use in this paper is that the preferences of the extremal agents in the organization (e.g., the ones with highest and lowest preferences) sum up to one. This simplification allows me to study all differences in triads' preferences with a two-dimensional figure at the expense of not being able to explore all possible "average" organizational preferences. In Appendix 6.1, I explain how this is done exactly and its limitations. I study organizations with three, five, or seven agents. To compare organizations with different number of agents, I create three summary measures: *preference range, bias,* and *polarization*. They are depicted in Figure 2 and summarize the preferences needed to describe an organization.

Figure 2: Preference diversity as range, bias, and polarization

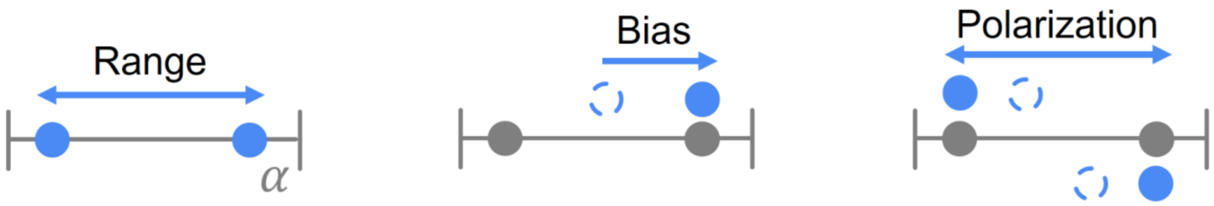

organizations need more curious members to explore as much as smaller organizations. To keep an exploration rate, at constant ϕ, one needs to increase the τ value in proportion to the square root of the number of agents.

3.3.1 Range: $\boldsymbol{R} = [\mathbf{0.0}, \mathbf{1.0}]$

Preference range (Figure 2 left) quantifies how distant the two extremal agents' preferences in an organization are from each other. A *preference range* of zero is sufficient to deem an organization homogeneous, as the two most distant preferences are the same. The higher the range, the more diverse the preferences of the triad can be.

3.3.2 Bias: $\boldsymbol{B} = [\mathbf{0.0}, \boldsymbol{R}/\mathbf{2}]$

Preference bias (Figure 2 middle) refers to the position of the median agent's preference in regards to the extremal agents. Preferences are unbiased if the median agent's preference is equidistant from the two extremal agents' preference. The highest *preference bias* appears when the median agent has the same preference as an extremal agent[12]. *Preference bias* can have the highest value, $B = R/2$, when the *preference range* is highest, $R = 1.0$. The utility function we use is symmetric in the preferences. This allows studying only organizations with positive bias without losing generality.

3.3.3 Polarization: $\boldsymbol{P} = [\mathbf{0.0}, \boldsymbol{R}]$

In organizations with more than three members, the agents' preferences can build coalitions of agents whose preferences are close to the extremal agents. *Preference polarization* is the measure of the distance of the most similar agents between the two coalitions, e.g., the blue agents' preference distance in the right panel of Figure 2. I study only organizations with an odd number of agents and with two symmetrical coalitions.

3.4 Universe of preferences

The goal of this paper is to explore how *preference diversity* affects organizational learning. To do this, we will compare organizations by their *preference range, bias,* and *polarization.* The first two measures can be plotted together as “pixels” in Figure 3. If the organization has more than three agents, we need to specify the polarization of the preferences and show multiple of

[12] Note that the median agent’s preference cannot go past an extremal’s preference. If that were the case, the agents would exchange their labels, and “previously extremal” agent would become the new median agent.

these graphs to present all results. In Appendix 6.1, I show how each pixel's position is translated to the preference vector of organizations with 3, 5, or 7 agents. For example, how the yellow pixel (R = 0.8, B = 0.15) in Figure 3 directly translates to $\vec{\alpha} = [0.9, 0.1, 0.65]$, for the case of a triad and $\vec{\alpha} = [0.9, 0.1, 0.65, P, 1.0 - P]$ in the case of 5-agent organization.

The three corners of Figure 3 are important reference points. Figure 4 presents a visual representation of each of these corner cases' preferences for the case of triad organizations[13]. The lower-left corner represents a *homogenous triad* whose members share the same preferences [0.5, 0.5, 0.5]. The lower right corner represents a triad whose agents have different preferences, and each difference between the preferences has the same value [1.0, 0.0, 0.5]. I call this the *diverse triad*, as the preferences are as different as they can be. Finally, the upper right corner has two agents with the same preferences and one agent whose preferences are completely different [1.0. 0.0, 1.0]. I call this the *biased triad* because even though each agent has a vote, the options that A and C propose tend to have more weight on the final choice, i.e., the choices are biased. As explained in Appendix 6.1, in organizations with more agents, the extra agents lie in between the extremal agents. Thus these three corner cases remain relevant.

Figure 3: Universe of organizations with different preferences

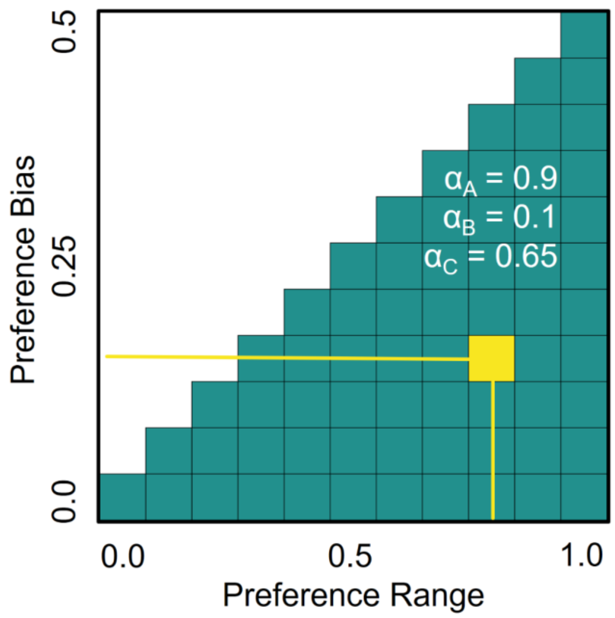

[13] Note that in triads, the preferences are technically polarized, as the "coalitions" are composed by only one agent.

Figure 4: Corner case for the case of triads

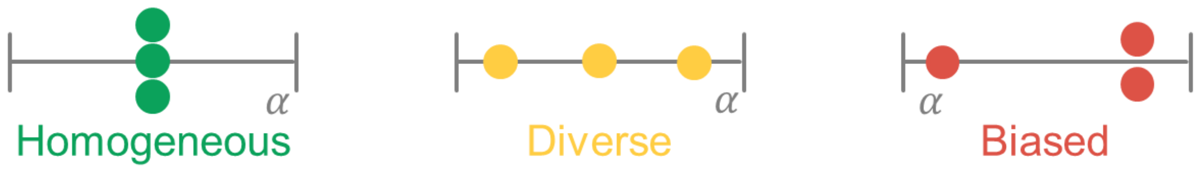

3.5 Environment

In Adner et al. (2014), options are drawn from an NK landscape with two performance dimensions. I do not do this. Instead, I will use the more general production productivity frontier employed by Porter (1996), i.e., a concave and semicircular frontier that represents the edge of what is possible for firms to produce and bring to the market[14]. In Adner et al. (2014) and Porter (1996), the firms were producers. However, this study takes the counter stance, with the organizations playing the role of customers of the environment's options. The utility function from Adner et al. (2014) was meant for this specific use, as it described the customers of the firms that supplied the market.

I define environments with N = 10 options available for the organization to select, the N-arms of the bandit. An option *"f"* is defined by the attribute vector: $\overrightarrow{x_f} = [x_{1_f}, x_{2_f}]$. The two attributes of the option, x_{1_f} and x_{2_f} are drawn from two uniform distributions with range [0, 1] constrained by $x_{1_f}^2 + x_{2_f}^2 \leq 1$. The constraint has the effect of limiting the options to be drawn from a frontier defined as one-quarter of a circle, as bounded by the dashed line in Figure 5.

Figure 5 shows an example of the 10 options in a learning environment. Every time an option "f", is chosen, each member of the triad receives feedback about the option's vector ($\overrightarrow{x_f} = [x_{1_f}, x_{2_f}]$) plus a noise vector ($\overrightarrow{u(t)} = [u_1(t), u_2(t)]$). The noise is added so that the agents benefit from learning and not just choosing each option once and then selecting the best final outcome. The noise values are taken from uniform distributions with the constraint that ${u_1}^2 + {u_2}^2 < 0.1^2$, that is, a circle of radius 0.1 around $\overrightarrow{x_\iota}$, the grey circle in Figure 5.

[14] Appendix 6.6 studies the effect of a convex production possibility frontier similar to a Cobb-Douglas production function as well as an environment where the options can freely have values between 0 and 1, square in form.

3.6 Simulation

As in Figure 3, I investigate triads that encompass all possible values of *preference range* $R \in [0,1]$, *preference bias* $B \in [0, 0.5]$, and *preference polarization*, $P \in [0,1]$. I follow Figure 3 using steps in *preference range* of 0.1 and *preference bias* of 0.05. In every simulation, an organization chooses between N = 10 options. The options are drawn from a semicircular productivity frontier. On each draw of an option *"f"*, noise is added so that the option appears to be drawn from a circle of radius 0.1 around the position of the option, $\overrightarrow{x_f}$ (see Figure 5).

I optimize the agents' learning parameters so that the *homogeneous triad* agents achieve the highest accrued utility. I can do this without making an "interpersonal comparison of utilities" as all homogeneous triad agents share their preferences. Thus, I do not value one agent's preferences over the other (Arrow, 1951:3; March 1962). For any other triad, the optimization process would need us to compare the agents' utilities, a process that is known to be problematic by social choice theorists. Unless directly stated, the simulations use ϕ = 0.3 and τ = 0.04 as learning parameters.

Figure 5: Example of a learning environment

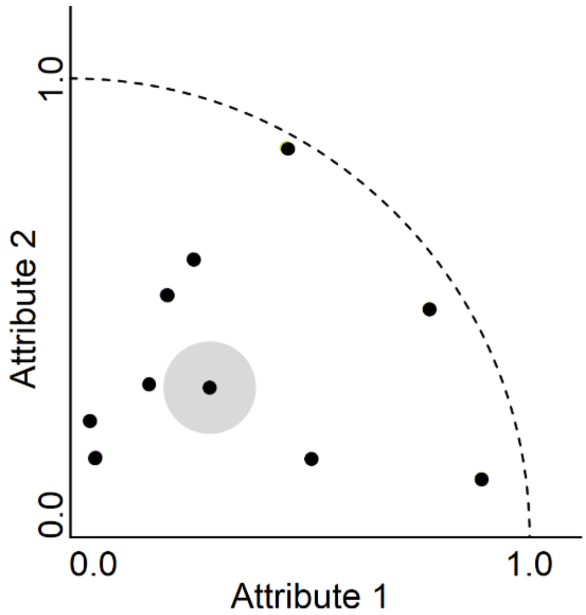

A learning epoch starts after the options are defined. A learning epoch encompasses the 100 periods under which an organization learns from its environment. At the start of the epoch, I initialize the agents' beliefs. I give as initial beliefs the average values of x_1 and x_2 for the current environment. The average is taken as the mean position of the 10 options in the environment. I do this so that the amount of updating for each option is not biased between trials[15]. During the learning epoch, the organization chooses options for a total of 100 periods. After the 100 periods, the beliefs of the agents are reset back to the average values. A new epoch begins, and the process is rerun. The positions of the bandits are maintained. The learning process is repeated for 100 epochs to understand better how each organization learns from each 10-arm bandit environment. Figure 6 shows in color-scale the log10 of the percentage of times each of the options is chosen; the lighter the option's color, the more frequently it is chosen[16]. Figure 6 represents the choices of a *homogeneous triad. T*herefore, the options closest to the frontier and furthest to the origin are chosen the most.

After the hundredth learning epoch of one environment, a new 10-arm bandit environment is drawn, and 100 learning epochs start for this environment. The process is repeated for 1000 different 10-arm bandit environments. If we put together the 1000 different versions of Figure 6, we get the aggregate results shown in Figure 7. After the process is finished, I store three main results for this organization: a) the number of times each bandit (e.g., option) was chosen, b) the number of exploration events (i.e., +1 if the option chosen in period t is different from the option chosen in period t-1) and c) the utility accrued by each agent. I then continue with the next organization preference configuration (e.g., the next “pixel” in Figure 3). This procedure is done for the full universe of organizations, i.e., the 66 pixels.

[15] Note that given that the initialization is not optimistic (i.e. initialize all beliefs at the maximum value), it often happens that not all options are explored. in one learning epoch. I have explored random initialization of the agent’s beliefs as well as optimistic initialization and the results are qualitatively the same except for an increased exploration level at the start of the simulation.
[16] If the triad chose only one option every time, the highest value would be zero. However, as more than one option is chosen the value is lower than 0

Figure 6: Log_{10} of the percentage of times an option is selected in one environment

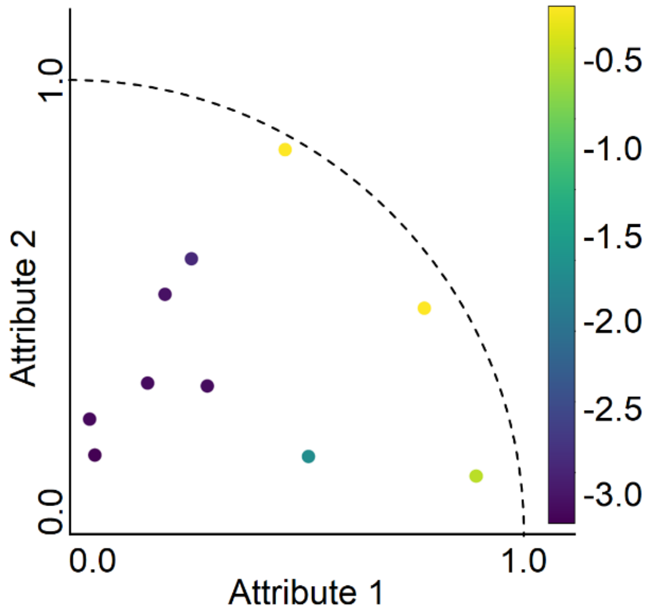

4. RESULTS

This study aims to understand how *preference diversity* affects how organizations explore options while learning from an uncertain environment. In this section, I present an answer to this question. The answer is presented in three parts—first, the effects of preference diversity in triads. Second, the effects in larger organizations. Finally, a summary of the effects.

4.1 Organizations with three agents

Preference diversity affects how often an organization explores different options and which options the organization uses the most. In this section, I present both results.

4.1.1 Use of the N-arm bandit's options

Figure 7 shows the logarithm base 10 of the percentage of times an option was chosen by a *homogeneous triad* (R = 0.0). Figure 7 collects the results of all the 1000 environments simulated, each for 100 epochs of 100 periods each. For each environment, there are 10 dots plotted, each with its own color. Similar to Figure 6, the color scale indicates how often an option is chosen. The lighter the color, the more often it is chosen. In Figure 7, the options furthest from the origin are chosen the most. There is also a spread of options chosen close to the productivity frontier. In contrast, options near the origin are chosen less often.

In general, these results show that the *homogenous triad* is acting rationally. This organization's agents give the same weight to both attributes and accrue the highest utility from options with the highest summed values. The results shown in Figure 7 are a good test to the model as the selection of options close to the frontier and furthest to the origin is what should be expected by Porter (1996) and Adner et al. (2014).

Figure 8 presents different results, the case of the *diverse triad (R = 1.0, B = 0.0)*. The *diverse triad*, just as the *homogeneous triad*, selects more options the furthest from the origin. However, there is a broader light region along the production possibility frontier, a sign of more variance in how options are chosen. The higher variance can be explained as in the *diverse triad. T*wo agents do not accrue the highest utility by selecting options furthest from the origin. Instead, agents A and B accrue the highest utility by selecting options with high values of Attribute 1 or Attribute 2, respectively. As the learning process develops, Agent C might be indifferent between two options, but agents A and B will not and will try to direct the triad's choices to the most valuable options. The *preference range* in the *diverse triad* leads to a broader set of options being chosen often and leads to the broader lighter area in Figure 8.

Figure 7: Log_{10} of the percentage of times an option is selected by the homogeneous triad

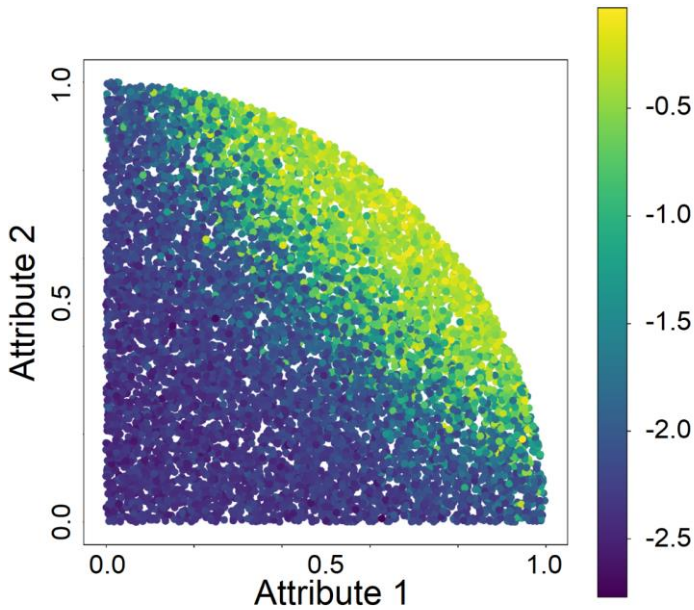

Figure 8: Log_{10} of the percentage of times an option is selected by the diverse triad

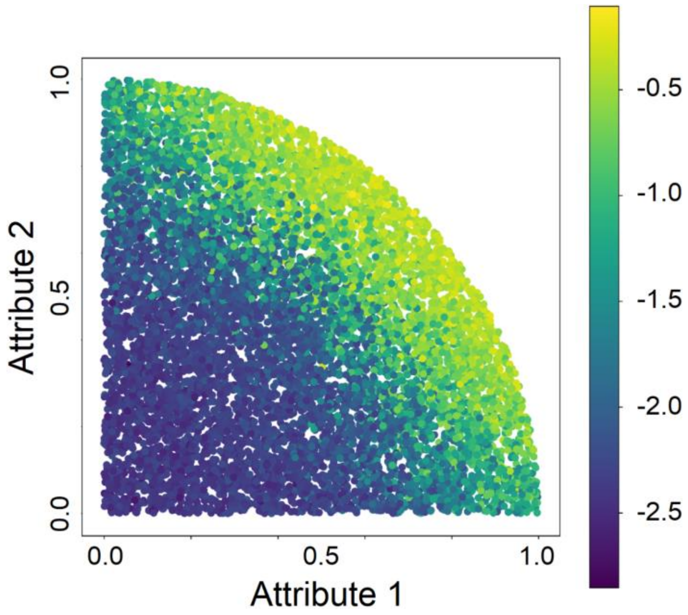

The *homogeneous triad* and the *diverse triad* are unbiased. That leads to options furthest away from the origin to be chosen the most. Figure 9 shows a very different case, the case of the *biased triad* (R = 1.0, B = 0.5). Two agents in this triad maximize their accrued utility by choosing options with high values of Attribute 1. The fact that most agents (two out of three) have a strong preference for Attribute 1 leads to this triad to choose options with the highest values of Attribute 1 without little regard for Attribute 2. Due to the utility production possibility frontier employed, the selection of options with high values of Attribute 1 is directly anticorrelated to options with low values of Attribute 2 (Adner et al., 2014). In Appendix 6.6, I show results for a convex possibility frontier.

The results in Figures 7, 8, and 9 explain how preferences affect where triads explore their environment, similar to the idea of Porter (1996) and Adner et al. (2014) that expects customers to choose options closer to the production possibility frontier and in the direction of their preferences. In the next section, I present results that show how much exploration the amount of exploration in each case.

Figure 9: Log_{10} of the percentage of times an option is selected by the biased triad

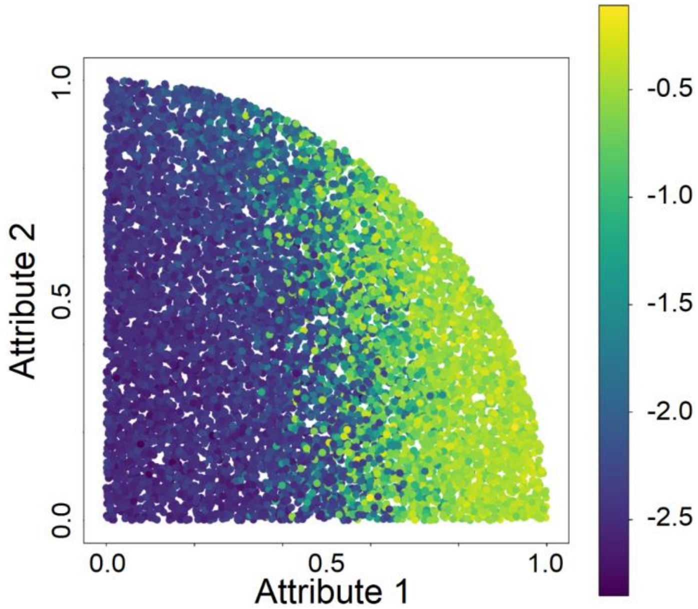

4.1.2 Effect on exploration rate

Figure 10 summarizes the main result of this study for the case of triads. Figure 10 plots the percentage of exploration events that each of the 66 different triads did while learning in an uncertain environment – the lighter the color, the higher the exploration rate[17]. The *diverse triad* explores the most, a total of 25.0% of the periods. The *biased triad* explores the least, only 5.3% of the periods. The *homogeneous triad* explores 17.2% of the time. *Preference diversity* can increase exploration by 52% but also decrease exploration by 68%.

Figure 11 plots the triads' exploration rate in the rightmost column of Figure 10 as a function of the preferences of Agent C[18]. The exploration rate is highest for the *diverse triad* (yellow dot), and the rate monotonically decreases for increasing values of *preference bias,* in a nonlinear manner. At low *preference bias* the exploration rate does not decrease much. However, already with a *preference bias of B = 0.18* ($\alpha_C = 0.68$), the exploration rate is exactly

[17] The marks in the color scale are accurate. However, the maximum and minimum values are not shown at the extreme values in the color scale. Instead, I quote the values in the text.
[18] The triads in the rightmost column of Figure 10 These triads have high *preference range, R = 1.0*. Agent A prefers Attribute 1 ($\alpha_A = 1.0$) and is indifferent to Attribute 2, the opposite is the case for Agent B ($\alpha_B = 0.0$). What separates the different triads in the column is the preference of Agent C. Agent C as the median voter has a strong effect on the choices of the triad (Holcombe, 2006).

the same as the exploration rate of the *homogeneous triad* shown as the green line in Figure 11. In the *biased triad* (the red point in Figure 11), the exploration rate is still lower, one fifth the exploration rate of the *diverse triads* and less than one third the exploration rate of the *homogeneous triad*[19].

Figure 10: Percentage of exploration by triads

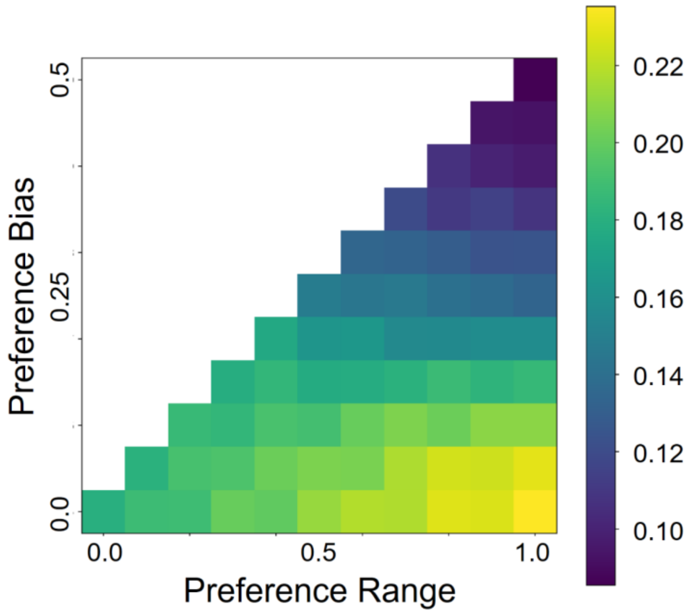

Figure 11: Percentage of exploration by triads with fully dispersed preferences, R = 1.0

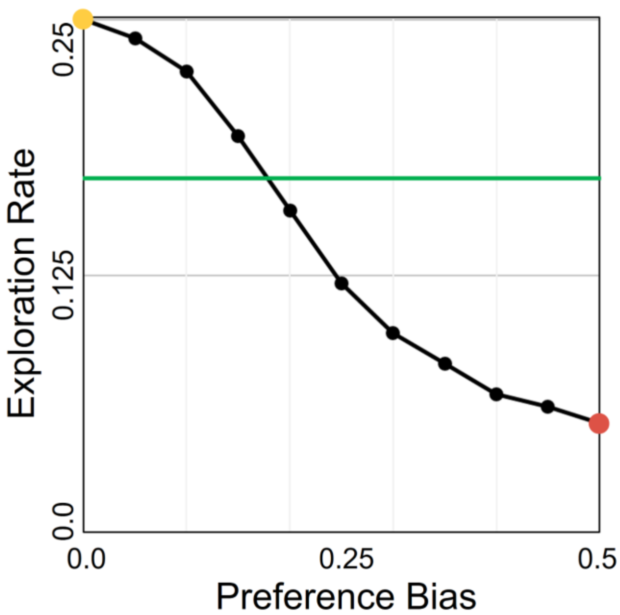

[19] The exact exploration rate depends on the number of periods in a learning epoch. As explained in Appendix 6.2, learning epochs with more periods require higher exploration rates, higher τ values. The results shown before are for a learning epoch of 100 periods. However, the results are qualitatively similar for any epoch length.

These results show that one cannot predict the exploration rate of a triad from just observing how different the preferences of the most extreme agents are, *preference range*. One needs to take into account as well the preferences of the third agent, i.e., the *preference bias*. Without accounting for the *preference bias* that comes from the third agent, any prediction about the exploration rate of a triad with people with different preferences can vary broadly; the triad might explore 50% more but also 66% less than a *homogeneous triad.*

I explore these results further in the Appendices. In Appendix 6.4, I present the results under a modified selection process. Appendix 6.6 employs an environment with a convex production possibility frontier, similar to a Cobb-Douglas production function. Appendix 6.7 discusses how the results are mostly unaffected by implementing learning processes similar to Denrell and March (2001) or Posen and Levinthal (2012).

4.2 Organizations with more agents

I simulated organizations composed of five-member that decide via majority voting. I follow the same simulation processes as in the simulations with the triads. The learning epochs have 100 periods, and there are 1000 environments simulated for each of the 66 pixels. However, the learning parameters need to be changed. As explained in Appendix 6.2, to achieve optimal performance, τ needs to increase in proportion to the square root of the number of agents in the organization. I find that a *homogeneous organization* with five members performs best in a 100-period learning epoch when its learning parameters are $\tau = 0.06$ and $\phi = 0.3$.

Figure 12 shows the exploration rate of a 5-member majority voting organization with full *preference polarization*, $P = 1.0$. The results are qualitatively similar to the ones in Figure 10. The *diverse organization* explores the most, a total of 28.1% of the periods. The *biased organization* explores the least, only 13.7% of the periods. The *homogeneous organization* explores 22.7% of the time. *Preference diversity* can increase exploration by 24% but also decrease exploration by 40%. Further analysis shows that just as in the triads, the organization with R = 1.0 and B = 0.18 has the same exploration rate as the *homogeneous organization*. Due to the increase in the τ learning parameter that I use to optimize the behavior of the 5-member organization, the exploration rates are higher in the 5-member organizations than in triads.

Figure 13 shows a non-polarized organization, $P = 0.0$, i.e., two agents are ambivalent between the two attributes of the options. In this case, I see that the effect of *preference bias* is almost irrelevant. The *diverse organization* explores the most, a total of 25.1% of the trials. The *homogeneous organization* explores the least, only 22.6% of the periods. The *biased organization* explores 25.0% of the time. *Preference diversity* can increase exploration by 11.2%. In contrast to the prior cases, *preference bias* has a much more limited effect, and it barely decreases the exploration rate. The *biased organization* has almost the same exploration rate increase as the *diverse organization*, 10.8%. In all other cases, the *biased organization* had a much lower exploration rate.

The results in Figure 13 are noisier than previous simulations. The increased noise appears as *preference bias* has a much-decreased effect on the organizations' exploration rate. Thus, the range of values represented by the color scale decreases significantly. The appearance of noise is explained more in detail in Appendix 6.1.3.

Figure 12: Percentage of exploration of 5-agent organizations with fully-polarized agents, P = 1.0

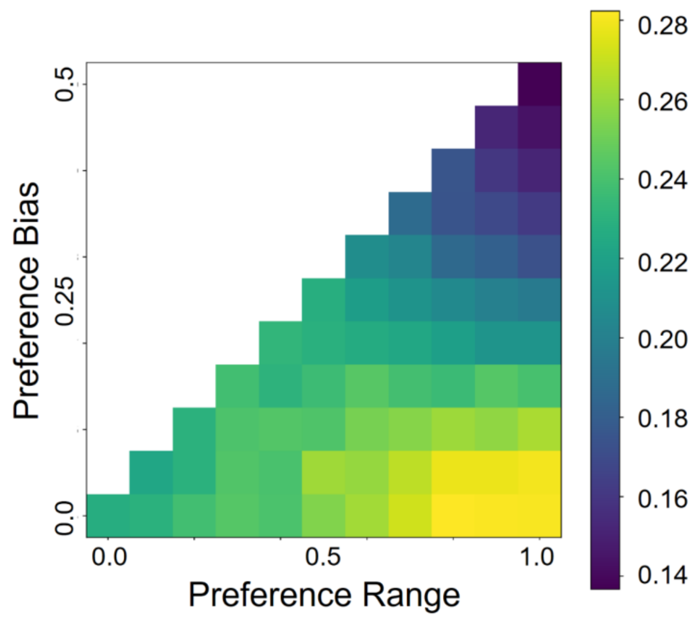

Figure 13: Percentage of exploration of 5-agent organizations with non-polarized agents, P = 0.0

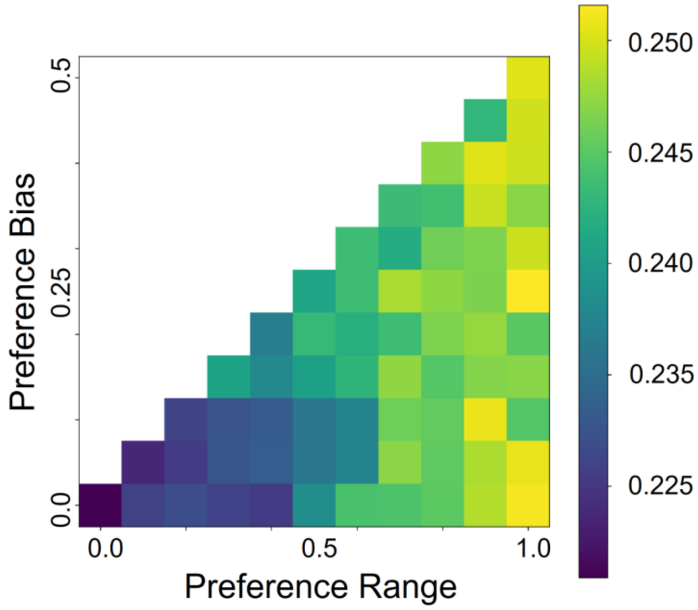

In non-polarized organizations, the effect of *preference bias* is of little importance. The *preference range* has a more direct effect on the organization's exploration rate. This has the effect of simplifying the effect of *preference diversity*. However, in non-polarized organizations, the maximum exploration rate is lower (25.1%) than in fully polarized organizations (28.1%). The exploration rate of the *homogeneous organizations* ranged between 22.4% and 23.0%. Therefore, a manager can reliably predict to increase the exploration rate by +11.2% if she promotes increased *preference range* and the decrease of *preference polarization* in the organization. If, instead, the manager decides to increase *preference polarization* and *range*, the organization could achieve double that increase, +23.8%. In this way, the manager needs to trade off predictability for a higher exploration rate.

The effect of *preference polarization* is equivalent to negatively moderate the effect of *preference bias.* At low levels of *preference polarization,* the *preference bias* has little effect on the organization's exploration rate; its effect increases only when polarization increases.

In Appendix 6.1.3, I include the case of a 7-member majority voting organization. I follow a similar procedure as when transitioning from three to five agents. The results of organizations with 7 agents do not differ from the results shown in this section.

4.3 Summary of the effect of preference diversity in exploration

I have shown how organizations with *polarized preferences* behave similarly to triads, whereas organizations with *non-polarized preferences* behave more straightforwardly. This happens because a triad can be seen as fully-polarized organizations whose coalitions are composed of just one member. In general, if an organization has *diverse preferences*, i.e., if at least two agents do not have the same preferences ($R > 0$), then we can predict how much the organization will explore with the two by two matrix of Table 1. The key combination is to understand how biased and polarized the preferences of an organization are. The plus and minus signs in Table 1 refer to a large or small change in the exploration rate when compared with organizations with homogeneous preferences, i.e., the change in exploration rate due to *preference diversity*.

Organizations that have high *preference polarization* achieve a much higher exploration level if the preference of the median agent is unbiased (i.e., $B = 0$). However, the exploration rate decreases sharply as the preference of the median agent deviates from its unbiased position to the point that any increase in the exploration rate due to *preference diversity* vanish entirely when the *preference bias* is higher than $B = 0.18$ ($\alpha_C = 0.68$). Similarly, organizations that have low *preference polarization* and high *preference range* achieve higher exploration rates than homogenous organizations. However, these organizations' exploration rate increase is about half as high as in organizations with polarized and unbiased preferences.

Table 1: How much would a diverse organization explore?

Effect of *Preference Diversity* on Exploration		Preference Bias	
		Low	*High*
Preference Polarization	*High*	**+ +**	**– –**
	Low	**+**	**+**

Note: Exploration rate when compared to a homogenous org.

5. DISCUSSION

Preference diversity has a strong and nontrivial effect on the exploration behavior of organizations learning under uncertainty. Organizations with *diverse preferences* can achieve higher but also lower exploration rates than organizations with *homogeneous organizations*. However, to account for these changes, managers need to acknowledge the complexity inherent in *preference diversity*. I present three measures that allow us to predict the exploration rate of organizations with *diverse preferences: preference range*, *preference bias*, and *preference polarization.* These three measures determine whether an organization will explore more or less when compared to a *homogeneous organization.*

For a manager who aims to create an ambidextrous organization, fostering *preference diversity* can be beneficial, but only if she implements adequate measures. I find two sets of measures that lead to higher exploration. First, the manager needs to promote differences in preferences but create structures to minimize their polarization. The second requires that the manager promotes polarization but minimizes the bias of preferences. Unfortunately, there are many more ways of creating organizations with *preference diversity* that explore less than organizations whose agents all share their points of view. However, by following these two routes managers can limit the downside of *preference diversity* and instead foster exploration.

In the past decades, the study of decision structures has dramatically improved our understanding of how to design more accurate decision-making organizations, to the point of "approaching perfection" (Christensen & Knudsen, 2010:77). This paper sounds a note of warning over these studies, as I find that organizations behave very differently when their agents have diverse preference. *Preference diversity* strongly affects the decision-making of a majority voting organization. Even though this study does not use the standard notion of Type I and Type II errors, the fact that exploration rates vary by up to 5x can justify imagining that the error rates are not stable across different preference diversity levels.

Since March's (1991) study, scholars have looked for ways of increasing exploration in the pursuit of creating ambidextrous organizations. Many solutions focus on separating the organization, be it in time, context, or structure (O'Reilly & Tushman, 2011). *Preference diversity* is a lever for the organization's exploration rate that does not require its members' separation. A manager can estimate the *range, bias,* and *polarization* of its agents' preferences and adjust accordingly. Just as in March (1991), turnover could be used to control these summary measures and keep the organization's exploration rate at the desired level.

This study follows the almost 60-year-old call from March (1962:666) to study organizations as composed of "a process by which decisions are reached without the explicit comparison of utilities". This was also a central argument within Arrow's (1951) formulation of social choice theory as the comparison of utilities is an ill-defined problem. An important ramification of avoiding comparison of utilities is that in this study, the agents have preferences, the agents make decisions, and the agents learn from the environment. The role of the organization is to set the rules for how to aggregate the decisions. The organization I model does not have goals, and it does not learn. Its agents do. Although uncommon, this is a direct implementation of Cyert and March's (1963) multi-goal organizations.

As organizations introduce machine learning to increase their prediction capabilities, an important research area is the inclusion of artificial agents as equal partners and decision-makers as humans (Murray, Rhymer, Sirmon, 2020). Machine learning can be a fruitful way for managers to further control their organization's exploration rate (Kellog, Valentine, Christin, 2020). The manager could give the artificial agents precisely determined preferences to minimize any bias in the organization (Raisch & Krakowski, 2020). A manager could go a step further and employ artificial agents to control the "direction" that the organization explores most often (e.g., invest in higher quality IP more often). However, these measures of control could backfire. For example, employees could try to all jointly follow the artificial agent's decisions. Employees once valued for the unique points of view now follow an artificial agent's preferences. A once diverse organization would become homogeneous.

An important caveat of this study is that agents' preferences are static. However, the preferences of people do change with time. An important extension of this work will be to understand how exploration is affected by diverse preferences that vary with time. Just as in the case of artificial agents, it might be that diversity cannot survive socialization pressures, and we need to return to March's (1991) use of turnover to maintain *diversity* in the organization.

One aspect that this paper does not touch is the effects of communication. In this, I follow prior literature in decision structures, but in organizations, people communicate their points of view while on meetings or committees. The model presented here then is still a simplification of organizational life even while extending prior models. The lack of clear conflict is the main drawback of the lack of communication. Further studies could investigate conflict by following top management teams as they make decisions, or conversely, implementing behavioral experiments to understand how conflict and conflict resolution affect how organizations with diverse preferences explore their environment.

This study's results hold for even larger organizations than the ones shown. However, as more members are added, a decision needs to be made regarding how to define the preferences of the new agents (see Appendix 6.1). To avoid this, a generalization can be made where instead of formally specifying each agent's preferences, one can instead use distributions to draw the agents' preferences and then estimate the three *preference diversity* parameters to study the behavior of larger organizations. The use of these methods could bridge the gap between this study's findings and research in social choice theory.

6. APPENDIX

The main result of this study is that *preference diversity* affects exploration, but to understand and control its effects, one needs to dig deeper and account for the *bias* and *polarization* of the organization's preference. The results have been replicated in a variety of other simulated learning conditions. These different conditions are explored more in detail in the Appendices. There I studied different learning processes (aspiration-level and ε-greedy search), different split minority rules (stability or change), utility functions (linear or social value orientation), and voting processes (synchronous or asynchronous). Results change in part when I draw the options from a convex environment (Cobb-Douglas production function), but they hold even in organizations with more agents (7, or more).

6.1 Specifying the agent's preferences

6.1.1 Triads

Each agent in the majority-voting triads has one vote, and given the selection process, all have the same possibilities to propose and select options. The triads are flat organizations, but the *preference diversity* makes it so that some agents might find it that they invest in the options they like more often than others. To do this, I assign every agent in the triad one preference that does not change throughout the simulation. The collection of preferences is stored in one vector that completely defines a triad: $\vec{\alpha} = [\alpha_A,\ \alpha_B, \alpha_C]$. Each dimension of the vector can range between 0.0 and 1.0. The three vectors span a unitary cube. Many positions in this cube are redundant. Therefore, I followed a simplification process to find the universe of triads that is both most relevant to this study and not redundant.

In order to remove redundancies and provide a clearer explanation, I perform two simplifications of the preferences studied in this paper. First, given that each agent has the same voting power, then the order of the preferences does not matter. A triad that has preferences defined by $[\alpha_A,\ \alpha_B,\ \alpha_C\ ]$ should have the same as a triad wit preferences defined by $[\alpha_B,\ \alpha_A,\ \alpha_C\ ]$ or any other combination of the three preferences. Given this symmetry, I can choose one constraint on the preferences and remove redundancies without losing any generality. Therefore, I select triads under the following preference constraint:

$$\vec{\alpha} = [\alpha_A,\ \alpha_B,\ \alpha_C\ ]\ where\ \alpha_B \leq\ \alpha_C \leq\ \alpha_A \tag{2}$$

This simplification reduces the space of preferences from one unitary cube to one-sixth of the cube. However, this paper's focus is to study the effect of differences in the preferences of the agents. The first simplification allows for every set of differences to be explored and every possible set of average preferences of the triad agents. The second simplification keeps all the possible differences but limits us to study only triads with an average preference of around 0.5.

Many of the triads that agree with the first simplification have agents with preferences that are to one another. However, these triads differ in the average value of the preferences. For example, [0.9, 0.7, 0.8] has an average preference of 0.8, and [0.4, 0.2, 0.3] has an average value of 0.2, but the inter-agent differences in the preferences are the same for both triads. In the second simplification, I equate these two triads. The simplification's net effect is to make a translation of the preferences so that the preferences of Agent A and Agent B mirror each other around 0.5, and thus the average preferences do not vary too much around 0.5. This simplification is operationalized by setting:

$$\mathrm{\alpha_B} = 1 - \mathrm{\alpha_A}\ and\ \alpha_A \in [0.5, 1.0] \tag{3}$$

The second simplification shifts all agents in a triad by increasing or decreasing the three agents' average preferences by a constant amount. For example, through this shift, the two triads from before get shifted to the triad [0.6, 0.4, 0.5]. Also triads like [1.0, 1.0, 1.0] will be translated to the same position as [0.0, 0.0, 0.0], that is to [0.5, 0.5, 0.5].

This second simplification narrows down the possible triads to be explored significantly. But it keeps all possible sets of differences in the preferences between two agents. This is the case as Equation 3 allows the differences between Agents A and B to have any range possible (i.e., the difference varies between 0 and 1), and Agent C then can have any value in between the two extreme agents. The symmetry in the utility function then allows us to bound the range of Agent to 0.5 and the preference of Agent A, and we are left with a full universe of triads that shows all possible sets of differences in a two-dimensional figure instead of a three-dimensional cube. Given that the study of differences in preferences is the focus of this study, this is an acceptable compromise on, and the motivation is to keep clarity and simplicity

After following these two simplifications, I can fully specify every possible set of differences in triads with two parameters: $\alpha_A \in [0.5, 1.0]$ and $\alpha_C \in [0.5, \alpha_A]$. The constraint in the second parameter leads to a half triangle that defines every possible difference in triads. The triangle of all possible triads that fit the two simplifications is shown in Figure 2. Note that only one of the pixels in Figure 2 denotes a triad with homogenous preferences, the lower-left corner pixel. All the other pixels represent triads with diverse preferences. Each pixel varies in the level of *preference range* and *preference bias.* Literature in organizational learning, so far, has focused on organizations with homogeneous preferences, and therefore in the next sections, the changes introduced by *preference diversity.*

6.1.2 Five agents

In this paper, I show that diverse preferences have an important and nuanced effect in the option exploration of a prototypical organization, a majority voting triad. Results suggest that the diversity and the bias of the agents' preferences jointly determine when the organizations will explore more or less than an organization whose agents all shared their preferences. In selecting a majority voting triad, I aimed at simplicity and generalizability. However, the framework I have presented allows us to study larger organizations. In this section, I show how to extend the results to organizations with five agents and motivate the generalizability of the results to organizations of any size.

The main change one needs to do to study 5-agent majority voting organizations is to determine how the agent preferences will be distributed. A simple way of doing this is to extend the prior organizational preference vector ($\vec{\alpha} = [\alpha_A, \alpha_B, \alpha_C]$) to five agents.

$$\vec{\alpha} = [\alpha_A, \alpha_B, \alpha_C, \alpha_D, \alpha_E,] \tag{4}$$

The new vector has two new parameters, α_D and α_E, but it maintains the simplifications used in the previous section (i.e. $\alpha_B = 1 - \alpha_A$ and $\alpha_C \in [0.5, \alpha_A]$). I specify three new conditions to the new parameters to guarantee that I can employ the previous section's data structures. The first is a simplification that specifies that the preferences of the new agents lie between the preferences of Agent A and Agent B:

$$\alpha_B < \alpha_E < 0.5 < \alpha_D < \alpha_A \tag{5}$$

Second, I define that the preferences of the two new agents mirror each other, similar to agents A and B so that:

$$\alpha_E = 1 - \alpha_D \tag{6}$$

Finally, I define a polarization variable that specifies how much the new agents approach the extreme agents (A and B). The agents will behave thus in a way similar to Agent C so that $\alpha_D \in [0.5, \alpha_A]$ *and* $\alpha_E \in [\alpha_B, 0.5]$. I define this condition by

$$\alpha_D = 0.5 + P \cdot (\alpha_A - 0.5) \tag{7}$$

Figure A.1 presents this in more detail. If $P = 0$, then the two new agents will be ambivalent about the attributes of an option ($\alpha_D = \alpha_E = 0.5$). In contrast, if $P = 1.0$ then the two new agents will share the preferences with Agents A and B ($\alpha_A = \alpha_D$ *and* $\alpha_B = \alpha_E$). The P parameter is a measure of polarization parameter as it defines how separate the two agents that prefer Attribute 1 (Agents A and D) are from the agents that prefer Attribute 2 (Agents B and E). If $P = 1.0$, they will be as distant as it is possible under the limitations of the *preference range* of the organization (i.e., Agents D and E cannot become more extreme than agents A or

B). If instead, $\beta = 0.0$, Agents D, and E will be indifferent about the two attributes and thus will act as a stabilizing factor in the organization.

Figure A.1: Diverse organization with different levels of preference polarization

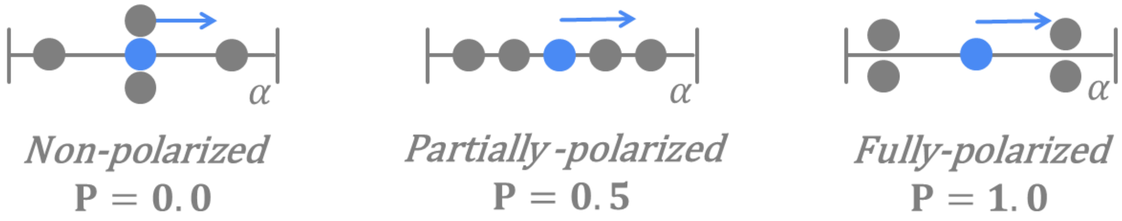

6.1.3 Seven agents

I presented simulations for 3 and 5 agents. I claim that the results generalize to larger organizations but that the exact definition of the preferences of the agents is complicated. In this section, I include the simulation of majority-voting organizations with 7 agents. This requires us to define the preferences for two more agents. I do this by placing the preferences of the two new agents, Agents F, and G in between the preferences of the agents of each coalition. So the new preferences are defined as:

$$\alpha_F = \frac{\alpha_A + \alpha_D}{2} \text{ } and \text{ } \alpha_G = \frac{\alpha_B + \alpha_E}{2} \tag{8}$$

There is a simple way of reducing this equation, and the preferences of these agents will depend solely on α_A and β. This process can be continued, and we can place as many agents in between Agents A and D and Agents B and E as we like in a relatively straight forward manner. However, ad infinitum, one would have two ranges of preferences $[\alpha_B, \alpha_E]$ and $[\alpha_D, \alpha_A]$ within which there will be a uniform distribution density of agents. Everything outside these ranges would be unpopulated except for the preference of Agent C. The idea that preferences behave in this way is highly implausible, and thus, the generalization presented in the discussion section is a better way forward for organizations with dozens of agents. Kamada and Kojima (2014) present another possibility. Both provide solutions for studying large organizations with diverse preferences without requiring a huge number of parameter definitions.

For seven agents, the specification of Equation 8 works well. I use this formulation to simulate a majority voting organization with 7 agents who learn in environments with 10 options, during learning epochs of 100 periods and 1000 environments. In this case, the optimal learning parameters are $\tau = 0.07$ and $\phi = 0.3$. I simulate 66 different combinations of the attributes according to the specifications explained before. The results are shown in Figure A.2 for the case of a fully-polarized organization and Figure A.3 for a non-polarized organization.

For the case of majority voting organizations with 7-agents fully polarized preferences, I find results that align with the case of fully-polarized organizations with 5 agents as well as triads. That is, all organizations with polarized preferences behave in a qualitatively similar fashion. The *diverse organization*, does the maximum exploration, 28.0% in this case. The least amount of exploration Is done by the *biased organization*, 16.6%. The *homogeneous organization* explores, on average 24.4% of the periods in a learning epoch. This means that the *diverse organization* explores 14.8% more and the *biased organization* 31.9% less than the *homogeneous organization—t*he same qualitative pattern as in Figure 10 and Figure 13.

Figure A.3 presents the case of a 7-agent whose agents are not polarized. In this configuration, the results resemble the ones found in Figure 14. A higher *preference range* between Agents A and B leads to more exploration independent of the preferences of Agent C. I find that in the non-polarized *homogeneous organization*, the exploration rate is 22.8%. The exploration rate of the rightmost column of Figure A.3, the column with the highest preference difference between agents A and B, has an exploration rate of 25.2%, an exploration rate 10% higher than the exploration rate of the homogeneous organization. However, the differences are small enough on each pixel of that column so that the highest exploration rate of 25.9% happens at $\alpha_C = 0.75$ and the lowest 24.5% in its neighbor pixel $\alpha_C = 0.75$. Longer simulations can lower this noise and show the maximum l. However, in general, the trend is that in non-polarized organizations, the effect of *preference bias* is much less relevant, even less so as the number of agents increases. There is a clear benefit in terms of predictability of exploration to having non-polarized organizations even though the exploration rate in the *diverse and non-polarized organizations* (25.2%) is just 90% of the exploration rate of a *diverse, unbiased, and fully-polarized* organization (28.0%).

Overall, I find that organizations with 7 agents behave in a qualitatively similar way as organizations with 5 agents. Additionally, polarized organizations with 3, 5, or 7 agents behave in qualitatively similar manners. Thus the results shown in the paper and summarized in Table 1 apply for larger organizations as well. I have simulated organizations with 9 agents and 11 agents with the logic I presented at the beginning of this section and the results do not change in any qualitative form. The simulations however take much longer, and the noise becomes harder to remove as the variation between the different pixels decreases. Therefore, seven agents are close to the limit of what the framework developed for this paper is capable of handling in an insightful manner. For larger organizations, it might be useful to think of mean-field approximations and study the effect of one agent in a sea of agents. This is the standard procedure in modeling atomic structures when the number of electrons increases (Opper & Saad, 2001).

Figure A.2: Exploration rate for 7-agent organizations with fully-polarized preferences

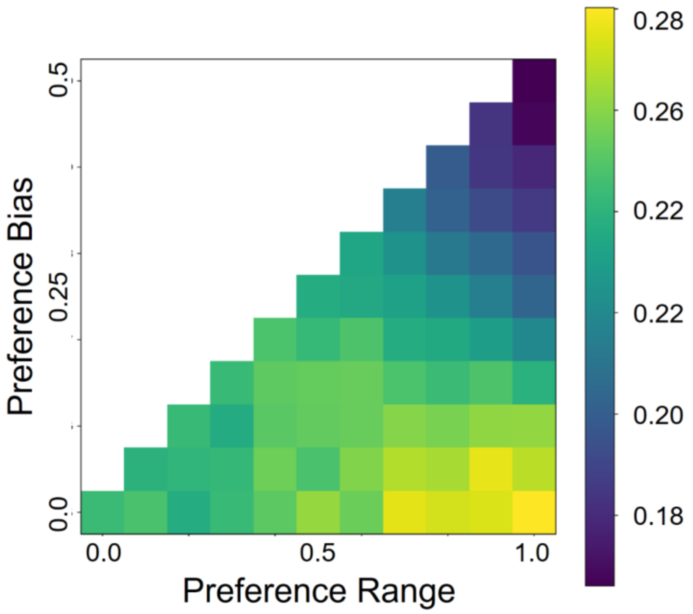

Figure A.3: Exploration rate for 7-agent organizations with non-polarized preferences

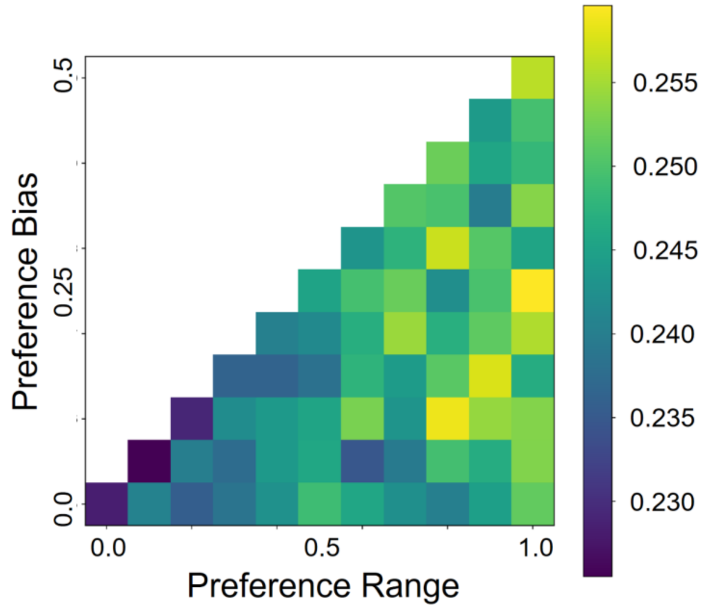

6.2 Optimization of learning parameters

I optimize the learning parameters of the majority-voting organizations by giving all agents the same preferences, i.e. $\alpha_i = 0.5 \; \forall \; i \in \{A, B, \dots C\}$. This describes a homogeneous organization where all agents act as if following one superordinate goal, given that all agents want to maximize the same accrued utility. Due to this homogeneity, I can explore different combinations of learning parameters (i.e., ϕ and τ) and determine the one that gives the highest accumulated utility after the simulation. The combination that achieves the highest accumulated utility is optimal because all agents agree on the preferences. This would no longer be the case in organizations with *diverse preferences* but in homogeneous organizations is adequate.

To simulate the parameter optimization, I employ N = 10 options per environment, 100 periods per learning epochs, and 1000 learning environments per pixel. A pixel, in this case, is a combination of ϕ and τ. I allow $\phi \in [0.05, 1.0]$ and $\tau \in [0.01, 0.2]$, for each, there are 20 steps. The organization simulated is a triad. The goal of all agents in the simulation is to propose options that give them the highest value while still exploring enough to avoid missing out on valuable unknown opportunities.

The results of the simulation are shown in Figure A.4. The figure has two axes, one for each learning rate. Lighter colors symbolize higher average accumulated utility per period by the agents in the organization. The value is the sum of the three agents' utility; thus, on average, the top-performing organizations accrue about 0.22 units of utility per decision. The worst organizations achieve a lower performance of about 0.18 units of utility per decision. The better the exploration rate of an organization matches the environment, the higher the performance.

Figure A.4 shows that for every ϕ value, there is an optimal level of τ that achieves close to the maximal performance in the simulation. That is, independent of how much the agents in the organization update their beliefs every period, there is an exploration level, τ, that will enable the agents to achieve as high a performance as if they had used another value of ϕ. Therefore the choice of ϕ can be seen as arbitrary as long as τ is well matched. In Figure A.1, we can also observe how the range of values that achieve high performance broadens as ϕ increases. For this reason, in the paper, I set $\phi = 0.3$. I do this because lower ϕ values make it complicated to determine accurately which τ value is best, and higher ϕ values lead to the agents updating a vast portion of their beliefs on every period.

Figure A.4: Optimization of learning parameters for triads

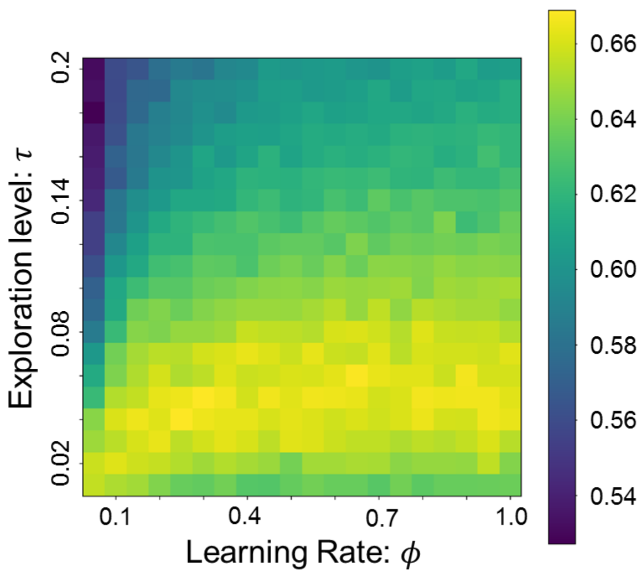

Figure A.5 shows the results of having increasing or decreasing the length of the learning epoch. The figure includes 50 different values of τ for the range $[0.01, 0.2]$ and twenty different lengths of the learning epoch, ranging from learning epochs of 10 periods to learning epochs of 200 periods—the lighter color mark higher average utility for the organization. Here we can see how the longer the learning epoch, the higher the performance. This is sensible as in longer simulations, as the agents benefit from exploiting a good option. We can also observe how a higher exploration level is needed for longer learning epochs.

In the case of very short learning epochs, the optimal exploration rate is also very low. That tells us that in short learning epochs, exploration is not a good strategy, and instead, the firms benefit from choosing their first option. However, the benefit is not very significant as the performance of these organizations is very low, close to the performance of an organization that chooses at random an option. For learning epochs with a higher number of periods, then the average accrued utility of the organizations increase. However, the more periods, the longer the simulation takes to finish.

Figure A.5: Optimization of $\boldsymbol{\tau}$ parameters for learning epochs of different learning epoch length

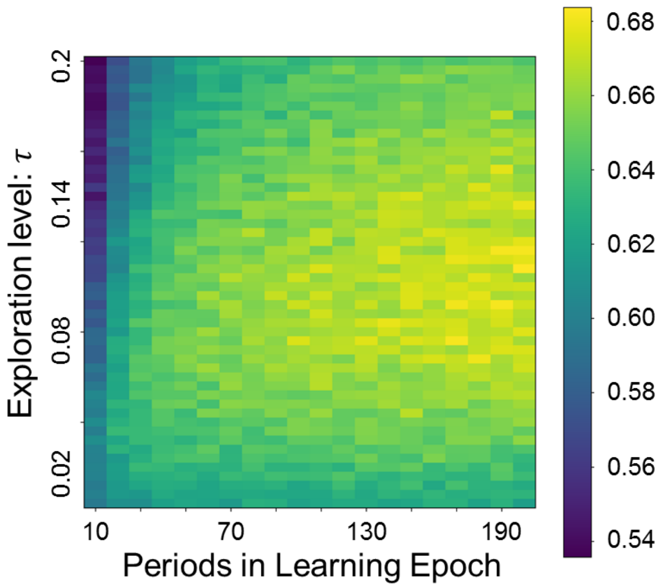

I choose to simulate the learning epochs of 100 periods. This decision is made as, in that case, we can achieve optimal performance close to the one achieve in longer learning epochs (0.66 at 100 periods and 0.68 at 200 learning periods) and limit the length of the simulation. A simulation as in Figure 10, takes around 8 hours of simulation time, and doubling the simulation time would limit the flexibility of the data collection process without achieving much better learning results.

The learning parameters also need to be optimized for organizations with more or fewer agents. Figure A.6 shows the results of simulating organizations with between 1 and 19 agents, in steps of two added agents per column. Same as before, I show the average accrued in color scale, but here I show the color that each agent individually. All agents in these simulations have the same preference ($\alpha = 0.5$), i.e., all get the same utility from each choice. However, like the columns in the simulation has a different number of agents, it is important to provide the accrued utility per member and thus normalize the color scale.

In its abstract, Csaszar (2013) stated that the study wanted to investigate "what is the effect of organization size on exploration?" The paper does this by showing that different organization structures achieve different commission and omission error rates and present these two error rates as microfoundations of exploration and exploitation in organizations. In doing this, the paper shows that there is no clear relationship between organizational size and exploration and exploitation, as even a small organization like a triad could have different exploration rates due to the structure of the agents.

However, one can find a relationship between organization size and exploration if one keeps the structure constant. This is not possible for every decision structure, but it should be possible for the three structures studied by Sah and Stiglitz (1988). Indeed in Figure A.6, I do the first step in that direction. I show how a majority-voting organization with different size require agents with different exploration levels τ, in order to explore the environment optimally.

I find a square root relationship between the number of agents and the optimal exploration level, τ. The more agents in a majority voting organization, the more explorative the agents in the organization should be. This is sensible as the more votes a decision requires, the less frequently that the organization will choose a "random" proposal. Given that learning under uncertainty requires exploration, to keep an optimal performance level, an organization will require agents that propose more to explore new options. I also find that as the number of agents increases, the need for matching the exploration level optimally decreases, there is a broader range of τ values that achieve high performance.

As in the case of the length of the learning epoch, here I choose triads as the main object of study in the paper as it limits the amount of time needed for a simulation to finish. However, as triads do not allow us to study polarization, I also employ organizations with more agents. I find that for a learning rate of 0.3, and a learning epoch of 100 periods, the optimal exploration level for a triad is 0.04, for an organization with 5 agents it is 0.06, and for an organization with 7 agents, 0.07.

Figure A.6: Optimization of learning parameters for organizations with N agents

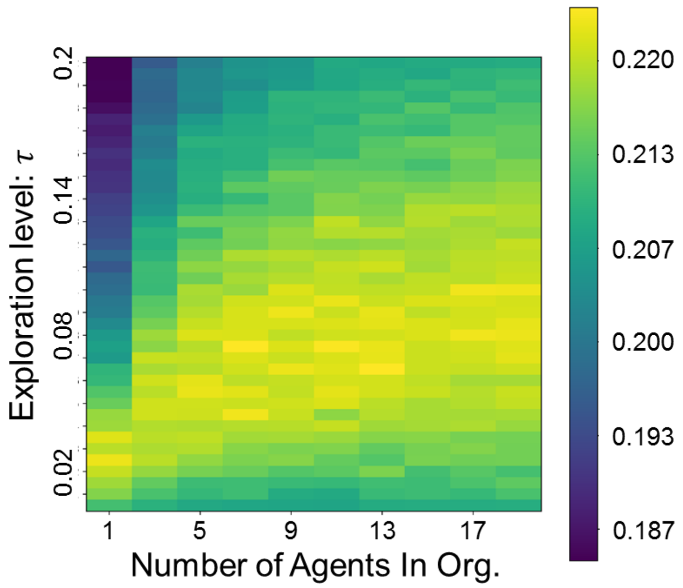

The results in Figures A.4, A.5, and A.6 present an interesting idea, that there is an optimal matching of learning parameters and organization size for learning under uncertainty that leads well-matched organizations to achieve optimal performance, i.e., equifinality between functionally equivalent organizations (Gresov & Drazin, 1997). I do not find any organization that achieves outstanding performance. In Figures A.4, and A.5, I find that there is one combination that accrues the maximum utility on each column. Thus, one can easily select the optimal number of agents and their learning parameters without losing anything other than the simulation time. The length of the learning epoch does affect the attainable performance. However, a given length can find organizations with any learning rate, ϕ, and the number of agents that achieve optimal performance, as long as they are matched with an appropriate exploration rate, τ.

6.3 Split minority outcomes

The results presented so far employed the chance rule for resolved *split minority outcomes* through a *chance rule, i.e.*, situations where, after proposing and voting for options, an option does not get the majority of votes. In these situations, with the *chance rule*, there is at least a 66% probability of choosing an option different from the one chosen last period. I implement another rule, the *stability rule*, that, when invoked, selects the prior choice instead of the options proposed by the agents. The stability rule makes an exploitation decision every time it is used. In Figure A.7, I show the *stability rule* results instead of the *chance rule* during learning.

The results with the *stability rule* resemble the ones with the *chance rule*. However, there are some differences. The biggest difference is that the maximum exploration does not happen at the highest levels of *preference range*. Instead, it happens when $R = 0.9$. The triad that explores the most (R = 0.9, B = 0.0) explores 17.8% of the periods. The *biased triad* explores the least, only 5.3% of the periods. The *homogeneous triad* explores 15.5% of the time. When employing the *stability rule*, *preference diversity* can increase exploration by 14.9% but also decrease exploration by 53%.

Figure A.7: Percentage of exploration in triads with the stability decision rule

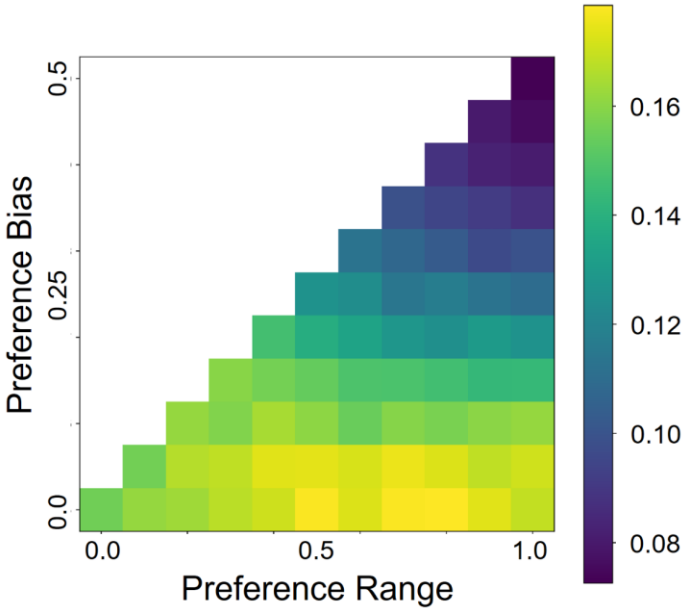

In Figures 10 and A.7, I showed the difference made by changing the rule that the organization employs when it reaches a split minority outcome. These simulations are run with learning epochs of 500 periods, but the learning parameters used are the ones of learning epochs for 100 periods. I simulated longer epochs to be able to see the long term trends of the variables plotted in the figures. Figure 10 shows organizations that employ the *chance rule*, and Figure A.7 organizations that use the *stability rule.* Although the qualitative results do not change if I swap the rules, organizations with a higher *preference range* explore more. Organizations with biased preferences explore less with the *chance rule*; the specifics do change. In this section, I explore the reason why these differences happen.

Figure A.8 plots the proportion of selection processes that get to the voting stage as the periods of the learning epoch progress. That is the proportion of trial in which there is no option proposed by a majority of the members of the triad. Figure A.7 plots this proportion for the three corner cases. We see how, in the case of the *homogeneous* and *biased triads*, the proportion of trials that get to the voting stage decreases rapidly and stabilizes at a very low value. In contrast, for the *diverse triad*, a much larger proportion of cases get to the voting stage.

The differences in the equilibrium voting proportions give a clue as to why the behavior changes when I use the *stability rule* instead of the *chance rule* for deciding on *split minority* cases.

Figure A.8: Percentage of voting as a function of time

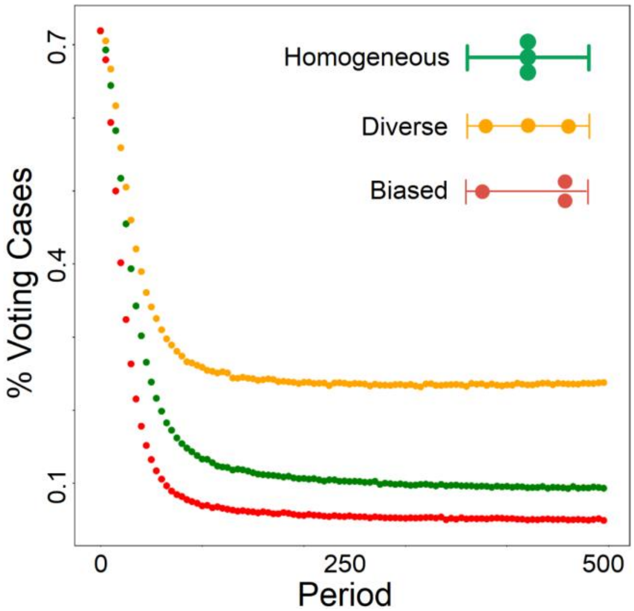

Figure A.9 plots the proportion of voting cases that yield a *split minority* outcome. Here I find that the proportion decreases rapidly and stabilizes at a low value in the *homogeneous triad* case. For the other two corner cases, the rate decreases early on but later starts to increase and then tends to an equilibrium level. As before, the *diverse triad* has a higher equilibrium value. Therefore, the *diverse triad* gets to a *split minority* outcome more often than the other corner cases. A large part of the increase in the *diverse triad's* exploration rate can be ascribed to the *chance rule*. However, as shown in Figure A.7, the *chance rule* is not the only reason why triads with diverse preferences explore more.

Figure A.9: Percentage of split minority outcomes in an environment

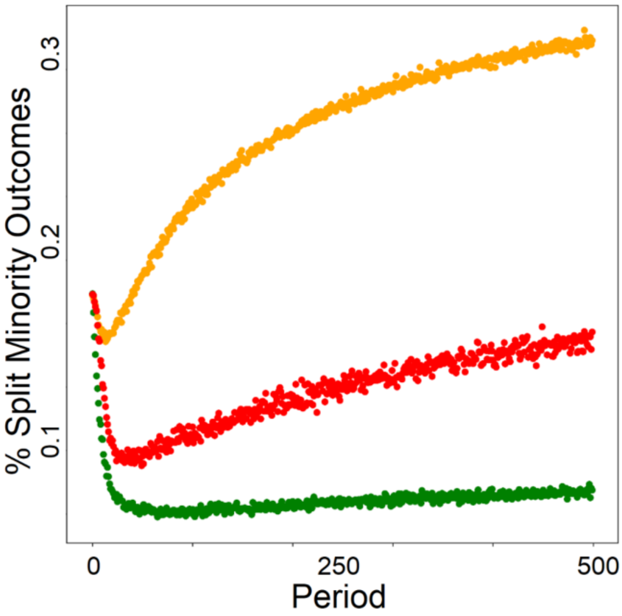

6.4 Asynchronous selection processes

I have presented just one type of option selection process. The selection shown in this paper was chosen because it was synchronous, i.e., it had all agents doing the same thing on each of its steps. For example, all agents had to propose an option and vote for options simultaneously. This gave all agents in the organization the same power to propose and decide, and thus, there was no imbalance in how the agents behaved. I have implemented several selection processes. In this section, I present one of them and discuss other possibilities.

Figure A.10 shows the exploration rate of organizations that employ an asynchronous selection process to decide which option to invest in every period[20]. The asynchronous selection process starts with two agents proposing options to invest. If the two agents proposed the same option, then this option is selected by the organization this period. If the two proposals are different, then the third agent is in charge of selecting the one it prefers. This selection process does not have *split minority* outcomes and does not need the *chance* or *stability* rule. However, it is hard to expand to organizations that are composed of anything other than three agents. For

[20] The asynchronous selection process I present next was originally proposed by Thorbjorn Knudsen in a presentation of this paper (Knudsen, 2019)

this reason, I did not choose this selection process for the main body of the paper. Importantly, the asynchronous selection rule achieves the same qualitative behavior as the synchronous rule shown in the manuscript.

I also tried a variation of this asynchronous process. Instead of two agents bringing an option, only one agent brings the option, and the other two agents need to vote whether to invest or not in the option. If the option receives one vote, then it is invested in. In that case, I count the proposer as giving one vote, and the vote received in the voting stage as a second vote. This method requires more processing after as when no vote is received, it is unclear what to do. One can use a *stability rule* and invest in the previous choice. One could instead try a new proposal from another agent and invest if that one gets a majority. In addition to different selection processes, I tried another rule for the *split minority outcome.* In this rule, one asks randomly one of the agents to cast a second vote when a *split minority outcome* is reached. In doing this, the agent would have a bit more power in that period, but all agents have the same power on average.

Figure A.10: Exploration rate for organizations with the asynchronous selection process

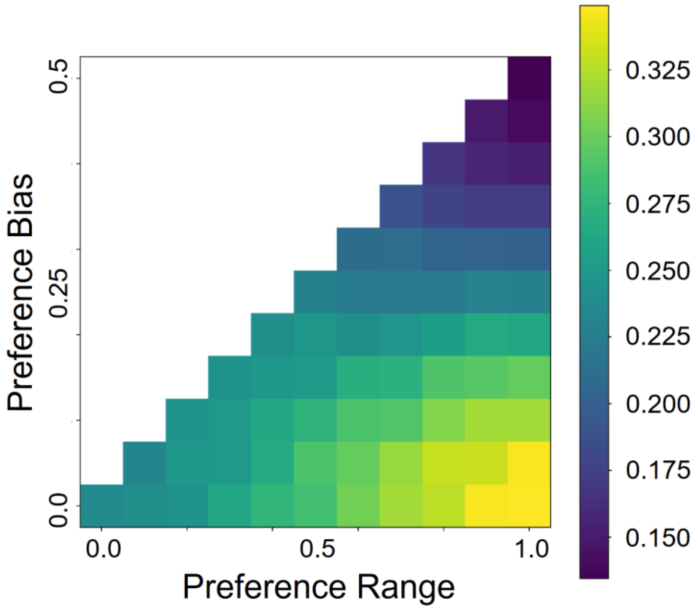

The results of this selection process resemble the results shown in Figure 10 and in Figure A.10 for the synchronous and asynchronous selection processes. Lastly, I tried a process where options come to the agents, and they decide to invest in them or not. A random number generator gives the options, and the agents vote to select them or not.

6.5 Linear utility function

I follow both the utility function and the environment choice of Adner et al. (2014). That paper had a utility function that weighted the contribution of each attribute in a linear way, but it used a logarithm base 10 to transform how an agent ascribed value to each attribute. This is common within the literature of the study but less common in other areas.

The qualitative results I show are not dependent on the utility function I use. Figure A.11 shows the use of a linear utility function instead of the one from Adner et al. (2014:2798). The function is linear in both the attributes and the preference each agent has for each attribute.

$$U(x_1, x_2, \alpha) = \alpha \cdot x_1 + (1 - \alpha) \cdot x_2 \tag{8}$$

Figure A.11: Percentage of exploration in triads with a linear utility function

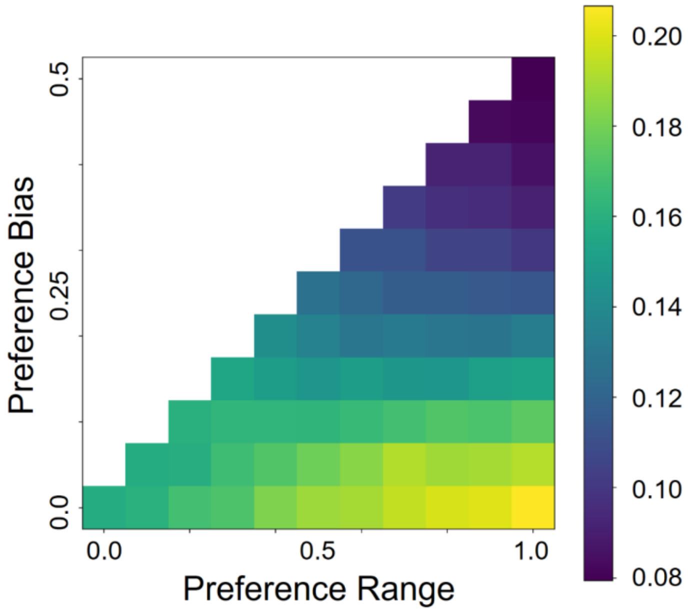

Figure A.11 shows the result of doing the same simulation as in Figure 10 but with the linear utility function of Equation 8. The changed utility function has no qualitative effect in the way triads with diverse preferences learn in an uncertain environment. I have explored other utility functions such as the one from social value orientation in (Murphy et al., 2011). However, in all cases, the results resemble the ones shown in Figure 10 and A.11.

6.6 Convex environment

So far, I have presented different robustness tests that show that the results presented in the paper are not affected by the selection process used nor by the utility function that the agents employ to evaluate the feedback of their decisions. In this section, I will present an aspect that does affect the exploration rate of organization: the convexity of the concavity of the environment from where the options are drawn.

The previous results had options drawn from a concave environment as the one in Figure 5. This environment is one defined by diseconomies of scale and is characterized as a partly circular production possibility frontier (Porter, 1996). However, in economics, convex production possibility frontiers are also used. Kamada and Kojima (2014) show a similar point in relation to the overreliance of voting models on concave utility functions. A convex production possibility frontier would look similar to Figure A.12, or Cobb-Douglas production functions. The options in the environment of Figure A.12 are defined by:

$$\vec{x} = [x_1, x_2] = [1 - y_1, 1 - y_2] \; where \; y_1^2 + y_2^2 > 1 \tag{9}$$

In general, what I do is to take an option that did not fit the criteria for a concave production function (i.e. $x_1^2 + x_2^2 \leq 1$). The options would fill the non-accessible part above the production possibility frontier. I then mirror these options in the x and y-axis, and the results are the equivalent of drawing options from the convex area shown in Figure A.12.

If I draw options from the gray area in Figure A.12, then there is a problem when initializing the agents' beliefs. Normally, I initialize the beliefs to the average values of each attribute of the options in the environment. The average value is accessible within the bounded region of a concave function, such as in Figure 5. That is an extension of what makes the area bounded concave. In the case of a convex function, the average does not necessarily have to be inside the bounded area. The fact that the initialization value is not necessarily inside the area makes it that the agents will need more rounds of selecting an option before they get an accurate overview of an option's value. Additionally, given that these agents tend to be biased in their preferences, the initialization will resemble an optimistic initialization of the bandits, and thus a higher exploration rate should be expected for these environments. There could be a way to avoid this, but it is unclear how to initialize the bandits in a sensible and unbiased way in a concave environment.

Figure A.13 shows the results of doing an equivalent simulation as in Figure 10 but with the options drawn from a convex environment, as in Figure A.11. One aspect is clearly different. The *homogeneous triad* explores the most in this environment. This is the case as this triad is trying to accrue high values of both attributes. This is not possible in this environment. However, given that the triad has no preference for either attribute, it will prefer an option with a high value of either attribute. This leads to a significant amount of exploration as the agents will find several options that have a high value.

Figure A.12: Convex production possibility frontier

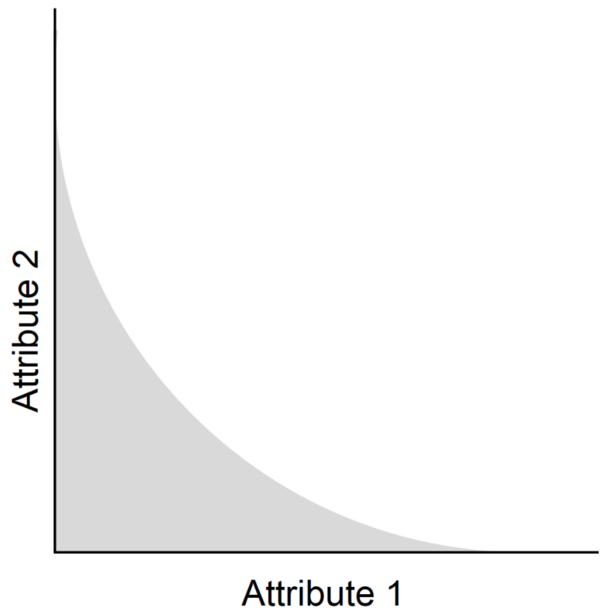

Figure A.13: Percentage of exploration in triads an environment with economies of scale

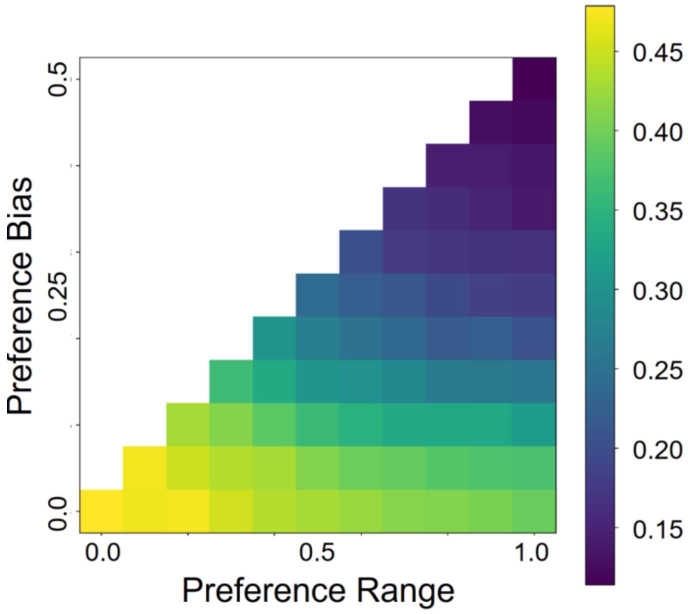

In contrast, when the agents have more diverse preferences, the symmetry of the attributes is broken. At least two agents have clear preferences for either the options that have high values of Attribute 1 or high values of Attribute 2. Give this asymmetry exploration goes down in the case of the *diverse triad.* An aspect that does not change is that the *biased triad* explores much less than any other triad.

In convex environments, *preference range* leads uniformly to lower exploration. The reason for it, the lack of differentiation of the homogeneous triad between the attributes, is clear, and therefore the finding does not nullify the results of Figure 10 or the core of the paper. However, it does add nuance. In the absence of knowledge of the environment's form (concave or convex), diversity could lead to even more variance as previously anticipated. In convex environments, an organization designer aiming to increase exploration has a simple job; it should deter diversity at any cost. Homogeneity in the convex environment will lead to more exploration.

The results of Figure A.13 happen only if the environment is convex. I have studied other concave environments, and the results always resemble Figure 10. The environments that

I have explored include an unbounded environment that allows x_1 and x_2 to vary between 0 and 1 without any other constraint. This environment looks like a square of area one. There, all agents would find the corner of [1,1] as giving the most valuable options. The environment is still single-peaked, but for most agents, the peak will be in the same place, the corner of [1,1]. This is very different from previous environments. However, the results in terms of the amount of exploration are equivalent to the ones shown in Figure 10. I also studied environments where the options are drawn only from the line defining the production possibility frontiers of Figure 5 and Figure A.12. The rationale is that in a competitive environment, only firms along the frontier should exist. There can be some fluctuations, as shown by Adner et al. (2014), but the ones closest to the frontier should be survive in the long term.

6.7 Different learning processes

Since the first study on reinforcement learning in organizational learning (Denrell & March, 2001), the field has used several different algorithms to operationalize the updating and the proposal stages of the reinforcement learning process. Denrell and March (2001) use an aspiration-level search process where agents update only one belief every period, their aspiration level. If the option outputs a payoff higher than the aspiration level, then the probability of choosing the option is increased. Instead, if the payoff was lower than the payoff, the probability of choosing the option is decreased. The aspiration level is also updated after the payoff is received, up if the payoff was higher than the aspiration, down otherwise. This update and proposal process has the benefit that it has just one belief, the aspiration level for the environment. It directly updates probabilities, so it does not require a function to propose an option. It can just draw from a Polya urn.

Figure A.14 shows the Denrell and March (2001) aspiration-level search process. The results change. Here, I find that *preference bias* still lowers exploration, but the *preference range* does not increase exploration. Similar to the case of a convex environment, an organization designer aiming to increase exploration would see a negative prospect in allowing diversity in the organization. This designer would then have a better choice in decreasing diversity in the firm as that will lead to high levels of exploration and avoid any requirements in managing the *preference bias* in the organization.

Figure A.14: Percentage of exploration in triads with the aspiration-based learning process

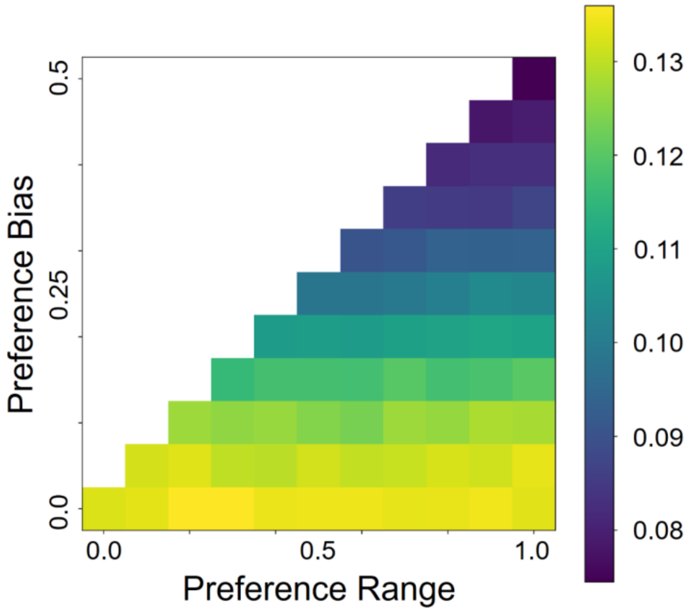

Denrell and March (2001) introduced the N-arm bandits to organizational learning. However, others have used it since. Posen and Levinthal (2012) introduce several different learning models in their seminal paper. They make two changes to the Denrell and March (2001) model. They have a new update function; in their case, there are beliefs about the value of every option, not just the aspiration level. The beliefs on each option are updated after the option is chosen and the payoff received. Posen and Levinthal (2012) do not use a constant learning rate, ϕ; instead, they have a learning rate that varies as ϕ = 1/(k+1). So that the first time an option is chosen, 50% of the beliefs are updated, the second time 1/3 of the beliefs, and so on. This should not be problematic if the environment is static, but if the options change, it should not work well (Laureiro-Martinez & Brusoni, Tata, & Zollo, 2019).

The second change in Posen and Levinthal (2012) was the proposal procedure; they introduce three different proposal procedures: greedy, ϵ-greedy, and softmax. In this paper, greedy algorithms do not work. If I initialize the beliefs in the same way, the organizations end up choosing at random even with long learning epochs. I can use the ϵ-greedy algorithm, and for doing this, I first need to find the optimal learning parameters. To find the optimal learning parameters, I did not use a variable learning rate as in Posen and Levinthal (2012); instead, I optimized both the ϵ and ϕ parameters. The highest performance of this algorithm comes when $\phi = 0.75$ and the greedy option is chosen 75% of the time. Using these two parameters with a ϵ-greedy selection process, I obtain the results of Figure A.15.

Figure A.15: Percentage of exploration in triads with the $\boldsymbol{\epsilon}$-greedy proposal process

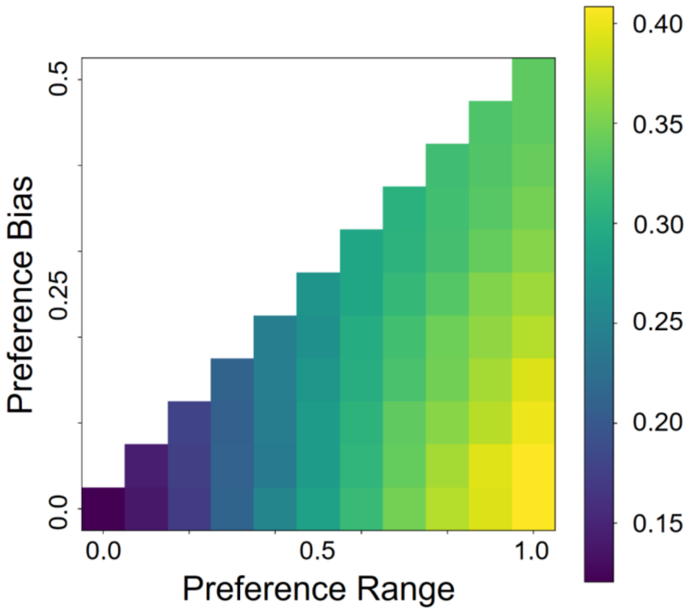

Figure A.15 shows a similar process as in the results shown in the main body of the paper. Namely, a higher *preference range* can lead to higher exploration, but *preference bias* lower these increases in exploration rate. However, the *preference bias* has a lower effect when the selection is made via an ϵ-greedy algorithm. Posen and Levinthal (2012) introduced the greedy and ϵ-greedy algorithms, but most of their paper focused on the softmax algorithm. The softmax algorithm, coupled with the constant learning rate updating function, has become predominant in reinforcement learning research within the organizational learning community (Puranam & Swamy, 2016). *For this reason, the constant-learning rate updating function and the softmax proposal algorithm is the combination I use in this paper.*

6.8 Interpersonal comparison of utilities

The comparison of utility between agents provides a weak foundation for social choice models (Arrow, 1951:3). In the paper, I do not compare utilities and present only the effects on the joint exploration of options as the study's main dependent variable. But as a robustness check, it is important to know whether the agents' actions are rational even in a limited way.

Figure A.16 shows the ratio of the total accrued utility of Agent A divided by the total accrued utility of Agent B. Note that both utilities are built differently as the agents prefer different attributes. However, from Figure A.16, we can observe that as the preference of Agents A and C get closer (i.e., to the diagonal in Figure A.16), the ratio starts to increase. In fact, the ratio reaches a maximum of 2.07 in the *biased triad.* This means that Agent A derives two units of utility for every unit of utility that Agent B derives from the learning process.

Figure A.16: Ratio of accrued utilities of Agent A divided by Agent B

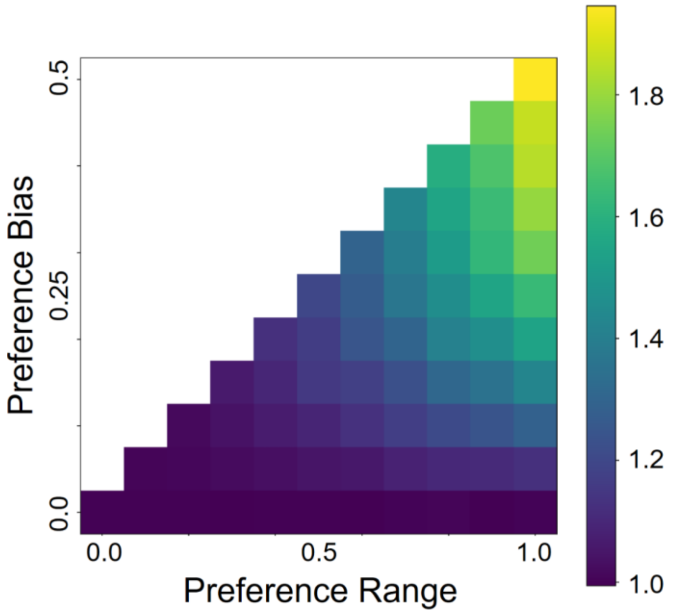

The total average utility of Agent B in the *biased triad is 18%* higher than the utility accrued by Agent B when it is part of a *homogeneous triad.* This is an important point as it shows that in some way, Agent A is acting in a way that gives them higher utility, and the behavior is not just detrimental to the other agent. This is shown more clearly in Figure A.17.

Figure A.17 presents the average sum of utilities of the three agents, i.e., organizational utility. That is the average utility that the agents in the triad accrued. Again, this makes the assumption that one unit of utility is valued the same amount by another agent. I highlight three aspects of Figure A.16. First, the average utility that the different organizations, i.e. pixels, accrue does not vary greatly. The highest organizational utility is 3.86% higher than the lowest organizational utility. Second, the organization achieves higher utility at lower levels of *preference range*. This is understandable as in those conditions, more agents are benefited from a common choice pattern. Finally, at a higher *preference range*, the triads that achieve the highest utility are the ones that have high *preference bias*. Unbiased triads (the ones along the bottom row) achieve a lower organizational utility than their more biased counterparts.

Figure A.17: The summed average accrued utility of all agents in the organization

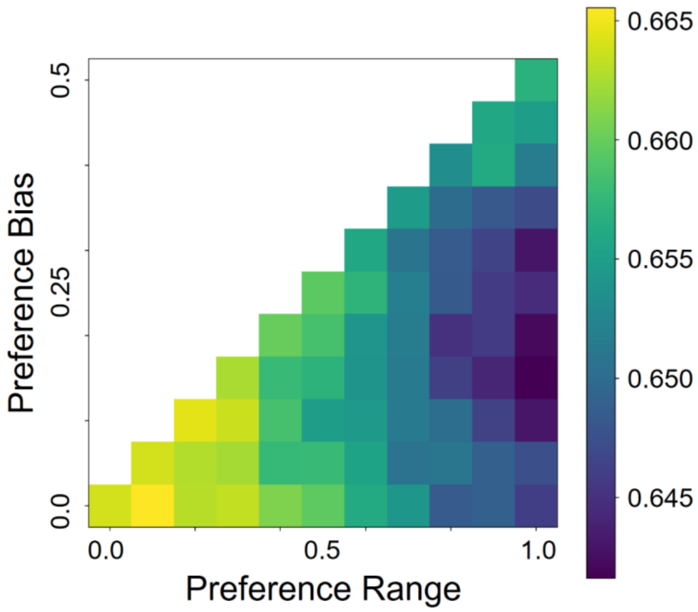

Note that this section's findings need to be taken with a grain of salt as they rely on comparisons of the agents’ utilities (Arrow, 1951:3). However, they serve as a good test to understand whether the agents' actions follow common sense. In general, organizational utility is not heavily affected by *preference diversity*. When possible, an agent will benefit the most from choosing its own selfish option than choosing a more balanced solution—both within the expectations of self-interested agents.

REFERENCES

Adner, R., Csaszar, F. A., & Zemsky, P. B. (2014). Positioning on a multi-attribute landscape. *Management Science*, 60(11), 2794-2815.

Anderson, P. W. (1972). More is different. *Science,* 177(4047), 393-396.

Argote, L. (2013). *Organizational learning*. Springer, Boston

Arrow, K. J. (1951). *Social Change and Individual Values*.

Baron, D. P. (1991). A spatial bargaining theory of government formation in parliamentary systems. *The American Political Science Review*, 137-164.

Baumann, F., Lorenz-Spreen, P., Sokolov, I. M., & Starnini, M. (2020). Modeling echo chambers and polarization dynamics in social networks. *Physical Review Letters*, 124(4), 048301.

Boone, C., & Hendriks, W. (2009). Top management team diversity and firm performance: Moderators of functional-background and locus-of-control diversity. *Management science*, 55(2), 165-180.

Butler, J., Morrice, D. J., & Mullarkey, P. W. (2001). A multiple attribute utility theory approach to ranking and selection. *Management Science,* 47(6), 800-816.

Chandler Jr, A. D. (1977). *The Visible Hand,* Cambridge, Mass. and London, England.

Christensen, M., & Knudsen, T. (2010). Design of decision-making org. *Management Scien*ce, 56(1), 71-89.

Congleton, R. D. (2004). The median voter model. In *The encyclopedia of public choice* (pp. 707-712). Springer, Boston, MA.

Csaszar, F. A. (2013). An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organization Science*, 24(4), 1083-1101.

Csaszar, F. A., & Levinthal, D. A. (2016). Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37(10), 2031-2049.

Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. Englewood Cliffs, NJ, 2(4), 169-187.

Dasgupta, P., & Maskin, E. (2008). On the robustness of majority rule. *Journal of the European Economic Association*, 6(5), 949-973.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12), 1704-1711.

Denrell, J., Fang, C., & Levinthal, D. A. (2004). From T-mazes to labyrinths: Learning from model-based feedback. *Management Science,* 50(10), 1366-1378.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science,* 12(5), 523-538.

Ely, R. J., Thomas, D. A. (2020). Getting Serious About Diversity: Enough Already with the Business Case. *Harvard Business Review,* November-December Issue

Friedman, M. (1953). *The methodology of positive economics. Essays in positive economics,* 3(3), 145-178.

Gaba, V., & Greve, H. R. (2019). Safe or profitable? The pursuit of conflicting goals. *Organization Science,* 30(4), 647-667.

Gresov, C., & Drazin, R. (1997). Equifinality: Functional equivalence in organization design. *Academy of management review,* 22(2), 403-428.

Ganz, S. C. (2020). *Conflict, Chaos, and the Art of Institutional Design*. Open Science Framework. Retrieved from: https://osf.io/qjn5y/

Greve, H. R. (2007). Exploration and exploitation in product innovation. *Industrial and Corporate Change,* 16(5), 945-975.

Holcombe, R. G. (2006). Public sector economics: The role of government in the American economy. Prentice Hall.

Kamada, Y., & Kojima, F. (2014). Voter preferences, polarization, and electoral policies. American Economic *Journal: Microeconomics*, 6(4), 203-36.

Kaplan, S. (2008). Framing contests: Strategy making under uncertainty. Organization Science, 19(5), 729-752.

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. Academy of Management Annals, 14(1), 366-410.

Knight, F. H. (1921). *Risk, uncertainty and profit (Vol. 31).* Houghton Mifflin.

Knudsen, T. (personal communication, November 28, 2019)

Kollman, K., Miller, J. H., & Page, S. E. (1997). Political institutions and sorting in a Tiebout model. *American Economic Review*, 977-992.

Krehbiel, K. (1998). *Pivotal politics: A theory of US lawmaking.* University of Chicago Press.

Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration-exploitation dilemma: An fMRI study of attention control and decision-making performance. *Strategic Management Journal*, 36(3), 319-338.

Laureiro-Martinez, D., Brusoni, S., Tata, A., & Zollo, M. (2019). The manager's notepad: Working memory, exploration, and performance. *Journal of Management Studies*, 56(8), 1655-1682.

March, J. G. (1962). The business firm as a political coalition. *The Journal of politics*, 24(4), 662-678.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71-87.

March, J. G. (2010). *The ambiguities of experience*. Cornell University Press.

Murphy, R. O., Ackermann, K. A., & Handgraaf, M. (2011). Measuring social value orientation. *Judgment and Decision making*, 6(8), 771-781.

Murray, A., Rhymer, J., & Sirmon, D. G. (2020). Human and Technology: Forms of Conjoined Agency in Organizations. *Academy of Management Review*, Forthcoming.

Ocasio, W., & Thornton, P. H. (1999). Institutional logics and the historical contingency of power in organizations: Executive succession in the higher education publishing industry, 1958-1990. *American Journal of Sociology*, 105(3), 801-843.

Opper, M., & Saad, D. (Eds.). (2001). *Advanced mean field methods: Theory and practice*. MIT press.

O'Reilly III, C. A., & Tushman, M. L. (2013). Organizational ambidexterity: Past, present, and future. *Academy of management Perspectives*, 27(4), 324-338.

Page, S. E. (2010). *Diversity and complexity (Vol. 2).* Princeton University Press.

Piezunka, H., Aggarwal, V. A., & Posen, H. E. (2020). *Learning-by-Participating: The Dual Role of Structure in Aggregating Information and Shaping Learning.* SSRN, Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3425696

Posen, H. E., & Levinthal, D. A. (2012). Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science*, 58(3), 587-601.

Porter, M. E. (1996). What is strategy?. *Harvard Business Review*

Puranam, P. (2018). *The microstructure of organizations.* Oxford University Press.

Puranam, P., Stieglitz, N., Osman, M., & Pillutla, M. M. (2015). Modelling bounded rationality in organizations: Progress and prospects. *Academy of Management Annals*, 9(1), 337-392.

Puranam, P., & Swamy, M. (2016). How initial representations shape coupled learning processes. *Organization Science*, 27(2), 323-335.

Raisch, S., & Krakowski, S. (2020). Artificial Intelligence and Management: The Automation-Augmentation Paradox. *Academy of Management Review*, Forthcoming.

Ramus, T., Vaccaro, A., & Brusoni, S. (2017). Institutional complexity in turbulent times: Formalization, collaboration, and the emergence of blended logics. *Academy of Management Journal*, 60(4), 1253-1284.
Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature reviews neuroscience*, 9(7), 545-556.
Rerup, C. (2009). Attentional triangulation: Learning from unexpected rare crises. *Organization Science*, 20(5), 876-893.
Rerup, C., & Feldman, M. S. (2011). Routines as a source of change in organizational schemata: The role of trial-and-error learning. *Academy of Management Journal*, 54(3), 577-610.
Rerup, C., & Zbaracki, M. J. (2020). The politics of learning from rare events. *Organization Science.* Forthcoming
Sah, R. K., & Stiglitz, J. E. (1986). The architecture of economic systems: Hierarchies and polyarchies. *The American Economic Review,* 716-727.
Salvato, C., & Rerup, C. (2018). Routine regulation: Balancing conflicting goals in organizational routines. *Administrative Science Quarterly*, 63(1), 170-209.
Shore, L. M., Chung-Herrera, B. G., Dean, M. A., Ehrhart, K. H., Jung, D. I., Randel, A. E., & Singh, G. (2009). Diversity in organizations: Where are we now and where are we going?. *Human resource management review*, 19(2), 117-133.
Simmel, G. (1995). Soziologie des Raumes. In Ders.*, Aufsatze und Abhandlungen 1901-1908,* Suhrkamp Verlag.
Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations: Volume One*. London: printed for W. Strahan; and T. Cadell, 1776.
Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning (Vol. 135).* Cambridge: MIT Press
Tata, A., & Niedworok, A. (2020). Is beauty in the eye of the beholder? An empirical study of how entrepreneurs, managers, and investors evaluate business opportunities at the earliest stages. *Venture Capital*, 22(1), 71-104.
Tolstoy, L. N. (1877). *Anna Karenina*
Weick, K. E. (1995). *Sensemaking in organizations* (Vol. 3). Sage.

PAPER 3

"What kind of a situation is this? What kind of a person am I?
What does a person such as I do in a situation such as this?"
– March and Olsen (2011)

In Search of Contrarian Opportunities from the Blind Spot of Majority Rule

Jose P. Arrieta[1] and Chengwei Liu[2]
[1]ETH Zürich, Switzerland, [2]ESMT Berlin, Germany

ABSTRACT

Superior profit usually depends on capturing opportunities that rivals fail to identify or utilize. A key challenge for strategists is how to be both different and viable. Prior research has tended to associate contrarian opportunities with rivals' behavioral failures. Herein, we argue that contrarian opportunities can emerge endogenously in an ecology whenever there is a dominant logic. We develop our argument in the context of organizational design, in which the majority voting rule is demonstrated to be an efficient and typically mainstream approach for screening alternatives. We formally demonstrate when antimajority—an unconventional screening rule where acceptance depends on the minority's approval and majority's disapproval—exploits the opportunities left behind by the majority rule. We illustrate how a contrarian niche emerges, and its scope conditions using the case of an antimajority voting venture capitalist firm together with an evolutionary model of competing rules. More generally, a contrarian niche emerges not necessarily because the dominant firms have been suboptimal or inefficient but because their homogeneity predicts an exploitable blind spot, preserving opportunities for strategists who can afford to be contrary.

Keywords: contrarian niche, organizational design, commission and omission errors, antimajority, ecology of competition, mixed methods

1. INTRODUCTION

Strategizing is about being different (Barney, 1991; Peteraf, 1993; Porter, 1996). However, identifying novel opportunities that are both different and viable is challenging (Pontikes & Barnett, 2017), partly because most good ideas are not new and most new ideas are not good (Levinthal & March, 1993). Strategy scholars have proposed various mechanisms that explain why some individuals or firms can identify profitable opportunities that others fail to see. For example, superior profit may result from having superior foresight and intelligence (Barney, 1986; Csaszar & Laureiro-Martínez, 2018; Levine, Bernard, & Nagel, 2017), superior understanding of resources' (re)combinatory values (Denrell, Fang, & Winter, 2003; Dierickx & Cool, 1989; Lippman & Rumelt, 2003), superior adaptive capability (Leiblein, Chen, & Posen, 2017; Teece, Pisano, & Shuen, 1997), superior learning and mental representation capacity (Csaszar & Levinthal, 2016; Gavetti, Levinthal, & Rivkin, 2005), or simply being at the right place at the right time (Denrell, 2004; Liu, 2019). The implication is that rivals' incompetence, behavioral failures, or bad luck are necessary for the presence of strategic opportunities (Denrell, Fang, & Liu, 2019; Gavetti, 2012).

Herein, we argue for an alternative source of superior profit without assuming rivals' suboptimalities. An effective strategy can create contrarian opportunities in an ecology of competing organizations when the strategy is overly subscribed. For example, a simple decision rule may be a smart heuristic that exploits environmental regularities (Davis, Eisenhardt, & Bingham, 2009; Gigerenzer & Todd, 1999). However, a wide adoption of such a rule to simplify complexity suggests that some opportunities will be overestimated (if recommended by the simple rules) while others will be underestimated (if overlooked by the simple rules). An "ecologically rational" simple rule can be successful in its own right while creating a contrarian niche because its efficiency attracts the majority to interpret the environment similarly (Gavetti & Porac, 2018; Porac, Thomas, Wilson, Paton, & Kanfer, 1995; Todd & Gigerenzer, 2012). More generally, we argue that a contrarian niche emerges whenever there exists a dominant logic of competing or organizing at the ecological level. This implies that strategists should search for strategic opportunities from not only the flaws in rivals' logic, as prior studies have suggested (Felin & Zenger, 2017), but also from situations where their logic is so impeccable that it generates a similar-minded majority.

We developed our theory in the context of organizational design, which is essential for balancing the inescapable errors of commission (accepting bad ideas) and errors of omission (rejecting good ideas) when screening proposals or candidates (Christensen & Knudsen, 2010). Previous research has demonstrated that hierarchical screenings (i.e., a project is rejected if anyone rejects it) tend to minimize commission errors but increase omission errors and vice versa for polyarchical screening (i.e., a project is accepted if anyone accepts it) (Sah and Stiglitz, 1986). Organizational structure with a majority rule can strike a balance between the two errors and is an efficient choice for organizational screening in many situations (Csaszar, 2013). Argote (2013) explains that when organizations face situations "without a demonstrably correct answer, a majority… decision scheme characterizes how groups make decisions" (p.134). Majority decision rules thus are a common way for organization to aggregate the different opinions of their members and find a way forward (Csaszar & Eggers, 2013). Similarly, in social choice theory, the majority voting rule is demonstrated as both the more efficient and more robust rule for aggregating the decisions of heterogeneous agents (Black, 1948; Dasgupta & Maskin, 2008). However, efficiency, commonality, and robustness can still allow imperfection (Arrow, 1951). The aggregation of omission errors by many majority voting firms can become an exploitable blind spot. The challenge is then how to search for these opportunities that the majority tends to overlook.

We formally demonstrate that an overlooked screening rule—antimajority—may help aspiring strategists to identify and exploit the blind spots of the majority rule. The antimajority rule privileges the minority view but also utilizes the majority view. Under antimajority, a candidate or proposal is accepted when the majority of committee members reject it, and a minority of committee members accepts it (hence, different from the polyarchical rule where the majority view is irrelevant). The way antimajority exploits the majority rule is explored using three approaches: formal analysis, a case study, and an evolutionary modeling. First, we build on and extend the canonical models in organizational design literature to formally demonstrate how antimajority differs from other screening rules (Christensen & Knudsen, 2010; Csaszar, 2013). If we assume that all evaluators have identical informational access as the prior models did, the antimajority rule does not have any advantage. Nevertheless, if we relax this assumption and allow candidates or projects to have multiple attributes while evaluators have differential access to these attributes, the antimajority rule can guide organizations to identify a very different set of opportunities that other screening rules overlook. Second, the VC case suggests that the antimajority rule works for organizations which: (a) have the goal to capture the upside extreme and can afford low probability of success; (b) operate in contexts with high uncertainty in evaluations; (c) have access to a very diverse pool of ideas; and (d) have members sharing aligned interests in finding a promising contrarian alternative, mitigating the risk of strategic voting. Finally, inspired by the VC case, we build an evolutionary model that allows simulated VCs with different decision rules to compete, in which those that manage to identify more profitable startups more likely to survive and reproduce. The simulation results show the scope conditions under which VCs with antimajority rule can invade and persist in an ecology dominated by other rules, including majority rule and consensus rule. Our results show that antimajority thrives because it guides firms to identify opportunities left behind by firms that follow the dominant logic. This also suggests a limit of antimajority: its growth ceiling is proportional to the popularity of the dominant logic. Stated differently, the antimajority rule is not for the majority but a decision model for a contrarian minority who are prepared to bear the risk when exploiting the majority.

Our findings have important implications for the search for strategic opportunities and entrepreneurship (Denrell et al., 2003; Felin & Zenger, 2017; Gavetti & Porac, 2018). The formal and evolutionary analyses suggest that antimajority is theoretically viable for exploiting opportunities overlooked by dominant screening rules. Our VC case illustrates why the antimajority is overlooked in both the literature and practice—several necessary conditions for its success are difficult to achieve in reality. These difficulties make the associative opportunities attractive precisely because they protect these opportunities from being discovered and exploited (Gavetti, 2012; Liu, 2020). More generally, our results shed light on an alternative source of profit: a contrarian niche emerges not necessarily because the dominant firms were suboptimal or inefficient but because their homogeneity predicts exploitable blind spots, preserving opportunities for entrepreneurs who can afford to be contrary.

The structure of the paper is as follows. We first review and extend the canonical models to formally demonstrate how the antimajority rule is different from alternative screening rules. A case from the VC industry is then utilized to illustrate how the antimajority helps to identify the next big startup. After formalizing the scope conditions of this unconventional screening rule derived from the case, an evolutionary model is presented to computationally examine these conditions under which the antimajority can successfully exploit alternative rules. We conclude by discussing the implications of our findings for organizational design, behavioral strategy, and entrepreneurship.

2. ANTIMAJORITY: AN OVERLOOKED BRANCH IN THE ORGANIZATIONAL DESIGN MODELS

Organizations employ specifically designed decision structures to minimize their commission and omission error rates when screening options (Csaszar, 2012, 2013; Romme, 2004). For this purpose, Christensen and Knudsen (2010:77) updated the Moore Shannon theorem and created a two-stage algorithm to minimize organizational decision errors to any arbitrary level; a process referred to as "approaching perfection" (Moore & Shannon, 1956). In the first stage, the Christensen-Knudsen algorithm removes the decision bias of the agents. In the second stage, decision errors are minimized through a majority voting rule.

In particular, the second stage of the Christensen-Knudsen algorithm hinges on the majority-voting triad. The majority voting triad is the smallest decision structure that allows both commission and omission error rates to be lower than the error rates of an individual rater. Larger majority voting structures can lower both errors further but are costlier to maintain and take longer to reach a decision (Sah & Stiglitz, 1988). In contrast, consensus structure decreases the commission error rate at the expense of a higher omission error rate.

2.1 Screening Functions with Homogeneous Mental Representations

The Christensen-Knudsen algorithm is built for homogeneous agents, i.e., agents evaluate options in exactly the same way. These agents achieve the same scalar acceptance probabilities ($x = f(\vec{q})$)[1] to each option ($\vec{q} = [q_1, q_2, \dots q_M]$). Figure 1 presents all the decision flows in a formulation equivalent to Christensen and Knudsen (2010). The circles represent agents; full lines represent options that are voted in favor, and dashed lines represent options voted against. The probability of any agent voting in favor of an option is given by x, and the probability of voting against is given by $1 - x$. From this, we can estimate the probability that an option will achieve any number of votes in its favor.

Figure 1: Graph of Decision Structures.

[1] We explain how the scalar acceptance probabilities are calculated in later section "Diverse Screening Functions".

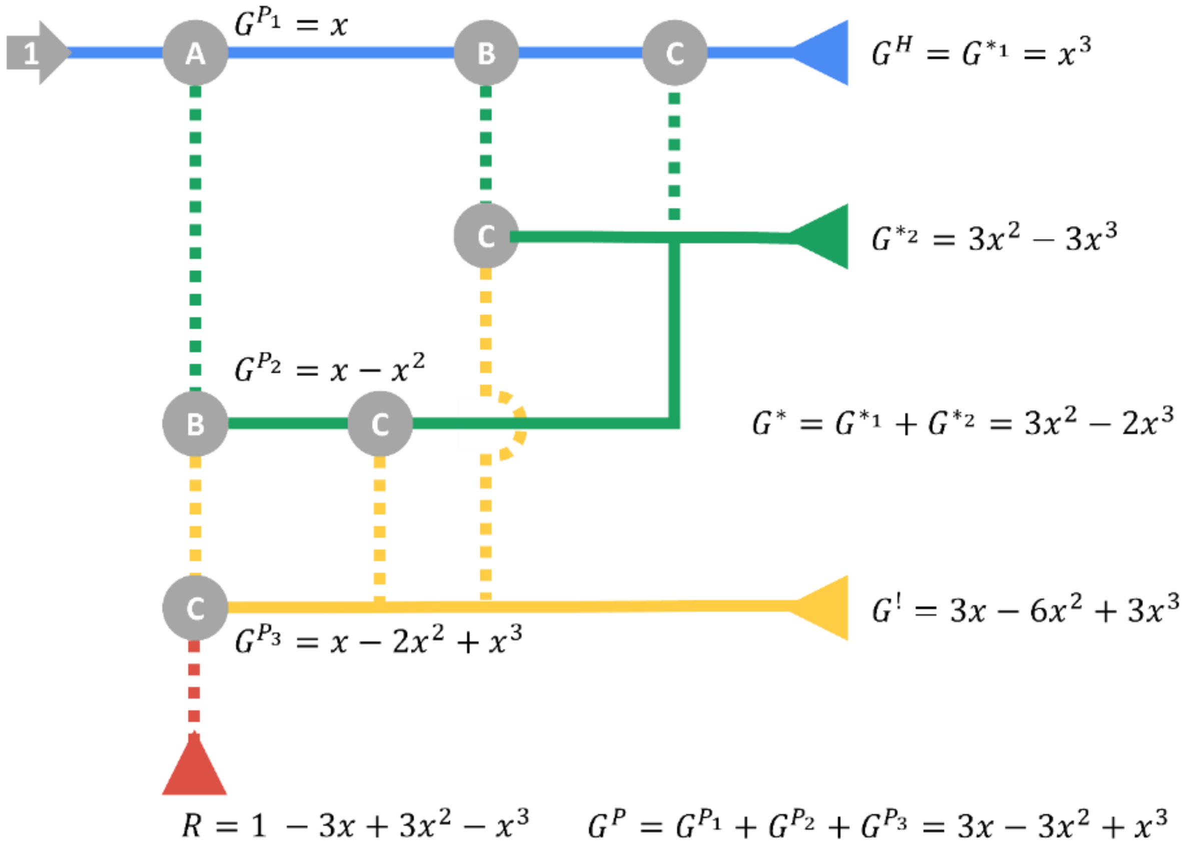

The blue line of Figure 1 follows a decision structure where options are accepted only if the three agents vote in favor. Next, the green lines represent the paths that lead to two agents voting in favor of an option. Further, the yellow lines represent the case where only one agent votes in favor of an option. Finally, the red line represents the case where no agent votes in favor of the option.

These lines and their combinations then determine the decision structure as well as the organization's acceptance probability (G^d, where d is a specific decision structure). A hierarchy decision structure (G^H) invests in an option only if all agents vote in favor (blue lines). Additionally, a majority voting decision structure (G^*) invests in an option if two or more agents vote in favor (blue plus green lines). In comparison, a polyarchical decision structure (G^P) will invest if one or more agents vote in favor (blue, green, and yellow lines combined). Comparatively, an antimajority voting structure ($G^!$) invests only if one agent votes in favor (yellow lines). Note that the investments of the polyarchical decision structure (G^P) are the sum of the accepted options from the majority voting decision structure (G^*) and the antimajority voting structure ($G^!$). Finally, the remaining options are left with a decision structure that invests only if everyone votes against the option (*R*, red line). Table 1 includes the polynomial screening function for each decision structure. It also includes the multinomial expansion of these screening functions, which allow each option to have multiple attributes to be evaluated, a point we will revisit shortly.

Figure 2 plots the five screening polynomials of Table 1. The *x*-axis of Figure 2 is the scalar acceptance probability of the agents, $x = f(\vec{q}) \epsilon [0,1]$. The grey line in Figure 2 plots this scalar acceptance probability. The other lines represent the screening polynomials of the other decision structures in Table 1. In black, Figure 2 also plots a Heaviside step function—the ideal screening function (Christensen and Knudsen, 2010). From Figure 2, we can observe clear features of why the different decision structures decrease or increase the error rate in comparison to an individual. For example, the polyarchy (G^P) has a higher acceptance rate at every point than one individual. This will lead polyarchies to invest in many more bad options (increased commission errors) but also to invest more often in good options (decreased omission error). The majority voting screening function (G^*) has an inflection point. When $x <$ 0.5, it lies below the individual screening function; when $x > 0.5$, it lies above. This property allows the majority voting to have both lower commission and omission errors than the individual.

Table 1: Polynomial and multinomial screening functions for decision structures with three agents.

Decision Structure	# of Votes in Favor	Polynomial	Multinomial
Hierarchy: G^H	3	x^3	xyz
Majority: G^*	> 2	$3x^2 - 2x^3$	$xy + xz + yz - 2xyz$
Polyarchy: G^P	> 1	$3x - 3x^2 + x^3$	$G^* + G^!$
Antimajority: $G^!$	= 1	$3x - 6x^2 + 3x^3$	$x + y + z - 2xy$ $- 2xz - 2yz + 3xyz$
Rest: R	= 0	$1 - 3x + 3x^2 - x^3$	$1 - G^P$

Figure 2: Plot of screening polynomials.

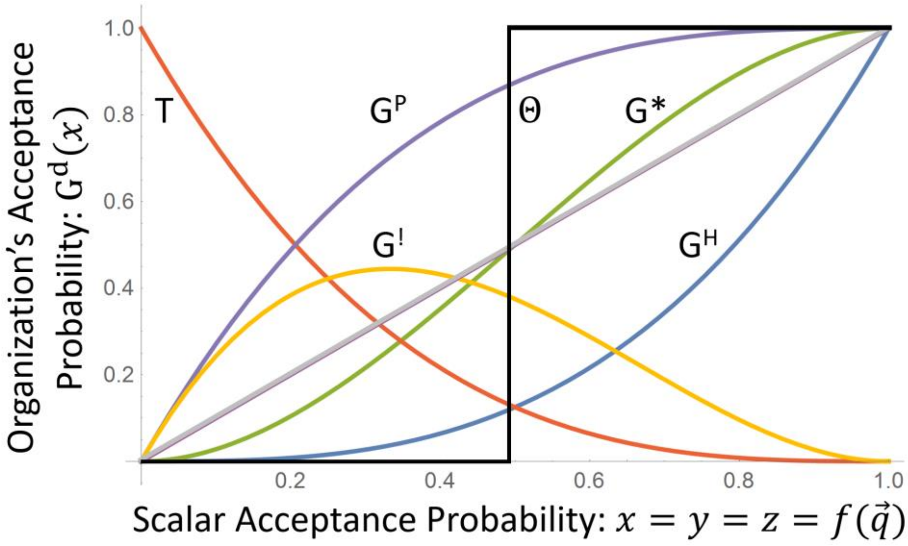

The antimajority voting screening function ($G^!$) differs from the other screening functions. It has a high probability for investment in options that a single agent would dismiss (low x) while it also has low probability of investing in an option in which an agent would invest most of the time (high x). In general, if options were one-dimensional, the antimajority would perform poorly, achieving both high commission and high omission error rates.

2.2 Screening Functions with Diverse Mental Representations

The advantage of the antimajority emerges when we relax the assumption of homogeneous mental representations. If we assume that all agents share their mental representations of the value accrued from each attribute of an option, the quality of an option can be simplified into one quality shared by all agents (Sah and Stiglitz, 1986; Csaszar, 2013). However, agents can vary in their mental representations of how quality maps to the different attributes of an option (Csaszar and Levinthal, 2016). In the extreme, if an option is described by M values, i.e., $\vec{q} = [q_1, q_2, \ldots q_M]$, then it is possible to create M mental representations that are independent of each other. In contrast, if the agents share part of their representations, then the number of independent representations will be lower than M. In this paper, we study decision structures composed of three agents, implying that at most, we will have three independent mental representations. These independent mental representations give rise to representational diversity, an aspect not considered when we assume that all agents have homogeneous mental representations.

If we now study the case of organizations with three agents, each with its unique mental representation, we can label their acceptance probabilities as: *x* for agent A, *y* for Agent B, and *z* for Agent C. The agents would agree on the value, $\vec{q}$, of each attribute of an option but disagree on how the values merit investment. The use of different scalar acceptance probabilities requires the use of multinomials to describe the screening functions of each decision rule. The main modification is that the transition probabilities between one rather and the next change, instead of having an x percent chance of being accepted on each circle in Figure 1, the probability depends on whether the decision is made by Agent A, B, or C (Christensen and Knudsen, 2010). The multinomial expression of the screening functions of each decision rule is shown in the right column of Table 1[2].

Commission and omission error rates exist even when the agents have diverse representations. The possibility of having independent scalar acceptance probabilities allows us to explore how this diversity can lead to blind spots for each decision rule. In later section on "Antimajority and its Scope Conditions", we will present how to operationalize diverse mental representations in detail. For now, we assume that we can impute agents with different mental representations.

Figure 3 plots contour maps that show when a screening function (majority, antimajority, or polyarchy) will have a high chance of investing in an option (lighter color) or lower chances of investing in an option (darker color). In this simulation, the first and third agents shared the same mental representation (namely x = y), which is different from that of

[2] To calculate the polynomials, we average the organizational acceptance probabilities of having each agent at every position. This implies that the order of A, B, and C are exchanged in all possible ways and then the acceptance probabilities averaged out. In some cases, as in the consensus, but its important in other but its important in others structures, such as the polyarchy.

the second agent (x = z ≠ y). This simplification allows us to plot the organizational acceptance probabilities in a two-dimensional space. The *x*-axis represents the probability of the first and third agents voting in favor of the option. The *y*-axis stands for the probability of the second agent voting in favor of the option.

Figure 3: Comparison of majority and antimajority screening multinomials with two agents out of three sharing their scalar acceptance probability

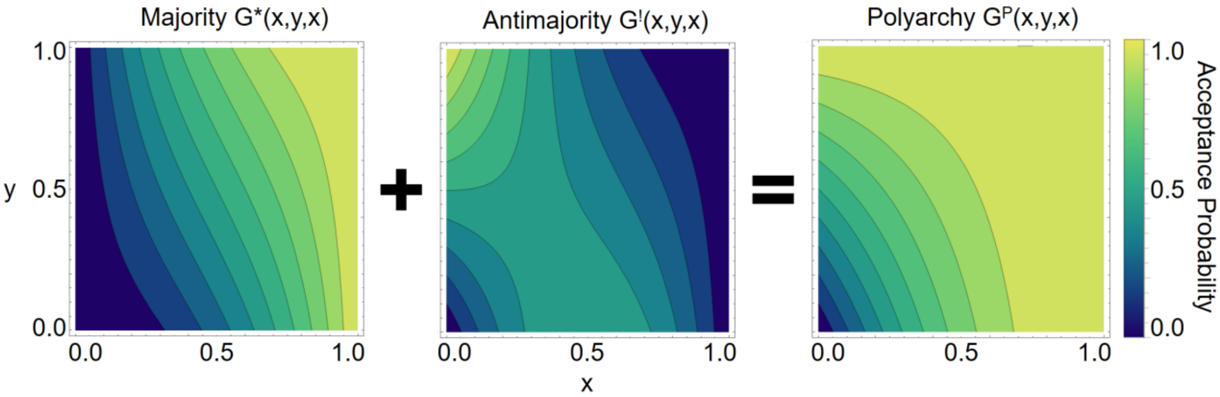

In the majority voting case (G*), high organizational acceptance probabilities appear in the upper right corner, as all agents there have a higher scalar acceptance probability of voting in favor of the option (i.e., x and y are high). We also find that high *x* alone is enough for the decision structure to invest in an option. The high investment probability emerges because two agents share the value of *x* and their vote is enough to get a simple majority.

The probability of investing in an option looks very different in the case of antimajority voting (G!). Antimajority voting leads to high probability of investing only when *x* is low and *y* is high, a corner case where majority voting has a low investment probability. The high investment probability happens because the top left corner is the only case where one single agent finds an option valuable. The corner on the right features multiple agents valuing the option, thus leading the organization to disregard the option. Additionally, when *y* is low, and *x* is around ½, the antimajority has about a 50% chance of investing in an option; here, the investment is not led by the second agent (who saw a low value of *y*) but by one of the other two agents who happened to vote positively. To avoid investing on a mediocre option, this is a danger zone the antimajority needs to avoid, a point we will revisit in the next section. Finally, the polyarchy organization (G^P) adds the results from the prior two screening functions. Polyarchy organizations invest when both *x* and *y* are high but also whenever one of them is high.

2.3 Antimajority as a Niche Strategy

Figure 3 highlights the difference between the majority and antimajority voting rules. These two screening functions are mostly anticorrelated—when one invests, the other mostly avoid investing. This separation allows us to posit that if polyarchy is too costly to implement and an industry is thereby composed of mostly majority voting decision structures, an antimajority voting structure might profit from options that are left behind.

Importantly, several conditions are required for antimajority to be a viable strategy. First, antimajority requires some independence between the mental representations of the evaluators. In the presence of homogeneous mental representations, the result regresses to the case of Figure 2, where both the commission and omission error rates of the antimajority are much higher than any other structure.

Antimajority also requires a higher amount of cooperation and trust between the agents in the organization. These requirements are not obvious from the analytical equations, but in organizations where agents interact (Lorenz, Rauhut, Schweitzer, & Helbing, 2011), social influence can destroy the independence condition that is necessary for antimajority and lead to strategic voting (Ganz, 2020; Ludwin, 1978). For example, if Agent A knows that Agent B will vote to invest in the option, and Agent A really dislikes the option, Agent A can veto the investment by also voting in favor of the option. Antimajority voting will only work if such “strategic voting” behavior is minimized. Consensus voting, majority voting, and polyarchies do not suffer from this limitation because, in all these decision rules, the agent cannot do better than showing their true intentions (Dasgupta & Maskin, 2008). The liability to strategic voting requires antimajority voting organizations to have agreements between the agents to trust the decisions of the others. This agreement is unnecessary in organizations that follow other decision rules.

These added requirements suggest that contrarian strategies are necessarily more complex than dominant strategies, making them harder to find but also harder to imitate if successful (Ethiraj & Levinthal, 2004). For example, in contrast to antimajority voting, a majority voting firm does not require high levels of trust between the agents. Independence between the agents can be useful if the option is novel, as it will accrue the Wisdom of the Crowd effects; however, if the option considered comes from a stable environment, the agents could share their views of the world without affecting the performance of the organization (Hong & Page, 2004; March, 1962). In the next section, we present a case of a VC firm that employs antimajority voting and computationally examine the key characteristics identified from the case in an evolutionary model that follows.

3. ANTIMAJORITY IN PRACTICE: A CASE ILLUSTRATION FROM THE VC INDUSTRY

Our formal analysis in the last section highlighted how antimajority differs from other rules, particularly against the majority rule. However, the analysis is silent about when an antimajority can be profitable. Draper Fisher Jurvetson (DFJ), a US VC firm based in Menlo Park, California, exemplifies the conditions under which an antimajority rule can be useful, such as enhancing the chance of identifying the next big thing (Liu, Vlaev, Fang, Denrell, & Chater, 2017). In particular, when DFJ entered the nanotechnology field in early 2000, they adopted an antimajority rule during the final stage of selection; DFJ invested in a startup when one partner out of three felt very passionate about the idea, whereas the remaining partners were reluctant if not against the idea. This strategy has led to impressive successes (Bohman, 2009), and there are at least four factors that enabled antimajority to work for the DFJ.

The decision to invest through antimajority voting was strategic and the output of deliberate choices by the founders of DFJ (Liu et al., 2017). The founders chose antimajority after evaluating how to be successful in the nanotechnology market. Antimajority emerged as a valuable solution for several reasons. First, the founders of DFJ chose to enter a field where the evaluation of business proposals is highly uncertain. In the early 2000s, nanotechnology had just begun to attract attention. No one was certain about which subfields would generate the next breakthrough or its commercialization potential. This implies that evaluations of opportunities tend to vary a lot, which is an important antimajority enabling condition. Proposals passed under the majority rule in such an occasion are more likely to indicate that the proposals are not extreme enough (with lower potential upside), or the evaluators are subject to homophily (only seeking ideas/entrepreneurs that appear to be similar to themselves). Utilizing antimajority when an evaluation is highly uncertain helps overcome damaging biases, such as homophily and bypass competitions.

A related mechanism is that the antimajority rule at DFJ suggests that the minority has to be passionate about the proposal. This avoids the case of the bottom area in Figure 3 (the antimajority case), where a proposal is accepted not because anyone received a strong signal but someone accepted a mediocre proposal. This reinforces the contrarian approach by focusing only on the proposals (top left in Figure 3, antimajority case) that rivals (who employ a majority rule) are unlikely to identify.

Second, entering an area with a potentially overlooked upside does not necessarily mean that one can identify these opportunities. One needs to allow these potential upsides to be available and to be assessed in the first place. DFJ publicized themselves as a leading investor in this field. They did so through high-profile activities, such as extensive blogging, media appearances, and speaker engagements. Note that this approach is the opposite of the conventional, secretive approach of most other venture capitalists. DFJ adopted an “attraction strategy"; as Steve Jurvetson, one of the partners of DFJ, put it: "We want to become a powerful magnet so the needles find us" (Liu et al., 2017: 49). This enabled the antimajority rule because the potential upside was more likely to be in DFJ's evaluation pool.

Third, DFJ partners were very aware of the consequences of antimajority. Antimajority makes strategic sense, but it does not guarantee success. In reality, it guarantees many failures, as most proposals invested under antimajority will not succeed. The fact that DFJ was successful during the dot-com boom (including high-profile successes such as Hotmail, Baidu, Skype, and Twitter) is important. The partners felt more relaxed in adopting such a high-risk strategy when exploring opportunities in an emerging field (nanotechnology), and their prior successes facilitated trust among the partners to such an extent that they allowed the investments to follow the minority instead of the majority view. The partners' aligned interests in searching for the next big startup attenuated the chance of strategic voting.

Fourth and finally, antimajority's success depends on a sufficient majority to adopt the majority rule. This is when the blind spot at the ecological level is great enough to promise an attractive contrarian opportunity (Liu, 2020). The puzzle is that many VC firms do not seem to evaluate the fitness of majority rule critically enough. In this industry, omission errors are usually costlier than commission errors (Stross & Karp, 2000). The difference in costs is created, as the amount invested in a startup acts as a floor to the commission error costs; however, omission errors (such as missing a startup that could become a "homerun") have no ceiling (Taleb, 2007). From this inequality in costs, one would expect that VC firms might "flock to" a structure that minimizes omission errors, like polyarchy decision structures (Csaszar, 2012:628). However, the majority or even hierarchical decision structure is still the default structure in many VC firms (Stross & Karp, 2000). This mismatch may be explained by factors such as resource constraints (polyarchical structure is too costly even for the most resourceful), reputational costs (many VCs cannot afford the social costs of failures), and complacency (successful VCs do not want to take risks in nascent fields). This is good news for the contrarian, such as DFJ, who is ready to exploit the predictable omission errors generated by rivals who uncritically follow the majority rule.

To summarize, our formal model illustrates how antimajority is different from other screening functions. The case of DFJ illustrates when such a difference promises interesting asymmetry that enables antimajority to exploit the blind spots of alternative screening functions. In the next section, we will integrate insights from both the analytical model and the DFJ case to formalize and computationally examine the scope conditions of antimajority rule.

4. ANTIMAJORITY AND ITS SCOPE CONDITIONS

In previous section, we established that antimajority invests in a very different way than other decision rules. The DFJ case in Section 3 suggests that antimajority may work in practice but with strict operating conditions. We cannot model all these conditions highlighted in the case, such as trust among partners and the risk tolerance thanks to their prior successes, but we can model other enabling conditions, such as VC partners having different access to information that leads them to diverse mental representations of an option. We can also model the idea that agents vote in favor of an option only when they are passionate about the option. We use these two characteristics to show some of the essential boundary conditions of when antimajority voting could work as a valid entrepreneurial strategy.

In this section, we compare the investment behavior of a majority voting firm (G*) and an antimajority voting firm ($G^!$). For each firm, we vary the level of homogeneity of the mental representations of the agents in the triad and the threshold upon which an agent transitions from not voting in favor of an option to voting for it. We call the first parameter similarity. Similarity ranges from zero (i.e., three agents have completely orthogonal mental representations) to one (i.e., the mental representations are identical among the three agents). The second parameter is the "biases" of the agent (Christensen & Knudsen, 2010): an agent with low bias will invest in firms they are not very passionate about, while an agent with higher bias will invest in options for which they really care.

4.1 Diverse Screening Functions

An agent determines whether to invest in an option in a two-stage process. The first stage is to estimate the quality they see in the option, and the second stage is to use this quality to estimate the scalar acceptance probability (e.g., the probability that an agent would vote in favor of an option). Here, we employ a weighted sum to operationalize the agent's mental representations. The weighted sum determines how an agent translates the attributes of an option into a quality measure (Csaszar & Levinthal, 2016). In particular, options in this simulation have three attributes. This is the minimum number required to achieve full independence between the three agents, and any higher number would be reduced to three due to the number of agents.

We give each agent one attribute that they care about the most. Their attention toward the other two attributes depends on the level of similarity we specify. The weights for the three agents ($\overrightarrow{w_a}, \overrightarrow{w_b}, and\ \overrightarrow{w_c}$) are specified as the rows in the following matrix:

$$\mathbb{W} = \begin{bmatrix} \overrightarrow{w_A} \\ \overrightarrow{w_B} \\ \overrightarrow{w_C} \end{bmatrix} = \left(1 - \frac{2}{3}s\right) \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} + \frac{s}{3} \cdot \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} \tag{1}$$

where $s \in [0,1]$ represents the similarity between each of the agent's weights.

We use the weight matrix to estimate the quality that each agent estimates for the options (Q_i). This quality is given by the product of the weight matrix ($\mathbb{W}$) and the vector of attributes of the option ($\vec{q}$). This multiplication gives us the three values used to create the scalar. The weights (are used by each agent to calculate the quality they see for each option. The quality estimates of the three agents are estimated by the matrix product, of the weight matrix and the option's attributes:

$$\overrightarrow{Q(s)} = \begin{bmatrix} Q_A \\ Q_B \\ Q_C \end{bmatrix} = \mathbb{W}\,\vec{q} = \begin{bmatrix} \left(1 - \frac{2}{3}s\right) q_1 + \frac{sq_2}{3} + \frac{sq_3}{3} \\ \frac{sq_1}{3} + \left(1 - \frac{2}{3}s\right) q_2 + \frac{sq_3}{3} \\ \frac{sq_1}{3} + \frac{sq_2}{3} + \left(1 - \frac{2}{3}s\right) q_3 \end{bmatrix} \tag{2}$$

If $s = 0$, then the three agents have independent mental representations and, therefore, independent quality evaluations,

$$\vec{Q}(s = 0) = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \end{bmatrix} \tag{3}$$

Whereas, if $s = 1$, all agents have the same representations and quality evaluations, the average of the three attributes of the option. When s=1, our model is equivalent to Christensen and Knudsen (2010). Namely,

$$\vec{Q}(s = 1) = \left(\frac{q_1 + q_2 + q_3}{3}\right) \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \tag{4}$$

The use weighted sums to operationalize mental representations is similar to the one employed by Adner, Csaszar, and Zemsky (2014) and Csaszar and Levinthal (2016).

We focus on situations where there is diversity of mental representations, i.e., s < 1. In the presence of diversity, the weights are different for each agent, and thus the quality evaluations (Q_i) differ as well. The quality evaluation is then inputted into a screening function equivalent in form to the one used in Christensen and Knudsen (2010:82). This function outputs the scalar acceptance probability of an agent, i.e., the probability that the agent will vote in favor of the option. The function is defined as:

$$f(Q_i) = \frac{1}{2}\left[1 + \tanh\left(\frac{Q_i - b}{d}\right)\right] \tag{5}$$

where q is the quality that the agent had computed for the option, b is the bias of the agent, and d is the width of the zone of uncertainty of the function (Christensen & Knudsen, 2010: 81). The effects of the bias and width of the zone of uncertainty are shown in Figure 4.

In both panels, we show the screening function of agents with three different biases. The bias is an operationalization of how passionate an agent has to be to invest in an option. The specific value of the bias indicates at what quality level the agent will have 50% probability of investing in an option. For example, if bias $b = 0$ (blue plots), an agent will have a higher than 50% probability of investing in the option if it has a positive quality. However, if the agent has a bias $b = 2$ (orange plots), higher quality values will be needed for the agent to achieve the same probability. An agent with a low bias (green plots) will have high scalar acceptance probabilities even if the option has low quality. Next, we also manipulate the width of the zone of uncertainty of the screening function, as regulated by d. The screening functions on the left panel have a wider zone of uncertainty ($d = 0.25$) than on the right panel ($d = 0.1$). The zone of uncertainty determines the steepness of the screening function, a narrower zone of uncertainty leads to a steeper function and, thus, a smaller deviation in quality will have a larger effect in the scalar acceptance probability of the agent.

We argue that an agent who invests in options only when their quality is high (high bias) can be seen as investing only when they are passionate about the option. In contrast, a less passionate agent (low bias) will invest even when the idea does not have a high value. This is as if an agent with a higher bias will have a higher standard that needs to be met in order to vote in favor of an option.

In the simulations, we varied the bias and the similarity of the mental representations of the agents. All agents within a firm had the same bias and similarity, and these two values varied between firms. The only difference between the agents within one firm was the weights they used to estimate the quality of an option. All simulations used the same zone of uncertainty value, $d = 0.25$, as in the left panel of Figure 4.

Figure 4: An illustration of how bias (b) and zone of uncertainty (d) moderate screening functions of one agent

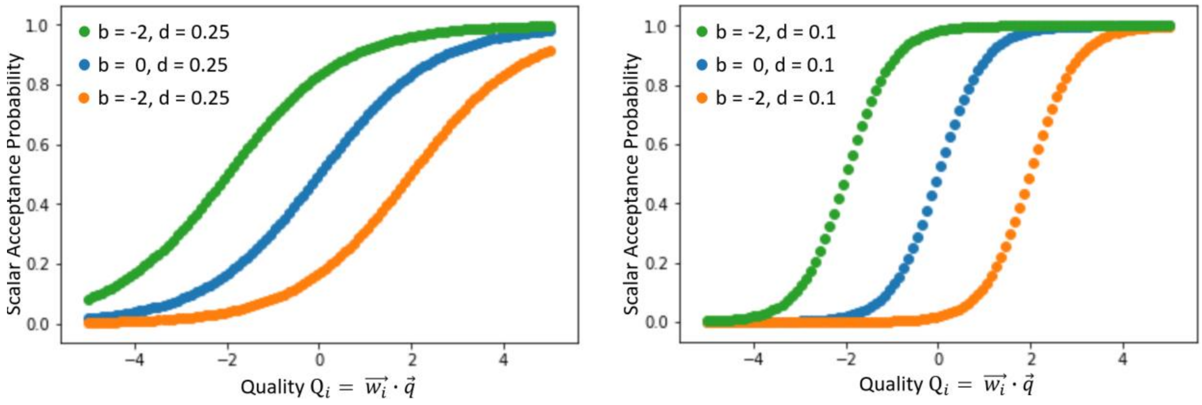

We now create a vector to store the scalar acceptance probabilities of the three agents ($\vec{V}$).

$$\vec{V} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} f(Q_A) \\ f(Q_B) \\ f(Q_C) \end{bmatrix} \tag{6}$$

The three scalar acceptance probabilities (x, y, and z) are then used to create the organizational acceptance probability, i.e., the probability that a firm will invest in an option. To calculate the organizational acceptance probability, we use the multinomial screening functions of Table 1.

4.2 Simulation

Armed with a way of simulating the voting behavior of firms that follow different decision rules as well as a way of creating options to screen, we explored how biases and similarity affect profitability. But before that, we need to specify a way of creating options and estimating their profitability.

We draw options from a multinomial log-normal distribution with a mean of zero and a standard deviation equal to one. Each draw gives the three attributes that define an option ($\vec{q} = [q_1, q_2, q_3]$). Log-normal distributions right skewed, options' attributes have a minimum value of zero but an unbounded right tail. The attributes of the distribution are not correlated, as any correlation would confound the effect of similarity between the mental representations of the agents[3].

We give each VC firm a cost for investing in an option, *k*. This cost was used to estimate the profits accrued from the option. The profit of an option is estimated as:

$$\pi(\vec{x}) = \sqrt{q_1^2 + {q_2}^2 + q_3^2} - k \tag{7}$$

that is the Euclidean distance from the origin minus the firm's cost of investment, *k*. An agent should then invest in options that they expect to generate profit for the firm. However, as agents might not necessarily know the exact costs the firm incurs in investing, we allowed the biases of each participant and the costs of the firm to vary independently of each other.

In Figure 5, we present VC the profit accrued by VC firms that differed in the similarity of the agents' mental representations (x-axis), and the agents' bias (y-axis). In this simulation, all firms had the same investment costs, a value that led to zero profits if firms invested at random (k = 3.56, the median value of the distribution).

Every pixel of the two panels in Figure 5 is a combination of similarity and biases, from low similarity and low biases on the bottom left corner to high biases and high similarity in the upper right. A pixel in row J and column K on the left panel in Figure 5 represents a firm that is exactly the same as the one of the right panel's row J and column K, except that the one on the left employs majority voting and the other uses the antimajority voting rule. Each figure has an independent color scale, where lighter colors show higher profitability and darker colors lower profitability. The two scales do not have the same ranges, so we can see that the profitability of majority voting firms is higher. However, both decision rules can make profitable investments. For each pixel in Figure 5, we have a firm screening 25,000 options.

[3] If the attributes of the option were collinear, it would be impossible for the agents to have quality measures independent from each other independent on how different their mental representations were (i.e. $\overrightarrow{w_a}$, $\overrightarrow{w_b}$, *and* $\overrightarrow{w_c}$). For example if $q_1 = q_2$ then $Q_1 = Q_2$ for any value of s. Decreasing the number of independent mental representations to just two, even if the agents are technically independent.

The profit is then calculating by multiplying the organizational acceptance probability by the profit, the firm would make if investing in the option and summing over all options. The width of the zone of uncertainty of the sigmoid function was 0.25 for all firms.

Figure 5: An illustration of the profit distribution of firms that follow the majority voting rule (left panel) and the antimajority rule (right panel) and how the profit is moderated by voting agents' similarity and biases.

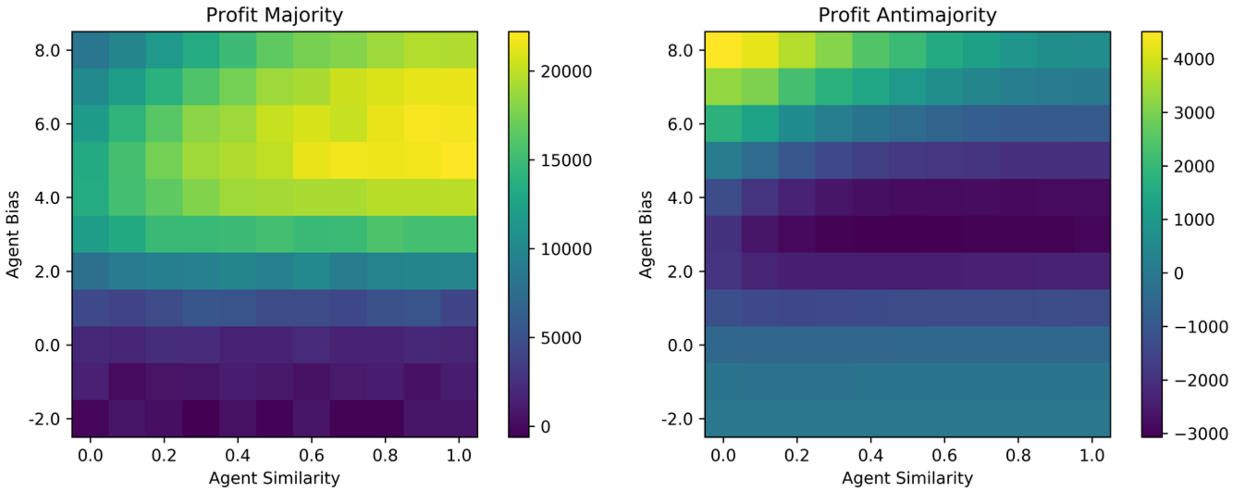

The main insight from Figure 5 is that the highest performing majority voting and antimajority voting firms differ in the level of similarity and the biases of their agents. We find that antimajority requires diverse agents (low similarity) to perform well, whereas the majority rule can perform well even if the agents are similar. Moreover, antimajority appears to necessitate agents with much higher biases than a majority voting firm. Although majority voting firms are always profitable, antimajority voting firms can actually have losses.

From this simulation, we replicate one of the key aspects of the DFJ decision process—only options that one founder was very passionate about were invested in. In our model, we see that the most profitable antimajority firms are the ones that have the highest biases. The biases determine the point at which an agent will transition from rejecting to investing in an option; the higher the bias, the more passionate an agent will need to be in order to invest.

We also verified the validity of the idea that antimajority works in situations where the signals taken from the environment are unclear. In the case of DFJ, they invested very early in nanotechnology at a time where it was unclear what would work and what would not. They used different points of view to select their options. In our model, we see that higher dissimilarity and more independent agents lead to higher profits for antimajority firms. Consistent with our formal analysis in Section 2, we found that the majority voting firms have much laxer conditions and achieve high profits under a broader set of firm configurations. This result reinforces that antimajority is very different from other rules. However, what if we allow them to compete and evolve? Can antimajority be a viable rule in the sense that it can invade and persist in an ecology dominated by other rules—primarily the majority rule?

5. ANTIMAJORITY IN AN ECOLOGY OF COMPETING FIRMS

There are many ways to model the relative advantage of antimajority. We chose to build an evolutionary model that allows alternative rules to compete and the successful ones to reproduce. The stabilized outcome will inform which rules can thrive in which conditions and the relative performances among competing rules. In this section, we present this evolutionary process in two ways. The first is descriptive, we present the characteristics of the surviving firms and how firms that use majority voting differ from firms that use antimajority voting rules. The second is normative. We show how the surviving firms differ in their average profitability and show how the descriptive characteristics, in all cases, appear due to an increase in profitability by the firms who employed them.

5.1 Evolutionary Model

As inspired by the case of DFJ, our model resembles the context of VCs: there are many startups seeking VC investment. VCs that choose the more profitable startups are more "successful" in the sense that they are more likely to reproduce in a competitive selection process. In particular, the bottom 10% of VCs will be selected out, and the newly joined VCs are more likely to follow the higher-performing VCs' decision rules. Additionally, in every period, 10% of the firms will have a random variation in their decision rule[4].

We allowed a "voting" process to unfold in each simulated VC. We assumed there were three partners in each VC, and as described, each obtained a signal from a startup to be evaluated. They voted "yes" if this signal was strong enough to pass their screening threshold; otherwise, they voted no. Whether a VC will invest in a startup depends on the VC's screening function, such as majority or antimajority. The profit earned by a VC is determined by the revenue of the startup the VC decides to invest in minus the investment cost incurred by the specific VC.

We simulated 100 startups to be evaluated by 100 VCs. The VCs were equivalent except for the decision rule they employed and that we drew the bias of the agents of the VC firm from a uniform distribution. Additionally, the mental representations of the agents in each VC firm were completely independent. In each period, each VC could invest in as many startups as they wanted. However, 10% of VCs that had the lowest average profitability were replaced by new VCs that preferentially replicated the rules of the firms with the highest average profitability. The process continued for 30 periods, and the results were averaged over 100 replications.

5.1.1 VC firms with varied bias

Figure 6 shows the first results from the evolutionary model. It presents the percentage of firms that used each decision rule and how it varied with the simulated periods. The left panel starts the simulation with all firms using majority voting; the middle panel displays the 50% using antimajority, and the 50% using majority at the start, and the right panel starts with only antimajority voting firms. In all cases, by period 30, around 20% of the firms still employed

[4] The random variation has no main effect in the results shown below other than increase the equilibrium level of antimajority voting firms. We added it for inclusions. What is required is the threshold under which firms are selected out, without this retention mechanism the results do not appear as the evolution would have no goal.

the antimajority voting rule. This explains that majority voting is a dominant strategy in this environment, but antimajority voting does have a niche that is not path dependent but that the niche is a part of this evolutionary environment.

We explore further what happens to the firms in these simulations. As shown in Figure 7, we find that the firms that end up adopting the antimajority voting rule are firms that have higher biases[5]. There were some majority voting firms with high biases, but there was no antimajority voting firm that survived that had low biases. This suggests that having high level of biases, or engaging options only when being passionate about them, is a necessary condition for antimajority.

Figure 6: How antimajority emerges as a contrarian strategy against the dominant majority rule.

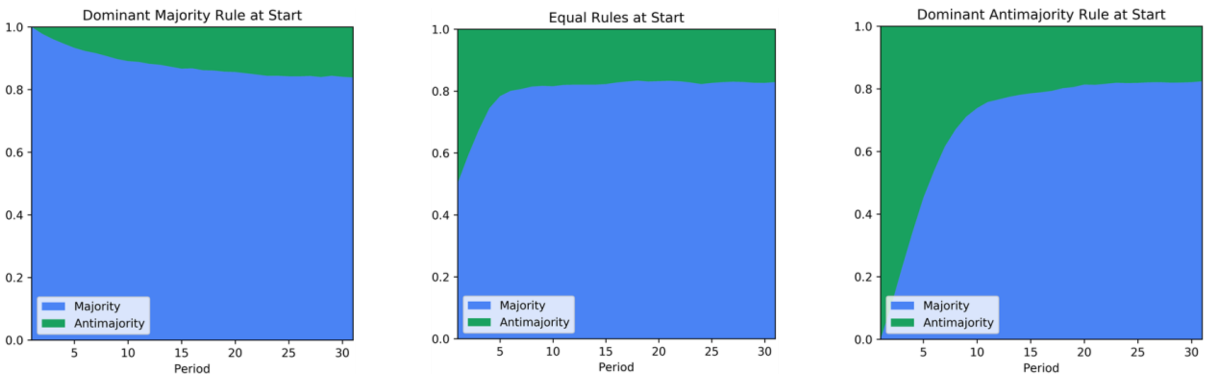

Figure 7: The bias characteristics of the surviving majority (left panel) and antimajority (right panel) firms in period 30. In this setup firms differed in their bias and the decision rule was allowed to evolve.

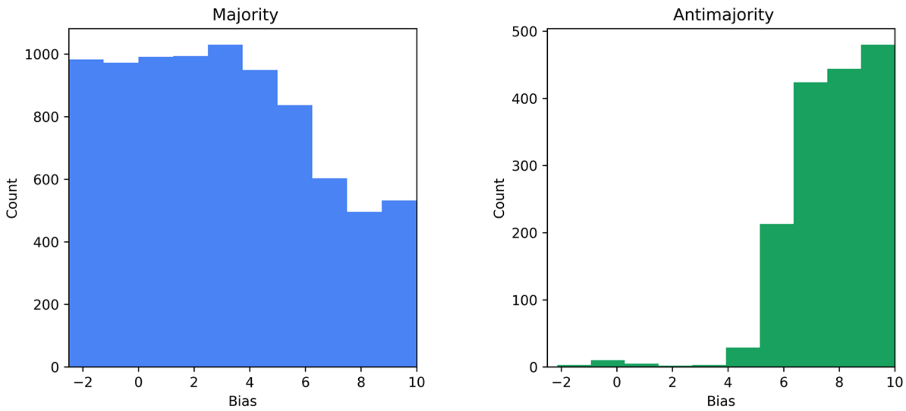

[5] The results in Figure 7 are based on the firms from the evolutionary process shown in the middle panel of Figure 6.

5.1.2 VC firms with varied bias and similarity

In Figures 8 and 9, we show a simulation where both the bias and similarity could vary in an evolutionary process. This simulation has 750 firms. We defined that half of the firms would start with majority voting rules and the other half with antimajority voting rules. The biases and similarities were drawn from uniform distributions. The results confirmed that antimajority can still emerge as a contrarian strategy against the dominant majority rule, but the stabilized proportion of antimajority drops from 20% to 15%. The firms that end up using antimajority voting have higher biases than the firms that use the majority voting rule (as shown in Figure 8). Moreover, Figure 9 shows that antimajority voting firms tend to have lower similarity in mental representations among the agents in the firm than firms with a majority voting rule. However, the selectivity of the similarity of the mental representations is not as strong as with the case of the biases.

Figure 8: The bias characteristics of the surviving majority (left panel) and antimajority (right panel) firms in period 30. In this setup firms differed in their bias and similarity and the decision rule was allowed to evolve.

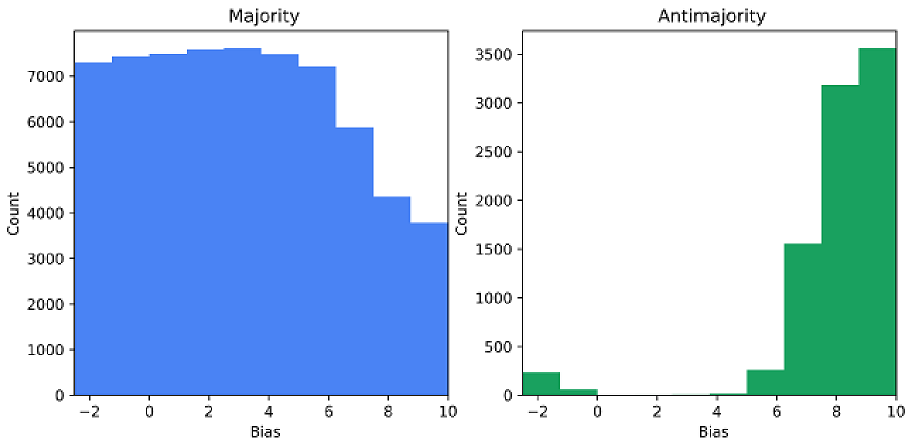

Figure 9: The similarity characteristics of the surviving majority (left panel) and antimajority (right panel) firms in period 30. In this setup firms differed in their bias and similarity and the decision rule was allowed to evolve.

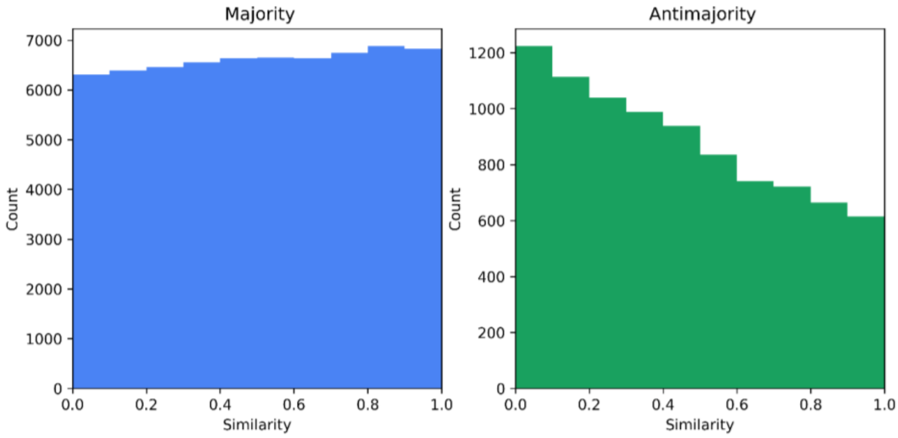

5.1.3 VC firms with varied bias, similarity, and cost

In Figures 10, 11, and 12, we show the results of including 100 VC firms that vary in their bias, similarity, and investment costs. We allowed these firms with varying characteristics to compete and examined how they evolved over time among the surviving firms. Consistent with the previous results, antimajority voting firms composed a minority of the firms in the population. Only about 28% of the VCs used antimajority voting at the end of the simulation. The VCs that used antimajority voting, though, tended to have much higher biases (Figure 10) and incurred lower investment costs (Figure 12) than the firms that used majority voting rules. The results regarding similarity (Figure 11) were less clear than in the previous simulation, but antimajority did have more VCs with low similarity than with high similarity.

With these results, we see validation of the idea that antimajority voting can create a stable niche in a market dominated by majority voting firms. We also found that, as in the case of DFJ, the antimajority voting firms possessed agents who were more passionate about the ideas, that is, having high biases, and who had agents who had the most dissimilar mental representations.

Figure 10: The bias characteristics of the surviving majority (left panel) and antimajority (right panel) firms in period 30. In this setup firms differed in their bias, similarity, and investment costs and the decision rule was allowed to evolve.

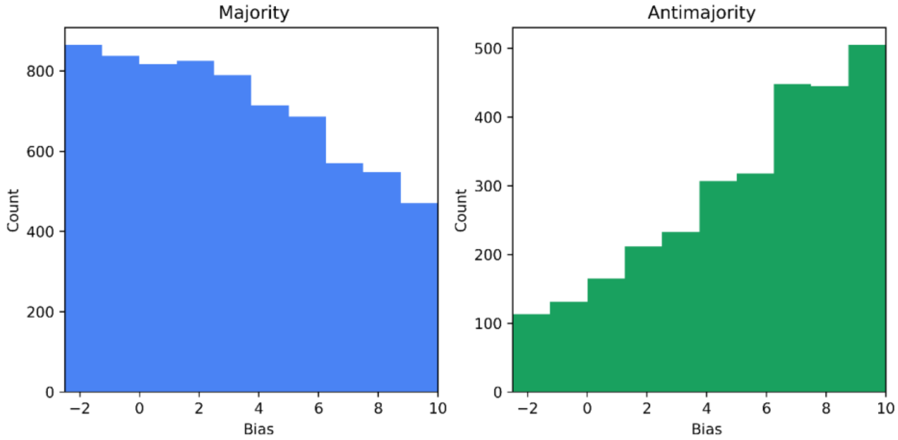

Figure 11: The similarity characteristics of the surviving majority (left panel) and antimajority (right panel) firms in period 30. In this setup firms differed in their bias, similarity, and investment costs and the decision rule was allowed to evolve.

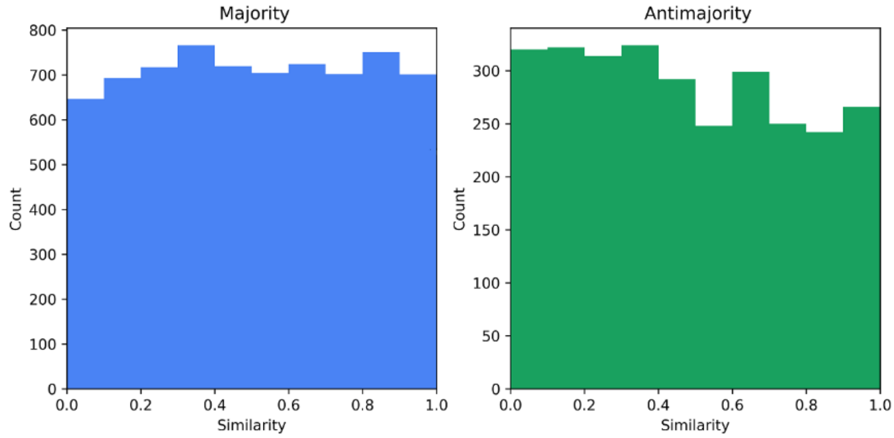

Figure 12: The investment cost characteristics of the surviving majority (left panel) and antimajority (right panel) firms in period 30. In this setup firms differed in their bias, similarity, and investment costs and the decision rule was allowed to evolve.

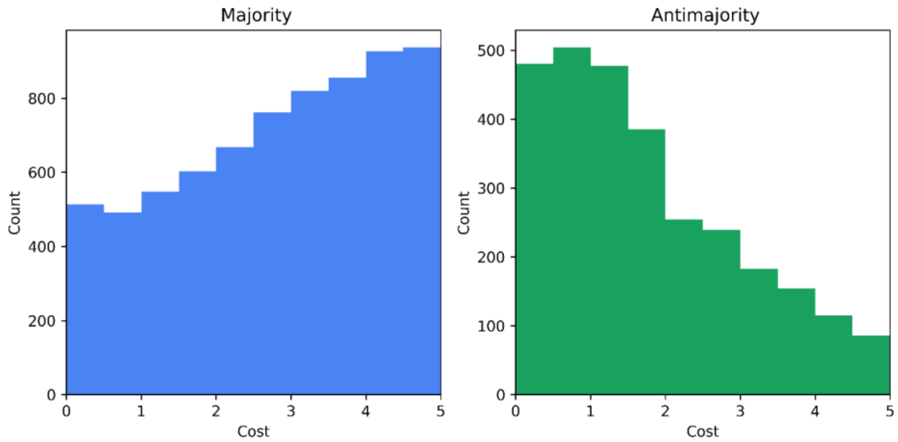

Interestingly, we also identified that the antimajority voting firms were predominantly firms that had low costs of investment. This likely explains why DFJ applied antimajority for nanotechnology in early 2000, where the uncertainty was so high in the field that the investment cost was lower. Early in a startup's development, the ideas are less well-rounded and clear; it is likely that only one agent might see the value of the idea, and only later on and after many pivots will more agents see this value. Therefore, antimajority might enable the earlier discovery of valuable opportunities. This would, in turn, provide lower-cost startups to the antimajority VC firms, increasing the profitability prospect of antimajority VCs. Our current model does not allow us to make a rigorous connection to the case of DFJ, but this point—antimajority voting allows VCs to find homerun startups earlier on in their lifetime—deserves future exploration.

5.2 Performance Analysis

The prior results show how competitive selections allow VCs with different rules and characteristics to survive and thrive. However, from these results, it is unclear which firms performed best and worst. To address this question, we did further analyses and created two dimensions to estimate the performance of the various types of VCs. The analysis is done in two steps. First, we have a large population of VCs, N=10,000, evolve together for 30 periods, on each period making the decision to invest in 1000 startups. The ten percent of firms with the lowest profitability are asked to choose a new decision rule (either majority or antimajority). These firms choose their new decision rule with a preferential weight towards the top-performing VCs in the previous period. This process is equivalent to the one shown in the previous steps except for the larger pool of VCs. After the evolution process is finished, the second step starts. We give the VCs 10,000 different firms to invest in and evaluate their investment probabilities and estimate the average profitability and average popularity of the investment decisions. We chose this approach of analysis mainly due to computational constraints. The results are consistent if we allow the above evolutionary processes to occur simultaneously but with fewer simulated VCs and startups. The current approach allows us to produce a large sample for regression analyses in order to demonstrate the marginal impacts of each characteristic under investigation.

The average profitability for an investment portfolio of VC j is given by

$$Average\ Profitability_j = \frac{\sum_{k=1}^{10000} Acceptance_j(k) \cdot \left[Revenue(k) - Cost_j\right]}{\sum_{k=1}^{10000} Acceptance_j(k)} \qquad (8)$$

where $Acceptance_i(j)$ is the acceptance probability of VC j for startup k. $Revenue(k)$ is the revenue accrued by startup k and $Cost_j$ is the investment cost that VC j incurs when investing in one startup. The average popularity of the investment portfolio of VC j is measured in two steps. In the first, we estimate the popularity of a startup as the proportion of VCs that invest in it. This is calculated as the sum of all VCs acceptance probability for each startup:

$$Popularity_k = \sum_{j=1}^{10000} \frac{Acceptance_j(k)}{10000} \qquad (9)$$

We then take the acceptance probability of each VC to create a measure of the VC's investment portfolio commonality:

$$Average\ Popularity_j = \frac{\sum_{k=1}^{10000} Acceptance_j(k) \cdot Popularity_k}{\sum_{k=1}^{10000} Acceptance_j(k)} \qquad (10)$$

The popularity measure estimates the percentage of other VCs who invested in the same startups as the VC in question. This serves as a proxy for the potential cost of investment in a startup, as more popular startups will achieve a higher price in the startup market.

Armed with the average profitability and average popularity measures, we estimate how the different VC characteristics (decision rule, bias, similarity, and investment cost) affect the performance of the VCs. Table 2 provides the descriptive statistics and the first-order correlations of these investments. Here, we can see how antimajority voting VCs were those that had a lower similarity among the mental representations of the agents, where the agents had higher biases, and the firms had lower investment costs. We also can how antimajority voting VCs have lower average profitability and lower popularity than majority voting firms. The lower average profitability is surprising, given the prevalence of antimajority voting in the evolutionary models.

Table 3 shows the ordinary least square (OLS) regressions on the two performance dimensions estimated: average profitability and average popularity. The models have three independent variables, i.e., similarity, bias, and cost of investment, as well as their interactions. The odd numbered models show the results of the OLS regressions using the antimajority voting firms that emerged after the evolutionary process, and the even numbered models show the results of majority voting VCs. Although the models are linear and many of the aspects used to create the VCs are nonlinear, the explained variance is very high in all models, with the lowest R-squared around 93%. Given that the results are simulated, and we can increase the sample size when required, we used stronger significance thresholds, with *** highlighting a one in a millionth chance, as well as 99% confidence intervals and not 5% as normally shown.

Table 2: The descriptive statistics and zero-order correlations of the simulated VCs

Variables	**1.**	**2.**	**3.**	**4.**	**5.**	**6.**
1. Antimajority	1					
2. Similarity	-0.036 (0.000)	1				
3. Bias	0.282 (0.000)	-0.006 (0.539)	1			
4. Cost of Investment	-0.315 (0.000)	0.000 (0.962)	-0.008 (0.448)	1		
5. Average Profitability	-0.176 (0.000)	0.010 (0.327)	0.634 (0.000)	-0.469 (0.000)	1	
6. Average Popularity	-0.409 (0.000)	-0.010 (0.322)	0.737 (0.000)	0.221 (0.000)	0.728 (0.000)	1
Mean	2920 of	0.495	3.768	2.498	2.317	5080.3
SD	10000	0.291	3.640	1.447	2.306	874.8

Note: *p*-value of the correlation shown in parentheses.

Table 3: A regression analysis of the performance of firms with varied characteristics

	Average Profitability		Average Popularity	
	(1)	(2)	(3)	(4)
	Antimajority	Majority	Antimajority	Majority
Bias	0.368^{***} (0.352, 0.383)	0.475^{***} (0.451, 0.498)	0.022^{***} (0.021, 0.022)	0.021^{***} (0.021, 0.022)
Similarity	-0.965^{***} (-1.126, -0.803)	-0.169 (-0.379, 0.042)	-0.049^{***} (-0.056, -0.042)	-0.004 (-0.009, 0.002)
Cost	-1.370^{***} (-1.440, -1.299)	-1.018^{***} (-1.056, -0.979)	-0.009^{***} (-0.012, -0.006)	-0.001 (-0.002, 0.0002)
Bias X Similarity	-0.021 (-0.048, 0.005)	0.160^{***} (0.119, 0.201)	0.004^{***} (0.003, 0.005)	-0.001 (-0.002, 0.001)
Bias X Cost	0.052^{***} (0.043, 0.062)	-0.001 (-0.009, 0.007)	0.001^{***} 0.001, 0.002)	-0.00004 (-0.0002, 0.0002)
Similarity X Cost	0.169^{**} (0.049, 0.290)	0.020 (-0.047, 0.087)	0.004 (-0.001, 0.009)	0.001 (-0.001, 0.003)
Bias X Sim. X Cost	-0.017^{*} (-0.033, -0.001)	-0.010 (-0.023, 0.003)	-0.0001 (-0.001, 0.001)	0.00001 (-0.0003, 0.0004)
Constant	2.017^{***} (1.921, 2.113)	3.796^{***} (3.676, 3.917)	0.347^{***} (0.343, 0.351)	0.469^{***} (0.466, 0.472)
Observations	2,920	7,080	2,920	7,080
R^2	0.947	0.933	0.976	0.951
Adjusted R^2	0.947	0.933	0.976	0.951
Residual Std. Error	0.230 (df = 2912)	0.258 (df = 7072)	0.155 (df = 2912)	0.222 (df = 7072)
F Statistic	$7.45 \times 10^{3***}$ (df = 7; 2912)	$1.42 \times 10^{4***}$ (df = 7; 7072)	$1.70 \times 10^{3***}$ (df = 7; 2912)	$1.95 \times 10^{4***}$ (df = 7; 7072)

Note: Odd numbered regressions show the Antimajority rule results, even numbered show the Majority rule results. 99% confidence intervals shown in parenthesis. * p<0.01, ** p<0.001, *** p $<10^{-6}$, Range of Similarity = [0, 1], Range of Cost = [0, 5], Range of Bias = [-2.5, 10]

The findings validate our expectation that antimajority voting firms have lower profitability than majority voting VCs; if the average profitability were higher, then antimajority would be the dominant logic in this environment, and thus the simulation would have been mis-specified.

The regression analysis allows us to understand why it is the case that antimajority voting firms emerge from the evolutionary process independent of the starting conditions. Specifically, for a decrease of one unit in the investment cost, antimajority voting firms increase their profitability by up to 1.37 units (compared to around one unit for majority voting firms.[6] This differential explains why antimajority voting firms tended to have low investment costs after the evolutionary process finished, and it links to the difference in average popularity of the investment portfolios of majority and antimajority voting firms (Models 3 and 4). Specifically, antimajority voting firms have investment portfolios that are much less popular (lower intercept in Models 3) than majority voting VCs (Model 4). This supports the idea that majority voting is indeed the dominant logic for this environment as highlighted by the prevalence and higher average profitability. However, antimajority can still emerge and stabilize as by focusing on less popular firms (hence lower investment costs).

Another way of observing these results is with the interaction plots shown in Figure 13. Here, we compare VCs that employ antimajority voting with firms that use majority voting for two values of bias, a high value shown in full lines, and a low value in dashed lines. Figure 13 plot the predictions of the regressions of Table 3 as a function of similarity for VCs and with zero cost of investment. We can observe how at low levels of similarity—where there are more antimajority voting firms—we obtain that at antimajority achieves equal or higher average profitability than the majority voting VCs. Additionally, the average popularity at every point for antimajority voting firms is much lower, and therefore we could expect that antimajority voting firms would be paying a lower price for their investment portfolio.

Figure 13 thus provides a useful connection to the success of DFJ: antimajority can emerge against the dominant majority rule but it is profitable when the firm can afford to have very diverse agents with high level of standard. Antimajority invests in ideas that are seen exceptionally strong by one of the agents but unimpressive by the others. In contrast, majority voting invests only if the startup is seen as good on average. This leads both rules to have

[6] This is a higher order effect due to the interactions between the different independent variables. The simpler version of Model 1 without interactions leads to a unit increase in average profitability for a unitary decrease in investment cost. The interactions link together to make antimajority voting benefit disproportionately from the lower investment costs. Note that the interactions are much more significant for antimajority voting firms than for the case of majority voting firms.

different investment choices. Given that the majority voting rule is dominant in the populations, the firms chosen by antimajority can be seen as lower cost than the ones that are invested in by majority voting. This is important for antimajority, as we found from the evolutionary model that the VCs that employed antimajority voting tended to have lower investment costs as well. From the performance analyses, we find that a lower investment cost can be warranted by investing in the startups disregarded by the majority, i.e., having a less popular investment portfolio. Finally, antimajority is profitable when a firm's agents have very different mental representations of the options and vote in favor of an option only when it sends a solid signal.

Figure 13: Interaction plots for average profitability and average popularity

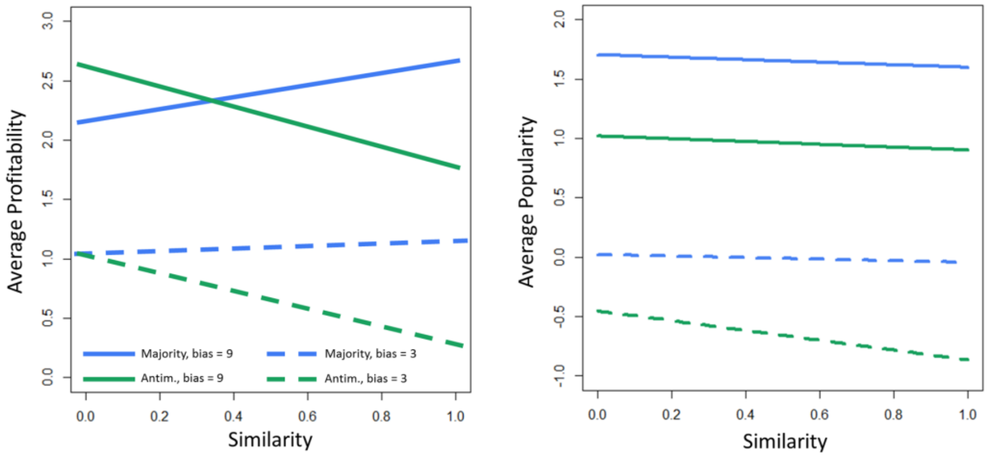

6. ROBUSTNESS

We reported the results where firms evolved their decision rules in an evolutionary process. In all cases, we allowed VCs that had agents with independent mental representations. We did this because if the agents have similar representations, the niche for antimajority voting disappears, and we are left with zero antimajority voting firms after the simulation ends. Similarly, if we simulate firms that have a uniform distribution of similarities but that are also equal in their biases and costs, the antimajority voting firms do not survive. It appears that dissimilarity in mental representation is not enough to create a niche for antimajority voting. However, if in the presence of variance in the VC's bias, then the niche can become populated by 20% of antimajority voting firms, and if investment cost is also allowed to vary then the niche can allow for up to 30% of the market.

Another boundary condition is the level of retention in the market. In our simulations, every period, we selected the bottom 10% of VCs to draw a new decision rule. However, as antimajority voting does not lead firms to achieve the highest performance, as this retention threshold increases, the percentage of antimajority voting VCs that appear in equilibrium decreases. We find that antimajority voting firms have an equilibrium population higher than 5% up to a 30% threshold and 2.5% up to the 50% threshold. If the market removes more than half its members every period, the remaining members will only have majority voting rules.

We replicated the results using consensus voting as the dominant logic and what we call anticonsensus—invest if one or two of the agents want to invest but not if all want to invest—and found equivalent results. We find that a contrarian niche emerges when the contrarian firms invest in the blind-spots of the dominant logic, either it being the majority or consensus rule.

We also simulated markets composed of three voting rules: majority voting, antimajority voting, and consensus voting. Here, both antimajority and consensus voting carve a niche from the majority voting rule. In the long term of the evolution, there were around 15% antimajority voting firms, 25% consensus voting firms, and 60% majority voting firms. We identified that consensus voting firms tended to have agents with “middle-levels” of bias. We found that the inclusion of consensus voting led the majority voting VCs with mostly low biases to survive, whereas VCs with high biases evolved to use antimajority voting rules, and the firms in the middle employed consensus rules.

We explored the effects of the variance of the distribution and found that antimajority voting tends to benefit from more skewed distributions. This is understandable as the more skewed a distribution, the higher the probability that if one value is high, the others will be lower. In that case, antimajority voting VCs should have an advantage in recognizing the value of these startups.

7. DISCUSSION AND IMPLICATIONS

This paper introduces an unconventional screening rule—antimajority—and outlines when it can help firms to identify a different set of opportunities that rivals who adopt conventional rules, such as the majority voting rule, tend to overlook. Rivals may overlook these opportunities not because they suffer from certain behavioral failures, as prior studies suggest. Instead, adopting conventional rules is sensible because they tend to be efficient and associated with decent performance. Nevertheless, a sensible rule also creates a contrarian niche when it becomes a dominant logic that the majority in an ecology adopts. This, in turn, creates profitable opportunities for doing things in a way contrary to the dominant logic. In particular, a majority voting rule is so efficient that it is usually the mainstream default when firms screen options or candidates. We show that this sensible default allows firms that adopt antimajority rule a chance to thrive because it helps to identify some opportunities left behind by firms that follow the majority rule.

However, firms that adopt antimajority rule are likely to fail even when it is viable in theory. Our VC case and the evolutionary model highlight the demanding scope conditions of antimajority; it requires firms to fulfill several necessary conditions before antimajority can guide them to viable contrarian opportunities. For example, antimajority requires that the evaluators in the firm have different mental representations on what constitutes a valuable option. If the evaluators do not have diverse mental representations, as the prior studies assume, antimajority will lead organizations to find some of the worst options. Antimajority also requires evaluators within the adopting firm to share the goal of finding valuable but overlooked opportunities. To do this, these evaluators have to agree in selecting only options that they are highly passionate and trust each other. If the evaluators do not trust each other or have misaligned interests, strategic voting will likely occur, and antimajority will again lead to worse options being accepted. Finally, antimajority should be used only if the firm can find options with lower costs. This should not be problematic as antimajority VCs tend to have investment portfolios that comprise less popular startups than majority voting firms. Antimajority can find these firms as the firms that they invest in are often overlooked by the majority, and thus should have a higher potential upside, i.e., lower cost.

These difficulties in implementing antimajority in reality partly explain why antimajority is rare in practice as well as overlooked in the organizational design literature. For antimajority rule to operate as a guide to unconventional opportunities, the firm that adopts it needs to be unconventional as well. As the case of DFJ illustrates, the firm needs to consciously pick a context that most sensible investors avoid; the partners need to be diverse but trust each other's judgment when they agree on the basis of disagreement; they also need to evaluate success not based on hit average but on the number homeruns. The implication is that antimajority is not a decision rule for the majority even though it requires a majority: the majority and their homogeneous approach of evaluating alternatives precisely creates a contrarian niche for the antimajority rule to thrive as a viable minority—fortune favor firms like DFJ who could afford to be contrary.

More generally, our findings suggest an alternative source of strategic opportunity and contribute to behavioral strategy. Recent studies have argued that attractive opportunities are likely protected by various behavioral failures that limit the majority of firms to sense, seize, integrate, or justify valuable resources (Denrell et al., 2019; Gavetti, 2012; Liu, 2020). Prior studies have explored several types of failures that promise contrarian opportunities, including gender biases in decision-making (Siegel, Pyun, & Cheon, 2018), convergence in mental models (Gavetti & Menon, 2016), constraints from institutional logic (Jonsson & Regnér, 2009), and resistance and inertia to changes (Fang & Liu, 2018). Here, we provide an alternative source of profit without assuming rivals' behavioral failures. It is well known that in competition, there is a negative frequency dependency in most strategies, i.e., a strategy is less profitable when many adopt the same strategy. Our theory suggests that a dominant logic, such as a popular strategy or best practice, not only reduces the expected profit of those who adopt it (i.e., red ocean) but also endogenously creates a contrarian niche (i.e., blue ocean). However, such a blue ocean strategy is bound to be a strategy for the minority because it hitchhikes on the success of the red ocean strategy, and a blue ocean strategy, per our theory, requires a strategist who understands its demanding scope conditions to implement it successfully.

Our findings also elucidate an important nuance to the organizational design literature. Prior studies have highlighted that the majority rule is an efficient design for simultaneously addressing both commission and omission errors (Csaszar, 2013). We demonstrated its inefficiency when we relaxed the assumption of homogeneous evaluators and considered the ecology of competition. The majority rule can produce systematic blind spots—omission errors at the ecological level. The wisdom of the antimajority emerges when “information is sufficiently heterogeneous and the well informed are not overly abundant” (Callander & Hörner, 2009: 1421). This increases the chance that the incumbents will be disrupted despite their local efficiency. The implication is that an optimal design can become suboptimal when too many organizations follow the same design.

The present results are relevant to the tension of balancing exploration and exploitation (March, 1991). An often-overlooked model in March (1991) is the competition for primacy model, where exploration is formalized as an increased variance in performance. Our model bridges organizational design and the literature on exploration and exploitation in the sense that the antimajority rule illustrates an alternative source of variance that enables organizations to explore the less traveled paths and potentially to win big. Our VC case illustration also suggests the caveats and risk when a contrarian searches for and exploits the blind spots of the majority. Antimajority is a design that exploits the under-explorations of rivals and a desirable strategy for firms that can afford to trade frequent losses with occasional primacy.

The scope conditions of antimajority also provide important caveats for entrepreneurs who aspire to become contrarian. Our results showed that a viable antimajority rule is very difficult to achieve. For example, one needs to select the relevant context (e.g., a highly uncertain field that deters resourceful incumbents) and work with trusting partners with differential access to information so that antimajority can guide one to identify viable contrarian opportunities. Our theory and findings reinforce that being different is necessary for entrepreneurial success, but being correct simultaneously is perhaps the more difficult challenge.

Finally, our paper may provide a methodological contribution. We applied mixed methods to triangulate our key theoretical argument. We first applied formal analysis to differentiate antimajority from other screening rules. The viability of antimajority rule was then illustrated by the case of DFJ, which also highlights, qualitatively, several important necessary conditions for antimajority to work in practice. We then computationally examined these scope conditions using an evolutionary model with competing rules. Each method we applied has its limitations; for example, formal analysis is rigorous, but its results may be irrelevant to reality. A case study may be inspiring, but it can be hard to generalize the finding. Our mixed methods approach allowed the three methods applied to complement each other, and together, they show when antimajority can be a viable rule that helps some firms to identify opportunities that are too unconventional for the majority to see.

REFERENCES

Adner, R., Csaszar, F. A., & Zemsky, P. B. (2014). Positioning on a Multiattribute Landscape. *Management Science*, 60(11): 2794–2815.

Argote, L. (2013). *Organization learning: A Theoretical Framework*. Boston, MA: Springer.

Arrow, K. J. (1951). *Social Choice and Individual Values, vol. 12*. New Haven, CT: Yale University Press.

Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1): 99–120.

Barney, J. B. (1986). Strategic factor markets: Expectations, luck, and business strategy. *Management Science*, 32(10): 1231–1241.

Black, D. (1948). On the rationale of group decision-making. *Journal of Political Economy*, 56(1): 23–34.

Bohman, C. (2009). *Attraction: A new driver of learning and innovation*. Stockholm School of Economics.

Callander, S., & Hörner, J. (2009). The wisdom of the minority. *Journal of Economic Theory*, 144(4): 1421–1439.

Christensen, M., & Knudsen, T. (2010). Design of decision-making organizations. *Management Science*, 56(1): 71–89.

Csaszar, F. A. (2012). Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management Journal*, 33(6): 611–632.

Csaszar, F. A. (2013). An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organization Science*, 24(4): 1083–1101.

Csaszar, F. A., & Eggers, J. (2013). Organizational decision making: An information aggregation view. *Management Science*, 59(10): 2257–2277.

Csaszar, F. A., & Laureiro-Martínez, D. (2018). Individual and Organizational Antecedents of Strategic Foresight. *Strategy Science*, 3(3): 481–553.

Csaszar, F. A., & Levinthal, D. A. (2016). Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37(10): 2031–2049.

Dasgupta, P., & Maskin, E. (2008). On the robustness of majority rule. *Journal of the European Economic Association*, 6(5): 949–973.

Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. (2009). Optimal structure, market dynamism, and the strategy of simple rules. *Administrative Science Quarterly*, 54(3): 413–452.

Denrell, J. (2004). Random walks and sustained competitive advantage. *Management Science*, 50(7): 922–934.

Denrell, J., Fang, C., & Liu, C. (2019). In search of behavioral opportunities from misattributions of luck. *Academy of Management Review*, 44(4): 896–915.

Denrell, J., Fang, C., & Winter, S. G. (2003). The economics of strategic opportunity. *Strategic Management Journal*, 24(10): 977–990.

Dierickx, I., & Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science*, 35(12): 1504–1511.

Ethiraj, S. K., & Levinthal, D. (2004). Bounded rationality and the search for organizational architecture: An evolutionary perspective on the design of organizations and their evolvability. *Administrative Science Quarterly*, 49(3): 404–437.

Fang, C., & Liu, C. (2018). Behavioral strategy: In search of an alternative source of profitability. *Advances in Strategic Management*, 38(1): 209–219.

Felin, T., & Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. *Strategy Science*, 2(4): 258–271.

Ganz, S. C. (2020). Conflict, Chaos, and the Art of Institutional Design, *Working paper*.

Gavetti, G. (2012). Toward a behavioral theory of strategy. *Organization Science*, 23(1): 267–285.

Gavetti, G., Levinthal, D. A., & Rivkin, J. W. (2005). Strategy Making in Novel and Complex Worlds: The Power of Analogy. *Strategic Management Journal*, 26(8): 691–712.

Gavetti, G., & Menon, A. (2016). Evolution cum agency: Toward a model of strategic foresight. *Strategy Science,* 1(3): 207–233.

Gavetti, G., & Porac, J. (2018). On the Origin of Great Strategies. *Strategy Science,* 3(1): 352–365.

Gigerenzer, G., & Todd, P. M. (1999). *Simple Heuristics that Make Us Smart.* Oxford, UK: Oxford University Press.

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences,* 101(46): 16385–16389.

Jonsson, S., & Regnér, P. (2009). Normative barriers to imitation: Social complexity of core competences in a mutual fund industry. *Strategic Management Journal,* 30(5): 517–536.

Leiblein, M. J., Chen, J. S., & Posen, H. E. (2017). Resource allocation in strategic factor markets: A realistic real options approach to generating competitive advantage. *Journal of Management,* 43(8): 2588–2608.

Levine, S. S., Bernard, M., & Nagel, R. (2017). Strategic intelligence: The cognitive capability to anticipate competitor behavior. *Strategic Management Journal*, *38*(12): 2390–2423.

Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal,* 14(8): 95–112.

Lippman, S. A., & Rumelt, R. P. (2003). A bargaining perspective on resource advantage. *Strategic Management Journal,* 24(11): 1069–1086.

Liu, C. (2019). *Luck: A Key Idea for Business and Society.* London, UK: Routledge.

Liu, C. (2020). Why do firms fail to engage diversity? A behavioral strategy perspective. Forthcoming at *Organization Science.*

Liu, C., Vlaev, I., Fang, C., Denrell, J., & Chater, N. (2017). Strategizing with biases: Engineering choice contexts for better decisions using the Mindspace approach. *California Management Review,* 59(3): 135–161.

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22): 9020–9025.

Ludwin, W. G. (1978). Strategic voting and the Borda method. Public Choice, 33(1): 85–90.

March, J. G. (1962). The Business Firm as a Political Coalition. *The Journal of Politics*, 24(4): 662–678.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science,* 2(1): 71–87.

Moore, E. F., & Shannon, C. E. (1956). Reliable circuits using less reliable relays. *Journal of the Franklin Institute*, 262(3): 191–208.

Peteraf, M. A. (1993). The cornerstones of competitive advantage: A resource-based view. *Strategic Management Journal*, 14(3): 179–191.

Pontikes, E. G., & Barnett, W. P. (2017). The non-consensus entrepreneur organizational responses to vital events. *Administrative Science Quarterly,* 62(1): 140–178.

Porac, J. F., Thomas, H., Wilson, F., Paton, D., & Kanfer, A. (1995). Rivalry and the industry model of Scottish knitwear producers. *Administrative Science Quarterly,* 40(2): 203–227.

Porter, M. E. (1996). What Is Strategy? *Harvard Business Review*, 74(6): 61–78.

Romme, A. G. L. (2004). Unanimity Rule and Organizational Decision Making: A Simulation Model. *Organization Science,* 15(6): 704–718.

Sah, R. K., & Stiglitz, J. E. (1986). The architecture of economic systems: Hierarchies and polyarchies. *The American Economic Review*, 76(4): 716–727.

Sah, R. K., & Stiglitz, J. E. (1988). Committees, hierarchies and polyarchies. *The Economic Journal,* 98(391): 451–470.

Siegel, J., Pyun, L., & Cheon, B. (2018). Multinational firms, labor market discrimination, and the capture of outsider's advantage by exploiting the social divide. *Administrative Science Quarterly*, 64(2): 1–28.

Stross, R. E., & Karp, J. (2000). eBoys: *The First Inside Account of Venture Capitalists at Work*. New York, NY: Crown Publishing Group.

Taleb, N. (2007). *The Black Swan: The Impact of the Highly Improbable*. New York, NY: Random House.

Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7): 509–533.

Todd, P. M., & Gigerenzer, G. (2012). *Ecological Rationality: Intelligence in the World.* Oxford, UK: Oxford University Press.

PAPER 4

"Creativity is the ability to introduce order into the randomness of nature."
– Erich Hoffer

On the Strategic Use of Product Modularity for Market Entry: Theory and Empirical Evidence

Jose P. Arrieta[1], Roberto Fontana[2], and Stefano Brusoni[1]
[1] ETH Zürich, Switzerland, [2] University of Pavia, Italy

ABSTRACT

In this study, we develop a formal model to account for the conditions under which firms should enter the market with multiple products, each of which is based on a different technological standard that competes for market dominance. Firms should enter the market with multiple products during the "era of ferment" of a technology cycle, but only if the early-mover advantages gained, and switching costs saved, offset the increase in development costs. We further outline the conditions under which product modularity can reduce development costs for a firm that wishes to introduce multiple products to the market. We find that product modularity broadens the conditions under which firms benefit from introducing multiple products to the market. Our model addresses how firms can use product modularity strategically to lower their risks by diversifying their product portfolio. We validate the model predictions with an empirical study and find further evidence that firms that follow the model's expectations outperform their competitors.

Keywords: Product modularity; Real Options; Standardization; Technology Cycles

1. INTRODUCTION

This paper explains how firms can use product modularity to enter markets in the presence of competing standards (Ulrich, 1994; Wiegmann, de Vries, & Blind, 2017). To do so, it builds upon and extends prior work on market-entry decisions under conditions of high uncertainty regarding the dominant standard (Folta & O'Brien, 2004). Previous work on firm entry focused on the uncertainty, irreversibility, and timing of market-entry decisions to explain post-entry performance (Folta, Johnson, & O'Brien, 2006). This paper extends this line of reasoning by exploring how a firm's product strategy (i.e., modular or integral product architecture) affects firms' entry and performance.

The presence of competing standards may represent an opportunity for firms, but it also complicates the market-entry decision (Suarez, 2004). In the presence of multiple standards, a firm can err by waiting too long to enter the market. However, it can also err by entering too early with the "wrong" standard (i.e., one that does not go on to become dominant) (Garud, Nayyar, & Shapira, 1997). We build a real-options model to outline how product modularity allows firms to reduce errors and hedge their bets by lowering the costs of introducing products based on multiple standards to the market.

Entry decisions are influenced by the type of product architecture firms decide to use. Modularity is a well-known strategy for organizing the development and improvement of new products—a strategy that firms could plausibly employ to support their entry decisions (Baldwin & Clark, 2000; Garud, Kumaraswamy, & Langlois, 2009). Scholars have explored in depth how modularity can enable firms to solve problems effectively (Schilling, 2000; Brusoni & Prencipe, 2001; Fang, Lee, & Schilling, 2010). They have also outlined the inherent limits to the use of modularity as an organizational strategy at times of technological uncertainty (Schilling & Steensma, 2001; Brusoni, Prencipe & Pavitt, 2001; Brusoni and Prencipe, 2006; Fang & Kim, 2018). Less is known about how modularity can help firms manage uncertainty in the composition of their product portfolio. We know that firms broaden or narrow their product portfolios and resource allocation to match the level of uncertainty in the environment; however, the use of product modularity has not yet been explored in conjunction with portfolio management (Klingebiel & Rammer, 2014; Klingebiel & Joseph, 2016). This paper builds upon and extends this line of work through formal modeling and empirical analysis of the Local Area Network (LAN) industry. This industry was of strategic importance in the emergence of the Internet as a foundational infrastructure for modern business. Modularity played a pivotal role in enabling firms to cope with uncertainty about which LAN standard would emerge as dominant.

A firm that decides to enter a market by employing product modularity faces a dilemma. Modularity enables market entry with products based on multiple standards (Ulrich, 1994); firms that choose to do so lower their chance of incurring switching costs and increase their chances of benefiting from early-mover advantages (Folta, O'Brien, & Johnson, 2006; Farell & Shapiro, 1988). Yet, they also increase their development costs, because they must develop interfaces to accommodate the various standards that their products are based uponuse. In contrast, interfaces are not necessary for an integral architecture that is designed to support only a single standard.

In this paper, we develop a real-options model to identify the boundary conditions under which product modularity can help a firm's market entry. We find that the tension between early-mover advantage, switching costs, and the development costs of modular interfaces explains whether and when product modularity is beneficial for a firm during market entry. Early-mover advantages and switching costs vary during the technology cycle. Early in the cycle, early-mover advantages can be high (Folta, Johnson, & O'Brien, 2006). Late in the cycle, firms incur switching costs if they have previously invested in the "wrong" standard (Garud, Nayyar, & Shapira, 1997; Suarez, 2004).

We test the model's propositions with data from a sample of firms and products from the LAN industry. The LAN industry experienced a full technology cycle (Anderson & Tushman, 1990) during the 1990s, when it moved from an "era of ferment," in which multiple standards battled for dominance, to an "era of incremental change," in which a single dominant standard was in place. The dataset allows us to identify the product-entry strategy and the time of market entry of each firm active in the industry during the period of interest. Differences in the strategy and timing of market entry allow us to test our propositions for firm entry and performance.

This paper is organized as follows. Section 2 introduces the prior literature that has studied market entry in the presence of competing technologies, and builds intuitions about how product modularity can be beneficial during market entry. Section 3 develops a real-options model to serve as a more formal and rigorous framework for our empirical exercise. Section 4 introduces the context and dataset for the empirical exercise, and motivates why the model's predictions can be validated in the real world. Section 5 tests the model's validity and extends these predictions to show how modularity can affect firm performance after market entry. Section 6 concludes the paper.

2. RELATED LITERATURE

In this paper, we study the process of market entry in the presence of competing standards (Folta & O'Brien, 2004). Within this context, we highlight the role of product modularity in helping firms achieve higher performance. In particular, we propose that firms can use product modularity as a tool to diversify and reduce the risks associated with entry into markets where a dominant standard has yet to emerge. This section introduces the relevant streams of literature and explains how we build on them in our study.

2.1. Market entry

Market entry is a key strategic move (Kalish & Lilien, 1986; Lilien & Yoon, 1990; Baron, 1995; Sundali, Rapoport, & Seale, 1995). Research in real-options theory explains how firms' decisions to enter a new market are strongly influenced by uncertainty (O'Brien, Folta, & Johnson, 2003; Folta, Johnson, & O'Brien, 2006). However, "the effect of uncertainty in entry is not monotonic." Market-entry decisions entail three crucial factors: the irreversibility of the decision, the potential for growth, and potential early-mover advantage (Folta & O'Brien, 2004: 121).

Irreversibility is essential, as substantial investments are necessary for entering a new market (Folta & O'Brien, 2004). Beyond that, the potential for growth and early-mover advantage further complicates the decision to invest. If entering a market does indeed enhance growth, a given firm might benefit from market entry even if the net present value of the investment is negative. The same holds when the entry decision will provide substantial early-mover advantages (Folta, Johnson, & O'Brien, 2006). The studies by Folta and colleagues allow us to estimate not only whether market entry is viable for a firm, but also whether an early entry is their best option.

A related stream of research studies adds an extra variable to Folta and colleagues' studies—namely, the effect of delayed product differentiation before market entry on a firm's performance (Lee & Tang, 1997; Swaminathan & Tayur, 1998). According to these studies, firms can create products that potentially serve multiple markets during periods of high uncertainty. These products are known as "vanilla boxes" (Swaminathan & Tayur, 1998). Developing "vanilla boxes" allows firms to bide their time, in the sense that they can decide to invest in a specific product type without the need to define the exact characteristics of the product. The firm can defer differentiation decisions until after the initial design has been developed. In so doing, the firm waits for uncertainty to decrease before making their market-entry decision.

In the context of product differentiation, such as the feature-phone industry, Klingebiel and Joseph (2016) present a different view of market entry. They argue that firms follow one of two different market-entry strategies: to delay differentiation so they can enter with "appropriate" products, or to enter early and hedge their bets by introducing multiple products. Klingebiel and Joseph (2016) find that firms make timing and differentiation decisions jointly, i.e., in parallel to deciding when to enter the market, the firm also needs to decide the breadth of products with which to enter. If one firm were to enter the market early, a broad entry strategy would be best, as it would increase the chances of entering with the "right" product choice. Conversely, if another firm decides to enter late, it would benefit from choosing only those products that are more likely to succeed, and thus, should enter with a narrower product line.

Klingebiel and Joseph's (2016) findings contrast with those of prior studies on real-options theory. However, they do relate to a conceptual study by Garud, Nayyar, and Shapira (1997), who provide a conceptual foundation for why firms may benefit from taking timing and differentiation decisions jointly. Garud and colleagues discuss entry timing in the context of technology cycles (Anderson & Tushman, 1990). They explain that early in the technology cycle, omission errors (i.e., failing to introduce the future dominant technology to the market) are costlier than commission errors (i.e., investing in a technology that fails to dominate the market in the future). In contrast, late in the technology cycle, introducing the "wrong" technologies to the market should be more costly than failing to introduce any technology at all.

In this study, we build on Klingebiel and Joseph's (2016) findings and the conceptual foundation from Garud et al. (1997) to develop a formal model that identifies the conditions under which a firm should enter with one product or with multiple products. We then show how product architecture—specifically, the use of product modularity—can significantly reduce costs when introducing multiple products to the market. We introduce the relevant literature on product modularity in the next section.

2.2. Product modularity

Modularity is a foundational concept of organization science (Simon, 1962; Henderson & Clark, 1990; Sanchez & Mahoney, 1996; Baldwin & Clark, 2000). "Modularity is a general set of design principles for managing the complexity of […] interdependent systems" (Ethiraj & Levinthal, 2004:161). Modularity involves "breaking up a complex system into discrete pieces—which can then communicate with one another only through standardized interfaces within a standardized architecture" (Langlois, 2002: 19).

In this study, we explore the role of one specific type of modularity: product modularity. Ulrich (1994) introduced the concept of product modularity to define how a product can be separated into distinct modules that communicate via well-defined interfaces. In this study, we treat product modularity as an attribute of the product—not as an attribute of the organization, as other studies have done (Sanchez & Mahoney, 1996; Schilling, 2000; Baldwin & Clark, 2000). We study the use of product modularity for market-entry decisions as a tool to broaden the conditions under which introducing multiple products is beneficial to the firm.

Product modularity allows firms to broaden their real-options portfolio for market entry. By employing product modularity, the firm might be able to afford to introduce multiple modular variations of the same product, targeting different market niches. Also, product modularity enables firms to delay product differentiation (Lee & Tang, 1997), hence gaining the time to acquire more information about superior solutions.

Lee and Tang's (1997) model captures the competitive dynamics of firms operating in established product markets e.g.,with clear architectures in place, where firms need to optimize their operational costs (e.g., inventory levels). Christensen, Suarez, and Utterback (1998: 207) look at sectors characterized by very rapid change, and argue instead that "firms that target new market segments with an architectural innovation will tend to be more successful than those that target existing markets." We leverage this idea and use it to identify and discuss the conditions under which the use of product modularity can be beneficial for a new market entrant. The key benefit of product modularity is the potential to lower the development costs of introducing multiple products to the market (Ulrich, 1994).

Developing a modular product is more costly than developing a single integral one, because it requires the design of well-defined interfaces in addition to the rest of the product. However, once such interfaces have been designed, the cost of creating a second or third module, and thus a second or third product, are lower than the costs of creating new integral products from the ground up. As we will outline, by lowering the costs of introducing multiple products to the market, product modularity allows firms to diversify their market-entry risk and increase their chances of gaining early-mover advantages whole avoiding switching costs.

2.3. Standardization

One context in which alternative technologies compete is that of a market that is still undergoing standardization. “Standardization aims to resolve situations where involved actors prefer a common solution to a problem, but have not yet agreed which option to choose” (Wiegmann, de Vries, & Blind, 2017: 1371). In a market that is undergoing standardization, each new entrant must decide when, and with which standards, to enter the market.

The process of standardization sometimes involves the transition from an old technology to a new and improved one. From this viewpoint, it can be associated with the eras that define a technology cycle (Anderson & Tushman, 1990; Tushman & Murmann, 2002). The cycle starts when an old technology that has been stable for some time progressively fails to resolve new “technological bottlenecks” (Rosenberg, 1969). Following this, a new set of technologies or insights emerges that changes what people regard as possible—the so-called “era of ferment,” as defined by Anderson and Tushman (1990). During this era, multiple alternative technologies appear as people find different ways of bringing new technology to market. Wiegmann et al. (2017) conceptualize these different technologies as multiple standards battling for dominance—which is, in their definition, a process of standardization. At some point, the battle for dominance ends, and a dominant design appears—i.e., a winning standard (Suarez, 2004; Fontana, 2008). The emergence of the winning standard ushers in a new era—the “era of incremental change”—and finalizes the standardization process. By this time, a single technology can be claimed as both the “dominant design” and the “winning standard,” in the respective terminologies of the two literature streams.

A firm that enters a market undergoing standardization must confront the risk of entering “too late,” but also the risk of entering with the “wrong” standard (i.e., one that will eventually fail to dominate the market). A firm that enters early with products based on the “right” standard will benefit from early-mover advantages (Folta, Johnson, & O’Brien, 2006), network effects (Farell & Klemperer, 1967), and high demand-side switching costs (Farell & Shapiro, 1988). However, it is not possible to predict with certainty which standard will ultimately become dominant. One way for a firm to increase its chances of introducing the “right” standard is to diversify its risks by entering the market with products based on multiple standards, and thus broadening its innovation portfolio (Klingebiel & Rammer, 2014). However, entering the market with products based on multiple standards is costly. One way of reducing these costs is through product modularity. Product modularity will incur an upfront development cost (i.e., designing interfaces and separating products into modules), but lower the risk of competing in several standards over time.

From the prior review, we note that product modularity could be beneficial for firms during market entry. In the next section, we present a real-options model to outline the conditions under which using product modularity to introduce products based on multiple standards to a market is the best option for a firm. We show how, in general, the condition depends on the interplay between early-mover advantage, switching costs, and development costs.

3. MODEL

This section presents a model that captures the fundamental entry choices available to a single firm when entering a market where multiple standards compete for domination. We assume that the firm knows that there are N competing standards ($\{r_1, r_2, \ldots r_N\}$) in the market. However, there is uncertainty over which standard will become dominant. The model will determine how many standards are most beneficial for the firm to introduce ($M \subseteq \{r_1, r_2, \ldots r_N\}$), as well as which type of product architecture (integral or modular) the firm should develop ($x \in \{integral, modular\}$).

The benefits of choosing a specific entry strategy are denoted by $P_X(M,t)$ and depend on three values: a) $D_x(M)$, the development costs of introducing M products, based on all its chosen standards, to the market; b) $S(M,t)$ the switching costs for the firm if the "wrong" (i.e., non-dominant) standard is chosen; and c) $E(M,t)$ the early-mover advantage that the firm can gain if it enters the market early with product(s) based on the (future) dominant standard.[1] The potential benefits of choosing a specific entry strategy can be written as:

$$P_x(M,t) \equiv E(M,t) - S(M,t) - D_x(M) \tag{1}$$

In our model, each firm adopts either an integral or a modular product architecture for market entry; for brevity, we refer to these two groups as "integral" and "modular" firms respectively. A firm that employs an integral product architecture develops products with an integrated structure that cannot be broken down further. In contrast, a firm that employs product modularity separates its products into different modules, such that modules of the same type or functionality are interchangeable. Both integral firms and modular firms can introduce products based on multiple standards to the market. To do this, the integral firm will need to design M products based on M different standards from the bottom up. In contrast, the modular firm can reuse modules, and hence reduce development costs, when introducing its M different standards to the market.

[1] Note that in the model, the development costs, $D_x(M)$, depend on the product's architecture, but the switching costs, $S(M)$, and early-mover advantage, $E(M)$, do not. This reflects our assumption that customers do not care about the product architecture *per se*, but only about the standard that the products embody.

3.1 Probability of winning the standards battle

For the sake of simplicity, we consider that modular products comprise only two modules. One module is used in all the *M* products developed by a modular firm. The firm needs to design this module once, but it can reuse it *M* times: once for every standard brought to the market. We call this module the "fixed module," as its design represents a fixed cost for the firm. The second module must be designed anew for each standard that the firm wishes to introduce to the market. We call this module the "standard-specific module."

An individual firm may enter the market with one or multiple products, and each firm introduces either integral products or modular products exclusively, according to its market entry strategy. Our model aims to define the conditions under which each one of these entry strategies can be optimal. To derive these conditions, we rely upon certain assumptions, which we describe below.

The first assumption states that at the point when the market-entry decision is taken, there are *N* standards ($\{r_1, r_2, \dots r_N\}$) that compete to dominate the market. When the standardization process is finished, however, only one standard out of the *N* will be dominant. Therefore, during the standardization process, each standard has a probability of dominating the market given by $p(r_i)$, and the sum of these probabilities adds up to 1:

$$\sum_{i \in M}^{N} p(r_i, t) = 1 \tag{2}$$

If a firm enters the market with products that comply with *M* out of the *N* standards ($M = \{r_1 \dots r_M\} \subseteq N$), then it will have a probability $p(M, t)$ of entering the market with the future dominant standard:

$$p(M, t) = \sum_{i \in M} p(r_i, t) \tag{3}$$

The higher the number of standards with which a firm enters the market, the higher the probability it will introduce one that embodies the future dominant standard. If the firm introduces products based on *all* the *N* standards to the market, it can be sure that it has invested in the future dominant standard. Any other combination of standards in its product range will lead to some degree of uncertainty.

The second assumption simplifies the comparison between entry strategies. We assume that during the standardization process (i.e., before the dominant standard emerges), it is entirely uncertain which alternative will dominate the market. That is, during the standardization process, we assume that each alternative standard has the same probability of ending up dominating the market. Namely:

$$p(r_i, before) = \frac{1}{N}\ \forall\ r_i \in \{r_1, r_2, \dots r_N\} \tag{4}$$

Using Equation (4), if a firm introduces products based on *M* standards to the market, we can compute the probability of introducing the future dominant standard by counting the number of standards introduced (i.e., the cardinality of the set).[2] Therefore:

$$p(M, before) = \frac{M}{N} \tag{5}$$

When the standardization process ends, we assume that only one standard will be dominant, and therefore uncertainty ends:

$$p(M, after) = \begin{cases} 1\ if\ the\ firm\ invested\ in\ the\ future\ dominant\ standard \\ 0\ if\ not \end{cases} \tag{6}$$

Many firms enter the market, some before the emergence of the dominant standard and some afterwards. We employ Equation (5) for the former and Equation (6) for the latter.

3.2 Early-mover advantage and switching costs

Relying upon the prior assumptions, we can define the potential benefits from Equation (1). The first step involves the definition of switching costs, *S(M)*, and early-mover advantages, *E(M)*, for a firm that decides to introduce products based upon *M* standards to the market. We

[2] To simplify the notation, we continue to refer to *M*, but now as a number rather than as a set of standards.

start with early-mover advantage. A firm that enters early (i.e., before the emergence of the dominant standard) and with the future dominant standard will enjoy an early-mover advantage:

$$E(M, before) = \sum_{r_k \in M} \mathrm{e}(r_k, t) \tag{7}$$

The more standards one firm introduces to the market, the higher its likelihood of benefiting from early-mover advantage. If one firm introduces products based on all *N* standards to the market, it will be certain to enjoy early-mover advantage. Thus:

$$E(N, before) = e \tag{8}$$

Following Equation (5), we can write the early-mover advantage of introducing *M* alternatives ($E(M)$) as:

$$E(M, before) = e \cdot p(M) = \frac{e \cdot M}{N} \tag{9}$$

In contrast, a firm that enters the market late (i.e., after the emergence of the dominant standard) will enjoy no early-mover advantage, regardless of which standard it introduces. Thus:

$$E(M, after) = 0 \tag{10}$$

Through a similar logic, a firm that enters before the emergence of the dominant standard, but fails to invest in it, will sustain switching costs for an amount of *s*. The more standards that firm introduces to the market, the lower its chances of incurring these switching costs. Therefore, we can define the switching costs of introducing *M* standards as:

$$S(M, before) = s\big(1 - p(M, before)\big) = s\left(1 - \frac{M}{N}\right) \tag{11}$$

A firm that enters after the emergence of the dominant standard will not incur any switching costs.[3] Therefore:

3 Here we assume that all industry players know about the emergence of the dominant standard. Therefore, firms would not introduce the "wrong" standard in the market after the "winning" standard had emerged. As detailed by Anderson and Tushman (1990), after a "dominant design" emerges in a market, firms aim at "elaboratin[g on] the

$$S(M, after) = 0 \tag{12}$$

3.3 Development-cost function

We define the development-cost function as a function that depends on the firm's product architecture at entry. We focus on two product architectures: integral and modular. Development costs vary depending on the chosen product architecture. However, we assume they are time-independent (i.e., they do not depend on whether entry occurs before or after the emergence of a dominant standard):[4]

$$D_x(M) = \begin{cases} D_{int}(M)\ if\ Integral \\ D_{mod}(M)\ if\ Modular \end{cases} \tag{13}$$

On the one hand, a firm that aims at introducing products based upon *M* different standards to the market by using an *integral product architecture* will need to design *M* products from the bottom up. We define development costs as:

$$D_{Int}(M) = \sum_{i \in M} d_{int}(r_i) \tag{14}$$

A firm that aims to introduce *M* different standards to the market by using a modular product architecture will still need to design *M* products. However, one benefit of modularity is that the products being developed can share some modules. Therefore, the modular firm does not need to design everything from the bottom up.

The simplest modular product consists of two modules and an interface that controls how the two modules interact. We call the module that is common to all the products based upon the *M* different standards the "fixed module," and define its (fixed) development cost as d_F. the development cost of the *M* standard-specific modules is defined as $D_S(M)$. The module

dominant design" by making more efficient products, instead of trying to introduce new standards to the market (p. 606). We follow this logic and assume that after the "winning" standard emerges, firms will only introduce this standard to the market. '

[4] We make this assumption because the task of designing a product entails similar activities for a firm independently of whether the product is designed "early" or "late" during the technology cycle. People need to be hired, interfaces created, modules designed—all these activities take time and resources, which are not directly time-dependent.

that has to be developed anew for every standard is called the “standard specific” module. Finally, the interface has development costs d_I. The total cost of developing *M* standards is thus given by:

$$D_{mod}(M) = d_F + d_I + D_S(M) = d_F + d_I + \sum_{i \in M} \mathrm{d_S}(r_i) \tag{15}$$

To compare the costs of developing *M* modular products with the costs of developing *M* integral products, we assume that the cost of developing each integral product or each standard-specific module is the same, regardless of the standard being developed for.[5] This implies that:

$$d_{int}(r_i) = d_{int}(r_j) = \mathrm{d_{Int}} \ \forall \ i \ and \ j \ where \ i \neq j \tag{16}$$

$$D_S(r_i) = D_S(r_j) = d_s \ \forall \ i \ and \ j \ where \ i \neq j \tag{17}$$

Equation (16) says that the cost of developing an integral product is independent of the standard that the product supports. Equation (17), in turn, does the same for the costs of developing the standard-specific modules. With this, Equations (14) and (15) can thus be simplified:

$$D_{Int}(M) = \mathrm{M} \cdot \mathrm{d_{Int}} \tag{18}$$

$$D_{mod}(M) = d_F + d_I + M \cdot d_s \tag{19}$$

Where $\mathrm{d_{Int}}$ is the cost of developing an integral product, d_F the cost of developing the fixed module, and d_s the cost of developing the standard-specific module for any of the different standards.

Equations (18) and (19) currently identify four sources of costs: d_{int}, d_F, d_I, and d_S. We need a further assumption to be able to compare the costs of integral and modular products.

[5] This assumption of “equality of costs of integral product development” can be justified through a counterfactual. If the contrary were true, and some standards were much more expensive to develop but achieved the same benefits, firms would choose to develop other, less costly, standards.

We assume that the cost of developing one integral product and one modular product differ only in the cost of developing the interface between the modules of the modular product.[6] Thus:

$$D_{mod}(1) \ = \ D_{Int}(1) + d_I \tag{20}$$

The total cost of developing *M* integral products can instead be defined as:

$$D_{int}(1) = M \cdot (d_F + d_S) \tag{21}$$

Using Equation (20), we can build a comparable cost function to develop *M* products with different architectures. We can rewrite Equation (13) as:

$$D_x(M) = \begin{cases} D_{int}(M) = M \cdot (d_F + d_S) \quad if\ integral \\ D_{mod}(M) = \ d_F + \mathrm{M} \cdot d_S + d_I \quad if\ modular \end{cases} \tag{22}$$

The above equations can now be employed to identify the conditions under which it is more beneficial to a) enter with just one product; b) enter with *M* modular products; or c) enter with *M* integral products.

3.4 Modular entry vs. integral entry

A firm will choose to enter the market with a modular product strategy when the cost of developing modular products is lower than the cost of developing integral products—that is, when the development costs in the bottom row of Equation (22) are lower than the development costs in the top row:

$$d_F + \mathrm{M} \cdot d_S + d_I < M \cdot (d_F + d_S) \tag{23}$$

Which can be simplified as follows:[7]

$$\mathrm{d_I} < (M - 1) \cdot d_F \tag{24}$$

Note that from Equations (5) and (6), the only appropriate number of standards to introduce after a single standard has become dominant is *M* = 1. Therefore, modular products

[6] This assumption of "equality of cost across standards" is tantamount to saying that development costs differ only in the cost of designing the interface. It can be motivated by the empirical observation that designing the interface is a lengthy process that requires significant discussion and deliberation.

[7] Note that if *M* = 1, the inequality never holds, and thus it is always less costly to design integral products. For this reason, when a firm enters with one single product, the product should employ an integral product architecture.

can only be economically viable during the "era of ferment." This is the case as d_I is always greater than zero. This leads to our first proposition:

Proposition 1a: Product **modularity** should be used as a market-entry strategy more often during the "era of ferment" of the technology cycle than during the "era of incremental change."

Proposition 1b: Product **integrality** should be used as a market-entry strategy more often during the "era of incremental change" of the technology cycle than during the "era of ferment."

3.5 Entry with multiple modular products

Having derived the condition under which a particular firm may decide to enter with either a modular or integral strategy, we focus on the decision to enter with multiple modular products. From Equation (22), a firm will find it beneficial to enter the market with multiple (*M*) modular products supporting different standards if the following inequality holds[8]:

$$E(M, before) - S(M, before) - D_{mod}(M) > E(1, \text{before}) - S(1, \text{before}) \ - D_{int}(1) \quad (25)$$

and rewritten as:

$$\frac{eM}{N} - s\left(1 - \frac{M}{N}\right) - (d_F + M \cdot d_S + d_I) > \frac{e}{N} - s\left(1 - \frac{1}{N}\right) - (d_F + d_S + d_I) \qquad (26)$$

Simplifying the expression, the firm will find it beneficial to enter with *M* standards and a modular architecture if the following inequality holds *before the emergence of the dominant standard*:

$$d_S < \frac{e + s}{N} \qquad (27)$$

Given that d_S has a positive value, inequality (27) can only hold during the "era of ferment"—specifically, in conditions where there are significant expected benefits to entering

[8] In Appendix A, we show the derivation for the boundary conditions under which a firm will find it beneficial to introduce multiple integral products. These conditions are always more stringent that the conditions for introducing multiple modular products.

the market early (i.e., when $(e+s)/N$ is high). The inequality will not hold in all conditions, but firms may find it beneficial to introduce multiple modular products in cases where there are significant expected benefits. This leads us to our second propositions:

Proposition 2a: Firms should enter the market with **multiple products** more often during the "era of ferment" of the technology cycle than during the "era of incremental change."

Proposition 2b: Firms should enter the market with **one product** more often during the "era of incremental change" of the technology cycle than during the "era of ferment."

Jointly, these two propositions predict that more firms should enter the market with multiple modular products during the "era of ferment," than during the "era of incremental change". In contrast, firms should enter the market more often with one product during the "era of incremental change" than during the "era of ferment".

4. EMPIRICAL EXERCISE

The real-options model proposed in the previous section identifies the optimal entry strategies available to firms wishing to enter an industry with either a modular or integral product architecture, and with one product or multiple, in a context characterized by technological uncertainty. Our model serves as a formal and rigorous framework for the empirical exercise described below. In this section, we present an industry context in which we will empirically test our propositions. This section presents a broader explanation of the different ways product modularity helps firms during market entry, in addition to those inferred from literature and formalized in our real-options model.

4.1 Context

The context analyzed is the Local Area Network (LAN) equipment manufacturing industry during a time of intense technological development: the decade spanning 1990 to 1999.[9] The

[9] LANs constitute the infrastructure that enables computers, other types of end-stations, and/or peripherals to be linked to form a network connecting different users within a relatively small area, such as a university campus or different buildings on a company site. The functioning of LANs can be described as follows. A computer wanting to transmit some information breaks the data into packets. The packets are sent to the LAN through adapter cards,

LAN industry during the 1990s represents an interesting case to investigate the effects that modular design strategies have on firms' entry strategies because, throughout the 1990s, LAN engineers struggled to eliminate one key technical bottleneck: network congestion. Firms followed two paths to achieve this aim. The first led to the development of high-speed standards. At the beginning of the 1990s, several high-speed standards (FDDI, Fast Ethernet, ATM/Asynchronous Transfer Mode, and 100VG-AnyLAN) were "battling" to become the main successor to Ethernet, the most widely adopted standard at the time. After a period of uncertainty, the battle ended in 1994 when Fast Ethernet became the *de facto* standard.[10] Hence, we can clearly identify a point at which a single standard emerged as dominant.

The second solution path led to changes and improvements in the hardware design of existing equipment. Modularity played an important role in facilitating these changes. The physical design of modular equipment revolved around a chassis (i.e., the fixed module) with several slots that housed the modules responsible for connectivity under a specific standard (i.e., the standard-specific modules).

Modularity also became part of the manufacturers' strategy to manage incremental changes and transform them into market opportunities. From users' viewpoint, modularity allowed existing investments to be enhanced by new technical features. First, when buying modular products, users could mix and match modules of different standards, thereby increasing the variety of technologies supported. Second, modular equipment increased users' flexibility to cope with small changes in LAN architectures and continue to benefit from previous investments in management software and training. Finally, when coupled with the deployment of high-speed standards, modularity could improve equipment performance. To

which physically connect the computer to the channel. Once sent to the channel, packets travel first either to a hub or to a switch. In early LANs, packets normally travelled to a hub, i.e., a device that sends the packets it receives to *all* users connected to a specific LAN segment, so that each user can "see" every packet. Switches are more sophisticated, because they can select the specific user to whom packets are to be sent.

[10] Fontana (2008) provides a detailed account of the events that led to the dominance of Fast Ethernet.

enjoy these benefits, buyers seemed willing to pay a relatively high price to purchase modular equipment (Fontana, 2007).

From the manufacturers' viewpoint, modularity enabled firms to master technical change while avoiding product cannibalization and keeping control of their existing customer base. In the LAN environment, standards were "open," and users could freely mix and match equipment that was manufactured by different firms but still supported the same standards. However, modules would only operate with the chassis from the same manufacturer. By commercializing modular equipment, firms could coordinate the migration of their installed base of users from one generation of equipment to the next.

Modularity also responded to an explicit strategy of "exploring" the product design space. This strategy had two aims. The first was to target new categories of users. Modular products generally offered users higher capacity and port density (i.e., more sockets on the hub) than integral products, and users with large LANs requested these features. The second aim was to restrict competition. A modular design allowed manufacturers to add new features and improve upon the existing design quickly. However, designing modular products was costly. Big incumbents (e.g., Ungermann Bass, Cabletron Systems, 3Com, Chipcom, DEC, etc.) could sustain the higher costs, while newer firms (e.g., Kalpana, Synernetics, Grand Junction, etc.) opted for integral products instead.

Finally, modularity played a central role in the introduction of high-speed standards. Indeed, product modularity was popular in the first half of the 1990s when multiple standards were competing. Modularity enabled manufacturers to introduce multiple standards to the market, thereby confronting existing uncertainty on which standard would finally prevail. In simple terms, given an installed base of LAN equipment, it was much easier to adopt a new standard if the equipment was modular. The new standard was implemented simply by sliding the standard-specific module into the slot provided on the product chassis (i.e., the fixed module). Such ease of upgrade was beneficial to both buyers and manufacturers, who were not required to commit to any one of the standards that were competing for supremacy until 1994—the year when Fast Ethernet emerged as the winning standard.

4.2 Hypotheses

From the above narrative, we can highlight some aspects that are crucial for testing our propositions. First, before 1994, when Fast Ethernet ultimately came to dominate, multiple standards battled for market domination, and firms frequently used modular product designs. However, we would expect the use of product modularity to decrease after 1994. Based on Propositions 1a and 1b, we can thus hypothesize:

Hypothesis 1a: Firms should enter the market more frequently with **modular products** before 1994 than afterwards.

Hypothesis 1b: Firms should enter the market more frequently with **integral products** after 1994 than beforehand.

Second, the *de facto* standardization of Fast Ethernet in 1994 represented a "watershed" in the LAN industry, reducing technological uncertainty and reorienting firms' effort toward developing product design rather than the proliferation of alternative solutions that had characterized the previous phase. We can thus identify an early phase in the industry (i.e., the pre-1994 "era of ferment"), characterized by high early-mover advantage and switching costs, and a late phase (i.e., the post-1994 "era of incremental change") in which both early-mover advantage and switching costs were low or nonexistent. The early-mover advantages and high switching costs, in addition to the noted benefits of product modularity for this industry, lead us to expand Propositions 2a and 2b as follows:

Hypothesis 2a: Firms should enter the market more frequently with **multiple products** before 1994 than afterwards.

Hypothesis 2b: Firms should enter the market more frequently with **one product** after 1994 than beforehand.

4.3 Data

4.3.1 Sample of products and firms

To test the real-options model's predictions empirically, we use a dataset from the LAN industry during the 1990s. The year 1999 is a good "cut-off point" because a reasonable mix of models supporting all the different standards was still present in the sample at this time. The dataset includes 1,068 LAN products manufactured by 85 firms that introduced at least one product supporting a high-speed standard during the study period.[11] Information on the products was collected from *Network World,* a specialized publication targeted at network professionals. Trade journals (*Network World* in particular) give extensive coverage to the introduction of new products by reporting product characteristics, prices, manufacturers, and dates of introduction. When possible, dates were double-checked against press communications and manufacturers' product announcements.

4.3.2 Variables

The variables we will use for the empirical analysis are as follows. MODULAR: this is a dichotomous variable equal to 1 if the firm entered the market with a modular product architecture, or 0 if the firm entered the market with an integral product. MULTIPLE STANDARDS: is equal to 1 if the first product introduced to the market by the firm supported more than one standard; 0 if not. BEFORE 1994: is a dichotomous variable with a value of 1 if the firm entered the market before 1994, the year Fast Ethernet became the *de facto* standard. If the firm entered from 1994 onwards, the value is 0. NUMBER OF PRODUCTS: the count of LAN products introduced by each firm each year. This is the only variable that is collected yearly; the others are time-invariant.

[11] Prior studies employing the same dataset are Fontana and Nesta (2009) and Fontana (2007).

5. RESULTS

Table 1 reports the descriptive statistics for this study's main variables, together with their correlation matrix.[12]

Table 1: Descriptive statistics and zero-order correlations

Variable	**1.**	**2.**	**3.**	**4.**
1. # of Products	1			
2. Multiple Standards	0.086 (0.433)	1		
3. Before 1994	0.236 (0.030)	0.371 (0.000)	1	
4. Modular	0.273 (0.011)	0.369 (0.001)	0.482 (0.000)	1
Mean	5.765	41 of 85	30 of 85	46 of 85
SD	8.417			

Note: p-value of correlation in parenthesis

5.1 Modularity and entry strategy

Eighty-five firms entered the market during the decade of study: 30 before 1994, during the "era of ferment," and 55 from 1994 onwards, during the "era of incremental change." Of the 85 firms, 41 entered the market with products that supported two or more standards.

Table 2 shows the number of firms that entered the market with modular products during each technology cycle era. From Hypothesis 1, we would expect that before 1994, most firms enter the market with modular products, and the opposite after 1994. We find evidence for this expectation. Almost twice as many firms entered the market during the "era of incremental change," and of these, more than two-thirds entered with integral products. During the "era of ferment," the situation was very different. Product modularity was the entry strategy chosen by almost 9 out of 10 firms early on in the technology cycle. The χ^2 test of the frequency table is significant with a p-value < 0.001 (χ^2=24.14). This result supports Hypothesis 1.

[12] Table 1 shows the sum of products introduced by a firm during the period of study. However, we have information on the number of products introduced every year: This information is used as the dependent variable in the regressions of Table 4.

Table 3 presents the number of firms that entered the market with one or multiple standards during each era of the technology cycle. We expect that before 1994, most firms would enter with multiple standards, whereas from 1994 onwards, firms would enter with just a single standard. Our results fall in line with this expectation. The number of entrants during the “era of incremental change” (i.e., from 1994 onwards) is higher overall than in the prior period (55 vs. 30), and the majority of firms entering do so with only one standard (36 vs. 19). On the contrary, during the “era of ferment” (i.e., before 1994), firms tend to enter with multiple standards rather than just one (22 vs. 8). All in all, this evidence is consistent with Hypothesis 2, which predicted that firms entering from 1994 onwards should be more likely to enter with single products than those entering before 1994. The χ^2 test is significant with a p-value of 0.001 (χ^2=11.70).

Table 2: Number of entrants by product architecture used at entry

	Integral	Modular	Total
Before 1994	4	26	30
From 1994	38	17	55
Total	42	43	85

Table 3: Number of entrants by number of standards supported

	Single standard	Multiple standards	Total
Before 1994	8	22	30
From 1994	36	19	55
Total	44	41	85

5.2 Modularity and firm performance

So far, we have presented empirical evidence concerning modularity and entry strategy for the firms in our sample. Our findings suggest that product modularity was particularly suitable as an entry strategy in the “era of ferment” (i.e., before 1994). Product modularity was associated with entry with multiple products instead of a single one by both new entrants and incumbent firms. Though important and consistent with our model’s prediction, this evidence is mostly descriptive and based on univariate analysis. It also considers only the effect of modularity on the pattern of firms' entry. What about the consequences of modularity for the post-entry performance of firms? For modularity to be strategically important, we would also expect firms that chose modular entry to perform better than others who chose to follow another strategy. In this section, we explicitly tackle these issues.

We can employ the argument developed above to speculate about the consequences of choosing a specific entry strategy. Indeed, in Section 4.1 we identified several advantages of modularity, ranging from greater mastery of technical change and better-coordinated migration of existing customers from one platform to the next to quicker exploration of the product-design space when looking to new solutions. These advantages indicate that firms that employed a modular strategy at entry should do relatively better than those who did not. We explore these ideas by employing the frequency of product introductions as a performance indicator.

We are interested in estimating the effects of modularity on a firm's product-introduction strategies. Specifically, we want to know whether these effects change in the transition from the "era of ferment" (pre-1994) to the "era of incremental change" (post-1994). To carry out our analysis, we employ a random effect Poisson estimator. More specifically, for each firm i, we estimate the following baseline equation:

$$E[y_{it}|X_t] = exp[\beta_0 + \beta_1 \cdot Before1994_t \ + \delta_t] \qquad (Model\ 1)$$

where y is the number of products produced by firm i at time t, $Before1994_{t.}$ is an indicator variable that captures the effect of *de facto* Fast Ethernet standardization in the industry, and δ_t is a full set of entry-year variables. If, in line with our previous findings, the change of regime from the "era of ferment" to the "era of incremental change" reduced uncertainty, we would expect β_1 to be positive and significant.

We then modify Model (1) as follows:

$$E[y_{it}|X_t] = exp[\beta_0 + \beta_1 \cdot Modular_i + \ \beta_2 \cdot Before1994_t + \delta_t] \qquad (Model\ 2)$$

and again:

$$E[y_{it}|X_t] = exp[\beta_0 + \beta_1 \cdot Modular_i \cdot Before1994_t + \beta_2 \cdot Modular_i + \beta_3 \cdot Before1994_t + \delta_t] \; (Model\ 3$$

where $Modular_i$ is an indicator variable equal to 1 if firm i entered with a modular product, and 0 if not. If modularity brought benefits in terms of the number of new product introductions, we would expect β_1 to be positive and significant in Model (2). Finally, if product modularity was particularly beneficial before the emergence of the dominant standard (i.e., during the "era of ferment"), we would expect β_1 to be positive and significant in Model (3).

Table 4 reports the results of the estimations. In Model (1), we look at the effect on product introduction depending on the timing of market entry. β_1 is positive and significant, suggesting that firms that entered the market before 1994 introduced more products than those that entered from 1994 onwards. Model (2) looks instead at the effect of the use of product modularity during market entry on the number of products introduced by the firm during the study period. Again, the positive and significant coefficient suggests that modular entry was associated with an increase in the number of product introductions. Finally, Model (3) presents the full specification together with the interaction of the previous two variables. In this case, β_1 is still significant and positive, suggesting that the use of product modularity during market entry did increase the number of product introductions. However, as we expected, this effect was limited to the "era of ferment" (i.e., the period before the emergence of the dominant standard). Interestingly, β_2 and β_3 are not significant in Model (3); this gives support to the idea that product modularity was beneficial to firms, but only during the "era of ferment." In summary, we find support for our expectation that the firms who employed product modularity during the "era of ferment" managed to outperform other active firms in this industry.

Overall, we find that firms employed product modularity predominantly during the "era of ferment," which gives support to Hypothesis 1. We also find that more firms introduced multiple products during the "era of ferment," supporting Hypothesis 2. Finally, we find that the firms that enter the market using product modularity during the "era of ferment" outperform other firms in the market by introducing more products. In the next section, we discuss these results and the implications of the proposed real-options model.

Table 4: Estimating the post-entry performance of firms in the LAN industry

	Number of products introduced: y_{it}		
	(1)	**(2)**	**(3)**
Before1994	0.784***		-0.0067
	(0.091)		(0.2381)
Modular		0.618***	-0.038
		(0.095)	(0.151)
Before1994 x Modular			0.913***
			(0.277)
Intercept	-3.714***	-3.704***	-3.703***
	(0.580)	(0.581)	(0.582)
Sigma	4,589 (p = 0.747)	4,589 (p = 0.747)	4,589 (p = 0.747)
	(14,210)	(14,210)	(14,210)
Year Dummies	Yes	Yes	Yes
Company Dummies	Yes	Yes	Yes
Log-likelihood	-949.05	-964.15	-939.86

Note: Standard errors in parenthesis, *** $p < 0.001$

6. DISCUSSION

In this paper, we have introduced a real-options model and an empirical example to explain how firms can use product modularity to enter markets in the presence of competing standards. In the model, we proposed that using modularity to bet on multiple standards allows firms to accrue early-mover advantages, and potentially reduce switching costs (Suarez, 2004; Folta et al., 2006; Farell & Shapiro, 1988), despite the higher development costs of modular product architectures. We also tested these claims and found empirical support for them. In the empirical analysis, we found that the firms did indeed employ product modularity during times of rapid change. We also found that the firms that did so performed better than firms that employed other product-entry strategies.

Prior studies have noted the joint decision to enter the market early, with multiple product features, or late, with a limited feature range (Klingebiel & Joseph, 2016). We validate their findings, as we find support for Hypothesis 2—i.e., firms that enter the market with multiple products do so during relatively more during the “era of ferment,” whereas late in the cycle, firms tend to enter with a single product. However, we extend the findings from Klingebiel and Joseph (2016) by showing how firms can use product architecture strategically, and reduce their exploration costs during the “era of ferment” by using product modularity.

Although we do not explore it empirically, the use of product modularity also allows firms to reduce their switching costs later on. A firm that introduced several modular products to the market will only need to develop a single new standard-specific module, rather than an entirely new integral product. Based on these lower switching costs, we can argue that product modularity is a way of delaying product differentiation while remaining active in the market. In our model, firms employ market strategies to diversify their risk while delaying their product differentiation; this stands in contrast to prior models, which modeled firms as using non-market strategies (e.g., deferring market entry) in order to delay product differentiation (Lee & Tang, 1997; Swaminathan & Tayur, 1998).

In the industry we studied, the standards employed by the firms were open. Firms did not need to incur the costs of developing the standards from scratch, and could freely choose the ones they wished to introduce to the market. However, in many industries, standards are not open (Suarez, 2004). When this is the case, our model’s propositions would still hold, but the costs of developing standards would increase significantly (d_S). The LAN industry was therefore a good setting for studying the implications of the model.

We found that firms benefit from using product modularity during the “era of ferment”—the most turbulent phase of the technology cycle. Brusoni, Marengo, Prencipe, and Valente (2007) predicted that a firm would benefit from making their products less modular in turbulent times, as turbulence requires the constant updating of designs. Yet, their simulation model is built on the assumption that any given product’s technical features drive competitive dynamics. In the LAN setting, standards played a pivotal role in explaining performance and survival. Standards could be adopted by our firms, but were not developed by them. The availability of open standards adds an extra layer of complexity to our model's entry decision, which was not in theirs. Hence, we argue, the different results. The use of product modularity that we put forth in this paper manages the market's turbulence by leading the firm to invest in multiple competing standards. By managing the market's turbulence, our model puts forth a new way in which modularity can be valuable during market entry.

The benefits of product modularity for organizational problem solving have been studied extensively (Brusoni et al. 2001; Baldwin & Clark, 2000; Garud et al., 2009). In this paper, we stepped away from this inquiry line and show how product modularity helps firms manage their environments' uncertainty during market entry.

Folta and colleagues had shown how market entry is a highly strategical decision, especially in an uncertain environment. Similarly, Klingebiel and Joseph (2016) had found how firms manage the environment's uncertainty by broadening or narrowing the range of features included in their products. This empirical result is in line with the intuition of Garud et al. (1997) that the costs of commission and omission errors vary during a technology cycle. Our study builds on these three streams of literature to show how firms can use product modularity strategically to take control of uncertainty and diversify their product portfolio during market entry. Uncertainty is a key strategic problem for any firm, and we find that product modularity can help to manage it and thus improve firm performance.

The literature on modularity has focused on its benefits for organizational problem-solving (Ethiraj & Levinthal, 2004; Brusoni et al., 2007; Fang & Kim, 2018). This paper contributes to broadening the strategic uses of product modularity by showing how it can be used to manage uncertainty through the diversification of risk. However, modularity is not a panacea; there are traps inherent in modularization (Brusoni, Prencipe, & Pavitt, 2001;

Chesborough & Kusunoki, 2001). Future work could extend the study of modularity traps by studying the limits inherent in the use of product modularity for problem-solving and managing uncertainty. Merging both streams could be a valuable addition to our understanding of how modularity influences new product development, market entry, and firm performance.

APPENDIX

Appendix A. Entry with multiple integral products

In the model, we explored the conditions under which the use of product modularity can be beneficial to firms—specifically, when it helps them introduce multiple modular products to the market. It is plausible that some environments could offer so much benefit from introducing multiple products that they might even warrant the introduction of multiple integral products. We explore these conditions in this appendix. From Equation (22), a firm will find it beneficial to enter the market with multiple (*M*) integral products supporting different standards if the following inequality holds:

$$E(M) - S(M) - D_{Int}(M) > E(1) - S(1) \ - D_{Int}(1) \tag{28}$$

which, from our previous calculations, implies that:

$$\frac{eM}{N} - s\left(1 - \frac{M}{N}\right) - M \cdot d_{int} > \frac{e}{N} - s\left(1 - \frac{1}{N}\right) - d_{int} \tag{29}$$

Simplifying, this inequality entails that a firm enters with multiple integral products if the switching costs and the early-mover advantage are higher than a threshold value given by:

$$d_{int} < \frac{e + s}{N} \tag{30}$$

Note that based on Equation (21), we can be sure that $d_I < d_S$, and thus inequality (27) holds true under a broader set of conditions than inequality (30). This implies that firms will find it beneficial to introduce multiple modular products more often than multiple integral products.

Appendix B: Firm Survival

In addition to the variables studied in Table 1, we have access to the firms' survival during the period of study. We codify survival as a dichotomous variable with a value of 1 if the firm remains active in the market at the end of the study, the year 2000, and a value of 0 otherwise. Adding survival to the other variables of the study, leads to Table A.1

Table A.1: Descriptive statistics and zero-order correlations for extended dataset

Variable	1.	2.	3.	4.	5.
1. Survived	1				
2. # of Products	0.203 (0.062)	1			
3. Multiple Standards	-0.171 (0.118)	0.086 (0.433)	1		
4. Before 1994	-0.185 (0.090)	0.236 (0.030)	0.371 (0.000)	1	
5. Modular	-0.110 (0.314)	0.273 (0.011)	0.369 (0.001)	0.482 (0.000)	1
Mean	30 of 85	5.765	41 of 85	30 of 85	46 of 85
SD		8.417			

Note: p-value of correlation in parenthesis

Table A.2 presents the number of firms that survived according to all possible market entry strategies and entry times. We observe that the highest survival rate is for firms that entered during the “era of ferment” with a modular product architecture but only one standard. These firms were relatively uncommon at the time. Only 6 of the 30 firms that entered the market during the “era of ferment” followed this entry strategy. In contrast, two-thirds of the firms employed the entry strategy predicted by the real-options model (i.e., multiple modular products). As we show in Table 4, the firms that followed predicted market entry strategies from the real-options model achieve the highest performance. However, from Table A.2, we can see that the highest probability of survival was for the minority of firms that employed product modularity for market entry but entered the market with one single product. Regression analyses of survival as a dependent variable are not robust due to the data set's small size.

Table A.2: Descriptive statistics and zero-order correlations

Era	Architecture	# of Standards	Total	Survived	% Survived
Ferment	Integral	Multiple	2	0	0.0
Ferment	Integral	Single	2	0	0.0
Ferment	Modular	Multiple	20	6	30.0
Ferment	Modular	Single	6	4	66.7
Integral	Integral	Multiple	11	3	27.3
Integral	Integral	Single	27	14	51.9
Integral	Modular	Multiple	8	3	37.5
Integral	Modular	Single	9	3	33.3

REFERENCES

Anderson, P., & Tushman, M. L. (1990). Technological discontinuities and dominant designs: A cyclical model of technological change. *Administrative science quarterly*, 604-633.

Baldwin, C. Y., & Clark, K. B. (2000). *Design rules: The power of modularity (Vol. 1).* MIT press.

Baron, D. P. (1995). Integrated strategy: Market and nonmarket components. *California management review,* 37(2), 47-65.

Brusoni, S. (2005). The limits to specialization: problem solving and coordination in 'modular networks'. *Organization Studies*, 26(12), 1885-1907.

Brusoni, S., Marengo, L., Prencipe, A., & Valente, M. (2007). The value and costs of modularity: a problem-solving perspective. *European Management Review*, 4(2), 121-132.

Brusoni, S., & Prencipe, A. (2006). Making design rules: A multidomain perspective. *Organization science*, 17(2), 179-189.

Brusoni, S., Prencipe, A., & Pavitt, K. (2001). Knowledge specialization, organizational coupling, and the boundaries of the firm: why do firms know more than they make?. *Administrative science quarterly*, 46(4), 597-621.

Chesbrough, H., & Kusunoki, K. (2001). The modularity trap: innovation, technology phase shifts, and the resulting limits of virtual organizations. *Managing industrial knowledge*, 202-230.

Christensen, C. M., Suárez, F. F., & Utterback, J. M. (1998). Strategies for survival in fast-changing industries. *Management scienc*e, 44(12-part-2), S207-S220.

Ethiraj, S. K., & Levinthal, D. (2004). Modularity and innovation in complex systems. *Management science*, 50(2), 159-173.

Fang, C., & Kim, J. H. (2018). The power and limits of modularity: A replication and reconciliation. *Strategic Management Journal*, 39(9), 2547-2565.

Fang, C., Lee, J., & Schilling, M. A. (2010). Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organization Science*, 21(3), 625-642.

Farell, J., & Klemperer, P. (1967). Coordination and Lock-In: Competition with Switching Costs and Network Effects. *Handbook of Industrial Organization*. MA and R. Porter, ed.

Farell, J., & Shapiro, C. (1988). Dynamic competition with switching costs. *The RAND Journal of Economics*, 123-137.

Folta, T. B., Johnson, D. R., & O'Brien, J. (2006). Uncertainty, irreversibility, and the likelihood of entry: An empirical assessment of the option to defer. *Journal of Economic Behavior & Organization*, 61(3), 432-452.

Folta, T. B., & O'Brien, J. P. (2004). Entry in the presence of dueling options. *Strategic Management Journal*, *25*(2), 121-138.

Fontana, R. (2007). Technical change, prices and communications technology: Insights from the Local Area Networking industry. *Technological Forecasting and Social Change,* 74(3), 313-330.

Fontana, R. (2008). Competing technologies and market dominance: standard "battles" in the Local Area Networking industry. *Industrial and Corporate Change*, 17(6), 1205-1238.

Fontana, R., & Nesta, L. (2009). Product innovation and survival in a high-tech industry. *Review of Industrial Organization*, 34(4), 287-306.

Garud, R., Kumaraswamy, A., & Langlois, R. (Eds.). (2009). *Managing in the modular age: architectures, networks, and organizations*. John Wiley & Sons.

Garud, R., Nayyar, P. R., & Shapira, Z. (1997). Technological choices and the inevitability of errors. Technological innovation: *Oversights and foresights*, 20-40.

Henderson, R. M., & Clark, K. B. (1990). Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative science quarterly*, 9-30.

Kalish, S., & Lilien, G. L. (1986). A market entry timing model for new technologies. *Management Science*, 32(2), 194-205.

Klingebiel, R., & Joseph, J. (2016). Entry timing and innovation strategy in feature phones. *Strategic Management Journal*, 37(6), 1002-1020.

Klingebiel, R., & Rammer, C. (2014). Resource allocation strategy for innovation portfolio management. *Strategic Management Journal*, 35(2), 246-268.

Langlois, R. N. (2002). Modularity in technology and organization. *Journal of economic behavior & organizatio*n, 49(1), 19-37.

Lee, H. L., & Tang, C. S. (1997). Modelling the costs and benefits of delayed product differentiation. *Management science*, 43(1), 40-53.

Lilien, G. L., & Yoon, E. (1990). The timing of competitive market entry: An exploratory study of new industrial products. *Management science*, 36(5), 568-585.

O'Brien, J. P., Folta, T. B., & Johnson, D. R. (2003). A real options perspective on entrepreneurial entry in the face of uncertainty. *Managerial and Decision Economics, 24(8), 515-533.*

Rosenberg, N. (1969). The direction of technological change: inducement mechanisms and focusing devices. *Economic development and cultural change*, 18(1, Part 1), 1-24.

Sanchez, R., & Mahoney, J. T. (1996). Modularity, flexibility, and knowledge management in product and organization design. *Strategic management journal*, 17(S2), 63-76.

Schilling, M. A. (2000). Toward a general modular systems theory and its application to interfirm product modularity. *Academy of management review*, 25(2), 312-334.

Schilling, M. A., & Steensma, H. K. (2001). The use of modular organizational forms: An industry-level analysis. *Academy of management journal*, 44(6), 1149-1168.

Simon, Herbert A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467-482.

Suarez, F. F. (2004). Battles for technological dominance: an integrative framework. *Research Policy*, 33(2), 271-286.

Sundali, J. A., Rapoport, A., & Seale, D. A. (1995). Coordination in market entry games with symmetric players. *Organizational behavior and Human decision processes*, 64(2), 203-218.

Swaminathan, J. M., & Tayur, S. R. (1998). Managing broader product lines through delayed differentiation using vanilla boxes. *Management Science*, 44(12-part-2), S161-S172.

Tushman, M., & Murmann, J. P. (2002). Dominant designs, technology cycles, and organizational outcomes. *Managing in the modular age: architectures, networks, and organizations, 316.*

Ulrich, K. (1994). Fundamentals of product modularity. In *Management of Design* (pp. 219-231). Springer, Dordrecht.

Wiegmann, P. M., de Vries, H. J., & Blind, K. (2017). Multi-mode standardisation: A critical review and a research agenda. *Research Policy*, 46(8), 1370-1386.

ACKNOWLEDGEMENTS

So much has changed in my life since I started as a research assistant at the TIM Chair in 2015. I have met great people, I have learned what social sciences is, and I have found my niche in this complex and shifting scientific world. None of this would have been possible without Stefano and Daniella taking a chance on me. I am grateful to them for taking this chance, and for giving me the freedom to define my own path. Few would have granted me such latitude—but looking back, it was the best thing that could have happened to me. I have learned immensely from both of you, and deeply enjoyed the research atmosphere you cultivate.

My colleagues at the chair helped me become who I am. Axel was the perfect starting mate, colleague, and friend. My research would be very different had we not started around the same time. I will always miss being able to find you next door and receiving answers to my questions from the depths of your knowledge. At the start of my PhD Amulya was a guiding light—another natural scientist trying to understand the meaning of terms like "entrepreneurship," "organizations," and "strategy" in the broader context of science.

It was through a conference organized by Amulya that I met Chengwei. We had a mentoring lunch together. While the lunch only lasted for a couple of hours, the mentoring has remained active in one way or another ever since. At that time, I had no idea that I would end up doing theoretical models. In physics and engineering, I was no theoretician—and five years ago, I would never have imagined that I would be one today. I say this, even though looking back now, it makes perfect sense.

A lot has happened at a personal level too. From the start, I had Barbara for support. When I lost home, you were my memory of it. I will always be grateful for this. You have supported me more than anyone—well, more than anyone other than Lena.

Lena, you came into my life and made it bigger, better, and more colorful. You defied my views of the world and in the process helped me learn to critique. I am grateful for who you have been to me, and I look forward to what the future will bring for us.

Finally, my family I am most grateful to all of you. You raised me to be the person I am, and gave me freedom to pursue my dreams. For me, the concept of distance took on a new meaning this year. I miss you dearly, and I am always grateful for your support and understanding.

I cannot name everyone here, but I can just say that all the coffees, meals, drinks, and conversations I have had around Weinbergstrasse 56/58 have made me who I am. I am thankful to all of you for the privilege of having met you.