Building Trust: The Costs and Benefits of Gradualism*

Melis Kartal[†]

Vienna University of Economics and Business & VCEE

Wieland Müller[‡]

University of Vienna & VCEE

Tilburg University, CentER & TILEC

James Tremewan[§]

University of Vienna & VCEE

May 15, 2017

Abstract

To study the evolution of trust, we examine a setting with (i) an infinite time horizon; (ii) uncertainty regarding the trustworthiness of the receiver; and (iii) various possible levels of trust. We show that equilibrium strategies are characterized by a gradual increase in the degree of trust shown to reciprocating receivers. In a series of experiments, we find strong evidence that senders use such gradualist strategies. However, when we compare this setting to another in which trust is an all-or-nothing decision, we find that the setting that enables the use of gradualist trust strategies is not conducive to higher cooperation (in contrast to our theoretical prediction) and may even reduce efficiency by crowding out trustworthiness. These findings are corroborated by the results of several follow-up treatments.

JEL classification numbers: C73, C91, C92, D82, D83.

Keywords: cooperation, trust, infinitely repeated games, gradualism, game theory, experiments

*We would like to thank Guillaume Fréchette, Steffen Huck, Jens Prüfer, Karl Schlag, Martin Salm, Andy Schotter and members of seminar audiences at the University of Konstanz, Tilburg University and the Berlin IO Day for helpful comments. The research reported in this paper is funded by University of Vienna.

[†]Department of Economics, Vienna University of Economics and Business, Welthandelsplatz 1, 1020 Vienna, Austria, Email: *melis.kartal@wu.ac.at*.

[‡]Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, E-mail: *wieland.mueller@univie.ac.at*.

[§]Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, E-mail: *james.tremewan@univie.ac.at*.

1 Introduction

How do partners in an economic relationship establish trust and cooperation when one of them does not fully know the other's motives and preferences? Is it better to build the relationship cautiously with small, "testing" levels of trade or could this type of strategy backfire and damage the relationship, for example because such cautiousness and testing may be perceived as being offensive? Our study addresses these questions theoretically and experimentally.

Many economic relationships lack formal enforcement and rely on trust for generating mutual gains from trade.[1] Such relationships may involve a substantial risk especially at the beginning since partners may not immediately observe each other's real motives and commitment to the relationship. Therefore, one major concern at the beginning of relationships is the prospect of facing an opportunistic partner who is around for short term gains. This is why a first-time buyer may be concerned about the credibility of a supplier, a lender about the repayment incentives of a first-time borrower, and a manager about the tendency of a new employee to shirk responsibilities. Thus, an important question arises: how to build trust and cooperation in the face of uncertainty about a partner's preferences?

To investigate this question theoretically and experimentally, we develop a model in which two players, a sender and a receiver, play a generalized version of the so-called *trust game* infinitely many times. The sender moves first by choosing a trust level. If a positive trust level is chosen, the receiver chooses whether to *return* or *default*, whereas if the sender chooses not to trust, the receiver has no choice to make, and both parties receive their outside options. The receiver has private information regarding his "type," which is either high or low. The high type receiver is "cooperative," whereas the low type is "opportunistic."

At the early stages of the repeated game, the sender may be inclined to choose relatively low trust levels because this limits the payoff loss of the sender from betrayal by a low type receiver. This is indeed our theoretical prediction: the equilibrium exhibits "gradualism". Gradualism means that the relationship starts off with small, testing levels of trust, and the sender increases the trust level over time as long as the receiver reciprocates

[1] This is the subject matter of a large literature on self-enforcing agreements and relational contracts (see, among others, Bull (1988), Macleod and Malcolmson (1989), Baker *et al.* (2002) and Levin (2003)).

the sender's trust. This enables the sender and the high type receiver to reach the efficient level of cooperation after a testing phase.

The theme of gradualism has appeared in a large economics literature with private information about players' preferences, for example Sobel (1985), Ghosh and Ray (1996), Kranton (1996), Watson (1999, 2002), Halac (2012) and Kartal (in press). These papers study in distinct environments how cooperation is established in the presence of private information and point out that gradualism is an effective remedy against problems caused by informational asymmetry.

While there is a huge theoretical literature on reputation formation and cooperation building in infinitely repeated games, there is no empirical evaluation of this literature.[2] Importantly, gradualist strategies are currently being used in practice, most notably, in microcredit lending. This practice is called "progressive lending" and involves gradually increasing loan sizes for borrowers who are in good standing.[3] The role of gradualism in relationships is not limited to market settings. For instance, many marriages are initiated after partners fulfill numerous smaller steps (dates, vacations, cohabitation, engagement, *etc.*). On the one hand, gradualist strategies can in theory alleviate private information problems in building trust. On the other hand, they are quite sophisticated (that is, people may "test" too little or too much relative to the optimal level) and might have unwanted effects due to the inherent testing of the receiver, calling into question the behavioral relevance of these strategies. Therefore, the lack of empirics regarding the use and effects of gradualist strategies is an important gap in the literature, which we try to fill with this paper.

For this purpose, we implement a trust game in which the sender chooses either no trust or one of three positive trust levels: low trust, medium trust or high trust. We call this game the "Gradual game." As mentioned above, the equilibrium is characterized by

[2] There is a growing experimental literature on infinitely repeated games. However, this literature has focused only on symmetric information games. There is also a fairly large experimental literature on reputation building in finitely repeated games with private information. However, it is typically not possible to build up a relationship in games with a fixed, finite duration. For example, in a finitely-repeated binary trust game, a sender is less likely to trust his partner over the course of the game even if the receiver has always returned in the past and has an impeccable record (see, for example, Camerer and Weigelt (1988)).

[3] Morduch (1999, p. 1582) writes: "[Microcredit] Programs typically begin by lending just small amounts and then increasing loan size upon satisfactory repayment. The repeated nature of the interactions—and the credible threat to cut off any future lending when loans are not repaid—can be exploited to overcome information problems and improve efficiency ...."

gradualism: the sender starts by choosing low trust and increases the trust level gradually as long as the receiver returns. Thus, in the Gradual game the sender uses what we refer to as "gradualist incentives", on top of "dynamic incentives" in the sense that default by the receiver is punished by no trust thereafter. Moreover, in every equilibrium of the Gradual game the sender and the high type receiver achieve full cooperation after a testing phase.

The experimental results provide strong support for our gradualism hypothesis. We show that a large majority of the sender behavior can be classified as belonging to one of six simple categories, the largest category comprising behavior consistent with equilibrium. Furthermore, we find that given the observed receiver behavior, equilibrium strategies make senders better off than non-equilibrium strategies.

In the Gradual game, the existence of an equilibrium which enables cooperation relies crucially on the availability of low and medium trust levels. In other words, if the sender had only two options, no trust and high trust, then the unique equilibrium would prescribe no trust in all periods, given our experimental parameters. Based on this observation, we have a second treatment in which the sender chooses between no trust and high trust—we call this game the "Binary game." The Binary game differs from the Gradual game only in that low and medium trust levels are absent; all other parameters are held constant. Hence, in the Binary game the sender can provide only dynamic but not gradualist incentives. While the sender *never* trusts in the unique equilibrium of the Binary game, the equilibrium in the Gradual game always entails positive levels of trust—unless the receiver defaults—and players eventually reach full cooperation provided that the receiver is a high type. Thus, the distinction between the two games and their respective equilibria displays in a stark way how gradualism is, in theory, indispensable to achieving high cooperation.

In contrast to the Gradual game, observed behavior in the Binary game diverges significantly from the theoretical prediction. In fact, a substantial share of senders choose to trust and many receivers choose to return—even when they are experienced—in clear contrast to the equilibrium prediction. As a result, payoffs in the Binary game are remarkably higher than our theoretical prediction and also higher than the payoffs in the Gradual game, again in clear contrast to our prediction. There are, however, two important additional observations. First, the increased payoffs in the Binary game stem from the fact that this game significantly boosts receiver payoffs. More precisely, sender payoffs do *not* differ across the two games. Second, the Gradual game provides senders with "insurance": the variance of sender payoffs is significantly lower in the Gradual game than in the Binary game, whereas the average sender payoff does not differ across the two treatments. Thus, the Gradual game makes *risk-averse* senders better off, whereas the Binary game makes receivers better off.[4]

Why does behavior in the Binary game deviate so much from the theoretical prediction? We provide evidence suggesting that many subjects choose to trust because they have a "homemade" belief that the proportion of high types is greater than the prior probability generated in the experiment (see Camerer and Weigelt (1988)). Senders' homemade beliefs may represent, for example, the percentage of receivers who are believed to be naturally altruistic so that they will always return rather than default. Moreover, such a belief is empirically justified given the observed receiver behavior. As it turns out, in the Gradual game the theoretical predictions are more robust to the effect of a homemade belief than they are in the Binary game. That is, a much higher homemade belief is required to change the predictions in the Gradual game than in the Binary game. Therefore, we ran follow-up treatments on both the Gradual and the Binary game in which we varied our baseline parameter values in a way that the Binary game prediction is more robust to the effect of a homemade belief. The results of this new pair of treatments were qualitatively very similar to what we observed before. This is true not only with respect to the frequent use of gradualist strategies in the new Gradual game, but, to our surprise, also with respect to the high levels of trust observed in the new Binary game. Again, joint, sender and receiver payoffs in the new Gradual game are not higher than in the Binary game.

We then ran a second pair of follow-up treatments, which varied our parameter values in a way that not only the Binary game is robust to the effect of homemade beliefs but also efficiency gains from cooperation are lower. We expected this to have a large effect on the Binary game in the form of increased no trust choices, and much less of an effect in the Gradual game due to the presence of intermediate levels of trust. But once again, results were qualitatively similar to what we found before: gradualist strategies were also commonly

[4] Although we assume risk-neutrality in the model, moderate levels of risk-aversion do not affect any of our theoretical predictions.

used in the new Gradual game, and we again observed a high frequency of trust choices in the new Binary game. Again, in this second set of new games, we never find payoffs in the Gradual game to be higher than in the Binary game.

To summarize, we find strong evidence for the gradualism hypothesis in the baseline Gradual game treatment as well as in the follow-up versions. This is our first contribution to the literature. To our knowledge, this is the first and the only experimental study of cooperation and reputation formation in an infinitely repeated setting. Our second contribution is the finding that the Binary game is (weakly) more efficient than the Gradual game in sharp contrast to the equilibrium predictions.[5] Thus, our experimental results suggest that when "dynamic incentives" are in place, using "gradualist incentives" on top of dynamic incentives is redundant and may even backfire.

In relation to this finding, we note that gradualism may create unintended behavioral consequences by "crowding out" the prosocial incentives of receivers. Conversely, the absence of testing strategies may "crowd in" prosocial incentives. A large literature has shown that certain variations in institutions, the choice set and other details of a game can crowd in or crowd out trustworthiness, honesty and prosocial behavior (Gneezy *et al.* (2011) provides an extensive survey of the literature). What is the underlying mechanism in our design? In our Binary game, testing strategies are absent and as a result, senders choose to trust a new partner most of the time rather than not to trust. This signals either that the behavioral norm is to return or that senders trust the receivers' intrinsic motivation to return (or both). As a result, trustworthiness is crowded in. However, in the Gradual game senders choose to test a new partner most of the time, which signals either that the prevailing norm of behavior is selfish or that senders do not trust the receivers' intrinsic motivation to return (or both). Thus, using gradualist incentives undermines receivers' prosocial motivations.

The crowding in and crowding out phenomena are intrinsically related to the homemade beliefs discussed above. In particular, if a sender has roughly correct homemade beliefs, then those beliefs should be higher in the Binary game than in the Gradual game because the Binary (Gradual) game crowds in (out) trustworthiness. This difference in homemade

[5] As discussed before, senders' payoffs are significantly less variable in the Gradual game and, therefore, risk-averse senders may prefer the Gradual game over the Binary game if they were given the choice.

beliefs across the two games will thus make the Binary game more conducive to cooperation than the Gradual game by encouraging initial sender trust, which in turn boosts receiver trustworthiness, justifying the higher homemade belief in the Binary game.

Our findings imply that an institution that uses only dynamic incentives may be more efficient than an institution that uses a combination of gradualist and dynamic incentives. This may be especially relevant for non-profit microcredit lenders that maximize borrower welfare subject to a break-even constraint. Given our findings regarding the effect testing strategies have on efficiency, it may be beneficial to evaluate the behavioral consequences of progressive lending in the field even if those policies may, in theory, be sensible and intuitive.[6]

2 Related Literature

There is an extensive reputation literature dating back to the seminal works by Kreps and Wilson (1982), Milgrom and Roberts (1982) and Kreps *et al.* (1982). While Kreps and Wilson (1982) and Milgrom and Roberts (1982) analyze an entry-deterrence game, Kreps *et al.* (1982) show that incomplete information about players' preferences can sustain cooperation in the equilibrium of a finitely repeated prisoners' dilemma game.[7]

Experimenters have implemented the model by Kreps *et al.* (1982) and its variants in the laboratory in order to test the sequential equilibrium predictions. While Andreoni and Miller (1993) study a finitely repeated prisoners' dilemma game as in Kreps *et al.* (1982), Camerer and Weigelt (1988) and the subsequent studies by Neral and Ochs (1992), Anderhub *et al.* (2002), Brandts and Figueras (2003), and Grosskopf and Sarin (2010) study whether sequential equilibrium is successful in predicting behavior in a finitely repeated binary trust game, varying certain dimensions in the design, such as the probability with which the receiver is a trustworthy type, the observability of past choices of the receiver to

[6] On a related note, many microcredit borrowers choose to become members of multiple institutions and take multiple loans (Vogelgesang, 2003). This is likely because the loan size from one institution is too small and taking loans from multiple institutions is one possible remedy. However, this is a costly remedy since dealing with multiple institutions brings about additional costs to the borrower, which subsequently increases the chances of a default. Thus, it seems that testing borrowers with very small loans may have *both* intrinsic and extrinsic costs.

[7] Suppose that each player is a type that is "committed" to a tit-for-tat strategy with some positive probability. Then, reputation concerns can motivate self-interested players to imitate a committed type for some time and generate cooperative behavior even with self-interested players.

the sender, whether the trustworthy type is computerized, *etc.* Brown and Serra-Garcia (in press) analyze a finitely repeated lending game in which a creditor can make one of several loan offers to a borrower to study the effect the threat of terminating the relationship has on the nature and functioning of implicit loan contracts. They find, among other things, that if the threat of termination is weak, the creditor-borrower relationship breaks down often and leads to smaller loans. Huck *et al.* (2016) use a binary trust game to study reputation building for quality on the part of a seller.

The experimental papers discussed above assume that there are only two choices available to players and that the game is finitely repeated, with the exception of Brown and Serra-Garcia (in press) who also use a finitely repeated game but allow for a bigger choice set on the part of the creditor. Given this game structure, it is not possible to build up a relationship. For example, the sender is less likely to choose trust over the course of the game even if the receiver has always returned in the past and holds a perfect record, which is due to the fixed duration of the game.[8] Moreover, many economic relationships have an indefinite time horizon rather than a fixed one. Therefore, we focus on an infinitely repeated trust game with variable trust levels. In this game, equilibrium strategies are gradualist, and they enable the sender and the high type receiver to reach the efficient level of trade after a testing period. This connects our paper to a branch of the theoretical reputation literature that draws a similar conclusion in diverse settings (see, for example, Sobel (1985), Ghosh and Ray (1996), Kranton (1996), Watson (1999, 2002), Halac (2012), Kartal (in press)).[9] Halac (2012) and Kartal (in press) study relational contracting with private information. A relational contract is analogous to a trust game between a principal and an agent. Halac (2012) analyzes relational contracts in a setting where the principal has private information regarding her outside option, whereas Kartal (in press) assumes that the principal is privately

[8] See the theoretical predictions in Camerer and Weigelt (1988).

[9] In Ghosh and Ray (1996), Kranton (1996), and Watson (1999, 2002), two partners simultenously choose the level of cooperation in every period in a prisoners' dilemma-like setting and decide individually whether or not to behave opportunistically. There is two-sided hidden information: "High" type players prefer to cooperate as long as their partners also cooperate whereas "low" types have an incentive to take advantage of the other player's trust. In Ghosh and Ray (1996) and Kranton (1996), agents are also allowed to break a partnership to form a new one. These papers show that there is an initial testing phase in which partners of high type build trust through their actions, and the stakes in the relationship grow as players trust each other more.

informed about her discount factor, similar to our setting.[10]

Finally, our paper relates to a burgeoning experimental literature on infinitely repeated games (see, for example, Palfrey and Rosenthal (1994), Engle-Warnick and Slonim (2004, 2006a, 2006b), Dal Bó (2005), Aoyagi and Fréchette (2009), Camera and Casari (2009), Duffy and Ochs (2009), Dal Bó and Fréchette (2011), and Fudenberg *et al.* (2012)). Engle-Warnick and Slonim (2004) is particularly related to our paper. They study an infinitely repeated binary trust game with *symmetric* information and show that the rates of trust and return are at substantial levels in the infinitely repeated game. Moreover, experienced senders are significantly more likely to trust and experienced receivers are significantly more likely to return in the infinitely repeated game than in the finitely repeated game. However, it is important to point out that, unlike in our Binary game, trust is an equilibrium outcome in their infinitely repeated game. Some other studies present weaker experimental evidence regarding the positive effects of infinite repetition. For instance, Palfrey and Rosenthal (1994) find that repetition leads to higher cooperation in a public good contribution game, but the magnitude of this increase is small even though they use a very high discount rate. Dal Bó and Fréchette (2011) analyze two types of repeated prisoners' dilemma games: one in which perpetual cooperation is an equilibrium and another one in which cooperation is not an equilibrium. They find that cooperation is not necessarily higher in the former and that even if cooperation is supported in equilibrium and is the risk-dominant action, the actual cooperation rate may be low.

To our knowledge, our paper is the first experimental study that incorporates private information in an infinitely repeated setting. Moreover, most of the experimental literature on infinitely repeated games focuses on simple stage games with two actions. Therefore, previous studies do not address the issues (i.e., gradualism) that we analyze in this paper.

[10]The discount factor of the principal matters in relational contracts because such contracts are sustained by the value of the future relationship, and this value increases in the discount factor of the principal.

3 Experimental Design

3.1 Experimental Game and Equilibrium Predictions

In the Online Appendix A.1, we present and analyze a general model of trust with private information. Our main treatment is a simplified version of this general model: We investigate an infinitely repeated game with the following features (the stage game is shown in Figure 1). At the beginning of each period $t = 0, 1, \ldots$, the sender chooses either no trust (N) or one of three positive trust levels: low trust (L), medium trust (M) and high trust (H). If a positive trust level is chosen, then the receiver chooses between return (R) and default (D). The receiver is privately informed regarding his type, which is either high or low. The low type receiver is "opportunistic," whereas the high type receiver is a "cooperative" commitment type who always returns the sender's trust. The prior probability of meeting a high type player is equal to 1/8, which is common knowledge.

Hereafter, the receiver refers to the low type receiver unless otherwise stated since the high type always chooses R. Figure 1 displays the stage game as well as the sender and receiver payoffs which we implement in the experiment. Payoff parameters are consistent with the typical trust game à la Berg *et al.* (1996). That is:

- The sender and the receiver have conflicting interests in the stage game: If a positive trust level is chosen, then the sender is better off if the receiver returns, whereas the receiver is better off if he defaults.

- The payoff from mutual cooperation (i.e., the sender chooses a positive trust level and the receiver returns) increases in the trust level for both parties.

- The higher the trust level, the lower (higher) the payoff of the sender (receiver) if the receiver defaults.

The sender has a discount factor of 0.75, whereas the low type receiver has a discount factor of 0.5. The low type is deemed opportunistic because with a discount factor of 0.5, the low type would *never* return the sender's trust in the infinitely repeated setting with symmetric information, and thus, the unique equilibrium would involve no trust. However,

Figure 1: The game tree of the Gradual game treatment

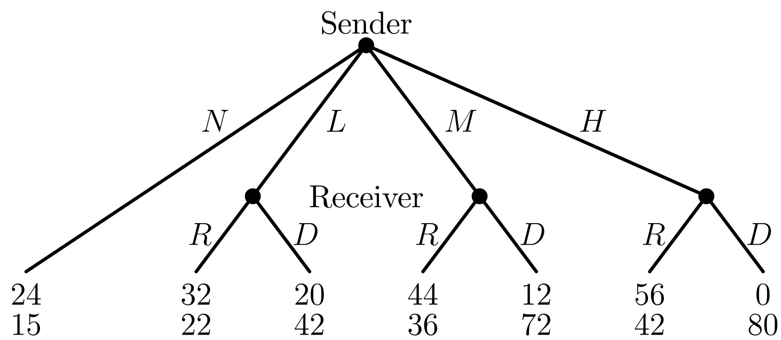

in the private information game, the low type may strategically choose R at lower trust levels in order to *imitate* the high type and build a reputation. This will be the case if returning in the current period will be reciprocated with a higher trust level in the future and defaulting at that higher trust level compensates the loss from returning in the current period.

Let us now discuss the perfect Bayesian Nash equilibria (PBE) of the infinitely repeated game in view of the parametrization discussed above and the payoffs in Figure 1. Our choice of experimental design and parameters serves our primary purpose of testing for the use of gradualist strategies. We find that every PBE of the repeated game is gradualist, and all PBE are qualitatively identical at $t = 0$ and $t = 1$. In every equilibrium, the sender starts the relationship with L, weakly increases the trust level as long as the receiver returns, and punishes default with perpetual no trust. In every equilibrium high trust is eventually reached after a testing phase if the receiver is a high type.[11] A formal and complete description of the PBEs of the experimental game as well as the related proofs can be found in Online Appendix A.2. Note that while we assumed risk neutrality in our analysis, the equilibrium predictions are qualitatively robust to moderate levels of risk-aversion and homemade beliefs (Camerer and Weigelt (1988)).

As discussed above, all PBEs of the game are similar. Therefore, and for the sake of simplicity in the main text we focus on a specific equilibrium, which we denote by E. After

[11] Note that the nonexistence of an equilibrium with no trust in all periods is due to the assumption that the cooperative type is a "commitment" type that always returns trust. Since the choice of L is not too risky for the sender and the prior of 0.125 is sufficiently high, the sender would always deviate from a no-trust equilibrium prescription and test the receiver by choosing L (see the details of the proof in the Online Appendix A.2).

explaining E, we briefly discuss how other PBE differ. Equilibrium E has the following features: It generates the highest possible payoff for each player, provides the fastest information revelation and involves the least amount of mixing (which is usually problematic in experimental studies as it cannot easily be observed in the data). Recalling that the receiver refers to the low type receiver unless otherwise stated, the equilibrium E is described as follows:

(i) *At $t = 0$, the sender chooses low trust (L), and the receiver randomizes between returning (R) and defaulting (D).*

(ii) *If the receiver defaults at $t = 0, 1, ...$, the sender chooses no trust (N) thereafter.*

(iii) *At $t = 1$, the sender randomizes between low trust (L) and medium trust (M) provided that the receiver chose R at $t = 0$.*

1. *If the outcome of this randomization at $t = 1$ is L, then the receiver returns (R) with probability one.*

2. *If the outcome of this randomization at $t = 1$ is M, then the receiver defaults (D) with probability one (whenever M is chosen, D is strictly optimal for the receiver).*

(iv) *At $t = 2$:*

1. *If the sender randomization at $t = 1$ resulted in L and the receiver chose R, then the sender chooses M and the receiver chooses D with probability one.*

2. *If the sender randomization at $t = 1$ resulted in M and the (high type) receiver returned (R), then the sender chooses H (H will be chosen in equilibrium E if and only if the receiver is a high type given part 2 in* (iii) *above).*

(v) *At the end of $t = 2$, the sender knows with certainty the type of the receiver given* (i)-(iv). *The sender chooses H at $t = 3, 4, ...$ if the receiver is high type and N otherwise.*

To reiterate, if M is chosen, it is strictly optimal for the receiver to default no matter what the sender's strategy prescribes thereafter (we opted for a design with this feature in order to stack the deck against equilibrium predictions).[12] Thus, once the receiver returns after M is chosen (whether at $t = 1$ or at $t = 2$), the sender learns that the receiver is high type with probability one and chooses H thereafter. This means that H is reached *only* with a high type receiver because M always precedes H in equilibrium, and the low type optimally defaults whenever M is chosen. Note that the low type receiver always returns with strictly positive probability if L is chosen unlike the case with M or H. This enables the low type to imitate the high type and "build a reputation" at the low trust level and to default later when the trust level increases.

The remarks above apply to every PBE, and the only distinction across various equilibria at $t = 0, 1$ is "numerical." That is, only the probability of choosing R (after the sender chooses L) at $t = 0$ and the probability of choosing L at $t = 1$ (assuming that L and R were chosen at $t = 0$) vary across equilibria. At $t = 2$, E prescribes M if L and R were chosen at $t = 0, 1$, whereas other PBE prescribe randomization between L and M. As a result, information revelation will (in expectation) take longer in equilibria other than E because the low type receiver's equilibrium strategy if L is chosen at $t = 2$ involves R with strictly positive probability. As explained above, choosing R at low trust enables the low type to build up his reputation and default later once the trust level increases.

Although our experimental game and equilibrium results may seem to be specific, in the Online Appendix A.1 we show that gradualism is a robust feature of equilibrium strategies in a much more general model.

Our main treatment implements the experimental game discussed above, which we call the "Gradual game" treatment. Note that the availability of L and M is crucial for the existence of an equilibrium that involves positive trust levels. If the sender had only N and H as options, then the unique equilibrium is such that the sender always chooses N because H does not pay off given the payoff parameters and the prior probability that we chose.[13]

[12] Note that a design in which defaulting after M is the dominant strategy for the receiver is only possible with receiver and sender payoffs that vary in a nonlinear fashion. Such payoff parameters enable us to have at least one equilibrium that involves gradualism with *at most* one period of mixing per player. Having more periods of mixing required by the equilibrium would make it easier to obtain evidence of equilibrium behavior, but we preferred to stack the deck against theoretical predictions. See also the discussion at the end of Section 4.1.1.

[13] The proof for this is straightforward. The expected sender payoff of choosing N at every t is $24/(1 - 0.75) = 96$. If the sender chooses H, then the low type will (optimally) default, and the sender will always

Figure 2: The game tree of the Binary game treatment

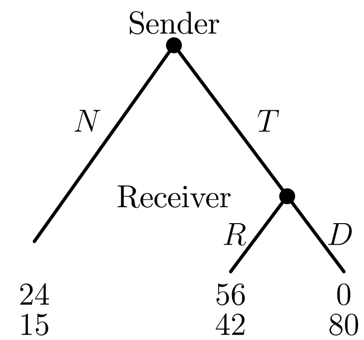

This theoretical observation underlies our secondary treatment, which we call the "Binary game" treatment. The Binary game treatment is identical to the Gradual game treatment with the exception that L and M are absent; all other parameters and design details are held constant across the two treatments (see the Binary game in Figure 2). The sender chooses between no trust (N) and trust (T) in the Binary game, and choosing T in the Binary game is identical to choosing H in the Gradual game in terms of payoff prospects for the sender and the receiver. Below, we summarize our main hypotheses derived from the equilibrium analysis of the two games.

Hypothesis 1 (Gradualism) Senders use the testing strategies prescribed in the PBE of the Gradual game (i.e., they start the game with L and gradually increase the trust level unless default is observed).

Hypothesis 2 (Efficiency of Gradualism) The Gradual Game is strictly more efficient than the Binary Game. To be more specific, the Gradual game makes both the sender and the receiver strictly better off than the Binary game.

In the experimental design, we chose a trust game with four possible levels of trust in order to balance two main concerns. On the one hand, we wanted the number of trust levels to be sufficiently high so that the sender uses testing levels of trust and *gradualism* is a viable equilibrium outcome. On the other hand, we wanted the number of trust levels to

choose N from then on. Thus, the expected payoff of choosing H at $t = 0$ and continuing with H in case the receiver returns is

$$\frac{7}{8}\left(0 + 0.75\frac{24}{1-0.75}\right) + \frac{1}{8}\left(\frac{56}{1-0.75}\right) = 91 < 96.$$

be relatively low so that we can have a very precise measure regarding Hypothesis 1, which aims at testing the consistency of observed sender behavior with equilibrium predictions. In addition, the game with four trust levels is intuitive and easy to explain to subjects, and a complete equilibrium analysis is feasible without restricting the strategy space to Markovian strategies, which is typically necessary in order to characterize equilibria in games with a larger action space (indeed, the theoretical analysis of our general model presented in the Online Appendix A.1 is only tractable with a restriction to Markovian strategies).

3.2 Experimental Protocol

Our design is between-subject: Each subject participated either in the Gradual game treatment or the Binary game treatment. In both treatments, subjects participated anonymously in a sequence of infinitely repeated games. At the beginning of each session, each subject was randomly assigned to be a sender or a receiver. Subjects remained in the same role throughout the session. In the experiment, the high type receiver is computerized and always returns a sender's trust if the sender chooses a positive trust level. We computerize the high type receiver in order to have more control over subjects' behavior.[14] Each session involved 15 "human" players: 8 senders and 7 receivers. The eighth receiver was the computerized player. We used neutral language, so that there was no mention of trust, return or default in the experiment.[15]

We refer to each repeated game as a "match" and each repetition of the game within a match as a "round." At the beginning of a match, each sender was randomly matched with a receiver, and a subject played either the Gradual game or the Binary game depending on the treatment. Senders knew that there was a 1/8 chance of being matched to a computerized receiver in each match but senders were never informed about the true type of the receiver with whom they were interacting.[16] Senders and receivers were randomly rematched after the end of each match. Each session ended either after 25 matches were completed, or at the end of the first match that was completed after 75 minutes passed, whichever occurred first.

[14] Numerous experiments on hidden information games involved computer players (see, for example, Andreoni and Miller (1993), Anderhub *et al.* (2002), Grosskopf and Sarin (2010), and Embrey *et al.* (2015)).

[15] No trust was labeled as A1, low trust as A2, and so on. Return was labeled as B1 and default as B2.

[16] Of course, senders can infer that a receiver who chooses to default is a human, not a computer.

The maximum of 25 matches was reached in all sessions but one, in which only 23 matches were completed.

We implemented an infinitely repeated game in the lab by using a random continuation rule. The probability of continuation after each round, δ was the same for all matches in both treatments and was equal to the discount factor of the sender; i.e., $\delta = 0.75$. Each human receiver was, by design, a low type receiver and had a discount factor of 0.5. This was implemented as follows.[17] Starting from the second round of a match, the receiver payoffs in Figures 1 and 2 were reduced by 1/3 in every round. For example, if the match continues to the second round and the sender chooses no trust in the second round, then the receiver obtains 10 points in the second round, rather than 15(since the continuation probability of 0.75 times 2/3 equals 0.5, the discount factor of the receiver is effectively equal to 0.5). This was clearly explained in the experimental instructions. Hence, our design for the low type's discounting combines the case where the future is less valuable than the present and the case where the interaction randomly terminates with a positive probability. Note that our design is a variation on an already existing experimental method in which pure payoff discounting is followed by pure random termination. To be more precise, this method involves a fixed, known number of rounds played with certainty, and payoffs in these rounds are discounted at a known discount rate of δ. After the rounds with certainty are played, payoffs are no longer discounted and in every round, the interaction continues for an additional round with probability δ (see Fréchette and Yuksel (in press) and the references therein).[18]

The experiment was conducted at the experimental laboratory of the Vienna Center for Experimental Economics (VCEE) at the University of Vienna. Subjects were recruited from the general undergraduate population. A total of 180 undergraduates participated in 12 experimental sessions. We conducted six sessions for each treatment, which translates to six independent observations per treatment. All sessions were conducted through computer terminals, using a program written in Z-Tree (Fischbacher, 2007). The experiment only started when all subjects had correctly answered a set of control questions that tested whether

[17] It would also be possible to implement the same discount factor of 0.5 for both the sender and the user and use only the random termination method but this would generate data predominantly from very short supergames, which is not desirable in a study of gradualism.

[18] Fréchette and Yuksel (in press) study different implementations of infinitely repeated games (with symmetric information) in the laboratory and find that this method generates the most stable cooperation.

they understood the instructions. At the end of each session, the total number of points earned by each subject was converted to Euros at the exchange rate of 100 points = 1 Euro and paid privately in cash. The average payoff per subject was about €22.60, which includes €2 for filling in a post-experimental questionnaire. The instructions for both treatments can be found in the Online Appendix G.

4 Results

In this section, we first discuss the behavior in the Gradual and the Binary games, and then present a payoff analysis of the two games. After a discussion of our findings in the Gradual and the Binary game treatments, we describe our follow-up treatments followed by concluding remarks. In our analysis, we focus on experienced play and discard the data of the first 10 matches unless otherwise stated.[19] For our nonparametric tests, the unit of observation is the session average, and for all tests we report two-sided significance levels.

4.1 Actions and Strategies

4.1.1 Behavior in the Gradual Game

We begin by discussing sender and receiver behavior in the first round of a match, as reported in Table 1. Low trust (L) is by far the most common choice: 67.1% of senders choose L in the first round of a match. Moreover, the proportion of L choices in the first round exceeds 57% in every session of the Gradual game treatment, whereas random play would result in 25% as the proportion since the sender has four possible actions in the stage game. Thus, senders' preference for L in the first round is far from being the result of random behavior according to a binomial test ($p < 0.001$).

The second most popular sender choice in the first round is medium trust (M), which is chosen only 16.5% of the time. Both a Wilcoxon signed-rank test and a sign test reject the equality of proportions of L and M choices ($p < 0.001$). Of course, this also holds when

[19] We observe that the percentage of behavior consistent with equilibrium predictions increases considerably during the first 10 matches. We elaborate on learning in the Gradual game in Online Appendix B, where we provide an analysis of behavior in the first 10 matches.

Table 1: Sender and receiver behavior in the first round of a match in the Gradual game

Sender	No Trust (N)	Low Trust (L)		Medium Trust (M)		High Trust (H)	
	51	467		115		63	
	(7.3%)	(67.1%)		(16.5%)		(9.1%)	
Receiver	–	Return	Default	Return	Default	Return	Default
	–	300	167	45	70	21	42
	–	(64.2%)	(35.8%)	(39.1%)	(60.9%)	(33.3%)	(66.7%)

Note: Numbers of observation and percentages in parentheses (including the computerized receiver).

we compare L with high trust (H) and no trust (N), which are chosen with a frequency of 9.1% and 7.3%, respectively. We conclude that there is strong support for our equilibrium prediction of L choice in the first round.[20]

How do receivers respond to senders' trust choices in the first round of a match? Table 1 shows that the rate of return is at its highest if the sender chooses L (64.2% return rate) and considerably lower if the sender chooses M (39.1%) or H (33.3%). These return rates include the computerized receiver. Excluding the data from the computerized player, the difference between the return rate if L is chosen and the return rate if M (or H) is chosen is statistically significant according to both a Wilcoxon signed-rank test and a sign test ($p < 0.001$ for both tests). This is consistent with our theoretical predictions. As discussed in Section 3.1, the receiver has in theory no willingness to return after M or H is chosen but if he returns after L, he can build up a reputation and default later when the trust level is higher. We gave subjects a post-experimental questionnaire in which we asked senders and receivers about their incentives to trust and return, respectively. Consistent with the theory, many receivers stated that if a low trust level was chosen, they imitated the computerized player with the hope that the sender would choose a higher trust rate in the coming rounds, which makes default more profitable.

Next, we classify senders' behavior based on the outcomes in the first two, three and four rounds of a match. Note that in order to provide a classification of sender behavior in the first $n \geq 2$ rounds, we must restrict the data to matches that last *at least* n rounds. In

[20] We also ran a random effects panel probit regression of the probability of choosing L in the first round of each match with the match number as the independent variable. The coefficient of the match number is insignificant, demonstrating that there is no time trend if we focus on experienced sender behavior.

view of this requirement, we limit our analysis to the first four rounds. Classifying sender behavior in the first five rounds, for example, requires restricting the data to matches that last at least five rounds, which results in substantial loss of data and thus small numbers of observations.[21]

We begin by identifying game histories that are "gradualist", as described below. These we further classify according to whether or not they are consistent with equilibrium predictions.

Gradualist: Sender behavior is gradualist if the sender starts the relationship with either L or M, (weakly) increases the trust level as long as the receiver returns, and punishes default with no trust thereafter.

Equilibrium: Sender behavior is equilibrium if it is consistent with the equilibrium predictions discussed in Section 4.1. The proportion of equilibrium behavior is our main variable of interest.

Non-equilibrium Gradualist: Sender behavior is gradualist, as described above, but is not consistent with equilibrium predictions.

Among game histories that are not consistent with our definition of gradualist, many can be classified as being consistent with four intuitive strategies which we identified after observing the experimental data:

Always no trust: The sender always chooses N in which case the receiver has no choice to make.

High trust until default: The sender starts by choosing H in the first round and chooses H if the receiver returned in all previous rounds, and N otherwise.

Lenient: The sender behavior is lenient if the sender starts with a positive trust level, (weakly) increases the trust level as long as the receiver returns, and punishes default leniently by either choosing a strictly lower but positive trust level or choosing N for a limited number of rounds followed by a positive trust level.

Hybrid: The sender behavior is hybrid if the sender starts with the strategy "always no

[21] Figure 5 in Online Appendix D.1 shows the frequency of matches with respect to the number of rounds played in a match. Classifying sender behavior in the first five rounds requires discarding more than 71% of the matches in the Gradual game treatment. We classify behavior for two, three, and four rounds because although the analysis based on fewer rounds is less informative, it is based on a greater number of observations.

trust" and then switches to an equilibrium strategy in the second, third or fourth round. This type of behavior suggests that the sender is indifferent between "always no trust" and "equilibrium" at the beginning of the game.

We define sender behavior that does not fit into any one of these categories as "unclassified." Of course, in the experiment we do not observe entire sender strategies; only histories of sender behavior, which we classify according to the categories above. Since the equilibrium strategies predict randomizing between two actions in some rounds, sender behavior satisfies the criterion of being "equilibrium" if either one of these actions is chosen. For example, focusing on match histories up to (and including) the sender's second decision, three are consistent with equilibrium: All three begin with low trust at $t = 0$, then involve playing either low or medium trust at $t = 1$ if the receiver returned at $t = 0$, and no trust if the receiver defaulted at $t = 0$. For further clarification, the reader may consult the Online Appendix D.2, which contains complete lists of the frequencies of the observed match histories up to the second, third, and fourth sender decisions as well as the assignment of these match histories to the categories described above.

The results are shown in the upper third of Table 2. (For the time being, please ignore the lower two thirds of this Table, which shows the results of the follow-up treatments that will be introduced in Section 4.4.) For $n = 2$, about 56% of sender behavior is consistent with equilibrium, and the three possible equilibrium histories up to (and including) the sender's second decision, which we described above, are the three most frequently observed outcomes in the data.[22] Moreover, the proportion of equilibrium behavior exceeds 40% in every session. If both players make choices randomly, then the expected proportion of equilibrium behavior that will be generated is 11.3%.[23] The prevalence of equilibrium behavior is far from being random according to a binomial test ($p < 0.001$). Finally, the proportion of gradualist behavior in the first two rounds is 75.2%.

[22] See the Online Appendix D.2.

[23] If the sender is matched with a human receiver, then there are 28 possible paths of play. Since a computerized receiver always returns, there are 16 possible paths of play with a computerized player. Only 1 out of 8 receivers is computerized, and three histories are equilibrium. Therefore, we have $3/(28 \cdot 7 + 16) \cdot 100 = 11.3\%$.

Table 2: Classification of Sender Behavior in the Gradual games in the first n rounds, $n = 2, 3, 4$

		Gradualist			Other				
		Equilibrium		Non-Equilibrium	Always No Trust	High Trust until Default	Lenient	Hybrid	Un-classified
Treatment	Matches that last at least n rounds	In Data	According to random play						
Gradual	$n = 2$	55.8%	11.3%	19.4%	2.1%	5.4%	3.7%	5.6%	13.6%
	$n = 3$	47.7%	3.3%	18.7%	2.4%	5.2%	9.1%	3.9%	16.9%
	$n = 4$	38.3%	0.8%	17.4%	1.9%	4.9%	16.3%	4.5%	21.2%
hr-Gradual	$n = 2$	46.3%	11.3%	9.1%	13.4%	6.3%	4.5%	8.4%	12%
	$n = 3$	41.3%	3.3%	9.6%	9%	6.1%	7.3%	8.1%	18.6%
	$n = 4$	41.3%	0.08%	7.2%	8.5%	5.5%	9.2%	5.2%	23.1%
le-Gradual	$n = 2$	48.3%	11.3%	9.2%	15.1%	13.8%	1.4%	4.7%	8.5%
	$n = 3$	36.7%	3.3%	16.1%	12.8%	10.3%	4.2%	7.8%	12.1%
	$n = 4$	30.2%	0.08%	18.3%	11%	9.6%	6.3%	7.7%	16.9%

Note: This table reports the results of the classification of sender behavior in the "Gradual" treatments as defined in the text. To classify individual sender behavior in the first $n \in \{2, 3, 4\}$ rounds of a match requires discarding data from matches that last shorter than n rounds. The column labeled "According to random play" indicates the expected proportion of equilibrium behavior that would be generated if both players would make choices randomly.

Considering sender behavior in the first three rounds of matches that last three rounds or longer, the six equilibrium histories are the six most frequently observed among 196 possible outcomes.[24] Table 2 shows that the proportion of equilibrium behavior among all 3-round histories is about 48%. Again, we find that the prevalence of equilibrium behavior in the data is far from being random ($p < 0.001$). We also find that more than 66% of behavior is gradualist.

When we investigate sender behavior in the first four rounds of matches that last four rounds or longer, we find that about 38% of the observed behavior is consistent with equilibrium, whereas the proportion of equilibrium behavior generated by random play would be smaller than 1%. Finally, about 56% of behavior is gradualist.

So far, we have compared the empirical percentage of equilibrium behavior with the percentage that would be generated if behavior was random. However, we must emphasize that while in theory there are fewer equilibrium histories than non-equilibrium gradualist histories for any $n \geq 2$, the percentage of the former in the data dominates the percentage of the latter.[25] For instance, in the 3-round analysis, there are 6 equilibrium histories (as discussed above) and 10 non-equilibrium gradualist histories. As can be seen in Table 2, the empirical percentage of those 10 non-equilibrium histories is merely 19% compared to 48%, the empirical percentage of the equilibrium histories. We observe a similar situation with $n = 2$ and $n = 4$. Thus, senders clearly lean towards equilibrium behavior rather than non-equilibrium gradualist behavior.

Table 2 demonstrates that the percentage of behavior that we classify as equilibrium and gradualist declines discernibly as the number of rounds that we focus on increases. A few remarks regarding this observation are in order. First of all, we observe that lenient behavior becomes more and more prevalent at the later stages of a match. This is natural because the longer the match, the more likely default becomes. Thus, it also becomes more likely to observe leniency in punishment at a later stage. This implies that behavior which we classify as equilibrium or gradualist based on the first two or three rounds may end up being classified as lenient in a subsequent round if the match continues. This is indeed observed

[24] See the Online Appendix D.2.

[25] Recall that the gradualist behavior category has two distinct components: equilibrium and non-equilibrium.

in our data.[26] Second, it is entirely conceivable that even rational subjects sometimes make mistakes.[27] Taking into account a positive decision error probability in every round implies that the percentage of equilibrium behavior naturally declines over time. Finally, our design minimizes equilibrium randomization, which stacks the deck against theoretical predictions especially in longer matches (see also Footnote 12). In particular, we designed the game in a way that the low type defaults once M is chosen, and therefore, randomization between M and H is never part of equilibrium, which limits the number of equilibrium outcomes in matches that last three rounds or longer.[28]

We close this section with a discussion of the average realized payoffs of various sender strategies outlined above and listed in Table 2. To be more specific, we compute the average expected return to various sender strategies (up to the fourth round, similar to what we did in the categorization of sender behavior) given the observed receiver behavior in the experimental sessions. Overall, equilibrium behavior provides the highest expected return followed by non-equilibrium gradualist behavior. The worst strategies are "high trust until default" and "always no trust." We present a detailed description of the payoff calculations and comparisons in the Online Appendix E.

4.1.2 Behavior in the Binary Game

Recall that the Binary game is obtained from the Gradual game by removing the trust choices L and M, and that perpetual no trust is the unique equilibrium. Therefore, both the receiver and the sender are predicted to be strictly better off in the Gradual game than in the Binary game.

However, behavior in the Binary game diverges significantly from the "no trust" prediction, as a substantial proportion of senders choose to trust, and many receivers choose to return in the first round of a match even when they are experienced. After dropping the first 10 matches as we did in the Gradual game, the average first-round trust rate is 66.3%,

[26] On a related note, our data shows that leniency might be an effective disciplinary tool as it sometimes makes receivers more cooperative.

[27] This is the reasoning behind some bounded rationality models, such as the QRE by McKelvey and Palfrey (1995), which is a statistical version of the Nash equilibrium that incorporates decision errors.

[28] Thus, for example, choosing L in the first round, and M in both the second and the third is not consistent with equilibrium.

and the average first-round return rate is 49.5% including the computerized receivers and 42.6% excluding them.

We again observe a divergence from the theoretical prediction when we classify senders' behavior in the Binary game.[29] We detect in the data the following categories of behavior: "trust until default", "always no trust", "always trust", "lenient" and "hybrid". "Always trust" means that the sender starts by choosing trust and always trusts regardless of what the receiver has chosen in the past.[30] We define sender behavior as "hybrid" if the sender's behavior suggests that she is indifferent between "always no trust" and "trust until default." That is, hybrid behavior involves one round or a few rounds of no trust followed by "trust until default." Finally, we define sender behavior as "lenient" if the sender starts with trust, chooses trust as long as the receiver returns and punishes default with *at least* one period of no trust but the punishment is lenient as the sender starts trusting again.

The classification of sender behavior in the Binary game is presented in Table 3. The main feature of the data is the prevalence of trusting behavior. For example, looking at the first four rounds of a match, 55% of subject behavior is consistent with either "always trust" or "trust until default". In contrast, only 12.1% play the equilibrium strategy of never trusting, which means that 87.9% trust at least once in the first four rounds. Note that this deviation from the theoretical prediction in the Binary game stems from the receiver side. Senders who use the strategy "trust until default" appear to be best-responding to the empirical distribution of receiver behavior since our sender payoff analysis based on the empirical distribution of receiver choices—analogous to what we did for the Gradual game—shows that the strategy "trust until default" results in the highest expected sender payoff among the strategies described in Table 3. We will discuss these observations and related issues in detail in Section 4.3.

[29] The list of observed outcomes and the associated frequencies can be found in the Online Appendix D.1.

[30] Note that "trust until default" and "always trust" do not overlap in our classification since we include in the "always trust" category only histories in which the receiver defaults at least once and the sender still *never* chooses N.

Table 3: Classification of Sender Behavior in the Binary Games in the first n rounds, $n = 2, 3, 4$

Treatment	Matches that last at least n rounds	Always no trust	Trust until default	Always trust	Lenient	Hybrid	Unclassified
Binary	$n = 2$	21.7%	52.2%	13.1%	–	12.9%	0.1%
	$n = 3$	16.4%	45.7%	13.6%	3.5%	13.8%	7%
	$n = 4$	12.1%	43%	12.5%	4%	15.1%	13.3%
hr-Binary	$n = 2$	24.5%	60%	4.6%	–	7.65%	3.3%
	$n = 3$	17.9%	56.5%	5.8%	2%	10.5%	7.4%
	$n = 4$	17.3%	53.3%	3%	2.7%	9.8%	14%
le-Binary	$n = 2$	43%	44.5%	0.8%	–	10.6%	1.1%
	$n = 3$	38.2%	40.7%	1.4%	0%	16.4%	3.4%
	$n = 4$	35%	36.1%	0.8%	0.8%	20.8%	6.4%

Note: This table reports the results of the classification of sender behavior in the "Binary" treatments as defined in the text. To classify individual sender behavior in the first $n \in \{2, 3, 4\}$ rounds of a match requires discarding data from matches that last shorter than n rounds.

The payoff analysis and comparisons presented in the next section corroborates the evidence we provided here regarding the divergence of the behavior in the Binary game from the theoretical prediction. In particular, we show that the Binary game generates higher payoffs than the Gradual game. Hence, the Binary game is more conducive to cooperation and more efficient than the Gradual game in stark contrast with our prediction.

4.2 Payoff Comparisons

4.2.1 Payoff comparison with equilibrium predictions

Since the length of a match is random in our experiment, and the match length is an important factor that affects payoffs, we would like to control for the number of rounds played in each match when we compare theoretical predictions with data. To that aim, we calculate the payoff-index $\pi^{obs}/\pi^{eq}(r;\theta)$, where π^{obs} represents the observed match payoff, and $\pi^{eq}(r;\theta)$ represents the expected equilibrium match payoff given the true receiver type θ and the realized number of rounds in the match, denoted by r.[31] Note that if the payoff-index is lower (greater) than one, this implies that the average realized payoff is lower (greater) than the prediction. We compute three payoff indices: one for joint payoffs, one for sender payoffs and one for receiver payoffs.[32] Table 4 reports the summary statistics for these indices by treatment.

First, we discuss our findings in the Gradual game treatment. As Table 4 shows, joint payoffs in the Gradual game are close to but statistically slightly higher than the prediction. This is not surprising given the fact that players conform to equilibrium predictions to a significant extent in this game. To be more specific, the average sender payoff is only 4% higher than the equilibrium prediction, which is insignificant, whereas receivers are 18% better off, which is significant (see Table 4). This is mainly due to the following. On the one hand, a nonnegligible proportion of senders choose medium and high trust levels earlier than predicted by theory. On the other hand, low type receivers do not fully exploit this; i.e., they return at elevated rates relative to the equilibrium prediction. Since choosing higher

[31] Although there are multiple equilibria in the Gradual game, they are all very similar. We use equilibrium E (see Section 3.1) in our payoff index computation.

[32] When the sender is matched to a computer player, the joint-payoff index is computed on the basis of the sender payoff since the computer has no earnings.

Table 4: Summary Statistics of Payoff Indices by Treatment

Treatment	Joint-payoff Index		Sender-payoff Index		Receiver-payoff Index	
Gradual	$1.105^{***,***}$	(0.016)	1.038	(0.024)	$1.180^{***,***}$	(0.017)
Binary	$1.662^{***,***}$	(0.044)	$1.161^{,*}$	(0.067)	$2.819^{***,***}$	(0.147)
hr-Gradual	$1.038^{*,*}$	(0.035)	1.004	(0.111)	$1.10^{***,**}$	(0.070)
hr-Binary	$1.45^{***,***}$	(0.140)	1.031	(0.202)	$2.636^{***,***}$	(0.383)
le-Gradual	$1.035^{*,**}$	(0.028)	$1.045^{***,**}$	(0.023)	$1.048^{*,**}$	(0.056)
le-Binary	$1.322^{***,***}$	(0.128)	1.016	(0.071)	$2.133^{***,***}$	(0.331)

Notes: This table reports averages of session averages of the payoff index $\pi^{obs}/\pi^{eq}(r;\theta)$ as defined in the text. Standard errors are in parentheses. Results are shown in the form $x^{a,b}$, where x is the summary statistic, a represents the significance level according to the nonparametric test, and b denotes the significance level according to the regression-based analysis. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

trust levels earlier results in mixed success for the senders, their payoffs are not significantly different than the equilibrium prediction, whereas receivers do better than the prediction, albeit to a limited extent because they do not fully exploit senders' goodwill.[33]

In the Binary game treatment, joint payoffs are both statistically and economically significantly higher than predicted by the theory. This is because a substantial proportion of senders choose to trust and many receivers choose to return even when they are experienced, as discussed in Section 4.1.2. As a result, receivers earn a remarkable 182% more, and senders 16% more than the equilibrium predictions. Note that the latter is significant only according to our parametric approach and at 10% level (see Table 4). Thus, while the receivers' return rate is high, we do not have strong evidence that choosing to trust improves sender payoffs relative to the PBE prediction.

4.2.2 Payoff comparison across treatments

For the purpose of comparing payoffs across the two treatments, we calculate the efficiency index

$$\frac{\pi^{obs} - \pi^{N}(r)}{\pi^{C}(r) - \pi^{N}(r)},$$

[33]There is no significant time trend in either joint payoffs nor the role-specific payoffs if we focus on experienced play (see our regression analyses in the Online Appendix F).

where π^{obs} denotes the observed match payoff, r denotes the match length, $\pi^N(r)$ denotes the payoff if the static Nash equilibrium is played repeatedly in all r rounds (i.e., the sender always chooses no trust), and $\pi^C(r)$ denotes the payoff from repeated full cooperation (i.e., in all r rounds the sender always chooses high trust in the Gradual game and trust in the Binary game, and the receiver always returns). Since the Binary and the Gradual games have different equilibria, the indices that we used in the previous section are not comparable. This is why we generate an index that uses the full cooperation payoffs and the static Nash equilibrium payoffs, which are of course identical across the two treatments. This index also allows us to control for the match length, similar to what we did in the previous section.

We compute three efficiency indices: one for joint payoffs, one for sender payoffs and the last one for receiver payoffs. These indices are meant to measure the cooperativeness of the play. In particular, all three indices equal one if players coordinate on the most cooperative outcome whereas they will be equal to zero if static Nash play prevails.

Table 5 reports the summary statistics of the three indices by treatment. The joint-efficiency index is clearly below one but well above zero for both treatments. A further look at the respective indices for senders and receivers reveals the following pattern: While the sender-efficiency index is quite close to zero in both games, the receiver-efficiency index is around 0.75 in the Gradual game and slightly greater than 1 in the Binary game. This difference between the sender and the receiver indices implies that while there is indeed some cooperation, many receivers gain at the expense of senders in both games by eventually defaulting.

The first column in Table 5 shows that the joint-efficiency index is 12.3 percentage points higher in the Binary game than in the Gradual game, which contrasts with our prediction. The difference is statistically significant ($p < 0.05$ according to a Mann-Whitney test, and $p < 0.01$ according to our parametric approach).[34]

Some remarks regarding this finding are in order. First, our efficiency comparison may seemingly imply that equilibrium (or gradualist) strategies are futile or even harmful given that the Binary game generates higher payoffs. However, as discussed in Section 4.1.1,

[34]We regress the efficiency index on a binary treatment variable, controlling for the match number, the number of rounds in a match, as well as possible interaction effects. Further details and statistical results are relegated to the Online Appendix F.

Table 5: Summary Statistics of Efficiency Indices by Treatment

Treatment	Joint-efficiency Index	Sender-efficiency Index	Receiver-efficiency Index
Gradual Game	0.333 (0.020)	0.080 (0.019)	0.751 (0.032)
	$<^{**,***}$	$\approx$	$<^{**,**}$
Binary Game	0.456 (0.031)	0.121 (0.050)	1.014 (0.081)
hr-Gradual Game	0.266 (0.019)	0.035 (0.034)	0.647 (0.033)
	$\approx$	$\approx$	$<^{*,**}$
hr-Binary Game	0.312 (0.040)	0.023 (0.062)	0.878 (0.083)
le-Gradual Game	0.378 (0.015)	0.046 (0.029)	0.594 (0.025)
	$\approx$	$\approx$	$\approx$
le-Binary Game	0.383 (0.058)	0.033 (0.058)	0.631 (0.075)

Notes: This table reports averages of session averages of the efficiency index $(\pi^{obs} - \pi^N)/(\pi^C - \pi^N)$ as defined in the text. Standard errors are in parentheses. Test results are shown in the form $x^{a,b}$, where $x =$ "$<$" indicates that the summary statistic in the line above is smaller than the summary statistic in the line below, $x =$ "$\approx$" indicates that there is no significant difference at the 10%-level, a represents the significance level according to a binomial test, and b denotes the significance level according to the regression-based analysis. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

equilibrium behavior generates the highest expected payoff for the sender in the Gradual game given the empirical receiver behavior, followed by non-equilibrium gradualist behavior. In particular, the expected payoff to the sender from the strategy "high trust until default" is considerably low (see the payoff table in Online Appendix E).

Second, when we study the efficiency indices for senders and receivers, we see that the efficiency difference across the treatments stems from the fact that receivers are remarkably better off in the Binary game, whereas this is not the case for senders. Both a Mann-Whitney test and a regression-based analysis show that the receiver-efficiency index differs across the two treatments but the sender-efficiency index does not.[35] Thus, the joint-efficiency index is higher in the Binary game because receivers are significantly better off in the Binary game, while sender payoffs do not differ across the two games.[36]

[35] Regarding the comparison of the receiver-efficiency index across treatments, we find that $p < 0.001$ according to a Mann-Whitney test and $p < 0.05$ according to our regression-based analysis, whereas $p > 0.1$ according to both tests regarding the comparison of the sender-efficiency index.

[36] One might wonder what levels the efficiency indices presented in Table 5 might have converged to had the experiment been conducted over a very long time horizon. In Online Appendix C we present an analysis that uses techniques proposed in Noussair *et al.* (1995) or Barut *et al.* (2002) that suggest that none of the tests results provided in Table 5 would have changed.

Finally, the Gradual game provides senders with some "insurance": While sender payoffs do not differ across the two treatments, the variance of sender payoffs is significantly lower in the Gradual game than in the Binary game ($p < 0.001$ according to a Mann-Whitney test). Thus, the Gradual game makes risk-averse senders better off, whereas the Binary game makes receivers better off.[37]

4.3 Discussion

We now summarize our results so far and discuss the findings that informed the parameter choices in the follow-up treatments that will be introduced in the next section. In our Gradual game treatment, we find strong empirical support for Hypothesis 1 (that is, senders choose testing levels of trust consistent the equilibrium predictions) and also show that equilibrium strategies make senders better off than other strategies. However, when we compare the payoffs in the Binary game and the Gradual game, we find that the payoffs in the former are larger than those in the latter, in contrast to Hypothesis 2. In particular, we observe a lot of non-equilibrium trust and return behavior in the Binary game (as if problems of hidden information and opportunistic attitudes of receivers were largely absent) but find substantial evidence for equilibrium strategies in the Gradual game (as if senders were indeed keenly aware of hidden-information problems and used testing strategies to overcome them). This raises a puzzle. We argue that homemade beliefs and crowding in/out of trust are prime candidates to explain this puzzle.

Homemade Beliefs and Crowding in/out In relation to our observation that many senders start a match with trust in the Binary game even after they gain experience, we note that senders most likely have "homemade beliefs" about receivers that differ from the prior probability that we generated with the computer receivers (see Camerer and Weigelt (1988)). Senders' homemade belief may represent, for example, the percentage of human receivers who are altruistic so that they return rather than default. If senders have a homemade belief in addition to the controlled probability that they meet a computerized

[37] Recall that moderate levels of risk-aversion do not affect our theoretical predictions in the Gradual game and the Binary game.

receiver, then this changes the theoretical prediction in the Binary game more easily than in the Gradual game. In the Gradual game, senders must believe that at least 22.2% of human receivers behave like a high type in order to change our predictions qualitatively.[38] The analogous figure for the Binary game is much smaller at only 3.8%. Thus, the Gradual game is theoretically more robust to the effect of a homemade belief than the Binary game. Therefore, we ran follow-up treatments that vary the parameter values in a way that the Binary game prediction is more robust to the effect of a homemade belief than in the baseline treatment. We discuss these treatments in the next section.

Holding a homemade belief and starting a match with trust in the Binary game treatment seems to be empirically justified given the high return rate of human receivers. As discussed before, a sender payoff analysis based on the empirical distribution of receiver choices shows that the strategy "trust until default" results in the highest expected sender payoff among the strategies described in Table 3. Thus, the deviation in the Binary game from the theoretical prediction stems from the receiver side; senders who use the strategy "trust until default" appear to be best-responding to the empirical distribution of receiver behavior, and their homemade beliefs are justified.

Why do receivers return at such high rates in the Binary game? Does this stem from prosocial incentives of receivers? To a certain extent, the answer is yes: in the questionnaire which we gave after the experiment ended, we asked receivers in the Binary game whether they returned after trust was chosen and if they returned, why they did so. Many receivers referred to prosocial motivations and willingness to cooperate as the reason for returning trust. In addition, the findings in Engle-Warnick and Slonim (2004) put our Binary game results into perspective. They analyze an indefinitely repeated trust game similar to our Binary game, albeit with symmetric information (i.e., there is no private information regarding the type of the receiver). They implement a discount rate of 0.8 for both receivers and senders and as a result, there exists an equilibrium in their game such that the sender trusts and the receiver returns, unlike in our Binary game. Engle-Warnick and Slonim (2004) find that the return rate in the first round of a match is staggeringly high at above 92% with

[38] For example, with a homemade belief of 22.2% it would be optimal for the sender to start the match with M rather than L.

experienced players (i.e., in the last 10 matches).[39] Their findings point to the power of dynamic incentives in sustaining trust and cooperation. We contribute to the literature by showing that dynamic incentives constitute a strong discipline device in infinitely repeated sequential games that makes second-movers cooperate *even if* their effective discount rate is relatively low and cooperation is not an equilibrium strategy for them (but, naturally, the trust and return rates are lower in our experiment than those in Engle-Warnick and Slonim (2004)). Note that our results on the Binary game treatment are novel as no prior results on this indefinitely repeated "workhorse" model with incomplete information exist in the literature.

An important observation regarding receivers' behavior is the following. The receiver is in an identical situation—in monetary terms—if the sender chooses trust in the first round of a match in the Binary game and H in the first round of a match in the Gradual game. However, we find that the first-round return rate in the Binary game (49.3%) is higher than the first-round return rate in the Gradual game after H is chosen (28.4%).[40] Thus, it seems that the Binary game *crowds in* (and the Gradual game *crowds out*) trustworthiness. This poses an interesting puzzle because a sender who starts with a choice of H when L and M are available is being "nice" by not testing the receiver. Thus, prosocial receivers could be more willing to return if H is chosen in the Gradual game than when trust is chosen in the Binary game. But we observe the opposite. Importantly, this divergence in the first round return rates after H is chosen in the Gradual game and trust is chosen in the Binary game is not present in the first few rounds.[41] Hence, it appears that in these two different games, different norms develop over time as receivers have different experiences: In the Gradual

[39] In a similar vein, the first-round trust rate is very high in the last 10 matches (almost 90%).

[40] We focus only on the first round behavior to keep the analysis simple because after the first round, history effects come into play. The difference across the two treatments is statistically significant according to a two-sided Mann-Whitney at 10% level both in the whole game (28.35% in the Gradual game vs 49.32% in the Binary game) and after 10 matches are dropped (23.64% in the Gradual game vs 42.62% in the Binary game). A random effects panel probit regression of the probability that the receiver chooses to return in the first round of a match after trust (high trust) is chosen in the Binary (Gradual) game as a function of the treatment variable, the match number and interaction term confirms that receivers are more likely to return in the Binary game ($p < 0.001$ with complete data and $p < 0.01$ after the first 10 matches are dropped). The interaction term captures a possible interaction between the treatment variable and match number. Errors are always clustered at the session level.

[41] To support this statement formally, we compared return rates in the two treatments averaged over the first n and the remaining matches. Nonparametric tests indicate that there is no statistical difference for early matches, but a (weakly) significant difference for later matches when $n = 1, 2, ...5$.

game, receivers are tested most of the time, whereas they are trusted most of the time in the Binary game. Therefore, we conjecture that over time the two games cultivate different behavioral norms or preferences.

This conjecture is consistent with the findings of a large experimental literature showing that certain variations in institutions, the choice set, and other rules of a game may *crowd in* or *crowd out* prosocial behavior and trustworthiness (we present a detailed discussion regarding this in Section 5). The crowding in and crowding out phenomena are intrinsically related to homemade beliefs discussed above. That is, if the Binary game crowds in and the Gradual game crowds out trustworthiness, and senders have roughly correct homemade beliefs given receivers' behavior as the experimental data suggests, then these beliefs will differ across the two games. In particular, senders should have higher homemade beliefs in the Binary game, and as a result, the Binary game may be more conducive to cooperation than the Gradual game.

4.4 Follow-up Treatments

Is our finding that the Binary game is more efficient than the Gradual game robust? As discussed above, senders most likely have a homemade belief about receivers, in addition to the prior probability that we generated in the experiment. Such a belief is empirically justified since many receivers choose to return if the sender chooses trust in the Binary game. Moreover, our findings suggest that the Binary game crowds in trustworthiness, which is likely to boost senders' homemade beliefs and increase trust.

As we explained in the previous section, the Gradual game is theoretically more robust to the effect of homemade beliefs than the Binary game. In order to make the Binary game more robust to the effect of homemade beliefs, we conducted two pairs of follow-up treatments.

First set of follow-up treatments Our first follow-up experiment modified some of the sender payoffs as follows. If the sender chooses trust (high trust) in the Binary (Gradual) game and the receiver defaults, then the payoff of the sender is -20 (instead of 0, the baseline payoff). In order to make the sender payoffs more balanced across L, M and H in the Gradual

Figure 3: The game trees of the follow-up treatments "*hr*-Gradual" and "*hr*-Binary"

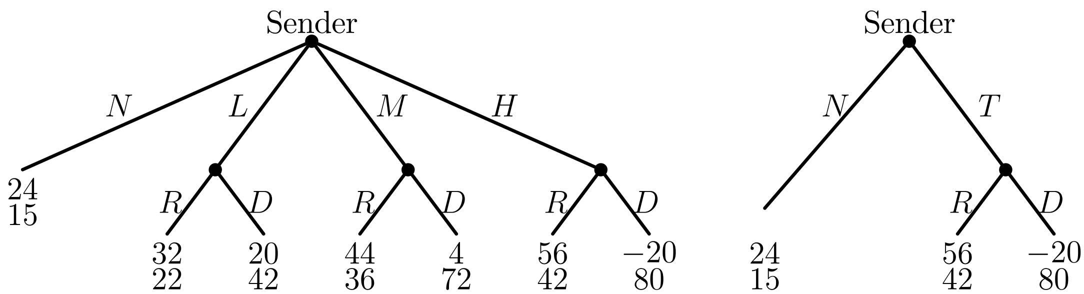

game, we made the following adjustment: If the sender chooses M and the receiver defaults, then the sender's payoff is 4 (instead of 12, the baseline payoff). We denote these new games as the *high-risk* Gradual game (hereafter, *hr*-Gradual game) and the *high-risk* Binary game (hereafter, *hr*-Binary game) and the new treatments, the *hr*-treatments. Figure 3 presents the new games.

Our equilibrium predictions are qualitatively unchanged across the baseline Gradual game and the *hr*-Gradual game: The sender starts the relationship with L, (weakly) increases the trust level as long as the receiver returns, and punishes default with no trust thereafter.[42] Of course, the expected equilibrium payoff for the sender is lower in the *hr*-Gradual game given the new parameter values. With these parameter values the Binary game is more robust to the effect of a homemade belief, as senders must believe that at least 15% of human receivers always return in order to change the equilibrium prediction to trust, compared to 3.8% in the original Binary game. We ran six sessions for each of the two new treatments, again conducted at the Vienna Center for Experimental Economics.

We now provide an overview of the behavior in the *hr*-treatments. As before, we drop the data from the first 10 matches in our analysis. Table 2 provides a classification of the sender behavior in the *hr*-Gradual game and reports the percentage of equilibrium behavior as well as other types of behavior, as we did in the baseline treatment. While the percentage of equilibrium sender behavior is higher in the baseline Gradual game than in the *hr*-Gradual game, this difference is not statistically significant. It must also be noted that the reduction in the sender payoff parameters in the *hr*-Gradual game translates to an

[42] In particular, the equilibrium description in Section 3.1 is still valid.

increased tendency among senders to choose N; that is, the strategy "always no trust" is more commonly chosen in the hr-Gradual game than in the baseline Gradual game.[43]

As in the baseline treatment, behavior in the hr-Binary game diverges significantly from our theoretical prediction of "always no trust". For example, a substantial proportion of senders choose to trust and many receivers choose to return in the first round of a match even after they gain experience. Table 3 presents the sender behavior classification in the hr-Binary game, which confirms the discrepancy between the equilibrium prediction and the observed behavior. The sender behavior is justified given the high receiver return rate in the data, just as in the baseline Binary game.

Regarding the efficiency comparison across the hr-Gradual game and the hr-Binary game, none of the results is remarkably different from the results in the baseline treatments. If anything, the payoff comparisons across the two types of games move slightly towards our equilibrium predictions. There is no statistically significant difference across the hr-Gradual game and the hr-Binary game in terms of sender payoffs or joint payoffs. Receivers still do better in the hr-Binary game than in the hr-Gradual game. However, the statistical evidence is weaker than in the baseline treatment—both a Mann-Whitney test and a regression-based analysis result in p-value < 0.1. Just as in the baseline treatment, we find that the variance of sender payoffs is significantly lower in the hr-Gradual game than in the hr-Binary game. This is also true with respect to joint payoffs.

Second set of follow-up treatments In both the baseline treatments and the hr-treatments, there are substantial efficiency gains from cooperation. Both players make large gains and the most efficient outcome is attained if the sender chooses to trust in the Binary game (high trust in the Gradual game) and the receiver returns. Such efficiency gains might be conducive to high cooperation. Our second set of follow-up treatments alters this feature of our previous treatments, in addition to making the Binary game setting more robust to the effect of a homemade belief. We expected a reduction in cooperation gains to have a large effect on the Binary game in the form of increased no trust choices, and a minimal effect in the Gradual game due to the presence of testing strategies.

[43] In the first round, the choice of N is more common in the hr-Gradual game than in the baseline ($p < 0.1$). Focusing on 2-round histories, always no trust is more commonly observed in the hr-Gradual game than in the baseline ($p < 0.05$). We do not detect a statistically significant difference in longer histories.

Figure 4: The game trees of the follow-up treatments "*le*-Gradual" and "*le*-Binary"

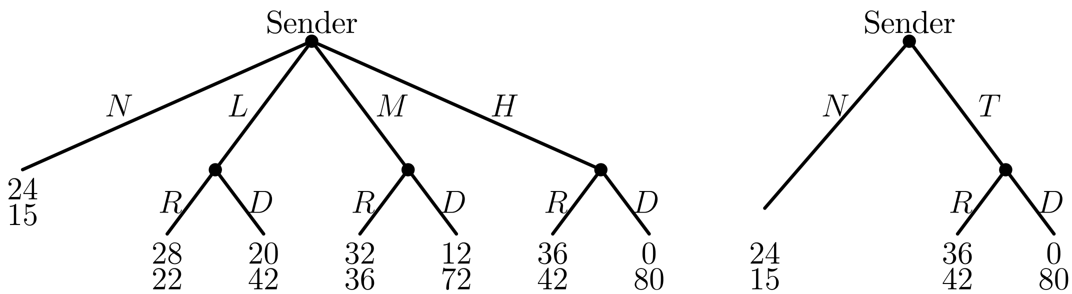

In order to reduce efficiency gains, we modified the sender payoffs as shown in Figure 4. If the sender chooses L in the Gradual game and the receiver returns, then the payoff of the sender is 28 rather than 32, the baseline payoff. If the sender chooses M in the Gradual game and the receiver returns, then the payoff of the sender is 32 rather than 44, the baseline payoff. Finally, if the sender chooses H in the Gradual game (trust in the Binary game) and the receiver returns, then the payoff of the sender is 36 rather than 56, the baseline payoff. Everything else is the same as in the baseline case. We denote these low-efficiency games as the *le*-Gradual game and the *le*-Binary game, and the treatments, *le*-treatments. With these parameter values, senders must believe that at least 20.8% of human receivers always return in order to switch the equilibrium prediction in the binary game from no trust to trust. We ran six sessions of each new treatment at the University of Konstanz, Germany.[44]

Overall, behavior in the *le*-Gradual game is quite similar to our previous Gradual game findings. Table 2 provides a classification of the sender behavior in the *le*-Gradual game. We note that the percentage of equilibrium sender behavior does not differ across any two Gradual games according to Mann-Whitney tests. Again, there is a significantly higher tendency among senders towards the strategy "always no trust" relative to the baseline, due to the reduction in the sender payoffs in the *le*-Gradual game.[45]

Behavior in the *le*-Binary game diverges significantly from our equilibrium prediction once again, despite reduced efficiency gains and despite making the Binary game more robust

[44] The change in subject pool is not problematic since our main focus is on comparisons between the *le*-treatments, and not with the earlier treatments.

[45] In the first round, the choice of N is more common in the *le*-Gradual game than in the baseline ($p < 0.01$). In 2-, 3-, and 4-round histories, always no trust is more commonly observed in the *le*-Gradual game than in the baseline Gradual game ($p < 0.01$).

to the effect of homemade beliefs. As before, a substantial fraction of senders choose to trust and many receivers choose to return in the first round of a match even after they gain experience. There is no statistically significant difference across the *le*-Gradual game and the *le*-Binary game in terms of sender payoffs, receiver payoffs, and joint payoffs. Hence, we move one step closer to our theoretical predictions but still find no evidence that the Gradual game improves efficiency. Finally, we again find that the variance of payoffs is significantly lower in the *le*-Gradual game than in the *le*-Binary game for sender payoffs and joint payoffs.

5 Concluding remarks

We study strategic behavior in a trust-game setting where (i) the time horizon is infinite; (ii) senders are uncertain regarding the trustworthiness of receivers; and (iii) there are various trust levels. We observe that this last feature of the game is essential to behavior as senders choose testing levels of trust in order to identify the "type" of the receiver they are matched with, and in turn, many receivers imitate the computerized player at low trust levels, as predicted. More generally, we find strong empirical support for Hypothesis 1, which predicts the use of testing strategies, and we also show that these equilibrium strategies make senders better off than non-equilibrium strategies in the Gradual games. These findings represent our first contribution to the literature. To our knowledge, our paper is the first and the only experimental study of cooperation and reputation formation in an infinitely repeated setting. Next, we compare the Gradual games to Binary games in which play is also infinitely repeated but trust is an all-or-nothing decision and senders are predicted to never trust. Note that a Binary game represents an institution that involves *only* dynamic incentives (i.e., default can be punished by no trust thereafter), whereas a Gradual game represents an institution with *both* dynamic and gradualist incentives (i.e., the sender can initially choose smaller trust levels and reward reciprocating receivers with increasing trust levels). In pairwise comparisons of the corresponding Gradual and Binary game treatments we never find the Gradual game to be more efficient than the Binary game (in contradiction to Hypothesis 2). This implies that infinitely repeated games, which already have dynamic incentives in place may not benefit from a gradualist "add-on" for improving trust and cooperation further.

In fact, the additional gradualist incentive may backfire and undermine efficiency as in our baseline treatments.

We now elaborate more on this point because testing strategies are not only theoretically appealing but also relevant in practice. For example, they are very common in microcredit. Since poor borrowers typically have no credit history or collateral, asymmetric information regarding their trustworthiness is an important concern. Therefore, microcredit institutions employ "progressive lending" and test borrowers with smaller loans at the beginning (Morduch (1999), Vogelgesang (2003)). To our knowledge, there has been no empirical study that evaluates the effectiveness of progressive lending, perhaps because progressive lending is an intuitive strategy in order to shield the organization against default. However, our findings suggest that evaluating the behavioral consequences of progressive lending in the field may be beneficial and especially so for non-profit organizations that aim to maximize borrower welfare subject to a break-even constraint.[46]

Our findings regarding the effects of gradual schemes may apply to other situations as well. For example, a principal who fears employee shirking, may choose to grant to a newly employed agent no decision power, limited decision power in matters of moderate importance, or more decision power in substantial matters. Our experimental findings then suggests that limiting the decision power of the new employees to a low level and to only mildly important matters may damage work morale. As dynamic incentives are already in place in employment relationships, granting decision power as an initial sign of trust may boost work morale and effort.

We argue that the presence of the intermediate trust levels L and M in the Gradual game and the fact that receivers are frequently tested by senders who choose these smaller trust levels crowd out trustworthiness. As discussed in Section 4.3, the return rate after T in the first round of a match in the baseline Binary game is higher than the return rate after H is chosen in the first round of a match in the baseline Gradual game. Moreover, this difference in receiver behavior develops over time due to their differing experiences in the two games. We observe the same significant divergence of receiver behavior in the hr-treatments.

[46] As discussed in Footnote 6, testing borrowers with very small loans may have not only intrinsic and but also extrinsic costs.

These findings connect our paper to a large literature which points out that certain variations in institutions, the choice set, and other details of the game can *crowd in* or *crowd out* trustworthiness, honesty and prosocial incentives. Gneezy *et al.* (2011) provide an extensive survey of the literature, and we will focus here on papers that are most related to our study. Bohnet *et al.* (2001) investigate a trust game in which the decision to return may be subject to legal enforcement with some probability. They find that receivers are less likely to betray the sender's trust when the enforcement probability is low than when it is medium. Thus, these authors conclude that trustworthiness is crowded in with weak and crowded out with medium enforcement. Falk and Kosfeld (2006) investigate a principal-agent setting in which the principal can control the agent by implementing a minimum effort requirement before the agent chooses the effort level. Most agents respond to principals who choose to control by reducing their effort. This means that exerting control entails "hidden costs" as it crowds out prosocial incentives of the agent to perform.[47] Also related is an observation by Andreoni and Miller (1993) and Anderhub *et al.* (2002). In these papers, the presence of a cooperative robot player appears to crowd out prosocial behavior by subjects.

What is the underlying crowding out mechanism in our design? We argue the following, in line with Falk and Kosfeld (2006), Sliwka (2007) and Gneezy et al. (2011). In our Binary game, testing strategies are unavailable, and senders frequently choose to trust the receiver. This behavior signals the receiver either that the behavioral norm is to return or that the sender trusts the receiver's intrinsic motivation to return (or both). As a result, trustworthiness is crowded in. In contrast, the frequent use of testing strategies in the Gradual game signals that the prevailing behavior (i.e., the norm) is to default and that the sender does not trust the receiver's intrinsic motivation to return, which is why testing strategies are being used. Thus, using gradualist strategies undermines receivers' prosocial motivations. Note that if the Binary game crowds in and the Gradual game crowds out trustworthiness, and senders have roughly *correct* homemade beliefs, then these beliefs will differ across the two games, and senders will (optimally) be more trusting in the Binary

[47] Falk and Kosfeld state that their results support the "self-fulfilling prophecy of distrust" because principals who choose to control the agent have more pessimistic beliefs regarding the agents' effort incentives than principals who trust. In turn, the reaction of the agents who are controlled by the principal confirms this belief.

game. Therefore, relationships in which incentives are predominantly dynamic may be more conducive to high cooperation than relationships that use a combination of dynamic and gradualist incentives.

However, another finding of our study is worth highlighting: While sender payoffs never differ across pairs of corresponding treatments, sender payoffs always have a lower *variance* in the Gradual game than in the corresponding Binary game as discussed before. Thus, the risk preferences of the sender is an important dimension to consider. If for example a sufficiently risk-averse lender would have to choose between a "Binary" or "Gradual" institution, this lender may prefer a gradualist one for risk management purposes despite the danger of crowding out. Hence, the endogenous choice of the institution is an important avenue for future research not only in the laboratory but perhaps also in the field.

5.1 References

Anderhub, V., D. Engelmann, and W. Güth. (2002). "An Experimental Study of the Repeated Trust Game with Incomplete Information." *Journal of Economic Behavior & Organization* 46, 197–216.

Andreoni, J. and J.H. Miller. (1993). "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence." *Economic Journal* 103 (418), 570–85.

Aoyagi, M. and G.R. Fréchette. (2009). "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence." *Journal of Economic Theory* 144 (3), 1135–1165.

Baker G., R. Gibbons, and K.J. Murphy. (2002). "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics* 117 (1), 39–84.

Barut, Y., D. Kovenock, and C. Noussair. (2002). "A Comparison of Multiple-Unit All-Pay and Winner-Pay Auctions Under Incomplete Information." *International Economic Review* 43, 675–707.

Berg, J., J. Dickhaut, and K. McCabe. (1995). "Trust, reciprocity, and social history." *Games and Economic Behavior*, 10(1), 122-142.

Bohnet, I., B.S. Frey, and S. Huck. (2001). "More order with less law: On contract enforcement, trust, and crowding." *American Political Science Review* 95, 131–144.

Brandts, J. and N. Figueras. (2003). "An exploration of reputation formation in experimental games." *Journal of Economic Behavior and Organization* 50 (1), 89–115.

Brown, M. and M. Serra-Garcia (in press): "The Threat of Exclusion and Implicit Contracting." *Management Science*, forthcoming.

Bull, C. (1987). "The Existence of Self-Enforcing Implicit Contracts." *Quarterly Journal of Economics,* 102 (1), 147–159.

Camera, C. and M. Casari. (2009). "Cooperation among Strangers under the Shadow of the Future." *American Economic Review* 99 (3), 979–1005.

Camerer, C. and K. Weigelt. (1988). "Experimental tests of a sequential equilibrium reputation model." *Econometrica* 56, 1–36.

Dal Bó, P. (2005). "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic Review* 95 (5), 1591–1604.

Dal Bó, P. and G. R. Frechette. (2011). "The evolution of cooperation in infinitely repeated games: Experimental evidence." *American Economic Review* 101, 411–429.

Duffy, J., and J. Ochs. (2009). "Cooperative behavior and the frequency of social interaction." *Games and Economic Behavior* 66(2), 785-812.

Embrey, M., G. R. Fréchette, and S. F. Lehrer. (2015). "Bargaining and reputation: An experiment on bargaining in the presence of behavioural types." *Review of Economic Studies* 82(2), 608-631.

Engle-Warnick, J. and R. Slonim. (2004). "The Evolution of Strategies in a Trust Game." *Journal of Economic Behavior and Organization* 55, 553–573.

Engle-Warnick, J. and R. Slonim. (2006a). "Learning to Trust in Indefinitely Repeated Games." *Games and Economic Behavior* 54, 95–114.

Engle-Warnick, J. and R. Slonim. (2006b). "Inferring Repeated Game Strategies from Actions: Evidence from Trust Game Experiments." *Economic Theory* 28, 603–632.

Falk, A. and M. Kosfeld. (2006). "The Hidden Costs of Control." *American Economic Review* 96 (5), 1611–1630.

Fischbacher, Urs. (2007). "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics,* 10(2), 171-178.

Fréchette, G.R. and S. Yuksel (in press): "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination." *Experimental Economics*, fothcomimg.

Fudenberg D, A. Dreber, and D.G. Rand. (2012). "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World." *American Economic Review* 102, 720–749.

Grosskopf, B. and R. Sarin. (2010). "Is reputation good or bad? An experiment." *American Economic Review* 100 (5), 2187–2204.

Ghosh, P. and Ray, D. (1996). "Cooperation in Community Interaction without Information Flows." *Review of Economic Studies* 63, 491–519.

Gneezy, U., S. Meier, and P. Rey-Biel. (2011). "When and Why Incentives (Don't) Work to Modify Behavior." *Journal of Economic Perspectives* 25 (4), 191–209.

Halac, M. (2012). "Relational Contracts and Value of Relationships." *American Economic Review* 102 (2), 750–779.

Huck, S., G. S. Lünser, and J.-R. Tyran (2016). "Price competition and reputation in markets for experience goods: an experimental study." *RAND Journal of Economics* 47 (1), 99–117

Kartal, M. (in press): "Honest equilibria in reputation games: The role of time preferences." *American Economic Journal: Microeconomics*, forthcoming.

Kranton, R. (1996). "The Formation of Cooperative Relationships" *Journal of Law, Economics and Organization* 12, 214–233.

Kreps, D., P. Milgrom, J. Roberts, and R. Wilson. (1982). "Rational cooperation in the finitely repeated prisoners' dilemma" *Journal of Economic Theory* 27 (2), 245–252.

Kreps, D. and R. Wilson. (1982). "Reputation and Imperfect Information." *Journal of Economic Theory* 27 (2), 253–279.

Levin, J. (2003). "Relational Incentive Contracts." *American Economic Review* 93 (3), 835–57.

MacLeod, W. Bentley, and J.M. Malcomson. (1989). "Implicit contracts, incentive compatibility, and involuntary unemployment." *Econometrica* 57, 447–480.

McKelvey, R. and T. Palfrey. (1995). "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10, 6–38.

Milgrom, P. and J. Roberts. (1982). "Predation, Reputation, and Entry Deterrence." *Journal of Economic Theory* 27 (2), 280–312.

Morduch, J. (1999). "The microfinance promise." *Journal of Economic Literature* 37(4), 1569-1614.

Neral, J. and J. Ochs. (1992). "The sequential equilibrium theory of reputation building: A further test." *Econometrica* 60, 1151–1169.

Noussair, C., C. Plott, and R. Riezman. (1995). "An Experimental Investigation of the Patterns of International Trade." *American Economic Review* 85(3), 462–91.

Palfrey, T. R. and H. Rosenthal. (1994). "Repeated Play, Cooperation and Coordination: An Experimental Study." *Review of Economic Studies* 61 (3), 545–65.

Sliwka, D. (2007). "Trust as a Signal of a Social Norm and the Hidden Costs of Incentive Schemes." *American Economic Review* 97 (3), 999–1012.

Sobel, J. (1985). "A Theory of Credibility." *Review of Economic Studies* 52 (4), 557–573.

Vogelgesang, U. (2003). "Microfinance in times of crisis: The effects of competition, rising indebtedness, and economic crisis on repayment behavior." *World Development*, 31(12), 2085-2114.

Watson, J. (1999). "Starting Small and Renegotiation." *Journal of Economic Theory* 85, 52–90.

Watson, J. (2002). "Starting Small and Commitment." *Games and Economic Behavior* 38, 176–99.

A Omitted Theory Materials

A.1 A General Model of Trust with Private Information

Two players, a sender and a receiver, interact repeatedly in periods $t = 0, 1, \ldots$. At the beginning of period $t \geq 0$, the sender makes a monetary transfer to the receiver by choosing a trust level, denoted by m_t, where $m_t \in [0, \bar{m}]$. Thus, the sender can choose from a continuum of positive trust levels and also has the option not to trust.

If the sender chooses $m_t > 0$, then the receiver chooses between "return" and "default." If $m_t > 0$ and the receiver chooses to *return*, then $r_S(m_t)$ and $r_R(m_t)$ denote the respective period-t payoffs of the sender and the receiver, whereas $d_S(m_t)$ and $d_R(m_t)$ denote the respective period-t payoffs of the sender and receiver if the receiver chooses to *default*. If the sender chooses $m_t = 0$, then the receiver has no choice to make, and both players end up with their outside options, denoted by $\bar{u}_S$ and $\bar{u}_R$ for the sender and the receiver, respectively. We assume that all the payoff functions are continuous at every $m \in (0, \bar{m}]$. Other key assumptions on the payoff functions are as follows:

- For every $m > 0$, $r_S(m) > d_S(m)$, whereas $d_R(m) > r_R(m)$. Thus, the sender and the receiver have a conflict of interest in the one-shot game: the sender is better off if the receiver returns whereas the receiver is better off if he defaults.
- If $m' > m$, then $r_S(m') > r_S(m)$ and $r_R(m') > r_R(m)$; that is, the payoff from mutual cooperation is strictly increasing for both parties.
- If $m' > m$, then $d_S(m) > d_S(m')$; that is, the higher the trust level, the lower the payoff of the sender if the receiver defaults.
- If $m' > m$, then $d_R(m') > d_R(m)$; that is, the higher the trust level, the higher the payoff of the receiver from default.

Hence, the stage game payoffs are in line with the typical trust game à la Berg *et al.* (1996). Our game involves one-sided private information. The sender's discount factor is $\delta > 0$, which is fixed and known, whereas the receiver's discount factor is δ_θ, where $\theta \in \{l, h\}$ is the receiver's private information, and $\delta_l < \delta_h$. The high type receiver is cooperative in the following sense: δ_h is sufficiently high so that for every $m \in (0, \bar{m}]$, there exists an equilibrium in the symmetric information game such that the sender always chooses $m > 0$ and the high type receiver returns, that is

$$\frac{r_R(m)}{1-\delta_h} > d_R(m) + \frac{\delta_h}{1-\delta_h}\bar{u}_R \tag{1}$$

holds for every $m \in (0, \bar{m}]$. However, δ_l is assumed to be such that

$$\frac{r_R(m)}{1-\delta_l} < d_R(m) + \frac{\delta_l}{1-\delta_l}\bar{u}_R \tag{2}$$

holds for every $m \in (0, \bar{m}]$. Thus, the low type receiver is myopic, and trust is never an equilibrium outcome in the symmetric information game with a low type. Even if (2) holds, the low type may strategically choose to return in the hidden information game in order to imitate the high type; this will be the case if returning in the current period will be reciprocated with a higher trust level in the future and defaulting at that higher trust level compensates the loss from returning in the current period.

We now study the perfect Bayesian Nash Equilibria (PBE) of this game. Let p_0 denote the prior belief of the sender that the receiver is low type and let p_t denote the posterior belief at the beginning of period t. Consider a "cooperative equilibrium" which we define as an equilibrium in which $m_t > 0$ at some $t \geq 0$, and the high type returns with probability one whenever $m_t > 0$ on the equilibrium path. A cooperative equilibrium surely exists if p_0 is not too low. Then, in every pareto-efficient equilibrium, the receiver's private information will be fully revealed in finite time; i.e., there exists a $T < \infty$ such that $p_t \in \{0, 1\}$ for all $t \geq T$. Moreover, the Pareto-efficient equilibrium prescribes $m_t = \bar{m}$ for all sufficiently large t provided that the sender is matched with a high type receiver (the proofs of these claims can be found below).

For our main result we assume that the sender follows a Markovian trust strategy because not only the set of PBE but also the set of Pareto-efficient PBE tend to be very large in repeated games with hidden information. In particular, we assume that the sender's trust strategy depends only on his belief p_t.[48] Our main result is as follows.

Proposition 1 *There is a $\underline{p} \geq 0$ such that a cooperative equilibrium exists for every $p_0 > \underline{p}$. Given $p_0 > \underline{p}$ every pareto-efficient equilibrium under the Markovian sender strategy assumption is such that either (i) $m_t = \bar{m}$ at every $t \geq 0$ or (ii) the equilibrium is gradualist; i.e., there is a finite $T > 0$ such that $m_t < m_{t+1}$ for every $t \in \{0, ..., T-1\}$ and $\bar{m}$ is chosen from T onwards as long as default is not observed. Moreover, there exists a nonempty, nonsingleton interval $(\underline{p}, \bar{p})$ such that every pareto-efficient equilibrium under the Markovian assumption must be gradualist for $p_0 \in (\underline{p}, \bar{p})$.*

[48] The receiver's strategy is not contrained to be Markovian.

Proof. First we will establish Claims 1 and 2 below.

Claim 1 *Suppose that a cooperative equilibrium exists. Then, in the pareto-efficient equilibrium, there exists a finite T such that $p_t \in \{0,1\}$ holds for all $t > T$; thus, equilibrium path beliefs are degenerate for all sufficiently large t.*

Proof. Let $m^* = \sup_{t \geq 0} m_t$, where m_t represents the equilibrium trust level chosen by the sender given that the receiver returned in all of $t-1$ previous periods. Either there exists $T < \infty$ such that $m_T = m^*$ or there exists a sequence $\{m_t\}_{t \geq 0}$ where $m_t < m^*$ for all t. This sequence has a convergent subsequence, which we also denote by $\{m_t\}_{t \geq 0}$ without loss of generality. Thus, $\{m_t\}_{t \geq 0}$ is such that $\lim_{t \to \infty} m_t = m^*$. In the first case in which there exists a finite T such that $m_T = m^*$, it is obvious that the low type receiver will have defaulted with probability one by the end of period T, and thus, $p_t \in \{0,1\}$ holds for all $t > T$. In the second case, there exists a T such that the sender chooses m_T if the receiver returned in all of $T-1$ previous periods, and m_T is such that

$$d_R(m_T) + \frac{\delta_R}{1-\delta_R}\bar{u}_R > r_R(m_T) + \delta_R d_R(m^*) + \frac{\delta_R^2}{1-\delta_R}\bar{u}_R$$

holds. Note that such $T < \infty$ exists by condition (2), continuity of $r_R(m)$ and $d_R(m)$ in m and due to the fact that $\lim_{t \to \infty} m_t = m^*$. This in turn implies that $p_t \in \{0,1\}$ for all $t > T$. Hence, the claim is proved.

Claim 1 holds regardless of the Markovian sender strategy assumption, and its proof does not depend on the assumption. Claim 2 also holds regardless of the Markovian assumption. However, its proof is less straightforward under the Markovian assumption, therefore we prove it assuming that the sender uses a Markovian trust strategy.

Claim 2 *If a cooperative equilibrium exists, then in the pareto-efficient equilibrium, $m_t = \bar{m}$ for all sufficiently large t provided that the receiver has never defaulted.*

Proof. Let $\{m_t\}_{t \geq 0}$ denote the equilibrium trust level sequence given that default is never observed. Assume that $m_t < \bar{m}$ for every $t \geq 0$, otherwise, the problem is trivial. Let $T-1$ be the smallest $t \geq 0$ such that the low type receiver chooses to default with probability one on the equilibrium path of the pareto-efficient equilibrium. Thus, $p_t \in \{0,1\}$ holds for every $t > T$. By Claim 1, T is a finite number. If the receiver has always returned until period T, then by our Markovian assumption and initial hypotheses, $0 < m_T = m_{T+1} = ... = \hat{m} < \bar{m}$. We now define a new sequence $\{\tilde{m}_T, \tilde{m}_{T+1}, ...\}$ as follows. Let $\tilde{m}_T$ be such that

$$r_R(m_T) + \delta_l d_R(\tilde{m}_T) + \frac{\delta_l^2}{1-\delta_l}\bar{u}_R = d_R(m_T) + \frac{\delta_l}{1-\delta_l}\bar{u}_R.$$

Note that the above equality can only hold with $\tilde{m}_T > m_T$. If the equality is not possible with any $\tilde{m}_T \le \bar{m}$, then we set $\tilde{m}_T = \tilde{m}_{T+1} = ... = \bar{m}$, and we are done. Our construction is an equilibrium, and both the sender and the high type receiver are strictly better off as the low type is still willing to default with probability one at T. If $\tilde{m}_T < \bar{m}$, then let $\tilde{m}_{T+1}$ be such that

$$r_R(\tilde{m}_T) + \delta_l d_R(\tilde{m}_{T+1}) + \frac{\delta_l^2}{1-\delta_l}\bar{u}_R = d_R(\tilde{m}_T) + \frac{\delta_l}{1-\delta_l}\bar{u}_R.$$

If the equality is not possible with any $\tilde{m}_{T+1} \le \bar{m}$, then we set $\tilde{m}_{T+1} = \tilde{m}_{T+2} = ... = \bar{m}$. If $\tilde{m}_{T+1} < \bar{m}$, then the procedure is analogous for $\tilde{m}_{T+2}$, and so on. In a finite K number of steps, the equality will no longer be possible for $\tilde{m}_{T+K-1}$ and $\tilde{m}_{T+K}$, and we will set $\tilde{m}_{T+K} = \tilde{m}_{T+K+1} = ... = \bar{m}$. To see why such finite K exists, suppose towards a contradiction that the sequence $\{m_0, m_1, ..., \tilde{m}_T, \tilde{m}_{T+1}, ...\}$ is such that the low type receiver is always indifferent between returning and defaulting. But this is impossible similar to what we argued in the proof of Claim 1. The sequence $\{\tilde{m}_T, \tilde{m}_{T+1}, ...\}$ is increasing and bounded, and therefore, it converges to some limit $m^* \le \bar{m}$. As a result, for sufficiently large K, $\tilde{m}_{T+K-1}$ and $\tilde{m}_{T+K}$ will be too close because they are both very close to m^*. Thus, by condition (2), the continuity of $r_R(m)$ and $d_R(m)$ in m and due to the fact that $\tilde{m}_t$'s are very close to m^* for sufficiently large t, there is a sufficiently large but finite K such that the low type strictly prefers defaulting when $\tilde{m}_{T+K-1}$ is chosen, a contradiction. Thus, there is a finite K such that $\tilde{m}_{T+K} = \tilde{m}_{T+K+1} = ... = \bar{m}$. Finally, for $K \ge 1$ we will complete the construction by setting the low type receiver's return rate to $\frac{1}{N}$ at every $t \in \{T-1, ..., T+K-2\}$ and to 0 at period $T+K-1$, where N is a large integer as we will make more clear below (for $K = 0$, $\tilde{m}_T = \bar{m}$ the desired result already follows as argued above). We need this positive but very small low type return rate until period $T+K-1$ for $K \ge 1$ due to the Markovian assumption, as this enables the posterior p_t to be extremely close to 1 but still increase at every $t \in \{T, T+1, ...T+K-1\}$ until $p_{T+K} \in \{0, 1\}$ (given the Markovian assumption, the sender has to choose the same trust level if the posterior remains the same). Note that by the way $\tilde{m}_t$'s and K are defined, the low type is indifferent between returning and defaulting at every $t \in \{T-1, ..., T+K-2\}$. Moreover, the high type strictly prefers returning as before.[49] The low type is indifferent whereas the high type is strictly better off in this construction. Moreover, it can easily be checked that if N is sufficiently large, then the sender is strictly better off and would not want to deviate from the choice of $\tilde{m}_t$ at any $t \ge T$. Hence, the claim follows.

Now we will prove the proposition. First, assume that a cooperative equilibrium exists and that $m_0 < \bar{m}$ in every pareto-optimal equilibrium (below, we will prove that if p_0 is not sufficiently high, then $m_0 < \bar{m}$ must hold in every cooperative equilibrium). We will show

[49] Naturally, an equilibrium in which the high type receiver randomizes between returning and defaulting is strictly inefficient and can be improved upon given condition (1).

that there exists a $T < \infty$ such that $m_t < m_{t+1}$ for every $t \in \{0, ..., T-1\}$ until $m_T = \bar{m}$, as long as default is not observed. That there is a finite T such that $m_t = \bar{m}$ for all $t \geq T$ in the pareto-efficient equilibrium follows from Claim 2. So, set $T = \min\{t \geq 0 | m_t = \bar{m}\}$. Thus, by definition of T, $m_{T-1} < m_T$ must hold. Next, we will show that $m_{T-2} < m_{T-1}$, and that $m_t < m_{t+1}$ for $t \in \{0, ..., T-3\}$ follows by induction. Suppose towards a contradiction that $m_{T-2} \geq m_{T-1}$. It follows that on the equilibrium path, the low type receiver will never default at $T-1$, since it is strictly better to default either at $T-2$ or at T. However, both are contradictory. In either case, it follows that $p_{T-1} = p_T$ (as the low type will not default at $T-1$ provided that he returned at $T-2$) and thus, $m_{T-1} = m_T$ must hold, which is a contradiction because $m_{T-1} < m_T$ by the definition of T. Hence, $m_{T-2} < m_{T-1}$ must hold. Now, assume that m_t is strictly increasing for $t \in \{\tau+1, ..., T-1\}$ by the induction hypothesis. We will show that $m_\tau < m_{\tau+1}$. Suppose towards a contradiction that $m_\tau \geq m_{\tau+1}$. It follows that on the equilibrium path, the low type receiver will never default at $\tau + 1$ since it is strictly better to default either prior to or after $\tau + 1$. However, this implies that $p_{\tau+1} = p_{\tau+2}$ (as the low type will not default at $\tau+1$ provided that he returned at τ) and thus, $m_{\tau+1} = m_{\tau+2}$ must hold, in contradiction to our initial induction hypothesis. Hence, $m_\tau < m_{\tau+1}$. It follows that $m_t < m_{t+1}$ for every $t < T$.

Next, we show that a cooperative equilibrium exists if p_0 is sufficiently high. To see why, note that there exists a $\tilde{p} < 1$ such that

$$\tilde{p}\left(d_S(\bar{m}) + \frac{\delta}{1-\delta}\bar{u}_S\right) + (1-\tilde{p})\frac{r_S(\bar{m})}{1-\delta} = \frac{\bar{u}_S}{1-\delta}$$

holds. Thus, a cooperative equilibrium exists for all $p_0 \geq \tilde{p}$. Even for some p_0 values such that $p_0 < \tilde{p}$ there exists cooperative equilibria, and in this case, pareto-efficient equilibria can only be gradualist as $m_0 < \bar{m}$ must hold. In particular, there exists a possibly large but finite N such that for every $p_0 \in (\tilde{p} - \frac{1}{N}, \tilde{p})$ pareto-efficient equilibria are gradualist. To see why, consider the following simple construction. Take $m_0 < \bar{m}$ such that if m_0 is chosen at $t = 0$ then the low type receiver is indifferent between returning and defaulting at $t = 0$ provided that $\bar{m}$ will be chosen at $t = 1$ after the receiver returns; i.e., m_0 is such that

$$r_R(m_0) + \delta_R d_R(\bar{m}) + \frac{\delta_R^2}{1-\delta_R}\bar{u}_R = d_R(m_0) + \frac{\delta_R}{1-\delta_R}\bar{u}_R.$$

Next, let $\gamma \in (0,1)$ denote the probability with which the low type returns at $t = 0$. Note that for $p_0 = \tilde{p}$ there exists a $\tilde{\gamma} < 1$ such that for all $\gamma > \tilde{\gamma}$, there exists a cooperative equilibrium in which the sender's expected payoff is *strictly* greater than $\frac{\bar{u}_S}{1-\delta_S}$. It follows that there exists a sufficiently large but finite N, such that cooperative equilibria exist and all pareto-efficient equilibria are gradualist for $p_0 \in (\tilde{p} - \frac{1}{N}, \tilde{p})$. Note that adding more intermediate trust steps will expand the possibilities of cooperation with even lower values of p_0.

A.2 Equilibria of the experimental game

As in the main text, the "receiver" refers to the low type receiver throughout this section, unless otherwise noted. We first formally describe equilibrium E, which was explained in the main text and show that it is indeed an equilibrium. To that aim, let $S_t \in \{N, L, M, H\}$ represent the action of the sender and let $R_t \in \{R, D\}$ represent the action of the receiver at $t \geq 0$ provided that $S_t \neq N$. Following, $h_t \in N \cup \{L, M, H\} \times R_t$ denotes the outcome at time t, and $h^t = (h_0, ..., h_t) \in \mathcal{H}^t$ denotes the outcome up to (and including) period t, where $\mathcal{H}^t$ represents the set of all possible date t histories with the understanding that $\varnothing \in \mathcal{H}^t$. Finally, $\mu(h^t)$ denotes the posterior belief that the receiver is low type given the outcome h^t, and $\mu_o \equiv \mu(\varnothing)$ represents the prior belief. Next, we define the strategies and beliefs associated with equilibrium E as follows.

$$\sigma_S(h^{t-1}) = \begin{cases} N & \text{if} \quad \mu(h^{t-1}) > \mu_o = 7/8 \\ \\ L & \text{if} \quad \mu(h^{t-1}) \in (21/32, 7/8] \text{ or if } \mu(h^{t-1}) = 21/32 \\ & \quad\;\; \text{and } h_{t-1} \neq (L, R) \\ L \sim M(34/43) & \text{if} \quad \mu(h^{t-1}) = 21/32 \;\&\; h_{t-2} \neq h_{t-1} = (L, R) \\ & \quad\;\; \text{or if } t = 1 \;\&\; h_0 = (L, R)) \\ M & \text{if} \quad \mu(h^{t-1}) \in (0.5, 21/32) \text{ or if } \mu(h^{t-1}) = 21/32 \\ & \quad\;\; \text{and } h_{t-2} = h_{t-1} = (L, R) \\ H & \text{if} \quad \mu(h^{t-1}) \leq 0.5, \end{cases}$$

where $X \sim Y(p)$ represents a mixed strategy such that X is chosen with probability p and Y with probability $(1-p)$ (in a similar vein, $X \sim Y$ represents a mixed strategy without specifying the probability of choosing X).

$$\sigma_R(h^{t-1}, S_t) = \begin{cases} R & \text{if} \quad S_t = L \text{ and } \mu(h^{t-1}) \leq 21/32 \\ R \sim D & \text{if} \quad S_t = L \text{ and } \mu(h^{t-1}) \in (21/32, 7/8] \\ D & \text{if} \quad S_t \in \{M, H\} \text{ or if } \mu(h^{t-1}) > 7/8 \end{cases}$$

where R in the mixed strategy $R \sim D$ is chosen with probability $\gamma(h^{t-1}, L)$ in such a way that $\mu(h^t) = 21/32$ holds; i.e.,

$$\frac{\mu(h^{t-1})\gamma(h^{t-1}, L)}{\mu(h^{t-1})\gamma(h^{t-1}, L) + 1 - \mu(h^{t-1})} = \frac{21}{32}.$$

Finally,

$$\mu(h^t) = \begin{cases} 1 & \text{if} \quad h_t = (S,d) \text{ where } S \in \{L,M,H\} \\ \mu(h^{t-1}) & \text{if} \quad h_t = N \text{ or } \mu(h^{t-1}) > 7/8 \text{ and } h_t = (S,R) \\ & \qquad \text{where } S \in \{L,M,H\} \\ \min\{\mu(h^{t-1}), 21/32\} & \text{if} \quad \mu(h^{t-1}) \leq 7/8 \text{ and } h_t = (L,R) \\ 0 & \text{if} \quad \mu(h^{t-1}) \leq 7/8 \text{ and } h_t \in \{(M,R),(H,R)\} \end{cases}$$

Notice that for consistency in belief updating, default *must* have taken place at some $\tau \leq t$ if $\mu(h^t) > 7/8$ because the high type receiver is a commitment type who always returns. Thus, $\mu(h^t) > 7/8$ automatically implies that $\mu(h^t) = 1$.

After proving that equilibrium E exists, we will prove that it is efficient and that every other PBE is very similar to E. To show that the sender and receiver strategies and the belief system specified above constitute an equilibrium, we check that it is immune to deviations for both the sender and the receiver.

First, consider the receiver side. We start by showing that whenever M or H is chosen the receiver chooses D with probability one. Note that choosing D once M or H is chosen is strictly optimal for the low type receiver regardless of the sender strategy and belief specification. To show this, it is enough to consider the decision of the receiver after M is chosen at an arbitrary $t \geq 0$, which will be followed by H at $t+1$ if the receiver chooses R after M is chosen at t (this enables the receiver to obtain the highest possible payoff from the return decision at t). From the parameters used in the experiment, it follows that $72 + 15 > 36 + 40 + \frac{15}{2}$. Thus, the receiver strictly prefers defaulting after M is chosen. This is also true if H is chosen; i.e., the receiver strictly prefers defaulting whenever H is chosen since $80 + 15 > 42 + 40 + \frac{15}{2}$ and 95 is the highest possible payoff that the low type receiver can achieve.

Next, we verify that given the above description of sender strategy and beliefs, the receiver is indifferent between R and D after L is chosen at $t = 0$ or at arbitrary $t > 0$ off-the-equilibrium path either with $\mu(h^{t-1}) \in (21/32, 7/8]$ or with $\mu(h^{t-1}) = 21/32$ and $h_{t-1} \neq (L,R)$. If the receiver chooses R at t, the sender strategy specifies randomization at $t+1$ such that L is chosen with probability $34/43$ and M with probability $9/43$. If L is the result of this sender randomization at $t+1$ and the receiver returns, then the sender chooses M at $t+2$. It can easily be checked that this leaves the receiver indifferent between R and D at t given our experimental parameters. Thus, randomization between R and D at t as prescribed by $\sigma_R(h^{t-1}, L)$ is feasible.

We now verify that the receiver strictly prefers returning after L is chosen at $t = 1$ with $h^0 = h_0 = (L,R)$ (or at $t > 1$ off-the-equilibrium path with $\mu(h^{t-1}) = 21/32$ and

$h_{t-2} \neq h_{t-1} = (L, R)$). This follows simply from the specification that $\sigma_S(h^{t-1}) = M$ for $\mu(h^{t-1}) = 21/32$ and $h_{t-2} = h_{t-1} = (L, R)$, and the inequality $22 + \frac{1}{2}72 + \frac{15}{2} > 42 + 15$.

Next, we analyze the sender side. First, we show that there exists a mixing probability $\gamma \equiv \gamma(\varnothing, L) \in (0, 1)$ for the receiver at $t = 0$ such that given the initial prior μ_o and the mixing probability γ, the sender strictly prefers choosing L at $t = 0$ and is indifferent between L and M at $t = 1$ provided that the receiver chooses R at $t = 0$ (using the same arguments as below, it can be argued that there exists a mixing probability $\gamma(h^{t-1}, L) \in (0, 1)$ off-the-equilibrium path at $t > 0$ with $\mu(h^{t-1}) \in (21/32, 7/8]$ or with $\mu(h^{t-1}) = 21/32$ and $h_{t-1} \neq (L, R)$ such that given the posterior $\mu(h^{t-1})$ and the mixing probability $\gamma(h^{t-1}, L)$, the sender strictly prefers choosing L at t and is indifferent between L and M at $t + 1$ provided that the receiver chooses R at t). We find the mixing probability γ as follows. Let

$$\mu_1 \equiv \mu(h^0) = \frac{\mu_o \gamma}{\mu_o \gamma + 1 - \mu_o}$$

where $h^0 = (L, R)$. Then, γ and μ_1 can be pinned down using above equation and the indifference condition

$$\frac{32}{1 - \delta_S} = \mu_1 (12 + \frac{\delta_S}{1 - \delta_S} 24) + (1 - \mu_1)(44 + \frac{\delta_S}{1 - \delta_S} 56)$$

where $\delta_S = 0.75$. It follows that $\mu_1 = \frac{21}{32}$ and $\gamma = 0.273$. The last equality ensures that the sender is indifferent between L and M at $t = 1$ (there is also a recursive element in this equality that prevents the sender from deviating to L at $t = 2$ following $h^1 = ((L, R), (L, R))$ instead of choosing M). We can verify that L is strictly preferred over M and H at $t = 0$ given our experimental parameters and given $\gamma = 0.273$:

$$\mu_o \left[\gamma \left(32 + \delta_S 12 + \frac{\delta_S^2 24}{1 - \delta_S} \right) + (1 - \gamma) \left(20 + \frac{\delta_S 24}{1 - \delta_S} \right) \right] + (1 - \mu_o) \left(32 + \delta_S 44 + \frac{\delta_S^2 56}{1 - \delta_S} \right) \geq$$

$$\max\{\mu_o \left(12 + \frac{\delta_S 24}{1 - \delta_S} \right) + (1 - \mu_o) \left(44 + \frac{\delta_S 56}{1 - \delta_S} \right), \mu_o \left(0 + \frac{\delta_S 24}{1 - \delta_S} \right) + (1 - \mu_o) \left(\frac{56}{1 - \delta_S} \right) \}$$

where $\delta_S = 0.75$. Note that H is dominated at $t = 1$ or $t = 2$ (even if there is no default at $t = 0$ and $t = 1$) because, given $\mu_1 = 21/32$ it is suboptimal to choose H over M (or L). This is true since $\mu_1 = \frac{21}{32}$ implies that

$$\mu_1 \left(12 + \frac{\delta_S 24}{1 - \delta_S} \right) + (1 - \mu_1) \left(44 + \frac{\delta_S 56}{1 - \delta_S} \right) > \mu_1 \left(0 + \frac{\delta_S 24}{1 - \delta_S} \right) + (1 - \mu_1) \left(\frac{56}{1 - \delta_S} \right)$$

where $\delta_S = 0.75$. In a similar vein, it can be checked that choosing N is strictly dominated at any $t \geq 0$ such that $\mu(h^t) \leq 7/8$ (that is, choosing N is strictly dominated as long as the receiver returns or at $t = 0$).

Similar to what we have shown, at $t > 0$ off-the-equilibrium path $\sigma_S(h^{t-1}) = L$ is optimal with $\mu(h^{t-1}) \in (21/32, 7/8]$ (or $\mu(h^{t-1}) = 21/32$ and $h_{t-1} \neq (L, R)$). L is optimal since $\sigma_R(h^{t-1}, L)$ prescribes in this case a randomization strategy $\gamma(h^{t-1}, L)$ for the receiver such that

$$\frac{\mu(h^{t-1})\gamma(h^{t-1}, L)}{\mu(h^{t-1})\gamma(h^{t-1}, L) + 1 - \mu(h^{t-1})} = \frac{21}{32}.$$

It follows that randomizing between L and M as prescribed by $\sigma_S(h^{t-1}, (L, R))$ is optimal. In particular, with $h_t = (L, R)$, it follows that $\mu(h^{t-1}, (L, R)) = 21/32$ and the sender is indifferent between L and M at $t + 1$ similar to what we have argued above. Moreover, H and N are dominated at $t + 1$ given $\mu(h^{t-1}, (L, R)) = 21/32$ similar to what we argued above. Using similar steps, it can be established that $\sigma_S(h^{t-1})$ is optimal given $\sigma_R(h^{t-1}, S_t)$ and $\mu(h^{t-1})$ in the remaining off-the-equilibrium path cases.

The following claims will prove that E is efficient, and that all other PBE have the same features as E at $t = 0$ and $t = 1$, and equilibria start differing from $t = 2$ onwards. In particular, in every PBE including E, the sender starts by choosing L at $t = 0$, the low type receiver randomizes between R and D after L is chosen at $t = 0$, and the sender randomizes between L and M at $t = 1$ if the receiver chose R at $t = 0$ (once a receiver defaults, the sender optimally chooses N thereafter in every PBE). If $h^1 = ((L, R), (L, R))$ then $\sigma_S(h^1) = M$ in equilibrium E, whereas in other PBE, $\sigma_S(h^{t-1}) = L \sim M$. The proofs of these claims and more can be found below.

Claim 1 *The sender does not randomize between L and H in any period in any PBE.*

Proof. First, note that H is strictly dominated at t; since $\mu_0 = 7/8$, M is better than H. Next, we show that there is no equilibrium such that the sender randomizes between L and H at $t > 0$. From now on, we will only consider histories such that the receiver always chose R after a positive trust level was chosen. Otherwise, N is obviously the optimal sender choice. For $L \sim H$ to be part of an equilibrium at t, $h_t = (L, R)$ is followed at $t+1$ by either

(i) $\sigma_S(h^t) = H$, or

(ii) $\sigma_S(h^t) \in \{N, L, L \sim N\}$, or

(iii) $\sigma_S(h^t) \in \{M, M \sim L, M \sim N\}$, or

(iv) $\sigma_S(h^t) \in \{H \sim M, H \sim L, H \sim N\}$

We will show that all of the above are contradictory. To that aim, let μ denote the belief of the sender at the beginning of period t (i.e., $\mu \equiv \mu(h^{t-1})$). For (i) to hold, $\mu(h^{t-1})$ must be such that

$$\frac{32}{1-\delta_S} = \mu(h^{t-1}) \left(0 + \frac{\delta_S}{1-\delta_S} 24 \right) + (1 - \mu(h^{t-1})) \left(\frac{56}{1-\delta_S} \right) \tag{3}$$

where $\delta_S = 0.75$ (note that $\mu(h^{t-1}) = \mu(h^t)$ since the low type strictly prefers returning after L given that $\sigma_S(h^t) = H$ for $h_t = (L, R)$). It follows that $\mu(h^{t-1}) = 12/19$. However, observe that if $\mu(h^{t-1}) > 0.5$, then M is *strictly better* than H at t, a contradiction. Likewise, (ii) cannot hold because the parameters are such that the receiver would default with probability one after L is chosen at t if $\sigma_S(h^{t-1}, (L, R)) \in \{N, L, L \sim N\}$. This in turn implies that $\mu(h^{t-1}, (L, R)) = 0$ and thus, $\sigma_S(h^{t-1}, (L, R)) = H$ is optimal, which is a contradiction.

(iii) is contradictory due to the following. Randomization between L and H at t implies that H is weakly preferred over M. This requires that the posterior belief at t is weakly lower than 0.5. But if the sender randomization at t results in L, and the receiver chooses R, then the sender either chooses M or randomizes between M and another trust level. This requires that the posterior is weakly greater than 0.5 so that M is weakly preferred over H at $t+1$. As a result, the posterior must be exactly 0.5 at both t and $t+1$. But similar to what we showed in part (i) and in equality (3) above, this is not possible given the parameters of our game. With a posterior belief of 0.5, the sender strictly prefers choosing H at t rather than randomizing between L and H.

Finally, we show that (iv) is impossible. If the posterior $\mu(h^{t-1})$ does not change after L is chosen at t because $\sigma_R(h^{t-1}, L) = R$, then we can treat this as in part (i). So, assume that $\mu(h^{t-1}) \neq \mu(h^t)$ if the receiver chooses R after L is chosen at t (in fact, $\mu(h^{t-1}) > \mu(h^t)$ must hold with $h_t = (L, R)$). Let $\gamma \equiv \gamma(h^{t-1}, L)$ denote the probability with which the low type receiver returns after L is chosen at t. By our initial hypothesis, the sender is indifferent between H and L at t, and if $h_t = (L, R)$, then the sender chooses H at $t+1$ with positive probability. From the indifference condition, we derive

$$224 - 152\mu(h^{t-1}) = \mu(h^{t-1})(\gamma 86 + (1-\gamma)92) + (1 - \mu(h^{t-1}))200$$

which gives

$$\mu(h^{t-1}) = \frac{24}{44 - 6\gamma}.$$

Thus, $\mu(h^{t-1})$ is an *increasing* function of γ, and equals 24/44 if $\gamma = 0$. As a result, $\mu(h^{t-1}) > 0.5$ for all $\gamma \in [0, 1]$. But this is a contradiction as M would then be strictly preferred over H at t.

Claim 2 *M or H is never chosen at $t = 0$ in a PBE.*

Proof. First, note that H is strictly dominated by M at $t = 0$ as $\mu_o > 0.5$. Next, suppose towards a contradiction that there is an equilibrium such that the sender chooses M at $t = 0$. We argue that the sender would rather deviate to L at $t = 0$ instead, and choose M at $t = 1$ if the receiver chooses R at $t = 0$ (i.e., $\sigma_S((L, R)) = M$). The deviation payoff that the sender can get is weakly greater than $191 - 99\mu_o$, which is precisely the payoff if $\sigma_R(\varnothing, L) = D$ (that is, the low type receiver defaults with probability one after L is chosen at $t = 0$).[50]

[50]Note that the payoff of this deviation strategy is increasing in the probability with which the receiver chooses R after L is chosen at $t = 0$.

It follows that choosing M at $t = 0$ is dominated because $191 - 99\mu_o > 212 - 128\mu_o$ given $\mu_o = 0.875$. More generally, M is dominated by L as long as $\mu_o > 21/29$.

Claim 3 *The sender does not choose N at $t = 0$ or at any $t > 0$ provided that the receiver has always returned in the past.*

Proof. Assume towards a contradiction that $\sigma_S(h^{t-1}) \in \{N, N \sim L, N \sim M, N \sim H\}$. First, note that if $h_t = N$, then $\sigma_S(h^t) \in \{L, N, L \sim N\}$; other strategies that involve M or H would be contradictory since (i) the posterior $\mu(h^{t-1})$ does not change after N is chosen at t; (ii) choosing M over L requires $\mu(h^{t-1})$ to be smaller than 21/29, as argued in the proof Claim 2; and (iii) choosing H over M requires $\mu(h^{t-1})$ to be smaller than 0.5. But with such a value for $\mu(h^{t-1})$, N is strictly dominated at t, a contradiction. Now, assume without loss of generality that $\sigma_S(h^t) \in \{L, L \sim N\}$ given $h_t = N$ (otherwise, there exists a $\tau > t$ such that $\sigma_S(h^n) = N$ for all $n \in \{t, ..., \tau - 1\}$ and $\sigma_S(h^\tau) \in \{L, L \sim N\}$, to which the proof below still applies). If L is chosen at $t + 1$, and the receiver returns (i.e., $h_{t+1} = (L, R)$), then the sender must choose M or H with positive probability, otherwise we would have a contradiction. (As we argued before, the parameters are such that the receiver would default with probability one at $t+1$ if L is chosen at $t+1$ and $\sigma_S(h^{t+1}) \in \{L, L \sim N\}$, where $h_{t+1} = (L, R)$. Thus, choosing R at $t + 1$ will imply that $\mu(h^{t+1}) = 0$, and we must have $\sigma_S(h^{t+1}) = H$, instead of $\sigma_S(h^{t+1}) \in \{L, N, L \sim N\}$, a contradiction.) First, assume that if $h_{t+1} = (L, R)$, then the sender chooses M at $t + 2$ with positive probability. The highest possible payoff for the sender is when the receiver returns with probability 1 after L (that is, $\sigma_R(h^t, L) = R$). Given these, the sender payoff equals $191 - 96\mu(h^{t-1})$. Next, assume that the sender chooses H (with positive probability) following $h_{t+1} = (L, R)$. In that case, the highest possible payoff is $200 - 108\mu(h^{t-1})$. Then, the overall payoff of choosing N at t is bounded *above* by

$$24 + 0.75 \max\{(191 - 96\mu(h^{t-1})), (200 - 108\mu(h^{t-1}))\}.$$

If L is not part of the sender strategy $\sigma_S(h^{t-1})$, the lowest deviation payoff that the sender can get (deviating to L at t rather than choosing N) is bounded *below* by

$$191 - 99\mu(h^{t-1}).$$

This is the payoff if M is chosen at $t+1$ after $h_t = (L, R)$ and the low type receiver defaults with probability one. To reiterate, $191 - 99\mu(h^{t-1})$ is a lower bound on the deviation payoff. This term is higher than

$$24 + \frac{3}{4} \max\{(191 - 96\mu(h^{t-1})), (200 - 108\mu(h^{t-1}))\}.$$

if $\mu(h^{t-1}) < 95/108$. But this is true because $\mu_o = 7/8$ and $\mu(h^{t-1}) \leq 7/8$ since by hypothesis the receiver has always returned until t. Thus, we can also rule out randomization between $\not{L}$ and N at t: The sender cannot be indifferent between L and N at t because, similar to what we argued above, the lowest possible payoff from choosing L at t is strictly higher than the highest possible payoff from choosing N at t.

Claim 4 *In every PBE, the sender chooses* L *at* $t = 0$.

Proof. Given Claims $1-3$, L is the only possible equilibrium choice for the sender at $t = 0$.

Claim 5 *The sender never randomizes between* M *and* H *in any PBE.*

Proof. At period $t > 0$ with arbitrary history h^{t-1}, randomizing between M and H can be optimal only if $h_{t-1} = (L, R)$ and $\mu(h^{t-1}) = 0.5$ (given Claim 3, N is never chosen at $t = 0$ or as long as the receiver has always chosen R). But since $\mu(h^{t-1}) = 0.5$ must hold, this is contradictory; from (3), it follows that $\mu(h^{t-1})$ must be weakly greater than $12/19$ for L to be chosen over H at $t-1$, a contradiction.

Claim 6 *In every PBE, the sender randomizes between* L *and* M *at* $t = 1$ *provided that* $h_0 = (L, R)$ *(by Claim 4, the sender chooses* L *at* $t = 0$ *in every PBE).*

Proof. Other possibilities are the pure strategies L, M and H but they all lead to a contradiction. For example, it cannot be the case that $\sigma_S(h^0) = M$ with $h_0 = (L, R)$. One can check that given $\sigma_S(h^0) = M$ with $h_0 = (L, R)$, the receiver would optimally choose R at $t = 0$, the posterior $\mu(h^0)$ would be the same as the prior μ_0 and thus, the sender would strictly prefer deviating to L at $t = 1$, similar to what we showed in Claim 2. The exact same argument also holds if $\sigma_S((L, R)) = H$. Finally, if the equilibrium prescribes a choice of L with probability one at $t = 1$ following $h_0 = (L, R)$, then only the high type receiver would choose R, implying that if $h_0 = (L, R)$, then $\mu(h^0) = 0$, and thus choosing H is optimal at $t = 1$ following $h_0 = (L, R)$, which is a contradiction.

Claim 7 (Reputation building) *In every PBE, the low type receiver randomizes between* R *and* D *at* $t = 0$ *(by Claim 4, the sender chooses* L *at* $t = 0$ *in every PBE). Moreover, the low type always chooses* R *with strictly positive probability after* L *is chosen at* $t > 0$ *provided that the receiver has never defaulted before and thus,* $\mu(h^{t-1}) \leq 7/8$*. Finally, the low type receiver defaults with probability one whenever* M *or* H *is chosen.*

Proof. Given Claim 6, the only possibility for the low type receiver is to randomize between R and D in such a way that the sender is indifferent between L and M at $t = 1$. If the low type receiver were to choose R with probability one, then the sender would never choose M and deviate to L at $t = 1$, which is a contradiction. If however the low type receiver were to choose D with probability one, then the sender would deviate to H at $t = 1$ following $h_0 = (L, R)$, which is another contradiction. Thus, the sender randomizes between L and M at $t = 1$ following $h_0 = (L, R)$ in such a way that the receiver randomizes between R and D. The second statement holds because otherwise, the decision to choose R at t would be a sure sign of high type and the sender would choose H at $t+1$. This in turn implies that the low type would strictly prefer choosing R at t, a contradiction. The last statement follows from the parameters of the game as we discussed before.

Claim 8 *In every PBE, L choice on the equilibrium path at period $t \geq 1$ is followed by either M or randomization between L and M as long as the receiver returns at every $\tau \leq t$, and $h_t = (L, R)$.*

Proof. All the possibilities other than the pure strategy choice of L or H are ruled out by Claims 1 and 3. Suppose towards a contradiction that L is chosen at $t+1$ following $h_t = (L, R)$ given the above stated assumptions. Then, similar to what we discussed in the proof of Claim 6, only the high type receiver would choose R at t, implying that if $h_t = (L, R)$, then $\mu(h^t) = 0$ and thus, choosing H is optimal at $t+1$ following $h_t = (L, R)$, which is a contradiction. Next, suppose towards a contradiction that H is chosen at $t+1$ following $h_t = (L, R)$ under the above stated assumptions. Then, similar to what we discussed in the proof of Claim 1, $\mu(h^{t-1})$ must be weakly greater than $12/19$ for L to be chosen over H at t. But then, $\mu(h^{t-1}) = \mu(h^{t-1}, (L, R)) \geq 12/19$ (because the low type optimally chooses R with probability one at t given the sender strategy) and M is strictly preferred over H at $t+1$, a contradiction.

Claim 9 *M choice in period t is followed by H if $h_t = (M, R)$ and $\mu(h^t) \leq 0.875$, and by N otherwise.*

Proof. This must hold due to the last statement of Claim 7.

Claim 10 *Equilibrium E is efficient.*

Proof. First, note that by Claims 3 and 7, the low type receiver cannot be better off in another PBE. We will now show that this is also true for the sender. First, recall that in every PBE the sender starts by choosing L at $t = 0$ and then randomizes between L and M provided that the receiver has chosen R at $t = 0$. We will show that the return rate after L is chosen at $t = 0$ is not higher in any other PBE, which will then prove that the sender cannot be better off in another equilibrium. Other PBE differ from equilibrium E in that the sender chooses M with probability one if $h^1 = ((L, R), (L, R))$ whereas in others $\sigma_S(h^1) = L \sim M$. In every PBE, the following must hold at $t = 1$:

$$\mu(h^0)(\gamma(h^0, L)95 + (1 - \gamma(h^0, L))92) + (1 - \mu(h^0))191 = \mu(h^0)84 + (1 - \mu(h^0))212,$$

where $h^0 = (L, R)$. It follows that

$$\mu(h^0) = \frac{21}{29 + 3\gamma(h^0, L)}.$$

Recall that in every PBE other than E, $\sigma_S(h^1) = L \sim M$ if $h^1 = ((L, R), (L, R))$. Put differently, the sender must be indifferent between L and M at $t = 2$. This requires that $\mu(h^1) \geq \frac{21}{32}$. Otherwise, it can easily be checked that $\mu(h^1)$ is sufficiently low so that M is strictly preferred over L. This is because the equilibrium payoff from L at $t = 2$ is

$$\mu(h^1)(\gamma(h^1, L)95 + (1 - \gamma(h^1, L))92) + (1 - \mu(h^1))191,$$

which is strictly smaller than $\mu(h^1)84 + (1 - \mu(h^1))212$ (regardless of $\gamma(h^1, L)$) provided that $\mu(h^1) < \frac{21}{32}$. Thus, $\mu(h^1) \geq \frac{21}{32}$. Finally, note that

$$\mu(h^1) = \frac{\mu(h^0)\gamma(h^0, L)}{\mu(h^0)\gamma(h^0, L) + (1 - \mu(h^0))} = \frac{21\gamma(h^0, L)}{24\gamma(h^0, L) + 8}$$

given what we have already shown above. It follows that $\frac{21\gamma(h^0,L)}{24\gamma(h^0,L)+8} \geq \frac{21}{32}$ must hold, and this in turn implies that $\gamma(h^0, L) = 1$. Thus, in every PBE (including E) $\gamma(h^0, L) = 1$ and $\mu(h^0) = \frac{21}{32}$. This in turn means that $\gamma(\varnothing, L)$ is identical in every PBE. Hence, the desired result.

B Learning in the Gradual Game

Here, we briefly compare sender behavior in the first 10 matches of the Gradual game with behavior after the first 10 matches. Table 6 demonstrates that senders learn to use equilibrium strategies over time. The percentage of behavior that we can classify as equilibrium based on the first two rounds of matches that last at least two rounds increases from 42.4% in the first 10 matches to 55.8% after the first 10 matches. Table 6 shows that the effect of experience is even more pronounced when we compare experienced 3-round and 4-round sender behavior with the behavior in the first 10 matches. Thus, we conclude that equilibrium behavior is much more common if senders are experienced.

Table 6: Learning to use Equilibrium Strategies

Matches that last at least n rounds	In the first 10 matches	After the first 10 matches
$n = 2$	42.4%	55.8%
$n = 3$	29.9%	47.7%
$n = 4$	23.4%	38.3%

Note: Shares of Equilibrium Behavior in the first n rounds, $n = 2, 3, 4$

C Projection of efficiency levels over a long time horizon

In this Appendix we assess to what joint, sender, and receiver efficiency levels play might have converged to had the experiment been conducted over a very long time horizon. For this purpose, we run the following two OLS regressions (similar to those suggested in Noussair *et al.* (1995) and Barut *et al.* (2002)):

$$\text{Eff}_{ijt} = \sum_{j=1}^{6} \alpha_j^{Bin} \times \frac{D_j^{Bin}}{t} + \beta^{Bin} \times \frac{(t-1)}{t} + \varepsilon \qquad (4)$$

and

$$\text{Eff}_{ijt} = \sum_{j=1}^{6} \alpha_j^{Grad} \times \frac{D_j^{Grad}}{t} + \beta^{Grad} \times \frac{(t-1)}{t} + \varepsilon, \qquad (5)$$

where Eff_{ijt} denotes observed efficiency index in group i, of session j in match t; D_j^{Bin} (D_j^{Grad}) is a session dummy for session j of treatment Binary (Gradual). The regressions were ran with all data and with observations clustered at the session level. As in Section 4 of the paper, the efficiency index is again defined as $(\pi^{obs} - \pi^N)/(\pi^C - \pi^N)$. Following the interpretation in Barut *et al.* (2002), the coefficient α_j^{Bin} (α_j^{Grad}) is an estimate of the efficiency index in match 1 of session j in treatment Binary (Gradual), whereas β^{Bin} (β^{Grad}) is the estimated efficiency index to which the time series converges if $t \to \infty$. Hence, this model allows the starting efficiency index to be different across the individual sessions of a treatment but assumes the efficiency index in all sessions of a treatment to converge to a common asymptote.

We are interested in whether in the very long run the efficiency indices of a pair of corresponding treatments would have converged to the same level. To test the hypothesis H_0: $\beta^{Bin} = \beta^{Grad}$, we compute the statistic

$$z = \frac{\widehat{\beta}^{Bin} - \widehat{\beta}^{Grad}}{\sqrt{se(\widehat{\beta}^{Bin})^2 + se(\widehat{\beta}^{Grad})^2}}, \qquad (6)$$

where $\widehat{\beta}^{Bin}$ and $\widehat{\beta}^{Grad}$ are the estimated coefficients and $se(\widehat{\beta}^{Bin})$ and $se(\widehat{\beta}^{Grad})$ the standard errors from equations (4) and (5), respectively. Table 7, which is organized like Table (7) in Section 4 of the paper, shows the results. The "<" and "≈" signs in Table 7 summarize the results of statistical tests based on the z-statistic defined in (6), with "≈" indicating no significant difference at the 10%-level. The superscripts indicate the level of significance. Our focus is on the vertical comparisons, testing for differences across pairs of corresponding Gradual and Binary treatments. Comparing the test results in Table 7 with those in Table 5 shows that all pairwise tests are qualitatively very similar. In particular, the pro-

Table 7: Estimated asymptotes and test results

$\frac{\pi^{obs}-\pi^N}{\pi^C-\pi^N}$	*Joint-efficiency Index*	*Sender-efficiency Index*	*Receiver-efficiency Index*
Gradual Game	0.341 (0.014)	0.079 (0.024)	0.766 (0.025)
	$<^{***}$	$\approx$	$<^{***}$
Binary Game	0.475 (0.029)	0.144 (0.050)	1.019 (0.081)
hr-Gradual Game	0.326 (0.048)	0.041 (0.032)	0.643 (0.029)
	$\approx$	$\approx$	$<^{**}$
hr-Binary Game	0.270 (0.021)	0.036 (0.060)	0.876 (0.092)
le-Gradual Game	0.379 (0.020)	0.053 (0.029)	0.610 (0.074)
	$\approx$	$\approx$	$\approx$
le-Binary Game	0.364 (0.059)	0.027 (0.056)	0.594 (0.032)

Notes: The table shows the estimated asymptotes $\widehat{\beta}^{Bin}$ and $\widehat{\beta}^{Grad}$ from equations (4) and (5). Standard errors in parentheses. The "<" and "≈" signs summarize the results of statistical tests based on the z-statistic defined in equation (6), with "≈" indicating no significant difference at the 10%-level. The superscripts indicate the level of significance, where *** and ** indicates significance at the 1% and 5% level, respectively.

jections provided in this section suggest that the Gradual game would never have generated higher efficiency levels than the binary game in case the experiments had lasted much longer (Hypothesis 2).

D Data Analysis Omitted in the Main Text

D.1 Match Frequency Distribution

Figure 5 shows a histogram of match lengths in the baseline Gradual game treatment.

Figure 5: Histogram of Match Lengths in the Gradual Game

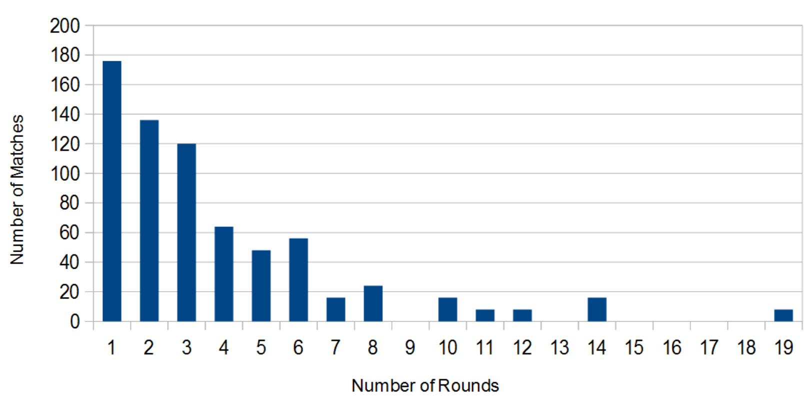

Note: First 10 matches dropped.

D.2 Sender Behavior in the Gradual Game

In this subsection we report frequencies of observed sender histories in the Gradual game up to the second, third, and fourth sender decision in a match. The results are presented in Tables 8-10. In these tables we indicate the histories, their observed frequencies and the type of history according to our classification introduced in Section 4.1.1. For instance, the first entries in Table 8 mean that the history up to the second sender decision "LD, N" in which the sender chooses L and the receiver chooses D in the first round of a match and the sender chooses N in the second round of a match, was observed 105 times, and that this history is consistent with an eq(uilibrium) strategy.

Table 8: Frequencies of observed two-round histories in the Gradual game

History	#(obs)	type	History	#(obs)	type
LD, N	105	eq	N-, N	11	nt
LR, L	97	eq	HD, H	9	un
LR, M	88	eq	MR, H	5	ng
LR, H	43	ng	HD, L	3	ln
MD, N	32	ng	MD, M	3	un
N-, L	29	hy	MD, H	3	un
MR, M	21	ng	LR, N	2	un
LD, L	21	un	MR, L	1	un
MD, L	16	ln	HR, L	1	un
HD, N	15	ht	HR, N	1	un
HR, H	13	ht	HR, M	1	un
			Total	520	

Notes: Sender histories are coded as follows: eq = equilibrium; ng = (non-equilibrium) gradualist; nt = always no trust; ht = high trust until default; ln = lenient; hy = hybrid; un = unclassified. For details on the precise definition of strategies, see Section 4.1.1 in the main text. First 10 matches dropped.

Table 9: Frequencies of observed three-round histories in the Gradual game

History	#(obs)	type	History	#(obs)	type
LD, N-, N	70	eq	HD, LD, N	2	ln
LR, LR, M	29	eq	HR, HD, L	2	ln
LR, MD, N	27	eq	N-, N-, L	2	hy
LR, MR, H	20	eq	HD, HR, H	1	un
LR, LR, L	19	eq	HD, HD, H	1	un
LR, LD, N	18	eq	MD, HD, H	1	un
MD, N-, N	16	ng	LD, LR, L	1	un
LR, HD, N	13	ng	HD, LD, L	1	un
HD, N-, N	12	ht	LD, N-, M	1	ln
MR, MR, H	12	ng	LR, HD, L	1	ln
LR, MR, M	10	ng	MD, LR, M	1	ln
MD, LD, N	9	ln	HR, LD, L	1	un
LR, LR, H	9	ng	LD, LR, M	1	un
N-, N-, N	9	nt	HR, N-, L	1	un
LD, N-, L	8	ln	LR, HR, L	1	un
N-, LD, N	7	hy	MD, LD, L	1	un
LD, LD, L	7	un	MD, MD, N	1	un
LR, HR, H	7	ng	LR, LD, M	1	un
LR, MD, L	6	ln	MR, HD, H	1	un
LD, LD, N	6	un	MD, MR, H	1	un
HD, HD, N	5	un	MR, MD, M	1	un
HR, HR, H	5	ht	N-, LR, N	1	un
N-, LD, L	4	un	MD, HD, N	1	un
MD, N-, L	4	ln	LR, HD, H	1	un
LR, MD, M	4	un	LR, N-, L	1	un
N-, LR, M	3	hy	HR, MR, H	1	un
N-, LR, L	3	hy	MR, MD, N	1	ng
LR, LD, L	3	un	LR, MR, L	1	un
HR, HD, N	3	ht	MR, MR, M	1	ng
MR, HR, H	3	ng	MD, N-, M	1	ln
		Total	384		

Notes: Sender histories are coded as follows: eq = equilibrium; ng = (non-equilibrium) gradualist; nt = always no trust; ht = high trust until default; ln = lenient; hy = hybrid; un = unclassified. For details on the precise definition of strategies, see Section 4.1.1 in the main text. First 10 matches dropped.

Table 10: Frequencies of observed four-round histories in the Gradual game

History	#(obs)	type	History	#(obs)	type
LD, N-, N-, N	44	eq	MD, HD, HD, H	1	un
LR, MR, HR, H	12	eq	LR, LD, LD, N	1	un
LR, LD, N-, N	11	eq	N-, LR, MR, H	1	hy
LR, MD, N-, N	11	eq	LD, LD, LD, N	1	un
MD, N-, N-, N	10	ng	N-, LR, MD, N	1	hy
LR, LR, MR, H	7	eq	LR, LR, LD, M	1	un
HD, N-, N-, N	7	ht	N-, LD, N-, L	1	un
N-, LD, N-, N	6	hy	N-, N-, N-, H	1	un
LR, MD, N-, L	6	ln	N-, LR, LR, L	1	hy
LD, N-, N-, L	6	ln	HR, HD, N-, N	1	ht
LR, LR, LR, H	6	ng	LR, LD, LD, H	1	un
N-, N-, N-, N	5	nt	LR, HD, LD, N	1	ln
LD, N-, LD, N	5	ln	N-, N-, LD, N	1	hy
LR, HD, N-, N	5	ng	MD, N-, LD, N	1	ln
LR, LR, HD, N	5	ng	MD, LR, MD, L	1	ln
LR, MR, HD, N	5	eq	HR, LD, LD, L	1	un
MD, LD, N-, L	4	ln	HD, LD, N-, M	1	ln
HR, HR, HR, H	4	ht	LD, LR, MD, M	1	un
LR, LR, MR, M	4	ng	HD, LD, N-, N	1	ln
LR, LR, MD, N	4	eq	LR, MD, MD, L	1	un
LD, LD, N-, N	4	un	MD, N-, LD, L	1	un
LR, LR, LR, L	3	eq	N-, LR, LR, M	1	hy
HD, HD, N-, N	3	un	LR, LD, MR, M	1	un
LD, LD, LD, L	3	un	MR, HD, HD, H	1	un
N-, LD, LD, L	3	un	LR, LR, HD, L	1	ln
MD, LD, N-, N	3	ln	MD, MR, HR, H	1	un
MR, HR, HR, H	3	ng	LD, LD, LR, M	1	un
LR, HR, HR, H	3	ng	LR, MR, MD, M	1	un
MD, N-, N-, L	3	ln	LR, LR, MD, L	1	ln
MR, MR, HD, N	3	ng	MR, MD, MD, N	1	un
LR, LR, LR, M	2	eq	N-, LR, N-, L	1	un
LR, LD, N-, L	2	ln	MD, HD, N-, N	1	un
LD, N-, LD, L	2	un	LR, HD, HR, M	1	un
LR, MD, MD, N	2	un	LR, N-, LR, H	1	un
LR, MD, LD, L	2	un	LR, LD, LR, L	1	un
LR, MR, MR, H	2	ng	HR, MR, HR, H	1	un
LD, LD, N-, L	2	un	MR, MR, MD, N	1	ng
LR, LR, LD, N	2	eq	MD, N-, MD, N	1	ln
MR, MR, HR, H	2	ng	HD, N-, N-, L	1	ln
LR, MR, MD, N	2	ng	HR, HR, HD, N	1	ht
LR, MD, LD, N	2	ln	N-, N-, N-, L	1	hy
LR, HD, N-, L	2	ln	LD, N-, N-, H	1	ln
			Total	264	

Notes: Sender histories are coded as follows: eq = equilibrium; ng = (non-equilibrium) gradualist; nt = always no trust; ht = high trust until default; ln = lenient; hy = hybrid; un = unclassified. For details on the precise definition of strategies, see Section 4.1.1 in the main text. First 10 matches dropped.

D.3 Sender Behavior in the Binary Game

In this subsection we report frequencies of observed sender histories in the Binary game up to the second, third, and fourth sender decision in a match. The results are presented in Tables 11-13. In these tables we indicate the histories, their observed frequencies and the type of history according to our classification introduced in Section 4.1.2. For instance, the first entries in Table 11 mean that the history up to the second sender decision "TR, T" in which the sender chooses T and the receiver chooses R in the first round of a match and the sender chooses T in the second round of a match, was observed 160 times, and that this history is consistent with the always trust strategy.

Table 11: Frequencies of observed two-round histories in the Binary game

History	#(obs)	type	History	#(obs)	type
TR, T	160	td	TD, T	70	at
TD, N	120	td	N-, T	69	hy
N-, N	116	nt	TR, N	1	un
			Total	536	

Notes: Sender histories are coded as follows: nt = always no trust; td = trust until default; at = always trust; ln = lenient; hy = hybrid; un = unclassified. For details on the precise definition of strategies, see Section 4.1.2 in the main text. First 10 matches dropped.

Table 12: Frequencies of observed three-round histories in the Binary game

History	#(obs)	type	History	#(obs)	type
TR, TR, T	76	td	TD, N-, T	13	ln
TD, N-, N	70	td	N-, TR, T	13	hy
N-, N-, N	60	nt	TD, TD, N	12	un
TR, TD, N	26	td	N-, TD, T	12	un
TD, TR, T	20	at	N-, N-, T	3	hy
N-, N-, T	20	hy	TD, TR, N	2	un
TD, TD, T	17	at	TR, N-, N	1	un
N-, TD, N	16	hy	TR, TR, N	1	un
TR, TD, T	14	at	Total	376	

Notes: Sender histories are coded as follows: nt = always no trust; td = trust until default; at = always trust; ln = lenient; hy = hybrid; un = unclassified. For details on the precise definition of strategies, see Section 4.1.2 in the main text. First 10 matches dropped.

Table 13: Frequencies of observed four-round histories in the Binary game

History	#(obs)	type	History	#(obs)	type
TD, N-, N-, N	49	td	TD, N-, TD, N	4	ln
TR, TR, TR, T	44	td	N-, TD, TR, T	4	un
N-, N-, N-, N	33	nt	TR, TD, TR, T	4	at
TR, TD, N-, N	14	td	N-, TD, TD, N	4	un
N-, N-, N-, T	12	hy	N-, N-, TD, T	3	un
TR, TR, TD, N	10	td	N-, TD, TD, T	3	un
N-, TD, N-, N	9	hy	N-, TR, TD, N	2	hy
TD, TR, TD, T	9	at	TD, TD, TR, T	2	un
N-, N-, TD, N	7	hy	N-, TD, N-, T	2	un
TD, TD, N-, N	6	un	TD, N-, TD, T	2	un
N-, TR, TR, T	6	hy	TR, TD, TD, N	2	un
TR, TR, TD, T	5	at	TR, TR, N-, N	1	un
TR, TD, TD, T	5	at	TR, N-, N-, N	1	un
TD, TR, TR, T	5	at	TD, TR, TR, N	1	un
TR, TD, N-, T	5	ln	N-, TR, TD, T	1	un
TD, TD, TD, N	5	un	TD, N-, N-, T	1	ln
N-, N-, TR, T	5	hy	TD, N-, TR, T	1	ln
TD, TD, TD, T	4	at	TD, TR, N-, N	1	un
			Total	272	

Notes: Sender histories are coded as follows: ant = always no trust; tud = trust until default; at = always trust; ln = lenient; hy = hybrid; un = unclassified. For details on the precise definition of strategies, see Section 4.1.2 in the main text. First 10 matches dropped.

E Expected Payoffs to Senders

To evaluate which strategies were *empirically* optimal, we compare the expected Sender payoffs of various strategies, given the actual behaviour of Responders in our experiment. To do so, we calculate the empirical return rate (including computerized Responders) after each relevant history. This recognizes that, for example, the return rate after H in the second round may depend on whether N or H was played in the first round. Our analysis is constrained by the large number of possible histories and the fact that many games terminate early on, meaning that the number of observations available for estimating return rates decreases rapidly. Recognising this trade-off between noise and learning about more complete strategies, we present results for one, two, and three round strategies, conditional on the game lasting at most one, two, or three rounds, respectively.

The following expression is the expected payoff to a Sender who first plays L, then M if the Receiver defaults, and N otherwise, conditional on the game lasting at most two rounds:

$$\frac{1}{1+\delta}\left[20(1-\hat{p}_L)+32\hat{p}_L\right]+\frac{\delta}{1+\delta}\left[(20+24)(1-\hat{p}_L)+\hat{p}_L((32+12)(1-\hat{p}_{LM})+(32+44)\hat{p}_{LM})\right]$$

where $\frac{1}{1+\delta}$ is the probability of a game lasting one round conditional on the game lasting at most two rounds; $\hat{p}_L$ is the observed proportion of R after L was played in the first round; $\frac{\delta}{1+\delta}$ is the probability of a game lasting two rounds conditional on the game lasting at most two rounds; and $\hat{p}_{LM}$ is the proportion of R following M in the second round after L was played and returned in the first round.

The results of our calculations can be found in Table 14. We abuse notation by omitting contingent elements of the equilibrium and non-equilibrium gradualist strategies, that is, we only describe the trust levels that would be chosen if the receiver returns in each round: it is always assumed default is punished with N thereafter.

Recall that in the first round of the Gradual game, every equilibrium strategy prescribes L. We find that, conditional on a match length of one round, L is the best choice given the observed receiver behavior and provides a payoff that is 13% higher than that of the second best choice, M. H is the worst choice. In a similar vein, conditional on a match length of at most two or three rounds, the expected return for every equilibrium strategy is strictly greater than the expected return to non-equilibrium gradualist strategies, as well as the return to strategies "H until default" and "always N."

F Parametric Analysis of Payoffs

F.1 Parametric Comparison of Payoffs with Theory

For the parametric comparison of players' observed payoffs with theory, we regress the relevant payoff index (for the definition, see Section 4.2.1 in the main text) on a constant and controls for the match number (i.e., time trend) as well as the number of rounds played within a match. We use the estimated coefficients of the independent variables to find out whether a payoff index is significantly different from one. Note that errors are always clustered at the session level. The estimation results are given in Table 15. Explanatory variables Match# and #Rounds refer to the match number and the number of rounds played within a match, respectively. Since we drop the first 10 rounds, the average of the match# is 17.770 and the average of the #rounds is 3.827. For the joint, the sender, and the receiver payoff index, we test the following null hypothesis:

H_0: $Constant + 17.77 \times match\# + 3.827 \times \#rounds = 1.$

The p-value given by the F-statistic is 0.0014 for the joint-payoff index. Thus, we reject the

Table 14: Expected Payoffs to Strategies

X	Type of Strategy	Expected Payoffs to Strategies that prescribe X for games that last **at most** 1 round
L	Equilibrium	27.71
M	-	24.52
N	-	24.00
H	-	18.67

XY	Type of Strategy	Expected Payoffs to Strategies that prescribe XY for games that last **at most** 2 rounds
LL	Equilibrium	39.28
LM	Equilibrium	38.90
MM	Non-equilibrium gradualist	37.39
MH	Non-equilibrium gradualist	36.42
LH	Non-equilibrium gradualist	36.41
NN	Always no trust	34.29
HH	High trust until default	30.45

XYZ	Type of Strategy	Expected Payoffs to Strategies that prescribe XYZ for games that last **at most** 3 rounds
LLM	Equilibrium	49.38
LMH	Equilibrium	49.29
LLL	Equilibrium	49.27
LMM	Non-equilibrium gradualist	49.14
MMH	Non-equilibrium gradualist	47.95
MHH	Non-equilibrium gradualist	47.03
LLH	Non-equilibrium gradualist	46.86
LHH	Non-equilibrium gradualist	46.69
MMM	Non-equilibrium gradualist	46.43
NNN	Always no trust	43.46
HHH	High trust until default	41.14

Note: This table reports the expected sender payoffs of various strategies, given the actual behaviour of responders in our experiment that is used to calculate the empirical return rate (including computerized responders) after each relevant history.

Table 15: Regression Analysis for the Gradual Game Treatment

	Dependent Variable					
	Joint-payoff Index		Sender-payoff Index		Receiver-payoff Index	
Constant	1.1254***	(0.026)	1.0252***	(0.041)	1.2197***	(0.065)
Match#	−0.0016	(0.002)	−0.0002	(0.002)	−0.0027	(0.005)
#Rounds	0.0020	(0.004)	0.0045	(0.006)	0.0079	(0.004)

Notes: (1) Errors are clustered at the session level. (2) Robust standard errors are given in parentheses. (3) The superscripts indicate the level of significance, where ***, **, and * indicates significance at the 1%, 5%, and 10% level, respectively.

null. As explained in the main text, this result is driven by the high receiver earnings in the data. We cannot reject the above null hypothesis for the sender-payoff index ($p = 0.1785$), but for the receiver-payoff index ($p = 0.0003$). Coefficients for the explanatory variables match# and #rounds are insignificant, as seen in Table 15. Thus, there is no time trend if we focus on experienced play (recall that the first 10 matches are dropped). None of our results would be affected if we drop the explanatory variables match# and #rounds and simply test for the null $H_0 : constant = 1$.

We, now, conduct the same type of analysis for the Binary game treatment. Table 16 shows the results of our regressions. We reject the null for both the joint-payoff index and the receiver-payoff index (p-value < 0.001 for both). However, we can reject it only at the 10% level for the sender-payoff index (p-value $= 0.064$). Recall that the binomial test failed to detect a difference between the sender-payoff index and one.

Table 16: Regression Analysis for the Binary Game Treatment

	Dependent Variable					
	Joint-Payoff Index		Sender-Payoff Index		Receiver-Payoff Index	
Constant	2.0890***	(0.162)	1.4331***	(0.188)	3.3421***	(0.299)
Match#	−0.0163	(0.008)	−0.0176*	(0.009)	−0.0150	(0.015)
#Rounds	−0.0380**	(0.012)	0.0129	(0.008)	−0.0503***	(0.011)

Notes: (1) Errors are clustered at the session level. (2) Robust standard errors are given in parentheses. (3) The superscripts indicate the level of significance, where ***, **, and * indicates significance at the 1%, 5%, and 10% level, respectively.

F.2 Parametric Comparison of Payoffs across Treatments

To parametrically test for efficiency differences across treatments, we regress the relevant efficiency index (for the definition, see Section 4.2.2 in the main text) on a binary treatment variable (denoted by Gradual) and controls for the match number (i.e., time trend), the number of rounds played within a match and possible interaction effects. We use the coefficients of the independent variables to find out whether the Gradual game treatment has an effect on payoffs. As before, errors are always clustered at the session level.

Table 17: Regression Analysis for Treatment Effects

	Dependent Variable					
	Joint-Efficiency Index		Sender-Efficiency Index		Receiver-Efficiency Index	
Gradual	−0.2941**	(0.108)	−0.2775*	(0.137)	−0.3040	(0.185)
Match#	−0.0113*	(0.005)	−0.0132*	(0.006)	−0.0086	(0.008)
#Rounds	−0.0244***	(0.007)	0.0097	(0.006)	−0.0295	(0.007)
Gradual×Match#	0.0082	(0.006)	0.0127*	(0.006)	0.0037	(0.010)
Gradual×#Rounds	0.0070	(0.008)	0.0010	(0.006)	0.0079	(0.004)

Notes: (1) Errors are clustered at the session level. (2) Robust standard errors are given in parentheses. (3) The superscripts indicate the level of significance, where ***, **, and * indicates significance at the 1%, 5%, and 10% level, respectively.

Explanatory variables Large×Match# and Large×#Rounds refer to the interaction variables. To find out whether the large game treatment has an effect on payoffs, we test for the following null hypothesis:

$$H_0\text{: } Gradual + 17.77 Gradual \times Match\# + 3.827 Gradual \times \#Rounds = 0.$$

The p-value given by the F-statistic for the joint-efficiency index is 0.0028. Thus, we reject the null. The p-value when the sender-efficiency index is the dependent variable is 0.3765. Therefore, we are unable to reject the null. Finally, we reject the null for the receiver-efficiency index with a p-value of 0.0104.

G Experimental Instructions

On the next pages, we reproduce the experimental instructions used in treatment "Gradual" and in the "Binary" game. For the other treatments the instructions were identical apart from the payoffs in the game trees.

INSTRUCTIONS

- Welcome to this experiment!
- During the course of the experiment, you will be asked to make a series of decisions. If you follow the instructions carefully, you can earn a considerable amount of money, depending on your decisions and those made by other participants. The money you earn will be paid to you in cash at the end of the experiment.
- Please remain silent during the experiment. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you.

Roles

- There are 15 people who participate in this experiment.
- Each of you will be randomly assigned to a role which could be either a so-called A or B participant. You will remain in the same role throughout the entire experiment. You will be informed about your role on your computer screen when the experiment begins.
- In particular, there will be 8 human participants assigned to the A role, and 7 human participants assigned to the B role. Decisions for an 8^{th} B role will be made by a computer (this role will be referred to as "the computerized B participant").

Matches

- The experiment is divided into various sequences, each of which consists of various rounds. Each sequence of rounds is referred to as a match.
- The length of a match is randomly determined. After each round of a match, there is a 3 out of 4 chance that the match will continue for another round. So, for instance, if you are in round 1 of a match, the chance that there will be a 2^{nd} round is 3 out of 4, and if you are in round 2 of a match, the chance that there will be a 3^{rd} round is also 3 out of 4, and so on.

Pairing

- At the beginning of a match, each A participant will be paired with a B participant and will remain with that B participant until the end of the match.
- The B participant may be human or computerized. The A participants will not be informed who they have been paired with and, in particular, they will not know whether they are paired with a human B participant or the computerized B participant. As there are 7 human B participants and 1 computerized B participant,

the chance of an A participant to be paired with a human B participant is 7 out of 8 and the chance to be paired with the computerized B participant is 1 out of 8.

- Once a match ends, each A participant will again be randomly paired with a B participant.

Within a Match

Each round of a match will proceed as follows:

- Each A participant begins the round by choosing one of four alternatives. These alternatives are labeled a1, a2, a3, and a4, and can be seen in Figure 1.
- The choice of an A participant is then shown to the B participant he/she is paired with. As you can see in Figure 1, if A's choice is a1, there is no decision to be made by the B participant. However, if the choice of the A participant is either a2, a3, or a4, then a human B participant can decide between choosing b1 or b2.
- **The computerized B participant is "programmed" to always choose b1 after a choice of a2, a3, or a4 of the A participant**.
- At the end of each round, each participant receives a payoff according to the decisions he/she and the participant with whom they were matched made, as shown in Figure 1.

Figure 1: Payoffs in Round 1 of a match

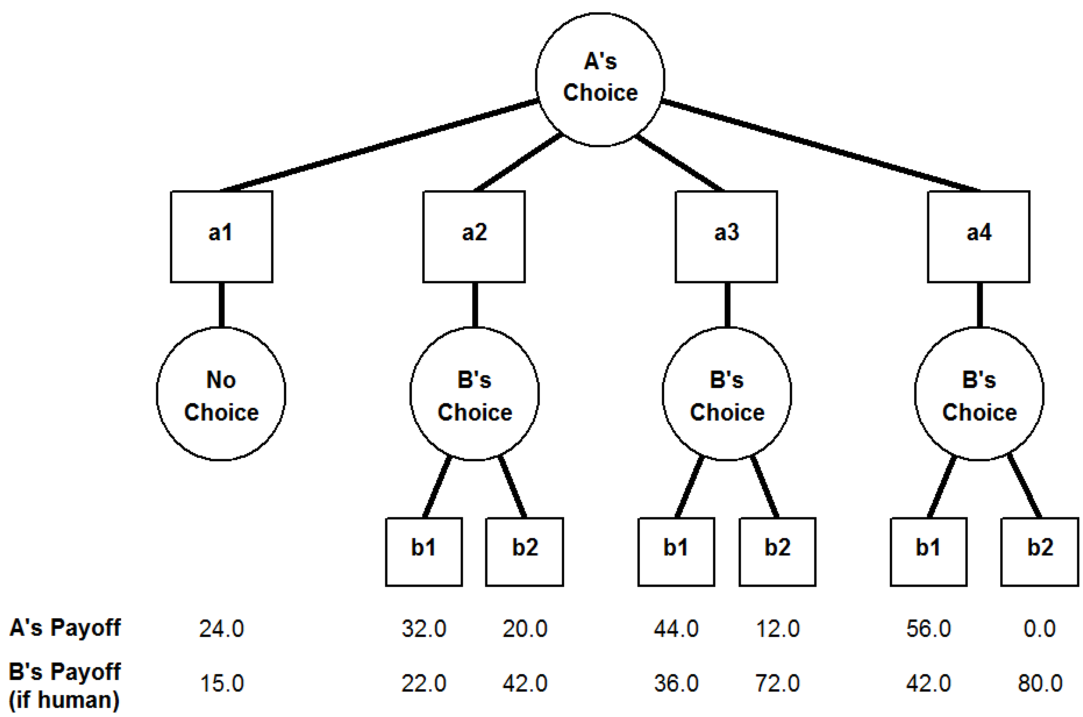

There is one more important rule concerning the payoffs:

- Starting in round 2 of a match, the possible payoffs as indicated in Figure 1 for a human B participant are reduced by a factor of 1/3 in each round.
- More precisely, in round 1 of a match, the payoffs for A and human B participants are as indicated in Figure 1. In Round 2 of a match, the payoffs for A and human B participants are as indicated in Figure 2, while in round 3 of a match, the payoffs for A and human B participants are as indicated in Figure 3, and so on.
- The payoffs will always be rounded to the nearest decimal place.
- In the decisions screen, which will be explained below, the payoffs in the current round of a match will always be indicated.
- The payoffs for A remain the same in every round of a match.
- Computerized B participants receive no payoff.

Figure 2: Payoffs in Round 2 of a match

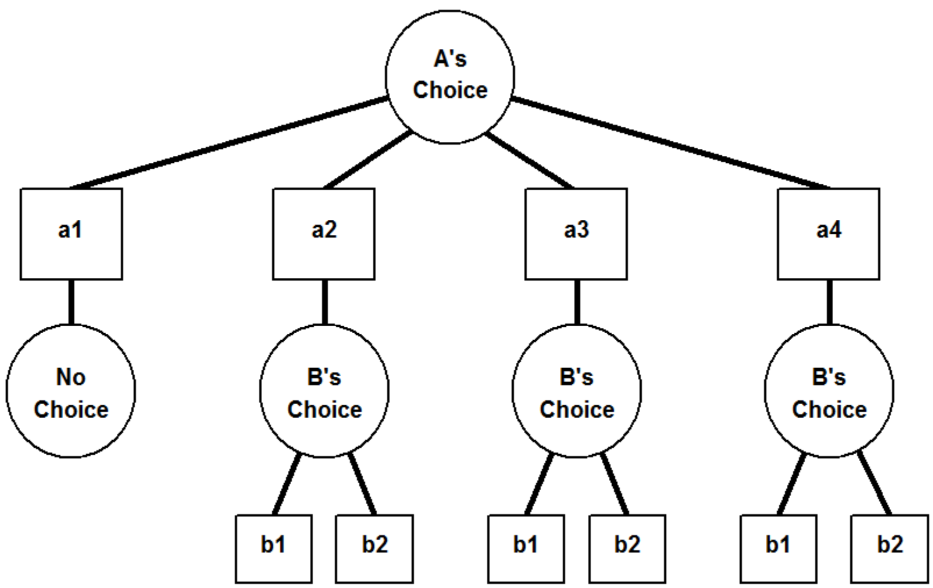

	a1	a2 b1	a2 b2	a3 b1	a3 b2	a4 b1	a4 b2
A's Payoff	24.0	32.0	20.0	44.0	12.0	56.0	0.0
B's Payoff (if human)	10.0	14.7	28.0	24.0	48.0	28.0	53.4

Figure 3: Payoffs in Round 3 of a match

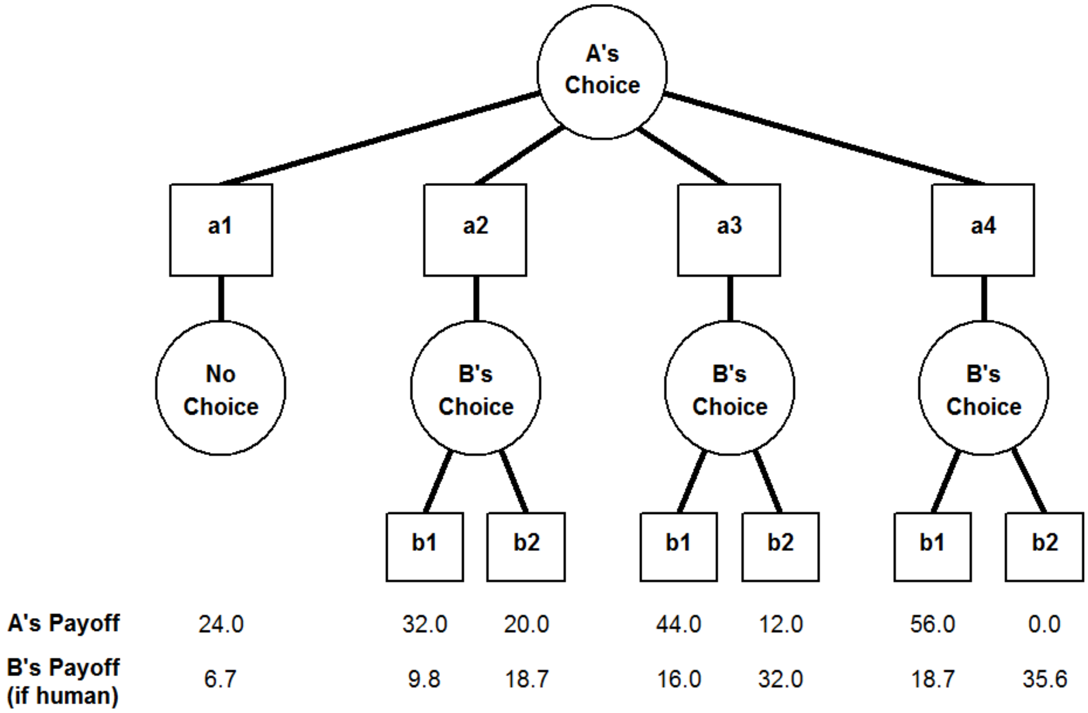

Information

- After each round of a match, all participants will be informed about their own payoff.
- Nobody will ever be informed about the identity of the participant they are (or were) paired with. In particular, A participants will not be informed about whether the B participant they are paired with is human or a computer program.
- However, when an A participant, after choosing a2, a3, or a4 in any round of a match, observes that a B participant reacted with choice b2, then the A participant can be certain that he/she is matched with a human B participant (as a computerized B participant will always choose b1).

The Decision and Information Screen

Figure 4 shows a screenshot of the decision and information screen that you will see and use during the experiment.

- On the right hand side of the screen, you will find the decision-making box. It shows a graphical representation of the experimental decision problem you will face in each round of a match.
  - The A participants make a choice by clicking on a1, a2, a3 or a4. The decision can be changed by clicking on a different choice. Choices are finalized by clicking the Continue button.

- After A participants have finalized their choices, B participants will be informed about this choice (B participants will see the choice of the A participant highlighted in their decision box).
- Human B participants then make a choice by clicking on b1 or b2 below the A participant's choice (except when the A participant has made the choice a1, in which case the B participant has no choice to make). B participants can change their decision by clicking on a different choice. Choices are finalized by clicking the Continue button.

- In the upper left hand corner of the screen, you find the information box. This box will indicate your role in the experiment (A or B role), the number of the current match, and the number of the round in the current match.
- In the lower left hand corner of the screen, you find the history box. As of round 2 of a match, the history of what happened in all previous rounds of the current match will be displayed here. In particular, this box will display the choice of the A participant, the choice of the B participant and your own earnings for every round of the current match.

Figure 4: The decision and information screen

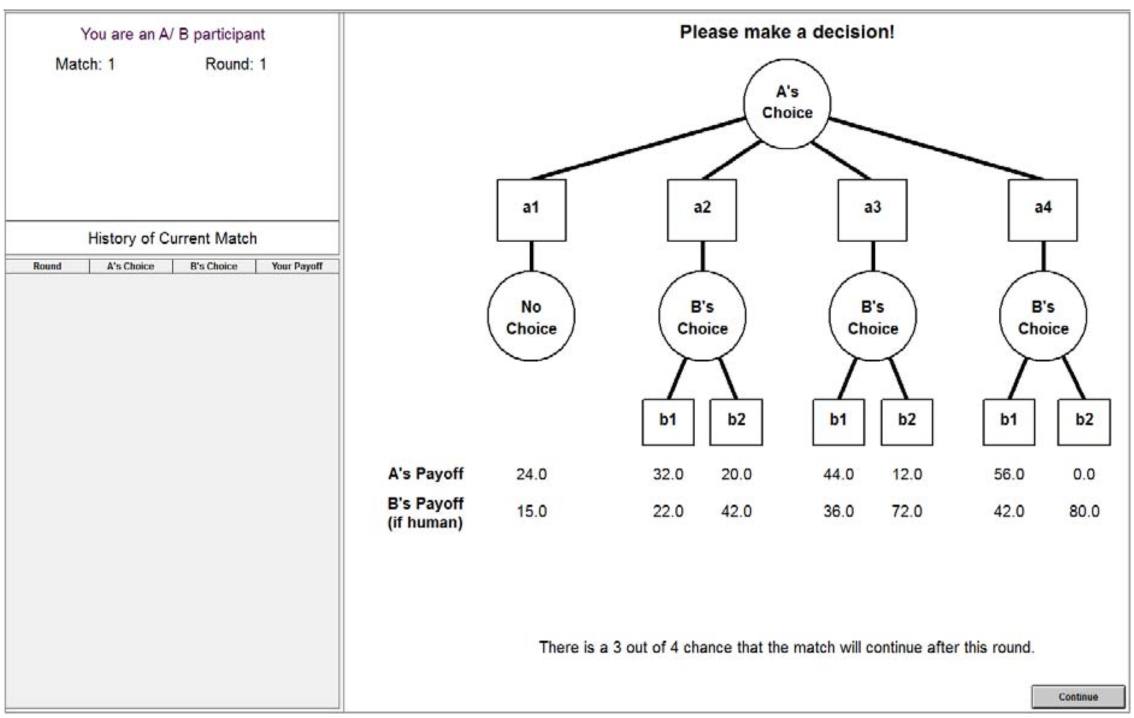

Number of matches

The experiment will end after the first match that is completed after 75 minutes have passed, or after 25 matches have been completed, whichever is the sooner.

Determination of your Earnings

Points will be converted to Euros at the following rate, (100 points ≡ 1 EURO). At the end of the experiment, your total points from all rounds in all matches will be converted into euros at the rate specified above, and will privately be paid to you in cash.

Summary

1. The experiment consists of various matches, each of which consists of a randomly determined number of rounds.
2. There are 8 A participants and 8 B participants. Of the 8 B participants, 7 will be human participants and 1 will be represented by a computer program.
3. Within a match:
    a. At the beginning of a match, each A participant is randomly paired with a B participant (either human or computerized). The A participants will not be informed about whether they are paired with a human or the computerized B participant.
    b. In each round of a match:
        i. The A participants make a choice first.
        ii. After learning the choice of the A participant they are paired with, the B participants make a choice.
        iii. All participants will be informed about what happened in their own pair.
    c. The payoffs available to B participants reduce by 1/3 every round.
4. After one match ends, all A and B participants are randomly paired again and a new match begins.

INSTRUCTIONS

- Welcome to this experiment!
- During the course of the experiment, you will be asked to make a series of decisions. If you follow the instructions carefully, you can earn a considerable amount of money, depending on your decisions and those made by other participants. The money you earn will be paid to you in cash at the end of the experiment.
- Please remain silent during the experiment. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you.

Roles

- There are 15 people who participate in this experiment.
- Each of you will be randomly assigned to a role which could be either a so-called A or B participant. You will remain in the same role throughout the entire experiment. You will be informed about your role on your computer screen when the experiment begins.
- In particular, there will be 8 human participants assigned to the A role, and 7 human participants assigned to the B role. Decisions for an 8^{th} B role will be made by a computer (this role will be referred to as "the computerized B participant").

Matches

- The experiment is divided into various sequences, each of which consists of various rounds. Each sequence of rounds is referred to as a match.
- The length of a match is randomly determined. After each round of a match, there is a 3 out of 4 chance that the match will continue for another round. So, for instance, if you are in round 1 of a match, the chance that there will be a 2^{nd} round is 3 out of 4, and if you are in round 2 of a match, the chance that there will be a 3^{rd} round is also 3 out of 4, and so on.

Pairing

- At the beginning of a match, each A participant will be paired with a B participant and will remain with that B participant until the end of the match.
- The B participant may be human or computerized. The A participants will not be informed who they have been paired with and, in particular, they will not know whether they are paired with a human B participant or the computerized B participant. As there are 7 human B participants and 1 computerized B participant,

the chance of an A participant to be paired with a human B participant is 7 out of 8 and the chance to be paired with the computerized B participant is 1 out of 8.

- Once a match ends, each A participant will again be randomly paired with a B participant.

Within a Match

Each round of a match will proceed as follows:

- Each A participant begins the round by choosing one of two alternatives. These alternatives are labeled a1 and a2, and can be seen in Figure 1.
- The choice of an A participant is then shown to the B participant he/she is paired with. As you can see in Figure 1, if A's choice is a1, there is no decision to be made by the B participant. However, if the choice of the A participant is a2, then a human B participant can decide between choosing b1 or b2.
- **The computerized B participant is "programmed" to always choose b1 after a choice of a2 of the A participant**.
- At the end of each round, each participant receives a payoff according to the decisions he/she and the participant with whom they were matched made, as shown in Figure 1.

Figure 1: Payoffs in Round 1 of a match

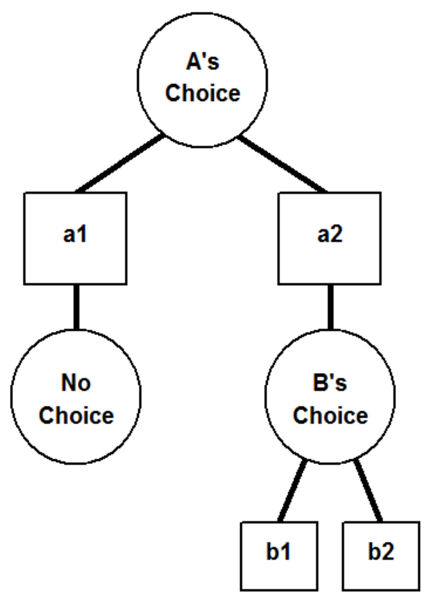

	a1	a2, b1	a2, b2
A's Payoff	24.0	56.0	0.0
B's Payoff (if human)	15.0	42.0	80.0

There is one more important rule concerning the payoffs:

- Starting in round 2 of a match, the possible payoffs as indicated in Figure 1 for a human B participant are reduced by a factor of 1/3 in each round.
- More precisely, in round 1 of a match, the payoffs for A and human B participants are as indicated in Figure 1. In Round 2 of a match, the payoffs for A and human B participants are as indicated in Figure 2, while in round 3 of a match, the payoffs for A and human B participants are as indicated in Figure 3, and so on.
- The payoffs will always be rounded to the nearest decimal place.
- In the decisions screen, which will be explained below, the payoffs in the current round of a match will always be indicated.
- The payoffs for A remain the same in every round of a match.
- Computerized B participants receive no payoff.

Figure 2: Payoffs in Round 2 of a match

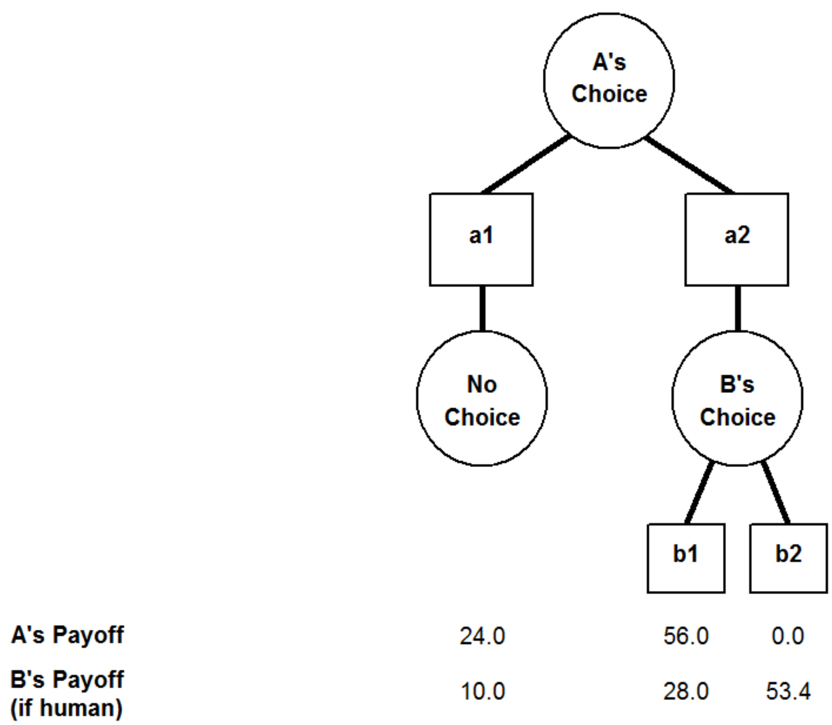

Figure 3: Payoffs in Round 3 of a match

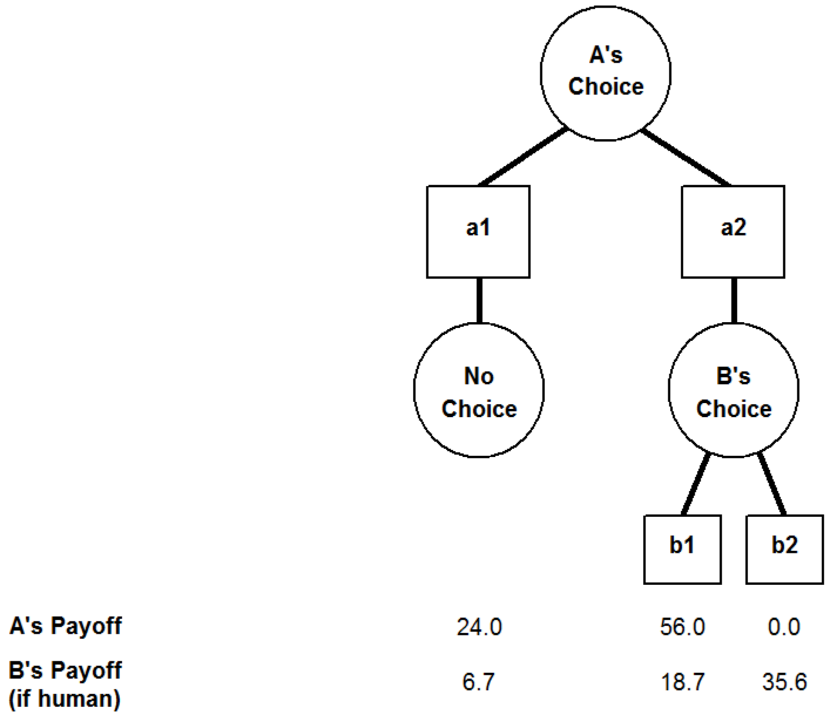

Information

- After each round of a match, all participants will be informed about their own payoff.
- Nobody will ever be informed about the identity of the participant they are (or were) paired with. In particular, A participants will not be informed about whether the B participant they are paired with is human or a computer program.
- However, when an A participant, after choosing a2 in any round of a match, observes that a B participant reacted with choice b2, then the A participant can be certain that he/she is matched with a human B participant (as a computerized B participant will always choose b1).

The Decision and Information Screen

Figure 4 shows a screenshot of the decision and information screen that you will see and use during the experiment.

- On the right hand side of the screen, you will find the decision-making box. It shows a graphical representation of the experimental decision problem you will face in each round of a match.
    - o The A participants make a choice by clicking on a1 or a2. The decision can be changed by clicking on a different choice. Choices are finalized by clicking the Continue button.

- o After A participants have finalized their choices, B participants will be informed about this choice (B participants will see the choice of the A participant highlighted in their decision box).
- o Human B participants then make a choice by clicking on b1 or b2 below the A participant's choice (except when the A participant has made the choice a1, in which case the B participant has no choice to make). B participants can change their decision by clicking on a different choice. Choices are finalized by clicking the Continue button.

- In the upper left hand corner of the screen, you find the information box. This box will indicate your role in the experiment (A or B role), the number of the current match, and the number of the round in the current match.
- In the lower left hand corner of the screen, you find the history box. As of round 2 of a match, the history of what happened in all previous rounds of the current match will be displayed here. In particular, this box will display the choice of the A participant, the choice of the B participant and your own earnings for every round of the current match.

Figure 4: The decision and information screen

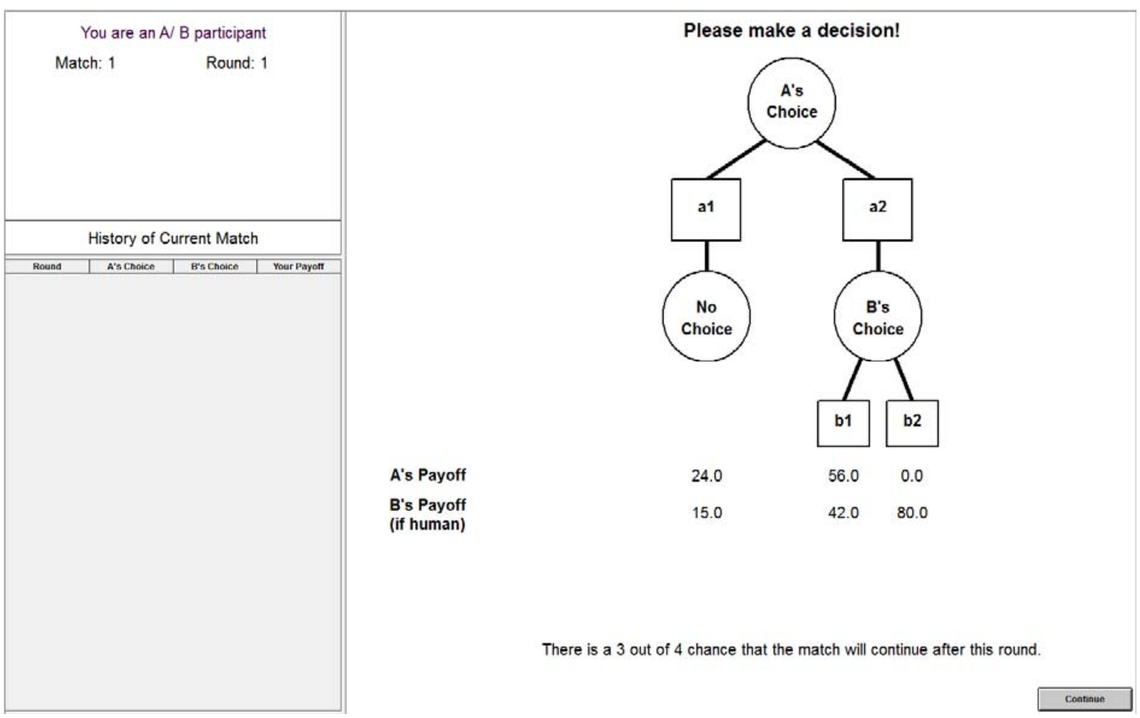

Duration of the Experiment

The experiment will end after the first match that is completed after 75 minutes have passed, or after 25 matches have been completed, whichever is the sooner.

Determination of your Earnings

Points will be converted to Euros at the following rate, (100 points ≡ 1 EURO). At the end of the experiment, your total points from all rounds in all matches will be converted into euros at the rate specified above, and will privately be paid to you in cash.

Summary

1. The experiment consists of various matches, each of which consists of a randomly determined number of rounds.
2. There are 8 A participants and 8 B participants. Of the 8 B participants, 7 will be human participants and 1 will be represented by a computer program.
3. Within a match:
    a. At the beginning of a match, each A participant is randomly paired with a B participant (either human or computerized). The A participants will not be informed about whether they are paired with a human or the computerized B participant.
    b. In each round of a match:
        i. The A participants make a choice first.
        ii. After learning the choice of the A participant they are paired with, the B participants make a choice.
        iii. All participants will be informed about what happened in their own pair.
    c. The payoffs available to B participants reduce by 1/3 every round.
4. After one match ends, all A and B participants are randomly paired again and a new match begins.