

Promises, Expectations, and Social Cooperation

Journal Article**Author(s):**

Mischkowski, Dorothee; Stone, Rebecca; Stremitzer, Alexander

Publication date:

2019-11

Permanent link:

<https://doi.org/10.3929/ethz-b-000396464>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

The Journal of Law & Economics 62(4), <https://doi.org/10.1086/706075>

Promises, Expectations, and Social Cooperation

Dorothee Mischkowski *University of Cologne*

Rebecca Stone *University of California, Los Angeles*

Alexander Stremitzer *ETH Zurich*

Abstract

Promising serves as an important commitment mechanism by operating on a potential cheater's internal value system. We present experimental evidence on why people keep their promises, identifying three motives. First, people feel duty bound to keep their promises regardless of whether promisees expect them to do so (promising per se effect). Second, they care about not disappointing promisees' expectations regardless of whether those expectations were induced by the promise (expectations per se effect). Third, they are even more motivated to avoid disappointing promisees' expectations when those expectations were induced by a promise (interaction effect). Clear evidence of some of these effects has eluded the prior literature because of limitations inherent to the experimental methods employed. We sidestep those difficulties by using a novel between-subject vignette design. Our results suggest that promising may contribute to the self-reinforcing creation of trust as expectations of performance encourage promise keeping and vice versa.

1. Introduction

Promises are ubiquitous in both private and commercial settings. Casual observation, introspection, and a string of recent empirical studies suggest that people are willing to keep promises even in the absence of third-party enforcement mechanisms like legal sanctions and second-party enforcement mechanisms that arise from repeated interactions (Ellingsen and Johannesson 2004; Charness and Dufwenberg 2006; Vanberg 2008). Promises therefore seem to serve as important mechanisms of commitment, enabling people to solve fundamental prob-

The authors are grateful to Nate Atkinson, Gary Charness, Martin Dufwenberg, Russell Korobkin, Ernesto Reuben, and seminar audiences at Harvard University, the University of Michigan, the University of California, Los Angeles, and the Yale Law School–University College London workshop Philosophy of Contract for helpful comments and suggestions. We thank Henry Kim, Zhipeng Liao, and Ben Nyblade for their advice on statistical methods and Camille Glaus, Rita Hsu, Henry Kim, Reto Seiler, and Danny Wong for excellent research assistance.

[*Journal of Law and Economics*, vol. 62 (November 2019)]

© 2019 by The University of Chicago. All rights reserved. 0022-2186/2019/6204-0023\$10.00

lems of social cooperation.¹ In particular, promises facilitate the processes of exchange over time by diminishing the holdup problem (for classic discussions of the holdup problem, see Williamson 1979, 1985; Grout 1984; Grossman and Hart 1986; Hart and Moore 1988).

A clear understanding of the determinants of promise keeping is important for institutional design, particularly given that a person's intrinsic reasons to keep her promises are likely to interact with the extrinsic incentives that are provided by second- and third-party enforcement mechanisms (Bénabou and Tirole 2003). Yet the empirical literature to date has only started to shed light on exactly why people keep their promises.

Three types of reasons why people keep their promises, which are not necessarily mutually exclusive, have been suggested by the experimental economics literature. First, a promisor may feel bound to keep his promises insofar as he fears disappointing the promisee's expectations of performance even though those expectations were not induced by his promise (Charness and Dufwenberg 2006). This could be because he recognizes that the promisee is likely to be harmed if her expectations are confounded, either because disappointment is intrinsically painful or because she is more likely to invest in reliance on a promise the more she believes that it will be kept. We call this the expectations per se effect. If this is the only reason why people keep their promises absent self-interested reasons to do so, then promises matter to promisors only derivatively through their effect on the promisee's expectations—that is, only insofar as they make the promisee more likely to anticipate performance. If a promisee formed the same expectation of performance absent a promise, the promisor would feel equally bound to perform.²

Second, a promisor may feel duty bound to keep his promises regardless of the promisee's expectations of performance (Vanberg 2008). We call this the promising per se effect.³ Such an effect is consistent with a Kantian conception of promising according to which a promise gives the promisee the moral right to demand performance regardless of whether the promisee would be harmed were the promise not kept (see, for example, Shiffrin 2008).

¹ Promises serve as first-party mechanisms of private governance. The distinguishing feature of a first-party system is that it operates on the potential cheater's internal value system, eliminating opportunistic behavior at the source, as opposed to eliminating it indirectly by creating extrinsic incentives as second-party and third-party enforcement systems do. See Dixit (2009). For third-party systems, see the vast economic literature on formal contracts beginning with Mirrlees (1976) and Holmström (1979). For second-party systems, see the literature on relational contracting (for example, Macaulay 1963; Klein and Leffler 1981; Bull 1987; Kreps 1996; MacLeod and Malcolmson 1989; Levin 2003).

² A party could, for example, come to expect another party's performance after hearing the latter make a mere statement of intent. See Battigalli and Dufwenberg (2007, 2009) for a formal model of what they refer to as "simple guilt," according to which people have a preference to avoid the guilt they experience when they let others down. Their model utilizes psychological game theory by incorporating others' beliefs into the utility function. See Dufwenberg and Gneezy (2000) for experimental evidence of guilt aversion (which they called "let-down aversion" at the time).

³ Further experimental evidence that is consistent with a commitment-based explanation for promise keeping can be found in Braver (1995), Ostrom, Walker, and Gardner (1992), Ellingsen and Johannesson (2004), and Ismayilov and Potters (2012).

Third, promises may enhance the effect of the promisee's expectations. A promisor might be especially concerned about not disappointing expectations that he brought about by his promise, such that the effect of promising is to make the promisor more sensitive to the promisee's expectations than he would have been had he not made a promise. We call this the interaction effect. An interaction effect is a feature of Scanlon's (1998, pp. 295–317) theory of promising according to which duties to perform arise only if the promisor intentionally brings about a promisee's expectations of performance, thus rendering him morally responsible for this expectation and so duty bound not to disappoint it.⁴

An interaction effect is also a feature of Ederer and Stremitzer's (2017) model of conditional guilt aversion. In that model, it takes on a strong form, as expectations matter if and only if they were generated by a promise (contrary to the expectations per se effect).⁵ Here we assume, more weakly, that a promise simply enhances the effect of expectations on the promisor's willingness to perform.

We employed a between-subject vignette study to explore these possible determinants of promise keeping. We asked subjects to imagine that they were a prospective buyer of a product who had told a seller that he would buy the product from her upon his return from a trip out of town. Some subjects were told that they had promised the seller that he would buy the product from her, while others were told that they had told the seller that they merely planned to do so, explicitly stating that they were not making a promise. Subjects were then told that the seller had formed a belief about the likelihood that the buyer would buy the product from her. Some were told that she was certain that he would buy the product from her, others were told that she believed that there was a 50 percent chance, while others were told that she was certain that he would not buy the product from her. Each subject was then asked how likely they thought it was that they would buy the product from the seller despite having learned that a second seller was selling the same product at a lower price. We ran a first study in July 2016 and a replication study on a different pool of subjects in 2017 and 2018.

We used a vignette study in which subjects were asked to imagine that the interaction took place as opposed to an experimental design in which the interaction played out with real monetary stakes, because it is difficult to manipulate subjects' expectations about the behavior of other subjects in a controlled fashion. We were able to sidestep this difficulty by simply telling subjects the seller's beliefs about the likelihood that the buyer would perform. This has enabled us to reproduce and clarify key results of the prior literature while also providing the

⁴ Indeed, Scanlon (1998, pp. 312–14) denies that there is a duty to keep a promise that has not caused the promisee to expect performance. Thus, Scanlon's Profligate Pal, who has consistently breached promises to repay loans he has received from his friends over the years such that none of them have any confidence that he will repay a future loan, cannot generate a promissory duty to repay such a loan according to Scanlon's theory. A purely Scanlonian promisor therefore would not exhibit the promising per se effect.

⁵ Stone and Stremitzer (forthcoming) propose an analogous theory of the effect of reliance on promise keeping.

first clear evidence of the interaction effect and the promising per se effect, all within the confines of a single experiment.⁶

We find evidence of all three hypothesized effects. Consistent with the promising per se effect, subjects were on average more inclined to perform when they promised to do so regardless of their counterparties' expectations. In particular, even if their counterparties were certain that they would not perform, subjects were more likely to perform when they promised they would do so. Thus, our subjects' behavior is not fully consistent with the behavior of someone who is acting in accordance with Scanlon's theory of promise keeping. Consistent with the expectations per se effect, subjects were on average more likely to perform the greater were their counterparties' expectations of performance even if they made no promises. Consistent with the interaction effect, we find that the average sensitivity of subjects' willingness to perform to their counterparties' expectations was higher when they made promises. Thus, our data support a weak version of the model of conditional guilt aversion proposed in Ederer and Stremitzer (2017).

Our findings have implications for understanding extralegal mechanisms of cooperation. They show that promising is a useful commitment mechanism and not simply because it creates expectations of performance in the promisee. Promising creates commitment independently of the promisee's expectations and enhances the commitment effect of those expectations.⁷

This raises the possibility of an interesting dynamic. Insofar as trust is associated with higher second-order expectations, both the expectation per se effect and the interaction effect suggest that even in the absence of promising, more trust leads to greater cooperation, and, conversely, the absence of trust undermines cooperation.⁸ This suggests that promising may be associated with self-reinforcing spirals of trust or distrust. When parties initially trust one another, possibilities for cooperation are enhanced, creating even more trust. Conversely, when parties initially distrust one another, cooperation is less likely, preventing the development of trust between them. The promising per se effect, however, gives parties a way of breaking out of a negative spiral of distrust. If promising is a way of creating commitment even in the absence of trust, promising can build trust where it is initially lacking, leading to positive instead of negative self-reinforcing dynamics.

The remainder of this paper is organized as follows. Section 2 develops a more formal model of the reasons why people may keep their promises. Section 3 de-

⁶ See Section 3.1 for a more thorough discussion of the pros and cons of the vignette method relative to a design with monetary incentives, real communication, and endogenous promising. See Section 5 for a discussion of our contributions to the literature.

⁷ Of course, there is a sense in which it is not surprising that promising could have this multitude of effects. In a world in which promising had no force independent of the expectations that promises create in promisees, promising would not cause rational promisees to expect performance. Our results underline that promising is an important mechanism of commitment, a fact that is taken for granted in the design of much of American contract law. See, for instance, *Restatement (Second) of Contracts*, sec. 1.

⁸ Reuben, Sapienza, and Zingales (2009) show that distrust is self-fulfilling in a setting in which promises are absent.

scribes the design and procedure of our experiment and derives the hypotheses we test. Section 4 reports our results. We discuss the results and their contribution to the existing literature in Section 5. Section 6 concludes.

2. Theory

Suppose that a buyer is thinking about buying a product from a seller. The buyer is unable to make the purchase immediately because he is out of town. But he informs the seller that he will purchase the product from her when he returns. He has a choice when communicating his intention to the seller: he can either promise her that he will purchase the product from her when he returns or simply tell her that this is his plan without making her any promises. After talking to the seller, however, he learns that a second seller is selling the same product at a lower price. And so he must decide whether to buy the product from the first seller at the higher price, as he said he would, or to purchase it from the second seller at the lower price. The first seller has formed an expectation about how likely the buyer is to purchase the product from her, and the buyer learns this expectation. To what extent does the expectation influence the buyer's willingness to buy the good from the first seller? And how is the buyer's willingness to buy the good from the first seller affected by whether or not he made the seller a promise?

Formally, the timeline is as follows. At the first stage, the buyer decides whether to make the first seller a promise, $p \in \{0, 1\}$. He can either tell the seller that he promises to buy the good when he returns, $p = 1$, or that while he plans to buy the good from her, he does not promise that he will do so, $p = 0$. At the second stage, the seller forms an expectation about the likelihood that the buyer will buy the product from her, $e \in [0, 1]$, and the buyer learns this expectation from a neutral third party.⁹ The buyer then learns that a second seller is selling the product at a lower price. Finally, the buyer chooses an action, $a \in \{0, 1\}$, either buying the good from the first seller, $a = 1$, or from the second seller, $a = 0$.

The buyer's behavior, and therefore his prediction about his own behavior, is determined by his preferences. In the experiment, we asked subjects to report their beliefs about the likelihood that they would buy from the first seller if they were in the position of the buyer in this scenario.¹⁰ Thus, it is helpful to envisage a penultimate stage of the timeline during which the buyer forms a prediction about his own future behavior, $b(p, e) \in [0, 1]$, the likelihood that he will end up buying the product from the first seller. To make sense of the idea that the buyer

⁹ We assume that the buyer learns the expectation in this way so that he has no reason to question the truth of the reported expectation. The seller (or any other interested party) would have an incentive to deceive the buyer about the seller's expectation if that expectation is likely to influence the likelihood that the buyer ultimately buys the product from her. Ellingsen et al. (2010) ask the recipient in a dictator game to report her expectations to the experimenter. Her expectations are then communicated to the dictator. This at least raises the possibility of strategic communication by the recipient and/or of the dictator discounting the recipient's reported expectations.

¹⁰ This mimics the question asked in Wilkinson-Ryan and Hoffmann (2015) for which our model provides an explicit justification. See Section 3.1 for a discussion of this design choice.

might be uncertain about his own future behavior, suppose that there are random components of his utility function $U(a, p, e)$, ε_a , that are realized only at the final stage. We assume that ε_0 and ε_1 are independently distributed, with $E[\varepsilon_a] = 0$. These random elements could result from uncertain subjective components of the buyer's utility or uncertain features of the environment that affect the utility of his options. They ensure that the buyer is uncertain about the content of his preferences, and hence his future decision, when he forms a prediction about the decision at the penultimate stage.

We suppose that the buyer's preferences are shaped by his attitudes toward promising and the seller's expectations—or, more precisely, his beliefs about the seller's expectations, but we elide this distinction by supposing that he learns the seller's expectations with certainty. In particular, we suppose that the buyer's preferences may be shaped by his moral beliefs. Thus, the buyer may prefer an action that runs counter to his material self-interest, because his moral beliefs may cause him to sacrifice some of his own welfare for the sake of duty or an altruistic desire to promote the well-being of another.¹¹

What attitudes toward promising and expectations are likely to be exhibited by people in the buyer's situation? We identify four plausible attitude types. People may display a combination of these attitudes.

Self-Interested. Classical economic theory predicts that the buyer cares only about his own material well-being. Such a buyer does not care about keeping his promises or about satisfying the seller's expectations. He will buy from the first seller only when it is in his self-interest to do so.

Compassionate. A buyer might not care about keeping his promises as such but nonetheless display compassion toward the seller by promoting the seller's well-being. Such a buyer will care about satisfying the seller's expectations if she is likely to be harmed in some way when her expectations are thwarted, because, for example, she has relied on the buyer performing or she experiences psychological pain as a result of disappointed expectations.

A buyer may not exhibit compassion toward the seller but nonetheless care about doing his duty such that he is inclined to keep promises that he perceives to be morally binding. There are two possibilities.

Kantian. First, the buyer might believe that promises are morally binding *per se*—that is, irrespective of the seller's expectations. The buyer might believe, for example, that the promise gives the seller a moral right to demand performance (see, for example, Shiffrin 2008).

Scanlonian. Second, in the spirit of Scanlon's theory, the buyer might believe he is obligated to honor expectations in others that he has created through his promises (see Scanlon 1998, pp. 295–327). Unlike the Kantian, he believes that his moral duty depends on the seller forming some expectation that he will perform. Unlike the compassionate buyer, he is indifferent toward the seller's expect-

¹¹ We use preferences to represent how the buyer chooses among the available actions and not necessarily to represent the buyer's welfare. As the sole purpose of our model is to make predictions, we can remain agnostic about the relationship between his preferences and his welfare.

tations if he does not believe himself responsible for them, as when he expressly disclaims such responsibility by telling her that he is not making any promises.

More formally, let $m(a)$ represent the material utility associated with action a , where $m(0) > m(1)$. Thus, a purely self-interested buyer cares only about maximizing m and so will always choose to buy the product from the second buyer. If the buyer exhibits compassion, then he cares about ensuring that his actions meet or exceed the seller's expectations regardless of whether he made a promise to do so. Formally, he cares about minimizing $j(e - a)$, where $j(x) > 0$, $j'(x) > 0$ if $x > 0$, and $j(x) = 0$ otherwise. If the buyer exhibits a Kantian disposition, then he cares about buying from the first seller, $a = 1$, whenever he has promised to do so. Formally, he cares about minimizing $pk(1 - a)$, where $k(x) > 0$ and $k'(x) > 0$ if $x > 0$ and $k(0) = 0$. Finally, if the buyer exhibits a Scanlonian disposition, then he cares about ensuring that his actions meet the seller's expectations only if he made a promise to do so. Formally, he cares about minimizing $pl(e - a)$, where $l(x) > 0$ and $l'(x) > 0$ if $x > 0$, and $l(x) = 0$ otherwise.

People might exhibit combinations of these attitudes. Thus, the utility function of buyer i can be expressed as

$$U_i(a, p, e) = m(a) - \alpha_i j(e - a) - p[\beta_i k(1 - a) + \gamma_i l(e - a)] + \varepsilon_a, \quad (1)$$

where $\alpha_i, \beta_i, \gamma_i \geq 0$ are parameters that describe the weight the buyer places on the aforementioned considerations. If $\alpha_i = \beta_i = \gamma_i = 0$, the buyer is self-interested. The strength of his compassionate instincts is measured by α_i , the strength of his Kantian instincts is measured by β_i , and the strength of his Scanlonian instincts is measured by γ_i .

The buyer will choose to buy from the seller whenever

$$\begin{aligned} \Delta U_i &\equiv U_i(1, p, e) - U_i(0, p, e) \\ &= \Delta m + \alpha_i j(e) + p\beta_i k(1) + p\gamma_i l(e) + \Delta \varepsilon \geq 0, \end{aligned} \quad (2)$$

where $\Delta m \equiv m(1) - m(0) < 0$ and $\Delta \varepsilon \equiv \varepsilon_1 - \varepsilon_0$, and therefore $E[\Delta \varepsilon] = 0$. A larger α_i increases the effect of the seller's expectations on the buyer's willingness to buy from her even if he made no promise. A larger β_i increases the effect of a promise to the seller on the buyer's willingness to buy from her regardless of her expectations. A larger γ_i further increases the effect of the seller's expectations on the buyer's willingness to buy from her if he made a promise to her.

Let $F(\cdot)$ be the cumulative distribution function of $\Delta \varepsilon$ and $f(\cdot)$ be the probability density function. It follows from equation (2) that, during the penultimate stage of the experiment, the buyer's belief about his own future behavior is given by

$$\begin{aligned} b_i(p, e) &= \Pr[\Delta \varepsilon \geq -\Delta m - \alpha_i j(e) - p\beta_i k(1) - p\gamma_i l(e)] \\ &= 1 - F(-\Delta m - \alpha_i j(e) - p\beta_i k(1) - p\gamma_i l(e)). \end{aligned}$$

To ensure that $b_i(p, e) \in [0, 1]$, we further make the technical assumption that $\Delta \varepsilon$ is uniformly distributed on $[-c, c]$, where

$$c > -\Delta m \quad \text{and} \quad -c < -\Delta m - \alpha_i j(1) - p\beta_i k(1) - p\gamma_i l(1). \quad (3)$$

There are three important implications. First, regardless of whether a promise was made, so long as the buyer exhibits some degree of compassion, $\alpha_i > 0$, b_i is increasing in the seller's expectation:

$$\frac{\partial b_i(1, e)}{\partial e} = f(\cdot)[\alpha_i j'(e) + \gamma_i l'(e)] \geq f(\cdot)\alpha_i j'(e) = \frac{\partial b_i(0, e)}{\partial e} > 0.$$

This is the expectations per se effect. Second, so long as the buyer exhibits a Scanlonian predisposition, $\gamma_i > 0$, the rate of increase of b_i in the seller's expectation e is higher when a promise was made than when no promise was made. This is the interaction effect. Third, so long as the buyer exhibits a Kantian predisposition, $\beta_i > 0$, b_i is always higher if a promise was made:

$$b_i(1, e) - b_i(0, e) = F(-\Delta m - \alpha_i j(e)) - F(-\Delta m - \alpha_i j(e) - \beta_i k(1) - \gamma_i l(e)) > 0.$$

This is the promising per se effect.¹²

3. Design, Hypotheses, and Procedure

3.1. Design

Subjects were asked to imagine that they were a prospective buyer of a good in a version of the scenario set out in Section 2. Specifically, they were asked to imagine that a seller, B, had offered to sell them a product that they were interested in buying for \$100 once they returned from a trip out of town but that just before returning from the trip they had learned that another seller, C, was offering to sell an equivalent product for \$85.

Subjects were then asked to report the likelihood that they would buy the product from C (instead of B) under one of six randomly selected conditions. The conditions differed according to whether the buyer promised the seller that he would buy the good from her ("I promise I will buy it from you") or simply informed the seller of his plans without making any promises ("All I can say is that I plan to buy it from you, though I cannot promise that I will do so") and according to the expectations that the seller formed about the likelihood that the buyer would in fact buy the good from her.¹³ That is, there were three promise conditions and three no-promise conditions, each one characterized by a particular expectation: 0 percent, 50 percent, or 100 percent. In the 0 percent conditions, subjects were

¹²The promising per se effect and the expectations per se effect arise regardless of the distribution of $\Delta \varepsilon$ so long as $f(\cdot) > 0$. The interaction effect arises so long as $\alpha_i j'(e)[f(-\Delta m - \alpha_i j(e) - p\beta_i k(1) - p\gamma_i l(e)) - f(-\Delta m - \alpha_i j(e))] + \gamma_i l'(e)[f(-\Delta m - \alpha_i j(e) - p\beta_i k(1) - p\gamma_i l(e)) > 0$. Many distributions will satisfy this condition over much or most of their range. A sufficient condition is that $f(\cdot)$ be non-increasing.

¹³We added the language disclaiming the promise in the no-promise conditions to remove potential ambiguity about whether the buyer had made a promise. We were concerned that without such language, subjects might, given the context, interpret a statement announcing their future plans as a promise despite the absence of the words "I promise."

told that the seller was sure that the buyer would not buy the product from her; in the 50 percent conditions, they were told that the seller thought that there was a 50 percent chance; and in the 100 percent conditions, they were told that the seller was sure that the buyer would buy the product from her.¹⁴

We ran two studies. Study 1 exclusively employed a between-subject design whereby, to minimize possible demand effects, each subject was exposed to only one of the six conditions.¹⁵ In each condition, subjects were asked to rate the likelihood that they would buy the product from the second seller, thereby not performing as they indicated they would to the first seller (for a similar design, see Wilkinson-Ryan and Hoffmann 2015). Specifically, subjects were asked to select one of seven options: no way, very unlikely, unlikely, 50:50, likely, very likely, and with certainty. The advantage of framing the question as a likelihood, as opposed to simply asking subjects whether they would buy from the seller, is that it gave us a more continuous measure of subjects' willingness to perform. Arguably, this framing also made it psychologically easier for subjects to reveal their preferences for not performing than would have been the case had they chosen between definitely performing and definitely not performing.¹⁶

Study 2 implemented the same between-subject design but with an additional within-subject component. After subjects' answers to the first condition were recorded, subjects were exposed to the other five conditions in a randomized order. Subjects were not told in advance that they would see all six conditions. Their answers to the first condition therefore constitute between-subject data, while their answers to all six conditions constitute within-subject data.¹⁷

3.2. Hypotheses

We are now in a position to formulate the hypotheses that flow from our theory. First, the promising per se effect entails that subjects will report that they will

¹⁴ See Appendix A for the text of the vignette and Online Appendix Section OA2 for screenshots of the instructions. There is arguably something incongruous about the 0 percent promise condition and the 100 percent no-promise condition. It might seem odd to suppose that the seller would feel certain that the buyer would not keep his promise given the limited information about him that the seller apparently has. Likewise, it might seem odd to suppose that the seller would feel sure that the buyer would buy the good from her when the buyer expressly told the seller that he was not making any promises. But people vary in their degrees of optimism and pessimism, so these conditions simply capture beliefs formed by people on the extreme ends of the optimism-pessimism spectrum. And presumably when people say they are sure that something will occur, what they mean to convey is that they believe it to be extremely likely, not that the alternative is literally impossible.

¹⁵ Had we instead asked each subject to respond to all six conditions, it would have become apparent to them that we were studying the effects of promising and expectations, and they might have distorted their answers to conform to their beliefs about our hypotheses or, more minimally, to create a false impression of consistency.

¹⁶ This may be particularly important given that subjects faced no monetary consequences for their decisions, so reporting a willingness to perform was a cheap way for subjects to feel good about themselves, which could have tempted them to distort their answers in the direction of performance.

¹⁷ As mentioned above, a within-subject design may create demand effects. However, it also allows us to see how a subject reacts when more attuned to the variables at issue as a result of being exposed to all conditions. Within-subject data also increase the study's power by increasing the number of observations while holding the sample size constant.

be more likely to buy from the initial seller B when a promise was made for all levels of B's expectation.

Hypothesis 1. Subjects' reported likelihood of buying from seller B in the promise conditions will exceed their reported likelihood in the no-promise conditions for all levels of seller B's expectation.

Explanation. The promising per se effect means that promising makes the buyer more willing to buy from seller B even when she has no expectations of performance. We should observe reported likelihoods consistent with hypothesis 1 so long as some subjects have a Kantian disposition because Kantian agents, as we have defined them, are inclined to do as they promised regardless of promisees' expectations while exhibiting no particular disposition to be nice to others if they have not made a promise.

Second, the expectations per se effect means that subjects will report that they are more likely to buy from seller B the higher is her expectation that they will do so, regardless of whether a promise was made.

Hypothesis 2. Subjects' reported likelihood of buying from seller B will be greater the higher are seller B's expectations in both the promise and no-promise conditions.

Explanation. The expectations per se effect means that the buyer becomes more willing to buy from seller B the greater are her expectations even if he made no promise. We should observe reported likelihoods consistent with hypothesis 2 so long as some subjects have a compassionate disposition, as compassionate agents do not like disappointing another person's expectations even if they made no promises.

Finally, the interaction effect entails that subjects will report a greater increase in their willingness to buy from seller B as her expectation increases when a promise was made.

Hypothesis 3. An increase in seller B's expectations will cause a greater increase in subjects' reported likelihood of buying from her in the promise conditions than in the no-promise conditions.

Explanation. The interaction effect means that the buyer becomes more sensitive to seller B's expectations when he made a promise. We should observe reported likelihoods consistent with hypothesis 3 so long as some subjects have a Scanlonian disposition because Scanlonian agents, as we have defined them, care about not disappointing the expectations of others but only when those expectations arose from their promises.

3.3. Procedure

Study 1 was conducted in July 2016 using vignettes programmed using Qualtrics. We recruited 169 subjects from the pool of master workers on Amazon Mechanical Turk (MTurk) and 614 from the general pool of MTurk workers with a human intelligence task approval rate of 95 percent or greater.¹⁸ Roughly 34 percent of subjects were MTurk master workers.¹⁹ We determined our sample size using G*Power analysis and a simulation based on pilot data.²⁰ All recruited subjects completed the survey. After subjects responded to the vignette, they were asked several control questions to ascertain whether they had understood the scenario. We also asked subjects other questions that were designed to assess how carefully and honestly they answered the survey questions. We used these questions to create additional robustness checks for our results. Finally, we asked subjects questions about their demographic characteristics.²¹ Before subjects proceeded to the main part of the experiment, we announced that the task would take around 5–7 minutes and that we would pay subjects \$1 for participating in the study. The announced hourly wage was therefore \$9–\$12, which is well above the current national minimum wage (\$7.25 per hour) and much higher than the wages paid in typical MTurk studies.²² On average, subjects took 4 minutes and 24 seconds, so the effective average hourly wage was \$14.

In study 2 we sought to replicate our results with a different pool of subjects. The study was conducted online using the server of the University of Hagen, and we finished collecting data in February 2018.²³ All 771 participants who com-

¹⁸ We also conducted three pilot studies on 50 subjects who were recruited on Amazon Mechanical Turk (MTurk). In the pilots, we recruited only MTurk master workers. According to Amazon, these are “an elite group of Workers with demonstrated accuracy across specific types of HITs [human intelligence tasks]” (Amazon Mechanical Turk 2012). Initially, we planned on restricting our final study to master workers who had not participated in the pilots. However, when we ran the study with this restriction, after an initial flurry of around 150 responses, the response rate slowed down to about one response per hour, which suggested that we had used most of the pool of master workers. Thus, after receiving 169 responses from master workers, we decided to eliminate the restriction and recruit 614 subjects from the general pool of MTurk workers with a HIT approval rate of 95 percent or greater.

¹⁹ Ninety-seven of the 614 workers from the general pool self-identified as MTurk master workers by answering affirmatively the following question in the postexperiment questionnaire, “Are you an MTurk Master Worker (your response to this question will have no effect on your payout)? (Yes/No/I don’t know what an MTurk Master Worker is).” We classified all subjects who answered yes as MTurk master workers.

²⁰ We conducted an a priori power analysis using G*Power (Faul et al. 2009) that revealed a required sample size of 713 to detect a conservatively estimated small effect ($f^2 = .02$) in a linear multiple regression (three predictors) with a high power of $1 - \beta = .90$ and $\alpha = .05$. We recruited a slightly larger sample of 783 subjects.

²¹ See Appendix B for the questions in the postexperiment survey and the subjects’ responses. Online Appendix Section OA2 contains the screenshots.

²² Horton and Chilton (2010) find a median hourly wage of \$1.38, and Mason and Watts (2010) report a typical payment of \$.01–\$.10 per HIT.

²³ When we decided to conduct the replication, we already had some data that had been collected in April 2017 for a related study ($N = 360$). For the purposes of the replication, we added to these in January and February 2018 ($N = 411$) to create a sample with a size similar to that of study 1, generating our final sample size of 771 subjects.

pleted the survey were members of the subject pool recruited by the Hagen Decision Lab in the Department of Cognitive Psychology at the University of Hagen. We used the same vignettes and questionnaire, programmed using Unipark, except that the vignettes were translated into German, and we exposed subjects to all six conditions in a randomized order. We aimed to recruit 783 subjects as in study 1 but stopped the study after we had recruited 771 subjects, because we seemed to have exhausted the active portion of the subject pool.²⁴ Ninety-nine percent of subjects who correctly entered their Decision Lab identification (ID) information completed the survey (779 subjects entered their ID information correctly and were thus allowed to take part). Subjects were sent an invitation using the email server of the Hagen Decision Lab. The email message did not announce that subjects would be expected to respond to multiple vignettes. However, participants were told at the outset the time we estimated it would take them to complete the survey (approximately 15 minutes) and the amount of money they would earn (€2). The announced hourly wage was therefore €8, which is in line with the average hourly earnings for online experiments in the Hagen Decision Lab.²⁵ On average, subjects took 10 minutes and 57 seconds. Thus, the effective average hourly wage was €11.35.

4. Results

Table 1 and Figure 1 summarize the means of the choice variables by treatment condition in study 1 and (replication) study 2.²⁶ Table 1 also reports the means of the combined sample. All reported results are based on between-subject data. We used the within-subject data for an additional robustness check.²⁷

As a descriptive matter, both sets of data are in line with our hypotheses. First, consistent with hypothesis 1 and the promising per se effect, the mean likelihood of performance is higher in the promise conditions than in the no-promise conditions for all levels of seller B's expectations. Thus, in Figure 1, the lines representing the mean likelihood of performance in the promise conditions are higher than the corresponding lines for the no-promise conditions. Importantly, this is

²⁴ At the time we ran the experiment, the subject pool contained approximately 1,200 members.

²⁵ The average hourly wage in the state of North Rhine–Westphalia was €17.83 before taxes in 2014. Importantly, participants do not have to pay taxes on their study earnings, so €8 is roughly in line with average after-tax hourly wages.

²⁶ Our reported statistical results are based on all observations in each sample. They remain qualitatively unchanged irrespective of whether we include or exclude participants who incorrectly answered the control questions. Indeed, all of our results hold irrespective of whether we include or exclude pilot data; whether, in study 1, we include or exclude MTurk master workers; or whether we include or exclude those who answered postexperiment survey questions in a way that makes us doubt their carefulness or seriousness.

²⁷ Results based on within-subject data may be affected by demand effects and to this extent are less convincing, which is why we use these data only for an additional robustness check. We find that the within-subject data from study 2 are fully consistent with the between-subject data from studies 1 and 2. All statistical tests are significant at significance levels higher than those reported for the between-subject data. This is not surprising given the greater number of observations ($N = 4,626$). Moreover, we can use the Wilcoxon sign-rank test for matched data, which has higher statistical power than the Wilcoxon rank-sum test for unmatched data.

Table 1
Mean Likelihood of Performance across Conditions

	0 Percent		50 Percent		100 Percent	
	Value	N	Value	N	Value	N
Study 1:						
No promise	.13	143	.22	117	.31	132
Promise	.38	119	.51	134	.67	138
Study 2:						
No promise	.18	205	.20	125	.33	43
Promise	.36	236	.47	129	.66	33
Combined:						
No promise	.16	348	.21	242	.32	175
Promise	.36	355	.49	263	.66	171

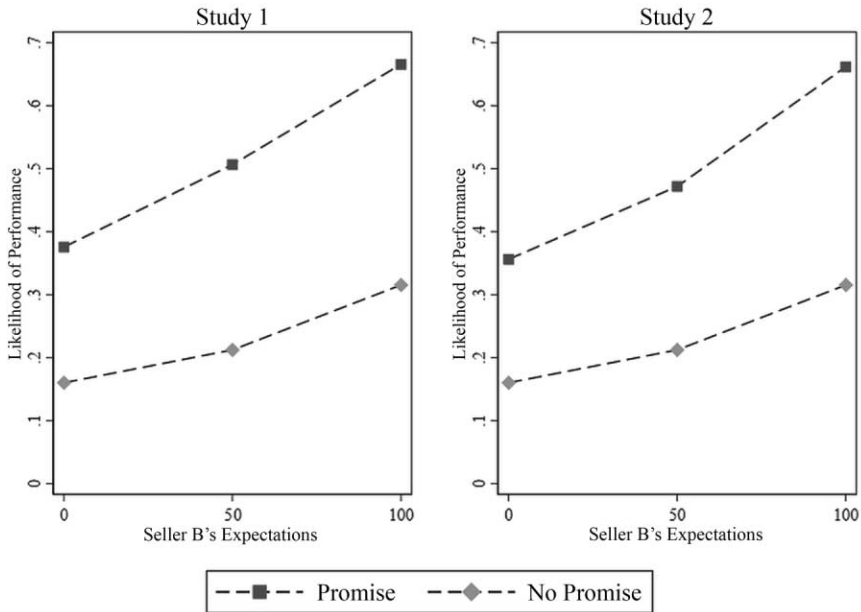


Figure 1. Mean likelihood of performance across conditions in study 1 and study 2

the case when expectations are 0. Thus, we see a pure effect of promising even when the promisee has no expectations of performance. Second, consistent with hypothesis 2 and the expectations per se effect, the mean likelihood of performance is higher the greater are seller B's expectations. This is true even in the no-promise conditions. Thus, expectations have an effect even when no promise was made. Accordingly, in Figure 1 the lines representing the mean likelihood

of performance in both conditions are upward sloping. Finally, consistent with hypothesis 3 and the interaction effect, the mean likelihood of performance increases more if there was a promise than if there was none. Thus, in Figure 1, the lines representing the mean likelihood of performance in the promise conditions are steeper than the corresponding lines for the no-promise conditions.

Figure 2 provides a more complete picture of our results. It shows the full distributions of the choice variable by condition for the combined sample.²⁸ Consistent with the promising per se effect, there is greater probability mass at higher ends of the distribution for the promise conditions than for the no-promise conditions. Consistent with the expectations per se effect, probability mass shifts to the right as expectations increase in both the promise and no-promise conditions. Finally, Figure 2 illustrates the nature of the interaction effect. In the no-promise conditions, more probability mass shifts to the right-hand tail as expectations increase, but the distributions remain skewed to the right. By contrast, in the promise conditions, the skewness of the distribution shifts from right to left.²⁹

We used a number of statistical tests to check our hypotheses. As a baseline, we used a standard ordinary least squares (OLS) regression model to test all three of our hypotheses:

$$b = \alpha \times \text{Promise} + \beta \times \text{Expectation} + \gamma \times \text{Promise} \times \text{Expectation} + \varepsilon. \quad (4)$$

Violations of normality assumptions about the distribution variables in a standard OLS model are not generally seen as a problem when, as here, the sample size is large. Nonetheless, we ran the nonparametric Wilcoxon rank-sum test as a robustness check to test hypothesis 1 (the promising per se effect) and hypothesis 2 (the expectations per se effect).³⁰ Finally, we performed a bootstrapping procedure that enabled us to employ a nonparametric test of hypothesis 3 (the interaction effect). Table 2 reports the results from the standard OLS model.³¹

We find that promising raises the average likelihood of performance by 24 percentage points in study 1 and 18 percentage points in study 2. This effect is statistically significant at the 1 percent level. Promising also increases the average likelihood of performance at each level of expectation including at 0 expectation (Wilcoxon rank-sum test, $p < .01$). This is strong evidence in support of our hypothesized promising per se effect as embodied in hypothesis 1.

Subjects are also more likely to perform if the seller has higher expectations

²⁸ The distributions of study 1 and study 2 taken separately reveal the same pattern.

²⁹ The differential skewness of the promise and no-promise conditions also means that outliers are likely to dampen the statistical significance of the interaction effect in statistical tests based on sample means.

³⁰ For our between-subject data, the Wilcoxon test for the interaction effect in hypothesis 3 is unavailable.

³¹ One concern about applying a standard ordinary least squares model to our data is that it treats expectations as interval data. While it makes sense to treat expectations in this way as a conceptual matter, we elicit subjects' responses for only three levels of expectation. Therefore, as a robustness check, we dummy coded the expectations variable and ran a categorical regression. We find that the joint test of the equality of each of the treatments, which is statistically equivalent to the analysis of variance F -test, is significant at (at least) the 5 percent level in both studies (see Table C1).

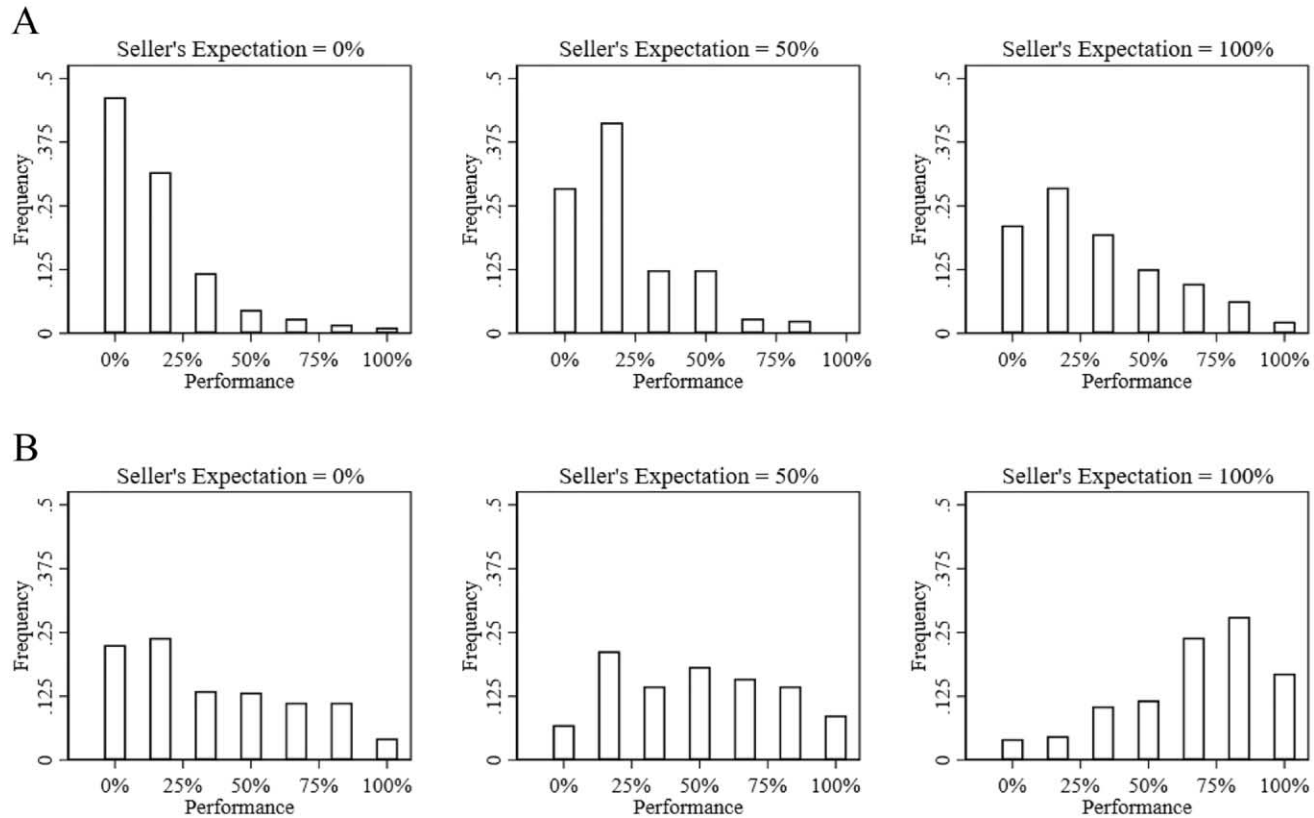


Figure 2. Distribution of responses across conditions for the combined sample. *A*, No-promise conditions; *B*, promise conditions

Table 2
Standard Ordinary Least Squares
Regressions for the Samples

	Study 1	Study 2	Combined
Promise	.24** (.03)	.18** (.02)	.20** (.02)
Expectations	.18** (.03)	.11** (.04)	.15** (.02)
Promise × Expectations	.11* (.04)	.16** (.06)	.15** (.03)
Cons	.13** (.02)	.17** (.02)	.15** (.01)
R ²	.33	.19	.27
N	783	771	1,554

Note. Standard errors are in parentheses.

* Significant at the 5 percent level.

** Significant at the 1 percent level.

even in the absence of a promise. Even in the no-promise condition, a shift in expectation from 0 to 100 percent increases the average likelihood of performance by 18 percentage points in study 1 and 11 percentage points in study 2. This effect is statistically significant at the 1 percent level in both studies. Moreover, in both the promise and, more crucially, the no-promise conditions, the average likelihood of performance increases with greater expectations for every pairwise comparison (Wilcoxon rank-sum test, $p < .01$).³² This strongly supports our hypothesized expectations per se effect as embodied in hypothesis 2.

Finally, in line with our hypothesized interaction effect as embodied in hypothesis 3, a shift in expectations from 0 percent to 100 percent increases the average likelihood of performance by 11 percentage points more in the promise conditions compared with the no-promise conditions in study 1 and by 16 percentage points more in study 2. This difference-in-difference effect is statistically significant at the 5 percent level for study 1 ($p = .015$) and at the 1 percent level for study 2 ($p = .005$). Our nonparametric test of hypothesis 3—the two-step bootstrapping procedure described in Online Appendix Section OAI—finds an interaction effect that is significant at the 5 percent level in both studies (in study 1, $p = .034$; in study 2, $p = .025$).

5. Discussion

Our results provide a more nuanced picture of the determinants of promise keeping than the experimental literature has done to date. We are, to our knowledge, the first to provide clear and direct evidence of a promising per se effect. The previous literature, notably Charness and Dufwenberg (2006) and Vanberg (2008), provides evidence that a dictator's promise to a recipient in a dictator

³² The only exception is the comparison of 0 percent and 50 percent in the no-promise conditions in study 2, which is significant at the 5 percent level (Wilcoxon rank-sum test, $p = .03$).

game increases the likelihood that she will perform. However, both papers document this effect for cases in which the recipient's performance expectations are positive. It would therefore be wrong to infer from the previous literature that promises matter irrespective of promisees' expectations. Indeed, if dictators care only about a recipient's expectations if they have made a promise, the increased performance levels that arise when there has been a promise could be explained entirely by the interaction effect. In other words, it is not clear from the previous literature that promises have any effect in the absence of positive expectations. A Scanlonian promisor will not feel obligated to keep his promises if the promisee is sure that he will not perform. To empirically isolate the promising per se effect, it is therefore crucial to show that promises matter even if expectations are 0, which we are able to demonstrate.³³

Our finding that expectations matter even in the absence of promising is in line with the theory proposed by Charness and Dufwenberg (2006) and results of Reuben, Sapienza, and Zingales (2009) and Regner and Harth (2014).³⁴ It runs counter to Vanberg (2008) and Ellingsen et al. (2010), who find evidence that expectations are not independently significant.³⁵

Our findings also contradict the strong version of the interaction effect that is embodied in Ederer and Stremitzer's (2017) theory of conditional guilt aversion. But we find that promising makes the promisor more sensitive to the promisee's expectations, an interaction effect that supports a weaker version of that theory. This finding accounts for the fact that the previous literature finds clear evidence for the relevance of expectations if there was a promise (Ederer and Stremitzer 2017) but at best mixed evidence for their relevance if there was no promise.³⁶

To our knowledge, we are the first to provide clear empirical evidence for the

³³ It is inherently difficult to create expectations of 0 in the recipient of a promise in a controlled fashion if this promise arises naturally out of the interaction of two experimental subjects. This is presumably the reason why this result has remained elusive until now. Our vignette study enables us to demonstrate the existence of a pure promising per se effect because we simply tell our promisors that the promisee's expectations are 0.

³⁴ However, Reuben, Sapienza, and Zingales (2009) employ a lost-wallet game in which arguably there is a preexisting duty to return the wallet, which could have a similar effect to a promise.

³⁵ Vanberg (2008) isolates expectations from promising by destroying the promissory link between the subjects but, in doing so, destroys any link between the two parties. Therefore, he really compares a relationship of promising with no relationship at all. Likewise, in Ellingsen et al. (2010) there is no relationship between the parties. Parties play a dictator game in which the experimenter communicates the recipients' elicited expectations to their dictators (unbeknownst to the recipients).

³⁶ It is possible that subjects make implicit assumptions about the extent to which the seller relies on the buyer's statement and that those assumptions vary systematically across our six conditions such that our results are driven by subjects' different beliefs about the seller's reliance rather than by the variables we are interested in. While we cannot rule out this possibility, we are doubtful that differential assumptions about reliance can explain the observed interaction effect or the promising per se effect (though it might be that the expectations per se effect is really an "expectations and reliance per se effect"). This is because a rational seller's reliance ought to be a function of her beliefs about the buyer's behavior, and while it is true that promising often makes a promisee more confident that the promisor will perform, our vignettes tell subjects what the seller's beliefs are. Holding constant those beliefs, the buyer has no reason to suppose that the seller will rely more in situations in which he made a promise, so long as the buyer believes that the seller is rational.

presence of an interaction effect. Indeed, Di Bartolomeo et al. (2019) find evidence inconsistent with the interaction effect, claiming that expectations matter only if there is no promise.³⁷ But their findings may largely be an artifact of their binary design of the performance decision, which forces dictators to choose between keeping and not keeping the promise. A dictator who has made a promise will believe that a promisee's greater expectations enhance his moral reasons for keeping his promise without this having any effect on his performance decision, if he believes that his promise alone gives him a sufficient reason to perform. Consider a promisor who in our continuous specification would choose "I am likely to perform" when the promisee's expectations are low and "I am very likely to perform" when the promisee's expectations are high. In the binary specification of the performance decision, this subject would likely just choose "perform" for both the low- and high-expectations scenarios, thus obfuscating the fact that he believes that there are stronger reasons to perform in the high-expectations scenario than in the low-expectations scenario. In line with this possibility, we see evidence of a ceiling effect in our own design, as the interaction effect for the third quartile in our quantile regressions is negative.³⁸ Of course, deciding whether or not to keep a promise is a binary decision. But, in many situations, as when performance is a continuous variable (Edlin and Reichelstein 1996), the promisor can choose between more and less egregious breaches of a promise. And even in settings where breach does not admit degrees, there are good instrumental reasons for measuring a subject's inclinations to keep a promise on a more sensitive scale, rather than the resulting binary decision whether or not to breach, particularly if one wants to extrapolate the findings to higher-stakes settings.

A notable feature of our design is that subjects were asked to imagine that the buyer communicated with the seller even when the buyer made no promise. In particular, we asked subjects to imagine that the buyer made a statement that he intended to buy the product while making clear that he did not promise that he would do so. This suggests that, holding expectations constant, it is the promise, and not merely the fact that the buyer and seller communicated with one another, that increases the buyer's willingness to perform in the promise condition. But of course we cannot determine whether expectations matter independently of communication. Nor can we determine whether our expectations per se effect is driven more specifically by the fact that the communication takes the form of a statement of intent. Subjects may have judged that the buyer's statement of intent

³⁷ Their findings contradict the theory and the experiment in Ederer and Stremitzter (2017), in which expectations matter if there was a promise, and in Vanberg (2008) and Ellingsen et al. (2010), in which there is no effect of expectations if there was no promise.

³⁸ In our design, performance is bounded above. Subjects cannot perform at a higher level than performing "for sure." So for subjects who feel duty bound to honor their promises no matter what—and who would therefore already be very willing to perform even for low expectations—there is not much scope for increasing their levels of performance as promisees' expectations increase. Our data are consistent with this explanation. A quantile regression run at the first and second quartiles has a positive coefficient for the interaction effect ($\gamma = .17$ and $\gamma = .33$, respectively; $p < .01$). However, the coefficient of the interaction effect for a quantile regression run at the third quartile is negative ($\gamma = -.17$; $p < .05$). See Table C2.

was made in order to alter the seller's expectation and hold the buyer responsible for it, albeit to a lesser extent than when the buyer made a promise. To determine whether expectations matter independently of communication, we would need to construct a vignette in which a party learns his counterparty's expectations of performance in the absence of any prior communication between the parties.³⁹ And to determine the extent to which the statement of intent is driving the results, we would need to construct a vignette with communication by the buyer that could not be construed as intended to convey information about his future actions.

A limitation of our design is that the scenario with which subjects were presented was purely hypothetical, and since subjects were paid a fixed fee for completing the survey, they had no pecuniary incentive to answer honestly. But it is difficult to study the effects of promising and expectations in a controlled manner in an incentivized experiment in which there is real interaction between subjects and subjects choose whether to make promises to one another, because it is difficult to manipulate subjects' expectations in a systematic fashion under such circumstances.⁴⁰ One possibility would be to configure the experiment so that different subjects ought rationally to form different expectations about some event by manipulating the exogenous uncertainty that they face (Ederer and Stremitzer 2017). But the event that is ultimately of interest here is not a move of nature. It is an action by a subject who may or may not have made a promise. And expectations about such actions will be informed by subjects' priors about the likelihood that people stick to their stated plans and keep their promises.

Moreover, we are ultimately interested not in a potential promisee's beliefs about the actions of the potential promisor (the seller's beliefs in our scenario) but rather the potential promisor's second-order beliefs about the expectations of the potential promisee (the buyer's beliefs about the seller's expectations). And if subjects' decisions had an effect on the payments they receive, we could not safely rely on the potential promisee to honestly report her expectations to the potential promisor. This is because she would have an incentive to overstate her expectations, if, as we hypothesize, the potential promisor is more likely to perform when he believes that she expects him to perform.

One advantage of our vignette design is that we can circumvent these problems by simply telling promisors what the promisee expects them to do in each scenario. Another is that we can cleanly study the effect of promises by comparing a scenario in which the potential buyer makes a promise with a scenario in which the potential buyer makes a statement of intention but explicitly disclaims a promise. Previous experiments compared scenarios in which parties have the opportunity to communicate and exchange promises with scenarios in which ei-

³⁹ The evidence presented in Reuben, Sapienza, and Zingales (2009) suggests that recipients' expectations in a dictator game might be important even absent communication.

⁴⁰ Ellingsen et al. (2010) use self-reporting that admits the possibility of strategic communication. Reuben, Sapienza, and Zingales (2009) use a multistage game in which the experimenter reports recipients' expectations from a previous stage of the game on the assumption that expectations stay constant over time.

ther the parties cannot communicate (Charness and Dufwenberg 2006) or the parties have not communicated with the parties they face when making their performance decisions (Vanberg 2008). Therefore, previous experiments could not cleanly disentangle the effects of communication from the effects of the promise, and it is not clear how a design that makes promising endogenous would be able to do so.⁴¹

6. Conclusion

Our paper provides a unified framework for studying the effect of promises and expectations on performance. We are able to document a promising per se effect, according to which promises matter regardless of the promisee's expectations, an expectations per se effect, according to which expectations matter even in the absence of a promise, and an interaction effect, according to which promising makes a promisor more sensitive to a promisee's expectations.

Our between-subject data do not allow us to determine whether our findings result from the presence of different pure types in the subject population (compassionate, Kantian, Scanlonian) or subjects who exhibit all three dispositions in a weighted combination. Subject to our earlier caveats about demand effects, our within-subject data provide us with information about the composition of the population of subjects.⁴² Fifty percent of our subjects conform to our model.⁴³ Of those subjects, 70 percent exhibit a compassionate disposition,⁴⁴ 79 percent exhibit a Kantian disposition,⁴⁵ and 55 percent exhibit a Scanlonian disposition over at least part of the range.⁴⁶ Each of the Scanlonian subjects also exhibit both of the other dispositions. We leave a more systematic inquiry into this heteroge-

⁴¹ Of course, one could imagine coding free-form communication more finely by asking coders to distinguish between promises and statements of intentions or giving subjects the possibility of choosing between two precoded messages, for example, "I promise to buy" and "I plan on buying but do not promise." However, such a design would introduce selection effects, as one could not be confident that subjects who make promises are similar to those who make mere statements of intent.

⁴² As noted above, presenting subjects with all six conditions makes salient to them the variables that we are studying, which raises the concern that their responses will be distorted. They might, for example, modify their answers to conform with their beliefs about the researchers' hypotheses or to create a false impression of consistency. See Charness, Gneezy, and Kuhn (2012).

⁴³ They exhibit a likelihood of performance that is weakly increasing in expectations, weakly greater when there is a promise, and increasing at a weakly greater rate in expectations when there is a promise. Most of the remaining subjects behave irregularly on at least one of the three dimensions, while the remainder are contrarian on at least one of the three dimensions by exhibiting a likelihood of performance that is either strictly decreasing in expectations over at least part of the range, strictly lower when there is a promise over at least part of the range, or increasing at a strictly lower rate in expectations when there is a promise.

⁴⁴ They exhibit a likelihood of performance that is always nondecreasing and strictly increasing in expectations over at least part of the range either when there is a promise or when there is no promise.

⁴⁵ Their likelihood of performance is always at least as great when there is a promise as when there is no promise and strictly greater when there is a promise over some of the range.

⁴⁶ They exhibit a compassionate disposition, and the rate of increase in the likelihood of performance as expectations increase is at least as great when there is a promise as when there is no promise and is strictly greater over at least part of the range.

neity of subject types and the possible relationship of those types to personality traits discussed in the psychology literature to future research.

Appendix A

Vignettes

A1. Promise Conditions

Seller B is offering a product for sale for \$100 that you are interested in buying. You are currently out of town for three days and therefore are unable to go to B's shop and buy the product immediately. But B may have the opportunity to sell the product to somebody else in the meantime, so you promise B that you will buy the product upon your return. The conversation proceeds as follows:

B says: "I would be willing to sell the product to you, but someone else might offer to buy it in the meantime. Why should I wait to sell the product to you?"

You say: "Well, I promise I will buy it from you upon my return."

B immediately concludes that there is a [0%/50%/100%] chance that you are going to keep your promise to buy the product from him. Imagine you know this. On the day you want to buy the product from B, you accidentally learn that another seller (C) is offering to sell an equivalent product at the price of \$85, which is \$15 less than the price that B is charging.

So the situation is this: C is offering to sell you the product at a lower price. You have made a promise to B to buy the product from him, and the product is still available. You also know that [B is sure that you will not/B thinks there is a 50% chance that you will/B is sure that you will] keep your promise to buy the product from him. How likely is it that you would choose to buy the product from the second seller C in this scenario?

A2. No-Promise Conditions

Seller B is offering a product for sale for \$100 that you are interested in buying. You are currently out of town for three days and therefore are unable to go to B's shop and buy the product immediately. But B may have the opportunity to sell the product to somebody else in the meantime. The conversation proceeds as follows:

B says: "I would be willing to sell the product to you, but someone else might offer to buy it in the meantime. Why should I wait to sell the product to you?"

You respond: "All I can say is that I plan to buy it from you, though I can't promise that I will do so."

B immediately concludes that there is a [0%/50%/100%] chance that you are going to buy the product from him. Imagine you know this. On the day you want to buy the product from B, you accidentally learn that another seller (C) is offering to sell an equivalent product at the price of \$85, which is \$15 less than the price that B is charging.

So the situation is this: C is offering to sell you the product at a lower price. You

have made no promise to B to buy the product from him, but the product is still available. You also know that [B is sure that you will not/B thinks that there is a 50% chance that you will/B is sure that you will] buy the product from him. How likely is it that you would choose to buy the product from the second seller C in this scenario?

Appendix B

Postexperiment Questionnaire

B1. Control Questions

1. In the given scenario, did you make a promise to B to buy her product? (Yes/No)

2. Please indicate the expectations B had whether you are going to buy her product. (0 percent/50 percent/100 percent)

3. Who made the better offer? (the first seller B/the second seller C).

In study 1, 78 of 783 subjects answered at least one of these control questions incorrectly. We use subjects' responses to these questions to construct a robustness check in our statistical tests below.

In study 2, 362 of 771 subjects answered one of the control questions incorrectly. This large number is mainly due to the second question. We have responses to the first question from only 769 subjects due to two missing values. Of those, 720 (93.63 percent) were correct and 49 (6.37 percent) were incorrect. The second question was answered correctly by 429 (55.64 percent) subjects and incorrectly by 342 (44.36 percent) subjects. The third question was answered correctly by 755 (97.92 percent) subjects and incorrectly by 16 (2.08 percent) subjects.

B2. Questions to Assess Subjects' Care and Honesty

1. I didn't take the scenario seriously. I just wanted to earn the \$1.00 fee as quickly as possible. (yes: 12 of 783)

2. I carefully read the instructions. (no: 4 of 783)

3. I chose my answers to make myself seem like a good person. (yes: 69 of 783)

4. This is the first time I have completed this survey. (no: 16 of 783)

We have no good explanation for the 16 subjects who self-reported having taken the survey before. We provided links to subjects that were good for only a single log in. We implemented filters preventing subjects (as identified by their MTurk IDs) from participating who had participated in pilots of our experiment or similar experiments we had run in the past. So the only reason for the 16 self-reported repeat takers could be that subjects have multiple MTurk IDs or mistakenly checked the wrong box.

We slightly changed the first and fourth questions in study 2. These questions were not forced choice, which is why we have a couple of missing values.

1. I only participated out of fun, without taking the questionnaire seriously. (yes: 46; no: 704; missing value: 21)
2. I carefully read the instructions. (yes: 758; no: 12; missing value: 1)
3. I chose my answers to make myself seem like a good person. (yes: 50; no: 700; missing value: 21)
4. This was the first study I ever participated in. (yes: 10; no: 739; missing value: 22)

As a general note, there is one Hagen Decision Lab ID that appears twice in the data set, which suggests a repeat participation. However, Unipark does not allow anyone to take part twice in the same study. We do not know whether the same person participated a second time from a different computer or whether a different person entered the same ID by mistake.

B3. Demographic Questions

1. What is your age? (study 1: range of 18–74, average of 35; study 2: range of 18–81, average of 34)
2. What is your gender? (study 1: 49 percent female; study 2: 63 percent female)
3. What is your highest level of schooling? (study 1: master's degree or more: 11 percent, bachelor's degree: 41 percent, associate's degree: 16 percent, vocational or technical certificate or diploma after high school [such as cosmetics]: 7 percent, high school diploma: 24 percent, I did not complete high school: 1 percent; study 2: master's degree: 150 [19.46 percent], bachelor's degree: 163 [21.14 percent], vocational or technical certificate or diploma after high school: 158 [20.49 percent], A-levels: 279 [36.19 percent], high school diploma: 13 [1.69 percent] [*Realschulabschluss*] + 5 [.65 percent] [*Hauptschulabschluss*], did not complete high school: 3 [.39 percent]).
4. Is English (study 2: German) your first language? (study 1: 98 percent yes; study 2: 92.48 percent yes).

Appendix C

Additional Tables

Table C1
Categorical Regressions for the Samples

	Study 1	Study 2	Combined
Expectations per se:			
0/100	.18**	.15**	.16**
0/50	.09**	.02	.05*
50/100	.09**	.12**	.10**
Joint test (p -value)	<.01**	<.01**	<.01**
Interaction effect:			
0/100	.11*	.16*	.15**
0/50	.04	.09*	.07*
50/100	.07	.07	.07*
Joint test (p -value)	.05*	.01*	<.01**
N	783	771	1,554

* Significant at the 5 percent level.

** Significant at the 1 percent level.

Table C2
Quantile Regressions by Quantile

	.25	.50	.75
Promise	.17** (.03)	.17** (.03)	.50** (.04)
Expectations	.17** (.04)	.00 (.04)	.33** (.06)
Promise \times Expectations	.17** (.06)	.33** (.05)	-.17* (.08)
Cons	.00 (.02)	.17** (.02)	.17** (.03)
Pseudo- R^2	.14	.21	.22

Note. Standard errors are in parentheses. $N = 1,554$.

* Significant at the 5 percent level.

** Significant at the 1 percent level.

References

- Amazon Mechanical Turk. 2012. Improving Quality with Qualifications: Tips for API Requesters. August 28. <https://blog.mturk.com/improving-quality-with-qualifications-tips-for-api-requesters-87eff638f1d1>.
- Battigalli, Pierpaolo, and Martin Dufwenberg. 2007. Guilt in Games. *American Economic Review* 97:170–76.
- . 2009. Dynamic Psychological Games. *Journal of Economic Theory* 144:1–35.
- Bénabou, Roland, and Jean Tirole. 2003. Intrinsic and Extrinsic Motivation. *Review of Economic Studies* 70:489–520.

- Braver, Sanford L. 1995. Social Contracts and the Provision of Public Goods. Pp. 69–86 in *Social Dilemmas: Perspectives on Individuals and Groups*, edited by David A. Schroeder. Westport, CT: Praeger.
- Bull, Clive. 1987. The Existence of Self-Enforcing Implicit Contracts. *Quarterly Journal of Economics* 102:147–60.
- Charness, Gary, and Martin Dufwenberg. 2006. Promises and Partnership. *Econometrica* 74:1579–1601.
- Charness, Gary, Uri Gneezy, and Michael A. Kuhn. 2012. Experimental Methods: Between-Subject and Within-Subject Design. *Journal of Economic Behavior and Organization* 81: 1–8.
- Di Bartolomeo, Giovanni, Martin Dufwenberg, Stefano Papa, and Francesco Passarelli. 2019. Promises, Expectations, and Causation. *Games and Economic Behavior* 113:137–46.
- Dixit, Avinash. 2009. Governance Institutions and Economic Activity. *American Economic Review* 99:5–24.
- Dufwenberg, Martin, and Uri Gneezy. 2000. Measuring Beliefs in an Experimental Lost Wallet Game. *Games and Economic Behavior* 30:163–82.
- Ederer, Florian, and Alexander Stremitzler. 2017. Promises and Expectations. *Games and Economic Behavior* 106:161–78.
- Edlin, Aaron S., and Stefan Reichelstein. 1996. Holdups, Standard Breach Remedies, and Optimal Investment. *American Economic Review* 86:478–501.
- Ellingsen, Tore, and Magnus Johannesson. 2004. Promises, Threats, and Fairness. *Economic Journal* 114:397–420.
- Ellingsen, Tore, Magnus Johannesson, Sigve Tjøtta, and Gaute Torsvik. 2010. Testing Guilt Aversion. *Games and Economic Behavior* 68:95–107.
- Faul, Franz, Edgar Erdfelder, Axel Buchner, and Albert-Georg Lang. 2009. Statistical Power Analyses Using G*Power 3.1: Tests for Correlation and Regression Analyses. *Behavior Research Methods* 41:1149–60.
- Grossman, Sanford J., and Oliver D. Hart. 1986. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy* 94:691–719.
- Grout, Paul A. 1984. Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach. *Econometrica* 52:449–60.
- Hart, Oliver, and John Moore. 1988. Incomplete Contracts and Renegotiation. *Econometrica* 56:755–85.
- Holmström, Bengt. 1979. Moral Hazard and Observability. *Bell Journal of Economics* 10:74–91.
- Horton, John J., and Lydia B. Chilton. 2010. The Labor Economics of Paid Crowdsourcing. Pp. 209–18 in *EC '10: Proceedings of the 11th Association for Computing Machinery Conference on Electronic Commerce*. New York: Association for Computing Machinery.
- Ismayilov, Huseyn, and Johannes J. M. Potters. 2012. Promises as Commitments. Discussion Paper No. 2012-064. Tilburg University, Center for Economic Research, Tilburg.
- Klein, Benjamin, and Keith B. Leffler. 1981. The Role of Market Forces in Assuring Contractual Performance. *Journal of Political Economy* 89:615–41.
- Kreps, David M. 1996. Corporate Culture and Economic Theory. Pp. 221–75 in *Firms, Organizations, and Contracts: A Reader in Industrial Organization*, edited by Peter J. Buckley and Jonathan Michie. New York: Oxford University Press.
- Levin, Jonathan. 2003. Relational Incentive Contracts. *American Economic Review* 93:835–57.

- Macaulay, Stewart. 1963. Non-contractual Relations in Business: A Preliminary Study. *American Sociological Review* 28:55–67.
- MacLeod, W. Bentley, and James M. Malcomson. 1989. Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment. *Econometrica* 57:447–80.
- Mason, Winter, and Duncan J. Watts. 2010. Financial Incentives and the “Performance of Crowds.” *ACM SigKDD Explorations Newsletter* 11:100–108.
- Mirrlees, James A. 1976. The Optimal Structure of Incentives and Authority within an Organization. *Bell Journal of Economics* 7:105–31.
- Ostrom, Elinor, James Walker, and Roy Gardner. 1992. Covenants with and without a Sword: Self-Governance Is Possible. *American Political Science Review* 86:404–17.
- Regner, Tobias, and Nicole S. Harth. 2014. Testing Belief-Dependent Models. Working paper. Max-Planck Institute of Economics, Jena.
- Reuben, Ernesto, Paola Sapienza, and Luigi Zingales. 2009. Is Mistrust Self-Fulfilling? *Economics Letters* 104:89–91.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Shiffrin, Seana Valentine. 2008. Promising, Intimate Relationships, and Conventionalism. *Philosophical Review* 117:481–524.
- Stone, Rebecca, and Alexander Stremtizer. Forthcoming. Promises, Reliance, and Psychological Lock-In. *Journal of Legal Studies*.
- Vanberg, Christoph. 2008. Why Do People Keep Their Promises? An Experimental Test of Two Explanations. *Econometrica* 76:1467–80.
- Wilkinson-Ryan, Tess, and David A. Hoffman. 2015. The Common Sense of Contract Formation. *Stanford Law Review* 67:1269–1301.
- Williamson, Oliver E. 1979. Transaction-Cost Economics: The Governance of Contractual Relations. *Journal of Law and Economics* 22:233–61.
- . 1985. *The Economic Institutions of Capitalism*. New York: Free Press.