Trust and Promises over Time

By FLORIAN EDERER AND FRÉDÉRIC SCHNEIDER*

Using a large-scale hybrid laboratory and online trust experiment with and without pre-play communication, we investigate how the passage of time affects trust. Communication (predominantly through promises) raises cooperation, trust, and trustworthiness by about 50 percent. This result holds even when three weeks pass between the time of the trustee's message/the trustor's decision to trust and the time of the trustee's contribution choice and even when this contribution choice is made outside of the lab. Delay between the beginning of the interaction and the time to reciprocate neither substantially alters trust or trustworthiness nor affects how subjects communicate.
JEL: C91, C72, D83
Keywords: trust, promises, persistence, trustworthiness, delay

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 reputation 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 belief in *future* actions.[1]

Casual observation suggests that, in a variety of settings, the more the promise and the corresponding act of trust recede into the past, the likelier it is for promisors to neglect their 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)

* Ederer: Yale School of Management, Yale University, 165 Whitney Avenue, New Haven, CT 06511, U.S.A., florian.ederer@yale.edu. Schneider: Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, United Kingdom, f.schneider@jbs.cam.ac.uk. 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, Kirby Nielsen, David Rand, Fiona Scott Morton, Andrew Schotter, Arjun Sengupta, Kelly Shue, Ted Snyder, Alexander Stremitzer, Roberto Weber, Kevin Williams, Jacqueline Würth, and seminar audiences at 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, Lucas Molleman for lioness programming support, and Decory Edwards, Andreas Klümper, and Maria Kogelnik for outstanding research assistance. The research was approved by the Yale University IRB and University of Zurich IRB. This article was previously circulated under the title "The Persistent Power of Promises."

[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.

argues that "most companies are simply not designed to remember and keep promises over time." On the other hand, contrary statements suggest that promises do (or at least should) have a lasting effect that is not diminished by the passage of time. For example, Alexander Hamilton argued that "[a] promise must never be broken" (Shea, 1879) and Immanuel Kant considered the duty to keep promises a categorical imperative (Byrd and Hruschka, 2006), that is an absolute, unconditional requirement that must be obeyed in all circumstances.[3]

This paper provides the first experimental investigation of the strength of trust, trustworthiness, and promises over time. We document that communication (primarily through the use of promises) 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 and without pre-play communication.[4] The main innovation of our paper is the introduction of a delay between the trustor's choice to trust and the trustee's decision 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 laboratory session*. These treatments allow 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 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 their promises.

Trust has been extensively studied in laboratory experiments, mostly using trust games beginning with the seminal work of Berg, Dickhaut and McCabe (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 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

[3] Echoing this view, the fictional character Frank Underwood in the TV series House of Cards pronounces that "[t]he nature of promises is that they remain immune to changing circumstances."

[4] 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).

trustworthiness. Better still, several recent studies offer experimental evidence that preplay 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, van Damme and Potters, 2013; Ismayilov and Potters, 2016; Ederer and Stremitzer, 2017; Di Bartolomeo et al., 2019; Bhattacharya and Sengupta, 2017; Casella et al., 2018).[5] 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.[6] In previous lab experiments using trust games (with and without communication), hardly any time passes during the period after the trustor's decision to trust or not (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 relative to everyday interactions.

First, people may be less likely to reciprocate another person's trust because they may feel less guilty as time passes. 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 distant 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 models an individual as 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 to reciprocate trust, most notably if induced by promises made in the past, are stronger over the short term than over the long term.[7]

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

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

[6]An extensive literature in economics and psychology has studied the delay of *outcomes*. For an introduction and comprehensive survey, see, for example, Frederick, Loewenstein and O'Donoghue (2002).

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

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

Several 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, Nielsen and Sengupta (2020) study whether the timing of communication (before or after actions are chosen) affects cooperation. They focus on a short time frame of four minutes in a lab setting and 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, Kuhn and Mironova (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.[9]

The remainder of the paper proceeds as follows. Section I describes our experimental design. Section II presents our experimental results. Section III offers concluding remarks.

I. 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 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.[10]

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

[9] Another sequence of papers study settings that feature implicit promises over time, but none of them vary or measure the effect of the delay. For example, Meyer and Tripodi (2018) investigate pledges to donate blood at some later point and find that almost all subjects renege on their pledge, but the authors 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 to donate right now, during the phone call, or at an undefined later time. Subjects who opted to pledge 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.

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

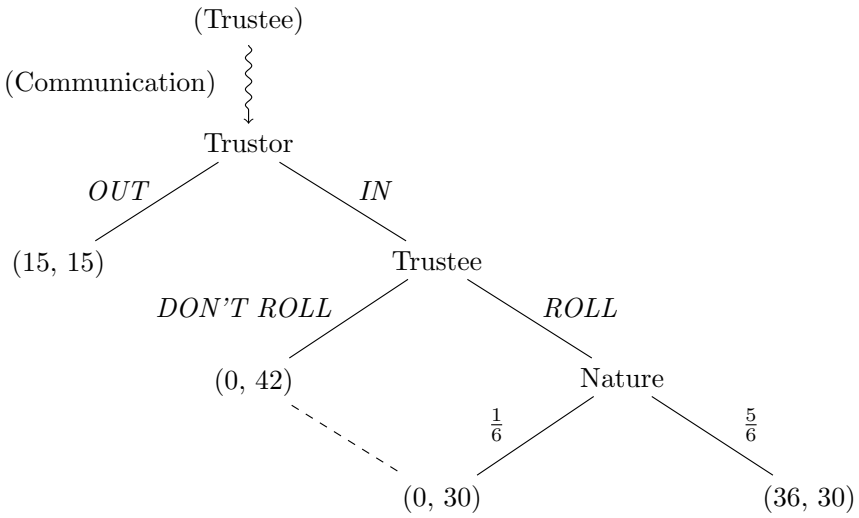

Figure 1. : Game Tree

Note: 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.[11] 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 modify the design of the Charness and Dufwenberg (2006) trust game such that subjects are presented with a menu of 10 IN/OUT (ROLL/DON'T ROLL) radio buttons as shown in Figure A1 of the online appendix. After the subjects have made their decisions, the computer selects one of the choices at random to count as the final decision. This modification allows 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.

[11] 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 be reminded of their message.

A. 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 had to make their decisions in Q1 while others had to decide later, in 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. Table 1 provides a detailed timeline of tasks and questions in the lab part and the questionnaire (Q1 and Q2) part of the experiment.[12]

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 all participants.[13]

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. 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 the six different conditions described in Table 2.

B. Procedures

We conducted 40 experimental sessions with a total of 707 student subjects at the University of Zurich and at Yale University.[14] Subjects were assigned to visually partitioned

[12] 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 to elicit trustees' ROLL choices *before* their beliefs.

[13] Trustors' IN/OUT decisions were always elicited in the initial lab session.

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

Part	All conditions	
	IMM/EARLY	LATE
Laboratory	Instructions Comprehension Questions [COMM: Message from Trustee to Trustor] *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*]	
	Belief elicitation: 1BG† Belief elicitation: 1BM† Belief elicitation: 2BM† Free-form explanation of IN/ROLL decision† Free-form explanation of MPL-Time decision Free-form explanation of MPL-Risk decision	
Q2		*Decision Trustee** [COMM: Message recall*] Belief elicitation: 1BG* Belief elicitation: 1BM* Belief elicitation: 2BM*
	Belief elicitation: 1BO Free-form explanation of 1BO	
		Free-form explanation of ROLL decision*

Table 1—: Order of tasks and questionnaires across conditions.

Note: 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.
Note: *Only trustees.
Note: †Except trustees in LATE.

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.[15] The laboratory part of the experiment was programmed and conducted with z-Tree (Fischbacher, 2007) and the online part with Qualtrics.

[15]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.

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

Table 2—: Treatment conditions, with number of subjects in parentheses.

II. Results

A. *Subject Behavior*

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 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.[16] These low attrition rates are due to our experimental design, which provided strong incentives to complete the additional surveys.

CHOICES. — We first study the behavior of trustees in our six treatment conditions. Recall that trustees made 10 decisions (i.e., 10 binary ROLL/DON'T ROLL radio buttons), 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 for 1 to 9 out of the 10 decisions.[17]

Panel (a) of Figure 2 shows the average number of times a trustee chooses ROLL. When there is no communication, trustees choose to roll on average between 35 percent to 38 percent of the time while this number increases to around 54 to 57 percent of the time when communication is possible. As expected, with immediate choices, communication substantially increases cooperation by trustees by about 50 percent, from 3.77 ROLL choices to 5.71 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

[16] 18 trustees did not complete Q2: 9 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 (less attrition in the LATE condition): 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.

[17] 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 Figure B1 of the online appendix. We tested whether the distributions differ across the three delay conditions, separately for trustors and trustees and by communication condition. None of the tests are significant: ROLL NOCOMM $p = 0.778$, ROLL COMM $p = 0.885$, IN NOCOMM $p = 0.507$, IN COMM $p = 0.911$, Kruskal-Wallis tests.

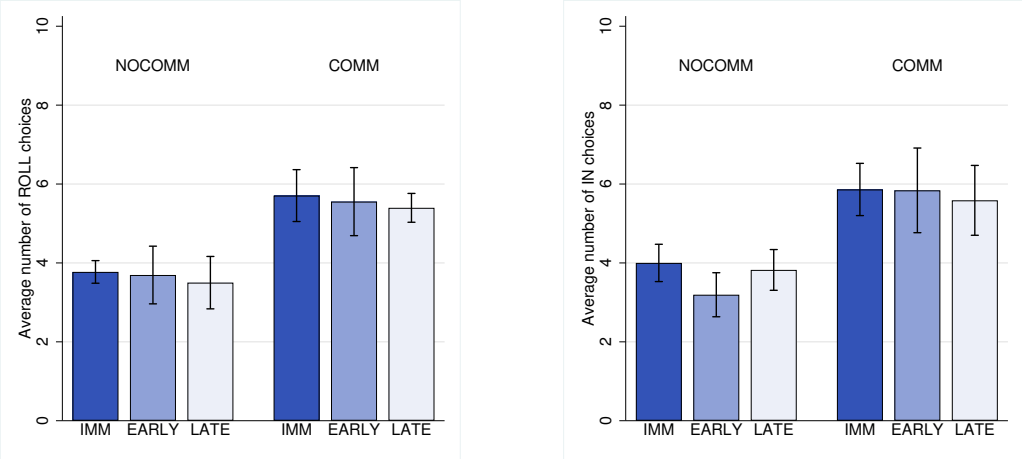

(a) Trustees: ROLL choices

(b) Trustors: IN choices

Figure 2. : (a) ROLL decisions of trustees; (b) IN decisions of trustors.

Note: Error bars indicate standard errors and are adjusted for clustering on the session level.

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 3 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 ($p = 0.006$).[18] 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:

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

	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

Table 3—: Behavior of Trustees and Trustors

Note: OLS regressions with 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.

$p = 0.005$, EARLY: $p = 0.005$, LATE: $p = 0.019$), and with the exception of the dip in the NOCOMM-EARLY treatment, we do not find any substantial differences across delay conditions.[19]

In these regressions, we included two controls for time and risk preferences, as well as age and gender of the participant.[20] We document that older participants in our experiment are both more trusting and more trustworthy. Furthermore, more risk-tolerant participants are also more trusting.

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

[20]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.

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, messages were classified according to whether they contained an explicit promise of choosing ROLL (e.g., "I will choose ROLL" or "I promise to choose ROLL") or whether they did not (e.g., "a singer in a smoky room, a smell of wine and cheap perfume").[21] The rate of promises is not higher in the EARLY and LATE conditions as one might expect 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. The rate of promises is not systematically different across the delay conditions ($p = 0.538$, χ^2 test). Furthermore, using the most recent Linguistic Inquiry Word Count (LIWC2015) model (Pennebaker, Booth and Francis, 2015), a widely used linguistic algorithm designed to detect deception in speech or writing, we do not detect any differences in the authenticity scores (Newman et al., 2003) of messages across the delay conditions.[22]

Recall that the COMM conditions have a 50 percent higher ROLL rate than the NO-COMM conditions (an average of 5.6 ROLL choices in COMM versus 3.7 in NO-COMM). 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 6.2 ROLL choices on average. In contrast, those who did not make a promise have a lower ROLL rate of 4.2.[23]

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 some of her ten choices. However, we do not see an increased use of combined ROLL and DON'T ROLL choices if trustees make a promise. On the contrary, if anything, the exclusive use of either one of the two choices 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.[24]

[21] Two hypothesis-blind research assistants classified the messages independently. See the online appendix for details about the procedure.

[22] We also explored whether the length of the messages differs between the delay conditions. First, 13 out of 172 COMM trustees chose to send an empty message (5 in COMM-IMM, and 4 each in COMM-EARLY and COMM-LATE, $p = 0.931$, χ^2 test). Second, the average number of words written was significantly greater ($p = 0.043$, Kruskal-Wallis test) in EARLY (36 words) than in LATE (28 words) and IMM (25 words).

[23] 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 ($p = 0.043$).

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

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 the part of the trustees and potentially contributed to the malleability of messages (Turmunkh, van den Assem and Van Dolder, 2019). 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. The remembered messages were then classified into explicit promises and other messages, using the same procedure described above.[25]

In general, the overlap between original and remembered messages is high. Among recalled messages, 92.3 percent are classified in the same category as the original message. Recall accuracy was similar for those who originally made a promise (93.4 percent) and those who did not (89.8 percent). Furthermore, accuracy does not decline over time, as messages from the LATE condition still have a category concordance of 90.2 percent with 88.5 percent of original messages containing a 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 II.A). Trustees who did not remember making a concrete promise had an average rate of 3.9 ROLL choices, while those who did remember doing so had a rate of 6.4 ROLL choices. In summary, we find little evidence for systematic distortions of memory, either intentional or unintentional.

B. 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 in order to avoid having the elicitation procedure interfere with decision-making.[26]

The first set of beliefs we elicited were first-order beliefs. We asked each trustor how many times his matched trustee would choose ROLL and we asked each trustee how many times her matched trustor would choose IN. We call these responses 1BM(ROLL) and 1BM(IN), respectively. To avoid hedging, we did not provide monetary incentives for 1BM. In addition, we asked each trustor about the aggregate percentage of ROLL decisions of the group of trustees and each trustee about the aggregate percentage of IN decisions of the group of trustors in their session. We call these beliefs about group behavior 1BG(ROLL) and 1BG(IN), respectively. We incentivized 1BG 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 2BM(ROLL). 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 trustors/trustees

[25] Due to a mistake with the message recall question, the recall data are missing for one session.

[26] 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 in time to their decision, but it also means that the questionnaire in which trustees' beliefs are elicited varies across delay conditions.

about the aggregate percentage of ROLL/IN decisions in these other delay conditions. We term these responses 1BO(ROLL) and 1BO(IN), respectively.[27]

Consistent with the effect of communication on behavior, the communication condition raises all of our measures of first- and second-order beliefs by about 20 percentage points, on average (Figure 3). Thus, our results suggest that communication increases trust, even over time horizons as long as three weeks. Although there is more variation in the positive effect of communication on first- and second-order beliefs than there is for ROLL and IN rates, the magnitudes are similar. All of the 21 pairwise comparisons of belief measures (NOCOMM vs COMM across 3 delay conditions for 7 belief measures) show a positive effect of communication and 16 of them are significant (two-tailed t-test, $p < 0.05$). Furthermore, 11 of these pairwise comparisons remain significant after applying a Holm-Bonferroni correction for the use of multiple belief measures.[28]

Table 4 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. As before, there is no evidence that greater delay has any significant effect, except for the statistically significant dip in the 1BG measure in the NOCOMM-EARLY condition that we previously documented for the IN choices.

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

III. 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 literature, beginning with Charness and Dufwenberg (2006), has documented the positive effect of pre-play communication (most notably promises) 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 after communicating with the trustors. In contrast, most real-word instances of trust involve a significant delay before the trustee has to decide whether or not to reciprocate the trust placed in her. In fact, one of the primary roles of promises is to facilitate

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

[28] In contrast, there are a total of 42 pairwise comparisons of delay effects (IMM vs EARLY, IMM vs LATE, EARLY vs LATE across 2 communication conditions for 7 belief measures). 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 t-test, $p < 0.05$), 2 of the negative effects are significant and 3 of the positive effects are significant. Finally, no effect remains significant after applying a Holm-Bonferroni correction.

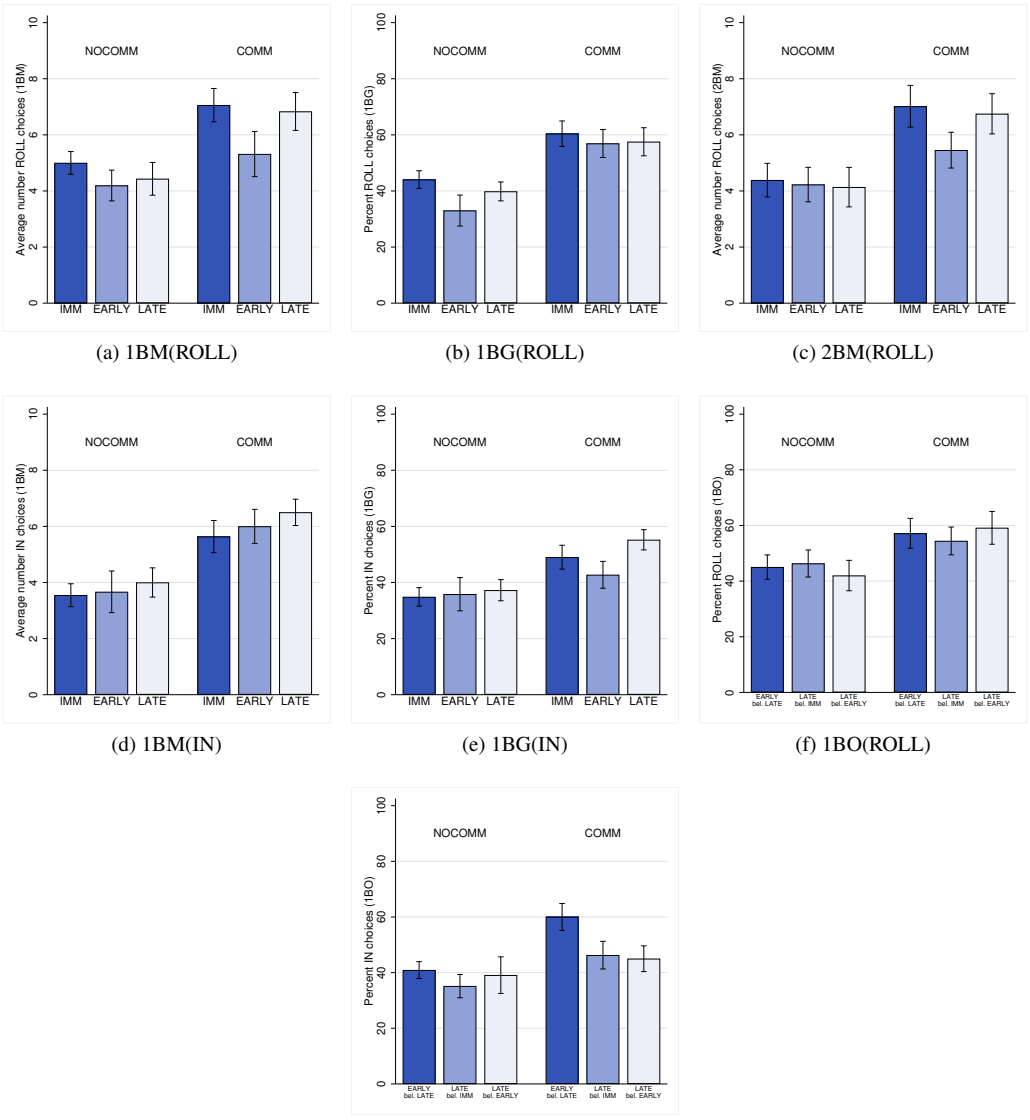

(a) 1BM(ROLL) (b) 1BG(ROLL) (c) 2BM(ROLL)

(d) 1BM(IN) (e) 1BG(IN) (f) 1BO(ROLL)

(g) 1BO(IN)

Figure 3. : Beliefs across conditions.

Note: (a) and (b) 1st-order beliefs about ROLL choices of matched trustee and group of trustees; (c) 2nd-order beliefs about ROLL choices of matched trustee; (d) and (e) 1st-order beliefs about IN choices of matched trustor and group of trustors; (f) and (g) 1st-order beliefs about ROLL and IN choices of groups in other delay conditions. Error bars represent standard errors from regressions of belief on treatment dummies and are adjusted for clustering on the session level.

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

	1BM(ROLL)	1BG(ROLL)	2BM(ROLL)	1BM(IN)	1BG(IN)	1BO(ROLL)	1BO(IN)
EARLY	-0.882	-1.168	-0.214	0.010	0.049	3.496	-3.276
	(0.525)	(0.526)	(0.659)	(0.715)	(0.602)	(6.316)	(3.920)
LATE	-0.790	-0.469	-0.255	0.448	0.268	7.445	-1.187
	(0.579)	(0.282)	(0.735)	(0.535)	(0.388)	(3.887)	(4.597)
COMM	1.959	1.608	2.606	2.008	1.382	10.688	8.447
	(0.592)	(0.432)	(0.766)	(0.537)	(0.413)	(3.878)	(2.912)
EARLY × COMM	-0.668	0.864	-1.343	0.384	-0.671	-5.304	9.356
	(1.014)	(0.726)	(0.834)	(0.981)	(0.772)	(8.259)	(5.665)
LATE × COMM	0.508	0.119	-0.053	0.356	0.335	6.101	4.370
	(0.913)	(0.569)	(0.987)	(0.742)	(0.488)	(5.709)	(5.933)
Time Preference	-0.038	-0.021	0.023	-0.009	0.019	-0.946	-1.598
	(0.075)	(0.062)	(0.098)	(0.074)	(0.059)	(0.687)	(0.605)
Risk Preference	0.208	0.302	0.002	0.136	0.102	2.696	0.401
	(0.130)	(0.095)	(0.148)	(0.138)	(0.099)	(0.758)	(0.927)
Age (years)	0.071	0.059	0.015	0.046	0.022	0.765	0.628
	(0.024)	(0.020)	(0.024)	(0.020)	(0.022)	(0.144)	(0.166)
Female	0.337	-0.157	-0.216	0.137	0.058	4.329	4.472
	(0.492)	(0.370)	(0.460)	(0.326)	(0.296)	(2.646)	(3.455)
Constant	2.046	1.233	3.966	1.504	2.076	4.179	39.101
	(1.039)	(0.783)	(1.492)	(1.025)	(0.971)	(7.329)	(9.807)
adj. R-squared	0.101	0.134	0.097	0.122	0.054	0.134	0.075
N	249	353	186	261	354	324	319
Clusters	36	40	32	36	40	40	40

Table 4—: Beliefs of Trustees and Trustors

Note: 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.

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 trustor's decision to trust and the trustee's decision to reciprocate trust, trustee behavior remained unchanged. Moreover, communication raised cooperation, trust, and trustworthiness by about 50% and this positive effect was as large as when choices immediately followed 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 persistence of trust and trustworthiness and the lasting effect of communication naturally raise 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. 2014. "From dual processes to multiple selves: Implications for economic behavior." *Journal of Economic Psychology*, 41: 1–11.

Andersen, Steffen, Uri Gneezy, Agne Kajackaite, and Julie Marx. 2018. "Allowing for reflection time does not change behavior in dictator and cheating games." *Journal*

of Economic Behavior & Organization, 145: 24–33.

Andreoni, James, and Marta Serra-Garcia. 2018. "The Pledging Puzzle: How Can Revocable Promises Increase Charitable Giving?" *CESifo Working Paper*.

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

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

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

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

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

Bhattacharya, Puja, Kirby Nielsen, and Arjun Sengupta. 2020. "Timing of Communication." *Economic Journal*.

Bicchieri, Cristina, and Azi Lev-On. 2007. "Computer-Mediated Communication and Cooperation in Social Dilemmas: An Experimental Analysis." *Politics, Philosophy, and Economics*, 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. 2017. "Registered Replication Report: Rand, Greene, and Nowak (2012)." *Perspectives on Psychological Science*, 12(3): 527–542. PMID: 28475467.

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

Byrd, B. Sharon, and Joachim Hruschka. 2006. "Kant on Why Must I Keep My Promise." *Chicago-Kent Law Review*, 81: 47.

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

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

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

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

Charness, Gary, and Martin Dufwenberg. 2011. "Participation." *American Economic Review*, 101(4): 1211–1237.

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

Di Bartolomeo, Giovanni, Martin Dufwenberg, Stefano Papa, and Francesco Passarelli. 2019. "Promises, expectations & causation." *Games and Economic Behavior*, 113: 137–146.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Newman, Matthew L, James W Pennebaker, Diane S Berry, and Jane M Richards. 2003. "Lying words: Predicting deception from linguistic styles." *Personality and social psychology bulletin*, 29(5): 665–675.

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

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

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

Pennebaker, James W, Roger J Booth, and Martha E Francis. 2015. *Linguistic Inquiry and Word Count: Operator's Manual.* Pennebaker Conglomerates.

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

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

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

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

Shea, George. 1879. *The Life and Epoch of Alexander Hamilton: A Historical Study.* Houghton, Mifflin.

Turmunkh, Uyanga, Martijn J van den Assem, and Dennie Van Dolder. 2019. "Malleable lies: Communication and cooperation in a high stakes TV game show." *Management Science*, 65(10): 4795–4812.

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

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

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