

DISS. ETH NO. 26910

**Towards the Economy of Things:
Insights from the Convergence of Blockchain
Technology and the Internet of Things**

A dissertation submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

Mathieu Chanson

MSc. ETH Physics

born on January 23, 1991

Citizen of Zurich

accepted on the recommendation of

Prof. Dr. Elgar Fleisch, examiner

Prof. Dr. Frédéric Thiesse, co-examiner

Ass.-Prof. Dr. Felix Wortmann, co-examiner

2020

Abstract

Over the last decades, the Internet of Things (IoT) has emerged as a vision of the computer of the 21st century, in which virtually any computing device could be connected to the Internet. Recently, this vision of ubiquitous computing has started to become reality and billions of connected things are deployed today, from smart home devices to cars to industrial machinery. For the coming years ahead, the number of these connected devices is projected to continue to grow at a rapid speed. The current stage of the IoT is characterized by the focus on things which are simply *connected* to the Internet and *generate data* which, in turn, enables new products, services and business models. Recently, technological progress has shifted the focus towards *interconnected* things which can *interpret data*, leading to an enhanced vision for the IoT called the Economy of Things (EoT). The EoT refers to a world in which IoT devices become increasingly intelligent things that are acting as autonomous agents in an economic system, instead of simply sensing, processing and exchanging data. In this context, devices are interconnected, exchanging data directly between each other in a peer-to-peer network, and intelligent in the sense that they can analyze available data and leverage this information to make decisions autonomously. A great showcase for an EoT application is the idea of robo-taxis, which autonomously drive around, spot passengers or parking options and settle payments with them, decide which workshops to go to for repairs and save money to buy an additional version of themselves, meaning another robo-taxi for the fleet.

The vision of the EoT is nurtured by technological developments which could build the infrastructural foundation for broader decentralization and disintermediation. In particular, a broad range of research scholars and practitioners alike argue that blockchain technology is a promising means to strengthen existing tendencies of disintermediation and build the technological foundation for the EoT. In this context, especially relevant properties of blockchains are the possibility to transact value (e.g., via cryptocurrencies) or incorporate business logic (e.g., via smart contracts) in an automated manner independent of intermediaries. In fact, the decentralized nature of blockchain technology allows trustless interactions and makes it possible to use a blockchain as a shared, secure computing platform between distrustful or anonymous agents. Ultimately, this enables commercial activities between unknown agents, humans or things, which is the fundamental basis of an EoT. Widely discussed use cases are, for example, machine-to-machine coordination and sensor data markets, where sensor devices sell data autonomously based on a set of economic rules. While blockchain technology has been widely

cited as a promising contender to pave the way towards an EoT, the existing body of knowledge that substantiates these claims is still in its infancy.

In this dissertation, we explore the role of blockchain technology as a potential infrastructure layer for such a visionary EoT. In particular, we concentrate on two fundamental and related aspects regarding this enquiry. First, we focus on the role of blockchain-based decentralized applications (DApps) for the EoT. The basis for all actions in the decentralized ecosystem of an EoT is trustworthy data. In particular, as things become more autonomous and base complex decisions on data they have acquired from other potentially unknown sources, it is essential that such data be reliable. As many IoT systems handle especially sensitive personal or enterprise information, preserving the privacy during storage and exchange of such data is particularly relevant. Consequently, we investigate how DApps can facilitate the reliable gathering and exchange of IoT data in a privacy-preserving manner and under what circumstances blockchain technology might or might not help to achieve this. For this, we conduct a comprehensive design science study, and develop an instantiation of a sensor data protection system which is tested in a four-month field test with 100 participating cars.

Second, we focus on the underlying system of decentralized finance (DeFi) which supports the development of DApps. While DApps might enable foundational applications for the EoT, only DeFi provides the financial infrastructure to turn things into economic agents. To date, DeFi has led to a number of basic and more complex financial products, of which blockchain-based fundraising has received the largest adoption. In fact, this mechanism is often used to fund the development of DApps and is responsible for a large amount of all investments made towards the whole blockchain ecosystem. Consequently, we examine the mechanism of blockchain-based fundraising and study particularly how investors leverage available information for their decisions. For this, we collect and analyze the financial data as well as information from discussion forums and microblogs for 95 successful fundraises. As such, this dissertation contributes to the emerging stream of literature on blockchain technology and its application in the context of the EoT. In particular, we develop a set of core design guidelines for DApps to facilitate the purposeful implementation of systems for IoT sensor data protection. Furthermore, we enrich these guidelines to form a holistic design theory and explicitly reflect the role of blockchain technology for the realization of individual components. Additionally, we present first analyses on the relevance of different information types for investors that fund the development of DApps. We synthesize our results and deduce relevant learnings for both research and practice and, finally, conclude with the limitations of this work and an outlook on future research.

Zusammenfassung

Über die letzten Jahrzehnte hat sich das Internet der Dinge, aus dem Englischen Internet of Things (IoT), als eine Vision des Computers des 21. Jahrhunderts herausgebildet, in der praktisch jedes Computer-Gerät mit dem Internet verbunden wird. In den letzten Jahren ist diese Vision Realität geworden und Milliarden von vernetzten Dingen werden heute eingesetzt, von intelligenten Heimgeräten über Autos bis hin zu Industriemaschinen. Für die kommenden Jahre wird prognostiziert, dass die Zahl dieser vernetzten Geräte weiterhin mit rasanter Geschwindigkeit wachsen wird. Die gegenwärtige Phase des IoT ist gekennzeichnet durch Dinge, die einfach *mit dem Internet verbunden* sind und *Daten erzeugen*, welche wiederum neue Produkte, Dienstleistungen und Geschäftsmodelle ermöglichen. In letzter Zeit hat der technologische Fortschritt den Schwerpunkt vermehrt auf *miteinander verbundene* Dinge verlagert, die *Daten interpretieren* können, was zu einer erweiterten Vision des IoT geführt hat, die im Englischen als Economy of Things (EoT) bezeichnet wird. Die EoT bezieht sich auf eine Welt, in der IoT-Geräte zu immer intelligenteren Dinge werden, die als autonome Agenten in einem Wirtschaftssystem agieren, anstatt einfach nur Daten zu erfassen, zu verarbeiten und auszutauschen. In diesem Zusammenhang sind die Geräte miteinander verbunden, tauschen Daten direkt untereinander in einem Peer-to-Peer-Netzwerk aus und sind in dem Sinne intelligent, dass sie verfügbare Daten analysieren und diese Informationen nutzen können, um autonome Entscheidungen zu treffen. Ein gutes Beispiel für eine EoT-Anwendung ist die Idee von Robo-Taxis, die autonom umherfahren, Passagiere oder Parkmöglichkeiten ausfindig machen und Zahlungen abwickeln, entscheiden, welche Werkstätten sie für Reparaturen aufsuchen, und Geld sparen, um eine zusätzliche Version von sich selbst zu kaufen, also ein weiteres Robo-Taxi für die eigene Flotte.

Die Vision des EoT wird von technologischen Entwicklungen getrieben, welche die grundlegende Infrastruktur für eine vermehrte Dezentralisierung bilden könnten. Ein breites Spektrum von Forschern und Praktikern argumentiert, dass die Blockchain Technologie ein vielversprechendes Mittel ist, um bestehende Tendenzen der Dezentralisierung zu verstärken und die technologische Grundlage für das EoT zu schaffen. In diesem Zusammenhang sind besonders relevante Eigenschaften von Blockchains die Möglichkeiten, Geld auszutauschen (z.B. über Kryptowährungen) oder Geschäftslogik abzuwickeln (z.B. über intelligente Verträge) und zwar automatisiert und unabhängig von Intermediären. Tatsächlich erlaubt die dezentrale Natur der Blockchain Technologie Interaktionen ohne Vertrauen in die beteiligten Parteien und macht es möglich, eine Blockchain als gemeinsame, sichere Computerplattform zwischen misstrauischen

oder anonymen Agenten zu nutzen. Letztendlich ermöglicht dies kommerzielle Aktivitäten zwischen unbekanntem Agenten, Menschen oder Dingen, was die fundamentale Grundlage einer EoT ist. Breit diskutierte Anwendungsfälle sind z.B. die Koordination von Maschinen untereinander oder Sensordatenmärkte, bei denen Sensoren auf der Grundlage einer Reihe wirtschaftlicher Regeln autonom Daten verkaufen. Während das Potential der Blockchain Technologie für die Infrastruktur einer EoT weithin als gross angesehen wird, sind wissenschaftliche Studien, welche diese Aussagen belegen, noch in den Anfängen.

In dieser Dissertation untersuchen wir die Rolle der Blockchain Technologie als potenzielle Infrastruktur für eine solche visionäre EoT. Insbesondere konzentrieren wir uns dabei auf zwei grundlegende und verwandte Aspekte. Erstens fokussieren wir uns auf die Rolle von Blockchain-basierten dezentralen Anwendungen, im Englischen Decentralized Applications (DApps) genannt, für die EoT. Die Grundlage für alle Handlungen im dezentralisierten Ökosystem einer EoT sind vertrauenswürdige Daten. Insbesondere je autonomer Dinge agieren und je komplexere Entscheidungen sie treffen, basierend auf Daten aus potenziell unbekanntem Quellen, desto wichtiger ist es, dass diese Daten vertrauenswürdig sind. Da viele IoT-Systeme besonders sensible persönliche oder geschäftliche Informationen verarbeiten, ist die Wahrung der Privatsphäre bei der Speicherung und beim Austausch solcher Daten besonders relevant. Daher untersuchen wir, wie DApps die zuverlässige Sammlung und den Austausch von IoT-Daten unter Wahrung der Privatsphäre ermöglichen können und unter welchen Umständen die Blockchain Technologie dabei helfen kann oder nicht. Dazu führen wir eine umfassende wissenschaftliche Design Studie durch und entwickeln ein Sensordaten-Schutzsystem, das in einem viermonatigen Feldtest mit 100 teilnehmenden Fahrzeugen getestet wird.

Zweitens konzentrieren wir uns auf das dezentrale Finanzsystem, aus dem Englischen Decentralized Finance (DeFi), das die Entwicklung von DApps unterstützt. Während DApps grundlegende Anwendungen für die EoT ermöglichen könnten, bietet nur DeFi die finanzielle Infrastruktur, um Dinge zu Wirtschaftsakteuren zu machen. Bis heute hat DeFi zu einer Reihe grundlegender und komplexerer Finanzprodukte geführt, von denen die Blockchain-basierte Unternehmensfinanzierung am meisten benutzt wurde. Dieser Mechanismus wird häufig zur Finanzierung der Entwicklung von DApps eingesetzt und ist für einen Grossteil aller Investitionen verantwortlich, die in das gesamte Blockchain-Ökosystem geflossen sind. Folglich untersuchen wir den Mechanismus der Blockchain-basierten Unternehmensfinanzierung und erforschen insbesondere, wie Investoren verfügbare Informationen für ihre Entscheidungen nutzen. Dazu sammeln und analysieren wir die Finanzdaten sowie Informationen aus online Diskussionsforen und Mikroblogs für 95 erfolgreiche Finanzierungsrunden. Damit leistet diese Dissertation einen Beitrag zur Literatur über die Blockchain Technologie und ihre Anwendung im Kontext des EoT.

Insbesondere entwickeln wir eine Reihe von Designrichtlinien für DApps, um die zielgerichtete Implementierung von Systemen zum Schutz von IoT-Sensordaten zu erleichtern. Darüber hinaus erweitern wir diese Richtlinien zu einer ganzheitlichen Designtheorie und reflektieren explizit die Rolle der Blockchain Technologie für die Realisierung einzelner Systemkomponenten. Zusätzlich präsentieren wir erste Analysen zur Relevanz verschiedener Informationstypen für Investoren, die die Entwicklung von DApps finanzieren. Wir fassen unsere Ergebnisse zusammen und leiten daraus relevante Erkenntnisse für Forschung und Praxis ab und schliessen mit den Einschränkungen dieser Arbeit und einem Ausblick auf die zukünftige Forschung.

Previous Publications

This dissertation contains contributions that have already been published previously as scientific articles in peer-reviewed journals or conference proceedings. Thus, some sections of this dissertation correspond literally to work previously published by me or bear strong similarities. Specifically, the following publications are included in parts, or in an extended version, in this dissertation (*for all publications the CORE 2018 Conference Ranking (ERA 2010 for journals) and the VHB JOURQUAL 3 ranking are included in brackets as the final item*):

Chanson, M., Bogner, A., Bilgeri, D., Fleisch, E., & Wortmann, F. (2019). Blockchain for the IoT: Privacy-Preserving Protection of Sensor Data. *Journal of the Association for Information Systems (JAIS)* (20:9), pp. 1274–1309. (A, A).

Chanson, M., Gjoen, J., Risius, M., & Wortmann, F. (2018). Initial Coin Offerings (ICOs): The Role of Social Media for Organizational Legitimacy and Underpricing. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. San Francisco, CA. (A*, A).

Chanson, M., Martens, N., & Wortmann, F. (2020). The Role of User-Generated Content for Blockchain-Based Decentralized Finance. In *Proceedings of the 28th European Conference on Information Systems (ECIS)*. Marrakesh, Morocco. (A, B).

Moreover, the following publications are part of my doctoral research, but are beyond the scope of this dissertation (*presented as above with CORE 2018 and VHB JOURQUAL 3 rankings*):

Bogner, A., **Chanson, M.**, & Meeuw, A. (2016). A Decentralised Sharing App Running a Smart Contract on the Ethereum Blockchain. In *Proceedings of the 6th International Conference on the Internet of Things*. Stuttgart, Germany. (-, -).

Chanson, M., Bogner, A., Wortmann, F., & Fleisch, E. (2017). Blockchain as a Privacy Enabler: An Odometer Fraud Prevention System. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)* (pp. 13–16). Maui, HI. (A*, -).

Chanson, M., Gahr, B., Dahlinger, A., Ryder, B., & Wortmann, F. (2018). Predicting Driver Stress with Connected Vehicle Driving Data. In *EPFL Applied Machine Learning Days*. Lausanne, Switzerland. (-, -).

Chanson, M., Risius, M., & Wortmann, F. (2018). Initial Coin Offerings (ICOs): An Introduction to the Novel Funding Mechanism Based on Blockchain Technology. In *Proceedings of the 24th Americas Conference on Information Systems (AMCIS)*. New Orleans, LA. (A, D).

Risius, M., French, R., & **Chanson, M.** (2020). Redefining the Customer Service Relationship through Blockchain. *Information & Management* (submitted). (A*, B).

Acknowledgements

This dissertation is the result of a wonderful journey of discovery and learning, immensely stimulating intellectually, intense at times, and always at the forefront of newest technological developments. I am extremely grateful for all the experiences I made throughout that journey, everything I learned from my supervisors, colleagues and friends during that time and would not want to miss any of it.

Above all, I thank my advisor Prof. Dr. Elgar Fleisch for putting his faith in me and giving me the opportunity to conduct my doctoral research at his chair at ETH Zurich, namely in the Bosch IoT Lab. His guidance in matters of research and beyond has been invaluable throughout this dissertation. Furthermore, I am especially grateful for the open and collaborative research environment he and his team established at the chair, which allowed me to focus relentlessly on my research, cooperate effectively with my peers and industry partners and experience many enriching moments beyond the specific target of excellent research output. Additionally, I thank my co-supervisor Prof. Dr. Frédéric Thiesse for supporting my dissertation in its final stages.

Furthermore, my sincere appreciation goes to Ass.-Prof. Dr. Felix Wortmann who co-directed the Bosch IoT Lab during my dissertation. Throughout this time, he proved to be a motivating and caring leader who undertook anything to facilitate successful research at the lab. Especially his scientific advice and the publishing directions he provided contributed meaningfully to my work. At this point, I also thank Timo Gessmann and Prof. Dr. Markus Weinberger who shared the direction of the Bosch IoT Lab with Felix Wortmann through periods of my PhD. In particular, Timo Gessmann and his tireless and attentive approach in managing contacts at major industry partners assisted the impact and success of this dissertation. Furthermore, I thank Elisabeth Vetsch-Keller from the University of St. Gallen, and Monica Heinz of ETH Zurich, for their outstanding assistance whenever it was needed.

In my dissertation I had the pleasure visit Columbia University in New York as a research scholar under the supervision of Prof. Dr. Miklos Sarvary, who I thank for welcoming me at Columbia University, for his interest in my research, and his support in obtaining any assistance needed.

Moreover, I thank Prof. Dr. Marten Risius for the excellent collaboration which resulted in a number of co-authored publications, his insightful career advice and enlightening discussions at various international conferences of the Association for Information Systems.

Finally, I thank my colleagues and friends at the Bosch IoT Lab and the chairs of Elgar Fleisch who created a welcoming and supportive atmosphere through all these years of research. In particular, I thank my co-authors from the lab Andreas Bogner, Arne Meeuw, Dominik Bilgeri, Bernhard Gahr, Andreas Dahlinger and Benjamin Ryder. Beyond that, I am grateful for the wonderful time I spent with colleagues and friends from the Bosch IoT Lab, including Thomas von Bomhard, Dominic Wörner, Shu Liu, Kevin Koch, and Claudio Lamprecht.

Table of Contents

Abstract	ii
Zusammenfassung	iv
Previous Publications	vii
Acknowledgements	viii
1 Introduction	1
1.1 Motivation.....	1
1.2 Research Objectives and Approach.....	5
1.3 Dissertation Outline.....	8
2 Foundations and Related Work	9
2.1 Internet of Things	9
2.2 Blockchain Technology.....	15
2.3 Economy of Things	29
3 DApps: Privacy-Preserving Protection of Sensor Data	33
3.1 Introduction	33
3.2 Conceptual and Theoretical Background	35
3.3 Research Study.....	41
3.4 Discussion	68
3.5 Conclusion.....	75

4	DeFi: Information from Social Media in Blockchain-Based Fundraising	79
4.1	Introduction	80
4.2	Conceptual and Theoretical Background	83
4.3	Research Study	90
4.4	Discussion	98
4.5	Conclusion.....	100
5	Discussion and Conclusion	105
5.1	Summary and Key Findings.....	105
5.2	Limitations and Future Research.....	114
5.3	Conclusion.....	116
	References	119

Table of Contents (detailed)

Abstract	ii
Zusammenfassung	iv
Previous Publications	vii
Acknowledgements	viii
1 Introduction	1
1.1 Motivation	1
1.2 Research Objectives and Approach.....	5
1.2.1 Decentralized Applications	6
1.2.2 Decentralized Finance.....	7
1.3 Dissertation Outline.....	8
2 Foundations and Related Work	9
2.1 Internet of Things	9
2.1.1 Existing Perspectives	9
2.1.2 Progress and Development.....	11
2.1.3 Technology Stack.....	12
2.2 Blockchain Technology.....	15
2.2.1 Blockchain Fundamentals	15
2.2.2 Decentralized Applications	23
2.2.3 Decentralized Finance.....	26
2.3 Economy of Things	29

2.3.1	Existing Perspectives	29
2.3.2	DApps and DeFi in the context of the IoT	30
3	DApps: Privacy-Preserving Protection of Sensor Data	33
3.1	Introduction	33
3.2	Conceptual and Theoretical Background	35
3.2.1	Internet of Things and Sensor Data.....	35
3.2.2	Security and Privacy in the Internet of Things.....	36
3.2.3	Existing Research on SDPSs and their Limitations	37
3.2.4	Blockchain Technology for SDPSs.....	39
3.3	Research Study.....	41
3.3.1	Methodology	41
3.3.2	Designing an IoT Sensor Data Protection System	45
3.3.3	Iterative Development of the Prototype	54
3.3.4	Ex-Post Evaluation of the Design	61
3.4	Discussion	68
3.4.1	SDPS Design Theory	68
3.4.2	Design Implications	71
3.4.3	Theoretical and Practical Contributions.....	74
3.5	Conclusion.....	75
4	DeFi: Information from Social Media in Blockchain-Based Fundraising	79
4.1	Introduction.....	80
4.2	Conceptual and Theoretical Background	83
4.2.1	ICOs and Underpricing	83
4.2.2	Legitimacy in ICOs	86
4.2.3	Social Media Driving Legitimacy in ICOs	87
4.3	Research Study.....	90
4.3.1	Sample and Data Analysis	90
4.3.2	Measures	91

4.3.3	Results.....	95
4.4	Discussion.....	98
4.5	Conclusion.....	100
5	Discussion and Conclusion	105
5.1	Summary and Key Findings.....	105
5.1.1	Decentralized Applications.....	107
5.1.2	Decentralized Finance.....	112
5.2	Limitations and Future Research.....	114
5.3	Conclusion.....	116
	References	119

List of Figures

Figure 1. IoT technology stack (adapted from Wortmann & Flüchter, 2015)	13
Figure 2. Core technological layers of a blockchain system.....	16
Figure 3 . Simplified structure of the bitcoin blockchain.....	17
Figure 4 . Merkle tree of transactions	18
Figure 5. Design cycles	43
Figure 6. Design requirements, principles, and features	51
Figure 7. Artifact architecture	53
Figure 8. A detailed view of the certification process.....	54
Figure 9. Prototype architecture	56
Figure 10. Sample Ethereum transaction	56
Figure 11. Transaction of Figure 10 on etherscan.io.....	57
Figure 12. First implementation of the verification process	58
Figure 13. The main screens of the smartphone application.....	60
Figure 14. ICO process	84
Figure 15. Research model for social media effects on organizational legitimacy.....	88
Figure 16. Summary of findings	98

List of Tables

Table 1 . Blockchain types	22
Table 2: General SDPS challenges and design requirements.....	47
Table 3: Overview and summary of the prototype evaluation	61
Table 4: Ex-post evaluation interview participants	62
Table 5: Components of an SDPS design theory	69
Table 6: Blockchain-based SDPS usage implications.....	72
Table 7: Design implications for blockchain-based SDPS	73
Table 8. Main coin utilities	85
Table 9. Descriptive statistics for complete sample	94
Table 10. OLS robust regressions	96
Table 11. Confidence interval bootstrapping	97

List of Abbreviations

DApp	Decentralized Application
DeFi	Decentralized Finance
DEX	Decentralized Exchange
DDoS	Distributed Denial-of-Service
DSR	Design Science Research
EoT	Economy of Things
ICO	Initial Coin Offering
IoT	Internet of Things
IPO	Initial Public Offering
IS	Information Systems
IT	Information Technology
MGC	Marketer-Generated Content
MIT	Massachusetts Institute of Technology
OLS	Ordinary Least Squares
PBFT	Practical Byzantine Fault Tolerance
PoS	Proof of Stake
PoW	Proof of Work
RFID	Radio-Frequency Identification
SPV	Simplified Payment Verification
UGC	User-Generated Content
VC	Venture Capital

Chapter 1

Introduction¹

“We always overestimate the change that will occur in the next two years and underestimate the change that will occur in the next ten. Don't let yourself be lulled into inaction.”

Bill Gates

1.1 Motivation

In the last decades, the Internet of Things (IoT) has evolved from a pure vision to reach reality. Originally, the idea of the IoT is rooted in the early 1990's and Mark Weiser's considerations of what the computer of the 21st century might look like (Mattern & Floerkemeier, 2010; Weiser, 1991). He envisioned that, eventually, computing technology would become ubiquitous, part of our environment and merge with it to become undistinguishable of the background (Weiser, 1991). Although the developments towards true ubiquitous computing are still at an early stage, an initial form of the IoT has in truth materialized since these first conceptualizations, as billions of things have been connected to the Internet to date (Gartner, 2018; Oberländer, Röglinger, Rosemann, & Kees, 2018). Projections assume that this number will further increase quickly and reach even higher order of magnitudes in the near future (Gartner, 2018; Oberländer et al., 2018). The types of devices that are being connected seem

¹ Parts of this chapter, which are not further demarcated in the text, were initially published in the context of the following academic publications: Chanson, Gjoen, et al. (2018), Chanson et al. (2019) and Chanson et al. (2020).

limitless and range from rather simple objects such as lightbulbs or video cameras to quite complex systems such as cars, industrial robots or medical devices (Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015; Oks, Fritzsche, & Möslin, 2017; Porter & Heppelmann, 2014). In fact, the range of objects encompassed under the IoT has widened to the point where it refers to “a vision that virtually any physical object can be connected to the Internet” (Bilgeri, Wortmann, & Fleisch, 2017, p. 2). As such, the IoT is often declared to be among the most relevant technological developments of current times (Atzori, Iera, & Morabito, 2010; Jeschke, Brecher, Meisen, Özdemir, & Eschert, 2017; Oberländer et al., 2018). Accordingly, discussions regarding the topic have been prolific since a long time, attracting both intense research efforts and attention from practice. The plethora of connected devices which form the IoT has vastly increased the amount of data gathered about our world and the human beings that inhabit it (Gubbi, Buyya, Marusic, & Palaniswami, 2013; Porter & Heppelmann, 2014; Sicari, Rizzardi, Grieco, & Coen-Porisini, 2015). In fact, the analysis of this data is the basis for many of the most prolific use cases of the IoT such as predictive maintenance, applications in supply chain management or medical diagnosis and monitoring (Iansiti & Lakhani, 2014; Magargle et al., 2017; Uckelmann, Harrison, & Michahelles, 2011). In summary, this stage of the IoT is characterized by the focus on things which are simply *connected* to the Internet and *generate data* which, in turn, enables new products, services and business models.

Recently, technological progress facilitating *interconnected* things which can *interpret data* has spurred interest in an enhanced vision for the IoT, which is sometimes referred to as the Economy of Things (EoT) (Beck, Stenum Czepluch, Lollike, & Malone, 2016; Davidsen, Gajek, Kruse, & Thomsen, 2019; Elsdén et al., 2018; Kouzinopoulos et al., 2018; Sousa, Antunes, & Martins, 2018). The EoT refers to a world in which IoT devices become increasingly intelligent things that are acting as autonomous agents in an economic system, instead of simply sensing, processing and exchanging data. In this context, devices are interconnected, exchanging data directly between each other in a peer-to-peer network, and intelligent in the sense that they can analyze available data and leverage this information to make decisions autonomously. Such an enhancement of devices to economic actors gives rise to completely new business models (Antonopoulos, 2016; D. Tapscott & Tapscott, 2016b). An excellent showcase for the EoT is the vision of robo-taxis, which has been extensively promoted by Tesla’s founder Elon Musk. The basic idea is that cars with full autonomous driving capability can be ordered to work as taxis by their owners. In practice, owners will be able to add their cars to a network similar to Uber, which bundles demands and distributes requests to different vehicles, and the cars execute the necessary driving autonomously. As such, idle time can be avoided whenever the owner of a car does not need the vehicle. As Elon Musk puts it, “people [will be able to] allow their car to earn money for

them as part of the Tesla shared autonomy fleet. [...] Tesla cars being made today will be able to do that for you.” (Musk, 2019). The longer a car is expected to act as a robo-taxi without human engagement, the more decisions beyond driving it has to make autonomously. When is the best time to take a break and recharge the batteries? If the car needs a repair, which workshop should it go to? How can it pay for the repair? While some of these cases should happen rarely enough that human intervention might be acceptable or no real-time action is necessary, the vision of the EoT is that a thing, such as a car, is able to make all relevant decisions without any human interference based on pre-defined high-level rules. For example, a car might choose a workshop based on the price and waiting time for a repair, which it can directly request from a workshop API. More interestingly, it could leverage outcomes of previous repairs that other cars have shared (similar to a review) after their maintenance at that workshop. As this data becomes valuable, cars might start to trade such information between each other for micropayments. These examples also illustrate special cases of some of the most widely discussed use cases of the EoT, namely machine-to-machine coordination and sensor data markets forming decentralized electronic marketplaces. The more independent these things become in their actions, we can even think of them as individual actors, that decide in their own will and are not necessarily controlled by any individual owner.

The vision of the EoT is nurtured by technological developments which could build the infrastructural foundation for broader decentralization and disintermediation. Specifically, the Internet has fueled first disintermediation, particularly in its early days, and gave birth to a large set of firms that enact directly with their customers (Antonopoulos, 2016; Beck, Müller-Bloch, & King, 2018; Risius & Spohrer, 2017; Swan, 2015). Similarly, a broad range of research scholars and practitioners alike argue that blockchain technology has the potential to strengthen tendencies of disintermediation and shape a future where things become economic actors and interact directly with each other (Beck, Müller-Bloch, & King, 2018; D. Tapscott & Tapscott, 2016a). For example, the possibility to transact value (e.g., via cryptocurrencies) or incorporate business logic (e.g., via smart contracts) in an automated manner independent of intermediaries are core technological building blocks to develop an EoT and enabled peculiarly by blockchain technology. In its general form, a blockchain is a “fully distributed system for cryptographically capturing and storing a consistent, immutable, linear event log of transactions between networked actors” (Risius & Spohrer, 2017, p. 386). Many blockchains also allow the implementation of smart contracts, which can be conceptualized as computer protocols that execute specific logic operations and transactions automatically without the participation of a third party (Beck, Müller-Bloch, & King, 2018; Chanson, Gjoen, Risius, & Wortmann, 2018). The networked actors participating in a blockchain system are also called nodes and share the blockchain as a database

and single source of truth (Egelund-Müller, Elsmann, Henglein, & Ross, 2017). The nodes reach agreement on the state of this event log through a consensus mechanism which is based on cryptographic and game theoretical fundamentals (Buterin, 2013; Nakamoto, 2008; Risius & Spohrer, 2017). In particular, the consensus mechanism is robust towards a certain share of faulty or malicious nodes and, thus, allows nodes whose identities and intentions may be unknown to collaborate, while ensuring the integrity of the blockchain (Buterin, 2013; Lamport, Shostak, & Pease, 1982; Nakamoto, 2008; Risius & Spohrer, 2017). Ultimately, this method of reaching consensus enables the system to be fully distributed. Consequently, using a blockchain as a shared, secure computing platform enables commercial activities between distrustful or anonymous agents (Nærland, Müller-Bloch, Beck, & Palmund, 2017). Widely discussed use cases are, for example, machine-to-machine coordination and sensor data markets, where sensor devices sell data autonomously based on a set of economic rules (Christidis & Devetsikiotis, 2016; Noyen, Volland, Wörner, & Fleisch, 2014; Subramanian, 2018; Wörner & von Bomhard, 2014; Zhang & Wen, 2017). Ultimately, this results in things that become increasingly autonomous at an individual or very low level, thus leading to a more decentralized economy (Antonopoulos, 2016; Beck, Müller-Bloch, & King, 2018; Swan, 2015). Consequently, a number of different blockchain-based IoT systems are currently under development, for example in the area of supply chains, energy markets and mobility (Curtis, 2015; Mengelkamp et al., 2018; Modum, 2018).

A cornerstone of the EoT, or generally a more decentralized economy, are decentralized applications (DApps). A DApp is a “web application that is built on top of open, decentralized, peer-to-peer infrastructure services” (Antonopoulos & Wood, 2018, p. 8). As such, DApps are typically built on blockchains, leverage smart contracts and other blockchain-enabled protocols such as decentralized storage or hosting services (Antonopoulos & Wood, 2018; Dannen, 2017; Raval, 2016). DApps have the potential to be more transparent and resilient than traditional software because they can inherit the core properties of the blockchain they are built on (Chanson, Bogner, Bilgeri, Fleisch, & Wortmann, 2019; Raval, 2016). A large range of use cases could be implemented by DApps, including products and services as different as (car) sharing applications, privacy-preserving data analysis platforms, stock exchanges and other marketplaces, funding platforms or platforms to manage intellectual property, for example of music songs (Beck, Müller-Bloch, & King, 2018; Bogner, Chanson, & Meeuw, 2016). As decentralization and the autonomy of devices are an essential paradigm of the EoT, DApps form the main archetype to build corresponding applications.

In turn, DApps rely heavily on various aspects of decentralized finance (DeFi). Essentially, DeFi is the decentralized version of a financial system and comprises, in particular, infrastructure related to payment and funding mechanisms (Aave, 2020; Adams, 2019; Chanson, Martens, &

Wortmann, 2020; Y. Chen & Bellavitis, 2020). In particular, DeFi is essential in enabling new innovative business models that are expected to emerge in the EoT and which rely on payment and financing infrastructure (Antonopoulos, 2016; Y. Chen & Bellavitis, 2020). While some DeFi services are inherently part of a blockchain, for example payments, more complex services, such as fundraising, have to be implemented through specific dedicated smart contracts (Chanson, Risius, & Wortmann, 2018; Fridgen, Regner, Schweizer, & Urbach, 2018). Depending on the complexity of a financial service and whether additional interfaces beyond a central smart contract are offered, such a service could be anything between a simple add-on to a blockchain (in the form of a smart contract) or a full-fledged DApp. Currently, DeFi offers a number of different types of equity and debt financing (Aave, 2020; Y. Chen & Bellavitis, 2020; Fridgen et al., 2018; MakerDAO, 2020). DApps mainly leverage DeFi to fund their development and to incentivize actors within the economic system of a DApp.

While the prospects of blockchain technology and, related, the EoT are promising and approached with enthusiasm from the research community, it is important to note that the corresponding academic research is still in its infancy (Avital, Beck, King, Rossi, & Teigland, 2016; Beck, Avital, Rossi, & Thatcher, 2017; Beck et al., 2016; Beck & Müller-Bloch, 2017; Lindman, Rossi, & Tuunainen, 2017). In particular, core challenges of blockchain technology, such as privacy, scalability, and potentially prohibitive transaction costs, remain to be addressed (Beck et al., 2016; Notheisen, Cholewa, & Shanmugam, 2017; Risius & Spohrer, 2017).

1.2 Research Objectives and Approach

Researchers and practitioners alike attribute great potential to blockchain technology as the foundational infrastructure enabling the transition to an EoT with new and sustainable business models (Beck et al., 2016; Elsdén et al., 2018; Kouzinopoulos et al., 2018; Sousa et al., 2018). However, as blockchain technology and the concept of the EoT are nascent, there is only little scientific evidence detailing these reports based on academic research. In an EoT, things gather and exchange data and base their decisions on that information. Hence, it is essential that such data be trustworthy and manipulations as limited as possible. Additionally, the cooperation between things in a complex ecosystem will not simply occur magically. Financial infrastructure is required to formalize cooperation incentives, in the short term such that things can compensate each other through payments (e.g., to compensate for data transmission), and in the long term such that innovation projects can be financed through investments (e.g., to create new DApps or large shared infrastructure). Especially the considerations of actors leading to long-term

investments is, while highly influential, still poorly understood. This leads to the following two research topics that this dissertation aims to investigate: First, how to build DApps to protect the gathering and exchange of IoT sensor data. Second, understanding novel DeFi financing mechanisms, in particular decision factors for investors. In the following sections, these two topics are introduced in more detail, leading to our overarching research questions and the according approaches taken to address these research questions.

1.2.1 Decentralized Applications

The vision of an EoT is fundamentally based on the availability of trustworthy data. As things become increasingly autonomous and base complex decisions on data they have acquired automatically from other potentially unknown sources, it is essential that such data be reliable – or at least it can be verified for critical decisions if the data is indeed reliable. However, with the rapid deployment of connected devices around the world, which form increasingly complex large-scale IoT systems, the number of attack vectors has steadily increased and these systems have become more and more attractive manipulation targets (Lee, Cho, & Lim, 2018; Newell & Marabelli, 2015; Sicari et al., 2015; Weber, 2010). As many IoT systems handle especially sensitive data from persons or companies, a privacy-preserving storage and exchange of such data is particularly relevant (Lowry, Dinev, & Willison, 2017; Porter & Heppelmann, 2015; Sicari et al., 2015). Consequently, considerable research effort has been put into creating applications which lead to more secure and privacy-protecting IoT systems (Atzori et al., 2010; Koliass, Kambourakis, Stavrou, & Voas, 2017; Lowry et al., 2017; Ronen, Shamir, Weingarten, & O’Flynn, 2017). However, there is a lack of actionable research that guides practitioners in their quest towards more secure IoT applications (Bélanger & Crossler, 2011; Pavlou, 2011; Smith, Dinev, & Xu, 2011). Additionally, while the potential of blockchain technology for such systems is deemed huge, it has not yet been reflected thoroughly in scientific studies (Glaser, 2017; Hyvärinen, Risius, & Friis, 2017; Nærland et al., 2017). In particular, research on DApps in the IoT is only just emerging (Bogner et al., 2016; Chanson et al., 2019). Consequently, we aim to shed light on how DApps could enable the reliable gathering and exchange of data and to what extent blockchain technology does or does not help to facilitate this. Against this background, this dissertation aims to answer the following overarching research question:

RQ 1: What design theory should guide the development of DApps that are able to protect IoT sensor data in a privacy-preserving manner?

To address this research question, we adapt the perspective of design science research (DSR) and set out to develop a design theory (Gregor & Jones, 2007; March & Smith, 1995). Within the

Information Systems (IS) community, the development of design knowledge, be it in the form of design theories, principles, or guidelines, is of high significance for both research and practice (Baskerville, 2008; Hevner, March, Park, & Ram, 2004; Winter, 2008) and continues to attract a great deal of interest (Baskerville, Kaul, & Storey, 2015; Gregor & Hevner, 2013; Rai, 2017). In this dissertation, we derive an artifact that consists of a set of interrelated design requirements, design principles, and design features. We demonstrate and refine our artifact on the basis of an instantiation that aims to prevent the fraudulent manipulation of car mileage data. Additionally, we provide an evaluation of the artifact and present our results in the form of a design theory.

1.2.2 Decentralized Finance

While DApps might enable foundational applications relevant for the EoT, such as the secure collection, storage and exchange of data, only DeFi provides the financial infrastructure necessary to integrate complex monetary incentive systems into these applications, finance innovation in decentralized ecosystems and, as such, ultimately turn things from *technical devices* into *economic agents*. In particular, DeFi enables things to exchange payments and finance their activities, which is vital for economic agents (Antonopoulos, 2016; Y. Chen & Bellavitis, 2020). While things could leverage DeFi in their interactions with each other and DApps, DApps in turn rely heavily on DeFi directly: Both as a financing mechanism for their development, and as a toolbox to incentivize different stakeholders of a DApp ecosystem (Chanson, Gjoen, et al., 2018; Y. Chen & Bellavitis, 2020). While DeFi has produced a number of important services, the largest adoption from practice has so far been geared towards equity-like blockchain-based fundraising of DApps. In fact, this mechanism is responsible for a large amount of all investments made towards the whole blockchain ecosystem (Coindesk, 2019). Traditionally, IS research on financing mechanisms has been directed in large parts towards exploring the role of information (R. Aggarwal & Singh, 2013; B. N. Greenwood & Gopal, 2016). These efforts are rooted in decades of research considering the impact of the availability and quality of information on markets (Akerlof, 1970). Recently, information from social media has become a focus of researchers (Lukyanenko et al., 2017; Mai, Shan, Bai, Wang, & Chiang, 2018). While in traditional financial markets, the relevance of social media is rather limited due to established information sources (e.g., audited financial reports or specialized news portals such as Bloomberg or Reuters) it is in fact an essential source of information in the context of decentralized blockchain-based fundraising (Chanson, Gjoen, et al., 2018; Mai et al., 2018). To comprehend the developments towards an EoT, it is essential to understand how investors leverage information to decide which innovation projects or DApps to support. Against this background, this dissertation aims to answer the following overarching research question:

RQ 2: How does information from social media relate to success in blockchain-based fundraising?

To address this research question, we adapt a similar approach as previous studies on the impact of information on markets and focus on empirical investigations (R. Aggarwal & Singh, 2013; B. N. Greenwood & Gopal, 2016; Mai, Bai, Shan, Wang, & Chiang, 2015). In particular, we analyze our data through a number of ordinary least squares (OLS) regression models. Such analyses can be of high relevance for both research and practice and continue to attract vivid interest among IS researchers, in particular regarding the investigation of social media data (R. Aggarwal, Gopal, Gupta, & Singh, 2012; Ghose, 2009; Li, van Dalen, & van Rees, 2018; Nishant, Teo, & Goh, 2017). In this dissertation, we gather a data set of 95 blockchain-based fundraises, related social media information and financial results. We include Twitter and various discussion forums in our investigations. We draw on theory of organizational legitimacy and apply it to the context of social media and blockchain-based fundraising (Dowling & Pfeffer, 1975; Lundmark, Oh, & Verhaal, 2017; Suchman, 1995; Zimmerman & Zeitz, 2002). We analyze our data set through a number of OLS regression models and present our results on the impact of different types of social media information on blockchain-based fundraising.

1.3 Dissertation Outline

The remainder of this dissertation is structured as follows. In the following chapter, the foundations and related work are presented, introducing major concepts which are subject of the investigations of this dissertation. In particular, key notions such as the *Internet of Things*, *blockchain technology*, and the *Economy of Things* are explained in detail. Following this, the first research question is investigated, focusing on the development of a design theory for DApps that protect IoT sensor data in a privacy preserving manner. Subsequently, the second research question is explored in the penultimate chapter, focusing on the role of information from decentralized origins in the fundraising success of DApps with DeFi services. Finally, the dissertation concludes with the last chapter, which reflects the key findings and their implications for research and practice, highlights related limitations of our work and presents promising avenues for future research to advance the topics covered in this dissertation.

Chapter 2

Foundations and Related Work

As a foundation for this dissertation, it is essential to reflect the fundamentals of the existing body of knowledge regarding the following three core topics of interest: First, the IoT, second, blockchain technology and, third, how these technologies are related to the emerging concept of the EoT. While this chapter introduces the existing fundamental research on each of these subjects, more specific discussions of the research background relevant for each Chapters 3, and 4, respectively, can be found in the according chapters.

2.1 Internet of Things

2.1.1 Existing Perspectives

While the IoT has long become an extensively discussed topic, attracting both intense research efforts and attention from practice, a clear and widely accepted definition of the term has yet to emerge (Atzori et al., 2010; Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015; Oberländer et al., 2018). Historically, the IoT emerged from the concept of ubiquitous computing, which was introduced by Mark Weiser in the beginning of the 1990s (Mattern & Floerkemeier, 2010; Weiser, 1991). The specific term IoT was coined in the context of research on networked radio-frequency identification (RFID) infrastructure in the late 1990s at the Auto-ID Labs at the Massachusetts Institute of Technology (MIT) (Wortmann & Flüchter, 2015). Specifically, the notion is often attributed to Kevin Ashton, co-founder and former executive director of the Auto ID Labs at MIT (Mattern & Floerkemeier, 2010; Suresh, Daniel, Parthasarathy, & Aswathy, 2014). Since the beginnings of the IoT, the extent of objects encompassed under this notion has grown to a wide range of devices other than RFID sensors, to the point where it refers to “a vision that virtually any physical object can be connected to the Internet” (Bilgeri et al., 2017, p. 2). Semantically, the

IoT can be interpreted as “a world-wide network of interconnected objects uniquely addressable, based on standard communication protocols” (Atzori et al., 2010, p. 2788). In turn, other approaches are more specific, only accounting for “the connectivity of physical objects equipped with sensors and actuators to the Internet via data communication technology” (Oberländer et al., 2018, p. 488). All these definitions feature the two core constituents of any intention to isolate the term IoT: A network related concept of connections and communication (i.e., the Internet aspect) as well as an object related concept of potentially physical dimensions (i.e., the Things aspect). The reason for the variance in definitions of the IoT is fundamentally rooted in different interpretations of these two concepts related to network and objects. On the one hand, scholars have reached no consensus about the requirements for specific *networking* technologies in the IoT, on the other hand, there are ongoing discussions about what type of *objects* should be included with reference to the aspect of things. Furthermore, the IoT is intimately associated with other technical notions such as ubiquitous and pervasive computing and, while there exist specific differences between these terms, they are not always used unambiguously (Gubbi et al., 2013; Oberländer et al., 2018; Uckelmann et al., 2011).

With regard to the concept of network, the notion of Internet can be interpreted narrowly as the technical ability to communicate with the Internet through TCP/IP or, in a much wider sense, as a general ability of communication over a variety of protocols and physical channels. Many academic definitions of IoT adapt rather wide-ranging perspectives which refrain from specifying individual protocols or networks in detail (Oberländer et al., 2018). For example, academic research often includes all wireless or wired communication technology as equally noteworthy in the regard of the network concept (O’Leary, 2013; Oberländer et al., 2018). Typical terms used to achieve such a comprehensive signification are, for example, “network technology” (Wortmann & Flüchter, 2015, p. 221) “wired and wireless networks” (Chui, Löffler, & Roberts, 2010, p. 1), or “telecommunication” (Atzori et al., 2010, p. 2787). However, there are also scholars that uniquely or primarily refer to wireless communication in the context of the IoT (Boos, Guenter, Grote, & Kinder, 2013; Mattern & Floerkemeier, 2010).

Concerning the concept of objects in turn, the notion of Thing can be interpreted rather widely, too. The only evident condition is that an object needs to comprise communication capacities in order to adhere to the concept of network. While the original conception of IoT started with RFID tags the range of objects considered quickly expanded to all sorts of sensors and actuators and more complex devices (Atzori et al., 2010; Suresh et al., 2014; Wortmann & Flüchter, 2015). Some scholars even argue that not only physical objects but also virtual objects can be considered things in the context of the IoT (Sundmaeker, Guillemin, Friess, & Woelfflé, 2010). In this regard, in particular the digital representation of physical objects, often called digital

twins, obtain a key role (Magargle et al., 2017; Sundmaeker et al., 2010; Tao et al., 2018). Such digital twins, which can be realized through cyber-physical systems, represent a prerequisite to leverage the rich set of data the IoT is expected to generate (Schilberg, Hoffmann, Schmitz, & Meisen, 2017). They can make detailed information of the physical objects they represent readily available anywhere in the world to be monitored or analyzed, for example for applications in supply chain management or predictive maintenance (Magargle et al., 2017; Schilberg et al., 2017; Uckelmann et al., 2011). However, most scholars restrict the type of objects for the IoT along two dimensions: First, regarding the importance of the aspect of connectivity for the object, and, second, concerning its ability in terms of computational power and smart decisions. With reference to the former dimension, objects that inherently depend on networking abilities for their primary functions are sometimes omitted from IoT objects, which leads for example to the exclusion of smartphones, tablets and personal computers (Mattern & Floerkemeier, 2010; Oberländer et al., 2018). In contrast, other researchers explicitly consider such objects, for example mobile phones (Atzori et al., 2010). Regarding the latter dimension, in some cases researchers focus primarily on everyday objects that that were upgraded with IoT technology (Mattern & Floerkemeier, 2010; Oberländer et al., 2018), whereas in other cases everyday objects such as household appliances are excluded (Atzori et al., 2010).

Reflecting on the preceding elaborations, we adapt the notion of the IoT that includes both physical and virtual objects that are connected to a worldwide network. In particular, no restrictions regarding the physical or virtual type of the connections apply, meaning that both wired and wireless connections through different types of protocols are included. In line with Mattern and Floerkemeier (2010) we focus on connected everyday objects which exist independently of communication technology. As such, we exclude smartphones, tablets and personal computers from our considerations.

2.1.2 Progress and Development

The development and according increasing importance of the IoT is rooted amongst others in the far-reaching phenomenon of Digitization (Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015; Mattern & Floerkemeier, 2010; Oks et al., 2017). Digitization is a broad term which designates the shift of information processing methods from analog to digital or, more generally, both the growing application of digital technologies as well as the accordingly arising effects of sociotechnical nature (Loebbecke & Picot, 2015; Nambisan, Lyytinen, Majchrzak, & Song, 2017; Negroponte, 1995). Digitization has already lead to a far-reaching expansion of the type and amount of information that can be gathered and analyzed today, in particular regarding the digital trace of humans in their everyday life, and this development is widely estimated to further

intensify (McAfee, Brynjolfsson, Davenport, Patil, & Barton, 2012; Newell & Marabelli, 2015). Through this increase in available data, the phenomenon of Digitization has fueled major trends such as Big Data, Artificial Intelligence and also the IoT (Loebbecke & Picot, 2015; Oks et al., 2017). In particular, Digitization has prompted an increase in connected devices, such as cars, industrial robots or home appliances, which are deployed all over the world (Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015; Oks et al., 2017; Porter & Heppelmann, 2014). Specifically, forecasts predict that the number of things connected to the Internet will grow to around 25 billion by 2021 (Gartner, 2018). This envisaged growth in the number of connected devices has also a strong positive effect on the related business potential. Namely, the combined markets of the IoT, including hardware, software, systems integration and data and telecom services, are predicted to reach over USD 500 billions in 2021 (Bosche, Crawford, Jackson, Schallehn, & Schorling, 2018). In conclusion, the number of connected, intelligent devices that will exist beyond the traditional separation of the physical and digital worlds is growing quickly, fostered among others by Digitization. It is this merger of two worlds, digital and physical, which is labelled as the IoT and which has received substantial attention in research and among practitioners in the last years (Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015; Oks et al., 2017).

2.1.3 Technology Stack

The sensory and computational abilities of IoT devices combined with their connectivity renders them into in so-called smart products and allows for new possibilities in terms of business models. Smart products refer to IoT solutions, that consist of smart objects and additional services which are offered on top of these objects. In turn, smart objects are objects with enhanced capabilities, such as the ability to sense, communicate and or interact (Noyen et al., 2014; Oberländer et al., 2018). While an exact definition of the capabilities of a smart object is lacking, it is clear that the increased potential of smart products mirrors itself in the related business potential (Iansiti & Lakhani, 2014; Porter & Heppelmann, 2014). However, creating sustainable business models in this new paradigm of the IoT can be challenging, as IoT solutions do not only consist in either a physical product or a digital service, but usually combine both aspects into one solution and, thus, a multifaceted effort by the providing company has to be compensated (Porter & Heppelmann, 2015; Wortmann, Bilgeri, Weinberger, & Fleisch, 2017). This is best illustrated by discussing a detailed view of the IoT technology stack, meaning the different technical and economical layers IoT solutions consist in (Porter & Heppelmann, 2015). Disentangling these layers for a given product is often insightful because smart products consist in numerous hardware and software components which are intertwined and at the same time fulfill distinct types of needs. In a rather business oriented approach, Fleisch, Weinberger and Wortmann (2015)

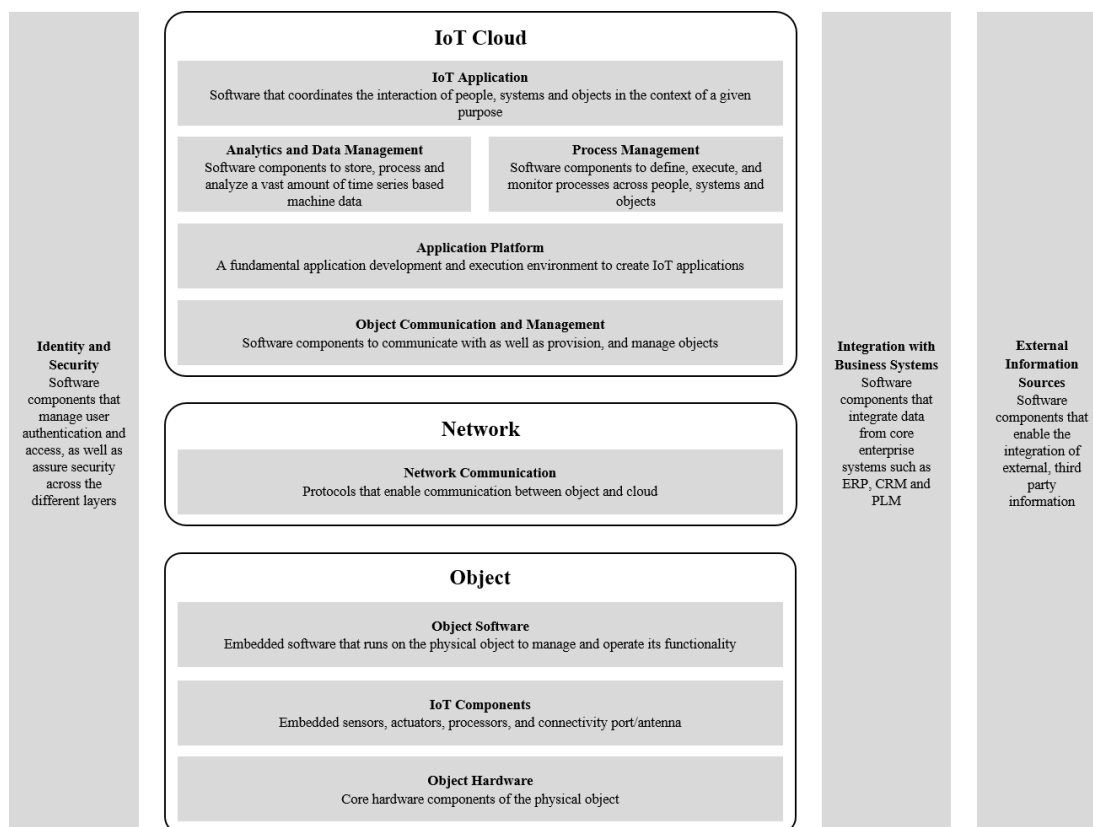


Figure 1. IoT technology stack (adapted from Wortmann & Flüchter, 2015)

distinguish five different technical or value layers: Layer 1 lies the base of the IoT technology stack and consists of physical objects, which are enhanced through sensors and actuators in Layer 2 and reach global connectivity through their access to the Internet in Layer 3. The data generated and diffused by the layers below can then be analyzed in Layer 4 to provide digital services in Layer 5, such as applications in supply chain management or predictive maintenance (Magargle et al., 2017; Schilberg et al., 2017).

Complementing the rather business oriented approach to distinguish different value layers of IoT products, one can also adapt a more technological viewpoint to identify the different layers in more detail (Atzori et al., 2010; Porter & Heppelmann, 2014). In particular, Wortmann and Flüchter (2015) have detailed a multi-faceted outline of the IoT technology stack which is originally based on the work of Porter and Heppelmann (2014). This technology stack is shown, in a variation adapted to this dissertation, in Figure 1 (Porter & Heppelmann, 2014; Wortmann & Flüchter, 2015). Accordingly, one can distinguish three core layers, namely the object layer, the network layer and the IoT cloud layer (Wortmann & Flüchter, 2015). In comparison to the approach of Fleisch et al. (2015) the object layer corresponds to the merger of Layer 1 and Layer 2, the network layer is similar to Layer 3 and the IoT cloud recombines Layer 4 and Layer 5.

Within the object layer it is possible to differentiate between the core hardware of the physical object (i.e., representing the thing as standalone without IoT functionalities), additional hardware components which uniquely enable the IoT capabilities of the and software that is embedded in the physical object to ensure its functionality (Wortmann & Flüchter, 2015). The network layer consists in the connectivity to the cloud which can be facilitated by numerous protocols, such as the general TCP/IP or more specific protocols such as MQTT(-SN) or 6LowPAN (Rüth, Schmidt, Serror, Wehrle, & Zimmermann, 2017; Wortmann & Flüchter, 2015). Finally, the layer of the IoT cloud handles the communication and management of the IoT devices and an independent application platform allows to develop IoT applications and run them in a dedicated execution environment (Wortmann & Flüchter, 2015). Crucially for the associated business benefits, analytics and data management software is used to store, process and analyze the raw data gathered by the IoT devices and extract according relevant consolidated information, which can assist, for example, in improving the product performance (Fleisch et al., 2015; Porter & Heppelmann, 2014). To gain deep insights from the available data, it is not only recorded and analyzed as is, but also used to detect patterns that allow to predict certain events, such as machine failures, and give recommendations on how to react to optimize the results given these predictions (Porter & Heppelmann, 2014). Besides, in the same layer, process management software allows to define, execute and monitor processes considering information from diverse sources combining all relevant objects, individuals and software (Wortmann & Flüchter, 2015). Finally, IoT application software manages the interplay of these relevant objects, individuals and software (Wortmann & Flüchter, 2015). Additionally, across all layers software components support identity and security services, such as user authentication, the integration with business systems, such as the incorporation of IoT data in enterprise systems such as ERP, CRM and PLM, and the integration with external information sources, such as weather data or energy prices (Porter & Heppelmann, 2014).

Considering the multifaceted IoT stack and the associated value layers it is clear that an IoT product is more complex than traditional manufactured items or services. This renders the development of IoT products more difficult, while it offers, at the same time, the potential for substantial revenue streams, which can originate from different layers (Porter & Heppelmann, 2015). However, the complexity of such IoT products also introduces issues regarding the distribution of this revenue within the offering company, as typically different business units are responsible for the product features in different layers (Bilgeri et al., 2017). Ultimately, two essential potentials of revenue can be differentiated through the aggregation of value layers: Along with the traditional value proposition of the physical product from the object layer, companies can also capitalize on IoT applications in the IoT cloud layer, which often are digital

services (Davenport, 2013; Wortmann et al., 2017). These services are usually relying radically on the data gathered by the IoT devices, which is why sensor data obtains such a key role in the IoT (C. C. Aggarwal, Ashish, & Sheth, 2013; Iansiti & Lakhani, 2014; O’Leary, 2013).

2.2 Blockchain Technology

2.2.1 Blockchain Fundamentals

The core concepts for blockchain technology were introduced under the pseudonym Satoshi Nakamoto in the Bitcoin whitepaper, which was released to The Cryptography Mailing List at metzdowd.com on November 1st, 2008 (Nakamoto, 2008; TheMailArchive, 2018). The whitepaper is entitled “Bitcoin: A Peer-to-Peer Electronic Cash System” and presents the technical foundations for such a system, which was later initiated in practice by the generation of the genesis (i.e., first) block on January 3rd 2009. As the title states, Bitcoin introduces the concept of an electronic currency that can be traded from peer-to-peer and does not rely on any trusted third party. In particular, in contrast to traditional currencies, no intermediaries for transactions nor a centralized institution controlling the money supply are necessary for the operation of the system. A key contribution to computer science of the Bitcoin whitepaper, which allowed for such a system to be created, is the solution of the double spending problem in a peer-to-peer network (Decker & Wattenhofer, 2013; Karame, Androulaki, & Capkun, 2012; Nakamoto, 2008). Preventing double spending means to preclude that any one owner of a coin can spend this coin multiple times. Historically, this has been prevented by central authorities, such as the issuer of an electronic currency, which usually have the power to mint new coins and or ban transactions they deem not proper (Conti, Kumar, Lal, & Ruj, 2018; Fridgen et al., 2018; Raval, 2016). Bitcoin, and in a similar way many subsequent blockchains, solve this problem through a combination of economic incentives and modern cryptography, which we discuss in more detail in the following paragraphs.

In its general form, a *blockchain system* is a “fully distributed system for cryptographically capturing and storing a consistent, immutable, linear event log of transactions between networked actors” (Risius & Spohrer, 2017, p. 386). These networked actors form a *peer-to-peer network* and are also called nodes. They reach agreement among each other on what data to append to this immutable log (or ledger) through a *consensus mechanism*. While the term *blockchain* can refer to the overall system, in its essence it represents the immutable ledger itself. From a technological perspective, one can consequently distinguish between three core layers of a blockchain system, as shown in Figure 2: The blockchain itself, the consensus mechanism and the peer-to-peer

network. Because of the trust provided by using a blockchain as a shared, secure computing platform, commercial activities between distrustful, even anonymous agents become possible (Nærland et al., 2017).

Related to these core definitions, additional terminology is used in the realm of blockchain technology: A *coin* or a *token* is a medium of exchange that can be transacted over such a blockchain. In turn, *cryptocurrency* refers to the system of coins or tokens of a blockchain. While numerous other blockchains have emerged since the inception of Bitcoin, offering a wide range of value propositions, the overwhelming part of them are based on the same structural mechanics as Bitcoin, which we are going to discuss in detail in the following.




Figure 2. Core technological layers of a blockchain system

Blockchain

The blockchain is the immutable ledger storing the information of the blockchain system, which typically consists in either simple transactions of cryptocurrency or more complex rule-based transactions (e.g., from smart contracts). This information is stored in separated batches, called blocks, which are appended to each other with progressing time as shown in Figure 3. Each block consists of two parts, a block header consisting in meta data and the rest of the block which contains the transactions. In case of the Bitcoin blockchain, the block header is built of meta data such as the hash of the header of the previous block which links each block to its predecessor, the nonce which plays a key role for the consensus mechanism and the merkle root which secures all the transactions of a block. Additionally, it also contains the Bitcoin version number, a UNIX timestamp and the difficulty target for the block. The rest of the block contains all the transactions included in the block, which prompt actual state changes of the blockchain. Note that as data can only be appended, the blockchain not only contains the current state but also the entire history of states that ever existed.

Each individual block is linked cryptographically to its preceding block by referencing the hash value of this preceding block in its header. This mechanism assembles the individual blocks into a chain and is a critical component for the immutability of a blockchain: It leads to the fact

that an attacker aiming to change a specific transaction in a block i would also have to change the entries of all subsequent blocks up to the present one. Namely, changing a transaction in a specific block i would result in a change of the merkle root of this block (see below) and thus alter the hash of the block header. Consequently, the attacker would also have to change the *hash of previous block header* in the next block $i + 1$, which in turn would lead to a need for change in the next block $i + 2$ for the same reasons, ultimately demanding changes in each block after the specific block i . Because adding a new block (or substituting a block through an updated version) is computationally very demanding (see Consensus Mechanism), such an attack becomes increasingly more difficult the more blocks follow in the chain after the targeted block i . As every block following the block containing a particular transaction is adding to the security of this transaction, the blocks are also called to confirm this transaction. The first confirmation of a transaction is the publishing of the valid block which contains the transaction itself, and every following valid block adds one confirmation. In practice, many professional services such as exchanges expect six confirmations in the Bitcoin blockchain to accept a payment as valid, which corresponds to a waiting time of around one hour (e.g., Kraken, 2019).

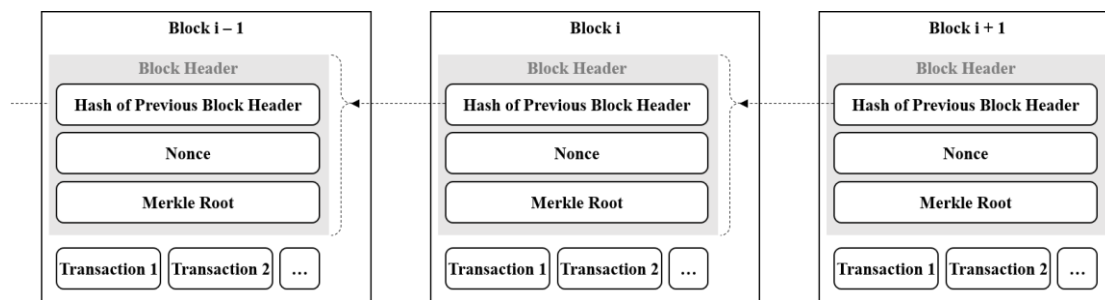


Figure 3 . Simplified structure of the bitcoin blockchain

The immutability of individual transactions is enabled by the fact that it is impossible to change a transaction without altering the corresponding block header, and thus all subsequent blocks. As elaborated above, a change in a specific block header leads to necessary adaptations in all subsequent blocks. Consequently, the goal is that all the transactions of a block become directly or indirectly part of this block header, which creates the link to the next block. The naive way of implementation would be to include all transactions in the block header. However, in order to support the quick verification of individual transactions by devices with limited bandwidth, storage and or computational capacity, an alternative approach is taken: Transactions are assembled in a merkle tree and only the merkle root is included in the header (Merkle, 1989). The structure of such a merkle tree is shown in Figure 4. The original leafs are constructed by hashing the transactions Tx_1 to Tx_n of all n transactions of the block. These hashes are then subsequently hashed again pairwise, which is continued until one single hash is left, which is called the root of

the three or merkle root. This merkle root is inserted in the block header as a reference for all transactions. As the merkle root is dependent on the hashes of all transactions, it fulfills the desired property that it changes given any single transaction is altered. Compared to the naive implementation, it bears the following advantage: Let us assume a device limited in its bandwidth, storage and or computational capacity needs to verify the validity of a given transaction Tx_3 in block i . To achieve this, it suffices that the device keeps track of the block headers and obtains some additional information for block i . Given it obtains the greyed values in Figure 4, it can verify that transaction Tx_3 is in fact included in the block by calculating the merkle root. By verifying that the block hash is included in the next block and the blockchain continues to grow, the device can verify that the overall network includes this transaction in the chain and, thus, considers it valid. Consequently, even though the device cannot verify the validity directly by inspecting preceding transactions or account balances, it can verify it indirectly by relying on the overall network. Naturally, this would also be possible if all transactions were included in the block heard, but this would increase the amount of data to be downloaded and inspected by the device tremendously for every single block. In Bitcoin this mechanism of indirect verification leading to reduced complexity is referred to as simplified payment verification (SPV) (Nakamoto, 2008).

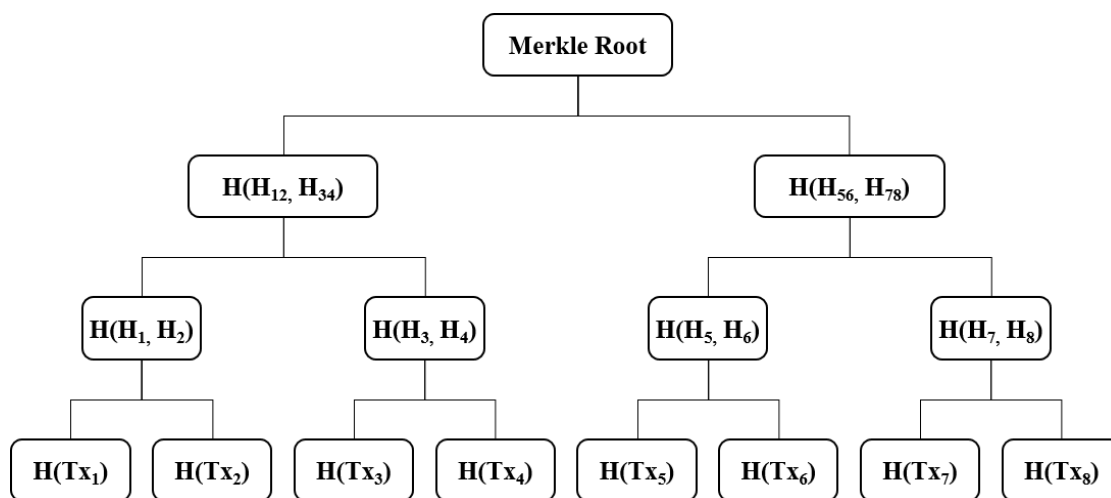


Figure 4 . Merkle tree of transactions

Consensus Mechanism

Adding a new block to the chain requires foremost that the nodes of the peer-to-peer network agree among each other on the specific content of this new block. In other words, the network

participants need to reach consensus, which is why the configuration of this whole process is referred to as consensus mechanism. In particular, consensus needs to be reached among nodes whose identities and intentions may be unknown and, thus, the consensus mechanism should be robust towards a certain share of faulty or malicious nodes. Historically, the problem of finding consensus among a group of networked actors that include faulty or malicious participants is referred to as the Byzantines General Problem and faulty or malicious network participants are called byzantine nodes (Lamport et al., 1982). It is a widely discussed problem in computer science and different solutions had been proposed before the emergence of Bitcoin, for example the influential Paxos algorithm (Lamport, 1998) which has been altered in many variations and can, for instance, be extended to obtain Practical Byzantine Fault Tolerance (PBFT) (Castro & Liskov, 1999). However, these approaches, while taking into account byzantine nodes, do assume additional properties concerning network participants, such as a maximum number of nodes, a fixed set of nodes, or even known identities. While this may be practical in some circumstances, a decentralized ecosystem such as Bitcoin needs the ability to constantly change the number of nodes and cannot assume to know anything about their identity. These properties are crucial to establish a true peer-to-peer network, allow anyone to partake in the process of adding blocks and thus prohibit censorship and increase the difficulty of distributed denial-of-service (DDoS) attacks.

In most blockchain systems, consensus is achieved on a round by round basis by randomly selecting a node to propose a block and economic incentives for this node that the proposed block adheres to the system rules (e.g., no double spend). Hereby, in each round one new block is added to the chain. This whole process is referred to as mining and participating nodes are called miners. Importantly, the selection of miners is not made by any particular entity or third party and instead happens in a completely decentralized manner. In the case of Bitcoin, the selection is achieved by letting nodes solve a mathematical puzzle which is cracked by trial and error (Nakamoto, 2008). Whoever completes the puzzle can propose a new block, after which a new round starts. Specifically, the puzzle to add a new block $i + 1$ to the chain of length i is as follows: A miner needs to adjust the nonce in the block header $i + 1$ in such a way, that the hash of this block header starts with a certain number of zero bits. The nonce is a figure that can be adjusted by the miner at his choice. It only serves this purpose of mining and has no other meaning. To illustrate the process of mining, we can think of a miner that starts calculating the block hash with the nonce equals zero. If this does not deliver a block hash with the necessary amount of zeros, the nonce is incremented iteratively by one, until a valid nonce is found. Subsequently, the new block is broadcast to the whole network and the mining of a new block is started. Note that different miners do not solve exactly the same problem, as the transactions they include in the blocks are different.

Given the properties of hash functions, the chance of any nonce delivering a correct result cannot be predicted and, thus, the chances of solving the puzzle can only be improved by increasing the amount of block hashes a miner can calculate (for different nonces) in a given time window. The computational capacity to calculate these hashes is measured through the hash rate (i.e., the amount of hashes calculated per second). The hash rate can refer to the overall network (i.e., the hash rate of all miners combined) or to individual miners or mining devices. Additionally, finding a solution for the puzzle can be made arbitrarily hard (i.e., by increasing the number of zero bits demanded), while checking that a proposed solution does in fact hold is computationally very easy and always of the same complexity, no matter how difficult the puzzle is made. In practice, the difficulty of the puzzle is adjusted regularly in such a way that the average time between two blocks (i.e., the average time needed for miners to solve the puzzle) stays constant at 10 minutes. Besides effectively selecting individual miners to propose blocks, the mining process also fulfills other roles. In particular, it prevents so-called Sybil attacks, in which an attacker creates a large number of replica nodes (called Sybil nodes) to subvert a reputation or voting based system. As in Proof of Work (PoW) the weight of a miner is solely based on its computational power Sybil attacks are prevented.

The economic incentives to participate faithfully in the mining process are tied to the block reward which is attributed to each new block. Namely, the miner of every block receives a number of newly minted Bitcoins as a compensation for participating in the mining. For these newly minted Bitcoins to be valid, the block proposed by the miner must be valid, too, and be included in the chain. As other miners will verify the validity of a newly proposed block before continuing with the next round, every miner has a high incentive to only propose a valid block. In case of an invalid block, other miners would simply discard it and propose a new block at the same height (i.e., at the same position of the blockchain). In case of the existence of several parallel chains, the longest one is considered to be valid. This set of incentives and rules effectively prevents double spending of Bitcoins. Note that while the Bitcoins from the block reward are created out of thin air, the amount of the block reward, which decreases exponentially, is specified in the initial Bitcoin implementation and lets any participant calculate exactly how many Bitcoins will be in circulation at any given time.

While the consensus mechanism outlined above is rather robust, it requires at least a certain share of nodes to behave faithfully, otherwise leaving open a number of attack possibilities on the system. Obviously, consensus can only be reached if at least some share of the miners act according to the system rules. A naive attack would consist in an attacker obtaining over half of the overall computational power of all miners (i.e., a so-called 51% attack). Statistically, such an attacker is able to produce more blocks than all other miners combined in a given time.

Consequently, this attacker could create the longest, and thus valid, chain all by himself and, for example, censor transactions. It would also allow the attacker to double spend coins by temporarily hiding this longest chain from the rest of the network. However, certain attacks are already possible controlling a share as low as 25% of the total hash rate (Eyal & Sirer, 2018).

Consensus mechanisms based on the concept described above are called PoW, because displaying the nonce in the block proofs that (computational) work was invested to find this nonce. This principle is originally inspired by Hashcash which was proposed by Adam Back to be used as a remedy to denial-of-service attacks (Back, 2002). There are other consensus mechanisms, such as Proof of Stake (PoS), where miners are randomly selected through another process and ensure their economic incentives by pledging a certain amount of cryptocurrency. This pledged amount can be removed should the miner not adhere to the system rules (i.e., so-called slashing conditions).

Peer-to-Peer Network

Blockchains serve as a source of truth within a peer-to-peer network of nodes, which use the ledger as a source of information and at the same time are responsible for keeping it alive. One can distinguish between *full nodes* and *light nodes*. In Bitcoin, full nodes download every block in its entirety, including all transactions, and verify that all consensus rules are adhered to in the whole blockchain (e.g., correct block reward, correctly signed transactions, no double spend, etc.). Often, full nodes store the complete blockchain, although only a fraction of it is actually needed for verification purposes (i.e., only unspent transaction outputs are relevant at any given time). In contrast, light nodes do not verify the consensus rules for the entire blockchain and instead rely on full nodes for that. Usually, light nodes are limited to download and verify block headers and rely on SPV to verify the existence of individual transactions. A special role in the network play miners, which are responsible for updating the ledger by appending new blocks. In principle, any node can become a miner, ensuring the decentralization of the network. Miners usually run full nodes, which is necessary to judge the validity of new incoming transactions when mining a new block. Currently, there are around 10,000 active full nodes supporting the bitcoin network (Bitnodes, 2019).

Blockchain Types

While the mechanics described in this section apply to a wide range of blockchains, it is worthwhile to note that there are also other types of blockchains. In particular, one can distinguish between public and private on the one hand, and permissioned and permissionless blockchains on the other hand, as shown in Table 1 (Beck, Müller-Bloch, & King, 2018). Public permissionless

blockchains allow anyone to become a node that can read and submit transactions, as well as mine the chain. In contrast, public permissioned blockchains allow only a selected set of nodes to participate in the mining of the chain. As a third option, private chains allow only a selected set of nodes in the first places. Mining is often reserved for another set of selected mining nodes. Private blockchains are in particular popular for prototyping in corporations (Iansiti & Lakhani, 2017). In this dissertation, the focus lies on public permissionless blockchains such as Bitcoin or Ethereum. The general introduction made in this section is valid for most of the public permissionless blockchains.

		Reading access	
		Public	Private
Mining access	Permissionless	Anyone can become a node and read and submit transactions. Additionally, all nodes are allowed to participate in the mining of the chain.	Not applicable
	Permissioned	Anyone can become a node and read and submit transactions. However, only selected nodes (i.e., permissioned nodes) are allowed to participate in the mining of the chain.	Only selected entities can become nodes and read and submit transaction. Specific access may differ for various types of nodes. Similarly, mining is restricted to a selected set of permissioned nodes.

Table 1 . Blockchain types

Challenges

While blockchain is a very promising technology it is also in its infancy and still faces numerous fundamental challenges that hinder adoption at the current point in time. Importantly, the scalability of blockchains towards the usage for worldwide and enterprise-grade solutions is still an open issue (Croman et al., 2016). Additionally, currently dominating PoW consensus mechanisms consume a lot of energy. Consensus mechanisms based on PoS could remedy this, and the Ethereum community is already working towards enabling this type of consensus on their network. Furthermore, as transactions in public permissionless blockchains are accessible to be read by anyone, the privacy of users is not always guaranteed, even though transactions can only be connected to pseudonymous identities directly. This issue has in fact led to a whole breed of privacy focused cryptocurrencies (e.g., Monero, Zcash or Grin) that make it increasingly difficult to relate inputs and outputs to individual pseudonyms or addresses and also disguise the amounts sent in transactions. Currently, these approaches cannot be easily transferred to Bitcoin or Ethereum though and research towards privacy-preserving mechanisms is still needed. Moreover, a pressing issue in the area of decentralized systems is their governance. Ensuring that

communities do not get stuck over disagreements, that the development of projects continues at all and contributors are incentivized to work, restricting the power of individuals or groups that might rise in influence, these are all issues of governance that have not yet been solved to a satisfying degree. Numerous new blockchain projects focus in particular on these issues of governance (e.g., Decred) and how collaboration in decentralized systems can be organized effectively. Finally, the development of blockchain systems has focused on solving hard technical problems in the past, while the interaction with regular users was not of much concern for a long time. Consequently, the usability of many use cases is not highly developed. If blockchain technology should be adopted by the masses, this is an issue that needs to be solved. Projects such as Metamask and Crypto Kitties have made tremendous progress in this directions in the last years.

2.2.2 Decentralized Applications

While Bitcoin and blockchain technology were originally built to create electronic cash, namely to send a certain amount of money from A to B, the community quickly realized that the possibility to reflect more complex transactions in a blockchain could reveal great opportunities. In fact, already the implementation of Bitcoin includes the option to create rather sophisticated transactions by leveraging Bitcoin script. Bitcoin script is a basic, stack-based scripting system which is used to structure transactions in Bitcoin. It enables, for example, so-called multisignature (or multisig) transactions which allow coins to be spent only if multiple owners sign the transaction. Multisig transactions can be further adjusted such that a minimal number out of a group of owners, for example three out of five, need to sign a transaction before it is deemed valid, where these three can be arbitrarily chosen out of the five co-owners. Such multisig transactions are for instance useful to share the responsibility of funds between various parties, prevent the loss of funds or enable deposit-like transactions. However, Bitcoin script is intentionally restrained in its functionality (e.g., no looping or recursion) to decrease security risks, as the potential of scripting could be used to attack the network in various forms. Consequently, it is not possible to run arbitrarily complex programs on the Bitcoin blockchain.

Going beyond the quest to process advanced financial transactions, the vision expanded to the objective of running full-fledged applications on a blockchain, so-called DApps. DApps are applications developed based on open, decentralized, peer-to-peer infrastructure services (Antonopoulos & Wood, 2018). This is an emerging type of software which is based on achieving consensus in a decentralized manner (Raval, 2016). As DApps are a recent phenomenon, the definition of what constitutes a DApp is still vague (Raval, 2016). Common features an application should exhibit to be considered a DApp are that the software is open source, consensus

is achieved in a decentralized manner and there is no central point of failure (Antonopoulos & Wood, 2018; Cai et al., 2018; Notheisen et al., 2017; Raval, 2016). To keep a DApp running in a decentralized manner some scholars argue that it should also feature a dedicated cryptocurrency to incentivize the necessary actors (Cai et al., 2018; Raval, 2016). However, it is also possible to achieve this without an additional currency (Adams, 2019). DApps are commonly thought of to have the potential to disrupt a wide range of industries and establish a new economic paradigm of decentralization (Antonopoulos, 2016; Swan, 2015; D. Tapscott & Tapscott, 2016b). The variety of use cases DApps could implement is large and includes products and services as different as voting systems, sharing applications, cryptocurrency exchanges, funding platforms or platforms to manage intellectual property, for example of music songs (Beck, Müller-Bloch, & King, 2018; Bogner et al., 2016). Compared to traditional software applications, DApps can be more transparent and resilient because they can inherit the core properties of the blockchain they are built on (Chanson et al., 2019; Raval, 2016). For example, a system that stores valuable sensor data and is built on an immutable public blockchain can leverage this immutability to ensure that the sensor data cannot be tampered with (Chanson et al., 2019). More generally, based on the trust-free interaction blockchains enable, DApps allow to form similarly trust-free systems which let diverse sets of agents interact and reach a common understanding leveraging the underlying blockchain as a shared source of truth (Notheisen et al., 2017). As a further advantage to traditional software, DApps facilitate the incorporation of incentive structures directly into the core product, typically through the usage of a cryptocurrency or a token specific to the DApp ecosystem (Raval, 2016). For example, Maker is building a DApp which can be used as a platform for credit issuing (Christensen et al., 2018). Users can request a credit in a cryptocurrency which is pegged to the U.S. dollar in exchange for depositing a collateral in another free-floating cryptocurrency. Another dedicated token called MKR is used to incentivize numerous actors to keep the system alive, for example to identify which credits are under-collateralized and liquidate according positions. Note that this system has no employees dedicated to the day-to-day business of the lending platform and that all necessary actions are executed by a distributed set of actors solely because of economic incentives implemented via MKR. All the involved cryptocurrencies are built on the same underlying blockchain. Similarly, all the logic of the Maker platform is also built on the same underlying blockchain. While the exact mechanics of how MKR is used to incentivize the actors in the Maker ecosystem is beyond scope here, the example illustrates well how closely the economic and technical worlds are intertwined in DApps.

Technically, DApps are based at their core on *smart contracts*. Smart contracts are computer protocols, which execute specific logic operations and transactions automatically without the participation of a third party (Beck, Müller-Bloch, & King, 2018; Chanson, Gjoen, et al., 2018).

Originally, the concept of smart contracts was introduced by Nick Szabo as a way to “formalize and secure relationships over computer networks” (1997, p. 1). , in a similar way as contract law formalizes and secures relationships between business partners in the paper-based world. The advent of blockchain technology enabled the creation of such contracts envisioned by Szabo such that their execution became independent of a trusted third party. Namely, smart contracts can be implemented as small programs which run on dedicated blockchains and are invoked by transactions on these blockchains. As such, the term smart contract has evolved to include any general purpose computation which is executed on a blockchain (Buterin, 2013). For example, such smart contracts can enable the exchange of one cryptocurrency for another between two parties. This simple exchange contract can then be further enhanced to a complete DApp which allows the exchange of multiple cryptocurrencies between an unlimited amount of users, the creation of limit orders and more complex transactions (Adams, 2019; AuroraLabs, 2019). Such DApps are referred to as decentralized exchanges (DEXs).

Ethereum was the first blockchain to launch with a built-in Turing-complete programming language and thus able to host general purpose smart contracts (Buterin, 2013; Egelund-Müller et al., 2017; Raval, 2016). As such, Ethereum is often described as a platform for DApps (Halaburda & Sarvary, 2016). Currently, it is based on a PoW consensus mechanism very similar to Bitcoin, however, plans exist to change to a PoS type consensus mechanism since a long time.

While the ability to process programs from a Turing-complete set of instructions gives a lot of power to creators of smart contracts, it also creates a breadth of security risks for these programs and the underlying blockchain. In Ethereum, a number of high-profile hacks that have exploited bugs in deployed smart contracts or related software have impressively showcased this fact, for example the DAO hack which was based on a reentrancy attack and ultimately led to the hard fork between Ethereum and Ethereum Classic (Madeira, 2019), or the loss of funds from Parity (Redman, 2017). It is interesting to note that bugs in smart contracts (e.g., the DAO hack) or the related programming language (e.g., DDOS attack) can also impact the stability of the whole blockchain. As such, an underpriced opcode in solidity, the Ethereum specific smart contract language, was used to launch a DDOS attack on Ethereum (Rush, 2018). For these reasons, Bitcoin has never extended the capabilities of its scripting language to a Turing complete set of instructions and newer approaches such as the Libra blockchain initiated by Facebook put a heavy emphasis on the security related to the programming language used to write smart contract and the contracts themselves (Amsden et al., 2019).

2.2.3 Decentralized Finance

The new possibilities of smart contracts have been leveraged to create a wealth of DApps, especially in the realm of financial services. As the original goal of blockchain technology was to create electronic cash, and cryptocurrencies often play a fundamental part in the incentivization of DApp users and contributors, it is not surprising that an initial focus of DApp development was on the finance industry. For example, DEXs were created to trade cryptocurrencies, new funding schemes were initiated to allow blockchain projects to fund themselves in a decentralized manner, and more complex financial products were created on top of other DApps, for instance credit platforms or markets for margin trading. The initial focus of DApps on financial applications is also a reason why research on blockchain technology has concentrated on financial applications for a long time (Risius & Spohrer, 2017; A. Tapscott & Tapscott, 2017; Wörner, Von Bomhard, Schreier, & Bilgeri, 2016). In fact, vivid research streams have formed concerning different aspects of DeFi, for example the influence of social media on DeFi (Chanson, Gjoen, et al., 2018; Mai et al., 2018; Xie, Chen, & Hu, 2017). This development of numerous financial DApps leads towards an ecosystem of decentralized financial services that interact and can be combined to create new and more sophisticated products. All these applications are combined under the umbrella term of DeFi.

One of the largest influences smart contracts and DApps have had so far is the creation of a new funding mechanism, which bears similarities to Initial Public Offerings (IPOs) (e.g., capital raise through security release), venture capital (VC) (e.g., early stage of investment) and crowdfunding (e.g., open to retail investors) (Chanson, Gjoen, et al., 2018). Because of the similarities to IPOs, a blockchain-based fundraise is often referred to as an Initial Coin Offering (ICO). ICOs can be conceptualized as a fundraising mechanism, in which new project-specific coins are sold to raise capital for a previously described project (Chanson, Gjoen, et al., 2018). Projects launching an ICO mainly aim to secure funding whereas the investors' objective is the possession of the project-specific coins and corresponding financial exposure to the success of the project. The exchange of capital from investors for the project-specific coins is completely automated through different smart contracts. As such, blockchain-based fundraising is a typical example of disintermediation (e.g., of investment banks). Blockchain-based fundraising is a complex and multi-faceted phenomenon, and under the umbrella term of ICO various different types of fundraising are referenced. These are sometimes also called token sale, token generation event or initial token offering. A clear definition and use of the terminology has yet to emerge in this novel field of research.

ICOs have developed tremendous influence within and beyond the blockchain sector over the last three years and are now the established funding form for blockchain-related startups (Chanson, Risius, et al., 2018). ICOs enabled the simple blockchain-based funding of new projects, which led to a tremendous increase in the numbers launched and the amount of capital invested in blockchain projects (Chanson, Gjoen, et al. 2018). In 2017, more capital was invested via ICOs than in traditional equity rounds of blockchain startups, adding to a magnitude comparable to all VC Internet investments in a typical quarter (Chanson, Risius, et al., 2018). In 2018 alone, over 600 projects were launched through ICOs (Coindesk, 2019). Furthermore, the potential of ICOs to disrupt traditional financing mechanisms in a wide range of industries was indicated by firms from outside the blockchain industry adopting the concept of ICOs, including established companies like Kodak and Telegram. It is also interesting to note that professional investors such as VCs have embraced the concept of ICOs and by now invest considerable shares of their funds through this mechanism (Wilson, 2018). Consequently, a new IS research stream has emerged in the last two years on the topic of ICOs (Chanson, Gjoen, et al., 2018; Chanson, Risius, et al., 2018; Fridgen et al., 2018; Oliveira, Zavolokina, Bauer, & Schwabe, 2018; Park & Yang, 2018). While this research stream is increasingly gaining traction, it is still in its infancy and leaves many open questions to explore.

Beyond ICOs DApps have developed influence in numerous aspects in the area of DeFi. A prominent example are DEXs which enable the trading of cryptocurrencies without ever giving a trusted third party (i.e., another entity than the users involved in a trade) control over the traded assets. The popularity and importance of DEXs is largely rooted in the loss of user funds through centralized exchanges, be it through a lack of security and resulting hacks or fraud. Such losses happened repeatedly over the last years, often involved the temporary or final loss of funds for users and, resulted, for example, in the famous bankruptcy of the Bitcoin exchange Mt. Gox (Decker & Wattenhofer, 2014). In principle, DEXs can be differentiated into custodial and non-custodial types (Glarner & Lindgren, 2018): In custodial DEXs the exchange of assets is executed via a smart contract which receives the assets and then redistributes it, meaning that the contract acts as a custodian for the assets. This type is implemented for example by Etherdelta, IDEX or Oasis Dex. In contrast, non-custodial DEXs allow for a direct exchange between the users, meaning that the assets never leave the wallet of a user until being directly transferred to another user. Such a type of DEX is implemented for instance by Kyber Network, Uniswap or Airswap. Note that while development on the Etherdelta project has been halted in the meantime, following a settlement of charges of the U.S. Securities and Exchange Commission against its founder Zachary Coburn, other DEXs such as Uniswap keep being developed and have also received major funding from reputable VCs (Chaparro, 2019; Glarner & Lindgren, 2018).

Besides trading with DEXs, DApps can also contribute to a decentralized financial system by creating much-needed stability through stablecoins. To be used in a financial system, a cryptocurrency needs to keep a certain stability such that it can act as a store of value (Halaburda & Sarvary, 2016). However, most cryptocurrencies fluctuate in value considerably at the moment. Therefore, dedicated cryptocurrencies called stablecoins were created with a focus on achieving price stability. Currently, three approaches of creating stablecoins are known, differentiating in the collateral the systems use: Fiat, crypto and algorithmic collaterals. The first and so far most successful approaches are fiat-backed, meaning that the system keeps an amount of fiat in reserve which can be exchanged for the stable cryptocurrency issued and, thus, guarantees its stability. This is very similar to traditional fiat currencies back in the time of the gold standard. The most prolific example of this approach is Tether, one of the first and still largest stablecoins (CoinMarketCap, 2018). Compared to the size of fiat-backed options, crypto-backed flavors are lagging behind. In this case, instead of fiat, cryptocurrency is held as a reserve. The largest and most well-known project of this type is Maker DAO. The third model of algorithmic stable coins does not rely on any reserve, and instead relies merely on an algorithmic system which keeps a currency stable, for example through the issuance and buyback of the currency itself and bonds tied to the currency. It has struggled to develop a considerable dimension, with the most prolific project Basis shutting down due to regulatory reasons, leaving only a smaller project active (CarbonUSD). In principle, all these approaches could be implemented as a DApp, however, without a reliable interaction between traditional financial systems and blockchains a fiat backed stablecoin still needs to rely on trusted third parties such as banks and custodians to provide the collateral and auditing firms for an according supervision. Consequently, at the moment only Maker DAO and other crypto-backed stablecoins are implemented as a DApp.

On top of basic applications of DeFi such as cryptocurrencies and tokens, DEXs or stablecoins, DApps that offer more complex financial products have been created. For example, Dharma has implemented a platform to lend and borrow cryptocurrency, supporting numerous different tokens (Hollander, 2017). Similarly, the Compound protocol allows users to borrow and lend crypto assets, where interests rates are set algorithmically based on supply and demand (Leshner & Hayes, 2019). In turn, dYdX offers an open trading platform with advanced options such as margin trading up to a fourfold leverage (Juliano, 2017). These are first examples of the emerging space of DApps that forms an ecosystem of DeFi and it is expected to grow considerably in the future as the technical foundations for building such DApps become more mature (Beck, Müller-Bloch, & King, 2018; Chanson, Gjoen, et al., 2018; Risius & Spohrer, 2017).

2.3 Economy of Things

2.3.1 Existing Perspectives

In the last decades, technology has fundamentally changed our economy. Specifically the Internet has fueled disintermediation and gave birth to new actors that interact directly with their customers (Antonopoulos, 2016; Beck, Müller-Bloch, & King, 2018; Risius & Spohrer, 2017; Swan, 2015). This development can be conceptualized as the emergence of a *decentralized economy*, which is enabled by technological artifacts such as web protocols or, more recently, blockchains, cryptocurrencies and smart contracts (Antonopoulos, 2016; Beck, Müller-Bloch, & King, 2018; Swan, 2015). A broad agreement among research scholars and practitioners alike persists that blockchain technology could take this change of disintermediation in our economy to a whole new level and shape the future of our economic organizations and interactions decisively (Beck, Müller-Bloch, & King, 2018; D. Tapscott & Tapscott, 2016a). Reputed venture capitalists such as Andreessen Horowitz or Union Square Ventures (USV) have been investing in the blockchain sector for years, similar to large corporations which have also launched countless innovation initiatives in that context (AndreessenHorowitz, 2019; Castillo, 2018; Sanderson, 2019; Wilson, 2018). Furthermore, leading technology companies such as IBM and Facebook are invested heavily in the topic through the development of their own blockchain technologies (Amsden et al., 2019; Risius & Spohrer, 2017; Sanderson, 2019). While the magnitude of expected effects is commonly seen as large, numerous different names have emerged to refer to such a newly organized economy. As the technology of blockchain is widely believed to be fundamental for the development of such a decentralized economy, some scholars refer to this vision directly as the *blockchain economy* (Beck, Müller-Bloch, & King, 2018; Manski, 2017; Mattila, 2016; Swan, 2015). Other experts and researchers name it the *Internet of money* or the *Internet of value*, reflecting that decentralized interaction is made predominantly via the Internet and that an important aspect of newly created DApps is their ability to directly include money or, more generally, digital assets in the form of cryptocurrencies in their inner workings (Antonopoulos, 2014, 2016; Lee Kuo Chuen, 2015; Peters & Panayi, 2016; Ravikant, 2013; D. Tapscott & Tapscott, 2016a). In turn, some researchers view blockchain as an enabling technology for a *smart economy*, which is also the stance of one of the leading smart contract platforms called NEO (NEO, 2016; Sun, Yan, & Zhang, 2016; Wang, 2018).

What all these terms unites is that they describe a fundamental shift in our economy from centralized to decentralized, from platforms with huge leverage over other stakeholders to systems with no single point of failure. The cornerstones of such a decentralized economy are reflected in

these different approaches of people to refer to the phenomenon. In short, the emerging decentralized economy is based on blockchain technology and, thus, disintermediation and the loss of relevance of trusted third parties. Besides, it is characterized by a new combination of technology and representation of value (i.e., digital assets) and the potential to automate a wide range of actions which need to be based on thorough checks through smart contracts. In general, technological artifacts such as blockchains, cryptocurrencies and smart contracts build the foundation for the emerging concept of a decentralized economy, in which technology replaces currently dominant intermediaries and enables new entities to engage directly in commercial transactions (Antonopoulos, 2016; Beck, Müller-Bloch, & King, 2018; Swan, 2015). In particular, only this new technology allows to build complex large-scale systems that do not rely on a trusted third party to function, which was the core invention of Bitcoin at launch (Nakamoto, 2008). Regarding the economic aspects, a core innovation is that the Internet as a medium of information exchange is enhanced to a medium of value exchange. Necessary for this is the invention of cryptocurrencies, which can easily be exchanged among users, and that these cryptocurrencies actually represent an asset similar to securities traded on stock exchanges (D. Tapscott & Tapscott, 2016a). The enhanced possibilities regarding automation are largely enabled by smart contracts. These contracts allow to enforce transactions automatically given some well-defined conditions the involved parties have agreed upon previously are fulfilled (Beck, Müller-Bloch, & King, 2018; Buterin, 2013). As such, smart contracts allow not only to automate actions but also mandatory verifications related to these actions, accordingly increasing the amount of actions possible to consider vastly. In the context of the IoT, the emergence of such a decentralized economy results in very particular applications which, combined, give rise to the EoT (Beck et al., 2016; Elsdén et al., 2018; Kouzinopoulos et al., 2018; Sousa et al., 2018). The EoT is sometimes also referred to as the *machine* or *machine to machine economy* (Brelie & Giehl, 2019; Panarello, Tapas, Merlino, Longo, & Puliafito, 2018; Strugar et al., 2018).

2.3.2 DApps and DeFi in the context of the IoT

In the context of the IoT, blockchain technology is widely seen as a foundational enabler of the EoT (Beck et al., 2016; Elsdén et al., 2018; Kouzinopoulos et al., 2018; Sousa et al., 2018). While the current form of the IoT is characterized by devices that are merely *connected* to the Internet and *generate data*, the EoT results from *interconnected* things which can *interpret data*. In particular, widely discussed use cases in the realm of the IoT are machine-to-machine coordination and sensor data markets, resulting in devices that become increasingly autonomous at an individual or very low level (Christidis & Devetsikiotis, 2016; Noyen et al., 2014; Subramanian, 2018; Wörner & von Bomhard, 2014; Zhang & Wen, 2017). All these examples

showcase, how simple IoT devices that mainly used to sense, process and exchange information become things that are autonomously acting as agents within an economic incentive system. This transfer of capabilities can lead to the development of a novel type of economy, for example by forming decentralized electronic marketplaces. The term EoT reflects this conversion from the IoT in an intuitive way and is therefore widely used to describe the underlying evolution (Beck et al., 2016; Elsdén et al., 2018; Kouzinopoulos et al., 2018; Sousa et al., 2018).

For example, blockchains could be used to decentralize firmware upgrades of IoT devices (Christidis & Devetsikiotis, 2016). Currently, devices receive these upgrades centrally from the manufacturer, which is consequently responsible to grant access to such updates for a long time. Based on a blockchain, this could be organized in a decentral manner instead. The manufacturer could create a smart contract storing all the hashes of the firmware updates, so that clients can verify the software no matter from whom it was received. This smart contract could also be used for queries regarding available updates and sources for that software. Initially, updates would need to be provided by the manufacturer, but as soon as enough third parties are in possession of an update, they could serve it instead, replacing the manufacturer as a source. To foster such behavior, the smart contract could implement incentives for nodes in the network to provide the latest software to other nodes, for example a small payment for each delivery. The amount of the payment could be decided by the manufacturer, or by the network of nodes. As an alternative, nodes could also be given the freedom to choose any amount they wish. The possibility to create incentives for third parties to participate in this networking is greatly simplified as blockchains enable cryptocurrencies and, thus, inherently have access to a highly automatable billing layer.

Another prominent example of EoT applications are peer-to-peer energy markets (Christidis & Devetsikiotis, 2016; Imbault, Swiatek, De Beaufort, & Plana, 2017; Meeuw, Schopfer, Ryder, & Wortmann, 2018; Münsing, Mather, & Moura, 2017). Especially with the rising deployment of decentralized energy resources, in particular renewable energy sources such as photovoltaic systems or wind turbines, the importance of microgrids for the local distribution of power increases historically (Imbault et al., 2017; Meeuw et al., 2018). Microgrids are localized distributed energy systems, often connecting small communities of neighbors who individually provide resources to the system to produce and store energy (e.g., solar panels or batteries) (Vatanparvar & Faruque, 2018). The main purpose of such a microgrid is to ensure that locally produced or stored energy can be consumed locally, without relying on a higher level grid and its associated costs (Vatanparvar & Faruque, 2018). Consequently, this is also a classic use case of the broader decentralized economy, reducing the dependence on an intermediary, namely the utility companies which typically own or manage conventional larger grids (Meeuw et al., 2018; Münsing et al., 2017). Blockchain technology is often cited as a potential remedy against

difficulties in the setup of microgrids such as integrating a billing process, keeping track of energy flows and providing transparency in a secure and automated manner while at the same time being deployed in unsupervised surroundings (Chanson et al., 2019; Meeuw et al., 2018). Accordingly, such microgrids leveraging blockchain technology have already been set up in large field tests in Brooklyn, NY and Walenstadt, Switzerland (Meeuw et al., 2018; Mengelkamp et al., 2018).

A plethora of other use cases in the context of EoT are prominently being discussed and have obtained first promising results from practice. For example the tracking of goods in international supply chains (Christidis & Devetsikiotis, 2016; Modum, 2017; Tian, 2016), the general provision of sensor data against cryptocurrency micropayments, similar to the principle of software as a service (Noyen et al., 2014; Wörner & von Bomhard, 2014), robo-taxis or the sharing of connected devices such as cars or home appliances (Bogner et al., 2016; Elsdén et al., 2018; Jarvenpää & Teigland, 2017; Musk, 2019). Consequently, some of the biggest IoT companies (e.g., Foxconn, Siemens, Bosch, Cisco) have formed a group called the *Trusted IoT Alliance* to foster the development of blockchain infrastructure which enables secure and scaled IoT ecosystems (TrustedIoTAlliance, 2019).

While many EoT applications are currently being envisaged, one of the most impactful changes could be the development of general-purpose data markets which can be accessed equally by all actors in a decentralized system (Noyen et al., 2014; Özyilmaz, Dougan, & Yurdakul, 2018; Panarello et al., 2018; Wörner & von Bomhard, 2014; Zhang & Wen, 2017). Currently, the business model of the Internet is based on leveraging centrally organized user data and this is also reflected in prominent applications envisaged for the EoT. In fact, many of the examples discussed before are based on leveraging the data gathered by sensors, for instance electricity consumption or the location of parts in global supply chains. However, these are special cases and the data gathered might not be directly monetized and only serve as an enabler of the overall function of the product. In contrast, there are also visions of unified data markets which could provide all sorts of data in a decentralized fashion (Noyen et al., 2014; Panarello et al., 2018; Zhang & Wen, 2017). Additionally, work is put into developing data markets that enable to analyze personal data without compromising the privacy of the users the data is referring to (Zyskind, Nathan, & Pentland, 2015). Such data markets could be of wide influence, especially in the context of machine learning and artificial intelligence (i.e., the training of algorithms) (Mattila, 2016; Özyilmaz et al., 2018). As IoT sensors gather a lot of information and the amount of sensor data will drastically increase in the future, the EoT could play a decisive role in the development of working data markets which provide enough resources for the training of artificial intelligence in an wide range of use cases (Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015; Noyen et al., 2014; Oks et al., 2017; Panarello et al., 2018; Zhang & Wen, 2017).

Chapter 3

DApps: Privacy-Preserving Protection of Sensor Data²

As we have introduced, trustworthy data is required as a fundamental basis for the EoT. In particular, as things become more autonomous and base complex decisions on data they have acquired automatically from other potentially unknown sources, it is essential that such data be reliable – or it can be verified at least for critical decisions if the data is indeed reliable. However, with the rapid deployment of connected devices around the world, which form increasingly complex large-scale IoT systems, the number of attack vectors has steadily increased and manipulation in such systems have become more common. As many IoT systems handle especially sensitive personal or company information, preserving the privacy during storage and exchange of such data is particularly relevant. In this chapter, we investigate how DApps can facilitate the reliable gathering and exchange of IoT data in a privacy-preserving manner and under what circumstances blockchain technology might support that process.

3.1 Introduction

In recent years, new forms of information technology (IT) (e.g., sensors and mobile devices) have dramatically expanded what can be measured and analyzed, thereby posing completely new challenges regarding security and privacy (Lee et al., 2018; Newell & Marabelli, 2015; Sicari et al., 2015; Weber, 2010). The potential for IS-related security and privacy issues to affect

² Parts of this chapter, which are not further demarcated in the text, were initially published in the context of the following academic publication: Chanson, Gjoen, et al. (2018), Chanson et al. (2019) and Chanson et al. (2020).

customers in their daily lives and private spheres makes these challenges top business priorities (Sicari et al., 2015). In fact, the widespread adoption of smart products might depend on the ability of organizations to offer systems that ensure adequate security levels while guaranteeing sufficient user privacy (Sicari, Cappiello, De Pellegrini, Miorandi, & Coen-Porisini, 2016). Such IoT systems, referring to a constantly growing pool of smart, connected devices, including cars, health applications, and industry machinery, offer adversaries a whole new range of attack vectors for manipulating IS (Lowry et al., 2017; Porter & Heppelmann, 2015). IoT systems are usually characterized by multi-party ecosystems, with data pipelines crossing organizational borders (C. C. Aggarwal et al., 2013; Roman, Zhou, & Lopez, 2013). In such systems, malicious adversaries can manipulate “data at various stages in the [processing] pipeline”, from sensor to service, making data integrity a key concern (C. C. Aggarwal et al., 2013, p. 419). The IS research community is well aware of these challenges and has specifically called for more design research that can facilitate secure and reliable data processing and exchange in multi-party ecosystems (Bélanger & Crossler, 2011; Pavlou, 2011; Smith et al., 2011).

Previous research has indicated that blockchain technology is a promising means to mitigate issues of data security arising in the IoT and has some decisive advantages over a conventional database system on central servers (Glaser, 2017; Hyvärinen et al., 2017; Nærland et al., 2017). More specifically, blockchains provide tamper-proof storage capabilities in the form of a distributed ledger that can be used to securely store and exchange IoT sensor data. However, core challenges, such as privacy, scalability, and potentially prohibitive transaction costs, remain to be addressed (Beck et al., 2016; Notheisen et al., 2017; Risius & Spohrer, 2017). While there are a variety of different blockchain-based IoT systems currently under development (Curtis, 2015; Mengelkamp et al., 2018; Modum, 2018), the corresponding academic research is still in its infancy (Avital et al., 2016; Beck et al., 2017, 2016; Beck & Müller-Bloch, 2017; Lindman et al., 2017).

In the IS community, privacy and security have been widely discussed as multidisciplinary, diverse concepts (Lowry et al., 2017; Oetzel & Spiekermann, 2014; Sicari et al., 2015). However, most studies do not provide actionable solutions. In this regard, Bélanger and Crossler (2011) note in their seminal literature review that scholars should “conduct design and action research with an eye towards actual implementation” (p. 1035). Similarly, Pavlou (2011) proposes that future IS security and privacy studies should adapt the design science perspective, “with emphasis on building actual implementable tools” (p. 980). While multiple technologies are available to realize IoT sensor data protection systems (SDPSs) (Ayoade, Karande, Khan, & Hamlen, 2018; Machado & Fröhlich, 2018; Margulies, 2015), limited prescriptive knowledge has been gathered to guide the development process of such systems. In addition, the potential of blockchain

technology in SDPSs is, to the best of our knowledge, not yet reflected in the literature. Against this background, we contribute to the IS literature by establishing theoretical insights into how to design an SDPS and by explicitly developing and evaluating a blockchain-based SDPS. More specifically, we aim to answer the following research questions:

RQ 1a: Which fundamental challenges arise in the context of IoT sensor data protection, and which requirements can be derived from these challenges for the design of information systems that facilitate IoT sensor data protection (i.e., SDPSs)?

RQ 1b: Which actionable guidelines in the form of design principles and design features address these design requirements and inform the development of SDPSs?

RQ 1c: What is the value proposition of blockchain technology in the realm of SDPSs, and what fundamental design implications of blockchain-based SDPSs must be considered?

Overall, our research is geared towards a design theory that guides the development of SDPSs that are able to protect IoT sensor data in a privacy-preserving manner. To answer our research questions, we follow the guidelines of design science (Gregor & Jones, 2007; March & Smith, 1995). Within the IS community, the development of design knowledge, be it in the form of design theories, principles, or guidelines, is of high significance for both research and practice (Baskerville, 2008; Hevner et al., 2004; Winter, 2008) and continues to attract a great deal of interest (Baskerville et al., 2015; Gregor & Hevner, 2013; Rai, 2017). We derive an artifact that consists of a set of interrelated design requirements, design principles, and design features. We demonstrate and refine our artifact on the basis of an instantiation that aims to prevent the fraudulent manipulation of car mileage data. Finally, we provide an ex-post evaluation of the artifact and present our results in the form of a design theory.

3.2 Conceptual and Theoretical Background

3.2.1 Internet of Things and Sensor Data

According to Atzori et al. (2010), the IoT refers to “a vision that virtually any physical object can be connected to the Internet”, a vision in which smart, connected devices generate unprecedented amounts of sensor data that can be classified as “big data” (H. Chen, Chiang, & Storey, 2012). Big data, in turn, is characterized by the ever-increasing volume, velocity, and variety of data combined with veracity-related challenges (Clarke, 2016; Goes, 2014; Schroeck, Shockley, Smart, Romero-Morales, & Tufano, 2012). This holds particularly for sensor data,

which is increasing extraordinarily both in the size and speed of data generation (Abbasi, Sarker, & Chiang, 2016; Brynjolfsson & McAfee, 2012). In addition, sensor data is available in a variety of formats and from disparate sources (Brynjolfsson & McAfee, 2012; Schroeck et al., 2012). Finally, veracity considers the varying degrees of reliability and credibility of sensor data sources (Abbasi et al., 2016). In light of these growing datasets and the corresponding technical and economic challenges, companies are increasingly relying on cloud solutions, which are typically operated by third parties (Lowry et al., 2017). In addition, more and more companies exchange and share sensor data to foster cross-organizational collaborations (Anderson, Baskerville, & Kaul, 2017).

3.2.2 Security and Privacy in the Internet of Things

The constantly growing pool of smart, connected IoT devices poses completely new challenges regarding security and privacy (Lee et al., 2018; Sicari et al., 2015; Weber, 2010). Companies are increasingly moving towards cloud solutions and sharing sensor data in multi-party ecosystems (Anderson et al., 2017; Lowry et al., 2017). However, distributed processing and sharing data with third parties is risky, as participating stakeholders (companies and end users) might misuse or lose control over data (Anderson et al., 2017; Moura & Serrão, 2016). Ultimately, the involvement of third parties significantly increases the risk of security and privacy breaches of IS systems (Lowry et al., 2017). In addition to intentional sharing in multi-party networks, unintentional access by malicious adversaries is a major security risk in the IoT, especially because of its “architecture of wireless transmitters and sensors that [...] connect into vast global networks” (Lowry et al., 2017, p. 556). For example, the Internet connectivity of IoT devices can enable malware to quickly infect large populations around the globe (Kolias et al., 2017). Even the networking capabilities of devices that are not connected to the Internet can be exploited to spread malware quickly and unobtrusively (Ronen et al., 2017). This is because IoT sensors are usually unsupervised when collecting data, leaving them particularly prone to various security threats (C. C. Aggarwal et al., 2013; Atzori et al., 2010; Ronen et al., 2017). The multilayered hardware and software stack of IoT solutions also makes these systems vulnerable to a variety of potential attacks (Sicari et al., 2016). For instance, malicious adversaries can manipulate “data at various stages in the [data processing] pipeline”, from sensor to service, making data integrity a key concern (C. C. Aggarwal et al., 2013, p. 419). Furthermore, many of the existing security principles that companies use to protect their systems, including routers, gateways, and firewalls, are not applicable to the IoT, as they “simply do not work for smaller and more mobile ‘things’” (Lowry et al., 2017, p. 556).

Against this background, the IoT fundamentally challenges the field of IS security and privacy, requiring the redefinition of well-established rules and organizational practices to protect sensor data (Fernandes, Rahmati, Eykholt, & Prakash, 2017; Singh, Pasquier, Bacon, Ko, & Eyers, 2016). The confidentiality and integrity of data are essential to security and privacy to ensure that personal data cannot be viewed or manipulated by objectionable third parties (Anderson et al., 2017; Baskerville & Siponen, 2002; Chellappa & Pavlou, 2002). Specifically, privacy is commonly defined as “the ability of the individual to personally control information about oneself” (Stone, Gueutal, Gardner, & McClure, 1983, p. 460). Westin (1967) refers, in particular, to the possibility of data generators to determine the manner, scope, and time in which data is collected by, and transferred to, third parties. The existing IS studies on privacy cover a wide range of aspects and perspectives (Dinev, Hart, & Mullen, 2008; Malhotra, Kim, & Agarwal, 2004; Xu, Dinev, Smith, & Hart, 2011). However, despite the existing body of knowledge, there is a lack of actionable solutions, as Bélanger and Crossler (2011) conclude in their seminal literature review. Specifically, they emphasize that beyond providing conceptual contributions towards the privacy debate, IS research should “conduct design and action research with an eye towards actual implementation” (p. 1035), developing tools to protect information privacy.

In summary, the IoT is advancing much faster than the related privacy and security measures and policies (Singh et al., 2016; Weber, 2010). The resulting security and privacy gaps are potentially dangerous loopholes that can be exploited by malicious actors to the detriment of consumers and organizations (C. C. Aggarwal et al., 2013; Lowry et al., 2017). In fact, the widespread adoption of IoT solutions might depend on organizations’ capabilities to offer systems that ensure adequate security levels while guaranteeing sufficient user privacy (Sicari et al., 2016). As such, the IoT, which is characterized by multi-stage data pipelines and big (sensor) data, has been identified by IS scholars as being “particularly compelling to security and privacy researchers”, as it carries “innate information and privacy risks” (Lowry et al., 2017, p. 546).

3.2.3 Existing Research on SDPSs and their Limitations

SDPSs, which aim to ensure the security and privacy of sensor data, are the subject of an extensive body of literature. In particular, the IS community has made considerable effort to investigate issues of security and privacy (Bulgurcu, Cavusoglu, & Benbasat, 2010; Chatterjee, Sarker, & Valacich, 2015; Y. Chen & Zahedi, 2016), which has resulted in various design theories (Heikka, Baskerville, & Siponen, 2006; Siponen & Iivari, 2006). A key research focus in the area of IS security is the use of organizational policies that define how the users of IS should prevent, identify, and react in security incidents (Anderson et al., 2017; Cram, Proudfoot, & D’Arcy, 2017; Moody, Siponen, & Pahlila, 2018; Niemimaa & Niemimaa, 2017). An excellent review of the

body of knowledge is provided by Cram et al. (2017), who analyzed 114 security policy-related journal articles. From this research stream, the study of Anderson et al. (2017) is especially relevant for our work. They combine discussions of security with those of information privacy, focusing on the risks and rewards of either sharing or retaining full control over data. Thus, they cover a topic that is also fundamental to SDPSs, namely, security, privacy, and the necessity, or economic benefit, of sharing information. However, similar to the approach of other literature on organizational policies, Anderson et al. (2017) deliberately refrain from providing actionable guidelines for the implementation of IS that would enable secure and privacy-preserving data exchange. Rather, they focus on how an organization and its personnel should behave in the vicinity of such systems. A lack of normative results can be similarly observed in most other examples of IS research on SDPSs (Crossler & Posey, 2017). This finding is in line with the seminal literature review by Bélanger and Crossler (2011) on information privacy, in which the authors conclude that “very few articles provide design and action contributions” (p. 1023). Moreover, the IS literature on privacy and security hardly addresses the specific design challenges that arise when processing IoT sensor data.

Beyond the domain of IS, there is a fruitful knowledge base of computer science literature that specifically addresses security and privacy issues in the IoT. Core insights from the latest research include the summation that “the task of affordably supporting security and privacy [in the IoT is] quite challenging” (Trappe, Howard, & Moore, 2015, p. 14) and the observation that while some known security principles should be adaptable to the IoT computing paradigm, “the nature of both physical processes and IoT devices lend themselves to the construction of new security mechanisms” (Fernandes et al., 2017, p. 83). Inspired by such statements, there has been an active stream of research developing specific solutions in the realm of SDPSs (Kolias, Stavrou, Voas, Bojanova, & Kuhn, 2016; Margulies, 2015; Ronen et al., 2017), including work on the potential value contribution of blockchain technology. Ayoade, Karande, Khan, and Hamlen (2018), for example, present a system for the management of IoT data in which all permissions for data access are enforced by smart contracts on a blockchain, which also ensures traceability by the logging of all data access requests. Liang, Zhao, Shetty, and Li (2017) present a system that leverages a public blockchain to ensure the integrity of data collected by drones and additionally secures the communication between the drone and its control system. Machado and Fröhlich (2018) propose a system that uses blockchain technology to enable the verification of the data integrity of IoT devices. More specifically, they present a proof of concept and evaluate the performance of the implemented data pipeline. While these studies contain detailed descriptions of specific prototypes, they lack both the codifications and the abstractions of the interrelated set of requirements that the system needs to fulfill, as well as the design principles

and features that address these requirements. Both types of research results, however, are necessary to allow for generalizability beyond a specific solution to a specific problem. The importance of such a thorough conceptualization has been extensively discussed among scholars and is a key aspect of DSR (Baskerville & Pries-Heje, 2010; Gregor & Hevner, 2013; Gregor & Jones, 2007; Meth, Mueller, & Maedche, 2015). Therefore, we suggest that the contributions of these existing studies could be expanded substantially by reflecting state-of-the-art DSR guidelines and providing a thorough conceptualization.

Taken together, there is a rich body of knowledge in the IS community on security and privacy. However, scholars have specifically called for studies that develop actionable guidelines to facilitate the design of practical tools. To the best of our knowledge, there are no examples of prior research dedicated to the design and actual implementation of SDPSs. Outside of the IS community, there is an active stream of research focused on the development of SDPSs, describing the technical design of prototypes in detail. However, these studies provide very specific solutions to equally specific problems. Thereby, they lack well-defined conceptualizations and thus generalizable results addressing an entire problem class. Finally, due to the novelty of blockchain technology, there has been a lack of reflection on the specific advantages and limitations of blockchain technology in SDPSs.

3.2.4 Blockchain Technology for SDPSs

A blockchain is a distributed transactional database that is cryptographically secured and controlled by a consensus mechanism (Beck et al., 2017). From an operational perspective, a blockchain comprises an event log storing transactions in such a way that they are immutable once submitted to the system (Moyano & Ross, 2017). Instead of storing the transactions on a central server, various copies of the data exist across different computers, otherwise known as nodes, that participate in the blockchain (Tschorsch & Scheuermann, 2016). This decentralization enables a distributed governance, with a “consensus mechanism between the participating nodes in the system” (Hyvärinen et al., 2017, p. 445), thus eliminating the need to trust other participants of the system (Egelund-Müller et al., 2017; Nakamoto, 2008; Notheisen et al., 2017). Blockchains only accept new entries if they obey a predefined protocol and are thus deemed valid (Nærland et al., 2017; Risius & Spohrer, 2017). Since the introduction of the initial blockchain application Bitcoin in 2009, different forms of distributed ledger technologies, or incarnations of blockchains, have emerged (Lindman et al., 2017; Nakamoto, 2008). Here, we focus on public permissionless blockchains that enable secure transactions in open ecosystems where the participants are not limited to known players, trust is not granted, and all participants are treated equally (Beck, Müller-Bloch, & Ling, 2018). In addition to the generic properties outlined above, this blockchain

type is characterized by a specific set of criteria. The protocols of public permissionless blockchains, such as Ethereum, allow anyone to see any transaction and every node to submit and validate transactions on the blockchain, “thus providing maximum transparency and replicability of transactions” (Hyvärinen et al., 2017, p. 444; Tschorsch & Scheuermann, 2016). Since they are open source, anyone can use these blockchains free of charge and legally (Nærland et al., 2017). In addition, as long as one follows the predefined protocol, there is no gatekeeper limiting access to the blockchain (Beck, Müller-Bloch, & Ling, 2018). Finally, permissionless blockchains are extraordinarily resistant to malicious attempts at manipulation, because the cryptographic logic driving the consensus mechanism and the storage of the transaction log both rely on a decentralized implementation (Gervais et al., 2016). Compared to traditional IS, public permissionless blockchains “avoid the need for copious, often duplicate documentation, third-party intervention, and remediation” (Underwood, 2016, p. 15). Against this background, blockchain technology is often perceived as groundbreaking and is predicted to fundamentally affect how business is conducted (e.g., Chanson, Gjoen, Risius, & Wortmann, 2018; Gomber, Kauffman, Parker, & Weber, 2018), as many industries depend on the fact “that individuals and organizations trust other entities to create, store, and distribute essential records” (Beck et al., 2017, p. 381).

The above-outlined blockchain properties are particularly useful for mitigating issues of data security arising in the IoT and have some decisive advantages over a conventional database system on central servers (Bogner et al., 2016; Glaser, 2017; Hyvärinen et al., 2017). Indeed, there are a variety of different blockchain-based IoT systems currently under development. Well-known examples address use cases in car leasing (Curtis, 2015), pharmaceutical supply chains (Modum, 2018), and energy markets (Meeuw et al., 2018; Mengelkamp et al., 2018). Applying blockchain to IoT use cases has the potential to ensure the “protection of critical infrastructure and data” (Hyvärinen et al., 2017, p. 443). More specifically, blockchains provide tamper-proof storage capabilities in the form of a distributed ledger that can be used to securely store IoT sensor data. In addition, they enable secure ledger access on the basis of well-defined protocols. Finally, blockchain solutions are not operated by one single party (Bogner et al., 2016); hence, they are neutral and particularly suitable in ecosystem settings with multiple parties and potentially diverging interests. However, recent research has often had a view of blockchain technology that is overly optimistic (Beck et al., 2017), and the core blockchain challenges in the field of IoT remain to be solved. First, simply writing IoT sensor data to a public permissionless blockchain is an unacceptable practice in light of the highly sensitive IoT data that is gathered across all areas of our lives (Beck et al., 2017; Lowry et al., 2017). The specific privacy challenges arising in the IoT (see Lowry et al., 2017; Sicari et al., 2015) require adequate countermeasures to ensure the

data privacy of public permissionless blockchain-based IoT systems (Beck et al., 2017; Fabian, Ermakova, & Sander, 2016). Second, public permissionless blockchains are known for their restrictions with respect to scalability as well as for their potentially prohibitive transaction costs (Beck et al., 2016; Risius & Spohrer, 2017). In summary, permissionless blockchain technology is a promising means to mitigate issues of data integrity and availability arising in the IoT. However, some core challenges, such as privacy, scalability, and the potentially prohibitive transaction costs, remain to be addressed.

3.3 Research Study

3.3.1 Methodology

Overall Research Design

We address the problems discussed in Section 3.2 through DSR (Gregor & Jones, 2007; March & Smith, 1995), and we base our specific research approach on the guidelines of Peffers, Tuunanen, Rothenberger, and Chatterjee (2007). Design science has its roots in the seminal work of Herbert Simon (Simon, 1969) and is anchored in many disciplines, such as engineering, architectural science, computer science, and economics (Baskerville, 2008; March & Smith, 1995). Within the IS community, the development of design knowledge is of high significance for both research and practice (Baskerville, 2008; Hevner et al., 2004; Winter, 2008) and continues to attract considerable interest (Baskerville et al., 2015; Gregor & Hevner, 2013; Rai, 2017). The focus of design science is on the creation of the artificial and accordingly the rigorous construction and evaluation of innovative artifacts. It aims to generate new knowledge about a specific and relevant problem class and corresponding solutions to that problem class (Hevner & Chatterjee, 2010). Hence, the creation of utility for practical application through the resulting artifact is one of the core goals of DSR (Hevner et al., 2004; Winter, 2008). While some scholars put their emphasis on the artifact and its relevance (Hevner et al., 2004; March & Smith, 1995), others stress the importance of contributions to theory (Gregor & Jones, 2007; Kuechler & Vaishnavi, 2008; Walls, Widmeyer, & El Sawy, 1992). However, it is widely agreed that impactful DSR arises through synergies between relevance and rigor, that is, the contributions to the application environment as well as to theory (Gregor & Hevner, 2013). We build upon this understanding and elaborate in the following on both the role of theory as well as the general design of the research process.

Concerning the role of theory, we draw on Gregor and Jones (2007), who extend the work of Walls, Widemeyer, and El Sawy (1992) and note that theorizing is a key goal in DSR that may culminate in establishing an IS design theory. On the one hand, existing theory can serve, in the form of kernel theories, as justificatory knowledge and inputs for design cycles (Gregor & Jones, 2007). In particular, the design principles derived from such kernel theories may guide the implementation of an artifact (Walls et al., 1992). On the other hand, design theorizing should contribute to a novel design theory with the aim of formalizing knowledge in DSR (Gregor & Hevner, 2013; Gregor & Jones, 2007). This type of theory provides instructions that link design principles and features with actions. It is prescriptive in the sense that it provides rules and actionable guidelines and hence belongs to the theories of type five in Gregor's taxonomy (Gregor, 2006; Gregor & Hevner, 2013). Communicating such a design theory can be enabled by an artifact instantiation that embodies the related design principles and features (Gregor & Jones, 2007). An ex-post evaluation, in which additional slices of data are gathered after the original design cycles and the corresponding evaluations and are then used in an evaluation process to generate further theoretical insight, can be an important and constructive step to reach a sufficient abstraction level and theoretical saturation (Beck, Weber, & Gregory, 2013).

Concerning the general design of the research process, there is wide agreement that an iterative procedure of well-defined steps is most applicable for DSR (Hevner et al., 2004; Nunamaker Jr, Chen, & Purdin, 1990; Takeda, Veerkamp, & Yoshikawa, 1990). Since the recognition of DSR in the mainstream of IS with the publication of Hevner et al.'s article (2004), the discourse within the IS community has been intense and ongoing regarding the specific structuring of this process. Many different approaches and improvements and derivatives thereof have been suggested by renowned scholars (Beck et al., 2013; Hevner, 2007; Peffers et al., 2007; Vaishnavi & Kuechler, 2015). Our project's research design is based on the guidelines of Peffers et al. (2007) and informed by the design approach of Meth et al. (2015). We extend Peffers et al.'s guidelines by considering an additional phase of ex-post evaluation (Pries-Heje, Baskerville, & Venable, 2008; Venable, Pries-Heje, & Baskerville, 2016) after finalizing the prototype, as suggested by Beck et al. (2013), which facilitates the generation of additional insight. Finally, to summarize the knowledge gathered, we follow Gregor and Jones (2007) and present our results in the form of a design theory.

Design Cycles

Based on the theoretical and procedural reflections above, we design our research project in three design cycles, each composed of five phases, which are followed by two final steps of evaluation and communication. This research design, the output of each phase, and the according

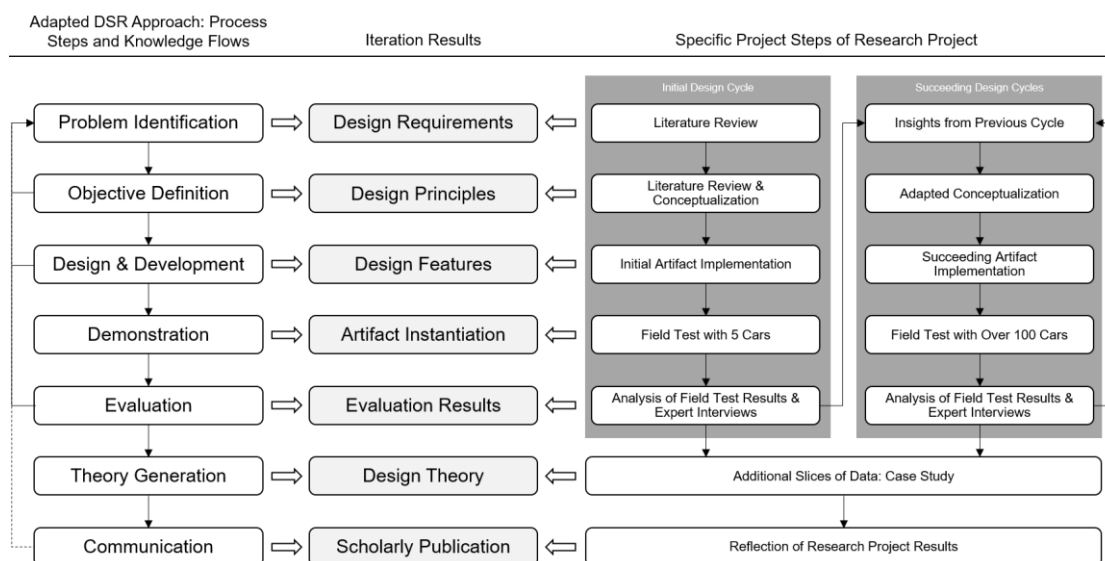


Figure 5. Design cycles based on Peffers et al. (2007), Beck et al. (2013), and Meth et al. (2015)

iteration between conceptualization, development, instantiation, and evaluation, is outlined in Figure 5. The first design cycle was initiated with an intensive literature review to identify the problem at hand and reflect on RQ1a. Our examination of the topic was triggered by a report of the prevalence of odometer fraud (TÜV Rheinland, 2015). Developing systems that are able to securely process and exchange odometer sensor data arose as a main challenge in this study. Our literature review quickly expanded to similar issues regarding IoT sensor data present in other industries, such as pharma (Modum, 2018) and energy (Mengelkamp et al., 2018). This initial literature review allowed us to develop the first preliminary requirements for the artifact to be built. We then conducted a second literature review to find reference points in theory and the extant body of knowledge to refine these preliminary requirements, deepening the findings concerning RQ1a. Based on this, we then derived design principles in the objective definition phase and identified the design features that are required to address these design principles, hence addressing RQ1b. All these steps focused on the generalized problem class. In the next step, we instantiated the developed design with respect to a specific use case (prevention of odometer fraud) and developed the first version of our prototype CertifiCar. We evaluated this initial version of CertifiCar in a field test with five cars as well as on the basis of expert interviews. We used the results of this evaluation to adapt the artifact design in the second design cycle and, based on these changes, implemented a new version of our artifact. Again, we evaluated the artifact in a field test and on the basis of expert feedback. We integrated these findings into the third design cycle, which was run similarly to the second design cycle and resulted in the final version of the artifact. The final version of CertifiCar was deployed in a field test with 100 cars, and the subsequent

evaluation was based on the results of this field test and expert interviews. During these loops of development and evaluation, we iteratively refined the design requirements, principles, and features, enhancing the results to RQ1a and RQ1b. Furthermore, the knowledge acquired in this phase built the foundation to approach RQ1c. Ultimately, we gathered additional slices of data for a detailed ex-post evaluation of the derived design requirements, principles, and features of the artifact (Beck et al., 2013; Pries-Heje et al., 2008). This helped to confirm the validity of our responses to RQ1a, RQ1b, and RQ1c and led to diverse additional insights into RQ1c.

In our conceptualization efforts, we follow three core design steps to derive the design requirements, principles, and features (Hevner & Chatterjee, 2010; March & Smith, 1995). In the first step, we develop design requirements based on the input from the problem identification step. The design requirements are generic requirements that should be met by any artifact aiming to create a solution for the underlying problem class. This notion of design requirements is closely related to the meta-requirements described by Walls et al. (1992) and the general requirements introduced by Baskerville and Pries-Heje (2010). In the second design step, we identify design principles based on the input of the suggestion step, for instance, by drawing on the extant information asymmetry literature. Our concept of design principles corresponds to the generic capabilities of an artifact through which the design requirements are addressed and relates these requirements indirectly with design features containing the technical specifics of the solution. This notion of design principles is closely linked to the meta-design introduced by Walls et al. (1992) and the relationship between general requirements and general components that Baskerville and Pries-Heje (2010) emphasize. In the third step, we derive design features on the basis of the design principles and implement them in an instantiation of the artifact. These design features capture the technical specifics of the solution and are closely related to the general components described by Baskerville and Pries-Heje (2010). A design principle that is instantiated by an explicit design feature can be understood as an explanation (design principle) of why a specified piece (design feature) leads to a predefined goal (design requirement) (Kuechler & Vaishnavi, 2012). These explanations will assist us in abstracting the results of the instantiation of our prototype (CertifiCar) to a more generalized level and in creating a better understanding of the conceptual foundation of the design theory we propose.

As we reported above, we attempted to ensure the appropriate grounding and viability of the proposed design and its corresponding artifact instantiation in multiple iterations of our research design. Thereby, we distinguish between the interim evaluations at the end of each design cycle and the ultimate ex-post evaluation after finalizing the artifact development. In practice, in each design cycle, we use the last two phases to demonstrate and evaluate the current instantiation of the prototype, as the guidelines of Peffers et al. (2007) suggest. This procedure is detailed in

Section 5, where we depict the iterative development of the prototype and the corresponding demonstrations and evaluations. Subsequently, we perform an additional ex-post evaluation (Pries-Heje et al., 2008), as suggested by Beck et al. (2013), to facilitate the generation of a novel theory. Specifically, we perform semi-structured interviews with nine experts on different security and privacy topics regarding IoT data to generalize and verify the viability of our proposed actionable guidelines, resulting in our final design theory. We only briefly discuss the interim evaluations and emphasize the ex-post evaluation because it focuses on the generalized problem class defined by the design requirements derived and, contrary to the interim evaluations, not on the specifics of the prototype implemented in this study.

3.3.2 Designing an IoT Sensor Data Protection System

Developing Design Requirements

To derive the specific design requirements for an SDPS that enables the process of IoT sensor data generation, processing, and exchange, we built upon practically motivated problems that are outlined in the existing literature. More specifically, as outlined in the foundations section, studies of interest include the following: (1) research regarding the IoT and sensor data (core key words: Internet of Things, IoT, cyber-physical systems, sensor data, big data, digital and Digitization³), (2) research regarding security and privacy (core key words: protection, security, secure, privacy, private, privacy-preserving, data, information and system¹), and (3) specific research focusing on systems that protect sensor data (core key words: Internet of Things, IoT, cyber-physical systems, sensor data, security, cybersecurity, attack, protection, privacy, private and privacy-preserving¹). To consolidate the existing research, we considered prestigious IS journals (i.e., the AIS basket of journals), international IS conferences (AMCIS, ECIS, ICIS, MCIS, PACIS), and high-quality journals with a specific focus on practical relevance (the Harvard Business Review, MIS Quarterly Executive, and MIT Sloan Management Review). Additional IS outlets were considered through the AIS eLibrary. With respect to research focused on systems that protect sensor data, we included the ACM Digital Library, as well as the IEEE Xplore Digital Library. Finally, we conducted a backward and forward search based on the gathered literature (Webster & Watson, 2002).

A core challenge in IoT is security and data manipulation (Lowry et al., 2017). The IoT creates new security challenges, for instance, that the data collection nodes are typically left unattended for long periods of time (C. C. Aggarwal et al., 2013; Ronen et al., 2017). In addition,

³ Using respective combinations

a data recipient cannot be sure if the received data is valid, because a malicious adversary, potentially the data owner himself, has the possibility to manipulate the data at several stages in the data pipeline (C. C. Aggarwal et al., 2013). Additional problems are introduced by the fact that the progress in deploying and developing the IoT is much faster than the accompanying security practices (Singh et al., 2016). Therefore, a recipient of IoT sensor data often encounters the problem that the data integrity cannot be taken for granted (Miorandi, Sicari, De Pellegrini, & Chlamtac, 2012; Sicari et al., 2015). Consequently, we derive the following design requirement:

DR1: Enable tamper-resistant data generation, processing, and exchange. *The process of IoT sensor data generation, processing, and exchange should be supported by systems that ensure tamper resistance throughout the whole data pipeline.*

A second challenge in the realm of IoT sensor data is privacy (Lee et al., 2018; Sicari et al., 2015). More specifically, there is a lack of well-established privacy-preserving mechanisms (Bélanger & Crossler, 2011). This is especially striking because IoT sensors often have access to very detailed personal data (Lowry et al., 2017). In addition, users are often not able to determine which data is recorded and transmitted (Davenport, 2013; Westin, 1967). Home assistance devices, such as Amazon Alexa and Google Home, are always on, although most of the time they are neither supposed to store nor transmit recorded information. Similar thoughts apply to other devices deployed inside the home of a user. Therefore, the goal of any data processing system in the realm of IoT is to preserve privacy (Alqassem & Svetinovic, 2014; Sicari et al., 2016). Consequently, we derive the following design requirement:

DR2: Enable privacy-preserving data generation, processing, and exchange. *The process of IoT sensor data generation, processing, and exchange should be supported by systems that are capable of preserving the privacy of the corresponding data owner.*

A third challenge is related to IoT and big data. As we have outlined, the technical transformation of information processing from analog to digital and the according merger of the physical and digital worlds are expected to generate unprecedented amounts of data (Lowry et al., 2017; Porter & Heppelmann, 2015). Hence, systems that enable tamper-resistant data generation and exchange must be able to cope with “big data” (H. Chen et al., 2012). To operate in such a context, a corresponding system should have sufficient throughput to handle the expected amounts of data the IoT will generate. This aspect becomes particularly relevant when using blockchain technology, as many of the existing blockchain technologies are still struggling with scalability problems (Hyvärinen et al., 2017; Tschorsch & Scheuermann, 2016). Consequently, we derive the following design requirement:

DR3: Enable large data volume throughput. *The process of IoT sensor data generation, processing, and exchange should be supported by systems that are capable of processing the large amounts of data that are typical of IoT applications.*

Finally, the advantages of IS must always be weighed against their disadvantages (Delone & McLean, 2003). In light of this fundamental economic principle, the IS-related costs are of particular importance in a business environment. Although this holds true for any IS, it is of special importance for solutions that rely on blockchain technology (Risius & Spohrer, 2017). As discussed above, the currently unsolved issues regarding the scalability of different blockchain technologies and high transaction costs have the potential to generate substantial financial expenditures (Beck et al., 2016; Hyvärinen et al., 2017). Consequently, we derive the following design requirement:

DR4: Ensure economic feasibility. *The process of IoT sensor data generation, processing, and exchange should be supported by systems that ensure economic feasibility.*

Summing up, based on the fundamental SDPS challenges, we derived four general design requirements (see Table 2). These design requirements determine our design theory's purpose and scope that the design principles and design features must address to overcome or reduce the existing challenges (see Figure 6).

ID	SDPS challenge	SDPS design requirement	Main corresponding literature
1	Adversaries have the possibility to manipulate sensor data at several stages in the processing pipeline, so data integrity cannot be taken for granted.	SDPS should ensure tamper resistance throughout the whole data pipeline.	(C. C. Aggarwal et al., 2013; Lowry et al., 2017; Sicari et al., 2015)
2	IoT sensors can capture detailed and very sensitive personal data.	SDPS should be capable of preserving the privacy of the data owner.	(Bélanger & Crossler, 2011; Davenport, 2013; Lee et al., 2018; Sicari et al., 2016)
3	IoT sensors are able to generate vast amounts of data.	SDPS should provide sufficient data throughput to process large amounts of data.	(H. Chen et al., 2012; Hyvärinen et al., 2017; Porter & Heppelmann, 2015)
4	The protection of IoT sensor data can require substantial resources and induce significant costs.	SDPS should ensure economic feasibility, that is, the protection benefits have to outweigh the protection costs.	(Beck et al., 2016; Hyvärinen et al., 2017; Risius & Spohrer, 2017)

Table 2: General SDPS challenges and design requirements

Deriving Design Principles

To address the design requirements, we build upon theory and the existing body of knowledge to derive design principles. With respect to DR1 (tamper-resistant data generation, processing, and exchange), theory on information asymmetry provides a fruitful basis to derive design principles. The (neo-)classical market model suggests that participants are fully informed about all goods (Albersmeier, Schulze, Jahn, & Spiller, 2009). However, business transactions are often characterized by fundamental information deficits (information asymmetries) that favor opportunistic behavior and restrict the smooth functioning of markets (Akerlof, 1970; Spence, 1976). To overcome these information deficits and avoid opportunistic behavior, certain measures such as certification, guarantees, or well-established brand names have been identified (Akerlof, 1970; Bond, 1982; Genesove, 1993).

With regard to the protection of sensor data, certification, in particular, appears to be a suitable measure to prevent opportunistic behavior (manipulation), as it is not restricted to companies that have high credibility or a strong brand name. Certification indicates the attainment of a certain quality level and is based on auditing (Akerlof, 1970). It most often relies on protection and investigation schemes that cover the whole supply (e.g., food business) chain or information (e.g., financial auditing) chain, as certain product and information qualities cannot be judged by inspections that are limited to the end of the chain (Albersmeier et al., 2009). This is particularly relevant for sensor data. Only in the case of very obvious manipulations is it possible to detect manipulated sensor data by means of a single inspection at a certain point in the information processing chain (e.g., when the odometer value of a car is equal to or even smaller than zero). Hence, the entire information chain from source (sensor) to sink (final data consumer) must be protected from manipulation, e.g., by applying an appropriate means of encryption. By protecting the data along the entire information chain, it can be certified that the data was not manipulated on the way from the source to the sink.

DP1: Sensor data is certified on the basis of source to sink protection.

If data is protected from source to sink, data producers can be made accountable for the data they provide. However, in the case of sensor data, even if the information chain is protected from source to sink, data manipulation can still occur. More specifically, the data producer can focus on the source and manipulate the sensor or its environment. For example, anecdotal evidence and a corresponding patent⁴ suggest that temperature sensors in cold chains are regularly covered with

⁴ <https://patents.google.com/patent/DE10228648A1/de>

insulation material to hide shorter periods of irregularities. In cars, as a second example, mileage sensors (odometers) are multi-component systems that are connected by cables so that manipulating devices (“CAN filters”, “CAN blockers”) can be placed between them. More specifically, small sensing units often do not have the computing power for encryption or processing and hence communicate their raw sensor values to more powerful control units over wires that can be intercepted. Therefore, sensors are not per se monolithic components that are well protected and cannot be manipulated. To account for the corresponding manipulation risk, additional means might be required to enable trustworthy certification. More specifically, cross-validation and plausibility checks are common means in auditing (Whittington & Pany, 2015) that might also be used with sensor data to reveal manipulations. In the case of car mileage manipulation, for example, GPS data can be used to cross-validate the mileage data of a car.

DP2: Sensor data is certified on the basis of cross-validation.

However, cross-validation and plausibility checks can only reduce the manipulation risk. Similar to financial auditing, a “detection risk” (Dong, Liao, & Zhang, 2018; Hogan & Wilkins, 2008) remains, which depicts the probability that manipulations are not detected. In summary, with the implementation of DP1, it can be certified that data was not manipulated on its way from the source to the sink, so that data producers can be made accountable for the data they provide. In addition, with the availability of cross-validation data and the implementation of DP2, it can be certified with an associated detection risk that the sensor or its environment were not manipulated.

With regard to DR2 (privacy-preserving data generation, processing, and exchange), we build upon Westin’s (1967) theory of privacy to derive a corresponding design principle. Westin’s theory is one of the best articulated and best supported theories of privacy (Margulis, 2011). A fundamental cornerstone of Westin’s theory is the existence of the following four states of privacy (Margulis, 2011): (1) solitude is about being free from observation by others, (2) intimacy is about the seclusion required to form close associations, (3) anonymity is about the condition of being unknown, and (4) reserve is about limiting disclosure to others. In essence, for Westin (1967, p. 7), privacy “is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others”.

At the core of Westin’s definition is the right of a data owner to have full control over the communication and use of her data. With respect to the exchange of sensor data, the data owner should therefore determine when and to what extent data is communicated and to whom. However, the means of exchange that determines the “how” is a software system. Hence, the data owner is limited in her privacy by the restrictions of the system. If the system restricts privacy too

much, though, the data owner still has the option to not use the system. In summary, in the context of sensor data exchange, we derive the following design principle that addresses DR2:

DP3: Data owners determine when and to what extent their certified data is communicated to others.

While the first two design requirements focus on what a sensor data exchange system should enable (tamper resistance and privacy), DR3 (large data volume throughput) and DR4 (economic feasibility) further qualify how the system should operate (scalable and thereby also cost efficient) and shift the focus from positive system outcomes (prevent manipulation, assure privacy) to possible negative outcomes (system costs). The existence of such positive and negative system outcomes is well reflected in IS theory. The DeLone and McLean Model of Information Systems Success captures the idea that the system impact has to reflect the balance of positive and negative impacts (Delone & McLean, 2003). The concept of “net benefits” depicts the rationale that “no outcome is wholly positive, without any negative consequences” (Delone & McLean, 2003, p. 22).

Applying the aforementioned rationale to tamper-resistant sensor data exchange, a potential solution has to ensure that the positive effects are not cancelled out by negative consequences. In respect to the design challenge at hand, the protection and certification of IoT sensor data can be resource-intensive and costly, especially in the context of large amounts of sensor data (Sicari et al., 2015). Hence, data has to be processed on a system architecture that is linearly scalable with respect to performance and costs. Thereby, the scalability captures how “well a particular solution fits a problem as the scope of that problem increases” (Schlossnagle, 2006). Linear scalability is an established concept that refers to the relationship between an input (e.g., amount of sensor data) and an output (e.g., performance or cost) (Bonvin, 2012). While the term defines a very specific type of relationship (linear), it is often used in a broader sense. In contrast to negative or sublinear scalability (Williams & Smith, 2004), linear scalability depicts the idea that the performance does not erode and the costs of a system do not explode at scale.

DP4: Data is certified on the basis of a linearly scalable system architecture.

Mapping Design Principles to Design Features

In the last step of the conceptualization, we map the identified design principles to design features. As we elaborated above, the design features are specific artifact capabilities designed to fulfil the design principles derived previously (Meth et al., 2015). An overview of these features, including the design principles and design requirements that we derived, is shown in Figure 6. The design features that we describe build on the fundamental premises (see Section 3.2.4) that

(1) permissionless blockchain technology is a fruitful means to address the issues of data security and privacy arising in the IoT and (2) the limitations of the existing blockchain technology, with respect to privacy, scalability, and costs, have to be addressed appropriately. In the following discussion, we introduce the design features along with three fundamental system capabilities, namely, capture data, store data, and provide data.

To implement the first design principle, that is, certify that the data was not manipulated on the way from the source to the sink, two features are needed. First, we have to collect the data (DF1) and, second, we need to preprocess the data in a way that prevents data manipulation from this point on (DF3). To achieve this, we follow existing practices (Ayoade et al., 2018; Nærland et al., 2017) and save only the hash of the data (i.e., the “digital fingerprint” of the data) in a public permissionless blockchain. We can later use this hash to check that the data, which is stored in raw format in a traditional database, has not changed by other parties since the transaction was signed. As only changes after the signature can be detected by this approach, it is essential to choose the earliest possible point in the data pipeline to create this signature and swiftly add the transaction to a blockchain.

The second design principle of cross-validation-based certification calls for two additional design features, namely, the collection of appropriate validation data (DF2) and a certification mechanism that performs the cross-validation (DF8). In the case of car mileage data, for example, GPS data can be collected for validation purposes in addition to odometer values. The GPS data can then be used to calculate the mileage data, which can be compared to the mileage values received from the odometer sensor.

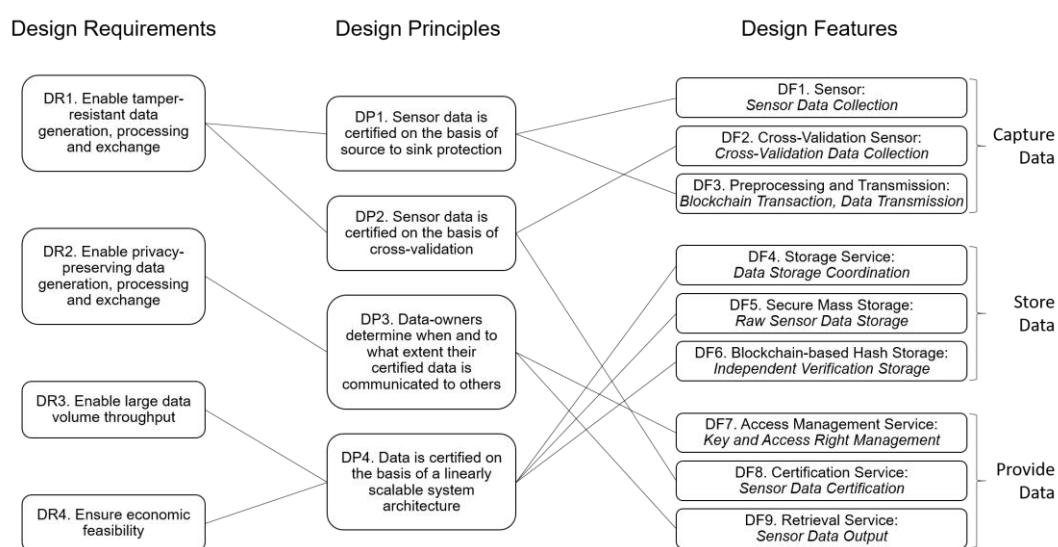


Figure 6. Design requirements, principles, and features

The third design principle postulates that data owners determine when and to what extent their data is communicated to others, which results in the implementation of two design features, namely, an access management service (DF7) and a data retrieval service (DF9). The access management service ensures that the raw data, which is stored in an encrypted form in a centralized mass storage, can only be decrypted by the owner of the data. The data retrieval service is implemented in such a way that, in accordance with the access management settings, only selected parts of the whole raw data can be transferred to the data-requesting party. Hence, in the odometer example, the data owner has the possibility to choose between only sharing the last odometer value or providing the full history of odometer values, e.g., in the form of a daily, weekly, or monthly history.

The fourth design principle, requiring a linearly scalable system architecture, needs three more design features, namely, a storage service (DF4) that writes into the raw data storage (DF5) as well as into an independent verification storage (DF6). In practice, the storage service saves the encrypted raw data in a cloud storage and propagates the signed transaction with the hash to the blockchain network. By implementing these three features, a “hybrid architecture” that addresses a central challenge of public permissionless blockchain technology is realized. It is well known that certain public permissionless blockchain technologies have severe technical and economic scalability issues, so that dedicated approaches have to be applied (Beck et al., 2016; Notheisen et al., 2017; Risius & Spohrer, 2017). More specifically, hybrid architectures that build upon blockchain-based “on-chain” transactions and non-blockchain-based “off-chain” transactions are known to cope with large amounts of data while preserving the key characteristics of distributed blockchain systems (Zyskind et al., 2015). In hybrid architectures, not all data is made available on a fully distributed blockchain. Instead, some data is stored centrally or shared only by a selected number of nodes. However, to enable trust and prevent manipulation, off-chain data has to be linked to on-chain transactions. In the case of IoT sensor data, sensor values can be stored in a central repository, and only the digital fingerprint (hash) of one or multiple records is recorded on-chain. Thereby, the data stored in the blockchain can dramatically be reduced while still ensuring data integrity.

To summarize, Figure 7 presents a general architecture for an SDPS, including all of the design features introduced above. In practice, different instantiations of this architecture are possible. In some cases, for example, the collected data itself might not be privacy-relevant, and hence a selected retrieval thereof would not be necessary (DF7, DF9). In other cases, the validation data might be publicly available or even has to be gathered manually by inspections, which would replace the validation sensor (DF2).

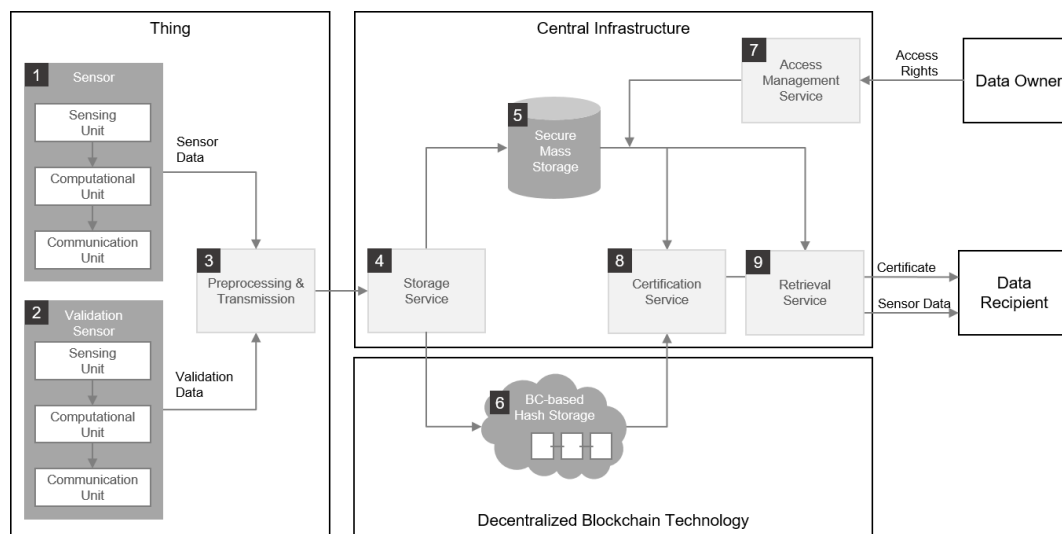


Figure 7. Artifact architecture

The architecture highlights that a sensor is not necessarily a monolithic component and that the data, which is preprocessed and incorporated in a blockchain transaction as a hash (DF3), has already been processed through several steps as follows: it is recorded by a sensing unit, then processed by a computational unit into meaningful information, and finally communicated to a receiver outside of the sensor through a communication unit. Hence, there are several attack vectors between the sensing unit of a sensor and the location where the blockchain transaction is actually signed. One important goal, therefore, is to build and sign the blockchain transaction (DF3) as close as possible to the sensing unit. In the future, one could imagine blockchain-enabled hardware that combine sensing and transaction management within one chip, similar to current hardware security modules. This would significantly reduce the attack vectors and ease the implementation of DP1. Storing only a hash in the blockchain supports several goals, in addition to the main objective of guaranteeing the immutability of the stored data. As opposed to storing the raw data in the blockchain, using only a hash additionally prohibits other participants from gaining useful, potentially privacy-related information, as the blockchain is public and accessible for everyone (DP3). Furthermore, the hash serves to reduce the amount of data that needs to be stored in the blockchain and therefore supports the scalability of the solution (DP4).

The details of the certification mechanism and its individual steps (DF8) are outlined in Figure 8. The process is initiated by the owner of the data by granting access (DF7). The raw dataset is decrypted and sent to the unit responsible for the certification (8.a). In the next step, the hashes of each raw data package (hashes can be calculated on the basis of single or multiple values) are calculated and stored (8.b). In parallel, for each raw data package, the corresponding transaction is looked up in the blockchain (8.c), and the saved hash is extracted (8.d). Then, the

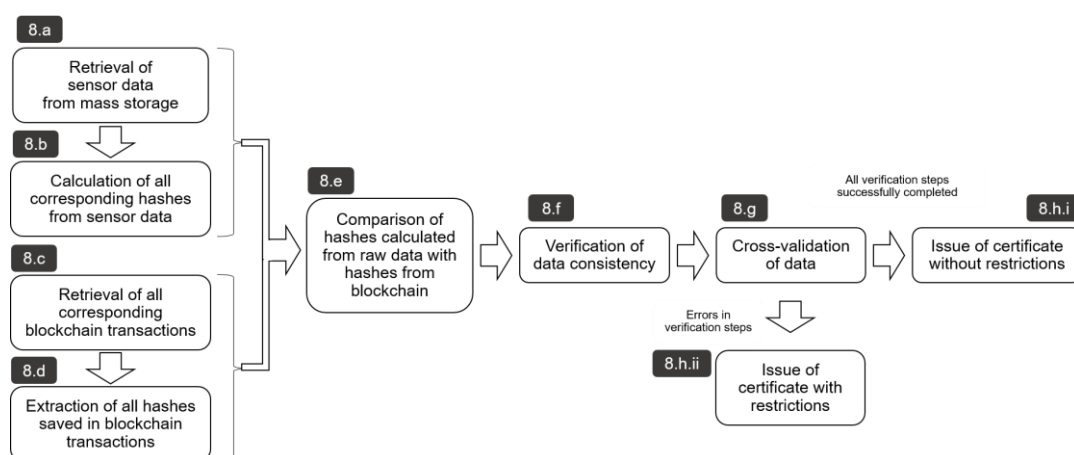


Figure 8. A detailed view of the certification process

algorithm compares the hashes calculated from the raw data with those retrieved from the blockchain (8.e). A match proves that the data package in question was not changed since the signature of the corresponding blockchain transaction. Hence, the data was not manipulated on its way through the processing pipeline, and the data owner can be made accountable for the data. Any mismatches are noted and inserted as warnings in the final certificate. In the next step, the data consistency is verified (8.f). Here, the verification logic depends on given domain rules and constraints. In the case of mileage data, for example, verification can rely on the simple fact that the odometer value increases with every trip; a decrease in mileage is thus a clear indicator of an irregularity or manipulation. Typically, the more interdependent the sensor values that are recorded, the more sophisticated the tests are that can be applied. In the final step of the verification, the validation data can be leveraged (8.g). In the case of odometer fraud, the increase in the mileage of a trip should, for example, be larger or equal to the shortest distance between the GPS coordinates of the start and the end of the trip (data which is available in connected cars as of today). Finally, the certificate is issued, either without restrictions, if all verification steps were passed successfully (8.h.i), or with restrictions and a detailed report on the issues (8.h.ii).

3.3.3 Iterative Development of the Prototype

One of the core goals of DSR is to create utility for practitioners. To succeed in this task, practitioners have to understand how to apply the abstract guidelines developed in the DSR process. As the implementation of such abstract guidelines is inherently ambiguous, scholars recommend describing the implementations of these guidelines, including the corresponding context, in detail and positioning the artifact in a natural setting, thus rendering these guidelines actionable (Baskerville, 2008; Chandra Kruse, Seidel, & Gregor, 2015; Chandra Kruse, Seidel, & Puro, 2016). Additionally, these descriptions enable researchers to establish the instantiation

validity of the implementation by showing how the abstract guidelines can be linked to specific features of an artifact (Lukyanenko, Evermann, & Parsons, 2015). Hence, in the following, we present the iterative problem solving process used to design and develop our prototype CertifiCar.

The aim of our prototype is to prevent odometer fraud. Odometer fraud prevention is a relevant IoT use case in which the integrity of data is of high value and privacy is desirable. Odometer fraud, i.e., the fraudulent manipulation of a car's mileage records, is a huge problem in many countries, which is why numerous governments, for example, in Belgium, New Zealand, and the USA, have fostered the creation of systems that impede manipulation, with according legal policies (Car-Pass, 2018; Carfax, 2018; CarJam, 2018). Germany is one of the largest car markets without a centralized prevention system, and it is estimated that odometer fraud in Germany affects one third of all resold cars, leading to an annual damage of almost 6 billion euros (TÜV Rheinland, 2015). Usually, odometer fraud is committed to increase a car's value by reducing the mileage. The procedure is extremely simple and inexpensive and can be performed within minutes. Detailed step-by-step instructions are available on YouTube, and corresponding devices can be ordered online for less than 100 euros.

The existing systems that fight odometer fraud, such as Carfax (USA) and CarJam (NZ), have several substantial challenges. They are not able to detect odometer fraud reliably, have severe privacy issues, and cannot support cross-country transactions. More specifically, new records are only captured occasionally, and the interval between two records can span months or even years, giving rise to considerable fraud potential. In addition, there is no cross-validation. This makes it very difficult to detect odometer fraud. Moreover, continuous odometer fraud enabled by specific hardware manipulation devices within the car cannot be detected at all. Finally, sensitive data is stored in central databases accessible to the public, and data acquisition is limited to the country of the respective service provider. The privacy problems in the approaches of the existing systems prohibit their application in countries with strict privacy laws, such as Germany.

Iteration 1: End-to-End Processing and Initial Verification

An overview of the prototype architecture in its final state is displayed in Figure 9. In the first iteration, we implemented an initial version of the end-to-end data pipeline. This included the recording of the odometer data in the car (DF1), the processing of this data in the application (DF3, DF4), and the subsequent storing of the encrypted raw data and hashes in the private cloud storage (DF5) and the blockchain (DF6), respectively.

We chose the Ethereum blockchain because it offered the best development support and a vibrant ecosystem at the time of the development of the prototype in the beginning of 2017

(Buterin, 2013). As a proof of principle, we used the public Ethereum blockchain for a set of transactions. In addition to individual sample transactions on the Ethereum main network, we set up a private instance of the Ethereum blockchain exclusively for the prototype, which was operated and used. The system has been in operation for over a year with only short interruptions. Additionally, a first version of the verification process (DF8) was implemented. This ensured that all data points were protected by a corresponding hash in the blockchain and were not manipulated (DF8.e). The verification process investigated if the mileage did not decrease at any point in time (DF8.f). To interact with the system seamlessly, a web-based user interface was added. We chose to record the data points on the trip level to ensure reasonable transaction costs while guaranteeing a resolution high enough to detect fraud reliably.

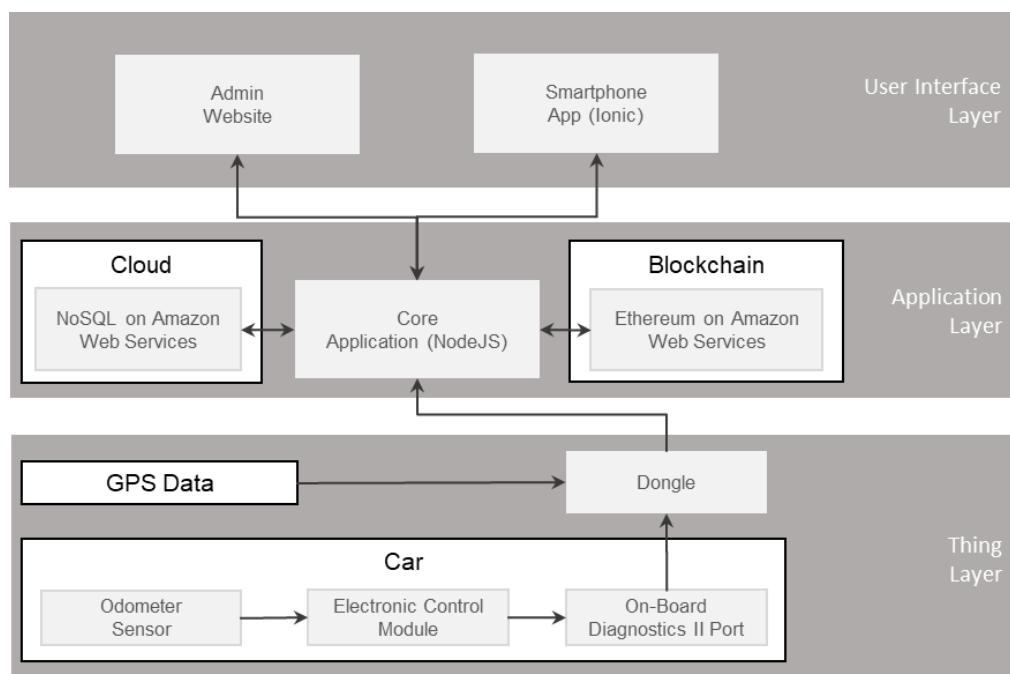


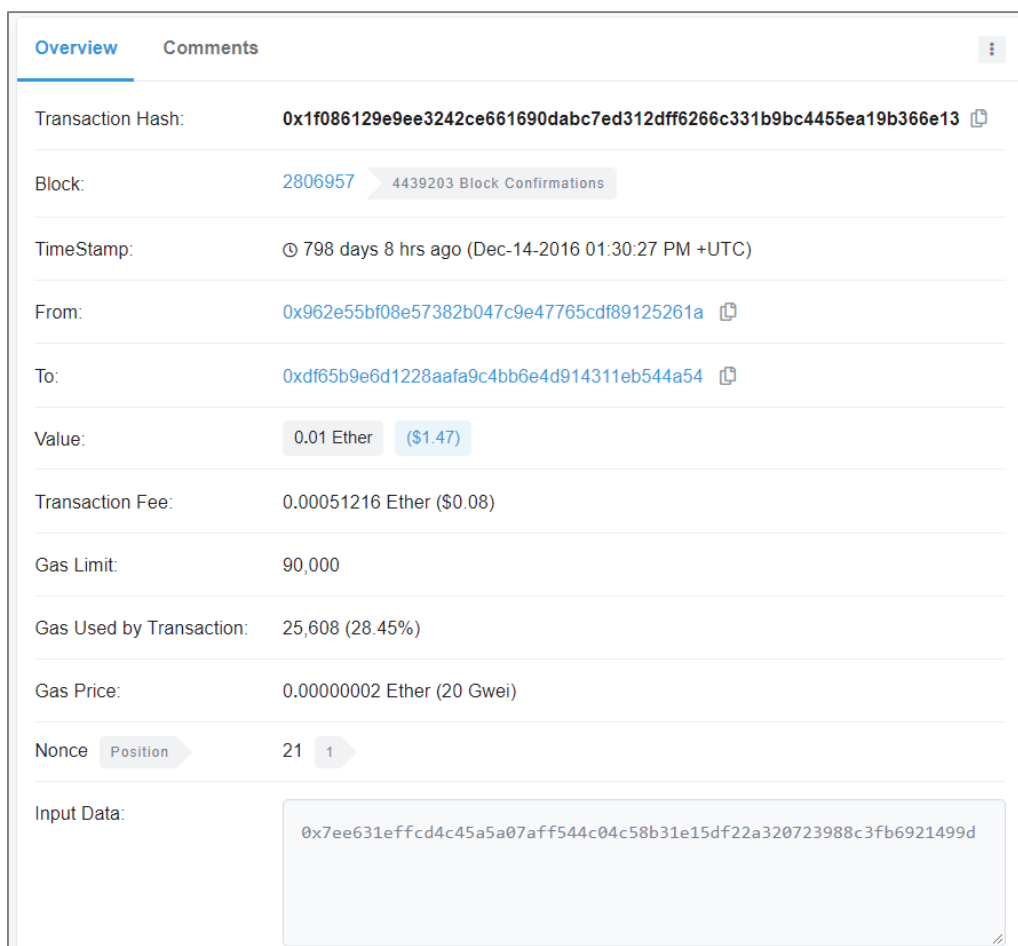
Figure 9. Prototype architecture

```
> web3.eth.getTransactionFromBlock(2806957, 1)
{
  blockHash: "0x10a31a62aead015458c0afb80e341a155b1f68785d10dd249d7ed9bda3a8d231",
  blockNumber: 2806957,
  from: "0x962e55bf08e57382b047c9e47765cdf89125261a",
  gas: 90000,
  gasPrice: 2000000000,
  hash: "0x1f086129e9ee3242ce661690dabc7ed312dff6266c331b9bc445ea19b366e13",
  input: "0x003700650065003600330031006500660066006300640034006300340035006100350061003000370061006600
660035003400340063003000340063003500380062003300310065003100350064006600320032006100330032003000370032
0033003900380038006300330066006200360039003200310034003900390064",
  nonce: 21,
  r: "0x227f73b1afaf96697abbad74e465b43d3ee63e00bfc8a3da51992cc42a51705f",
  s: "0x6cbfa5161d650f81159b28f8eb3bdebd0bb155783bbd8bfa5724c68046bfd9b",
  to: "0xdf65b9e6d1228aafa9c4bb6e4d914311eb544a54",
  transactionIndex: 1,
  v: "0x26",
  value: 1000000000000000
}
```

Figure 10. Sample Ethereum transaction

Every iteration of the creative and heuristic design as a search process should generate a representation of the artifact that is being demonstrated and evaluated (Hevner et al., 2004; Peffers et al., 2007). We tested this iteration with five cars that were driven daily for several hours for two weeks. This showed that the prototype was running without any major issues.

A sample Ethereum transaction of the prototype at this stage, written to the public blockchain, is shown in Figure 10 and Figure 11. Figure 10 shows the view of the transaction in the JavaScript command line interface of the geth client, the official Go implementation of the Ethereum protocol. Please note that the hash of the sensor data is labeled “input”, while the value depicted as “hash” is the hash value of the overall blockchain transaction. In the example at hand, the corresponding transaction is the first transaction (“transactionIndex”) in the 2,806,957th block (“blockNumber”). To prove that the transaction was submitted to the Ethereum mainnet, Figure 11 shows a view from etherscan.io, where one can recognize the stored hash (“Input Data”).



The screenshot displays the transaction details for block 2806957 on etherscan.io. The interface includes tabs for 'Overview' and 'Comments'. The transaction hash is 0x1f086129e9ee3242ce661690dabc7ed312dff6266c331b9bc4455ea19b366e13. The block number is 2806957, with 4439203 block confirmations. The transaction timestamp is 798 days 8 hours ago (Dec-14-2016 01:30:27 PM +UTC). The sender (From) is 0x962e55bf08e57382b047c9e47765cdf89125261a, and the recipient (To) is 0xdf65b9e6d1228aafa9c4bb6e4d914311eb544a54. The transaction value is 0.01 Ether (\$1.47), with a transaction fee of 0.00051216 Ether (\$0.08). The gas limit is 90,000, and the gas used by the transaction is 25,608 (28.45%). The gas price is 0.00000002 Ether (20 Gwei). The nonce is 21, and the position is 1. The input data is 0x7ee631effcd4c45a5a07aff544c04c58b31e15df22a320723988c3fb6921499d.

Transaction Hash:	0x1f086129e9ee3242ce661690dabc7ed312dff6266c331b9bc4455ea19b366e13
Block:	2806957 4439203 Block Confirmations
TimeStamp:	🕒 798 days 8 hrs ago (Dec-14-2016 01:30:27 PM +UTC)
From:	0x962e55bf08e57382b047c9e47765cdf89125261a
To:	0xdf65b9e6d1228aafa9c4bb6e4d914311eb544a54
Value:	0.01 Ether (\$1.47)
Transaction Fee:	0.00051216 Ether (\$0.08)
Gas Limit:	90,000
Gas Used by Transaction:	25,608 (28.45%)
Gas Price:	0.00000002 Ether (20 Gwei)
Nonce	21
Position	1
Input Data:	0x7ee631effcd4c45a5a07aff544c04c58b31e15df22a320723988c3fb6921499d

Figure 11. Transaction of Figure 10 on etherscan.io

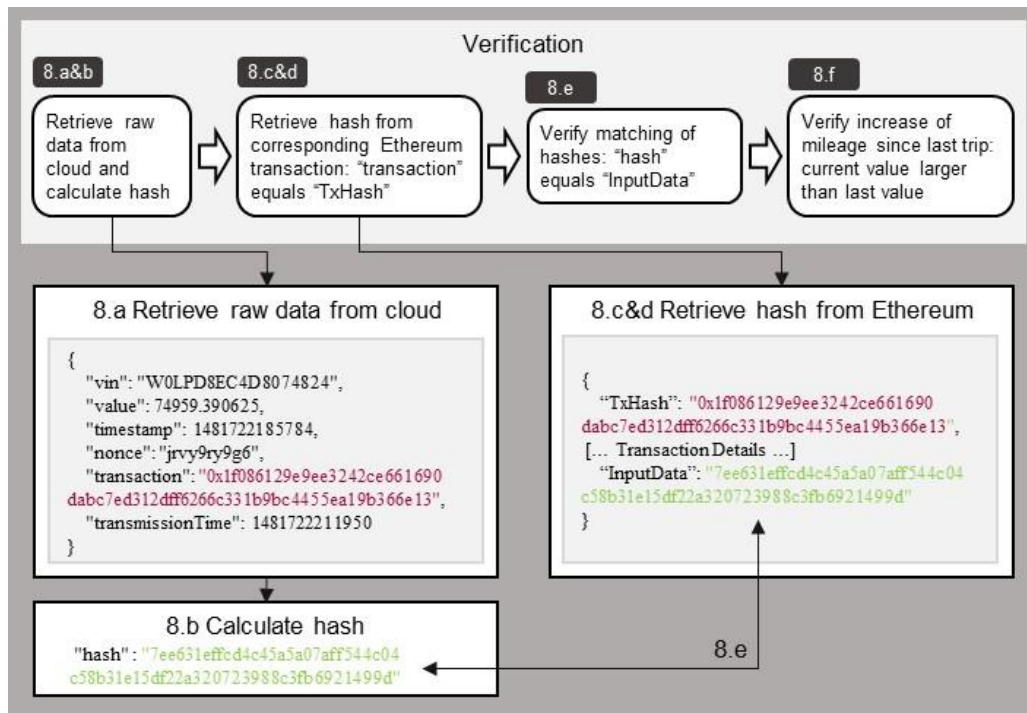


Figure 12. First implementation of the verification process

For this Ethereum transaction, Figure 12 shows the verification process at that point in time and how it links to the respective design feature (DF8) and its sub-processes. Note that at this stage of the prototype, the cross-validation (DF8.g, see Figure 8) was not yet implemented.

Finally, we conducted a series of workshops and semi-structured interviews with automotive and IT industry experts. This revealed that a special case of so-called “continuous odometer fraud”, previously unknown to us, was impossible to detect with the existing system. In the case of continuous odometer fraud, the mileage of the car is continuously recorded at a lower-than-actual rate, i.e., only a certain percentage of the mileage actually driven is added to the odometer, for example, 80%. This is achieved by installing additional hardware, a so-called “CAN filter”, in the car. Such odometer filters are readily available on the Internet, for example, on eBay, for less than 50 USD. Our solution after this first iteration, however, was only focusing on odometer mileage reduction to uncover potential fraud; an increase at a lower rate could not be related to fraudulent behavior. We addressed this issue in our next iteration by adding a cross-validation feature (DF2, DF8.g).

Iteration 2: Cross-Validation and Scalability

To address the problem of continuous odometer fraud, we leveraged GPS data (start and end coordinates of a trip) in the second iteration (DF2). To use the GPS data to enhance fraud prevention, the verification process needed a substantial update. In addition to verifying the

increase of the odometer value, we also checked that the trip distance calculated on the basis of the odometer mileage exceeded the distance between the GPS points from the start and the end of the trip (DF8.g).

Furthermore, we addressed the scalability of the solution in this iteration (DP4). The internal processes of the application were optimized and structured by several queues to enable the fault-tolerant processing of data from a larger fleet of cars. For the evaluation of this iteration, 100 cars were deployed in a field test. These cars were supplied by one of the leading German car manufacturers, whom we contacted for the evaluation of the initial iteration of the prototype. Supplying the whole fleet with dongles would have been very costly and out of the scope of this study, which is why, as of this iteration, the data was routed over the internal backend of the car manufacturer, where it was sent directly by the connected cars used for the field test.

This version of the prototype was tested over twelve weeks with 100 cars that were used on a daily basis. We also conducted another series of workshops and interviews. The test revealed that the processing and verification of the enriched dataset, including the GPS values, worked as intended. By manipulating the sensor data from the administrator interface, the usage of odometer filters was simulated. The cross-validation procedure thus reliably detected the simulated continuous odometer fraud. Even minor manipulations (e.g., a continuous reduction in the mileage by 10%) could be consistently identified after 15 trips.

The evaluation also revealed stability problems with the underlying infrastructure, specifically in respect to the Ethereum integration. Issues such as clients losing connection to the blockchain network or cloud servers running out of storage could easily be fixed. Other problems were more severe. Especially delicate was the fact that the Ethereum client responded to the sending of a signed transaction to the blockchain network with a valid transaction hash, even if the transaction itself had not necessarily been successfully processed by the network. Hence, additional logic was necessary to assure that a transaction had successfully been processed by the blockchain network. These issues were addressed in a third iteration, leading to the final prototype.

Iteration 3: Stability and Usability

In the third iteration, we addressed the stability problems observed in the evaluation step of iteration two and implemented a smartphone app for end users to interact with the CertifiCar system. We improved the stability of the system with several measures. First, we started relevant processes via a daemon to ensure their uptime and introduced an additional queue to check that a

transaction had been successfully inserted in the blockchain. Additionally, we set up an infrastructure monitoring tool (Nagios) to reduce the response time to system problems.

To improve the usability, we provided a smartphone app, which is shown in Figure 13. It includes an overview screen and a history of the driven distance, as well as a screen that allows the creation of a certificate that can be sent to a receiving party via email. The smartphone application, and in particular the process of creating a certificate, was tested by a focus group of 16 people. The feedback led to a simpler design, specifically with respect to data sharing and certification. The data owner has the option to share only the current odometer value, for example, with a potential buyer. Hence, no detailed car-usage data is revealed. However, data owners might want to share historic data to increase trust and ultimately the sales price. Therefore, they can also share the odometer history on a monthly, weekly, or daily basis.

Overall, the final iteration resulted in a prototype with increased stability and an intuitive interaction possibility through the smartphone application. The robustness and the reception by users was encouraging, resulting in the final clearance for a larger field test, which is now ongoing.

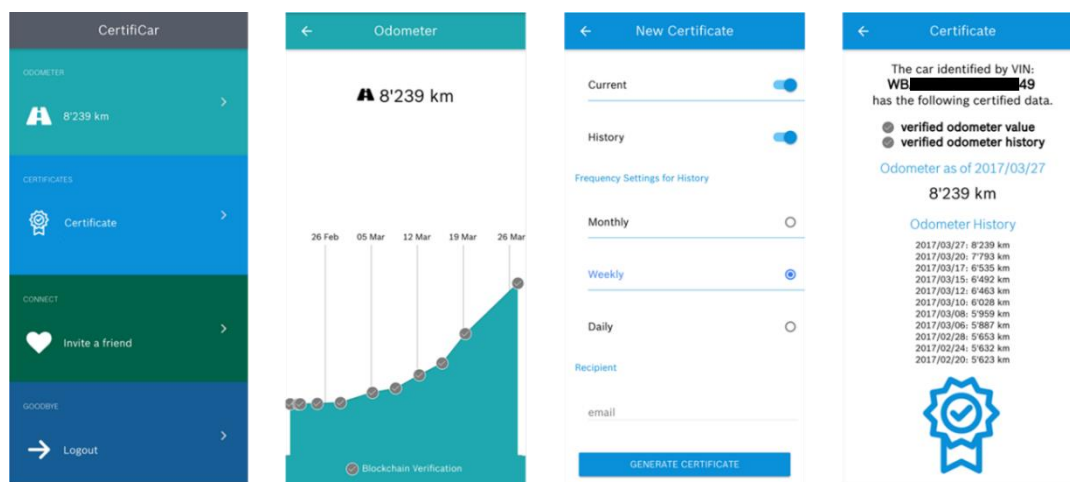


Figure 13. The main screens of the smartphone application

Prototype Evaluation

We have continuously evaluated the implementation against practical results from an accompanying field test that consisted eventually of 100 cars and expert feedback from workshops and semi-structured interviews. Overall, we held six workshops between December 2016 and September 2017, each comprising two to six experts and three to four researchers (in total, 22 evaluators participated) and lasting between three to five hours. Additionally, we interviewed sixteen experts between January and August 2017 for 45 minutes to 1 hour each. We prompted

the participants for specific feedback and related that to the corresponding design decisions to adapt the design principles and features. Among the experts in the workshops and interviews were engineers from a German car manufacturer, a data protection law expert from a German car manufacturer, specialists from a German technical certification provider, an online car sales platform CEO, and engineers from a German car supplier. An overview of the evaluation is shown in Table 3.

	Iteration 1	Iteration 2	Iteration 3
Core Developments and Improvements	Initial end-to-end prototype First implementation of verification process to detect odometer fraud	GPS-based cross-validation to address continuous odometer fraud Queue management for scalability and reliability	Transaction queue to assure reliable blockchain transaction processing Smartphone app for end users
Evaluation	Focus on fraud detection, scalability, and reliability Field test with 5 cars 5 interviews and 2 workshops (lasting 3.5 and 4.5 hours and with 2 and 5 participants, respectively)	Focus on continuous odometer fraud detection, scalability, and reliability Field test with 100 cars and simulated continuous odometer fraud 7 interviews and 2 workshops (4 hours each, 3 people each)	Focus on smartphone application, particularly the process of certificate creation, and system reliability Field test with 100 cars and focus group of 16 people 4 interviews and 2 workshops (lasting 4 and 3 hours and with 6 and 3 people, respectively)
Core Results	Initial verification procedure detects odometer reductions but not continuous odometer fraud Limited scalability and fault tolerance	Cross-validation procedure reliably detects continuous odometer fraud Successful blockchain transaction processing is not guaranteed, occasional loss of transactions	Stable prototype Well-accepted smartphone app Clearance for larger field test

Table 3: Overview and summary of the prototype evaluation

3.3.4 Ex-Post Evaluation of the Design

The scope of the evaluation in DSR reaches beyond the question of whether an artifact works and fulfils the design requirements proposed. Additionally, DSR should thoroughly explore how and why an artifact works (Pries-Heje et al., 2008). Therefore, we conducted an additional ex-post evaluation (Beck et al., 2013; Pries-Heje et al., 2008) to address these questions and to investigate

to what extent the proposed guidelines are actionable and help to create a solution for the underlying problem class. In addition to the case of odometer fraud that we have investigated in detail for the development of our artifact, we included two other use cases for this ex-post evaluation. Thereby, we want to go beyond a single prototype evaluation and gear the evaluation more towards the overall problem class. As suggested by Beck et al. (2013), to reach a higher level of abstraction, we collected additional slices of data and discussed the viability of our proposed guidelines with a purposive group of domain experts. Therefore, we selected two additional cases in which IoT sensor data needs to be protected and that are discussed extensively as fruitful blockchain use cases, namely, supply chain management (Pilkington, 2016; Tian, 2016; Underwood, 2016) and energy microgrids (Imbault et al., 2017; Mengelkamp et al., 2018; Münsing et al., 2017). The first case relates to cold chains, where it must be ensured that the temperature along the supply chain stays within a certain range (Modum, 2018). In the second case, we considered an energy microgrid with participating consumers and prosumers, where it is essential to protect the readings of smart meters for a well-functioning peer-to-peer market (Exergy, 2017b, 2017a).

Participant name	Role	Industry	Case
BC Dev, Manufacturing	Blockchain developer	Manufacturing and engineering	Supply chain
PM BC, Manufacturing	Project manager blockchain	Manufacturing and engineering	Automotive
BC Sol Arch, Energy	Blockchain solution architect	Energy	Energy
PM BC, Energy	Project manager blockchain	Energy	Energy
BC Dev, Software	Blockchain developer	Software consulting	Supply chain
Sol Arch, Automotive	Solution architect	Automotive	Automotive
PM Innovation, Automotive	Project manager innovation	Automotive	Supply chain
PM Innovation, Manufacturing	Project manager innovation	Manufacturing and engineering	Energy
Certification Expert, Inspection	Certification expert	Inspection and product certification	Automotive

Table 4: Ex-post evaluation interview participants

As we had already developed a real-world instantiation of an artifact for the odometer fraud case, we were already in contact with several experts from the automotive and IT certification industries. These relations helped us to recruit a purposive sample of interview participants (Miles & Huberman, 1994; Robinson, 2014) with expertise in the IoT domain, including dedicated experts on subjects such as IoT sensor systems, blockchain technology, odometer fraud, supply chain management, and energy microgrids (see Table 4). We conducted a total of nine interviews (three per use case: odometer, cold chain, and microgrid), each of which lasted 45 to 70 minutes (four face-to-face and five via phone). The conversations were semi-structured, fully recorded, transcribed, and analyzed. We opted for the format of semi-structured interviews to decrease the risk of biasing participants to concrete answers and to allow a more free way of expression, especially as the interviewees often had more expertise in the specific subject matter than the interviewer (Myers & Newman, 2007; Wengraf, 2001). Below, we provide evidence from the transcripts of the nine interviews regarding the efficacy of the proposed design principles and corresponding features to address the design requirements defining our problem class.

DP1: Sensor data is certified on the basis of source to sink protection.

A majority of the participants deemed DP1 to be of “utmost importance” (*PM BC, Manufacturing; BC Sol Arch, Energy; Certification Expert, Inspection*) or even “the most important” (*BC Dev, Manufacturing; PM Innovation, Automotive*), independent of the use case. One participant mentioned that DP1’s “implementation applies to all use cases” and that DP1 is a “necessary basis to guarantee the validity of sensor data” (*BC Dev, Manufacturing*). However, interviewees agreed (*BC Dev, Manufacturing; BC Sol Arch, Energy; PM BC, Energy; BC Dev, Software Consulting; PM Innovation, Automotive*) that “in practice, it is difficult to comply 100% with [the DP]” (*PM BC, Manufacturing*), especially “in the fragmented ecosystem of the IoT, where sensors are built by one company, deployed by another, and a third runs a service on top of that infrastructure” (*PM BC, Manufacturing*). With regard to future developments, it was articulated that the implementation of DP1 could become easier, for example “if sensors can communicate directly with the blockchain” (*PM Innovation, Automotive*) or at least “sign transactions” (*BC Dev, Manufacturing*). One participant additionally noted that “[for a scalable solution] a sensor that is able to sign transactions would be sufficient” (*BC Dev, Manufacturing*), although a sensor that is able to directly communicate (bidirectional) with the blockchain “would open fascinating new possibilities, as it could directly interact with smart contracts and, instead of a one-way communication, a dialogue could be realized” (*BC Dev, Manufacturing*), which would allow the sensor to also receive coins and instructions from the blockchain.

Many interviewees expect blockchain-enabled hardware (*BC Dev, Manufacturing; PM BC, Manufacturing; PM Innovation, Automotive*), for example “sensors similar to hardware security modules” (*PM BC, Manufacturing*) that ease the implementation of DP1 to be available in the future. However, one participant noted that “currently, the software specifications of blockchains [e.g., of signature algorithms] are still evolving [and therefore] the development of sensor ASICs still has to wait” (*BC Dev, Manufacturing*). The advantage of application-specific integrated circuits (ASICs) would rather lie in a “more energy efficient processing than in increased speed” (*BC Dev, Manufacturing*). While the implementation specifics are expected to change, the “basic concept of blockchains as a record of an immutable shared truth [is not]” and, therefore, the usage of blockchain transactions as in DF3 to fulfil DP1 “should continue to make sense” (*BC Sol Arch, Energy*).

While generally source to sink protection through a blockchain transaction was appreciated as a sound measure to hinder data tampering, several participants agreed (*BC Dev, Manufacturing; Certification Expert, Inspection*) that there “will probably never be a way to ensure a completely tamper-proof solution” (*PM Innovation, Automotive*). For example, “one could simply manipulate the surrounding of the sensor – in the case of a cold chain, for example, by putting a cooling element or ice on top of the temperature sensor” (*BC Dev, Manufacturing*). One participant concluded that “while it makes sense to aim for a tamper-proof solution, it is sufficient to ensure tamper resistance that is strong enough to make it economically unprofitable to commit fraud, similar to proof of work [a blockchain mining mechanism]” (*PM Innovation, Automotive*), which is in line with DR1. Overall, DP1 and its implementation (corresponding design features) were strongly supported by all nine interviewees, and the participants provided fruitful insights into how blockchain technology might evolve to enable DP1.

DP2: Sensor data is certified on the basis of cross-validation.

Most participants acknowledged that an implementation of DP2 would be needed, as either DP1 could not uniquely prevent all data tampering or it could not be implemented to the full extent. As such, one participant noted that “it is good that the dependence on DP1 is reduced by the introduction of DP2” (*BC Sol Arch, Energy*), and another stated that “[some kind of] cross-validation is always necessary because already the reading of the sensor could be influenced [in a manipulative way]” (*Sol Arch, Automotive*). Relating to future developments, a participant noted that “increasing the security by implementing DP2 is probably faster and more economically viable than perfecting the implementation of DP1, possibly with future technology” (*PM Innovation, Automotive*).

Participants also noted that DP2 is “rather use case specific” (*PM BC, Manufacturing*), in contrast to DP1, and speculated that “in some cases it might be difficult to find appropriate data for cross-validation” (*BC Dev, Software Consulting*). Regarding the cold chain case, an interviewee suggested that “weather data could be combined with cooling power consumption data of the truck to detect anomalies” (*BC Dev, Software Consulting*). With respect to the microgrid case, they proposed to use data of “a transformer station supplying several houses with electricity” and “weather data in combination with power data from the installed solar panel” (*BC Sol Arch, Energy*) for cross-validation. In the case of odometer fraud, the “cross-validation could be expanded considerably with service and maintenance data”, for example, by “validating that the exchange of brake disks occurs after roughly 50,000 kilometers” (*PM Innovation, Manufacturing*). In essence, all participants supported DP2 and highlighted its context dependency as well as the interlinked nature of DP1 and DP2.

DP3: Data owners determine when and to what extent their certified data is communicated to others.

Generally, the participants stated that the privacy-preserving mechanisms introduced through DP3 are very strong. According to one interviewee, “the propagation of information is organized well in the system and occurs in a very safe way” (*BC Sol Arch, Energy*). Three participants mentioned the upcoming General Data Protection Regulation (GDPR) of the EU (European Commission, 2018) and noted that the most important parts thereof are covered in DP3 and its features (*PM BC, Manufacturing; PM BC, Energy; PM Innovation, Manufacturing*). One participant stressed additionally that “there is also an obligation to inform the data owner about how her data will be used by the receiving party” (*PM BC, Energy*), and another stressed that “there should be a possibility to revoke the sharing of data any time after the data has been sent for the first time” (*PM Innovation, Manufacturing*).

Regarding the importance of privacy, it was noted that it is highly dependent on the gathered data in the specific case and, importantly, on the perception of the data owner towards sharing this data (*BC Dev, Manufacturing; BC Dev, Software Consulting*). For example, people are “used to sharing their electricity consumption data with their energy supplier” (*PM BC, Energy*), and in a cold chain, “a driver might not perceive the sharing of temperature data as very sensitive” (*BC Dev, Software Consulting*). Therefore, a participant argued, “it might actually be a challenge to convince users that data privacy is valuable in their case” and raised the question of “how do you want to raise awareness for that?” (*PM BC, Manufacturing*). This statement is in line with the comment of another participant that “at the moment, privacy is typically driven by regulatory decisions [in Europe] and not by customer demand”, concluding that “currently, it is often not

essential for flourishing businesses [to provide privacy-preserving solutions], but it will probably become a core feature in the future” (*PM Innovation, Automotive*). In line with this last comment, one participant noted that new technology enables gathering and transmitting data at a more granular level, possibly changing users’ perceptions as follows: “If you start sharing your electricity consumption on a minute basis, instead of delivering a quarterly or annual meter reading, you might get more uncomfortable” (*PM BC, Energy*). In summary, the participants appreciated DP3 and the corresponding design features. They also emphasized that privacy is becoming increasingly important as the technological performance and the ability to collect detailed data increases.

DP4: Data is certified on the basis of a linearly scalable system architecture.

Several participants noted that DP4 is, together with DP1, essential for any solution trying to solve the problem of data protection and certification (*BC Dev, Manufacturing; PM Innovation, Automotive*). One participant with a strong business background said that this “needs to be fulfilled right away” (*PM Innovation, Automotive*). In general, the participants noted that the scalability provided by the proposed principles and features is indeed sufficient for real-world applications like, for example, the processing of the majority of all cars in the EU. One participant noted that the “scalability properties of blockchain-based solutions strongly depend on the use case at hand and the specific implementation”, continuing that “the often-heard statement that anything involving blockchain technology does not scale and costs a lot is simply not true [...] as, for example, CertifiCar and the OpenTimestamps project reveal” (*PM Innovation, Manufacturing*).

The hybrid approach of using both decentralized and traditional infrastructures was deemed appropriate by all interviewed blockchain experts, independent of the cases discussed. “Currently, such a solution can only be built on the basis of a hybrid approach” (*PM BC, Energy*), noted one participant, while another added that “taking into account the current state of the blockchain ecosystem, this approach definitely makes sense” (*BC Dev, Manufacturing*). However, considering future developments, many participants speculated (*PM BC, Manufacturing; BC Dev, Software Consulting*) that “it might be possible to build the entire system on a decentralized infrastructure in a far future” (*BC Dev, Manufacturing*), as already hinted before, and it was also noted that already “many players are working, for example, towards decentralized storage solutions with throughput and scalability for enterprise environments” (*BC Dev, Software Consulting*).

Regarding the question of whether a scalable protection system is better built without blockchain technology, i.e., disregarding DF6, many interviewees agreed (*PM BC,*

Manufacturing; Sol Arch, Automotive) that “technically, this would be possible” (*BC Dev, Software Consulting*). However, different considerations in favor of the usage of blockchain technology were made. One participant noted that “using a blockchain to store the hashes makes sense whenever the certification happens in an environment with a multitude of parties with [partially] conflicting interests” (*PM BC, Manufacturing*). For example, in the case of odometer fraud, “the owner of the car, a potential buyer of the car, the car manufacturer, associated and independent workshops, and even different departments within a car manufacturer have different interests regarding odometer fraud” (*PM BC, Manufacturing*). Therefore, establishing a central database for all participants that is operated by just one of the involved parties is a major challenge. Uninvolved third parties can take over the responsibility to run such a system. It was also noted that “new business models based on other sensor data that is shared in a multi-party system” (*Certification Expert, Inspection*) will increase in importance. In principle, it “might be possible to find a traditional database provider [for this role]” (*BC Dev, Software Consulting*); however, it could be costly and potentially difficult to reach an agreement between all parties involved. “A blockchain provides a viable alternative in such a case, with no need to trust a third party” (*BC Dev, Software Consulting*).

In addition, the participants noted that the “overhead of the blockchain is small – really expensive are [hardware] sensors and connectivity” (*BC Dev, Manufacturing*). The blockchain “can even reduce costs”, as its security is less dependent on third-party certification, which is costly and time-consuming (*BC Dev, Manufacturing*). This is especially important for smaller companies, which might not have the resources and processes to deploy highly secure databases. An expert in the research department of a multinational company stated that “the business side clearly does not see the need for a blockchain-based solution yet”, as they think that “a secure and trustworthy database can also be provided by the company itself and its brand name” (*Sol Arch, Automotive*). In line with that, several participants noted that when a blockchain is used, the trust question is transferred to “technology” or “engineering”, while in traditional systems, it is addressed with “brand names” and “company processes” (*BC Dev, Manufacturing; BC Sol Arch, Energy*).

An additional interesting point was made regarding the standardization potential of a solution relying on blockchain technology. An expert from the energy sector noted that individual energy suppliers “might be more willing to accept a solution as an industry standard if its cornerstone is based on blockchain technology, and this decreases the dependence on another company” (*PM BC, Energy*). In contrast, “if a solution’s core is in control of another energy supplier or technology provider, the adoption as a standard would be very difficult” (*PM BC, Energy*).

In essence, the interviewees highlight the importance of DP4 and agree that the proposed features are indeed appropriate to address this design principle. Furthermore, they provide several reasons why a blockchain-based SDPS might be superior to a traditional solution in particular situations. First and foremost, they highlight the potential of blockchain technology in cases where sensor data protection has to be assured in ecosystems with multiple parties with conflicting interests.

In summary, the nine interviews provided additional evidence of the usefulness of our proposed design. The participants reinforced the core considerations and major design decisions of the SDPS design. In addition, the interviews revealed new insights, for example, with respect to the evolution of blockchain technology and its specific business potential. The results also correspond to the findings from the development and evaluation of our prototype. However, by building upon additional slices of data (Beck et al., 2013), they go beyond a “one instance evaluation” of the design.

3.4 Discussion

3.4.1 SDPS Design Theory

After the ex-post evaluation, we integrate our findings and formulate a design theory as summarized in Table 5. Thereby, we follow the seminal work of Gregor and Jones (2007), who laid out six fundamental components of a design theory. Finally, we discuss our findings in light of their theoretical and practical implications.

According to Gregor and Jones (2007), the first component of a design theory is its purpose and scope. The aim of our artifact is to develop a system that protects IoT sensor data generation, processing, and exchange in a privacy-preserving and efficient manner. With respect to the boundaries of the design, we want to highlight that the development of the guidelines was clearly focused on the processing of IoT sensor data and the corresponding challenges, such as big data, multistage data processing pipelines, and distributed data processing across organizational boundaries or multi-party ecosystems. This problem class covers a wide range of relevant issues, which is in stark contrast to existing studies on SDPSs (e.g., Ayoade et al., 2018; Liang et al., 2017; Machado & Fröhlich, 2018) that focus on specific solutions to very specific problems. The generalizability within our wide problem class constitutes an important foundation for our theoretical contribution.

1	Purpose and scope	The aim is to develop a system that protects IoT sensor data generation, processing, and exchange in a privacy-preserving and efficient manner.
2	Constructs	<ul style="list-style-type: none"> • Tamper resistance • Privacy • Scalability • Economic feasibility • Certification
3	Principles of form and function	Design principles (DP1-4) to support the protection of IoT sensor data and corresponding design features (DF1-9) are presented.
4	Artifact mutability	SDPSs have to be mutable, specifically with respect to the amount of data they can handle. DR2 and DR3 articulate this fundamental thought, and DP4 subsequently poses a linearly scalable system. SDPS can be used with benefit by different organizations. However, they need to be adapted particularly with respect to cross-validation. The cross-validation data and the certification procedure are highly dependent on the context.
5	Testable propositions	<ul style="list-style-type: none"> • P1: The artifact enables tamper-resistant IoT sensor data generation, processing, and exchange • P2: The artifact enables privacy-preserving IoT sensor data generation, processing, and exchange • P3: The artifact is capable of processing large amounts of IoT sensor data • P4: The positive effects of the artifact are not negated by artifact development and operation costs
6	Justificatory knowledge	Design requirements are based on the literature on IoT, security, and privacy. Design principles are derived from theory on information asymmetry, privacy, and IS success. Design features build upon blockchain literature.

Table 5: Components of an SDPS design theory

The second component that Gregor and Jones (2007) depict is constructs, which represent core entities of interest in the design. The core constructs we propose are tamper resistance, privacy, scalability, and economic feasibility, which are reflected in our design requirements. These constructs capture the impact of an SDPS and may therefore serve as dependent variables in efforts to investigate SDPS success. In addition, the theory on information asymmetry (Akerlof, 1970) suggests that certification is a core concept and means to overcome information deficits and avoid opportunistic behavior, such as intentional data manipulation. We build upon these insights and base our design on certification. Therefore, certification is a fundamental, independent construct of our work.

Regarding the third component of a design theory, we present principles of form and function that may serve as a blueprint for the construction of IoT sensor data protection systems. To this end, we identify the SDPS design requirements (DR1-4), derive design principles (DP1-4) to

support the protection of the IoT sensor data and depict corresponding design features (DF1-9) (see Figure 6). The requirements, principles, and features constitute actionable guidelines, which highlights a core difference between our work and the extant research. Thereby, we reflect the various calls in the IS literature to support the development of implementable tools to increase security and privacy, especially in the IoT (Bélanger & Crossler, 2011; Lee et al., 2018; Medaglia & Serbanati, 2010; Pavlou, 2011).

To account for the special nature of IS artifacts, Gregor and Jones (2007) call for explicitly addressing the mutable nature of these artifacts as a fourth component. In the case of SDPSs, we reflected the importance of mutability specifically with respect to the amount of data they can handle. DR2 and DR3 articulate this fundamental thought, and DP4 subsequently poses a linearly scalable system. However, the design that we derived is not universally applicable, nor is it “one-size-fits-all”. While SDPSs can be used with benefit by different organizations, they need to be adapted particularly with respect to cross-validation. The cross-validation data and the certification procedure are highly dependent on the context, as the development of the instantiation that we presented clearly indicates.

The fifth component of a design theory comprises testable propositions. These propositions might be presented as “if a system or method that follows certain principles is instantiated, then it will work, or it will be better in some way than other systems or methods”. Following this argumentation, we can deduce propositions from the presented design requirements. The design requirements disentangle the “it will work, or it will be better” into specific, contextualized needs that must be addressed by the artifact. Propositions postulate that these needs have been successfully addressed and serve as a basis for assessing the impact of the artifact. Applying this rationale to DR1-4, we deduce the following four propositions: the artifact enables tamper-resistant IoT sensor data generation, processing, and exchange (P1). The artifact enables privacy-preserving IoT sensor data generation, processing, and exchange (P2). The artifact is capable of processing large amounts of IoT sensor data (P3). The positive effects of the artifact are not negated by the artifact development and operation costs (P4). These propositions might be helpful in developing test cases for future instantiations.

Finally, Gregor and Jones (2007) encourage scholars to provide the justificatory knowledge of their design. We base our design requirements on insights from the literature on IoT, security, and privacy (see Section 4.1). The design principles are mainly derived from theory on information asymmetry, privacy, and IS success (see Section 4.2). Ultimately, the design features build primarily upon the blockchain literature (see Section 4.3). This theoretical grounding enabled us, in close interplay with insights from practice, to derive a set of purposive guidelines

for the design of SDPSs in the form of DRs, DPs, and DFs. Gregor and Jones (2007) emphasize the importance of explanatory theory as a “linking mechanism for a number, or all, of the other aspects of the design theory” (p. 327). We reflect this role of explanatory theory by explicitly deriving design principles that serve as a link between design requirements and design features. This thorough conceptualization of the problem is a key distinction from previous literature (e.g., Ayoade et al., 2018; Liang et al., 2017; Machado & Fröhlich, 2018), and it facilitates the generalizability of our findings, which enables our theoretical contribution.

3.4.2 Design Implications

Our research has important design implications for SDPSs that address IoT-related security and privacy challenges (Ayoade et al., 2018; Crossler & Posey, 2017; Liang et al., 2017), specifically with respect to the value proposition of blockchain technology. Blockchain-based SDPSs inherit core characteristics of blockchain technology (Notheisen et al., 2017) and therefore are particularly useful in certain scenarios (see Table 6). While SDPSs are used to protect simple data pipelines, for example, to secure data transfer from sensors to one single intra-organizational system, they are also leveraged in the case of multi-stage data pipelines that cross organizational boundaries and involve a potentially large ecosystem of players, as our prototype case reveals. In the latter case, blockchain-based SDPSs are particularly valuable because they can protect sensor data even in large ecosystems with conflicting interests through the use of a shared, immutable ledger. In addition, a blockchain-based SDPS is a decentralized system. Hence, the involved parties are peers, and no single party controls the overall system (Beck, Müller-Bloch, & King, 2018). As our ex-post evaluation reveals, such a system is often perceived as “neutral” and might be accepted as an industry standard much faster than a centralized system. Finally, important security and protection technology, such as public-key cryptography, is already built into blockchain technology (Buterin, 2013; Noyen et al., 2014). Additionally, the infrastructure to use these protocols is readily provided by a decentralized set of actors (e.g., miners), who are typically incentivized through the economics of cryptocurrencies. Essentially, blockchain technology offers a ready-to-use set of well-defined security protocols. For smaller companies, in particular, that do not have cryptography specialists or corresponding technology available, blockchain-based SDPSs offer the opportunity to leverage state-of-the-art security technology that is usually license-free and often designed for rapid adoption.

Blockchain characteristic	Related advantages	SDPS usage implications
Shared, immutable ledger	<ul style="list-style-type: none"> • Blockchain integrates the advantages of distributed databases and crypto technology • Well-managed data redundancy across different parties • Secure data processing that fosters data integrity 	“using a blockchain to store the hashes makes sense whenever the certification happens in an environment with a <i>multitude of parties with [partially] conflicting interests</i> ” (PM BC, Manufacturing)
Decentralized system	<ul style="list-style-type: none"> • No central authority • All parties are peers with the same rights • No single party controls the overall system 	“[members of an ecosystem] might be more willing to <i>accept a solution as an industry standard</i> if its cornerstone is based on blockchain technology and this <i>decreases the dependence on another [single] company</i> ” (PM BC, Energy)
Ready-to-use set of well-defined security protocols and infrastructure	<ul style="list-style-type: none"> • Private and public key cryptography stack built into blockchain • Infrastructure readily provided by a decentralized set of actors incentivized through economics of cryptocurrencies • Security does not rely on third-party certification, which is costly and time-consuming • Even smaller companies with no dedicated cyber-security or cryptography specialists can leverage state-of-the art security technology 	“overhead of the blockchain is small – really expensive are [hardware] sensors and connectivity” (BC Dev, Manufacturing), “the <i>blockchain can reduce costs</i> ” (BC Dev, Manufacturing)

Table 6: Blockchain-based SDPS usage implications

However, as our design theory reveals, blockchain-based SDPSs have to be carefully designed. Blockchain technology is not a universal solution that addresses the derived design requirements out of the box. The fundamental design implications must be considered to address the derived design requirements (see Table 7). With respect to DP1 (sensor data certified on the basis of source to sink protection), it is important to note that, as of today, sensors cannot communicate directly with the blockchain. Therefore, the data must be protected as early as possible in the processing chain by building and signing blockchain transactions as close as possible to the sensing unit. In the future, blockchain-enabled sensors could drastically simplify this and might allow for signing within the sensing unit. In addition, DP2 (sensor data certified on the basis of cross-validation) has to be carefully addressed. More specifically, system designers have to realize that blockchain technology generally cannot assure “tamper-proof” processes, and the additional cross-validation of the sensor data is necessary to enable effective tamper resistance. Thereby, a nondetection risk of fraud remains. With respect to DP3 (data owners

DP1	Sensor data is certified on the basis of source to sink protection	
	Prototype design & eval.	<ul style="list-style-type: none"> • Data must be protected as early as possible in the processing chain • In the prototype, we collected odometer data and preprocessed it as soon as possible in a way that data manipulation from that point on was prevented, and we built and signed the blockchain transaction as close as possible to the odometer sensing unit • However, in the prototype, we could only do this rather late in the processing chain, as a blockchain cannot be directly integrated into the odometer sensor
	Ex-post evaluation	<ul style="list-style-type: none"> • “[source to sink protection] is a necessary basis to guarantee the validity of sensor data” (BC Dev, Manufacturing) • “in practice, it is difficult to comply 100% with [source to sink protection]”, especially “in the fragmented ecosystem of the IoT” (PM BC, Manufacturing) • Implementation of DP1 could become easier, for example “if sensors can communicate directly with the blockchain” (PM Innovation, Automotive) or at least “sign transactions” (BC Dev, Manufacturing)
DP2	Sensor data is certified on the basis of cross-validation	
	Prototype design & eval.	<ul style="list-style-type: none"> • Blockchain technology cannot assure “tamper-proof” processes per se, so additional cross-validation is necessary to enable effective tamper resistance, and a nondetection risk of fraud remains • Initial prototype verification procedure detects odometer reductions but not continuous odometer fraud • Prototype cross-validation procedure finally reliably detects continuous odometer fraud
	Ex-post evaluation	<ul style="list-style-type: none"> • “[blockchain] will probably never be a way to ensure a completely tamper-proof solution” (PM Innovation, Automotive) • “[some kind of] cross-validation is always necessary because already the reading of the sensor could be influenced [in a manipulative way]” (Sol Arch, Automotive) • DP2 is “rather use case specific” (PM BC, Manufacturing)
DP3	Data owners determine when and to what extent their data is communicated to others	
	Prototype design & eval	<ul style="list-style-type: none"> • Blockchain technology cannot assure data privacy per se, so privacy must be implemented on top of the blockchain in the form of an access management service • Feedback of 16 prototype users that fine-grained sharing mechanisms have to be implemented • Clearance of app for large field test that included user feedback & legal compliance check
	Ex-post evaluation	<ul style="list-style-type: none"> • “there should be a possibility to revoke the sharing of data any time” (PM Innovation, Manufacturing) • “the propagation of information is organized well [in the proposed design] and occurs in a very safe way” (BC Sol Arch, Energy)
DP4	Data is certified on the basis of a linearly scalable system architecture	
	Prototype design & eval.	<ul style="list-style-type: none"> • Hybrid blockchain architecture necessary to enable scaling • Odometer sensor values are stored in a central repository, and only the digital fingerprint (hash) of the records is recorded on-chain • System for 100 cars was deployed on the basis of two low-performance standard Amazon EC2 instances, and there were no performance issues during the evaluation
	Ex-post evaluation	<ul style="list-style-type: none"> • “scalability properties of blockchain-based solutions strongly depend on the use case at hand and the specific implementation” (PM Innovation, Manufacturing) • “the often-heard statement that anything involving blockchain technology does not scale and costs a lot is simply not true” (PM Innovation, Manufacturing) • “Currently, such a solution can only be built on the basis of a hybrid approach” (PM BC, Energy)

Table 7: Design implications for blockchain-based SDPS

determine when and to what extent their data is communicated to others) it should be noted that a blockchain is not a universal remedy that can guarantee privacy (Conti et al., 2018; Fabian et al., 2016; Goldfeder, Kalodner, Reisman, & Narayanan, 2018; Kumar, Fischer, Tople, & Saxena, 2017). In the context of sensor data sharing specifically, privacy mechanisms have to be implemented on top of the blockchain in the form of an access management service. In addition, by relying on a hybrid blockchain approach, there must be assurances that the sensor data itself is not stored in a public permissionless blockchain and that data integrity can be maintained. Finally, regarding DP4 (data certified on the basis of a linearly scalable system architecture), specific blockchain architectures have to be implemented. With the current state of technology, hybrid blockchain architectures (Ayoade et al., 2018; Zyskind et al., 2015) are necessary to enable scaling. Therefore, viable systems store sensor values in a central repository, and only the digital fingerprint (hash) of the sensor values is recorded on the blockchain.

3.4.3 Theoretical and Practical Contributions

In summary, the proposed SDPS design theory is the key theoretical contribution of our work. We synthesize our design into a conceptual solution that addresses a whole problem class. Notably, the codification and abstraction of our design, including the design requirements, design principles, and design features, enables generalizability beyond a particular problem. The provision of actionable guidelines based on such a thorough conceptualization is, to the best of our knowledge, a novel contribution, which was specifically called for (e.g., Bélanger & Crossler, 2011). Thereby, we add to the literature on IoT and IoT-related security and privacy challenges, as well as to the literature on blockchain technology.

More specifically, our investigation of the problem class confirms and conceptualizes earlier evidence from the literature (C. C. Aggarwal et al., 2013; Lowry et al., 2017) that the distributed, multilayered nature of IoT systems, as well as IoT ecosystems with multiple parties and potentially diverging interests, introduces very specific and particularly serious challenges. The derived design requirements can serve as a basis for future research, for example, investigating how their fulfillment affects the adoption of IoT IS. Furthermore, we base the design principles, in particular, on the theory of information asymmetry, which has been used before as a fruitful basis in the design of IS that enables the reliable exchange of data (e.g., Notheisen et al., 2017). In contrast to the existing SDPS-related literature, we specifically focus on certification as a well-known means of overcoming information asymmetries. As such, we leverage deep insights from the existing body of knowledge on information asymmetries (Bond, 1982; Genesove, 1993; Spence, 1976), and certification in particular (Akerlof, 1970; Albersmeier et al., 2009), which we strongly believe represents a useful basis for other design research in the realm of SDPSs.

Finally, we discuss the design features and the design implications of our research on the usage of blockchain technology in detail. Notably, we shed light on both the advantages as well as the potential problems of using a blockchain for SDPSs. We elaborate how the proposed design can address the widely discussed shortcomings of blockchains, such as scalability and privacy. We do this by building upon the existing research on hybrid blockchain architectures (Ayoade et al., 2018; Zyskind et al., 2015) and thereby encourage design researchers to specifically reflect the latest developments in this domain.

With regard to practical contributions, we first of all provide a blueprint that guides the development of SDPSs. Furthermore, we address emerging blockchain concerns that more and more practitioners share, namely, blockchains have no scalability, they induce high costs, and they cannot assure privacy. Our design – and more specifically the prototype – reveals that these concerns can be addressed with existing technology. This might inspire practitioners to overcome their concerns and start leveraging blockchain technology for their enterprises. In addition, in line with the existing research (Beck et al., 2016; Christidis & Devetsikiotis, 2016), our evaluation reveals where the use of blockchains might be particularly helpful in practice. Ecosystems with a multitude of parties with potentially conflicting interests often rely on an intermediary to ensure reliable data exchange and trust. In these cases, blockchain technology might serve as such an intermediary. Additionally, blockchain-based solutions might facilitate the establishment of industry standards. Finally, in light of ever-increasing regulation, blockchain-based solutions might serve as a cost-efficient complement to third-party certification. Smaller companies, in particular, might benefit from the ready-to-use security protocols and corresponding infrastructure that the blockchain provides. In the realm of IoT, however, physical devices must be blockchain-enabled. As of today, the data pipeline too often remains unprotected directly after the sensing unit of such devices.

3.5 Conclusion

The study at hand uses a DSR approach to propose a design theory for a sensor data protection system (SDPS). More specifically, we derive design requirements, design principles, and design features for a blockchain-based SDPS. In addition, we design and develop an instantiation of an SDPS (CertifiCar) on the basis of three iterative cycles. Our prototype prevents the fraudulent manipulation of car mileage data. Finally, we provide an ex-post evaluation of our design theory considering two additional use cases in the realms of pharmaceutical supply chains (Modum, 2018) and energy microgrids (Mengelkamp et al., 2018). The findings of our evaluation suggest

that the proposed design ensures the tamper-resistant gathering, processing, and exchange of IoT sensor data in a privacy-preserving, scalable, and efficient manner.

The results of this study should be assessed in light of its limitations. We derive design principles on the basis of specific theoretical lenses. Building upon an alternative selection of theoretical lenses, we might have identified different or additional design requirements and principles (see Meth et al., 2015). However, the chosen theories are well accepted and undisputed and represent a reliable and stable basis for analysis. In addition, our evaluation confirms that our design principles are concise and independent of current technology and upcoming technology developments, as well as applicable to the chosen problem class across different use cases. A second limitation refers to the design features that are grounded in the capabilities of today's blockchain technology. Blockchain technology is in an early stage of development (Beck et al., 2017), and, in particular, new on-chain/off-chain approaches are still emerging (Ayoade et al., 2018; Machado & Fröhlich, 2018; Zyskind et al., 2015). Therefore, the proposed design features might change with future, potentially disruptive blockchain breakthroughs. However, we want to highlight the fact that we build upon the latest blockchain research at the forefront of technology, and our features reflect latest on-chain/off-chain architecture approaches that provide a viable tradeoff between security and scalability (Ayoade et al., 2018; Zyskind et al., 2015). A third limitation is related to the evaluation of our design theory. We developed and evaluated CertifiCar and investigated two additional use cases to reflect our design. While a quantitative and broader evaluation is desirable and encouraged, we want to emphasize that at this point in time, corresponding systems and domain experts are not widely available.

Beyond the aforementioned opportunities, there are many other possible extensions to our work. We contribute to an emerging literature stream that aims to advance the theoretical understanding of blockchain technology. We hope that our study serves as a fruitful basis for further research on how blockchain technology facilitates new modes of ecosystem collaboration, for example, by establishing security, privacy, and trust. More specifically, we encourage scholars to investigate and compare the various blockchain-based data protection approaches that are currently emerging with respect to their business potential (Risius & Spohrer, 2017). Finally, while there are several industry initiatives, such as the Trusted IoT Alliance, and many companies are currently developing promising use cases, we see an absence of design and theory to bridge the gap between technology and business. Blockchain technology is rapidly evolving, but its business potential still remains vague. It is not only researchers who have been too optimistic about the potential of blockchain technology (Beck et al., 2017). In practice, blockchain technology is still overhyped, and discussions are either very technology-focused or business-driven without reflecting the actual capabilities and restrictions of the current technology. In line

with Bélanger and Crossler's (2011) call for more actionable solutions, we encourage design science researchers to fill the articulated gap and link (business) problem classes to blockchain technology and corresponding applications.

Chapter 4

DeFi: Information from Social Media in Blockchain-Based Fundraising⁵

As we have seen in Chapter 3, DApps might enable foundational applications relevant for the EoT, such as the secure and privacy-preserving collection, storage and exchange of data. However, only DeFi provides the financial infrastructure necessary to integrate complex monetary incentive systems into these applications and, thus, ultimately turn things from *technical devices* into *economic agents*. Besides more complex financial products, DeFi has led to fundamental financial infrastructure such as crypto-payments pegged to fiat currencies, the issuance of loans and also equity-like fundraising mechanisms. To date, the DeFi mechanism with the largest adoption has been equity-like fundraising. In fact, this mechanism is responsible for a large amount of all investments made towards the whole blockchain ecosystem. As such, many companies developing EoT solutions have been financed through blockchain-based fundraising. In this chapter, we examine the mechanism of blockchain-based fundraising and focus particularly on how investors leverage available information for their decisions. Studying the role of information in financial decisions has a long tradition in IS and is relevant to understand what innovations prevail. Here, we focus particularly on data from social media as this is considered the most relevant information channel in the case of blockchain-based fundraising.

⁵ Parts of this chapter, which are not further demarcated in the text, were initially published in the context of the following academic publications: Chanson, Gjoen, et al. (2018), Chanson et al. (2019) and Chanson et al. (2020).

4.1 Introduction

Virtually every successful IT company was dependent on external financing over the whole course of its existence, with a particular vulnerability in its first years (Dos Santos, Patel, & D'Souza, 2011; Kim, Mithas, & Kimbrough, 2017). External capital plays a key role in the formation and development of IT ventures, as these companies often have to finance costly growth strategies while no significant revenue sources are present (R. Aggarwal et al., 2012). As such, even the formation of entire new technological ecosystems depend on external financing (Breznitz, Forman, & Wen, 2018). Consequently, external financing has been studied intensely by the IS community. For example, scholars have examined the influence of blogs on different aspects of VC (R. Aggarwal et al., 2012; R. Aggarwal & Singh, 2013), the role of VC for the formation of new technological ecosystems (Breznitz et al., 2018), or the influence of media coverage on the geographic focus of venture capitalists (B. N. Greenwood & Gopal, 2016). Besides VC, IS researchers have also studied other funding mechanisms such as IPOs (e.g., Ceccagnoli et al. 2012; Lundmark et al. 2017), crowdfunding (e.g., Burtch and Chan 2019; Hong et al. 2018), or debt (e.g., Kim et al. 2017; Kim and Mithas 2011).

Recently, blockchain technology has led to the emergence of a blockchain-based system of DeFi, which has received only little attention from research so far. DeFi introduces a number of decentralized versions of equity and debt financing. In particular, different blockchain protocols offer the issuance of collateralized (MakerDAO, 2020) or uncollateralized loans (Aave, 2020), the trading of these loans (Uniswap, 2019) and also fundraising methods resembling an equity-based capital raise. The coins serving as the currency of DeFi as well as the core processes of DeFi mechanisms are based on computer protocols on a blockchain, called smart contracts, which automate the key constituents of the funding procedures. Consequently, DeFi is highly relevant for IS researchers both because of the fundraising capabilities as well as the technological foundations. A particular segment of DeFi, namely a new method of equity-like capital raise referred to as ICO, has attracted tremendous interest from practice. An ICO is a novel financing mechanism in which coins are released on a blockchain in exchange for funding, transferred in the form of cryptocurrencies (Chanson, Gjoen, et al., 2018). More capital was raised through ICOs than through traditional VC in 2017 in the blockchain industry, totaling to a similar size as *all* VC Internet investments in a typical quarter (Chanson, Risius, et al., 2018). To be precise, Coindesk (2018) reports \$5.4B ICO funding in 2017. As such, ICOs are the DeFi mechanism with the highest level of adoption in history up to this date in spring 2020. While a fruitful research stream on ICOs is emerging in the IS community, it is still in its infancy and lacks other aspects of DeFi so far (Fridgen et al., 2018; Guske & Bendig, 2018; Park & Yang, 2018).

Considerable IS research concerning external financing has been directed towards the informational power of different mediums such as news articles, blogs, microblogs, etc. (e.g., Aggarwal et al. 2012; Aggarwal and Singh 2013; Greenwood and Gopal 2016). These efforts are rooted in decades of research considering the impact of the availability and quality of information on markets, which have even led to a Nobel Prize (Akerlof, 1970). With the emergence of the web 2.0, researchers have started to investigate the effect of online information that is disseminated in social media, for example in comments, blogs, videos and reviews (R. Aggarwal & Singh, 2013; Lukyanenko et al., 2017; Mai et al., 2018). We differentiate between social media in interactive (e.g., discussion forums), and non-interactive contexts (e.g., blogs or microblogs) (Hansen, Lee, & Lee, 2014; Lukyanenko et al., 2017). Social media research has often focused on the implications for financial markets (R. Aggarwal & Singh, 2013; Lukyanenko et al., 2017). However, in traditional financial markets, the relevance of social media is rather limited in light of established information sources such as audited financial reports, statements of company representatives or specialized news portals (e.g., Bloomberg or Reuters). In contrast, in the context of blockchain-based fundraising, social media is in fact an essential source of information (Chanson, Gjoen, et al., 2018; Mai et al., 2018). Recently, IS researchers have reflected the distinct role of social media in the realm of blockchain: For example, Mai et al. (2018) use tweets and forum posts to predict price movements of Bitcoin. Furthermore, considering ICOs, Guske and Bendig (2018) find a relation between the number of twitter messages and the amount of funding raised.

One of the most important dimensions for equity-based corporate fundraising is the change of valuation a capital raise fuels (Braun & Fawcett, 2006; Lundmark et al., 2017). Both investors and companies have an interest in how the valuation changes after a fundraising event (Allen & Faulhaber, 1989; Booth & Chua, 1996; Loughran & Ritter, 2004). Often, the price at which investors buy shares during an IPO is substantially lower than the price of the same shares when they are traded on exchanges shortly afterwards. This phenomenon is referred to as underpricing and received considerable attention from research in the case of IPOs (Allen & Faulhaber, 1989; Booth & Chua, 1996; Loughran & Ritter, 2004). There are numerous theories why it is in the interest of organizations to foster underpricing of their stocks in IPOs (e.g., exploit signaling effects to stress the company's quality, or increase the diversity of ownership). However, independent from the specific explanation it holds that more organizational legitimacy leads to stronger underpricing (Lundmark et al., 2017). Legitimacy constitutes the perception of an entity's behavior to be socially desirable (Suchman, 1995). In practice, legitimacy is not directly observable and therefore gauged with proxy measures (Zimmerman & Zeitz, 2002). Numerous loci of legitimacy are identified in organizational literature in general (Aldrich & Fiol, 1994) and

for the specific case of capital raising (Pollock & Rindova, 2003; Zuckerman, 1999). The recent study of Lundmark et al. (2017) introduced social media as a source of organizational legitimacy. They show how the strategic use of Twitter by organizations can confer legitimacy, using underpricing in IPOs as a proxy measure.

However, the extant knowledge from traditional finance and initial blockchain-related studies leaves open a multi-fold literature gap regarding the impact of information from social media on the outcome of fundraising events. First, the relevance of various forms of non-interactive social media (e.g., blogs, reviews and microblogs) for financial markets has been addressed in numerous studies, however, the impact of *interactive social media* (e.g., discussion forums) for venture financing has not been studied in depth before. Second, extant research predominantly focuses on addressing traditional financial markets or cryptocurrency prices, while research on *blockchain-based fundraising* is virtually inexistent. We argue, however, that especially IS-scientists with their work at the intersection between business, technology, and people (Hevner et al., 2004) can inform the entrepreneurial concerns regarding this blockchain-dependent funding mechanism. To address this multi-fold literature gap, we formulate the following overarching research questions:

RQ 2a: How does information from interactive and non-interactive social media directly relate to the extent of underpricing in blockchain-based fundraising?

RQ 2b: How is information from interactive and non-interactive social media intertwined in its relation to the extent of underpricing in blockchain-based fundraising?

To address this, we explain the fundamental principles of ICOs and introduce them as a blockchain-based funding mechanism. We draw on theory of organizational legitimacy and apply it to the context of social media and ICOs. We include Twitter and discussion forums in our investigations. We gather a data set of 95 ICOs to address our research questions. We find that discussion forum activity has a direct impact on underpricing and that the effect of microblogs on underpricing is mediated by this activity.

The remainder of this chapter is structured as follows. In Section 2, we introduce the foundations of ICOs, underpricing, legitimacy and social media, and derive a number of hypotheses regarding the impact of social media on underpricing. In Section 3, we evaluate these hypotheses empirically by analyzing the data we gathered. Section 4 provides a critical reflection of our results and research approach, while Section 5 concludes with the contributions and limitations of our study.

4.2 Conceptual and Theoretical Background

4.2.1 ICOs and Underpricing

In an ICO investors support a project with funding and receive newly generated project-specific coins in return. The main goal of projects launching an ICO is to secure funding while investors aim at owning a stake in such a project via the possession of the project-specific coins. Both the payment of capital by investors and the distribution of coins as a return occur automated through a blockchain. Typically, the entire ICO is conducted on one specific blockchain which serves as an ICO platform, for example the Ethereum blockchain. In this case, the project-specific coins are issued by a smart contract on the platform blockchain and are called tokens. Smart contracts are computer protocols, which automatically perform specific transactions without the involvement of a third-party after execution criteria have been met (Beck et al., 2017; Szabo, 1997). Currently, Ethereum is by far the most commonly used platform for ICOs, although others are emerging (e.g., Neo or Qtum). The process of an ICO on such a platform is depicted in Figure 14. Before the launch of the ICO, the funds seeking project creates two smart contracts which define the key parameters of the ICO and the tokens to be distributed: For instance, the amount of money going to be accepted maximally (i.e., the *hard cap*), the time frame when the ICO will happen, the prize of the project-specific coins and how many of these coins will exist. After these smart contracts are deployed on the blockchain, investors can participate in the ICO by paying capital to the ICO Smart Contract. Notably, the capital is not paid directly to the project itself. After the payment of investors, the following part of the process is completely automated according to the pre-defined rules in the smart contracts. The project receives access to the capital paid into the ICO Smart Contract and investors receive their share of tokens from the Token Smart Contract. Thus, the core machinery of the ICO process – the exchange of capital for tokens – is a fully automated system running on a blockchain and can, as such, be viewed as an artifact of IT. In the future, as already existing DApps on blockchains, for example concerning identity verification, are maturing, it is probable that also peripheral processes of a funding, like “know-your-customer” and “anti-money-laundering” verifications are possible to integrate in this automated processing of smart contracts.

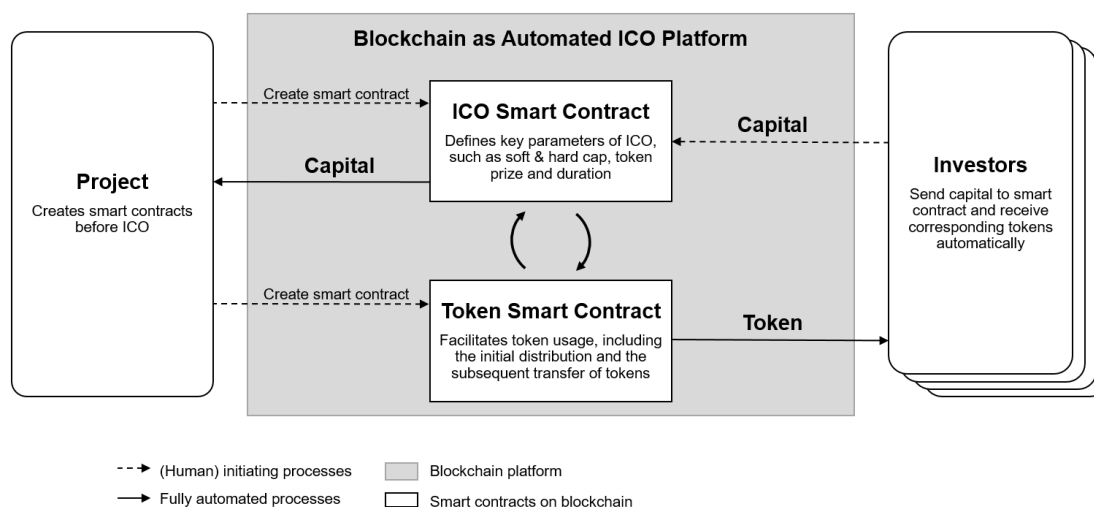


Figure 14. ICO process

Besides these technical aspects, it becomes apparent that the application of this technology in the form of an ICO (i.e., a generally applicable funding mechanism) is of importance for human organizations and their management, as the funding of projects is one of the core tasks in the establishment of an enterprise. For these reasons, studying the phenomenon of ICOs meets the core interest of the IS discipline to advance knowledge about the use of IT in organizations and their management (Hevner et al., 2004).

In contrast to stocks in an IPO, the utility of the project-specific coins returned to investors can vary a lot and is defined individually for each ICO project. Commonly, the utility is distinguished into three core components: The ability (1) to transfer value, (2) to access a service, and (3) to receive a profit share of the project (FINMA, 2018). Bitcoin, in its original sense of a currency, is a good example for (1): Possession of a Bitcoin essentially allows to easily transfer value worldwide over the Bitcoin blockchain. Ethereum and its coin Ether is a good example for (2): Ethereum provides infrastructure for a computer, the Ethereum Virtual Machine, which is accessible worldwide for anyone. In order to use this computer, for example to deploy smart contracts, a fee uniquely payable in Ether is due. Hence, only the possession of Ether allows access to the service provided by the Ethereum computer. Sharing profits of a project (3) is very similar to the payments of dividends in stocks and is implemented, for instance, by Modum or NEX. Individual coins of ICOs typically possess one or a combination of these three benefits, although many more can be linked to a coin at discretion, for example the ability to vote on important project decisions or the earning of more coins through provision of core infrastructure or supervisory services.

Coin utility	Description	Examples
Value transfer	Value can be transferred by exchanging coins with other people just like with traditional currencies	Bitcoin, Monero, Dash
Service access	A special service provided by the blockchain project, such as smart contract hosting or storage services, can only be accessed by paying a fee in the native coin	Ethereum, Filecoin, Gnosis
Profit share	Coin holders are entitled to receive a certain share of the profit the blockchain project generates	NEX, Modum

Table 8. Main coin utilities

Underpricing is the phenomenon often observed in IPOs that stocks issued by the company prior to the listing are sold at a lower price than they are later traded on the stock exchange after the listing. In the case of ICOs, we define underpricing as the difference between the issuance price of the coins in the ICO and the closing price after a workweek on at least one publicly accessible exchange. This is equivalent to the return for ICO investors, however we keep using the term underpricing, as it is established in the management and finance literature. This literature offers numerous explications why underpricing is often observed in IPOs. One line of argument is that underpricing is not desirable by the issuing company, because it can be interpreted that more money could have been raised issuing the same amount of stock at a higher price. In this case, underpricing is typically explained by misaligned incentives introduced by the underwriting third party investment bank, which suppresses the normal interplay of supply and demand in the setting of the offer price (Ritter & Welch, 2002, p. 1803). However, other explanations indicate that underpricing is actually in the long term interests of the issuing company. According to Allen and Faulhaber (1989) signaling effects of underpricing are used deliberately by firms to underline their high quality and benefit from better conditions in subsequent financing rounds after a listing, ultimately maximizing their yield. Booth and Chua (1996) show that the issuer's preference for a broad ownership dispersion incentivizes underpricing. Aggarwal et al. (2002) claim that firms intentionally underprice to increase demand for the stock after the listing. Loughran and Ritter (2004) and Cliff and Denis (2004) demonstrate that underpricing is used by firms to compensate highly ranked analysts for their future coverage of the stock. In the case of ICOs the decision on the issuance price is essentially made by a small team of founders based on future expectations. They operate with a lot of freedom because ICO projects are typically very early stage and there is no operating performance to base the valuation on. As the capital market can be accessed directly over the blockchain, no underwriters are needed and the influence of third parties on the process is very limited. Additionally, the core team typically holds a major share of 10% to 30% of all project-specific coins and controls another substantial portion of coins as an endowment to the project, which can later be used as project funding. For all these reasons, we assume that

underpricing in the case of ICOs will be in the long term interest of the issuing project. Therefore, and in line with scholars investigating underpricing in IPOs, we argue that organizations which are viewed as more legitimate will yield a higher level of underpricing, as these audience perceptions will be reflected in the closing price after five days of trading (Lundmark et al., 2017). In the following, we will elaborate that the use of social media poses one of the few but effective means for ICO projects to establish and manage their legitimacy.

4.2.2 Legitimacy in ICOs

Suchman (1995) defines legitimacy in his seminal article as a “generalized perception or assumption that the actions of an entity are desirable, proper, or appropriate within some socially constructed system of norms, values, beliefs and definitions” (p.574), adopting an inclusive and broad-based notion of the term. Notably, the conferral of legitimacy ultimately is a reaction of observers *perceiving* an organization and, hence, legitimacy can differ between viewers (Lamin & Zaheer, 2012) and “is possessed objectively, yet created subjectively” (Suchman, 1995, p. 574). Extant literature has identified a variety of potential loci of legitimacy, which Deephouse and Suchman summarize, in their review of the sources of legitimacy (2008, pp. 54–56), as society-at-large, interorganizational relations and the media. Evidence from journalism and mass communications strongly advocates that the media not only serve as a proxy measure but indeed confer legitimacy independently (Deephouse, 1996). Extant research distinguishes between different types of legitimacy. Aldrich and Fiol (1994) distinguish between sociopolitical and cognitive legitimacy, where the latter simply reflects the public knowledge about a new venture and is especially important for entrepreneurs (Shepherd & Zacharakis, 2003). Dacin et al. (2007) emphasize the difference between investment legitimacy, market legitimacy, relational legitimacy, and social legitimacy. Recently, in a first study Lundmark et al. (2017) investigated the impact that social media can have on the accrual and management of legitimacy. Their results indicate, using Twitter as the focal social medium, that organizations’ activity on microblogs could actually confer legitimacy prior to the launch of IPOs.

The conferral of legitimacy has important consequences for organizations, as it can enhance survival prospects through better access to stakeholders and increased resource flow (Aldrich & Fiol, 1994; Hampel & Tracey, 2017). Legitimation protects a firm’s perceived relevance by justifying to a peer or super ordinate system its right to exist (Dowling & Pfeffer, 1975). Under conditions of uncertainty and complexity, accruing and managing legitimacy is of particular relevance (Kostova & Zaheer, 1999). Investing in early-stage companies with often unproven technologies, as it is typically the case with ICOs, carries inherent risk and uncertainty. Hence, it seems reasonable to assume that the perceived degree of organizational legitimacy from investors

plays a vital role in the price developments of any ICO. Although there is no literature on legitimacy and ICOs, there exist some studies on legitimacy in traditional capital markets. Scholars have highlighted the importance of external authorities such as security analysts (Zuckerman, 1999), reputed underwriters or renowned media (Pollock & Rindova, 2003). While this is probably also true to a certain extent in the case of ICOs, contrary to more mature capital markets, the ecosystem of potential external authorities is very limited. There is not a whole industry of professional analysts, underwriters or professional media outlets that observers could refer to for their perception of legitimacy. However, there is very high activity of all stakeholders involved (i.e., project initiators, investors, influencers, developers, etc.) on social media, which were also identified as an origin of legitimacy by Lundmark et al. (2017). Besides this study, information on legitimacy effects of social media in IPOs is virtually absent, although there are studies on price predictions with social media data in stock markets (C. Oh & Sheng, 2011) and the Bitcoin market (Mai et al., 2018). We aim to supplement to the knowledge on how social media serves as a mechanism for conferring legitimacy in two ways. First, we include social media data of individual user-generated content (UGC) in our study, specifically from the online discussion forums Bitcointalk and Reddit. This is supplementary to existing research which includes only marketer-generated content (MGC), namely tweets of the same companies whose legitimization is investigated. Second, we focus on legitimacy effects in the context of rather small and new companies that launch ICOs, in contrast to existing research which focuses on organizational legitimacy of established enterprises that launch an IPO.

4.2.3 Social Media Driving Legitimacy in ICOs

Kaplan and Haenlein (2010) define social media as “a group of Internet-based applications that build on the ideological and technical foundations of Web 2.0 and that allow the creation and exchange of UGC” (p. 61). Current examples of social media applications include microblogging sites like Twitter, social or professional networking sites like Facebook and LinkedIn, discussion forums like Reddit and many more. Social media allows firms to communicate directly to large audience groups like customers or small-scale investors and to interact with this audience in a swift way. The information exchanged over social media can vary extensively between firms regarding the topic focus and the level of details provided. Firms launching an IPO use these channels mainly to confer cognitive legitimacy and raise awareness, for example by referral to articles of authorities discussing the launch of the IPO (Lundmark et al., 2017). In the case of ICOs, however, the social medium Twitter is often the main, and in the beginning the only, public communication channel of the project besides their website. Therefore, sometimes also vital information is announced via Twitter, for example the launch of an ICO, the start of trading of a

coin on an exchange or the publishing of a whitepaper providing project details. This further increases the importance of and the dependency of investors on tweets in the ICO realm, compared to IPOs. For example, Nebulas announced a change of price terms on Twitter *during* the ICO, with retrospective effect: “New Pricing Rules for NAS Pre-Sale Early Bird Participants! Nebulas will make settlements with all the early bird pre-sale participants according to the highest ETH price at 6:00am EST during the early bird period. <https://t.co/H8KMG1PiEc>”, and Iota announced the trading launch of their coin: “IOTA Exchange Launch <https://t.co/fgVANntbV6> #iota #tangle #blockchain #exchanges”.

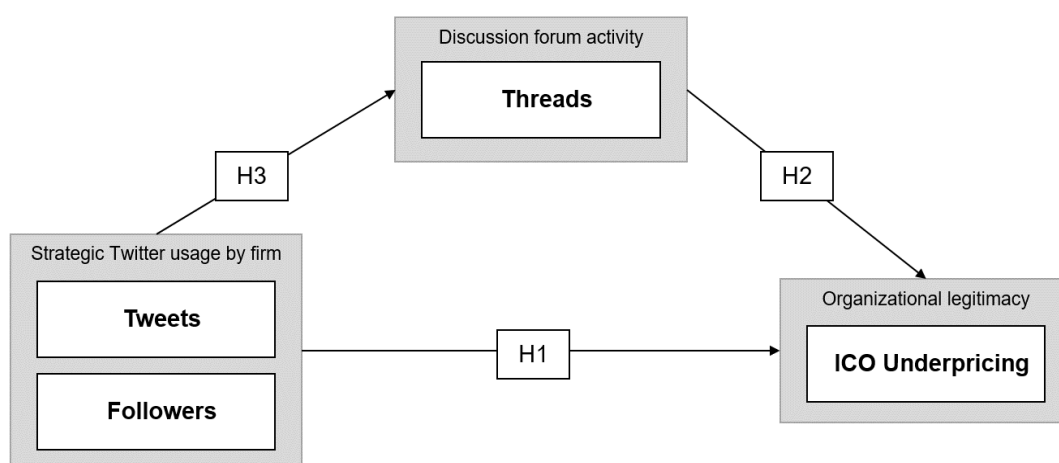


Figure 15. Research model for social media effects on organizational legitimacy

In this study, we elaborate how social media can increase organizational legitimacy by opening a direct information channel to individual investors. An overview of our research model is shown in Figure 15. As legitimacy is not directly observable (Zimmerman & Zeitz, 2002), we follow the approach of Lundmark et al. (2017) and use underpricing as our dependent variable, serving as a proxy for organizational legitimacy. We use Twitter as a focal medium because of its importance in the ICO community, as discussed above, and existing studies indicating that, specifically, Twitter can confer legitimacy in the context of traditional capital markets (Lundmark et al., 2017). If it is possible to confer legitimacy via Twitter we expect that a more intense usage of the medium increases this effect (Miller & Tucker, 2013; Risius & Beck, 2015). As legitimacy is a perception, the conferral of legitimacy should also increase with the size of the audience (Clark & Melancon, 2013; Miller & Tucker, 2013). This leads to two predictions about the mechanisms that confer legitimacy and, ultimately, may increase the level of underpricing observed in an ICO.

Hypothesis 1a: *The higher the activity of a firm on social media prior to their ICO, the higher the level of underpricing.*

Hypothesis 1b: *The higher the reach of a firm on social media prior to their ICO, the higher the level of underpricing.*

Companies use microblogs, such as Twitter, as a means to disseminate information directly to their follower network, whether this be potential investors, consumers, partners, or society-at-large. This is primarily a one-directional communication system that effectively spreads MGC to a targeted audience. The crypto environment has, however, formed online communities where communication can flow multi-directional, for example in public online discussion forums such as Reddit and Bitcointalk. On these platforms, millions of users are discussing different facets of the cryptocurrency industry, hence creating UGC on these networks. Pollock and Rindova (2003) provided “preliminary evidence about the difference between the impact of firm-provided and media-provided information in the IPO market”, stating that “it is media-provided, rather than company-provided, information that has the credibility and/or reach necessary to influence investor behavior systematically” (p. 640). Antweiler and Frank (2004) were the first ones to demonstrate the general impact of UGC on stock price movements. Meanwhile, Goh et al. (2013) highlighted that UGC drives sales immediately while effects of company generated content depend on the addressed audience, and Mai et al. (2018) specify in the context of cryptocurrency values that messages on an Internet forum have a stronger impact compared to tweets. Given these findings, we expect that the discussion on such online forums could develop a decisive impact on the legitimacy of an ICO project. Including UGC in the study of social media and legitimacy might offer a novel understanding of how social media can drive legitimacy. So far, to the best of our knowledge, studies regarding the effects of discussion forums on organizational legitimacy have been virtually absent. Additionally, although numerous studies have found systematic relationships between microblogging and the stock market (C. Oh & Sheng, 2011), the documented effects from discussion forum activity is weak (Antweiler & Frank, 2004; Das & Chen, 2007; Tumarkin & Whitelaw, 2001). As the cryptocurrency market, and hence the investor community, is inherently digital, the digital activity from the audience might provide a previously unavailable opportunity to discover the legitimation effect of discussion forums in capital markets. This leads to an additional prediction about the mechanisms that confer legitimacy and, ultimately, may increase the level of underpricing observed in an ICO.

Hypothesis 2: *The higher the amount of UGC on discussion forums mentioning a company before its ICO, the higher the level of underpricing.*

In an early phase of the existence of a company, as it is the case previous to an ICO, participants on such discussion forums rely to a big part on the publication of information by the company itself, which can then be discussed among the forum members. As such, we expect that

a more active communication on Twitter as well as an increase in people receiving this communication has a positive effect on the activity on such discussion forums. This leads to two predictions about mechanisms that increase discussion forum activity and, possibly, may increase the level of underpricing observed in an ICO in a mediated indirect way.

Hypothesis 3a: *The higher the activity of a firm on social media prior to their ICO, the higher the amount of UGC on discussion forums.*

Hypothesis 3b: *The higher the reach of a firm on social media prior to their ICO, the higher the amount of UGC on discussion forums.*

Referring, again, to the findings that communication from other parties than the firm under investigation develops stronger influence (Goh et al., 2013; Pollock & Rindova, 2003), and that the use of Twitter may increase discussion activity (Dunlap & Lowenthal, 2009), and considering the previous hypotheses, we expect that the influence of Twitter is actually mediated by discussion forum activity. This leads to the final two predictions about the mechanisms that confer legitimacy and, ultimately, may increase the level of underpricing observed in an ICO.

Hypothesis 4a: *The effect of a firm's activity on social media prior to their ICO on ICO underpricing is mediated by the amount of UGC on discussion forums.*

Hypothesis 4b: *The effect of a firm's reach on social media prior to their ICO on ICO underpricing is mediated by the amount of UGC on discussion forums.*

In this study, we operationalize the hypotheses formulated above using the number of tweets as a measure of social media activity, the number of followers on Twitter as a measure of reach and the number of threads on selected subreddits as a measure for the amount of UGC on discussion forums.

4.3 Research Study

To investigate these hypotheses regarding the effects of social media-dependent organizational legitimacy on the genesis of ICO underpricing, we collected a comprehensive sample of ICO prices and related social media communication.

4.3.1 Sample and Data Analysis

The ICO sample was drawn from the independent ICO database ICODrops on March 8th 2018. To guarantee economic feasibility of the manual data collection, a random sample of 212 ICOs was drawn from the overall sample of 340 successfully completed ICOs since the database's

start in May 2017. Of these, 72 did not have registered trading data for the five days necessary to calculate underpricing, mostly because they were not yet listed on an exchange. Another 12 ICOs were removed from the sample because other necessary data to perform the analyses, such as twitter or discussion forum information, was not available. To ensure a minimal quality of the unregulated ICOs we removed all that raised less than \$1M (i.e., two) and those that achieved less than 50% of their funding target (i.e., 31). This resulted in a final sample of 95 ICOs.

For all the 95 ICOs we collected data from several websites: For price information of the coins we used Coinmarketcap (coinmarketcap.com). The tweets were manually gathered directly from the official accounts of the ICOs on Twitter (twitter.com), totaling to 4,188 unique tweets. The number of followers was derived from snapshots of the Twitter page of projects registered by the Internet Archive (web.archive.org). The discussion forum activity was retrieved from Bitcointalk (bitcointalk.org) and selected subreddits on Reddit (reddit.com). In total, we crawled 33,784 mentions in threads. Data concerning the control variables was gathered from ICODrops (icodrops.com), project webpages, whitepapers and LinkedIn.

To test the proposed hypotheses, we conducted a number of OLS regressions. Multicollinearity was assessed by examining the variance inflation factors. With VIF scores around one, there was no indication of multicollinearity among the independent and control variables. Next, the leverage and influence of outliers were assessed through the Cook's distance. Filecoin's ICO is the observation with the most leverage and influence, raising \$257M during the ICO. As such outliers are considered relevant, no further sample treatment was conducted. Furthermore, normality and linearity were assessed through visual inspection of a normality and residual Q-Q plot. Although White's test showed no significance for heteroscedasticity, Heteroscedasticity-consistent standard errors (HC1) are applied as a conservative measure in order to ensure consistent estimates of standard errors (Long & Ervin, 2000).

To test for mediation effects, we performed a formalized test, namely *confidence interval bootstrapping*, to measure whether the mediation effect is statistically significant (Hayes, 2017). The test is run on a macro developed for SPSS by Preacher and Hayes (2004). A number of 5,000 bootstrap samples were made, and the same control variables and robust standard errors were applied to test all hypotheses.

4.3.2 Measures

Dependent variable

Underpricing is the dependent variable and serves as a proxy for organizational legitimacy. In this, we follow the approach of Lundmark et al. (2017), because legitimacy is impossible to

observe directly (Zimmerman & Zeitz, 2002). ICO underpricing is understood as the difference between the issuance price of the coins in the ICO and the closing price after a workweek of trading on at least one publicly accessible exchange. We chose five days because of the extreme volatility of ICOs often observed within the first days after a coin is listed. This is typical for an initial phase of trading, which is dominated by uncertainty about the market value of a company. However, in contrast to IPOs, several additional factors increase the high volatility in the case of ICOs: Exchanges often cannot handle the high traffic in the initial phase of a coin listing, essentially shutting down for some of the users, preventing them from reacting to price changes and adjusting orders. Additionally, the exact start of trading is typically not announced to the public in advance. Together with the fact that coins often can only be transferred to an exchange after trading has started and it takes time for investors to move their coins with blockchain transactions, the volume at trading start can be very low. The issuance price is the price investors pay for the coins when participating in the ICO. Prices were determined in the denomination Ether, as the reference currency for ICO investments, for two reasons. First, the great majority of recent ICOs launch on the Ethereum network and, second, Ether is the common currency that an ICO accepts as payment. Accordingly, we calculate underpricing for ICOs as the following percent change: $(5^{\text{th}} \text{ day closing price} - \text{ICO issue price}) / (\text{ICO issue price}) \times 100$.

Independent variables

To address hypotheses 1a-3b, we followed the classical mediator analysis approach (Baron & Kenny, 1986) by testing the following sets of models. Models 1 and 2 test the effect of strategic Twitter management and discussion forum activity on ICO underpricing (i.e., *Tweets*, *Followers* and *Threads*). Model 3 tests the effect of strategic Twitter management on discussion forum activity as the dependent variable.

Tweets represents the total number of tweets (including re-tweets, i.e. sharing, or “re-tweeting” someone else’s Tweet) posted 30 days before the ICO date. This 30-day interval does not include the ICO date itself in order to avoid any biases, such as reverse causality. The data is gathered by going through a firm’s Twitter feed, subsequently counting the total number of tweets posted in the 30-day time interval before the ICO.

Followers represents the total number of followers on the Twitter account of a given organization the day before the ICO, measured in thousands of followers. Hoffman and Fodor (2010) consider this variable a measure of popularity given that people follow the organization to obtain information for investing decisions. As historical data on the number of followers of an account is not available directly on twitter we approximate the total number of followers using the Internet Archive (web.archive.org) which provides past snapshots of companies’ Twitter

pages. If no snapshot was provided for the day before the ICO, we conducted a linear approximation in relation to the deviating days from the ICO date, and the creation of the Twitter account. Snapshots taken before the ICO date were preferred over snapshots taken shortly after the ICO date in order to minimize potential biases in the data, such as reverse causality. This resulted in a median deviation of 16 days which we consider satisfactory.

Threads represents the (logarithmic) total number of threads that a firm was mentioned on selected online discussion forums in the 30-day interval prior to its ICO date. Cha et al. (2007) suggests that mentioning infers acknowledgement, addressivity and attribution. Since actively mentioning a certain company on discussion forums requires cognitive effort, a company that is mentioned frequently might be associated with stronger influence. A multitude of online discussion forums exist on cryptocurrency topics and the largest communities are Bitcointalk and Reddit. With millions of users, these two forums arguably serve as a satisfactory proxy for online discussion forum activity on cryptocurrencies. Bitcointalk is a forum solely for the purpose of discussing cryptocurrency topics, hence the forum as a whole was included in the search. However, Reddit is a general forum with no restrictions in terms of discussion content. Hence, only a subset of Reddit forums, called subreddits, were included in the search, according to the three following conditions. First, the forum should have a considerable number of subscribers, such that posts are likely to play a part in the legitimation process. Specifically, we defined a minimal number of subscribers of 100,000. Second, a relevant number of discussions about ICOs should exist on the given subreddit. Third, the subreddit should evolve around a topic related to the cryptocurrency environment. This ensures that posts are seen by users within the target group of potential investors. This resulted in the inclusion of 5 subreddits. We then searched the according forums with Google and used Google syntax to limit the search results to the given online domains and the correct time interval. We searched for either the name of the project or its trading ticker symbol.

Control variables

To respect potentially confounding effects beyond those hypothesized above, eight control variables are incorporated in our models. Following the approach of renowned scholars investigating IPO underpricing (e.g., Pollock and Rindova 2003), we control for essential loci of legitimacy in the context of investments, specifically quality indicators of the individual projects and case-specific variations of the ICO process (Ibbotson & Ritter, 1995), as well as more general influences like the accompanying media coverage.

Variable	Mean	Median	Min	Max	Std
ICO Underpricing	110.97	42.54	-73.23	1730.78	245.28
Tweets	44.08	33.00	0.00	235.00	42.64
Followers	6.89	4.88	0.00	46.22	7.14
Threads	5.36	5.46	2.08	8.31	1.04
Crypto news	3.25	2.00	0.00	20.00	4.21
Firm age	16.51	11.00	0.00	107.00	15.56
Raised funds	92.12	51.78	1.91	1154.63	154.03
Oversubscribed	0.79	1.00	0.00	1.00	0.41
ICO duration	14.11	7.00	1.00	84.00	16.20
Valuation	267.23	118.31	4.76	8667.79	893.53
Min cap	0.14	0.00	0.00	2.00	0.38
Max cap	0.14	0.00	0.00	1.00	0.35

Table 9. Descriptive statistics for complete sample

Crypto news represents the total number of unique articles a project is mentioned in on cryptocurrency news webpages 30 days before its ICO date. Previous research showed clear effects of media-provided content on organizational legitimacy (Pollock & Rindova, 2003). The set of news webpages was found by incrementally adding search terms related to cryptocurrency news in Google’s “News” tab (e.g., “Crypto news”, “ICO news”, etc.), until further searches did not add any more additional news sites. The resulting 50 websites were examined to filter out websites that do not frequently publish cryptocurrency-related content. Out of the 18 relevant news pages, Similarweb (similarweb.com) was used to exclude webpages with less than 1 million visits per month. The remaining 9 websites were included in a Google search in a similar fashion to *threads*, limiting the search to the relevant web domains and time span. We did not include any traditional media because they provide almost no information on projects before the ICO. Specifically, a Factiva search of all the mainstream news sites with over 100 million monthly visits (i.e., Business Insider, Forbes and Bloomberg) revealed only six articles on our full sample of 95 ICOs. *Firm age* represents a company’s age in months, measured as the date of foundation subtracted from the ICO date. Firm age might affect underpricing as older companies have had more time to develop legitimacy, both actively and passively (Lundmark et al., 2017). Hence, it is expected that, for instance, a three-year-old company presents itself more credibly and legitimately in the ICO process compared to a company that was founded only three months before. *Raised funds* describes the total amount raised during the ICO, measured in thousands of Ethers. It is reasonable to expect that a company’s ICO that raised substantially more than its

peers did, received more attention and is considered more “desirable and proper”. **Valuation** represents the implied valuation of the focal company, calculated as the amount of funds raised divided by the share of coins for sale. Similar to *Raised funds*, this variable is measured in thousands of Ethers. Applying the same reasoning as above, a company with a higher valuation might be perceived more “desirable” than its peers, and might therefore confer more legitimacy in the cryptocurrency community. **Oversubscribed** is a dummy variable describing whether the ICO reached the maximum cap, which is given the value 1 if this is the case. When a company needs to stop the ICO before the pre-announced end date because the maximum cap is reached the ICO generated excess demand. If the investors active in ICOs are also active on exchanges the coin is listed on, it is reasonable to assume that the excess demand will manifest itself in increased underpricing once the coin starts trading. Furthermore, the very fact that a ICO was oversubscribed can spark interest with investors in general, because it is clear that the coins are “desirable”, which is a strong signal of legitimacy. **ICO duration** represents the total amount of days the ICO lasted. A long ICO duration can signal that the demand is relatively lower for a given coin, meaning it is perceived as less desirable by the market. Conversely, ICOs ending within the first day often generate attention and might be perceived as more legitimate. Hence, it is expected that increasing ICO duration has a negative impact on ICO underpricing. **Min cap** represents the minimum investment required by investors to participate in the ICO, measured in Ether. If the minimum limits are substantial, it is reasonable to assume that certain investors are excluded from the ICO. If this is the case for interested investors, there exists excess demand for the coin that might result in higher underpricing once the coin starts trading on exchanges. **Max cap** is a dummy variable that gives the value 1 if the company has set an upper investment limit in their ICO. Several factors indicate that setting a maximum investment might increase underpricing. In case some investors were prohibited from investing the entire amount they wished, excess demand is created that might result in higher underpricing. Furthermore, a maximum cap might lead to a more fragmented pool of coin holders. As the value of a coin is likely to increase as a function of how many coin holders exists, a successful ICO with a maximum cap might be perceived as more desirable. Hence, it is expected that the maximum cap might increase the underpricing.

4.3.3 Results

The results of the regression models addressing hypotheses 1-3 are summarized in Table 10. In addition, the results of the confidence interval bootstrapping are displayed in Table 11.

Dependent variable	ICO Underpricing		Threads
	Model 1	Model 2	Model 3
Tweets	0.0195 (0.4265)		-0.0011 (0.0023)
Followers	6.4155 (4.0839)		0.0263 (0.0120) **
Threads		57.1669 (20.3703) ***	
Crypto news	-3.8299 (5.3352)	-5.8737 (4.5823)	0.033 (0.0221)
Firm age	-2.0819 (1.3219)	-1.7610 (1.1664)	-0.0055 (0.0061)
Raised funds	-0.1689 (0.2198)	-0.1098 (0.1879)	0.0005 (0.0006)
Oversubscribed	89.6081 (39.3218) **	33.5220 (36.0724)	1.0000 (0.3244) ***
ICO duration	0.1829 (1.5639)	-1.0023 (1.5936)	0.0167 (0.0087) *
Valuation	0.0156 (0.0195)	0.0139 (0.0189)	-0.0000 (0.0000)
Min cap	180.159 (174.292)	181.857 (183.787)	-0.0340 (0.1354)
Max cap	104.19 (70.1130)	84.1594 (73.4662)	0.4209 (0.2351) *
Intercept	12.0124 (74.5026)	-189.573 (88.0862) **	4.1013 (0.3916) ***
N	95	95	95
F	2.25	2.59	3.12
Adjusted R ²	0.1490	0.1783	0.1048
Unstandardized coefficients are reported. Robust standard errors (HC1) in parenthesis. * p<0.1, ** p<0.05, *** p<0.01			

Table 10. OLS robust regressions

Model 1 tests hypothesis 1, i.e. whether strategic Twitter management is systematically associated with ICO underpricing. However, it does not show support for these hypotheses. Model 2 addresses hypothesis 2 that higher discussion forum activity is related to higher ICO underpricing. It provides support for this hypothesis. *Threads* is significant at a 1% level, indicating that more discussion forum activity increases a firm's legitimacy, ultimately driving up ICO performance. Note that the variable is logarithmically transformed, implying that the marginal value of additional discussion forum activity is positive, but decreasing. Hence, while being mentioned on discussion forums might drive legitimacy, this effect decreases with more online attention. Examining hypothesis 3, Model 3 provides the regression that answers our hypotheses regarding the relationship between internal and external social media mechanisms. Hypothesis 3a suggests the more tweets a firm posts prior to its ICO, the more the firm will be mentioned on discussion forums. However, Model 3 does not offer any support for this

hypothesis. Hence, it does not seem like Twitter activity in itself is sufficient to spark online dialogue on discussion forums. Conversely, Model 3 shows that the total number of Twitter followers of a company is indeed systematically associated with more online discussion forum activity. Hence hypothesis 3b is supported, concluding that companies with a large follower base on Twitter do receive more exposure on online discussion forums.

Independent variable: Followers^a	Total effect	Direct effect	Indirect effect
Effect	6.4155	5.0850	1.3306
Standard error (HC1)	4.0839	4.2661	-
Bootstrap standard error	-	-	0.9413
T-statistic	1.5709	1.1919	-
p-value	0.1200	0.2367	-
Lower level confidence interval	-0.3768	-2.0114	0.0804
Upper level confidence interval	12.2078	12.1813	3.0101
^a Number of bootstrap samples for percentile bootstrap confidence intervals: 5,000			

Table 11. Confidence interval bootstrapping

To test hypothesis 4, we consider if discussion forum activity has a mediating effect on the relationship between strategic Twitter management and ICO underpricing. Indeed, although there was seemingly no direct effect between the two, the presumed effect might work through discussion forum activity. In the formal mediation test, *Followers* is selected as the independent variable as it was the only variable that predicted discussion forum activity. Is it so that having a large follower base has an indirect effect on ICO underpricing through discussion forum activity? The output from SPSS is summarized in Table 11. As shown previously, the total effect of the number of Twitter followers is not statistically associated with ICO underpricing, confirmed with a p-value of 12%. We also indicate that the significance strongly weakens when the mediator is introduced, showed under the ‘Direct effect’ column in Table 11. The bootstrapping macro produced a 90% confidence interval based on the sorted values of the estimated indirect effects (i.e., the difference between the total and direct effect). Based on the values of the confidence interval, it can be concluded that the indirect effect is statistically different from 0. Hence, hypothesis 4b is supported. The legitimacy of a large follower base influences online attention and activity, which again confers legitimacy in the blockchain community, ultimately measured through ICO underpricing. The summarized results are visualized in Figure 16.

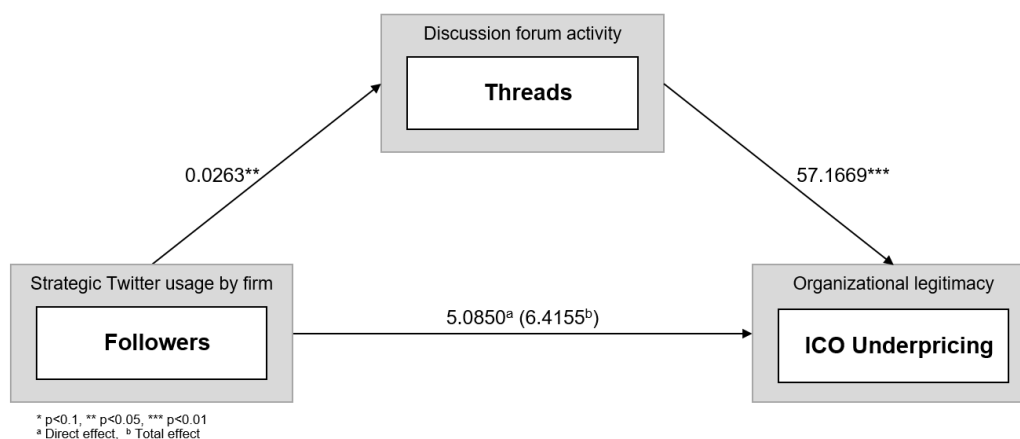


Figure 16. Summary of findings

4.4 Discussion

The results of our investigation reveal four key findings regarding ICO underpricing through social media enabled organizational legitimacy. Our first hypothesis that the strategic use of Twitter (i.e., tweets and followers) increases underpricing in ICOs directly is not supported. This is in contrast to previous findings that show a significant correlation between these measures and underpricing in the case of IPOs (Lundmark et al., 2017). Possible explanations for our result include that the mechanisms conferring legitimacy in the context of ICOs and recently established ventures might differ more than expected from the setting of established companies and IPOs. However, a close consideration of social media related research shows that effects of company generated messages depend on the communication style. Namely, marketer generated messages only drive company performance when specifically addressing individual users (Goh et al., 2013). Thus, company messages have been found to be fully mediated through the appealed users.

In line with this argumentation, we find a highly significant correlation between the activity on public discussion forums (i.e., threads) and ICO underpricing, confirming our second hypothesis. This suggests that social media confers legitimacy through the provision of content by third parties in an environment not directly controlled by the legitimacy-accruing entity. This relates to earlier key IS contributions that UGC has more impact on increasing purchases than MGC (Goh et al., 2013) and that forum contributions are more influential than tweets in driving the price of Bitcoin (Mai et al., 2018). Additionally, this result can be viewed in perspective of the key contribution of Pollock and Rindova (2003) on organizational legitimacy, which found that media-provided content is more important than firm-provided information in influencing

investor behavior. The fact that the control variable “crypto news” is not significant in any of our models, prompts us to speculate that social media actually replaces, to a certain extent, the influence of traditional media in the context of ICOs. We further elaborate on this interpretation below.

The third hypothesis, that the usage of Twitter (i.e., tweets and followers) influences the activity on discussion forums, is partially supported by our results. Specifically, we find a systematic correlation between followers and threads (i.e., hypothesis 3b), but no significance for the number of tweets (i.e., hypothesis 3a). Thus, we note that the reach of the company providing information is important, whereas the frequency of these updates alone seem to have no significant effect on the activity on discussion forums. Hence, we conclude that the strategic use of twitter may influence the activity on discussion forums. This is a further indication that the mediation effect of hypothesis 4 might be valid.

Hypothesis 4 regarding the mediation effect is indeed supported by our analysis. This indicates that it is actually rather the activity on discussion forums than the usage of Twitter by the projects which confers legitimacy. Additionally, it suggests the possibility that the activity on discussion forums, and with it the legitimacy of an organization, could be influenced through strategic Twitter management by that same organization. We stress that our formal analysis itself conveys no justification regarding the causality of the effects observed. However, our argumentation leading to the hypothesis 4 delivers several explications that make a causal relationship plausible. Additionally, Dunlap and Lowenthal (2009) also suggest that Twitter use may increase discussion activity. This would, in turn, mean that organizations can actually use Twitter strategically for establishing legitimacy by influencing the behavior of the crowd on a multitude of social media and, hence, have major implications.

We highlight two additional observations regarding the control variables. First, “oversubscribed” is significant in model 1 and 3, however, not in model 2. The significance is very intuitive, as a good (i.e., in our case project-specific coins representing a stake in the project) for which more demand than supply exists is clearly “desirable” in the sense of organizational legitimacy. However, “oversubscribed” is not of much informational value in practice, as it is only available after the ICO has finished. Additionally, it is typically not possible to measure the magnitude of oversubscription in an ICO because of its fully automatic nature. No formal offerings need to be made before the ICO as the payment to the ICO smart contract itself directly leads to the successful participation in the ICO without any prior action of the investor, contrarily to IPOs. Interestingly, the significance of “oversubscribed” vanishes when the effect of discussion forums is introduced in model 2. This suggests that the activity on discussion forums *before* the

ICO is a more accurate predictor for the extent of underpricing than the information if the ICO was oversubscribed.

Second, the variable “crypto news” is not significant in any of the models. This is unexpected, as previous research on organizational legitimacy in the market of IPOs shows clear effects of media-provided content (Pollock & Rindova, 2003). The variable “crypto news” represents the number of articles published regarding an ICO before its launch on specialized news sites covering cryptocurrency topics. Mainstream media, such as Forbes or Bloomberg, were excluded because they provide almost no such articles. We speculated above, that in the market of ICOs social media, and especially UGC on discussion forums, might have replaced this influence of traditional media-provided content. On the one hand, our assumption might be valid only temporarily, as long as mainstream media keep providing almost no coverage of projects prior to their ICO. On the other hand, it might be a more fundamental change, that in the inherently digital and decentralized space of cryptocurrencies and ICOs, individual users and their contributions on key forums develop more influence than centrally curated articles of traditionally organized publishers. This view is supported by the fact that there are already established news sites specialized on cryptocurrency topics, which do include information about projects before the launch of their ICO and are, in this study, integrated into the control variable “crypto news”, which never reaches significance. However, the non-significance of this control variable should not be over-interpreted, as it was not specifically designed to test the influence of media-provided content in our research setup but rather to control for its effects.

4.5 Conclusion

This study set out to address how the use of social media relates to organizational legitimacy, which is not directly observable in practice, by investigating the effects on ICO underpricing. This is motivated by a three-fold literature gap. First, social media have previously been introduced as a source of legitimacy (Lundmark et al., 2017), however, there is a lack of information on the influence of *different types of social media* on organizational legitimacy, especially those that are beyond the control of the legitimacy-accruing entity. Second, extant research focuses on the influence of social media on the legitimacy of established organizations that are mature enough to conduct an IPO, while research considering *young ventures* is lacking. Third, research considering the *phenomenon ICO*, and consequently also the important enquiry of the mechanisms related to organizational legitimacy in this context, is virtually absent in IS research, although it is a key IS concern. As such, the main goal of this study is to further the

understanding of how social media affects organizational legitimacy, especially in the unexplored context of young ventures in the context of ICOs.

Our study provides a number of theoretical contributions that are empirically validated. First, based on extant IS-research on social media and studies on organizational legitimacy, we derive possible legitimation effects of user-driven social media, namely discussion forums, which are supported by the results of our empirical study. As such, this is the first work that stresses the importance of UGC on organizational legitimacy and substantiates this claim with empirical results. Second, we pose that different types of social media interact in their effects on organizational legitimacy. Namely we find that strategic Twitter management by the legitimacy-accruing entity is mediated in its effects on organizational legitimacy by the activity on discussion forums. Contrary to previous studies, we cannot confirm direct legitimacy effects of strategic Twitter management. Third, we are the first to investigate legitimacy effects of social media on young ventures as opposed to extant research on bigger corporations. We lay out the special importance of social media in this context and provide results that point towards the parallel decrease of importance of traditional media and, as such, a potential shift in loci of legitimacy in the context of new ventures and ICOs.

Various effects of social media have been an important and ongoing topic in IS (Dewan & Ramaprasad, 2014; Ge, Feng, Gu, & Zhang, 2017; Mai et al., 2018; O. Oh, Agrawal, & Rao, 2013; Qiu, Tang, & Whinston, 2015; Schlagwein & Hu, 2017; Scott & Orlikowski, 2014). We apply key findings of the IS community on social media to the context of organizational legitimacy and confirm the validity of these findings in the new context. We expand the previous investigations of Lundmark et al. (2017) on legitimacy effects of social media in three ways. We include additional types of social media in our study, namely UGC generating discussion forums, observe the effects on new ventures, as opposed to more mature ones, and focus on the special context of ICOs. In addition to the insights on the legitimacy effects of social media, we are the first to introduce the fundamentals of the important phenomenon ICO as an IT-based funding mechanism to IS research. We position ICOs at the core interest of the IS community as a highly automated funding mechanism based on a single IT artifact with important consequences for humans, organizations and their management.

Lastly, this study also contains important implications for practice. By showing that social media-induced legitimacy can affect a project's financial performance by inspiring cross-platform discussions, we inform management launching an ICO on important engagement areas. These insights could also inform other industries, indicating that the level of discussion activity is an important performance metric for the social media channels of organizations and, as such, could

contribute to the ongoing discussion about measuring the return of social media marketing (Hoffman & Fodor, 2010). Our study also provides data that could enable investors to estimate the level of underpricing of an ICO in advance and help to assess risk factors. Our findings show that it is important for prospects to consider the forum discussion activity for assessing the ICOs legitimacy as part of their due diligence before making investment decisions.

While our study is a first approach to investigate the effects of a variety of social media and their interactions on organizational legitimacy, there are some important concerns to our research. First, our sampling procedure was limited to a specific set of social media, namely Twitter as a means to broadcast news by the organization, as well as Reddit and Bitcointalk as discussion forums. While these are the most prolific ICO related social media platforms, investigating additional social media sources might produce different results. However, as our hypotheses were derived from general research insights on social media and we chose the most prevalent broadcasting channel with Twitter and only the most active discussion forums for our data collection, we strongly believe that our results are reasonably robust and generalizable. Still, investigating additional social media (e.g., Telegram or Medium), especially those that enable interaction between independent users and the legitimacy accruing entity, might be an interesting avenue for future research. Second, our study merely considered the bare number of tweets and forum posts and no other features such as the nature of their content or user reactions (e.g., retweets or replies to forum posts). Exploring the specific content of these messages might lead to further insightful discoveries. We intend to apply text mining techniques, such as sentiment analysis and latent dirichlet allocation, to substantiate our understanding regarding the specific effects of different forum discussions and company communication. Third, our study deliberately limits the loci of legitimacy to the realm of social media. The influences of other legitimacy sources like traditional media or third party authorities are only briefly touched upon. Additionally, organizational legitimacy is not directly observable (Zimmerman & Zeitz, 2002) and underpricing can only assist as a proxy for legitimacy, although a very established one. It would be interesting to further investigate the influence of other sources of legitimacy in the context of ICOs and especially compare their competing relevance to that of social media, potentially using other proxies. Fourth, additional control variables (e.g., circulating supply or additional media outlets) might be able to further explain some of the underpricing effects observed in our study. Considering the recent rise in popularity of ICO vesting schemes and, thus, limited circulating supplies, the relevance of this variable is particularly rising. An additional limitation is the sample size of 95 ICOs. Although this represented a considerable amount of all ICOs ever performed at the time of data collection, in a future study more ICOs could be included in the sample, as more ICOs finish every month and the potential sample size is growing quickly.

This would also allow to integrate additional control variables without compromising generalizability and to verify the conclusions of our study. Lastly, it needs to be stressed that our empirical findings are correlative in nature. While the referenced legitimacy theory substantiates the assumed effective direction, the empirical analysis only demonstrates the association of these constructs. Thus, as often called for but seldom realized (Aral, Dellarocas, & Godes, 2013), experiments are needed to ultimately confirm causality.

Chapter 5

Discussion and Conclusion⁶

This dissertation has started out by elaborating the motivation for this work in its initial chapter and subsequently presenting the foundations and related work for the two research topics in the following chapter. The focus of these topics have been the privacy-preserving protection of IoT sensor data and the role of information in blockchain-based fundraising. The details of our investigations into these topics, including the conceptual and theoretical background grounding our research and the results found, have been presented in two separate succeeding chapters. This chapter reiterates with a short introduction to the general context and motivation of our research. It then reviews and recapitulates the key findings of the two research studies and outlines their implications for research and practice. Additionally, the main limitations of our work as well as promising avenues for future research are presented. Finally, this chapter ends with some concluding remarks on the work presented.

5.1 Summary and Key Findings

Over the last years, the IoT has seen massive adoption both from consumers as well as businesses and billions of things are connected to the Internet today (Gartner, 2018; Oberländer et al., 2018). While the growth of the IoT has been impressive, it is still at its early stages and researchers and practitioners alike expect the number of connected devices and their relevance for humankind to keep increasing quickly in the near future (Gartner, 2018; Oberländer et al.,

⁶ Parts of this chapter, which are not further demarcated in the text, were initially published in the context of the following academic publications: Chanson, Gjoen, et al. (2018), Chanson et al. (2019) and Chanson et al. (2020).

2018). Consequently, the IoT is often counted among the most important contemporary developments in technology (Atzori et al., 2010; Jeschke et al., 2017; Oberländer et al., 2018). Currently, the IoT is predominantly built leveraging devices that are merely *connected* to the Internet and *generate data*. Based on these devices and the information they collect, new products, services and business models can be created. However, recent technological advancements have shaped the vision of networks of more powerful IoT devices which are directly *interconnected* between each other and can *interpret data* they receive independently (Beck et al., 2016; Davidsen et al., 2019; Kouzinopoulos et al., 2018). Instead of simply sensing, processing and exchanging data, independent IoT devices are enabled to make their own decisions based on information they interpret and, thus, become autonomous agents in an economic system. This enriched vision of the IoT is often referred to as the EoT. It gives rise to completely new business models and challenges current conceptions of relationships between humans and things, as illustrated for example through the popular case of robo-taxis (Antonopoulos, 2016; Musk, 2019; D. Tapscott & Tapscott, 2016b). As these enhanced IoT devices become increasingly autonomous, we might start to think of them as independent actors, that have their own will and are not necessarily controlled by a distinct owner.

The vision of the EoT is fueled by innovations which could lead to the technical infrastructure necessary for networks of economically acting things. In particular, blockchain technology is put forward by a wide range of research scholars and practitioners as a potential foundation for the infrastructure that enables things to become economic actors and interact directly with each other (Beck, Müller-Bloch, & King, 2018; D. Tapscott & Tapscott, 2016a). For example, the possibility to transact value (e.g., via cryptocurrencies) or incorporate business logic (e.g., via smart contracts) in an automated manner independent of intermediaries are core technological building blocks to develop an EoT and are enabled distinctively by blockchain technology. Ultimately, the decentralized nature of blockchains facilitates commercial activities between distrustful or anonymous agents, which is essential for the EoT (Nærland et al., 2017). Accordingly, a number of research projects and companies currently explore the usefulness of blockchain technology for different IoT systems and build extended prototypes or full-fledged applications, for example in the area of supply chains, energy markets and mobility (Curtis, 2015; Mengelkamp et al., 2018; Modum, 2018).

In this dissertation, we examine the role of blockchain technology as a potential infrastructure layer for such a visionary EoT. Two aspects are particularly relevant in this investigation, namely DApps, the applications that form this new world, and DeFi, the financial infrastructure that enables these DApps to be built. DApps are blockchain-based applications and, because they can inherit key features of the blockchain technology they leverage, DApps have the potential to be

more transparent and resilient than conventional centralized software (Chanson et al., 2019). DeFi is the blockchain-based financial infrastructure that enables DApps to be funded and unlocks complex financial services such as the issuing of debt or currency trading to these applications. In the past, DApps have mainly leveraged DeFi to fund their development and operations and to create complex incentive mechanisms for participants in the ecosystem of a DApp. However, as the whole space is still in its very early beginnings, we can expect many variations of such use cases to emerge with time.

While blockchain technology is widely credited with a large potential for the infrastructure layer of the EoT, this dissertation identified a lack of research to substantiate these claims. In particular, there is a research gap regarding actionable guidelines for the development of secure IoT applications, especially regarding those applications that leverage blockchain technology (Bélanger & Crossler, 2011; Pavlou, 2011; Rossi, Mueller-Bloch, Thatcher, & Beck, 2019). Furthermore, novel DeFi services are relevant for the development of DApps and build a cornerstone of the emerging EoT, however, because of their novelty there is only very little research considering these new financial services. Consequently, we have explored both the topics of DApps and DeFi in this dissertation. Regarding our investigations of DApps, we focused on systems that enable the protection of sensor data. Regarding the realm of DeFi, we concentrated on the equity-like funding of DApps. The main results of these investigations are discussed in the following.

5.1.1 Decentralized Applications

To address the research objective outlined in Section 1.2.1 and explore what design theory should guide the development of DApps that are able to protect IoT sensor data in a privacy-preserving manner, we conduct a comprehensive DSR study. The advancement of design knowledge has a long tradition among IS researchers and is considered highly relevant for both research and practice by senior scholars (Baskerville, 2008; Hevner et al., 2004; Winter, 2008). In this dissertation, we base our specific DSR approach on the guidelines of Peffers, Tuunanen, Rothenberger, and Chatterjee (2007). We derive an artifact that consists of a set of interrelated design requirements, design principles, and design features. Furthermore, we demonstrate and refine our artifact on the basis of an instantiation that aims to prevent the fraudulent manipulation of car mileage data. We build this instantiation on the basis of three iterative development cycles. In particular, this instantiation is exposed in several field tests, including one that ran over 4 months with 100 participating cars. We provide an evaluation of this instantiation and the overall artifact and extend Peffers et al.'s (2007) guidelines by considering an additional phase of ex-post evaluation, reflecting two additional use cases in the fields of pharmaceutical supply chains and

energy microgrids. Finally, we summarize the insights generated by presenting our results in the form of a design theory.

We start out by reflecting the results concerning our first research question:

RQ 1a: Which fundamental challenges arise in the context of IoT sensor data protection, and which requirements can be derived from these challenges for the design of information systems that facilitate IoT sensor data protection (i.e., SDPSs)?

To address this research question, we lay out the practical issues which are identified in the existing literature and which stimulated the launch of this study. We review the related literature regarding security and privacy in the IoT, with a special focus on sensor data protection and existing research on blockchain technology for SDPSs. Based on this body of knowledge, we extract some of the most relevant challenges associated with SDPSs and transform these into system requirements. Thus, we provide the groundwork for understanding shortcomings in current SDPS architectures and enable the direction of our research towards potential improvements. Based on these analyses the following key finding, and its implications for research and practice is highlighted:

IoT systems face a set of fundamental challenges concerning the protection of data, and these challenges can be conceptualized as four design requirements defining the problem class of SDPSs. The first challenge we derive from literature is that data integrity cannot be guaranteed in most IoT systems because adversaries can manipulate sensor data at several stages in the processing pipeline. We translate this challenge into DP1 which states that SDPS should ensure tamper resistance throughout the whole data pipeline. The second major challenge we review is that the data gathered by IoT sensors is typically of very high resolution and often very personal and sensitive. This challenge is reflected in DP2 which calls for SDPSs that should be capable of preserving the privacy of the data owner. The third key challenge of IoT systems is that they can generate huge amounts of data, which is especially pronounced if a novel technology with scalability issues such as blockchain is considered as part of a solution. We convert this challenge into DP3 noting that SDPSs should provide sufficient data throughput to process large amounts of data. Finally, we discuss the challenge that the protection of IoT sensor data can demand considerable resources both from a technical as well as economical perspective. Consequently, we formulate DP4 that SDPSs should ensure economic feasibility, that is, the protection benefits have to outweigh the protection costs. Note that our study of SDPS challenges confirms earlier evidence from literature that the distributed, multilayered nature of IoT systems, as well as IoT ecosystems with multiple parties and potentially diverging interests, introduces very specific and particularly serious challenges regarding data protection (C. C. Aggarwal et al.,

2013; Lowry et al., 2017). By defining the four requirements described above, we establish a relevant and broad problem class on which we focus our investigations. As such, we contribute to the literature by explicitly conceptualizing the challenges of IoT data protection as a set of holistic design requirements. The generalizability within our wide problem class constitutes an important foundation for our theoretical contribution. To ensure such a generalizability of our results, the design requirements serve as the basis to deduce the testable propositions of our design theory. Furthermore, these requirements can inform future research, for example, investigating how their fulfillment affects the adoption of IoT IS. As such, we also support practitioners and enable novel relevant insights by highlighting a problem class central to the success of IoT applications. In particular, the ability of organizations to offer systems that ensure adequate security levels while guaranteeing sufficient user privacy seems highly relevant for the adoption of smart products (Sicari et al., 2015).

We continue by considering the results regarding our second research question:

RQ 1b: Which actionable guidelines in the form of design principles and design features address these design requirements and inform the development of SDPS?

To address this research question, we base our investigations on the design requirements derived before and leverage theory of information asymmetry and certification to derive initial design principles and features (Akerlof, 1970; Albersmeier et al., 2009; Bond, 1982; Genesove, 1993; Spence, 1973). As our study progresses, these design principles are continuously refined and become increasingly informed by the knowledge gathered from our iterations of DSR and the instantiation and according evaluation of the artifact. The deep theoretical grounding combined with insights from practice enable us to derive a set of purposive guidelines for the design of SDPSs. Based on this work the following two key findings, and their implications for research and practice are highlighted:

Blockchain technology is not a universal remedy for data protection, instead, considering four key design principles is essential to build purposeful SDPSs. The first fundamental implication that should be reflected (DP1) is that sensor data should be certified on the basis of source to sink protection. Currently, this represents an ideal condition which is not yet realizable per-se. As sensors are not yet blockchain-enabled (i.e., able to communicate with a blockchain) there is always going to be a certain part of the data pipeline unprotected by blockchain immutability, namely between the sensor that records the data and the device that signs the blockchain transactions. However, this logical distance can be designed to minimize the attack vectors or, more figuratively, be kept as short as possible. In the future, blockchain-enabled chips that simultaneously sense and communicate with a blockchain might simplify this process

considerably. The second fundamental implication (DP2) is that additional secondary data should be leveraged to cross-validate the original information of interest. On the one hand, this is tied to the realization that complete compliance with DP1 is difficult to achieve, on the other hand, even DP1 cannot prevent a manipulation that induces a sensor to read technically correct but misleading data (e.g., cooling just the surroundings of a sensor instead of an entire package in a cold chain application). As such, it is important to note that even utmost care with the application of blockchain technology alone cannot guarantee tamper-proof processes. Additionally, we stress that the data gathered for cross-validation and, correspondingly, the certification procedure are highly case-specific and require major creative adaptations for different use cases. However, as IoT sensors are usually not deployed alone, we expect that IoT systems usually should be able to gather enough diverse datasets to enable cross-validation. Third, it should be addressed that the data owners determine when and to what extent their data is communicated to others (DP3). With respect to this design principle, it is crucial to reflect that blockchains are, in principle, open immutable databases that can be read by anyone and have no possibility of deletions. Consequently, system designers have to be very careful about what type of data is saved in the blockchain and also what type of metadata these transactions generate (e.g., addresses used, frequency of transactions, etc.). In particular, sensor data itself cannot be stored in a public permissionless blockchain and a hybrid approach should be followed instead. As such, an off-chain access management service needs to be implemented to enable basic privacy mechanisms. Finally, the fourth principle insists on a linearly scalable system architecture which is especially challenging regarding the integration of blockchain technology in the architecture. Currently, scalability is still a major issue of all major public permissionless blockchains and, consequently, hybrid architecture approaches combining blockchain and traditional databases are necessary for industry relevant applications. This leads to a standard system architecture which stores raw data in traditional distributed databases and saves only a digital fingerprint (hash) of the original data on the blockchain to enforce the immutability of that data.

Our theory grounded DSR approach working towards a set of interrelated design requirements, principles and features has proven to generate relevant practical insights and enabled us to theorize beyond a single artifact instantiation in a meaningful way. In particular, the initial version of all our design requirements, principles and features are grounded in relevant theory. The design requirements are informed by the literature on IoT, security, and privacy. The design principles are based on theory on information asymmetry, privacy, and IS success. Finally, the design features leverage the literature on blockchain technology. We believe that this literature provides an informative basis for other researchers contributing to the body of knowledge on SDPSs. Combined with insights from practice, which originate from field tests and

their evaluation and an additional ex-post evaluation, this theoretical grounding allowed us to derive appropriate design guidelines for the development of SDPSs. Ultimately, the generalizability of our results beyond a specific artifact instantiation is enabled by the codification and abstraction of our research into these specific design guidelines, namely design requirements, design principles and design features. With this thorough conceptualization we also hope to inform future DSR studies and contribute to the further acceptance of this relevant research stream in IS.

Finally, we conclude by reflecting the insights regarding our third research question:

RQ 1c: What is the value proposition of blockchain technology in the realm of SDPSs, and what fundamental design implications of blockchain-based SDPSs must be considered?

To address this research question, we focused particularly on the literature of blockchain technology to inform the development of our design features. As such, major features of our artifact leverage blockchain technology and were thus investigated in practice in our field tests and the subsequent evaluations of the system. Furthermore, we put additional emphasis on the value proposition of blockchain technology and its design implications in our ex-post evaluation, which contributed in a major form to a number of our key insights. Based on this work the following two main findings, and their implications for research and practice are highlighted:

Well-designed blockchain-based IoT systems, including SDPSs, can inherit core characteristics of blockchain technology, such as immutability, accessibility or censorship resistance. The development of our design guidelines shows that if certain key design implications are followed, a rather complex SDPS can obtain some of the key features of the underlying blockchain technology it is built upon. First of all, one of the most relevant properties SDPSs can inherit from blockchain technology includes the immutability of gathered sensor raw data. Additionally, the decentralized nature of blockchains and their censorship resistance can also be transferred to the overall system to a certain extent, which reduces platform risks and renders an SDPS more acceptable in the perception of third parties. More generally, SDPSs can leverage the trust-free interaction blockchains enable to create similarly trust-free systems which leverage the underlying blockchain as a reference for a shared source of truth. However, we need to stress again that blockchain technology is not a one size fits all solution. The relevant properties are not simply adopted by the overall system because of the bare use of blockchain technology. Instead, SDPSs need to be designed carefully, reflecting the principles and features of our design theory, in order to reach their full potential.

Blockchain-based systems feature considerable disadvantages over traditional IT and are thus mainly effective in specific scenarios, where their advantages are especially relevant. One relevant parameter is the number of individual parties involved in an ecosystem. Although SDPSs can be used to protect simple data pipelines within single entities, they are most effective in complex multi-party ecosystems where data transverses multiple organizations and stakeholders. In the context of the IoT, the sensing systems itself often integrates hardware and software from multiple active stakeholders, creating a multi-party ecosystem even before considering the end customer. In such circumstances, blockchain-based systems can be specifically valuable, acting as a trusted shared source of ground truth being available to all parties equally. Ultimately, the equal access of all stakeholders to a blockchain is rooted in its decentralized structure. These aspects are also the reason why a blockchain-based system is often perceived as neutral and might therefore be more readily accepted as an industry standard than other systems as our investigations show. Finally, blockchains have inherently integrated high security provisions regarding the immutability and accessibility of data. While it is, in principle, possible to build systems with similar guarantees on a centralized basis, it requires considerable and very specialized resources. In contrast, blockchains offer these qualities out of the box and are ready to be used by any developer. Consequently, we believe that blockchain technology is especially valuable for rapid prototyping and small teams who are constrained in resources but require state-of-the-art security technology with corresponding high guarantees for the immutability and availability of data, or the correct execution of code in the case of smart contracts.

5.1.2 Decentralized Finance

To address the research objective outlined in Section 1.2.2 and explore how the information from social media relates to success in blockchain-based fundraising, we conduct a number of quantitative analyses of finished fundraises. In this dissertation, we focus on the effects in ICOs, which have seen the widest adoption among all DeFi applications beyond payments in the past. In particular, we leverage the existing body of knowledge concerning the legitimacy of organizations, in the context of fundraising and social media, as well as deep insights from the practice of blockchain-based fundraising, to develop a core set of hypotheses regarding the relation of social media and underpricing in ICOs. We collect a data set of 95 fundraises, including financial data and social media data from various sources, to test these hypotheses. Specifically, we include highly interactive discussion forums and rather non-interactive microblogs in our analyses. As the dependent variable we focus on the amount of underpricing observed in an ICO, which serves as a proxy for organizational legitimacy, which is in itself not directly observable.

We discuss our results in the context of the emerging stream of literature on blockchain-based fundraising and conclude with the limitations of our study and an outlook on future research.

We start out by reflecting the results concerning our first research question:

RQ 2a: How does information from interactive and non-interactive social media directly relate to the extent of underpricing in blockchain-based fundraising?

To address this research question, we focus, on the one hand, on the main public communication channel of firms for their dominantly one-directional communication towards their investors, which is Twitter. In particular, we record the number of tweets and the number of followers of a firm prior to their fundraising activity. On the other hand, we include interactive social media by integrating discussion forums in our investigations. Specifically, we record the number of threads that a firm was mentioned in on selected online discussion forums, namely Reddit and Bitcointalk. We then investigate, how these measurements individually relate to the extent of underpricing in ICOs. Based on these analyses the following key finding, and its implications for research and practice is highlighted:

Our studies find a highly significant correlation between the activity on public discussion forums and underpricing in blockchain-based fundraises. This confirms our hypothesis that increased activity of investors on interactive social media coincides with higher underpricing. Furthermore, this substantiates the relevance of UGC, meaning information produced by average users rather than the organization raising funds or marketers hired on its behalf. Note that our investigations could not confirm the hypothesis that also the strategic emission of information by the fund-seeking firm on microblogs increases observed underpricing. As such, our results relate to earlier key IS contributions which found that UGC is more relevant than MGC in driving purchases (Goh et al., 2013) or that the impact of forum contributions on the price of Bitcoin is larger than the effect of microblogs (Mai et al., 2018). It is especially compelling to study such and associated effects in the realm of DeFi because the novelty of the sector leads to a very distinct media environment. In particular, social media are a major source of information and relevant signal in the context of blockchain-based fundraising. In contrast, the relevance of social media is limited in mature traditional financial markets, which feature established information sources such as audited financial reports, statements of company representatives or specialized news portals (e.g., Bloomberg or Reuters). For practitioners, our results indicate that investors judging the prospects of an opportunity mainly value information from peers, rather than one-directional communication dominated by marketers.

We continue by considering the results regarding our second research question:

RQ 2b: How is information from interactive and non-interactive social media intertwined in its relation to the extent of underpricing in blockchain-based fundraising?

To address this research question, we perform a number of additional analyses beyond the investigation of the direct relationship of tweets, followers and threads to the level of underpricing. Specifically, we examine how the strategic usage of Twitter by firms influences the activity on discussion forums. Additionally, we consider if discussion forum activity has a mediating effect on the relationship between strategic Twitter management and underpricing in blockchain-based fundraises. Based on these analyses the following key finding, and its implications for research and practice is highlighted:

The effect of the strategic usage of Twitter of a firm prior to its fundraise is mediated by the activity from investors on discussion forums. Thus, our results indicate that it is in fact the activity of investors on discussion forums driving the underpricing in a fundraise, rather than the usage of Twitter by the firm which seeks to raise those funds. However, our analysis indicates that the firm can leverage its Twitter management to influence the relevant activity of third parties on discussion forums and, in turn, indirectly influence the underpricing. While our empirical analysis can demonstrate the association of the underlying constructs, it is strictly correlative in nature and cannot confirm any causalities. However, the theoretical grounding leading to the hypotheses of our study as well as our insights from practice render a causal relationship plausible. For practitioners this could indicate key consequences, for example a focus of the strategic organizational use of Twitter towards establishing legitimacy by influencing the behavior of investors on other social media.

5.2 Limitations and Future Research

The promising findings and implications of this dissertation should be diligently assessed in the light of their limitations, which have been discussed in detail in Chapters 3 and 4. This section reexamines some of the major challenges of the research setting and the corresponding limitations of our work in both the realm of DApps and DeFi. Furthermore, additional limitations of this dissertation are emphasized and promising avenues to advance the research topics introduced in this dissertation are presented.

With regards to the work of this dissertation on DApps, our research towards developing a comprehensive design theory for SDPSs presented in Chapter 3 comprises a number of limitations, which encourage future research on this topic. First of all, blockchain technology, which is at the center of our investigations and the design theory we develop, is a novel technology

that is experiencing rapid innovation in many different contexts. However, there are still open questions if and how some of the most fundamental concerns regarding this technology can be solved. While we leverage the most recent state-of-the-art knowledge on the technology available, and develop a design theory that copes with the current extensive limitations of blockchains, the proposed approach for SDPSs will likely need to be adjusted to novel technological developments in the near future. Especially the work around the scalability of blockchains could fundamentally impact some of the design principles we established. However, based on the current state of technology we have made careful tradeoffs between security, privacy and scalability, which is reflected in our design requirements, principles and features, in particular concerning the approach regarding a hybrid on-off-chain architecture. A second important limitation concerns the evaluation of our design theory. While we developed and evaluated an instantiation of our artifact, CertifiCar, in several field tests and performed an additional ex-post evaluation based on further prominent IoT data protection use cases, a broader quantitative evaluation is desirable.

With regards to the work of this dissertation on DeFi, our research on the relation of information from social media and the success of blockchain-based fundraising presented in Chapter 4 features additional limitations, which motivate future research related to this subject. First of all, we acknowledge that the specific DeFi mechanism investigated in this dissertation, namely the ICO, has lost a large share of the momentum it held when our examinations were conducted. However, the historical magnitude of interest in ICOs is still unmatched by any other DeFi mechanism at any point in time up to date. As such, ICOs emerge as a core theme of study of early DeFi services. Additionally, DeFi services have united much of the innovation in the blockchain space in the last two years and many of the most used DApps are currently of financial nature. Consequently, it still holds that DeFi is one of the most relevant spaces in the context of DApps and blockchain technology. We suppose that many of our fundamental investigations regarding the effects of information on fundraising success can also inform interactions around other DeFi services. As such, we believe that our early studies pose fruitful basis for further research targeting recently emerging services such as decentralized debt financing or the blockchain-based issuance of more complex synthetic financial products. Furthermore, all our empirical analyses are only able to reveal correlative relations. While the theoretical grounding leading to our hypotheses substantiate the assumed effective directions, ultimately, only experiments could truly confirm causality.

Finally, we recognize further potential for promising extensions to our work in several strategic areas beyond these specific limitations of our studies. As our studies on both DApps and DeFi reflect, a major advantage of blockchain technology is that it can facilitate collaboration beyond traditional ecosystem borders. As such, we hope that this dissertation serves as an

inspiration and impulse for additional investigations on how blockchain technology enables new modes of collaborations across organizational boundaries, for example, by establishing security, privacy and trust. In particular, we call for research examining the business potential of various emerging blockchain-based data protection approaches (Risius & Spohrer, 2017). Furthermore, we extend this call to generally motivate more research that bridges the gap between technology and business in the area of blockchain-based IT solutions. While the technology is undergoing rapid progress and various industry initiatives, such as the Trusted IoT Alliance or R3, are emerging, the business potential of the technology still remains vague. Both researchers and practitioners have been overly optimistic regarding the short and medium business potential of blockchain technology in the past (Beck et al., 2017). While the overall hype around the technology has decreased with less performant cryptocurrency markets, discussions are still often either predominantly technology-focused or merely business-driven without an appropriate reflection of the actual capabilities and limitations of the technology. Consequently, we call IS researchers whose expertise is centered on this intersection of business and technology to fill this research gap articulated and link business problems to blockchain technology and according applications.

5.3 Conclusion

Over the last decade, the IoT has become a reality with billions of connected devices deployed today and strong growth projected ahead for the next few years. While current IoT systems are predominantly based on devices which are merely connected to the Internet and generate data, an enhanced vision of the IoT has emerged based on interconnected things that can interpret data and take economic decisions autonomously. This enhanced vision is often referred to as the EoT. While researchers and practitioners alike have made it a key purpose to develop first functional EoT systems, the technological foundation for such systems is only just emerging. Blockchain technology is often named as a potential basis for the foundational infrastructure of the EoT. In fact, the decentralized nature of blockchain technology allows trustless interactions and makes it possible to use a blockchain as a shared, secure computing platform between distrustful or anonymous agents. Ultimately, this enables commercial activities between unknown agents, humans or things, which is the fundamental basis of an EoT. However, as blockchain technology and the concept of the EoT are both nascent, there is only little scientific evidence detailing the potential of blockchain for the EoT based on academic studies.

In this dissertation, we explore the role of blockchain technology as a potential infrastructure layer for the EoT. In particular, we investigate how both DApps and DeFi could be leveraged as foundational technology for the EoT. Regarding DApps, we focus on how data can be gathered and exchanged securely and in a privacy-preserving manner. In particular, we develop specific actionable guidelines that can facilitate the development of SDPSs, even with the apparent current restrictions of blockchain technology. Our results show that blockchain technology is not a universal remedy for data protection, but that it is in fact possible to develop DApps that balance the need for security, privacy, scalability and cost-efficiency in a practical way. We find that blockchain technology is especially effective in specific scenarios, where it might facilitate the collaboration over organizational boundaries and pave the way for industry standards that are perceived to be neutral. Regarding DeFi, we focus on how innovation geared towards the EoT is funded. In particular, we show that fundraising activities of DApps are uniquely related to information divulged on discussion forums and that fundraising firms might be able to stimulate that discourse via microblogs. Ultimately, further research is needed to explore how the rapid development of blockchain technology might change the strategies to build effective DApps and if our results on blockchain-based fundraising apply to the vivid broader developments in DeFi.

References

- Aave. (2020). Aave Protocol Whitepaper. Retrieved February 28, 2020, from https://github.com/aave/aave-protocol/blob/master/docs/Aave_Protocol_Whitepaper_v1_0.pdf
- Abbasi, A., Sarker, S., & Chiang, R. H. L. (2016). Big data research in information systems: Toward an inclusive research agenda. *Journal of the Association for Information Systems*, 17(2), i–xxxii.
- Adams, H. (2019). Uniswap Whitepaper. Retrieved July 17, 2019, from https://hackmd.io/C-DvwDSfSxuh-Gd4WKE_ig
- Aggarwal, C. C., Ashish, N., & Sheth, A. (2013). The internet of things: A survey from the data-centric perspective. In C. C. Aggarwal (Ed.), *Managing and mining sensor data* (pp. 383–428). Berlin, Germany: Springer Science+Business Media.
- Aggarwal, R., Gopal, R., Gupta, A., & Singh, H. (2012). Putting money where the mouths are: The relation between venture financing and electronic word-of-mouth. *Information Systems Research*, 23(3), 976–992.
- Aggarwal, R. K., Krigman, L., & Womack, K. L. (2002). Strategic IPO underpricing, information momentum, and lockup expiration selling. *Journal of Financial Economics*, 66(1), 105–137.
- Aggarwal, R., & Singh, H. (2013). Differential Influence of Blogs Across Different Stages of Decision Making: The Case of Venture Capitalists. *MIS Quarterly*, 37(4), 1093–1112.
- Akerlof, G. A. (1970). The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500.
- Albersmeier, F., Schulze, H., Jahn, G., & Spiller, A. (2009). The reliability of third-party certification in the food chain: From checklists to risk-oriented auditing. *Food Control*, 20(10), 927–935.
- Aldrich, H. E., & Fiol, C. M. (1994). Fools rush in? The institutional context of industry creation. *Academy of Management Review*, 19(4), 645–670.
- Allen, F., & Faulhaber, G. R. (1989). Signalling by underpricing in the IPO market. *Journal of Financial Economics*, 23(2), 303–323.

- Alqassem, I., & Svetinovic, D. (2014). A taxonomy of security and privacy requirements for the Internet of Things (IoT). In *Proceedings of the 2014 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)* (pp. 1244–1248). Bandar Sunway, Malaysia.
- Amsden, Z., Arora, R., Bano, S., Baudet, M., Blackshear, S., Bothra, A., ... Zhou, R. (2019). The Libra Blockchain. Retrieved July 18, 2019, from <https://libra.org/en-US/white-paper/>
- Anderson, C., Baskerville, R. L., & Kaul, M. (2017). Information Security Control Theory: Achieving a Sustainable Reconciliation Between Sharing and Protecting the Privacy of Information. *Journal of Management Information Systems*, 34(4), 1082–1112.
- AndreessenHorowitz. (2019). a16zCrypto: Crypto Portfolio. Retrieved July 22, 2019, from <https://a16z.com/crypto/#vertical-landing-portfolio>
- Antonopoulos, A. M. (2014). *Mastering Bitcoin: Unlocking Digital Cryptocurrencies*. Sebastopol, CA: O'Reilly Media.
- Antonopoulos, A. M. (2016). *The Internet of Money - Volume 1*. Seattle, OR: Merkle Bloom.
- Antonopoulos, A. M., & Wood, G. (2018). *Mastering Ethereum: Building Smart Contracts and DApps*. Sebastopol, CA: O'Reilly Media.
- Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of internet stock message boards. *The Journal of Finance*, 59(3), 1259–1294.
- Aral, S., Dellarocas, C., & Godes, D. (2013). Introduction to the special issue—social media and business transformation: a framework for research. *Information Systems Research*, 24(1), 3–13.
- Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805.
- AuroraLabs. (2019). IDEX: A Real-Time and High-Throughput Ethereum Smart Contract Exchange. Retrieved July 18, 2019, from <https://idex.market/static/IDEX-Whitepaper-V0.7.6.pdf>
- Avital, M., Beck, R., King, J., Rossi, M., & Teigland, R. (2016). Jumping on the Blockchain Bandwagon: Lessons of the Past and Outlook to the Future. In *Proceedings of the 37th International Conference on Information Systems (ICIS)*. Dublin, Ireland.
- Ayoade, G., Karande, V., Khan, L., & Hamlen, K. (2018). Decentralized IoT Data Management Using Blockchain and Trusted Execution Environment. In *Proceedings of the 2018 IEEE*

- International Conference on Information Reuse and Integration (IRI)* (pp. 15–22). Salt Lake City, UT.
- Back, A. (2002). Hashcash - A Denial of Service Counter-Measure. Retrieved July 10, 2019, from <http://www.hashcash.org/papers/hashcash.pdf>
- Baron, R. M., & Kenny, D. A. (1986). The moderator--mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, *51*(6), 1173.
- Baskerville, R. (2008). What design science is not. *European Journal of Information Systems*, *17*(5), 441–443.
- Baskerville, R., Kaul, M., & Storey, V. (2015). Genres of Inquiry in Design-Science Research: Justification and Evaluation of Knowledge Production. *MIS Quarterly*, *39*(3), 541–564.
- Baskerville, R., & Pries-Heje, J. (2010). Explanatory Design Theory. *Business & Information Systems Engineering*, *2*(5), 271–282.
- Baskerville, R., & Siponen, M. (2002). An information security meta-policy for emergent organizations. *Logistics Information Management*, *15*(5/6), 337–346.
- Beck, R., Avital, M., Rossi, M., & Thatcher, J. B. (2017). Blockchain Technology in Business and Information Systems Research. *Business & Information Systems Engineering*, *59*(6), 381–384.
- Beck, R., & Müller-Bloch, C. (2017). Blockchain as Radical Innovation: A Framework for Engaging with Distributed Ledgers. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)* (pp. 5390–5399). Waikoloa, USA.
- Beck, R., Müller-Bloch, C., & King, J. L. (2018). Governance in the Blockchain Economy: A Framework and Research Agenda. *Journal of the Association for Information Systems*, *19*(10), 1020–1034.
- Beck, R., Müller-Bloch, C., & Ling, L. J. (2018). Governance in the Blockchain Economy: A Framework and Research Agenda. *Journal of the Association for Information Systems*, in press.
- Beck, R., Stenum Czepluch, J., Lollike, N., & Malone, S. (2016). Blockchain - The Gateway to trust-free cryptographic Transactions. In *Proceedings of the 24th European Conference on Information Systems (ECIS)*. Istanbul, Turkey.
- Beck, R., Weber, S., & Gregory, R. W. (2013). Theory-generating design science research.

- Information Systems Frontiers*, 15(4), 637–651.
- Bélangier, F., & Crossler, R. E. (2011). Privacy in the digital age: a review of information privacy research in information systems. *MIS Quarterly*, 35(4), 1017–1042.
- Bilgeri, D., Wortmann, F., & Fleisch, E. (2017). How Digital Transformation Affects Large Manufacturing Companies' Organization. In *Proceedings of the Thirty-Eighth International Conference on Information Systems (ICIS)*. Seoul, South Korea.
- Bitnodes. (2019). Global Bitcoin Nodes Distribution. Retrieved July 11, 2019, from <https://bitnodes.earn.com/>
- Bogner, A., Chanson, M., & Meeuw, A. (2016). A Decentralised Sharing App Running a Smart Contract on the Ethereum Blockchain. In *Proceedings of the 6th International Conference on the Internet of Things*. Stuttgart, Germany.
- Bond, E. W. (1982). A direct test of the "Lemons" model: The market for used pickup trucks. *The American Economic Review*, 72(4), 836–840.
- Bonvin, N. (2012). *Linear Scalability of Distributed Applications*. École Polytechnique Fédérale de Lausanne, Thèse No. 5278.
- Boos, D., Guenter, H., Grote, G., & Kinder, K. (2013). Controllable accountabilities: the internet of things and its challenges for organisations. *Behaviour & Information Technology*, 32(5), 449–467.
- Booth, J. R., & Chua, L. (1996). Ownership dispersion, costly information, and IPO underpricing. *Journal of Financial Economics*, 41(2), 291–310.
- Bosche, A., Crawford, D., Jackson, D., Schallehn, M., & Schorling, C. (2018). Unlocking Opportunities in the Internet of Things. Retrieved June 17, 2019, from <https://www.bain.com/insights/unlocking-opportunities-in-the-internet-of-things/>
- Brau, J. C., & Fawcett, S. E. (2006). Initial public offerings: An analysis of theory and practice. *Journal of Finance*, 61(1), 399–436. <https://doi.org/10.1111/j.1540-6261.2006.00840.x>
- Brelie, J. von der, & Giehl, S. (2019). Internet of Things: Die neuen Geschäftsmodelle der Machine to Machine Economy. Retrieved July 26, 2019, from <https://www.zuehlke.com/blog/die-neuen-geschaeftsmodelle-der-machine-to-machine-economy/>
- Breznitz, D., Forman, C., & Wen, W. (2018). The Role of Venture Capital in the Formation of a new Technological Ecosystem: Evidence from the Cloud. *MIS Quarterly*, 42(4).

- Brynjolfsson, E., & McAfee, A. (2012). *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*. Lexington, MA: Digital Frontier Press.
- Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: an empirical study of rationality-based beliefs and information security awareness. *MIS Quarterly*, 34(3), 523–548.
- Burtch, G., & Chan, J. (2019). Investigating the relationship between medical crowdfunding and personal bankruptcy in the United States: Evidence of a digital divide. *MIS Quarterly*, 43(1), 237–262.
- Buterin, V. (2013). Ethereum White Paper. Retrieved September 28, 2017, from <https://github.com/ethereum/wiki/wiki/White-Paper>
- Cai, W., Wang, Z., Ernst, J. B., Hong, Z., Feng, C., & Leung, V. C. M. (2018). Decentralized applications: The blockchain-empowered software system. *IEEE Access*, 6, 53019–53033.
- Car-Pass. (2018). Car-Pass is your guarantee of an accurate odometer. Retrieved January 5, 2018, from <https://www.car-pass.be/en/about-car-pass>
- Carfax. (2018). Buying Used American Cars? Check the Carfax Report. Retrieved January 13, 2018, from <https://www.carfax.eu/de>
- CarJam. (2018). CarJam. Vehicle Facts, History, Money Owing and more. Retrieved January 17, 2018, from <https://www.carjam.co.nz/>
- Castillo, M. del. (2018). Big Blockchain: The 50 Largest Public Companies Exploring Blockchain. Retrieved July 23, 2019, from <https://www.forbes.com/sites/michaeldelcastillo/2018/07/03/big-blockchain-the-50-largest-public-companies-exploring-blockchain/#1148c0c62b5b>
- Castro, M., & Liskov, B. (1999). Practical Byzantine Fault Tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation* (pp. 173–186). New Orleans, LA.
- Ceccagnoli, M., Forman, C., Huang, P., & Wu, D. J. (2012). Co-creation of value in a platform ecosystem: The case of enterprise software. *MIS Quarterly*, 36(1), 263–290.
- Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.-Y., & Moon, S. (2007). I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system. In *Proceedings of the 7th ACM Conference on Internet Measurement* (pp. 1–14). San Diego, CA.

- Chandra Kruse, L., Seidel, S., & Gregor, S. (2015). Prescriptive knowledge in IS research: Conceptualizing design principles in terms of materiality, action, and boundary conditions. In *Proceedings of the 48th Hawaii International Conference on System Sciences (HICSS)* (pp. 4039–4048). Kauai, USA.
- Chandra Kruse, L., Seidel, S., & Purao, S. (2016). Making Use of Design Principles. In *Proceedings of the International Conference on Design Science Research in Information Systems and Technology (DESRIST)* (pp. 37–51). St. John's, Canada.
- Chanson, M., Bogner, A., Bilgeri, D., Fleisch, E., & Wortmann, F. (2019). Blockchain for the IoT: Privacy-Preserving Protection of Sensor Data. *Journal of the Association for Information Systems (JAIS)*, 20(9), 1274–1309.
- Chanson, M., Gjoen, J., Risius, M., & Wortmann, F. (2018). Initial Coin Offerings (ICOs): The role of Social Media for Organizational Legitimacy and Underpricing. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. San Francisco, CA.
- Chanson, M., Martens, N., & Wortmann, F. (2020). The Role of User-Generated Content for Blockchain-Based Decentralized Finance. In *Proceedings of the 28th European Conference on Information Systems (ECIS)*. Marrakesh, Morocco.
- Chanson, M., Risius, M., & Wortmann, F. (2018). Initial Coin Offerings (ICOs): An Introduction to the Novel Funding Mechanism Based on Blockchain Technology. In *Proceedings of the 24th Americas Conference on Information Systems (AMCIS)*. New Orleans, LA.
- Chaparro, F. (2019). Paradigm backs decentralized exchange protocol Uniswap. Retrieved July 19, 2019, from <https://www.theblockcrypto.com/2019/04/23/paradigm-backs-decentralized-exchange-protocol-uniswap/>
- Chatterjee, S., Sarker, S., & Valacich, J. S. (2015). The behavioral roots of information systems security: Exploring key factors related to unethical IT use. *Journal of Management Information Systems*, 31(4), 49–87.
- Chellappa, R. K., & Pavlou, P. A. (2002). Perceived information security, financial liability and consumer trust in electronic commerce transactions. *Logistics Information Management*, 15(5/6), 358–368.
- Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business Intelligence and Analytics: From Big Data to Big Impact. *MIS Quarterly*, 36(4), 1165–1188.
- Chen, Y., & Bellavitis, C. (2020). Blockchain Disruption and Decentralized Finance: The Rise of Decentralized Business Models. *Journal of Business Venturing Insights*, 13, e00151.

- Chen, Y., & Zahedi, F. M. (2016). Individuals' Internet Security Perceptions and Behaviors: Polycontextual Contrasts Between the United States and China. *Mis Quarterly*, 40(1), 205–222.
- Christensen, R., Mushegian, N., Brockman, D., Rowe, K., Milenius, A., & Zurrer, R. (2018). The Dai Stablecoin System. *Whitepaper*.
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and Smart Contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303.
- Chui, M., Löffler, M., & Roberts, R. (2010). The Internet of Things. *McKinsey Quarterly*, 2010(2), 1–9.
- Clark, M., & Melancon, J. (2013). The influence of social media investment on relational outcomes: A relationship marketing perspective. *International Journal of Marketing Studies*, 5(4), 132–142.
- Clarke, R. (2016). Big data, big risks. *Information Systems Journal*, 26(1), 77–90.
- Cliff, M. T., & Denis, D. J. (2004). Do initial public offering firms purchase analyst coverage with underpricing? *The Journal of Finance*, 59(6), 2871–2901.
- CoinDesk. (2019). ICO Tracker. Retrieved February 26, 2019, from <https://www.coindesk.com/>
- CoinDesk. (2018). ICO Tracker. Retrieved April 22, 2018, from <https://www.coindesk.com/ico-tracker>
- CoinMarketCap. (2018). Cryptocurrency Market Capitalizations. Retrieved March 8, 2018, from <https://coinmarketcap.com>
- Conti, M., Kumar, S., Lal, C., & Ruj, S. (2018). A survey on security and privacy issues of bitcoin. *IEEE Communications Surveys & Tutorials*, in press.
- Cram, W. A., Proudfoot, J. G., & D'Arcy, J. (2017). Organizational information security policies: a review and research framework. *European Journal of Information Systems*, 26(6), 605–641.
- Croman, K., Decker, C., Eyal, I., Gencer, A. E., Juels, A., Kosba, A., ... Wattenhofer, R. (2016). On Scaling Decentralized Blockchains (A Position Paper). In *Proceedings of the International Conference on Financial Cryptography and Data Security* (pp. 106–125). Christ Church, BB.
- Crossler, R. E., & Posey, C. (2017). Robbing Peter to pay Paul: Surrendering privacy for security's sake in an identity ecosystem. *Journal of the Association for Information Systems*,

- 18(7), 487.
- Curtis, S. (2015). Visa uses Bitcoin's blockchain technology to cut paperwork out of car leasing. Retrieved March 18, 2018, from <https://www.telegraph.co.uk/technology/news/11961296/Visa-uses-bitcoins-blockchain-technology-to-cut-paperwork-out-of-car-rental.html>
- Dacin, M. T., Oliver, C., & Roy, J.-P. (2007). The legitimacy of strategic alliances: An institutional perspective. *Strategic Management Journal*, 28(2), 169–187.
- Dannen, C. (2017). Cryptoeconomics Survey. In *Introducing Ethereum and Solidity: Foundations of Cryptocurrency and Blockchain Programming for Beginners* (pp. 139–147). Berkeley, CA: Apress.
- Das, S. R., & Chen, M. Y. (2007). Yahoo! for Amazon: Sentiment extraction from small talk on the web. *Management Science*, 53(9), 1375–1388.
- Davenport, T. H. (2013). Analytics 3.0. *Harvard Business Review*, 91(12), 64–72.
- Daidsen, M., Gajek, S., Kruse, M., & Thomsen, S. (2019). Empowering the Economy of Things. Retrieved March 8, 2020, from http://weeve.network/documents/Weeve_Technical_Whitepaper_2019.pdf
- Decker, C., & Wattenhofer, R. (2013). Information Propagation in the Bitcoin Network. In *IEEE P2P 2013 Proceedings*. Trento, IT.
- Decker, C., & Wattenhofer, R. (2014). Bitcoin Transaction Malleability and MtGox. In *Proceedings of the European Symposium on Research in Computer Security* (pp. 313–326). Wroclaw, PL.
- Deephouse, D. L. (1996). Does isomorphism legitimate? *Academy of Management Journal*, 39(4), 1024–1039.
- Deephouse, D. L., & Suchman, M. (2008). Legitimacy in organizational institutionalism. In R. Greenwood, C. Oliver, R. Suddaby, & K. Sahlin (Eds.), *The SAGE Handbook of Organizational Institutionalism* (pp. 49–77). Thousand Oaks, CA: SAGE Publications.
- Delone, W., & McLean, E. R. (2003). The DeLone and McLean model of information systems success: a ten-year update. *Journal of Management Information Systems*, 19(4), 9–30.
- Dewan, S., & Ramaprasad, J. (2014). Social media, traditional media, and music sales. *MIS Quarterly*, 38(1), 101–121.
- Dinev, T., Hart, P., & Mullen, M. R. (2008). Internet privacy concerns and beliefs about

- government surveillance - An empirical investigation. *The Journal of Strategic Information Systems*, 17(3), 214–233.
- Dong, W., Liao, S., & Zhang, Z. (2018). Leveraging Financial Social Media Data for Corporate Fraud Detection. *Journal of Management Information Systems*, 35(2), 461–487.
- Dos Santos, B. L., Patel, P. C., & D'Souza, R. R. (2011). Venture Capital Funding for Information Technology Businesses. *Journal of the Association for Information Systems*, 12(1), 57–87.
- Dowling, J., & Pfeffer, J. (1975). Organizational legitimacy: Social values and organizational behavior. *The Pacific Sociological Review*, 18(1), 122–136.
- Dunlap, J. C., & Lowenthal, P. R. (2009). Tweeting the night away: Using Twitter to enhance social presence. *Journal of Information Systems Education*, 20(2), 129–135.
- Egelund-Müller, B., Elsmann, M., Henglein, F., & Ross, O. (2017). Automated Execution of Financial Contracts on Blockchains. *Business & Information Systems Engineering*, 59(6), 457–467.
- Elsden, C., Manohar, A., Briggs, J., Harding, M., Speed, C., & Vines, J. (2018). Making sense of blockchain applications: A typology for HCI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Montreal, CA.
- European Commission. (2018). Data protection in the EU. Retrieved February 20, 2018, from https://ec.europa.eu/info/law/law-topic/data-protection_en
- Exergy. (2017a). Electric Power Technical Whitepaper. Retrieved March 28, 2018, from <http://exergy.energy/wp-content/uploads/2017/11/Exergy-WhitePaper-v5.pdf>
- Exergy. (2017b). Exergy Business Whitepaper. Retrieved March 28, 2018, from <https://exergy.energy/wp-content/uploads/2017/12/Exergy-BIZWhitepaper-v5.pdf>
- Eyal, I., & Sirer, E. G. (2018). Majority is not Enough: Bitcoin Mining is Vulnerable. *Communications of the ACM*, 61(7), 95–102.
- Fabian, B., Ermakova, T., & Sander, U. (2016). Anonymity in Bitcoin? The users' perspective. In *Proceedings of the 37th International Conference on Information Systems (ICIS)*. Dublin, Ireland.
- Fernandes, E., Rahmati, A., Eykholt, K., & Prakash, A. (2017). Internet of things security research: A rehash of old ideas or new intellectual challenges? *IEEE Security & Privacy*, 15(4), 79–84.
- FINMA. (2018). ICO Guidelines. Retrieved from

- <https://www.finma.ch/en/news/2018/02/20180216-mm-ico-wegleitung>
- Fleisch, E., Weinberger, M., & Wortmann, F. (2015). Geschäftsmodelle im Internet der Dinge. *Schmalenbachs Zeitschrift Für Betriebswirtschaftliche Forschung*, 67(4), 444–465.
- Fridgen, G., Regner, F., Schweizer, A., & Urbach, N. (2018). Don ' t Slip on the Initial Coin Offering (ICO) - A Taxonomy for a Blockchain-enabled Form of Crowdfunding. In *Proceedings of the 26th European Conference on Information Systems (ECIS)*. Portsmouth, UK.
- Gartner. (2018). Gartner Identifies Top 10 Strategic IoT Technologies and Trends. Retrieved June 17, 2019, from <https://www.gartner.com/en/newsroom/press-releases/2018-11-07-gartner-identifies-top-10-strategic-iot-technologies-and-trends>
- Ge, R., Feng, J., Gu, B., & Zhang, P. (2017). Predicting and deterring default with social media information in peer-to-peer lending. *Journal of Management Information Systems*, 34(2), 401–424.
- Genesove, D. (1993). Adverse selection in the wholesale used car market. *Journal of Political Economy*, 101(4), 644–665.
- Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. Vienna, Austria.
- Ghose, A. (2009). Internet exchanges for used goods: An empirical analysis of trade patterns and adverse selection. *MIS Quarterly*, 33(2), 263–291.
- Glarner, A., & Lindgren, A. (2018). What have we learnt from the case of EtherDelta? Retrieved July 18, 2019, from https://www.mme.ch/fileadmin/files/documents/MME_Compact/2019/190307_What_hav_e_we_learnt_from_the_case_of_EtherDelta.pdf
- Glaser, F. (2017). Pervasive Decentralisation of Digital Infrastructures: A Framework for Blockchain enabled System and Use Case Analysis. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*. Waikoloa, HI.
- Goes, P. B. (2014). Editor's comments: big data and IS research. *MIS Quarterly*, 38(3), iii-viii.
- Goh, K.-Y., Heng, C.-S., & Lin, Z. (2013). Social media brand community and consumer behavior: Quantifying the relative impact of user-and marketer-generated content. *Information Systems Research*, 24(1), 88–107.

-
- Goldfeder, S., Kalodner, H., Reisman, D., & Narayanan, A. (2018). When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies. *Proceedings on Privacy Enhancing Technologies*, 2018(4), 179–199.
- Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the Fintech Revolution: Interpreting the Forces of Innovation, Disruption, and Transformation in Financial Services. *Journal of Management Information Systems*, 35(1), 220–265.
- Greenwood, B. N., & Gopal, A. (2016). Ending the Mending Wall: Herding, Media Coverage, and Co-Location in IT Entrepreneurship. *MIS Quarterly*, 41(3), 989–1007.
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611–642.
- Gregor, S., & Hevner, A. (2013). Positioning and Presenting Design Science Research for Maximum Impact. *MIS Quarterly*, 37(2), 337–355.
- Gregor, S., & Jones, D. (2007). The Anatomy of a Design Theory. *Journal of the Association for Information Systems*, 8(5), 312–335.
- Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660.
- Guske, N., & Bendig, D. (2018). Cutting Out the Noise-Costly vs. Costless Signals in Initial Coin Offerings. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. San Francisco.
- Halaburda, H., & Sarvary, M. (2016). *Beyond Bitcoin: The Economics of Digital Currencies*. London, UK: Palgrave Macmillan.
- Hampel, C. E., & Tracey, P. (2017). How organizations move from stigma to legitimacy: The case of cook's travel agency in Victorian Britain. *Academy of Management Journal*, 60(6), 2175–2207.
- Hansen, S. S., Lee, J. K., & Lee, S.-Y. (2014). Consumer-generated ads on YouTube: Impacts of source credibility and need for cognition on attitudes, interactive behaviors, and eWOM. *Journal of Electronic Commerce Research*, 15(3), 254–266.
- Hayes, A. F. (2017). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. New York City, NY: Guilford Publications.
- Heikka, J., Baskerville, R., & Siponen, M. (2006). A design theory for secure information systems design methods. *Journal of the Association for Information Systems*, 7(11), 31.

- Hevner, A. (2007). A Three Cycle View of Design Science Research. *Scandinavian Journal of Information Systems*, 19(192), 87–92. Retrieved from <http://aisel.aisnet.org/sjis>
- Hevner, A., & Chatterjee, S. (2010). *Design research in information systems: theory and practice*. Berlin, Germany: Springer Science+Business Media.
- Hevner, A., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- Hoffman, D. L., & Fodor, M. (2010). Can you measure the ROI of your social media marketing? *MIT Sloan Management Review*, 52(1), 41–49.
- Hogan, C. E., & Wilkins, M. S. (2008). Evidence on the audit risk model: Do auditors increase audit fees in the presence of internal control deficiencies? *Contemporary Accounting Research*, 25(1), 219–242.
- Hollander, N. (2017). Dharma Protocol in a Nutshell. Retrieved July 18, 2019, from <https://blog.dharma.io/dharma-protocol-in-a-nutshell-a7abcc716429>
- Hong, Y., Hu, Y., & Burtch, G. (2018). Embeddedness, pro-sociality, and social influence: Evidence from online crowdfunding. *MIS Quarterly*, 42(4), 1211–1224.
- Hyvärinen, H., Risius, M., & Friis, G. (2017). A Blockchain-Based Approach Towards Overcoming Financial Fraud in Public Sector Services. *Business & Information Systems Engineering*, 59(6), 441–456.
- Iansiti, M., & Lakhani, K. R. (2014). Digital Ubiquity: How Connections, Sensors, and Data Are Revolutionizing Business. *Harvard Business Review*, 92(11), 90–99.
- Iansiti, M., & Lakhani, K. R. (2017). The truth about blockchain. *Harvard Business Review*, 95(1), 118–127.
- Ibbotson, R. G., & Ritter, J. R. (1995). Initial public offerings. *Handbooks in Operations Research and Management Science*, 9, 993–1016.
- Imbault, F., Swiatek, M., De Beaufort, R., & Plana, R. (2017). The green blockchain: Managing decentralized energy production and consumption. In *Proceedings of the 2017 IEEE International Conference on Environment and Electrical Engineering and 2017 IEEE Industrial and Commercial Power Systems Europe (EEEIC/I&CPS Europe)*,. Milan, Italy.
- Jarvenpaa, S., & Teigland, R. (2017). Trust in Digital Environments : From the Sharing Economy to Decentralized Autonomous Organizations. *Proceedings of the 50th Hawaii International Conference on System Sciences*, 5812–5816.

- Jeschke, S., Brecher, C., Meisen, T., Özdemir, D., & Eschert, T. (2017). Industrial Internet of Things and Cyber Manufacturing Systems. In S. Jeschke, C. Brecher, H. Song, & D. B. Rawat (Eds.), *Industrial Internet of Things* (pp. 3–19). Berlin, DE: Springer.
- Juliano, A. (2017). dYdX: A Standard for Decentralized Margin Trading and Derivatives. Retrieved July 18, 2019, from <https://whitepaper.dydx.exchange/>
- Kaplan, A. M., & Haenlein, M. (2010). Users of the world, unite! The challenges and opportunities of Social Media. *Business Horizons*, 53(1), 59–68.
- Karame, G., Androulaki, E., & Capkun, S. (2012). Two Bitcoins at the Price of One? Double-Spending Attacks on Fast Payments in Bitcoin. *Cryptology EPrint Archive*, 2012(248).
- Kim, K., & Mithas, S. (2011). How does bond market view IT investments of firms? An empirical evidence of bond ratings and yield spreads. In *Proceedings of the 32th International Conference on Information Systems (ICIS)*. Shanghai, China.
- Kim, K., Mithas, S., & Kimbrough, M. (2017). Information Technology Investments and Firm Risk Across Industries: Evidence from the Bond Market. *MIS Quarterly*, 41(4), 1347–1367.
- Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and other botnets. *Computer*, 50(7), 80–84.
- Kolias, C., Stavrou, A., Voas, J., Bojanova, I., & Kuhn, R. (2016). Learning internet-of-things security “hands-on.” *IEEE Security & Privacy*, 14(1), 37–46.
- Kostova, T., & Zaheer, S. (1999). Organizational legitimacy under conditions of complexity: The case of the multinational enterprise. *Academy of Management Review*, 24(1), 64–81.
- Kouzinopoulos, C. S., Spathoulas, G., Giannoutakis, K. M., Votis, K., Pandey, P., Tzovaras, D., ... Nijdam, N. A. (2018). Using blockchains to strengthen the security of internet of things. In *Proceedings of the 1st International ISCIS Security Workshop* (pp. 90–100). London, UK.
- Kraken. (2019). Cryptocurrency deposit processing times. Retrieved July 9, 2019, from <https://support.kraken.com/hc/en-us/articles/203325283-Cryptocurrency-deposit-processing-times>
- Kuechler, W., & Vaishnavi, V. (2008). On theory development in design science research: anatomy of a research project. *European Journal of Information Systems*, 17(5), 489–504.
- Kuechler, W., & Vaishnavi, V. (2012). A framework for theory development in design science research: multiple perspectives. *Journal of the Association for Information Systems*, 13(6), 395–423.

- Kumar, A., Fischer, C., Tople, S., & Saxena, P. (2017). A traceability analysis of Monero's blockchain. In *Proceedings of the 22nd European Symposium on Research in Computer Security (ESORICS)* (pp. 153–173). Oslo, Norway.
- Lamin, A., & Zaheer, S. (2012). Wall Street vs. Main Street: Firm strategies for defending legitimacy and their impact on different stakeholders. *Organization Science*, 23(1), 47–66.
- Lamport, L. (1998). The part-time parliament. *ACM Transactions on Computer Systems*, 16(2), 133–169.
- Lamport, L., Shostak, R., & Pease, M. (1982). The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems*, 4(3), 382–401.
- Lee, J. K., Cho, D., & Lim, G. G. (2018). Design and Validation of the Bright Internet. *Journal of the Association for Information Systems*, 19(2), 63–85.
- Lee Kuo Chuen, D. (2015). *Handbook of Digital Currency: Bitcoin, Innovation, Financial Instruments, and Big Data*. Cambridge, MA: Academic Press.
- Leshner, R., & Hayes, G. (2019). Compound: The Money Market Protocol. Retrieved July 18, 2019, from <https://compound.finance/documents/Compound.Whitepaper.pdf>
- Li, T., van Dalen, J., & van Rees, P. J. (2018). More than just noise? Examining the information content of stock microblogs on financial markets. *Journal of Information Technology*, 33(1), 50–69.
- Liang, X., Zhao, J., Shetty, S., & Li, D. (2017). Towards data assurance and resilience in IoT using blockchain. In *Proceedings of the 2017 IEEE Military Communications Conference (MILCOM)* (pp. 261–266). Baltimore, MD.
- Lindman, J., Rossi, M., & Tuunainen, V. K. (2017). Opportunities and risks of Blockchain Technologies in payments – a research agenda. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)* (pp. 1533–1542). Waikoloa, Hi.
- Loebbecke, C., & Picot, A. (2015). Reflections on societal and business model transformation arising from digitization and big data analytics: A research agenda. *The Journal of Strategic Information Systems*, 24(3), 149–157.
- Long, J. S., & Ervin, L. H. (2000). Using heteroscedasticity consistent standard errors in the linear regression model. *The American Statistician*, 54(3), 217–224.
- Loughran, T., & Ritter, J. (2004). Why has IPO underpricing changed over time? *Financial Management*, 33(3), 5–37.

- Lowry, P. B., Dinev, T., & Willison, R. (2017). Why security and privacy research lies at the centre of the information systems (IS) artefact: proposing a bold research agenda. *European Journal of Information Systems*, 26(6), 546–563.
- Lukyanenko, R., Evermann, J., & Parsons, J. (2015). Guidelines for Establishing Instantiation Validity in IT Artifacts: A Survey of IS Research. In *Proceedings of the International Conference on Design Science Research in Information Systems and Technology (DESRIST)* (pp. 430–438). Dublin, Ireland. <https://doi.org/10.1007/978-3-319-18714-3>
- Lukyanenko, R., Wiersma, Y., Huber, B., Parsons, J., Wachinger, G., & Meldt, R. (2017). Representing crowd knowledge: Guidelines for conceptual modeling of user-generated content. *Journal of the Association for Information Systems*, 18(4), 297.
- Lundmark, L. W., Oh, C., & Verhaal, J. C. (2017). A little Birdie told me: Social media, organizational legitimacy, and underpricing in initial public offerings. *Information Systems Frontiers*, 19(6), 1407–1422.
- Machado, C., & Fröhlich, A. A. M. (2018). IoT Data Integrity Verification for Cyber-Physical Systems Using Blockchain. In *Proceedings of the 2018 IEEE International Symposium on Real-Time Distributed Computing (ISORC)* (pp. 83–90).
- Madeira, A. (2019). The Dao, the Hack, the Soft Fork and the Hard Fork. Retrieved July 17, 2019, from <https://www.cryptocompare.com/coins/guides/the-dao-the-hack-the-soft-fork-and-the-hard-fork/>
- Magargle, R., Johnson, L., Mandloi, P., Davoudabadi, P., Kesarkar, O., Krishnaswamy, S., ... Pitchaikani, A. (2017). A simulation-based digital twin for model-driven health monitoring and predictive maintenance of an automotive braking system. In *Proceedings of the 12th International Modelica Conference* (pp. 35–46). Prague, CZ.
- Mai, F., Bai, Q., Shan, J., Wang, X. S., & Chiang, R. H. L. (2015). The impacts of social media on Bitcoin performance. In *Proceedings of the 36th International Conference on Information Systems*. Fort Worth, TX.
- Mai, F., Shan, Z., Bai, Q., Wang, X., & Chiang, R. H. L. (2018). How does social media impact Bitcoin value? A test of the silent majority hypothesis. *Journal of Management Information Systems*, 35(1), 19–52.
- MakerDAO. (2020). The Maker Protocol: MakerDAO's Multi-Collateral Dai (MCD) System. Retrieved February 28, 2020, from <https://makerdao.com/en/whitepaper>
- Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns

- (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4), 336–355.
- Manski, S. (2017). Building the blockchain world: Technological commonwealth or just more of the same? *Strategic Change*, 26(5), 511–522.
- March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems*, 15(4), 251–266.
- Margulies, J. (2015). Garage door openers: An internet of things case study. *IEEE Security & Privacy*, 13(4), 80–83.
- Margulis, S. T. (2011). Three theories of privacy: An overview. In S. Trepte & L. Reinecke (Eds.), *Privacy online: Perspectives on privacy and self-disclosure in the social web* (pp. 9–17). Berlin, Germany: Springer Science+Business Media.
- Mattern, F., & Floerkemeier, C. (2010). From the Internet of Computers to the Internet of Things. In *From active data management to event-based systems and more* (pp. 242–259). Berlin, DE: Springer.
- Mattila, J. (2016). *The Blockchain Phenomenon: The Disruptive Potential of Distributed Consensus Architectures*. Berkeley Roundtable on the International Economy (BRIE) Working Paper Series.
- McAfee, A., Brynjolfsson, E., Davenport, T. H., Patil, D. J., & Barton, D. (2012). Big Data: The Management Revolution. *Harvard Business Review*, 90(10), 60–68.
- Medaglia, C. M., & Serbanati, A. (2010). An overview of privacy and security issues in the internet of things. In D. Giusto, A. Iera, G. Morabito, & L. Atzori (Eds.), *The Internet of Things* (pp. 389–395). New York: Springer.
- Meeuw, A., Schopfer, S., Ryder, B., & Wortmann, F. (2018). LokalPower: Enabling Local Energy Markets with User-Driven Engagement. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. Montreal, CA.
- Mengelkamp, E., Gärtner, J., Rock, K., Kessler, S., Orsini, L., & Weinhardt, C. (2018). Designing microgrid energy markets: A case study: The Brooklyn Microgrid. *Applied Energy*, 210, 870–880.
- Merkle, R. C. (1989). A Certified Digital Signature. In *Proceedings of Crypto 89: Conference on the Theory and Application of Cryptology* (pp. 218–238). Santa Barbara, CA.
- Meth, H., Mueller, B., & Maedche, A. (2015). Designing a requirement mining system. *Journal*

- of the Association for Information Systems, 16(9), 799–837.*
- Miles, M. B., & Huberman, A. M. (1994). *Qualitative data analysis: An expanded sourcebook*. Thousand Oaks, CA: SAGE Publications.
- Miller, A. R., & Tucker, C. (2013). Active social media management: the case of health care. *Information Systems Research, 24(1), 52–70.*
- Miorandi, D., Sicari, S., De Pellegrini, F., & Chlamtac, I. (2012). Internet of things: Vision, applications and research challenges. *Ad Hoc Networks, 10(7), 1497–1516.*
- Modum. (2017). Data integrity for Supply Chain Operations, Powered by Blockchain Technology - Whitepaper. Retrieved March 28, 2018, from <https://assets.modum.io/wp-content/uploads/2017/08/modum-whitepaper-v.-1.0.pdf>
- Modum. (2018). Data integrity for supply chain operations powered by blockchain. Retrieved March 22, 2018, from <https://modum.io/>
- Moody, G. D., Siponen, M., & Pahlila, S. (2018). Toward a Unified Model of Information Security Policy Compliance. *MIS Quarterly, 42(1), 285–311.*
- Moura, J., & Serrão, C. (2016). Security and privacy issues of big data. In N. Zaman, M. E. Seliaman, M. F. Hassan, & F. P. G. Marquez (Eds.), *Handbook of Research on Trends and Future Directions in Big Data and Web Intelligence* (pp. 20–51). Hershey, PA: IGI Global.
- Moyano, J. P., & Ross, O. (2017). KYC optimization using distributed ledger technology. *Business & Information Systems Engineering, 59(6), 411–423.*
- Münsing, E., Mather, J., & Moura, S. (2017). Blockchains for decentralized optimization of energy resources in microgrid networks. In *Proceedings of the 2017 IEEE Conference on Control Technology and Applications (CCTA)* (pp. 2164–2171). Mauna Lani, HI.
- Musk, E. (2019). Twitter Reply to Marcel Feldkamp. Retrieved April 16, 2020, from <https://twitter.com/elonmusk/status/1113990464984813568?s=20>
- Myers, M. D., & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization, 17(1), 2–26.*
- Nærland, K., Müller-Bloch, C., Beck, R., & Palmund, S. (2017). Blockchain to Rule the Waves - Nascent Design Principles for Reducing Risk and Uncertainty in Decentralized Environments. In *Proceedings of the 38th International Conference on Information Systems (ICIS)*. Seoul, South Korea.
- Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System (White Paper). Retrieved

- September 29, 2017, from <https://bitcoin.org/bitcoin.pdf>
- Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital innovation management: Reinventing innovation management research in a digital world. *MIS Quarterly*, *41*(1), 223–238.
- Negroponte, N. (1995). *Being Digital*. New York City, NY: Alfred A. Knorpf.
- NEO. (2016). NEO: A distributed network for the Smart Economy. Retrieved July 21, 2019, from <https://docs.neo.org/docs/en-us/basic/whitepaper.html>
- Newell, S., & Marabelli, M. (2015). Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of datification. *The Journal of Strategic Information Systems*, *24*(1), 3–14.
- Niemimaa, E., & Niemimaa, M. (2017). Information systems security policy implementation in practice: from best practices to situated practices. *European Journal of Information Systems*, *26*(1), 1–20.
- Nishant, R., Teo, T. S. H., & Goh, M. (2017). Do Shareholders Value Green Information Technology Announcements? *Journal of the Association for Information Systems*, *18*(8), 542.
- Notheisen, B., Cholewa, J. B., & Shanmugam, A. P. (2017). Trading Real-World Assets on Blockchain. *Business & Information Systems Engineering*, *59*(6), 425–440.
- Noyen, K., Volland, D., Wörner, D., & Fleisch, E. (2014). When Money Learns to Fly: Towards Sensing as a Service Applications Using Bitcoin. Retrieved from <http://arxiv.org/abs/1409.5841>
- Nunamaker Jr, J. F., Chen, M., & Purdin, T. D. M. (1990). Systems development in information systems research. *Journal of Management Information Systems*, *7*(3), 89–106.
- O’Leary, D. E. (2013). “Big Data”, the “Internet of Things” and the “Internet of Signs.” *Intelligent Systems in Accounting, Finance and Management*, *20*, 53–65.
- Oberländer, A. M., Röglinger, M., Rosemann, M., & Kees, A. (2018). Conceptualizing business-to-thing interactions--A sociomaterial perspective on the Internet of Things. *European Journal of Information Systems*, *27*(4), 486–502.
- Oetzel, M. C., & Spiekermann, S. (2014). A systematic methodology for privacy impact assessments: a design science approach. *European Journal of Information Systems*, *23*(2), 126–150.

- Oh, C., & Sheng, O. (2011). Investigating Predictive Power of Stock Micro Blog Sentiment in Forecasting Future Stock Price Directional Movement. In *Proceedings of the 32nd International Conference on Information Systems (ICIS)*. Shanghai, China.
- Oh, O., Agrawal, M., & Rao, H. R. (2013). Community intelligence and social media services: A rumor theoretic analysis of tweets during social crises. *MIS Quarterly*, 37(2), 407–426.
- Oks, S. J., Fritzsche, A., & Möslin, K. M. (2017). An Application Map for Industrial Cyber-Physical Systems. In S. Jeschke, C. Brecher, H. Song, & D. B. Rawat (Eds.), *Industrial Internet of Things* (pp. 21–46). Berlin, DE: Springer Science+Business Media.
- Oliveira, L., Zavolokina, L., Bauer, I., & Schwabe, G. (2018). To Token or not to Token: Tools for Understanding Blockchain Tokens. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. San Francisco, CA.
- Özyilmaz, K. R., Dougan, M., & Yurdakul, A. (2018). IDMoB: IoT Data Marketplace on Blockchain. In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)* (pp. 11–19). Zug, CH.
- Panarello, A., Tapas, N., Merlino, G., Longo, F., & Puliafito, A. (2018). Blockchain and IoT integration: A systematic survey. *Sensors*, 18(8), 2575.
- Park, J.-W., & Yang, S.-B. (2018). An Empirical Study on Factors Affecting Blockchain Start-ups' Fundraising via Initial Coin Offerings. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. San Francisco, CA.
- Pavlou, P. A. (2011). State of the information privacy literature: Where are we now and where should we go? *MIS Quarterly*, 35(4), 977–988.
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, 24(3), 45–77.
- Peters, G. W., & Panayi, E. (2016). Understanding Modern Banking Ledgers through Blockchain Technologies: Future of Transaction Processing and Smart Contracts on the Internet of Money. In P. Tasca, T. Aste, L. Pelizzon, & N. Perony (Eds.), *Banking Beyond Banks and Money* (pp. 239–278). Berlin, DE: Springer Science+Business Media.
- Pilkington, M. (2016). 11 Blockchain technology: principles and applications. In F. X. Olleros & M. Zhegu (Eds.), *Research Handbook on Digital Transformations* (pp. 225–253). Cheltenham, UK: Edward Elgar Publishing.
- Pollock, T. G., & Rindova, V. P. (2003). Media legitimation effects in the market for initial public

- offerings. *Academy of Management Journal*, 46(5), 631–642.
- Porter, M. E., & Heppelmann, J. E. (2014). How smart, connected products are transforming competition. *Harvard Business Review*, 92(11), 64–88.
- Porter, M. E., & Heppelmann, J. E. (2015). How smart, connected products are transforming companies. *Harvard Business Review*, 93(10), 96–114.
- Preacher, K. J., & Hayes, A. F. (2004). SPSS and SAS procedures for estimating indirect effects in simple mediation models. *Behavior Research Methods, Instruments, & Computers*, 36(4), 717–731.
- Pries-Heje, J., Baskerville, R., & Venable, J. R. (2008). Strategies for Design Science Research Evaluation. In *Proceedings of the 16th European Conference on Information Systems (ECIS)* (pp. 255–266). Galway, Ireland.
- Qiu, L., Tang, Q., & Whinston, A. B. (2015). Two formulas for success in social media: Learning and network effects. *Journal of Management Information Systems*, 32(4), 78–108.
- Rai, A. (2017). Editor's comments: diversity of Design Science Research. *MIS Quarterly*, 41(1), iii–xviii.
- Raval, S. (2016). *Decentralized applications: harnessing Bitcoin's blockchain technology*. Sebastopol, CA: O'Reilly Media.
- Ravikant, N. (2013). Bitcoin – The Internet of Money. Retrieved July 21, 2019, from <https://nav.al/bitcoin-the-internet-of-money>
- Redman, J. (2017). Ethereum's Parity Users Lose Millions in a Multi-Sig Hack. Retrieved July 18, 2019, from <https://news.bitcoin.com/ethereums-parity-client-users-lose-millions-multi-sig-hack/>
- Risius, M., & Beck, R. (2015). Effectiveness of corporate social media activities in increasing relational outcomes. *Information & Management*, 52(7), 824–839.
- Risius, M., & Spohrer, K. (2017). A Blockchain Research Framework: What we (don't) know, where we go from here, and how we will get there. *Business & Information Systems Engineering*, 59(6), 385–409.
- Ritter, J. R., & Welch, I. (2002). A review of IPO activity, pricing, and allocations. *The Journal of Finance*, 57(4), 1795–1828.
- Robinson, O. C. (2014). Sampling in interview-based qualitative research: A theoretical and practical guide. *Qualitative Research in Psychology*, 11(1), 25–41.

- Roman, R., Zhou, J., & Lopez, J. (2013). On the features and challenges of security and privacy in distributed internet of things. *Computer Networks*, 57(10), 2266–2279.
- Ronen, E., Shamir, A., Weingarten, A.-O., & O’Flynn, C. (2017). IoT goes nuclear: Creating a ZigBee chain reaction. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP)* (pp. 195–212). San Jose, CA.
- Rossi, M., Mueller-Bloch, C., Thatcher, J. B., & Beck, R. (2019). Blockchain Research in Information Systems: Current Trends and an Inclusive Future Research Agenda. *Journal of the Association for Information Systems*, 20(9), 1390–1405.
- Rush, T. J. (2018). Defeating the Ethereum DDos Attacks. Retrieved July 18, 2019, from <https://medium.com/@tjayrush/defeating-the-ethereum-ddos-attacks-d3d773a9a063>
- Rüth, J., Schmidt, F., Serror, M., Wehrle, K., & Zimmermann, T. (2017). Communication and Networking for the Industrial Internet of Things. In S. Jeschke, C. Brecher, H. Song, & D. B. Rawat (Eds.), *Industrial Internet of Things* (pp. 317–346). Berlin, DE: Springer Science+Business Media.
- Sanderson, H. (2019). Ford to use blockchain in pilot to trace cobalt mined in Congo. Retrieved July 22, 2019, from <https://www.ft.com/content/d5ba0434-1979-11e9-9e64-d150b3105d21>
- Schilberg, D., Hoffmann, M., Schmitz, S., & Meisen, T. (2017). Interoperability in Smart Automation of Cyber Physical Systems. In S. Jeschke, C. Brecher, H. Song, & D. B. Rawat (Eds.), *Industrial Internet of Things* (pp. 261–286). Berlin, DE: Springer Science+Business Media.
- Schlagwein, D., & Hu, M. (2017). How and why organisations use social media: five use types and their relation to absorptive capacity. *Journal of Information Technology*, 32(2), 194–209.
- Schlossnagle, T. (2006). *Scalable internet architectures*. Indianapolis, IN: Sams Publishing.
- Schroeck, M., Shockley, R., Smart, J., Romero-Morales, D., & Tufano, P. (2012). Analytics: The real-world use of big data. Retrieved February 19, 2018, from <https://www-935.ibm.com/services/us/gbs/thoughtleadership/ibv-big-data-at-work.html>
- Scott, S. V., & Orlikowski, W. J. (2014). Entanglements in practice: performing anonymity through social media. *MIS Quarterly*, 38(3), 873–893.
- Shepherd, D. A., & Zacharakis, A. (2003). A new venture’s cognitive legitimacy: An assessment by customers. *Journal of Small Business Management*, 41(2), 148–167.

-
- Sicari, S., Cappiello, C., De Pellegrini, F., Miorandi, D., & Coen-Porisini, A. (2016). A security- and quality-aware system architecture for Internet of Things. *Information Systems Frontiers*, 18(4), 665–677.
- Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76(2015), 146–164.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Singh, J., Pasquier, T., Bacon, J., Ko, H., & Eysers, D. (2016). Twenty security considerations for cloud-supported Internet of Things. *IEEE Internet of Things Journal*, 3(3), 269–284.
- Siponen, M., & Iivari, J. (2006). Six design theories for IS security policies and guidelines. *Journal of the Association for Information Systems*, 7(7), 445–472.
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: an interdisciplinary review. *MIS Quarterly*, 35(4), 989–1016.
- Sousa, P. R., Antunes, L., & Martins, R. (2018). The Present and Future of Privacy-Preserving Computation in Fog Computing. In A. M. Rahmani, P. Liljeberg, J.-S. Preden, & A. Jantsch (Eds.), *Fog Computing in the Internet of Things: Intelligence at the Edge* (pp. 51–69). Berlin, DE: Springer.
- Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374.
- Spence, M. (1976). Informational aspects of market structure: An introduction. *The Quarterly Journal of Economics*, 90(4), 591–597.
- Stone, E. F., Gueutal, H. G., Gardner, D. G., & McClure, S. (1983). A field experiment comparing information-privacy values, beliefs, and attitudes across several types of organizations. *Journal of Applied Psychology*, 68(3), 459–468.
- Strugar, D., Hussain, R., Mazzara, M., Rivera, V., Lee, J. Y., & Mustafin, R. (2018). On M2M micropayments: a case study of electric autonomous vehicles. In *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)* (pp. 1697–1700). Halifax, CA.
- Subramanian, H. (2018). Decentralized blockchain-based electronic marketplaces. *Communications of the ACM*, 61(1), 78–84.
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571–610.

- Sun, J., Yan, J., & Zhang, K. Z. K. (2016). Blockchain-based sharing services: What blockchain technology can contribute to smart cities. *Financial Innovation*, 2(26).
- Sundmaeker, H., Guillemin, P., Friess, P., & Woelfflé, S. (2010). *Vision and challenges for realising the Internet of Things. Cluster of European Research Projects on the Internet of Things*. Brussels, BE: European Commission, Information Society and Media DG.
- Suresh, P., Daniel, J. V., Parthasarathy, V., & Aswathy, R. H. (2014). A state of the art review on the Internet of Things (IoT) history, technology and fields of deployment. In *IEEE 2014 International Conference on Science Engineering and Management Research (ICSEMR)*. Chennai, IN.
- Swan, M. (2015). *Blockchain: Blueprint for a new economy*. Sebastopol, CA: O'Reilly Media.
- Szabo, N. (1997). Formalizing and securing relationships on public networks. *First Monday*, 2(9).
- Takeda, H., Veerkamp, P., & Yoshikawa, H. (1990). Modeling design process. *AI Magazine*, 11(4), 37–48.
- Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018). Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology*, 94(9–12), 3563–3576.
- Tapscott, A., & Tapscott, D. (2017). How blockchain is changing finance. *Harvard Business Review*, 1(9).
- Tapscott, D., & Tapscott, A. (2016a). *Blockchain Revolution: How the Technology Behind Bitcoin is Changing Money, Business, and the World*. London, UK: Penguin.
- Tapscott, D., & Tapscott, A. (2016b). The impact of the blockchain goes beyond financial services. *Harvard Business Review*, May 10.
- TheMailArchive. (2018). Bitcoin P2P e-cash paper. Retrieved June 27, 2019, from <https://www.mail-archive.com/cryptography@metzdowd.com/msg09959.html>
- Tian, F. (2016). An agri-food supply chain traceability system for China based on RFID & blockchain technology. In *Proceedings of the 13th International Conference on Service Systems and Service Management (ICSSSM)*. Kunming, China.
- Trappe, W., Howard, R., & Moore, R. S. (2015). Low-energy security: Limits and opportunities in the internet of things. *IEEE Security & Privacy*, 13(1), 14–21.
- TrustedIoTAlliance. (2019). Trusted IoT Alliance: Our Mission. Retrieved July 25, 2019, from <https://www.trusted-iot.org/about>

- Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies. *IEEE Communications Surveys & Tutorials*, 18(3), 2084–2123.
- Tumarkin, R., & Whitelaw, R. F. (2001). News or noise? Internet postings and stock prices. *Financial Analysts Journal*, 57(3), 41–51.
- TÜV Rheinland. (2015). Das Problem Tachomanipulation. Retrieved September 29, 2017, from https://www.arvato.com/content/dam/arvato/%0Adocuments/financial-solutions/PK_%0ATachomanipulation_TÜV_Rheinland.pdf
- Uckelmann, D., Harrison, M., & Michahelles, F. (2011). An Architectural Approach Towards the Future Internet of Things. In D. Uckelmann, M. Harrison, & F. Michahelles (Eds.), *Architecting the Internet of Things* (pp. 1–24). Berlin, DE: Springer Science+Business Media.
- Underwood, S. (2016). Blockchain beyond bitcoin. *Communications of the ACM*, 59(11), 15–17.
- Uniswap. (2019). Uniswap Docs. Retrieved February 28, 2020, from <https://github.com/Uniswap/docs>
- Vaishnavi, V. K., & Kuechler, W. (2015). *Design science research methods and patterns: innovating information and communication technology*. Boca Raton, FL: CRC Press.
- Vatanparvar, K., & Faruque, M. A. Al. (2018). Control-as-a-Service in Cyber-Physical Energy Systems over Fog Computing. In A. M. Rahmani, P. Liljeberg, J.-S. Preden, & A. Jantsch (Eds.), *Fog Computing in the Internet of Things: Intelligence at the Edge* (pp. 123–144). Berlin, DE: Springer.
- Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: a framework for evaluation in design science research. *European Journal of Information Systems*, 25(1), 77–89.
- Walls, J. G., Widmeyer, G. R., & El Sawy, O. A. (1992). Building an Information System Design Theory for Vigilant EIS. *Information Systems Research*, 3(1), 36–59.
- Wang, F.-Y. (2018). Blockchain Intelligence: Cornerstone of the future smart economy and smart societies. In *Proceedings of the 2nd World Intelligence Congress*. Tianjin, CN.
- Weber, R. H. (2010). Internet of Things - New security and privacy challenges. *Computer Law & Security Review*, 26(1), 23–30.
- Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, xiii--xxiii.

-
- Weiser, M. (1991). The Computer for the 21 st Century. *Scientific American*, 265(3), 94–105.
- Wengraf, T. (2001). *Qualitative research interviewing: Biographic narrative and semi-structured methods*. Thousand Oaks, CA: SAGE Publications.
- Westin, A. F. (1967). *Privacy and freedom*. New York City, NY: Atheneum.
- Whittington, R., & Pany, K. (2015). *Principles of Auditing & Other Assurance Services*. New York City, NY: McGraw-Hill Education.
- Williams, L. G., & Smith, C. U. (2004). Web Application Scalability: A Model-Based Approach. In *Proceedings of the International Computer Measurement Group Conference (CMG)* (pp. 215–226). Las Vegas, USA.
- Wilson, F. (2018). Investing In Token Focused Funds. Retrieved November 22, 2018, from <https://www.usv.com/blog/investing-in-token-focused-funds>
- Winter, R. (2008). Design science research in Europe. *European Journal of Information Systems*, 17(5), 470–475.
- Wörner, D., & von Bomhard, T. (2014). When Your Sensor Earns Money : Exchanging Data for Cash with Bitcoin. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 295–298.
- Wörner, D., Von Bomhard, T., Schreier, Y.-P., & Bilgeri, D. (2016). The Bitcoin Ecosystem: Disruption beyond financial Services? In *Twenty-Fourth European Conference on Information Systems (ECIS)*. Istanbul, Turkey.
- Wortmann, F., Bilgeri, D., Weinberger, M., & Fleisch, E. (2017). Ertragsmodelle im Internet der Dinge. *Schmalenbachs Zeitschrift Für Betriebswirtschaftliche Forschung, Special Issue*, 71(17), 1–28.
- Wortmann, F., & Flüchter, K. (2015). Internet of things. *Business & Information Systems Engineering*, 57(3), 221–224.
- Xie, P., Chen, H., & Hu, Y. J. (2017). Network structure and predictive power of social media for the Bitcoin market.
- Xu, H., Dinev, T., Smith, J., & Hart, P. (2011). Information privacy concerns: Linking individual perceptions with institutional privacy assurances. *Journal of the Association for Information Systems*, 12(12), 798–824.
- Zhang, Y., & Wen, J. (2017). The IoT electric business model: Using blockchain technology for the internet of things. *Peer-to-Peer Networking and Applications*, 10(4), 983–994.

- Zimmerman, M. A., & Zeitz, G. J. (2002). Beyond survival: Achieving new venture growth by building legitimacy. *Academy of Management Review*, 27(3), 414–431.
- Zuckerman, E. W. (1999). The categorical imperative: Securities analysts and the illegitimacy discount. *American Journal of Sociology*, 104(5), 1398–1438.
- Zyskind, G., Nathan, O., & Pentland, A. (2015). Decentralizing privacy: Using blockchain to protect personal data. *Proceedings of the 2015 IEEE Security and Privacy Workshops (SPW)*, 180–184.

