Trust and Promises over Time*

Florian Ederer† Frédéric Schneider‡

Yale University University of Cambridge

January 3, 2020

Abstract

Using a large-scale hybrid laboratory and online trust experiment with pre-play communication this paper investigates how the passage of time affects trust, trustworthiness, and cooperation. We provide evidence for the persistence of both trustworthiness and promises over time. Even when three weeks pass between the decision to trust and the message of the trustee, and actual choices and even when these choices are made outside of the lab, communication (predominantly through the use of promises) raises cooperation, trust, and trustworthiness by about 50 percent. Delays between the beginning of the interaction and the time to reciprocate neither substantially alter trust or trustworthiness nor affect how subjects choose to communicate.

Keywords: trust, promises, persistence, trustworthiness, delay

JEL codes: C91, C72, D83

*We thank Antonio Arechar, David Cesarini, Gary Charness, Stefano DellaVigna, Eugen Dimant, Martin Dufwenberg, Guillaume Fréchette, Jana Gallus, Nadja Jehli, Emir Kamenica, Navin Kartik, Corinne Low, Song Ma, Barry Nalebuff, David Rand, Fiona Scott Morton, Andrew Schotter, Kelly Shue, Ted Snyder, Alexander Stremitzer, Roberto Weber, Kevin Williams, Jacqueline Würth, and seminar audiences at University of Basel, Cambridge Judge, NYU Economics, UPenn Wharton, Yale SOM, the Behavioral Personnel Economics Conference, the ESA Asia-Pacific Conference, the Nordic Conference on Behavioral and Experimental Economics, the Tsinghua BEAT Conference, and the UPenn Norms and Behavioral Change Workshop for helpful comments as well as Lucas Molleman for lioness programming support.

†Yale School of Management, Yale University, 165 Whitney Avenue, New Haven, CT 06511, U.S.A., florian.ederer@yale.edu.

‡Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, United Kingdom, f.schneider@jbs.cam.ac.uk.

1 Introduction

Trust is the foundation for many social and economic interactions and acts as "an important lubricant of the social system" (Arrow, 1974). In particular, when contractual or reputational incentives are weak, the mutual trust that each party will hold up its end of the bargain becomes particularly important. Time plays a crucial role in this context: Trust often has to be repaid later, and promises have to be made good only after some time has passed. Indeed, the very purpose of promises is to facilitate production and exchange *over time* and to foster the belief in *future* actions.[1]

Casual observation suggests that in a variety of settings the more the act of trust, or the promise, recedes into the past, the likelier it is for the promisor to neglect her obligation and to behave opportunistically. For example, a stereotype about politicians is that they promise everything to get elected, but over time conveniently "forget" their promises. The German post-war chancellor Konrad Adenauer famously uttered, "What do I care about my chitchat from yesterday?" (Weymar, 1955). Proverbs from various cultures, such as "Evening promises are like butter: morning comes, and it's all melted," or "Promises are like the full moon—if they are not kept at once they diminish day by day," also encapsulate this view.[2] In an organizational context, Paine (2004) argues that "most companies are simply not designed to remember and keep promises over time." This paper provides experimental evidence for the robustness of trust, trustworthiness, and promises over time. We document that communication has a large effect on trust, trustworthiness, and cooperation that remains remarkably persistent over time.

Specifically, to investigate how the passage of time affects trust, trustworthiness, and the power of communication (and promises in particular) we analyze the behavior of over 700 subjects in a one-shot, two-person trust game with pre-play communication.[3] The main innovation of our paper is the introduction of delay between the trustee's message/the trustor's choice to trust and the trustee's decision whether to reciprocate. Trustees in our experiment made their choice either (i) *immediately* during the laboratory session (as is common in all previous trust game studies), (ii) in a 24-hour window *after leaving the laboratory*, or (iii) in a 24-hour window *21 days after the*

[1] The word's Latin roots, "pro" (forward) and "mittere" (send) also reflect this crucial timing aspect.

[2] See Hourani (2012) and Braude (1965) for these Moroccan and German proverbs, respectively.

[3] In a trust game, a "trustor" can send an amount of money to a "trustee." The amount of money is multiplied so that sending money is socially efficient. The trustee, however, is free to return money to the trustor or not, introducing an element of moral hazard. The trust game is seen as a vehicle to study trust (amount sent by trustor) and trustworthiness (amount sent back by the trustee).

laboratory session. This allows us to analyze the differential impact of time on trust, trustworthiness, and promise-keeping, and to compare our results to standard laboratory scenarios without any time delay.

We find that pre-play communication raises cooperation by about 50 percent. More importantly, we document that the increase in cooperation, trust and trustworthiness resulting from communication is not diminished even when three weeks pass before the trustee's actual decision. Thus, we provide evidence for the lasting power of communication and of promises in particular. In addition, we show that the behavior in trust games without pre-play communication is essentially unaffected by the additional delays introduced as part of our experimental design. Taken together, our results suggest that the findings of previous laboratory studies on trust and communication extend to more externally valid scenarios in which subjects choose actions outside the lab and do so long after they made promises.

Trust has been extensively studied in laboratory experiments, mostly using trust games beginning with the seminal work of Berg et al. (1995). Most studies find that trustees behave, on average, trustworthily: they honor the trust that trustors put in them by sending back a substantial part of the money. A recent meta-analysis (Johnson and Mislin, 2011) counted 162 replications of the original game, all showing the same pattern. This appears to be good news, as it suggests that many real-world problems involving hidden information or hidden action can be mitigated through trust and trustworthiness. Better still, several recent studies offer experimental evidence that pre-play messages by trustees (e.g., promises), even if they come in the form of mere cheap talk, considerably enhance trust and subsequent levels of cooperation in trust and dictator games (Ellingsen and Johannesson, 2004; Charness and Dufwenberg, 2006; Vanberg, 2008; Ben-Ner and Putterman, 2009; Charness and Dufwenberg, 2011; Serra-Garcia et al., 2013; Ismayilov and Potters, 2016; Ederer and Stremitzer, 2017; Di Bartolomeo et al., 2017; Bhattacharya and Sengupta, 2017; Casella et al., 2018).[4] Although these messages are non-binding, they seem to sharpen the trustees' sense of obligation or the amount of guilt that is attached to letting the trustor down.

However, all these studies are limited in one important dimension: time.[5] In previous lab

[4]Notable contributions to the broader literature on promise-keeping in political science and social psychology include Ostrom et al. (1992), Kerr and Kaufman-Gilliland (1994), Sally (1995), and Bicchieri and Lev-On (2007).

[5]The delay of *outcomes* has long been studied. For an introduction to this literature, see, for example, Frederick et al. (2002).

experiments using trust games (with and without communication), hardly any time passes during the period after the trustor's choice (and, if possible, the trustee's communication) and before the trustee's choice (in particular, the choice to deliver on any previously sent promise). It is thus possible that the very short time horizon in lab situations between trustor choices (and trustee messages) on the one hand and trustee choices on the other hand leads to an overestimation of the extent of reciprocity and trustworthiness that we would expect in everyday interactions. We can think of a couple of reasons why this may be the case.

First, people may be less likely to honor another person's trust as time passes as they may feel less guilty. The feeling of obligation may be strongest immediately after the trustor has decided to trust the trustee. As time passes, and the triggering event becomes more abstract in the trustee's memory, the feeling of obligation may become weaker, and, by extension, the feeling of guilt associated with letting the trustor down. This is closely related to the multiple-selves literature that sees an individual is a sequence of short-term selves (Fudenberg and Levine, 2006; Jamison and Wegener, 2010; Alós-Ferrer and Strack, 2014). Because one's present self is more closely tied to yesterday's or tomorrow's self than to one's self in the more distant past or future, an individual may remember or anticipate more about, and empathize more with, selves from the recent past or near future than those from the more distant past or future (Bénabou and Tirole, 2004). Parfit (1973) argues that, because psychological connectedness between selves diminishes over time, moral obligations of promises are stronger over the short term than over the long term.[6]

Second, delays may introduce the potential for moral wiggle room, as memories become vague over time, and thus can be more prone to self-serving bias. For example, Li (2013) finds that trustees who defected on their trustor rate their own defection as less unkind as time passes: if asked immediately after their defection, they rate this act as very unkind (9.6 on a Likert scale from 0, very kind, to 10, very unkind); if asked after 43 days, that rating drops to 6. Similarly, Kouchaki and Gino (2016) find that subjects tend to forget their unethical behavior, a phenomenon that the authors call "unethical amnesia."[7]. This is especially important for promise-keeping. We conjecture that subjects may form a biased memory of the promises they made, which could make it easier

[6] Although some psychological studies find behavioral effects consistent with the theory of psychological connectedness between multiple selves (Bartels and Rips, 2010; Van Gelder et al., 2013), the only existing test in economics finds no explanatory power of psychological connectedness on intertemporal choice in individual decision-making (Frederick, 2003).

[7] Carlson et al. (2018); Saucet and Villeval (2019) find similar results in dictator games.

for them to renege on them. The misremembering of promises should be easier as time passes and memory becomes vaguer.

Some related studies investigate the effect of time pressure and very short delays (up to 15 minutes) on decisions in social dilemmas (Cone and Rand, 2014; Neo et al., 2013; Rand et al., 2014; Bouwmeester et al., 2017). Bhattacharya et al. (2019) study whether the timing of communication (before or after actions are chosen) affects cooperation. Focusing on a short time frame of four minutes in a lab setting they provide some evidence that individuals are most cooperative closest to the time of communication. Two studies allow for a one-day delay for reflection: Imas et al. (2016) make customers wait for one day before they can redeem coupons, which leads them to make more patient purchasing decisions; Andersen et al. (2018) delay participants' decisions in a dictator game by one day and find no significant difference in the amount given. Other studies delay *payments*, but not decisions in experimental social dilemmas (Kovarik, 2009; Dreber et al., 2016). In these latter studies, however, all decisions are still made immediately, during the laboratory session, and effects on behavior are mixed.[8]

The remainder of the paper proceeds as follows. Section 2 describes our experimental design. Section 3 presents our experimental results. Section 4 offers concluding remarks.

2 Experimental Design

We use the one-shot, two-person trust game introduced by Charness and Dufwenberg (2006). Figure 1 shows the game tree. The trustor ("he") has two options: choosing a safe outside option ("OUT") with equitable but relatively low payoffs ($15 for each player); or choosing to continue the game ("IN"). If the trustor chooses IN, the trustee ("she") has two options: one that generates a high payoff for herself ($42) but no money at all for the trustor ("DON'T ROLL"); and one where she receives $30, and the trustor's payoff depends on a lottery (or virtual die roll) with a 5/6 chance of gaining $36 and a 1/6 chance of receiving nothing ("ROLL"). Thus, the selfish decision for the

[8]Some studies include a kind of delayed promise but without measuring the effect of the delay. Meyer and Tripodi (2018) investigate pledges to donate blood at some later point and find that almost all subjects renege on their pledge; they do not specify or vary the time horizon of the pledge. Fosgaard and Soetevent (2018) study a donation drive by the Danish Refugee Council. Subjects were contacted via phone and could (endogenously) decide whether donate right now, during the phone call, or at an undefined later time. Subjects who pledged to pay at a later point were less likely to follow through. Andreoni and Serra-Garcia (2018) similarly find that endogenous pledges are rarely followed up.

trustee is to choose DON'T ROLL. If the trustor nonetheless chooses IN, he has to trust that the trustee deviates from the selfish choice and instead chooses ROLL.[9]

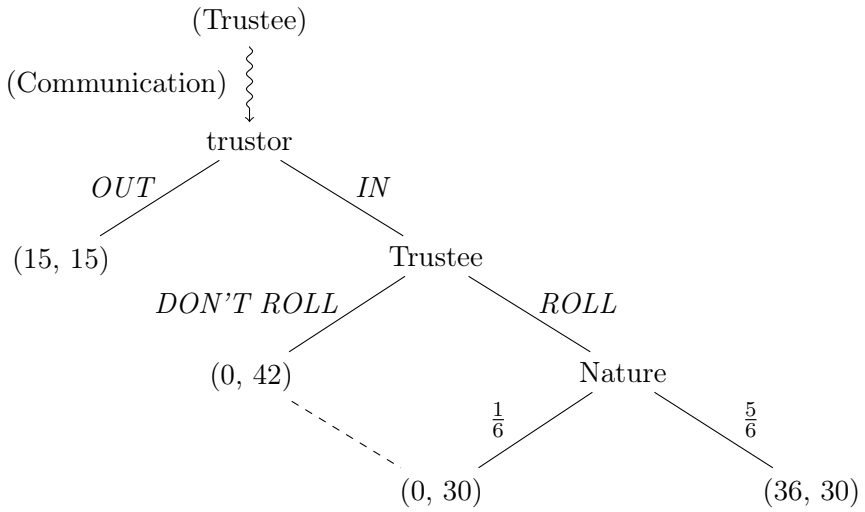


Figure 1: Game Tree. Note the information set for the trustor between "Don't Roll" and the negative die roll outcome. In both cases, the trustor receives 0. In COMM conditions, the trustee can send a message before the trustor makes his decision. In the NOCOMM conditions, the trustee cannot send a message.

Following Charness and Dufwenberg (2006), our first set of treatment conditions adds a communication stage before the trustor's choice. Depending on the treatment condition, the trustee can either write a free-form message to the trustor before the trustor makes his decision (COMM) or not (NOCOMM). COMM gives the trustee the opportunity to communicate any message she desires.[10] For example, she may promise that she will choose ROLL. The trustor, however, only observes his own payoff, which means that he remains uncertain whether the trustee chose ROLL or DON'T ROLL since even in this case there is a 1/6 chance of the trustor receiving nothing.

The trustee makes her decision not knowing whether the trustor has chosen IN or OUT. This allows us to record choices for all trustees, not just for those trustees paired with a trustor who chose IN.

Finally, we adapt the original design so that subjects make the above decision 10 times instead of once. Specifically, subjects are presented with a menu of 10 IN/OUT (ROLL/DON'T ROLL)

[9] Instructions were framed neutrally with the trustor referred to as "person A" and the trustee as "person B."

[10] Trustees were *not* reminded of their message before they made their ROLL decision but were asked how well they recalled the message *after* they made their decision. Neither trustees nor trustors were informed about the fact that trustees would not reminded of their message.

radio buttons. After they have made their decisions, the computer selects one of the choices at random to count as the final decision. We do this to allow subjects to submit a more nuanced strategy space, with the probability of IN (ROLL) ranging from 0 to 1 in increments of 0.1, instead of a stark binary choice. In particular, this creates the possibility for moral wiggle room when a trustee promises to choose ROLL.

2.1 Delayed Decision

Our primary experimental manipulation—and critical departure from the design of Charness and Dufwenberg (2006)—introduces a delay before the trustee's decision. All participants had to complete two short, web-based questionnaire forms, Q1 and Q2, in addition to the laboratory session. The primary purpose of these two questionnaires was to implement different conditions that varied the delay after which the trustees made their decisions: some trustees would decide in Q1 while others decided Q2. To hold the structure of the experiment constant across conditions and subject roles, every participant had to complete Q1 and Q2, which slightly differed in content.[11]

Subjects could access Q1 *in a 24-hour window after leaving the laboratory* and Q2 in a 24-hour window *21 days after the laboratory session*. To minimize attrition, we created strong monetary incentives to complete the online part of the study. At the end of the laboratory session, subjects just received their show-up fee. They received all the remaining payments only after completing both online questionnaires. Depending on the condition, the trustee chose ROLL/DON'T either just before filling out Q1 (EARLY) or just before filling out Q2 (LATE). In an additional condition, we administered Q1 *at the end of the laboratory session*, so that trustees made their decision while still in the laboratory (IMM). Importantly, the timing of the decisions was common knowledge among participants.

The crucial identification of the delay effect is the comparison between the EARLY and LATE conditions. We conducted the IMM condition to compare our results to the usual experimental setting of trust games, where subjects make all their decisions at once, during the laboratory session. We kept the decision-making process as similar to the EARLY and LATE conditions as possible.

[11] Q1 elicited subjects' first- and second-order beliefs and asked for free-form explanations of their laboratory decisions; Q2 revealed the existence of the other delay condition and asked for subjects' beliefs about the other delay. We were careful not to elicit trustee's ROLL choices before their beliefs. See Appendix A for a detailed timeline of tasks and questions in the lab part and the questionnaire (Q1 and Q2) part of the experiment.

Table 1: Treatment conditions, number of subjects in parentheses

	Immediate	Early	Late
No Communication	NOCOMM-IMM (114)	NOCOMM-EARLY (124)	NOCOMM-LATE (125)
Communication	COMM-IMM (116)	COMM-EARLY (112)	COMM-LATE (116)

Most importantly, in the IMM condition, subjects filled out the same browser-based questionnaire they would have seen in the other conditions, they still had to complete the Q2 questionnaire three weeks later, and they were paid only after completion of Q2. However, other aspects of the decision environment cannot be kept constant: in IMM sessions, trustees were still in the laboratory and made their decisions with the trustors and the experimenter still in the same room, which in itself may increase their propensity to choose ROLL.

Our 2×3 between-subjects design yields six different conditions, see Table 1.

2.2 Procedures

We conducted 40 experimental sessions with a total of 707 student subjects at the University of Zurich and at Yale University.[12] Subjects were assigned to visually partitioned computer terminals. At each terminal they found paper instructions, which were also read aloud by the experimenter. Questions were answered individually at the subjects' seats. Subjects interacted with another randomly chosen participant from the same session. All subjects were paid after 21 days with Amazon gift vouchers.[13] The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

3 Results

3.1 Subject Behavior

3.1.1 Attrition

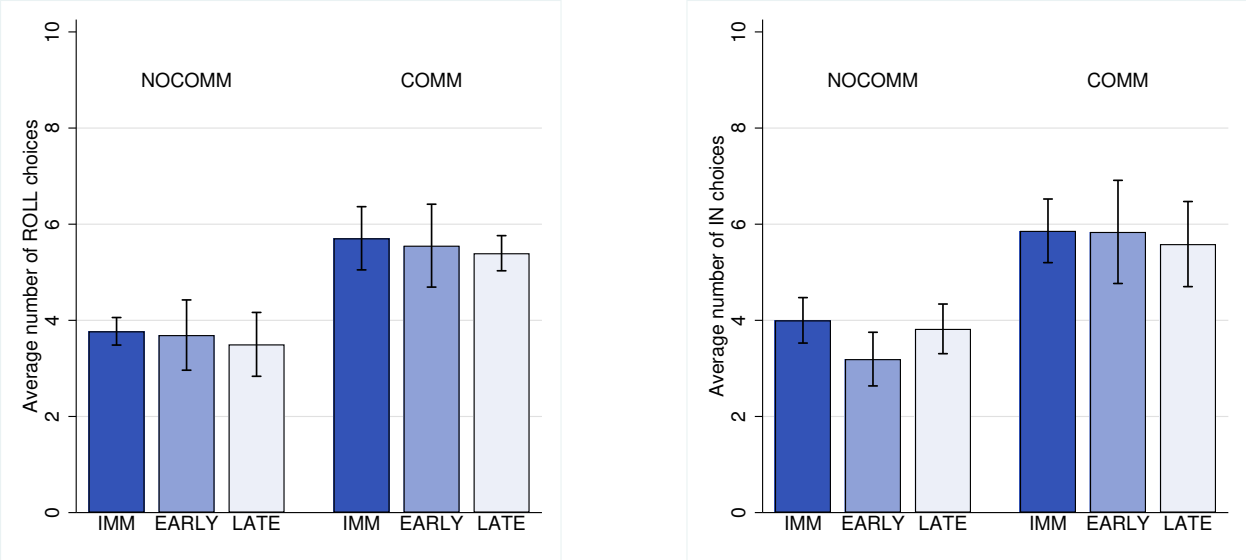
Despite the significant time delay between the two parts of the experiment, attrition rates are very low. Only 0.8 percent of trustees (2.3 percent of trustors) did not complete Q1 (the first of the

[12] Three subjects participated twice in this experiment due to walk-ins at the Yale lab. We removed their second appearances from the data.

[13] For the Zurich sessions, the money was mailed in cash to the recipients. The US dollar amounts were paid out at a ratio of 1:1 in Swiss francs. The exchange rate at the time of the experiment was about 0.966 CHF/USD.

two surveys). 5.1 percent of trustees (7.4 percent of trustors) did not complete Q2 (the second survey, administered three weeks after the initial session). By design, the Q1 attrition rate is zero for trustees in the IMM conditions. For trustees in the EARLY delay condition, the Q1 attrition rate is 1.7 percent, whereas for trustees in the LATE delay, the Q2 attrition rate is 3.3 percent. We find no systematic differences across delay conditions for trustees.[14] These low attrition rates are due to our experimental design which provided strong incentives to complete the additional surveys.

3.1.2 Choices



(a) Trustees: ROLL choices

(b) Trustors: IN choices

Figure 2: Roll Decisions; In Decisions. Error bars indicate standard errors, allowing for clustered errors on the session level.

We first study the behavior of trustees in our six treatment conditions. Recall that trustees made 10 decisions, one of which was randomly implemented. Trustees predominantly chose exclusively one option: 42 percent of trustees never chose ROLL, 33 percent always chose ROLL, and 25 percent chose ROLL between 1 and 9 out of the 10 decisions.[15]

[14] 18 trustees did not complete Q2: 13 out of 115 in IMM, 5 out of 118 in EARLY, and 4 out of 121 in LATE, $p = 0.25$, χ^2 test. The attrition rates are a bit more skewed for trustors: 13 out of 115 in IMM, 9 out of 118 in EARLY, and 4 out of 120 in LATE, $p = 0.06$, χ^2 test

[15] Trustors made less extreme decisions. 32 percent only chose OUT, 29 percent only chose IN, and 39 percent chose a combination of IN and OUT. We report the cumulative distributions in the appendix, Figure 4. We tested whether the distributions differ across the three delay conditions, separately for trustors and trustees and by communication condition. None of the tests is significant: ROLL NOCOMM $p = 0.778$, ROLL COMM $p = 0.885$, IN NOCOMM

Panel (a) of Figure 2 shows the average number of times a trustee chooses ROLL (recall that subjects made their decision 10 times). When there is no communication, trustees choose to roll on average between 35 percent to 38 percent of the time while this number increased to around 55 to 58 percent of the time when communication was possible. As expected, with immediate choices, communication substantially increases cooperation by trustees by about 50 percent, from 3.65 ROLL choices to 5.55 ROLL choices, consistent with previous studies on trust games with pre-play communication. More surprisingly, however, this positive effect of communication persists over time and is essentially unchanged across the three delay conditions. In contrast, within both NOCOMM and COMM conditions, the time delay between the initial session and the roll decision has only a negligibly small effect on roll rates.

Panel (b) shows that trustors are more trusting when communication is possible. In the treatments with communication, trustors increase their IN choices by roughly 20 percentage points, that is, a large increase comparable in size to the increase in ROLL choices. This effect is again almost uniform across the different delay treatments and mirrors the pattern shown in Panel (a). Communication raises the rate of opt-in choices in all three delay conditions, but within both NOCOMM and COMM conditions there is no substantial difference in behavior due to different time delays.

Figure 2 also suggests that neither trustworthiness nor trust decline substantially across the different delay conditions. All twelve pairwise comparisons between delay conditions are statistically insignificant.

Table 2 reports the regression analysis estimates which confirm the visual impression. The omitted category is IMM. Communication raises the ROLL rate by 1.80 ROLL choices in the IMM condition ($p = 0.007$), by 1.91 in EARLY ($p = 0.089$), and by 1.77 in LATE ($p = 0.006$).[16] Again, delay has no discernible effect. There is no statistical difference between the IMM and EARLY conditions (NOCOMM: $p = 0.719$, COMM: $p = 0.859$), IMM and LATE ($p = 0.658$, $p = 0.565$), or EARLY and LATE ($p = 0.983$, $p = 0.854$).

The same pattern emerges from the second column, where we analyze trustors' IN decisions. Communication significantly raises trust in all three delay conditions (IMM: $p = 0.005$, EARLY: $p = 0.005$, LATE: $p = 0.019$), and with the exception of the dip in the NOCOMM-EARLY treatment,

$p = 0.507$, IN COMM $p = 0.911$, Kruskal-Wallis tests.

[16] We report Wald tests whenever we test linear combinations of coefficients.

Table 2: Behavior of Trustees and trustors

	ROLL	IN
EARLY	-0.281	-0.912*
	(0.774)	(0.527)
LATE	-0.299	-0.276
	(0.671)	(0.480)
COMM	1.800***	1.828***
	(0.638)	(0.618)
EARLY × COMM	0.105	0.997
	(1.282)	(1.135)
LATE × COMM	-0.035	-0.084
	(0.874)	(0.936)
Time Preference	0.106	-0.050
	(0.080)	(0.075)
Risk Preference	0.169	0.599***
	(0.163)	(0.100)
Age (years)	0.084**	0.093***
	(0.034)	(0.030)
Female	-0.480	0.290
	(0.489)	(0.428)
Constant	-0.005	-2.014
	(1.418)	(1.313)
adj. R-squared	0.054	0.151
N	353	353
Clusters	40	40

* p<0.10, ** p<0.05, *** p<0.01
OLS regressions, standard errors adjusted for clustering on the session level. Risk preference: number of risky choices in MPL-R; Time preference: number of later choices in MPL-T.

we do not find any substantial differences across delay conditions.[17]

In these regressions, we included two controls for time and risk preferences, as well as age and gender of the participant.[18] Older participants tend to be both more trusting and more trustworthy. Furthermore, more risk tolerant participants tend to be more trusting.

[17]IMM vs. EARLY: NOCOMM $p = 0.091$, COMM $p = 0.932$; IMM vs. LATE: $p = 0.569$, $p = 0.659$; EARLY vs. LATE: $p = 0.076$, $p = 0.699$.

[18]We presented subjects with two multiple-price lists (MPLs) at the end of the laboratory session, one measuring risk preference (MPL-R) and the other time preference (MPL-T). The MPL-R task asked for subjects' certainty equivalent for a 50-50 lottery between \$15 and \$0. The MPL-T task asked for their preference between a fixed \$10 early payment and a late payment that varied between \$10 and \$15.

3.1.3 Promises

Our previous analysis documents that ROLL rates differ across communication treatments, but do not differ across delay conditions. We now investigate how trustees in the communication conditions use their freeform messages to induce higher levels of cooperation. In order to do so, we classified the messages into whether they contained an explicit promise of choosing ROLL (e.g., "I will choose ROLL" or "I promise to choose ROLL") and into whether they did not (e.g., "a singer in a smoky room, a smell of wine and cheap perfume").[19] The rate of promises is not higher in the EARLY and LATE conditions as one could suspect if subjects were trying to prevent a potential decay of trust with the time delay: 66 percent of trustees in COMM-IMM make a promise, 75 in COMM-EARLY, and 69 in COMM-LATE.[20]

Recall that the COMM conditions have a 50 percent higher ROLL rate than the NOCOMM conditions (an average of 5.6 ROLL choices in COMM versus 3.7 in NOCOMM). We can now decompose this effect into those trustees who made an explicit promise versus those who did not. Those who made promises have an even higher ROLL rate of an average 6.2 ROLL choices. In contrast, those who did not make a promise have a lower ROLL rate of 4.2.[21]

Recall that subjects saw ten radio buttons for their choice, one of which was randomly implemented. Only 25 percent of trustees choose a mixture of the two choices while 75 percent choose either ten times ROLL or ten times DON'T ROLL. The option to make ten different choices introduces the possibility of moral wiggle room in the communication conditions. A trustee could potentially stick to the letter of her promise "I will choose ROLL" while only choosing ROLL for a part of her ten choices. However, we do not see an increased use of mixing ROLL and DON'T ROLL choices if trustees make a promise. On the contrary, if anything, pure-strategy use is higher if the trustee made a promise (77 percent) than if she did not make a promise (63 percent), consistent with subjects sticking to the spirit of the promise and not exploiting loopholes in the wording.[22]

[19] Two hypothesis-blind research assistants classified the messages independently. For details about the procedure, see Appendix C.

[20] The rate of promises is not systematically different across the delay conditions, $p = 0.538$, χ^2 test.

[21] If we regress number of ROLL choices on a communication condition dummy and an additional dummy that is one if and only if the trustee made a promise, adjusting for clustered standard errors on the session level, we obtain an insignificant coefficient for communication condition without promise ($p = 0.525$) and a significant additional effect of actually making a promise ($p = 0.043$).

[22] However, this difference is only marginally significant ($p = 0.074$, χ^2 test).

3.1.4 Message Recall

We did not remind trustees of the messages they sent before they made their ROLL choices. This created an opportunity for biased recollection or motivated forgetting of the message on behalf of the trustees. To test how well trustees recalled the content of the message, we asked them to re-write their message as accurately as possible directly *after* their ROLL decisions. We did not provide monetary incentives for accuracy. We then classified the remembered messages into explicit promises and other messages, as before.

In general, the overlap between original and remembered messages is high. Among recalled messages, 78.1 percent are classified in the same category as the original message. When the original message contained a promise, the concordance was even higher—85.7 percent of trustees remembered that they made a promise. Furthermore, accuracy does not decline over time, as messages from the LATE condition still have a category concordance of 78.0 percent with 80.0 percent of original messages containing a promise.

Among the messages that do not contain an explicit promise, many took the form of an appeal to efficiency, in the generic form of "the best thing we can do to maximize our earnings is to go IN and ROLL." We did not classify these appeals as explicit promises. However, it seems that trustees actually wanted to make a promise when they wrote such messages. 29.2 percent of trustees who originally wrote about striving for the efficient outcome later recalled a message that we classify as an explicit promise.

Finally, we obtain virtually the same ROLL rates when we split trustees into those who recalled making a promise and those who did not as we did when we split trustees according to their original message (see Section 3.1.3). Trustees who did not remember making a concrete promise had an average rate of 4.7 ROLL choices, while those who did remember doing so had a rate of 6.5 ROLL choices. In summary, we find little evidence for systematic distortions of memory, either intentional or unintentional.

3.2 Beliefs

In addition to subjects' behavior, we also elicited their beliefs. We made sure to elicit any beliefs only *after* subjects had made their choice to avoid the elicitation procedure to interfere with decision-

making.[23]

The first set of beliefs we elicited were first-order beliefs about how many times their matched trustee (trustor) would choose ROLL (IN). We call these responses FOB1. We did not provide monetary incentives for FOB1 to avoid hedging. In addition, we asked each trustor about the aggregate percentage of ROLL decisions and each trustee about the aggregate percentage of IN decisions in their session. We call these responses FOB2. We incentivized FOB2 with a bonus of $5 if the guess fell within 10 percentage points of the actual value.

We also elicited trustees' *second-order beliefs* about the matched trustor's first-order beliefs. We call these second-order beliefs SOB. At the end of Q2, we also elicited first-order beliefs about the *other* delay condition. That is, at the end of the study, we revealed the existence of the other delay conditions and asked trustees (trustors) about the aggregate percentage of IN (or ROLL) decisions in these other delay condition. We term these responses FOB-O.[24]

Our measures for first-order beliefs FOB1 and FOB2 are significantly higher in the communication treatments. Thus, our results suggest that communication increases trust, even over time horizons as long as three weeks. The only exceptions are second-order beliefs in the early delay condition where the positive effect of communication is marginally significantly weaker.

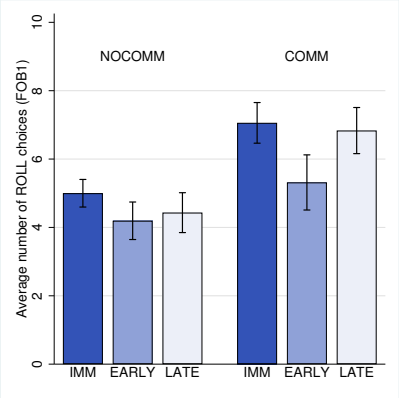
Indeed, Figures 3a, 3b, and 3c show that communication consistently raises first- and second-order beliefs about ROLLS. Although there is more variation in the positive effect of communication on first- and second-order beliefs than there is for ROLL rates, the magnitudes are quite similar. On average, communication raises beliefs by about 20 percentage points.

Table 3 again confirms the visual impression of a large and persistent effect of communication on trust and trustworthiness. Trustors have higher expectations about the ROLL rates of trustees and therefore behave more trustingly when communication is allowed, but there is no evidence that greater delay has any significant effect, except for the marginally significant dip in the FOB2 measure in the NOCOMM-EARLY condition that we previously documented for the IN choices.[25]

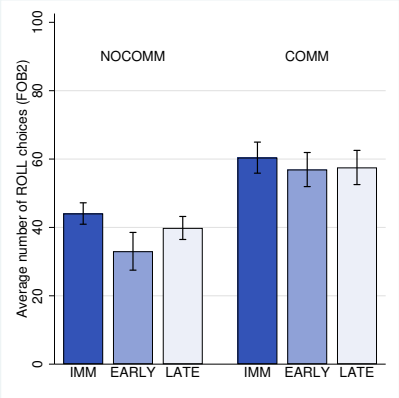
[23]We elicited trustees' beliefs directly after they made their decisions (i.e., in Q1 for IMM/EARLY and in Q2 for LATE). This has the advantage that their beliefs are elicited closely to their decision but also means that the questionnaire in which trustees' beliefs are elicited varies across delay conditions.

[24]Note that we did not elicit all beliefs in all sessions, resulting in varying sample sizes for these beliefs below.

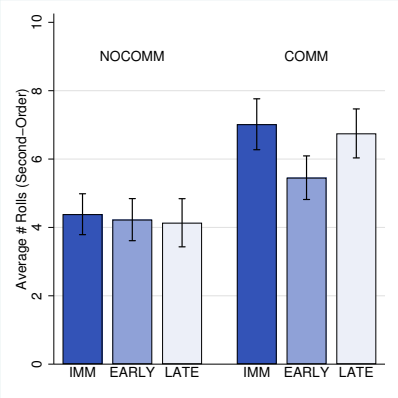
[25]There are a total of 42 pairwise comparisons of delay effects (IMM-EARLY, IMM-LATE, EARLY-LATE across 7 measures and 2 communication conditions). Of these, 18 are negative, indicating a decline of beliefs with increasing delay, and 24 are positive, indicating an increase with increasing delay. If we only consider statistically significant effects (two-tailed $p < 0.05$), 2 of the negative effects are significant and 4 of the positive effects are significant. Finally, no effect remains significant after applying a Holm-Bonferroni correction.



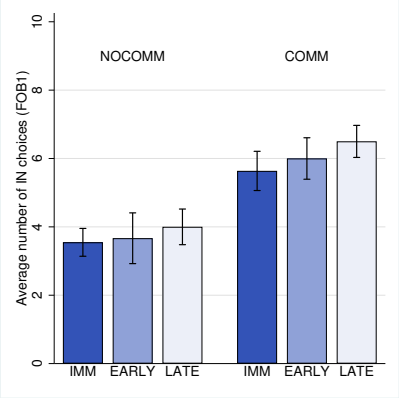
(a) FOB1:ROLL



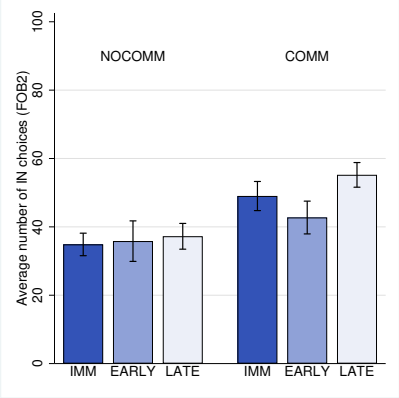
(b) FOB2: ROLL



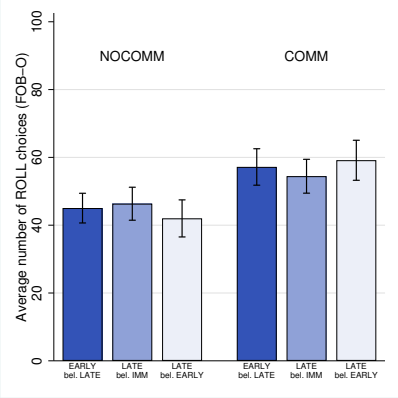
(c) SOB: ROLL



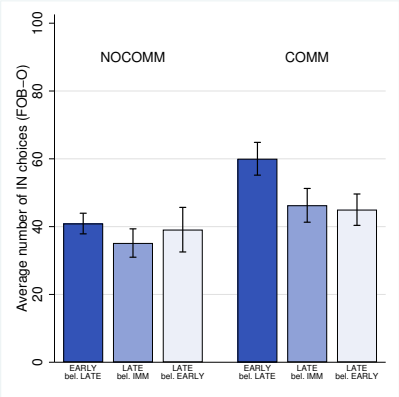
(d) FOB1: IN



(e) FOB2: IN



(f) FOB-O: ROLL



(g) FOB-O: IN

Figure 3: 1st-order beliefs about population (incentivized); 2nd-order beliefs about partner (Trustees only); 1st-order beliefs about other delay condition; error bars represent standard errors from regressions of belief on treatment dummies, adjusted for clustering on the session level.

Taken together, our data on both actual behavior and beliefs suggest a surprising persistence over time of trustworthiness in general and promises in particular. Even when trustees make their

Table 3: Beliefs of Trustees and trustors

	FOB1: ROLL	FOB2: ROLL	FOB1: IN	FOB2: IN	SOB: ROLL
EARLY	-0.882	-1.168**	-0.214	0.010	0.049
	(0.525)	(0.526)	(0.659)	(0.715)	(0.602)
LATE	-0.790	-0.469	-0.255	0.448	0.268
	(0.579)	(0.282)	(0.735)	(0.535)	(0.388)
COMM	1.959***	1.608***	2.606***	2.008***	1.382***
	(0.592)	(0.432)	(0.766)	(0.537)	(0.413)
EARLY × COMM	-0.668	0.864	-1.343	0.384	-0.671
	(1.014)	(0.726)	(0.834)	(0.981)	(0.772)
LATE × COMM	0.508	0.119	-0.053	0.356	0.335
	(0.913)	(0.569)	(0.987)	(0.742)	(0.488)
Time Preference	-0.038	-0.021	0.023	-0.009	0.019
	(0.075)	(0.062)	(0.098)	(0.074)	(0.059)
Risk Preference	0.208	0.302***	0.002	0.136	0.102
	(0.130)	(0.095)	(0.148)	(0.138)	(0.099)
Age (years)	0.071***	0.059***	0.015	0.046**	0.022
	(0.024)	(0.020)	(0.024)	(0.020)	(0.022)
Female	0.337	-0.157	-0.216	0.137	0.058
	(0.492)	(0.370)	(0.460)	(0.326)	(0.296)
Constant	2.046*	1.233	3.966**	1.504	2.076**
	(1.039)	(0.783)	(1.492)	(1.025)	(0.971)
adj. R-squared	0.101	0.134	0.097	0.122	0.054
N	249	353	186	261	354
Clusters	36	40	32	36	40

* p<0.10, ** p<0.05, *** p<0.01
OLS regressions, standard errors adjusted for clustering on the session level. Risk preference: number of risky choices in MPL-R; Time preference: number of later choices in MPL-T.

decisions outside of the lab and as long as three weeks after they communicated with trustors, they still behave in essentially the same way as when they decide immediately in the lab. In any time frame, they behave more generously toward trustors when communication (and thus promises) are possible. Trustors anticipate both the relevance of promises (increasing their trust greatly when promises are made) and the irrelevance of the time frame (exhibiting the same rate of trust across all time frames).

4 Conclusion

A vast literature has used the laboratory trust game as a vehicle to study economic transactions involving trust and trustworthiness (Johnson and Mislin, 2011). An important strand of this lit-

erature, beginning with Charness and Dufwenberg (2006), has documented the positive effect of pre-play communication on cooperation, trust, and trustworthiness in these experiments. However, the external validity of these studies remains limited because they force trustees to take their actions immediately the trustors. In contrast, most real-word instances of trust involve a significant delay before the trustee has to deliver. In fact, one of the primary roles of promises is to facilitate production and exchange *over time*.

Using a hybrid lab and online experiment we provided evidence for the persistent power of communication over time. In our trust experiment, trustees chose how much to return to trustors either (i) immediately in the lab, (ii) in a 24-hour window after they left the lab, or (iii) three weeks after they left the lab. Even when three weeks passed between the communication stage and actual choices, trustee behavior remained unchanged. Moreover, communication raised cooperation, trust, and trustworthiness by about 50% and this positive effect is as large as when choices immediately follow communication. Our results suggest that the findings of previous laboratory studies on trust in general, and promises in particular, extend to more externally valid scenarios in which subjects choose actions outside the lab and do so long after they made promises.

The surprising persistence of trustworthiness, and the lasting effect of communication naturally raises a number of additional questions about the role of communication. For example, at what point does the persistence of communication diminish? What institutional arrangements could enhance potentially decaying trust over time? We leave these and other interesting questions for future research.

References

Alós-Ferrer, Carlos and Fritz Strack, "From dual processes to multiple selves: Implications for economic behavior," *Journal of Economic Psychology*, 2014, *41*, 1–11.

Andersen, Steffen, Uri Gneezy, Agne Kajackaite, and Julie Marx, "Allowing for reflection time does not change behavior in dictator and cheating games," *Journal of Economic Behavior & Organization*, 2018, *145*, 24–33.

Andreoni, James and Marta Serra-Garcia, "Pledges as Screening Devices," Technical Report, Unpublished Working Paper 2018.

Arrow, Kenneth J., *The Limits of Organization* Fels Center of Government Series, Norton, 1974.

Bartels, Daniel M and Lance J Rips, "Psychological connectedness and intertemporal choice," *Journal of Experimental Psychology: General*, 2010, *139* (1), 49–69.

Bartolomeo, Giovanni Di, Martin Dufwenberg, Stefano Papa, and Francesco Passarelli, "Promises, Expectations, and Causation," *Sapienza University of Rome Working Paper*, 2017.

Ben-Ner, Avner and Louis Putterman, "Trust, communication and contracts: An experiment," *Journal of Economic Behavior and Organization*, 2009, *70* (1), 106 – 121.

Berg, Joyce, John Dickhaut, and Kevin McCabe, "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, 1995, *10* (1), 122–142.

Bhattacharya, Puja and Arjun Sengupta, "Promises and Guilt," *OSU Working Paper*, 2017.

_ , Kirby Nielsen, and Arjun Sengupta, "Timing of communication," *SSRN Working Paper 2940770*, 2019.

Bicchieri, C. and A. Lev-On, "Computer-Mediated Communication and Cooperation in Social Dilemmas: An Experimental Analysis," *Politics, Philosophy, and Economics*, 2007, *6* (2), 139–168.

Bouwmeester, S., P. P. J. L. Verkoeijen, B. Aczel, F. Barbosa, L. Bègue, P. Brañas-Garza, T. G. H. Chmura, G. Cornelissen, F. S. Døssing, A. M. Espín, A. M. Evans,

F. Ferreira-Santos, S. Fiedler, J. Flegr, M. Ghaffari, A. Glöckner, T. Goeschl, L. Guo, O. P. Hauser, R. Hernan-Gonzalez, A. Herrero, Z. Horne, P. Houdek, M. Johannesson, L. Koppel, P. Kujal, T. Laine, J. Lohse, E. C. Martins, C. Mauro, D. Mischkowski, S. Mukherjee, K. O. R. Myrseth, D. Navarro-Martínez, T. M. S. Neal, J. Novakova, R. Pagà, T. O. Paiva, B. Palfi, M. Piovesan, R.-M. Rahal, E. Salomon, N. Srinivasan, A. Srivastava, B. Szaszi, A. Szollosi, K. Ø. Thor, G. Tinghög, J. S. Trueblood, J. J. Van Bavel, A. E. van 't Veer, D. Västfjäll, M. Warner, E. Wengström, J. Wills, and C. E. Wollbrant, "Registered Replication Report: Rand, Greene, and Nowak (2012)," *Perspectives on Psychological Science*, 2017, *12* (3), 527–542. PMID: 28475467.

Brandts, Jordi, David J Cooper, and Christina Rott, "Communication in laboratory experiments," in "Handbook of Research Methods and Applications in Experimental Economics," Edward Elgar Publishing, 2019.

Braude, Jacob Morton, *Proverbs, epigrams, aphorisms, sayings, and bon mots*, Prentice Hall, 1965.

Bénabou, Roland and Jean Tirole, "Willpower and Personal Rules," *Journal of Political Economy*, 2004, *112* (4), 848–886.

Carlson, Ryan W, Michel Marechal, Bastiaan Oud, Ernst Fehr, and Molly J Crockett, "Motivated misremembering: Selfish decisions are more generous in hindsight," *PsyArXiv Preprints*, 2018.

Casella, Alessandra, Navin Kartik, Luis Sanchez, and Sébastien Turban, "Communication in context: Interpreting promises in an experiment on competition and trust," *Proceedings of the National Academy of Sciences*, 2018.

Charness, Gary and Martin Dufwenberg, "Promises and partnership," *Econometrica*, 2006, *74* (6), 1579–1601.

_ **and** _ , "Participation," *American Economic Review*, June 2011, *101* (4), 1211–1237.

Cone, Jeremy and David G Rand, "Time pressure increases cooperation in competitively framed social dilemmas," *PLoS one*, 2014, *9* (12), e115756.

Dreber, Anna, Drew Fudenberg, David K. Levine, and David Rand, "Self-Control, Social Preferences and the Effect of Delayed Payments," *SSRN Working Paper*, 2016.

Ederer, Florian and Alexander Stremitzer, "Promises and expectations," *Games and Economic Behavior*, 2017, *106*, 161–178.

Ellingsen, Tore and Magnus Johannesson, "Promises, Threats and Fairness," *Economic Journal*, 2004, *114* (495), 397–420.

Fischbacher, Urs, "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 2007, *10*, 171–178.

Fosgaard, Toke and Adriaan R Soetevent, "Promises undone: How committed pledges impact donations to charity," Technical Report 2018/044-VII, Tinbergen Institute 2018.

Frederick, Shane, "Time preference and personal identity," *Time and decision*, 2003, pp. 89–113.

_ , George Loewenstein, and Ted O'Donoghue, "Time discounting and time preference: A critical review," *Journal of economic literature*, 2002, *40* (2), 351–401.

Fudenberg, Drew and David K. Levine, "A Dual-Self Model of Impulse Control," *American Economic Review*, 2006, *96* (5), 1449–1476.

Gelder, Jean-Louis Van, Hal E Hershfield, and Loran F Nordgren, "Vividness of the future self predicts delinquency," *Psychological science*, 2013, *24* (6), 974–980.

Hourani, Fayek S, *Daily Bread for Your Mind and Soul: A Handbook of Transcultural Proverbs and Sayings*, Xlibris, 2012.

Imas, Alex, Michael Kuhn, and Vera Mironova, "Waiting to Choose," *SSRN Working Paper*, November 2016.

Ismayilov, H. and J. Potters, "Why do promises affect trustworthiness, or do they?," *Experimental Economics*, 2016, *19*, 382–393.

Jamison, Julian and Jon Wegener, "Multiple selves in intertemporal choice," *Journal of Economic Psychology*, 2010, *31* (5), 832–839.

Johnson, Noel and Alexandra A. Mislin, "Trust games: A meta-analysis," *Journal of Economic Psychology*, 2011, *32* (5), 865–889.

Kerr, N. L. and C. M. Kaufman-Gilliland, "Communication, Commitment and Cooperation in Social Dilemma," *Journal of Personality and Social Psychology*, 1994, *66* (3), 513–529.

Kouchaki, Maryam and Francesca Gino, "Memories of unethical actions become obfuscated over time," *Proceedings of the National Academy of Sciences*, 2016, *113* (22), 6166–6171.

Kovarik, Jaromir, "Giving it now or later: Altruism and discounting," *Economics Letters*, 2009, *102* (3), 152–154.

Li, King King, "Asymmetric memory recall of positive and negative events in social interactions," *Experimental Economics*, 2013, *16* (3), 248–262.

Meyer, Christian and Egon Tripodi, "Image Concerns in Pledges to Give Blood: Evidence from a Field Experiment," *Available at SSRN 3132289*, 2018.

Neo, Wei Siong, Michael Yu, Roberto A Weber, and Cleotilde Gonzalez, "The effects of time delay in reciprocity games," *Journal of Economic Psychology*, 2013, *34*, 20–35.

Ostrom, E., J. Walker, and R. Gardner, "Covenants With and Without a Sword: Self-Governance Is Possible," *American Political Science Review*, 1992, *86* (2), 404–417.

Paine, Lynn Sharp, *Value Shift: Why Companies Must Merge Social and Financial Imperatives to Achieve Superior Performance*, New York: McGraw-Hill, 2004.

Parfit, Derek, "Later Selves and Moral Principles," in Alan Montefiore, ed., *Philosophy and Personal Relations*, London: Routledge & Kegan Paul, 1973, pp. 149–160.

Rand, David G, Alexander Peysakhovich, Gordon T Kraft-Todd, George E Newman, Owen Wurzbacher, Martin A Nowak, and Joshua D Greene, "Social heuristics shape intuitive cooperation," *Nature communications*, 2014, *5*, 3677.

Sally, David, "Conversation and Cooperation in Social Dilemmas: A Meta-Analysis of Experiments from 1958 to 1992," *Rationality and Society*, 1995, *7* (1), 58–92.

Saucet, Charlotte and Marie Claire Villeval, "Motivated memory in dictator games," *Games and Economic Behavior*, 2019.

Serra-Garcia, Marta, Eric van Damme, and Jan Potters, "Lying About What You Know or About What You Do?," *Journal of the European Economic Association*, 2013, *11* (5), 1204–1229.

Vanberg, Christoph, "Why do people keep their promises? An experimental test of two explanations," *Econometrica*, 2008, *76* (6), 1467–1480.

Weymar, Paul, *Konrad Adenauer: Die autorisierte Biographie*, Kindler, 1955.

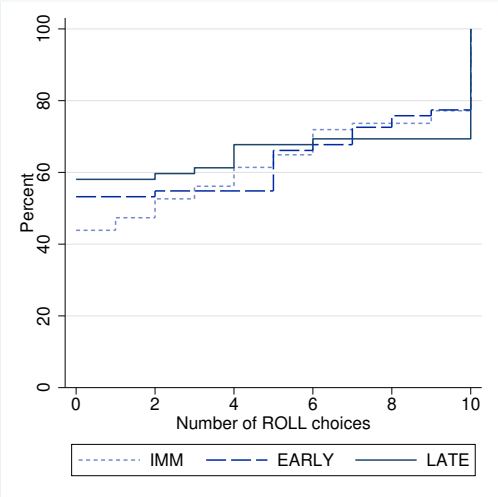
A Structure of Lab Session and Online Questionnaires

Table 4: Order of tasks and questionnaires across conditions. Q1 was available online for 24 hours after the laboratory session in EARLY and LATE conditions, and was administered in the laboratory immediately following the "Laboratory" part in the IMM condition.

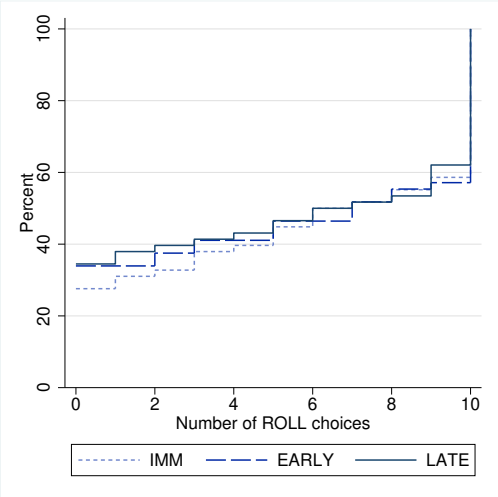
Part	IMM/EARLY	All conditions	LATE
Laboratory		Instructions Comprehension Questions [COMM: Message from Trustee to trustor] Comprehension Questions *Decision trustor* Multiple Price List - Time Multiple Price List - Risk Questionnaire: socioeconomic data Questionnaire: attitudes (time, risk, trust) Questionnaire: Machiavellianism	
Q1	*Decision Trustee** [COMM: Message recall*]	First-Order Belief 2† First-Order Belief 1† Second-Order Belief† Free-form explanation of IN/ROLL decision† Free-form explanation of MPL-Time decision Free-form explanation of MPL-Risk decision	
Q2		First-Order Belief Other Free-form explanation of FOB Other	*Decision Trustee** [COMM: Message recall*] First-Order Belief 2* First-Order Belief 1* Second-Order Belief* Free-form explanation of ROLL decision*

* Only trustees † Except trustees in LATE

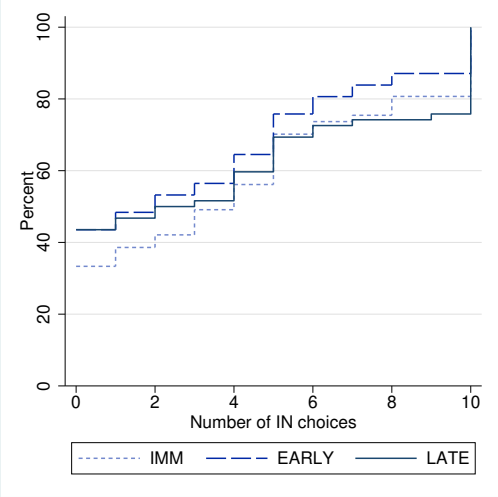
B Cumulative Distributions of ROLL and IN Choices



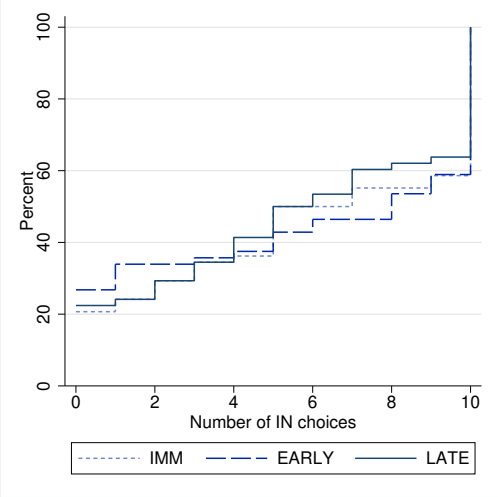
(a) Trustees, NOCOMM



(b) Trustees: COMM



(c) Trustors: NOCOMM



(d) Trustors: COMM

Figure 4: Cumulative choice distributions across conditions.

C Message Classification

The main classification, which we fixed ex ante, contained the categories "Promise," "Other Content," and "No Message."[26] We then gave the experimental instructions and the messages to two research

[26] After sighting the messages, we decided to break down the "Other Content" category into "Appeal to Efficiency," "Appeal to Fairness," and "Other Content." We did this because we realized that some participants seemed to use

assistants who were blind to our hypotheses. The assistants then coded the messages independently according to the categories described above and received a monetary incentive for indicating the same categories as the other coder.[27]

We allowed three levels for the coding: 0 (category not present in message); 1 (category somewhat present in message); 3 (category strongly present in message). The correlation between the two research assistants was high (for the Promise category, Pearson correlation is 0.87 and Cronbach's $\alpha = 0.93$). We classify a message as a promise if both assistants coded the message as such.[28]

other, non-promise, types of arguments to convince the trustor to choose IN. The same message could be classified into more than one category (e.g., both as Promise and as an Appeal to Efficiency). Only 13 messages (8 percent) contain an appeal to fairness but 84 messages (49 percent) contain an appeal to efficiency. The latter is highly correlated with the presence of a promise (Pearson correlation 0.75).

[27]This is a standard procedure for message coding, see Brandts et al. (2019).

[28]That is, if both coders scored the message at least with one. The results remain qualitatively unchanged if we apply more stringent definitions of a promise.

D Supplementary Material: Subject Instructions

Instructions

Thank you for participating in today's study.

I will read through a script to explain to you the nature of today's experiment as well as how to navigate the computer interface with which you will be working. I will use this script to make sure that the information given in all sessions of this study is the same.

In addition to a $10 payment that you receive for your participation, you will be paid an amount of money that you accumulate from the decision task that will be described to you in a moment. The total amount you receive will be determined during the experiment and will depend on your decisions as well as the decisions of others.

If you have any questions during the experiment, please raise your hand and wait for an experimenter to come to you. Please do not talk, exclaim, or try to communicate with other participants during the experiment. Participants intentionally violating the rules may be asked to leave the experiment with only their participation payment.

Timeline of the experiment

As already mentioned in the invitation e-mail, this study requires you to answer **two** questionnaires, besides participating in today's laboratory session. Each questionnaire will take no more than 10 minutes to complete. The link for the second questionnaire will be sent to you via e-mail.

- The first questionnaire must be completed at the end of this session. We will come to your computer station and switch the program to an internet browser window for this.
- The second questionnaire must be completed **within a 24-hour window, starting 3 weeks from today** (on Friday, October 27, 12pm – Saturday, October 28, 12pm).

Important: You will only be paid for this study if you complete this laboratory session and **both questionnaires**. *At this point, please let the experimenter know if you will be unable to complete the online questionnaire three weeks from now.*

Payment

The experiment comprises 10 decision tasks. In each of these decisions, you have the opportunity to earn money. At the end of the experiment, we will pay you the amount of money produced by **one randomly selected decision**. Each of the 10 decisions is equally likely to be chosen for payment.

Moreover, you will have the possibility to earn an additional amount of money in other tasks. Your final payment at the end of the experiment will consist of the money obtained in one randomly chosen decision, the amount obtained from the additional tasks, and the participation payment.

Your participation payment ($10) will be paid out in cash at the end this session. Your remaining payment will be paid to you as an Amazon voucher three weeks from now. It will be sent to your email address after the conclusion of the second questionnaire. Depending on your decision in an additional task, you may receive a part of your payment sooner, also as an Amazon voucher. If you receive money from this task, the voucher will be sent to you by tomorrow. No participant will be told the payoffs of other participants.

Decision Task

At the beginning of the experiment you will be randomly matched with another participant. This will be the only participant with whom you will interact. All interactions will be anonymous and will take place through a computer. This means that you will never know the identity of the other participant with whom you are matched, and this participant will never know your identity.

In each pair, one participant will have the role of Person A and the other will have the role of Person B. Roles will be assigned randomly, by the computer. The amount of money you will earn depends on the decisions made in your pair.

At the beginning of the decision task, each Person A will indicate whether he or she wishes to choose *IN* or *OUT*. If A chooses *OUT*, the task ends and both, A and B, receive $15.

Next, each Person B will indicate whether he or she wants to choose *ROLL* or *DON`T ROLL*. Note that when Person B makes her decision, she will not know about A's previous decision. However, the decision of Person B will only make a difference when A has actually chosen *IN*. Therefore, for the purpose of making this decision, Person B should presume that the paired Person A has chosen *IN*, since this is the only case in which Person B's decision will matter.

Assuming that A has chosen *IN*, then:

- If B chooses *DON'T ROLL*, A's payoff is $0 and B's payoff is $42.
- If B chooses *ROLL*, then the payoffs for the two people in the pair are determined by the (computerized) roll of a six-sided die. If the die roll is 2, 3, 4, 5 or 6 then A's payoff is $36 and B's payoff is $30. If the die roll is 1, then A's payoff is $0 and B's payoff is $30. This information is summarized in the following chart:

Payoffs

	Payoff of A	Payoff of B
A chooses *OUT*	$15	$15
A chooses *IN* and B chooses *DON'T ROLL*	$0	$42
A chooses *IN* and B chooses *ROLL*, die = 1	$0	$30
A chooses *IN* and B chooses *ROLL*, die =2, 3, 4, 5, 6	$36	$30

You will make 10 decisions in the task just described. In each decision you will be interacting with the same other person. Specifically, each Person A in a pair will make the decision whether to choose *OUT* or *IN* ten times and each Person B in a pair will make the decision whether to *ROLL* or *DON'T ROLL* ten times. Remember that **only one** of these 10 decisions in a pair will be selected to count for payment.

Person B's decision will be made at the end of this session. B will choose whether to *ROLL* or *DON'T ROLL* when answering the first questionnaire.

The experiment consists of 5 steps that are described below in more detail.

Step 1: Enter your email address

Before the decision task starts, you will be prompted on the screen to provide your email address. After the conclusion of the entire experiment, that is, after you complete both questionnaires in three weeks, the final payoff from the above decision task, as well as additional tasks, will be emailed to this address as an Amazon voucher.

To receive your payments, it is very important that you **provide the correct email address**. Your email address is treated confidentially and is only used to send the questionnaire link and to transfer your payments.

Step 2: Role assignment

At the beginning of the decision task, you will be randomly paired with another participant by the computer. However, no participant will ever know the identity of the person with whom he or she is matched. In each pair, one person will then be randomly assigned the role of Person A and the other the role of Person B. It is equally likely that you will be assigned to either role. Your role will be displayed on the computer screen. You will remain in the same role throughout the experiment.

Step 3: Communication Phase

After role assignment, but before participants make their 10 decisions in the decision task, Person B will have an option to send a written message to Person A, if desired. This message can be seen by A before she will have to make a decision.

Important: You are not allowed to reveal your identity. (That is, you are not allowed to reveal your name, or any other identifying feature such as gender, appearance etc.) The experimenter will monitor the messages. Violations of these rules will result in the exclusion from the experiment and all payments for B. In this case, the paired Person A will then get the participation fee as well as the payoff that would arise if he or she chose OUT, i.e. $15.

Also, please refrain from using profane or offensive language.

Other than these restrictions, person B is free to send any message.

Step 4: Decision by Person A

Each Person A will indicate whether to choose *IN* or *OUT*. Since there are ten decisions, Person A will decide for each decision separately whether to choose *IN* or *OUT*. For the decisions in which A chooses *OUT*, Person B's decisions are irrelevant and both receive payoffs of $15. For the decisions in which A chooses *IN*, Person B's decisions of whether to ROLL or DON'T ROLL, possibly along with the outcome of a die roll, will determine payoffs.

Step 5: Decision by Person B

This step will be conducted at the end this session, while you are completing the first questionnaire. When completing the questionnaire, each Person B will decide whether to choose *ROLL* or *DON'T ROLL* for each of the ten decisions. These choices will determine the payoffs for those decisions in which Person A chose *IN*.

As mentioned above, B will not be told about A's decisions before submitting his or her decisions. For the decisions in which B chooses *DON'T ROLL*, the task ends at that point. For the decisions in which B chooses *ROLL*, a separate computerized die roll determines the payoffs for A and B for each decision.

For each pair, one of the 10 decisions will be randomly selected to count. For this decision, Person A's choice of *IN* or *OUT*, and Person B's choice of *ROLL* or *DON'T ROLL*, which is made during the first questionnaire, at the end of this session, will determine payments.

Remember that you will only be paid in three weeks, after you have completed **the online questionnaire**. You will not be affected if you complete the study but the person with whom you are matched fails to complete both questionnaires.

Additional tasks:

In addition to the above task, you will be presented with some additional tasks, at the end of this session and in the later questionnaire. Instructions for these tasks will be displayed on your screen before you begin the respective task. For some of these tasks, you may accumulate additional money. However, the choices that you make in these additional tasks will never affect the outcome of the 10 decisions described above.

Please raise your hand if you have any questions.

We will now proceed to a few comprehension questions to ensure that the instructions are clear to everybody.