Experimental Examination of Behavior in a Sequential versus Simultaneous Trust Game

Maroš Servátka [1] , Steven Tucker [1] and Radovan Vadovič [2]

[1] Department of Economics, University of Canterbury, New Zealand
[2] Department of Economics, Instituto Tecnologico Autonomo de Mexico

Keywords: *Experimental economics; Trust; Signaling; Guilt Aversion*

EXTENDED ABSTRACT

We conduct an experiment to examine the strategic use of trust in an environment similar to Berg, Dickhaut, and McCabe (1995) investment game. The environment differs in that the second mover is restricted to the binary choice of returning half of the tripled amount (fair split) or zero (selfish split). We use the theory of guilt aversion to explain the behavior in strategic and non-strategic environments represented by playing the game sequentially and simultaneously respectively. We find that in the sequential treatment first movers invest significantly more than when the transfer decisions are conducted simultaneously. Moreover, in line with the theoretical prediction, 91% of subjects who invested the entire endowment received half of the surplus. On the other hand only 5% of subjects who invested anything less than the entire endowment received half. In the simultaneous treatment the proportions yield 11.1% and 32.1% respectively. These allocations along with the beliefs collected in a salient manner are consistent with the predictions of guilt aversion.

1. INTRODUCTION

A number of studies have shown that trust plays a critical role in economic relationships. Consider a simplest form of a trust (or a hold-up) situation (Williamson, 1975). Initially, an investor has an option to invest towards a project which an entrepreneur will carry out. If the investor invests, a surplus is created on the side of the entrepreneur. The question is how she will divide this surplus between herself and the investor. A unique subgame-perfect equilibrium paints a gloomy picture: the entrepreneur keeps the whole surplus and the investor does not invest, i.e., a relationship does not form. However, in real life, these types of situations occur quite frequently and parties to a relationship are able to overcome the incentive problem even in the absence of contracts. Berg et al. (1995) have studied the trust game in laboratory conditions. They found significant deviations from classical equilibrium behavior: often it is the case that the investor invests and the entrepreneur chooses a fair division of the surplus. A leading explanation is trust. The investor trusts the entrepreneur to divide the surplus fairly, which prompts her to invest. The entrepreneur, in turn, rewards her mind-set (trusting behavior) with a fair division of the surplus. This evidence suggests that trust could play an important role in fostering relationships.

The objective in this paper is to explore strategic implications of trust. The investor in the trust game could decide to invest for two different reasons. First, she could be a generally "trusting" (type of) person who believes that most people are fair, and accordingly, she expects the entrepreneur to split the surplus fairly. Her belief about the proportion of fair individuals in population is purely subjective and summarizes her own experiences and biases. Secondly, she could not be a trusting person; she believes there are sufficiently many people in a population who are opportunistic (selfish) and would keep the whole surplus. Her beliefs about the average response of the entrepreneur are not high enough on their own to induce investment. However, she is aware of the fact that some of the opportunistic entrepreneurs could interpret investment as a sign of trust and reward this trust with a fair split. Thus, if she invests her chances of receiving a fair split rise due to the proportion of entrepreneurs who reward trust when they see it. This may give the investor sufficient incentives to invest. The main difference between the first and the second motivation is that in the first the investor's degree of trust is directly tied to her inherent initial belief about the amount of fairness in the population of investors. In the second, on the other hand, she exploits the signaling value of investing as a way of communicating trust to the entrepreneur. The investor counts on the fact that the entrepreneur rewards her with a fair split. This raises a following question: to what extend is trusting behavior motivated by the investor's inherent belief and to what extend is it a strategic decision relying on the value of signaling trust. It is this question that we hope to settle in this paper.

Although there is currently a large body of literature exploring various aspects of trust, there is no uniformly accepted theory about what trust is and how it originates. Several models attempt to explain trust as a product of rational behavior. For example, Dufwenberg (2002), Dufwenberg and Gneezy (2002), Charness and Dufwenberg (2006), and Battigalli and Dufwenberg (2007) rely on the theory of guilt aversion; Falk and Fischbacher (2006) make argument based on reciprocity; and Sliwka (forthcoming) defines conformist-types of individuals who trust only when they believe others do. Here we adopt the view of Battigalli and Dufwenberg (2007) who interpret trust in the framework of guilt-aversion. This story has already received some empirical support in experimental studies of Dufwenberg and Gneezy (2002), Charness and Dufwenberg (2006), Schnedler and Vadovič (2007), and Dufwenberg et al. (2007), and it allows for introducing strategic considerations of trust in a tractable way. The main idea is that if the entrepreneur is guilt-averse then she will experience a disutility from feeling guilty whenever she "hurts her opponent," i.e., returns the investor less than what was expected. Notice that the entrepreneur's utility depends on her belief about what the investor expects her to do. This feature of the model makes it an example of (dynamic) psychological game. To avoid guilt, a sufficiently guilt-averse entrepreneur will optimally split the surplus in a way that matches her belief about the expectations of the investor. Therefore, in this framework, the investor will trust the entrepreneur to split the surplus fairly only if she is confident that the entrepreneur holds sufficiently high belief about her expectations.

The model of guilt-aversion helps us illustrate the benefits of signaling trust. Initially, the investor has a subjective belief about the share of the surplus she expects to receive. If this belief is high enough, then she optimally invests and if it is small, then she does not. However, we argue that even when the initial belief is small the investor may want to invest. The reason is that objective of the entrepreneur is not to match her initial belief, but the updated belief that takes into account the amount invested. Because the entrepreneur moves second, she observes the amount invested prior to making her decision. Her belief about the share expected by the investor

should then depend on the size of the investment. Higher investment signals a stronger belief in receiving a fair share of the surplus. Hence, the incentives of the guilt-averse entrepreneur to split the surplus fairly should grow in the amount invested.

Notice that the crucial element in this logic is that the investment is used as a (credible) commitment device. The greater the investment the greater the loss to the investor if the entrepreneur decides to keep everything. Because of such credible exposure, it should be unambiguous that the investor has a high expectations. Hence, the entrepreneur should revise her belief upwards and then split the surplus fairly to avoid feeling guilty. This reasoning is also known as psychological forward induction (Duwfenberg and Battigalli, 2007). Thus, what we call "trust signaling" could be understood in the context of guilt-aversion as follows: higher investment credibly signals higher expectations of the investor that are matched by the guilt-averse entrepreneur.

To illustrate the effect of a strategic use of trust on a relationship we run an experiment with two treatments. In both treatments players A and B play a modified investment game. Player A decides how much to invest, i.e., she chooses amount t from the interval between 0 and 10. The invested amount is tripled by the experimenter. Player B then decides whether to return a fair split, 3t/2, or a selfish split, 0, back to the player A. Notice that the difference between this game and the classical investment game is that here B's decision is binary, that is, no splits other than fair and selfish are possible.

The treatments vary in how information is displayed. In the first treatment players A and B play the game sequentially. B makes her decision only after having observed what A has done. In this treatment we expect both the inherent belief and the strategic signaling of trust to matter. In the second treatment, both players make their decisions simultaneously. Therefore, player B must make her decision without knowing the investment decision of player A. This eliminates the possibility of A signaling trust via her high investment. In this treatment, we expect that displayed trusting behavior purely reflects subjects' inherent (initial) beliefs about the proportion of fair individuals in the population. The difference between these two treatments will measure the importance of the strategic use of trust.

Our game differs from the previous experimental literature on guilt-aversion in two major ways. First, our main focus is on the behavior of player A in terms of signaling her beliefs. In contrast, Dufwenberg and Gneezy (2000) measure the correlation between the outside option of player A and the amount player A receives from B in a lost wallet game. Surprisingly, they find there is none. However, their results show a positive correlation between how much player B allocates to player A and B's expectations of how much A expects B to allocate. Charness and Dufwenberg (2006) examine the effects of promises made by player B on the decision of player A. Unlike our game, their game includes a chance move so that it is not detectable whether player B defaulted on her promise or whether it was just a bad state of the world. They find that pre-play communication might influence the motivation and behavior of the subjects by affecting beliefs about beliefs. Dufwenberg et al. (2007) create an environment where subjects can strike an informal agreement about how much should be returned upon investment. If an agreement is made the beliefs of both players coincide. The experimental data reveal that players A are influenced by agreements and invest more often, however a large fraction of players B behaves opportunistically and defaults. Dufwenberg et al. find that only about 1/3 of players B honor the informal agreements. Second, our game is different in structure. It is the player A, who has a rich action set in comparison to rich action set of player B in experiments by Dufwenberg and Gneezy (2000) and Dufwenberg et al. (2007). Charness and Dufwenberg (2006) use a game where choices of both players and nature are binary.

The rest of the paper is organized in the following manner. Section 2 presents the formal model of trust signaling. Section 3 describes the experimental design and procedures. In section 4 we test the hypotheses and discuss the experimental results. Section 5 concludes.

2. Trust Signaling

There are two players A and B. Player A moves first and decides how much of their endowment they want to send to player B, $t \in [0,1]$, and how much to keep for themselves $(1-t)$. The amount sent is tripled when it reaches player B. Player B must decide whether to return a "fair split" of $3t/2$ or a "selfish split" of 0. We study behavior in two different versions of this game with timing of play as our treatment variable. In the first version, the game is played sequentially, and thus B does observe *t* before making her decision. In the second version, the game is played simultaneously, and thus player B does not observe *t*.

Let us suppose that the player B is averse to guilt. This means that her utility depends on

what she believes player A expects her to do. If she falls short of A's expectations, i.e., she chooses the selfish split when A expected to receive a fair amount, then she experiences a feeling of guilt which is proportional to the difference of what was expected and what was received by A. To be a little more precise, denote α as the initial belief that player A assigns to what B is going to do, i.e., $\alpha = \Pr(\text{"fair split"})$. And, let α_B denote B's estimate of α. If B chooses a selfish split, then she disappoints A (lets A down) in the amount $\alpha(3t)/2$ and experiences guilt in the magnitude $\theta_B \alpha_B (3t)/2$, where θ is the guilt-aversion parameter and $\theta \in [0,1]$. On the other hand, if she sends $3t/2$ back to A, then she avoids feeling guilty but incurs a monetary cost in the amount $3t/2$. Hence, for a given belief α_B, B will choose a fair split if

$$\theta_B \alpha_B (3t)/2 \geq (3t)/2$$
$$\alpha_B \geq 1/\theta_B,$$

that is, if α_B is sufficiently high. If α_B is low, she will keep everything. The decision of the player A is based on her own belief, α, of what B is going to do. If $\alpha \geq 1/\theta_B$ (and A assumes that α_B is a correct estimate of α) then A should be confident that B will chose the fair split and hence she should send $t = 1$. If $\alpha \geq 1/\theta_B$ then she should send 0.

We focus first on what happens in the sequential game when *t* is observed by B before making her decision. Now, player A may be able to "communicate" though her action, *t*, how confident she is about receiving $3t/2$. Notice that given her belief, α, A's payoff is given by

$$\alpha\left(1 - t + (3t)/2\right) + (1-\alpha)(1-t).$$

A quick look at this expression reveals that her payoff is increasing in *t* if

$$\alpha \geq 2/3,$$

and decreasing otherwise. Hence, if A trusts B sufficiently to choose the fair split, then A will maximize her payoff by sending $t = 1$ to B. On the other hand, if A is doubtful about receiving a fair share, i.e., $\alpha < 2/3$, then she should send nothing to B. However, since *t* is now observable, B can use it to infer more information about α and revise her own belief α_B. Let $\hat{\alpha}_B = \alpha_B(t)$ be the updated belief of B after having observed *t*. In particular, if B observes $t = 1$, then she should revise her belief to $\hat{\alpha}_B \geq 2/3$ and, vice versa, if $t = 0$ was observed, then B's belief should be $\hat{\alpha}_B < 2/3$. Next, if we assume that θ_B is sufficiently high, i.e., $\theta_B \geq 3/2$, then this implies that after observing $t = 1$, it must be that $\theta_B \geq 3/2 \geq 1/\hat{\alpha}_B$. This is sufficient to induce B to choose the fair split. But then, player A should be confident to get the fair return after sending $t = 1$. As a result she will always have an incentive to send $t = 1$. Thus, when *t* is observable, our theory predicts a single outcome ($t = 1$, Fair split). In other words, player A signals her high expectations to player B who will then match these expectations by choosing the fair split.

Next let us examine what happens in the simultaneous game where *t* is not observed. Both players will face some uncertainty about their respective beliefs. Player B bases her decision on her own guess of A's expectations, α_B. On the other hand, player A's expectations, α, will correspond to her guess about what B is planning to do. Each player's belief is subject to her own experiences and biases. Because of this, the beliefs of both players will most likely not be identical in the environment of one-shot games, as in our design, where it is impossible to learn/identify behavioral traits of their counterparts. If we restricted ourselves to only analyze equilibrium behavior in which the beliefs coincide, i.e., $\alpha = \alpha_B$, then we may falsely reject a true hypothesis. Therefore, we will focus on whether the decisions and beliefs of the individual are consistent with the theory of guilt aversion, and not specifically whether they are consistent with their counterpart's decision and beliefs. When players act optimally subject to their own beliefs, our theory predicts four different kinds of outcomes: ($t = 1$, Fair split), ($t = 1$, Selfish split), ($t = 0$, Fair split), ($t = 0$, Selfish split).

3. Procedures and Experiment

The general structure of the trust game is similar to Berg et al. (1995). In the first stage of each trust game, players A were endowed with $10NZ. They had to decide how much of this endowment they wanted to keep for themselves and how much to transfer to their anonymous player B counterpart. This was done by circling one of the whole numbers ranging from zero to ten on their decision sheet. It was common knowledge that any amount transferred by player A would be tripled by the experimenter. That is, players B would receive three times the amount that their

player A counterpart transferred to them. In the second stage, players B must decide how much of the tripled amount they want to keep for themselves and how much to transfer back to their player A counterpart. This decision is restricted to a binary choice of either half or zero. Just as for players A, this decision was done by circling one of the two choices on their decision sheet.

We have two treatments in the experiment, i.e., sequential (SEQ) and simultaneous (SIM) play of the trust game. Four sessions in total were conducted for each treatment. The sequence of events in a session was the following. (1) A coin was flipped to determine player types. (2) The instructions were read aloud for the subjects, who followed along with their own copy. To assist in their understanding, a copy of the instructions was also placed on an overhead and any decisions sheets, tables, etc… were illustrated specifically. The subjects were encouraged to ask questions relating to the rules of the game at any time. (3) Both player types completed the belief elicitation task. (4) The experimenter collected the belief decision sheets and distributed the trust game decision sheets. (5) The sequence of events differed slightly between sessions implementing the sequential and simultaneous trust games. In the sequential trust game sessions, players A first made their transfer decision to players B. All decision sheets were collected and the amount transferred from players A were copied to their counterpart players' B decision sheets, which were then returned to players B. Presented with the decision of their player A counterpart, players B made their decision on whether to return half or zero. The experimenter collected all decision sheets, transferred the decision information of players B to their player A counterparts' decision sheet, and returned the decision sheets to all players to reveal their earnings. In the simultaneous trust game sessions, both player types of participants made their transfer decisions simultaneously. The experimenter collected all decision sheets, transferred the decision information each decision sheet to their counterparts', and returned the decision sheets to all players to reveal their earnings. (6) Subjects completed a short survey on the experiment and general demographic information for which they were paid $5 instead of a show up fee. This was not announced to the subjects at the start of the experiment. (8) Subjects were privately paid their earnings for the session.

4. Results

The presentation of results is divided into two parts. In section 4.1 we discuss the findings using the decisions made by subjects in our experiment. Section 4.2 uses beliefs elicited during the experiment to ascertain whether there is further empirical support for trust signaling as explained by guilt aversion theory.

4.1. Decisions

The neoclassical subgame perfect equilibrium for both SEQ and SIM is for all players A to send $t = 0$ and all players B to return *Zero*. The number of decisions consistent with this equilibrium is very different for players A and B across treatments. Players A sent $t = 0$ only 5 out of 41 (12%) instances in SEQ and 4 out of 37 (11%) instances in SIM. On the other hand, players B returned *ZERO* 21 out of 41 (51%) instances in SEQ and 27 out of 37 (73%) instances in SIM. Nevertheless, the subgame perfect equilibrium predictions for self-regarding players do not find much support in our data.

We now explore behavioral patterns in the data given the theory of trust signaling proposed in section 2. Among the 41 pairs in SEQ, 21 (51%) of the players A sent $t = 1$ and 20 (49%) of the players B returned *HALF*. Note that this analysis does not do justice to the theory of trust signaling since player A must send a high *t* in order to induce player B to return *HALF*. Therefore, if we focus only on those 21 pairs in which player A sent $t = 1$, 19 (91%) of the player B counterparts returned *HALF*, and thus there is strong support of the theory.

Since (i) two of these outcomes involve player A sending $t = 0$ and (ii) signaling is not possible due to simultaneous decision making, we should observe players A sending greater *t* in SEQ than SIM. The mean *t* sent by players A in SEQ and SIM was 6.59 and 5.22 respectively. A one-sided Mann-Whitney test indicates that they are significantly different at the 5% level (p=0.046). This result is also supported by comparing across treatments the number of players A who *signal trust* by sending $t = 1$. In SEQ, 21 out of 41 (51%) players A sent $t = 1$ versus 9 out of 37 (24%) in SIM. According to a 1-sided Fisher's exact test, the fraction of players A sending $t = 1$ is significantly higher in SEQ at the 5% level (p=0.013).

We now turn our attention to the behavior of players B. We expect that significantly more players B return *HALF* in SEQ than in SIM. If this were not the case, then players A are making a serious mistake and would be much better off if they deviated to $t = 0$. First, we examine the frequency of players B returning *HALF* on the aggregate level. In SEQ, players B returned *HALF* 49% (20 out of 41) of the time, whereas in SIM

players B returned *HALF* only 27% (10 out of 37) of the time. According to a 1-sided Fisher's exact test, the frequency of *HALF* in SEQ is significantly greater than SIM at the 5% level (p=0.040).

Second, we study the behavior of players B conditional on her paired player A's decision. As discussed previously, 19 out of 21 (91%) of players B returned *HALF* in SEQ when player A sent $t = 1$. In only 1 out of 20 instances (5%) did player B return *HALF* when player A sent $t < 1$. Obviously in SEQ, the decision of player B of whether to return *HALF* or *ZERO* depended heavily upon the observed decision of player A. This clear pattern is not present in the SIM data where t is not observable to players B before they make their decisions. When player A sent $t = 1$ in SIM, only 1 out of 9 (11%) players B returned *HALF* compared to 9 out of 28 (32%) when player A sent $t < 1$. Since we are not limiting ourselves to the theoretical predictions for this treatment, we use a 2-sided Fisher's exact test. The test does not detect a significant difference between players' B decision when players A sent $t = 1$ or $t < 1$ in the SIM treatment (p=0.393).

4.2. Beliefs

In this section, we first check whether the beliefs of players A on average match the observed behavior and the beliefs of players B about players' A beliefs. In SEQ, 51% of players B returned HALF. Players A on average believed that 51% of players B would return HALF, and players B on average believed that players A believed that 51% of players B would return HALF. Thus in SEQ data, players' A beliefs are not significantly different from the actual choices (p=0.78) nor from players' B beliefs (p=0.81).

Along the same lines, 27% of players B in SIM returned HALF. Players A on average believed that 46% of players B would return HALF, and players B on average believed that players A believed that 37% of players B would return HALF. Players' A beliefs are significantly higher from the actual choices of players B (p<0.01). On the other hand, players' A beliefs are not significantly different from B's beliefs (p=0.13), although as the relatively low p-value indicates, this result is not as strong as in SEQ.

Next, we compare the beliefs across treatments to verify whether our conclusions from the previous section are supported. Recall that players A choose on average higher t in SEQ than in SIM. Given that, our theory predicts that the beliefs of players A in SEQ should be higher than those in SIM. Similarly, the beliefs of players B in SEQ should be higher than those in SIM.

The average belief of players A in SEQ is higher than in SIM. We test whether difference is significant using 1-sided Mann-Whitney test. The hypothesis does not receive considerable support in the data (p=0.264). However, the direction is correct. On the other hand, the hypothesis that the beliefs of players B is higher in SEQ than in SIM is supported at p<0.01 significance level.

The two most appropriate tests of trust signaling using beliefs are the comparison of beliefs of players A who chose $t = 1$ versus those who chose $t < 1$ and of players B who returned ZERO versus those who returned HALF in treatment SEQ. The theory requires the following two statements to be true:

$$\text{Beliefs } A^{SEQ}\big|t = 1 > \text{Beliefs } A^{SEQ}\big|t < 1$$

$$\text{Beliefs } A^{SEQ}\big|\text{HALF} > \text{Beliefs } A^{SEQ}\big|\text{ZERO}.$$

The average belief of players A who chose $t = 1$ in treatment SEQ is equal to 64.19 where as of those who chose $t < 1$ is equal to 36.40. A Mann-Whitney test detects a significant difference between the two samples at a p<0.01 level.

The average belief of players B who chose HALF and ZERO in treatment SEQ is 53.95 and 46.67, respectively. The data includes a pair of outliers. Since the beliefs are bounded by 0 from below and 100 from above, we accepted a data point to be an outlier if it is more than two standard deviations from the mean. We find one such outlier among the players B who returned HALF (Belief = 15) and another among players B who returned ZERO (Belief = 95). After removing them from the sample, the averages become 56.00 and 45.30, respectively. A 1-sided Mann-Whitney shows that the beliefs of players B who returned HALF are higher at p=0.058 significance level. We conclude that the guilt aversion prediction for the presence of trust signaling in supported by the beliefs of both types of players in our experiment.

5. Discussion

We set out to study trust in two environments that allow different degrees of strategic behavior. In the first environment players A and B make decisions in a trust game simultaneously and in the second they make decisions sequentially. The main difference is that in the first game the trust of the first mover A in B's fair response derives purely from A's personal, subjective (inherent) belief about the likelihood of fair response in a population. In the second environment B observes whether he was trusted or not before he makes his choice and hence his response will likely depend on A's action. We find

there is more to trust than meets the eye. Our model, which relies on guilt-averse preferences, predicts that if agents are sufficiently sensitive to trust (i.e., guilt), then B will always behave fairly if he observed he was trusted. This is not necessarily so when B does not observe what A has done. Hence, we expect that in a sequentially played trust game A invests more often in order to induce B to behave fairly than if the game is played simultaneously. In other words, A uses trust strategically (or signals trust) in the former and not in the latter environment.

The results of our experiments for the most part confirm our conjectures. We find that the first movers invest significantly more in the sequential than in the simultaneous trust game. Similarly the second movers respond fairly in significantly higher proportion in the sequential versus simultaneous trust game. In our model changes in behavior are driven by the heterogeneity in players' beliefs. To verify that our model fits the data we elicit beliefs. As expected, players have on average higher beliefs in the sequentially than in the simultaneously played trust game. Our results indicate that in the environment that allows players to credibly communicate trust, i.e., where one player is able to signal trust in another player, the relationship has higher chances of forming than in the environment where this is not true. Our results are relevant from theoretical standpoint but also from he point of view of designing institutions. We found that in certain environments (i.e., that allow trust-signaling) the two involved parties can create trust and hence form relationship on their own. Thus we advocate designing institutions that allows full or partial display of trusting behavior because this would, in line with our findings, enhance the efficiency of exchange.

6. References

[1] Battigalli, P. and M. Dufwenberg, "Dynamic Psychological Games," University of Arizona working paper, 2005.

[2] Battigalli, P. and M. Dufwenberg, "Guilt in Games," University of Arizona working paper, 2007.

[3] Berg, J., J. Dickhaut, and K. McCabe, "Trust, Reciprocity, and Social History," *Games and Economic Behavior*, July 1995, Vol. 10(1), 122-42.

[4] Brandts, J. and G. Charness, "Hot and Cold Decisions and Reciprocity in Experiments with Sequential Games," *Experimental Economics*, 2000, Vol. 2(3), 227-238.

[5] Charness, G. and M. Dufwenberg, "Promises & Partnership," *Econometrica*, 2006, Vol. 74(6), 1579-1601.

[6] Cox, J. C., "How to Identify Trust and Reciprocity," *Games and Economic Behavior*, 2004, Vol. 46, 260-281.

[7] Cox, J. C., "Implications of Game Triads for Observations of Trust and Reciprocity," University of Arizona discussion paper, September 1999, revised 2000.

[8] Cox, J. C. and C. A. Deck, "On the Nature of Reciprocal Motives," *Economic Inquiry*, 2005, Vol. 43(3), 623-635.

[9] Dufwenberg, M., "Marital Investment, Time Consistency & Emotions," *Journal of Economic Behavior and Organization*, 2002, Vol. 48, 57-69.

[10] Dufwenberg, M. and U. Gneezy, "Measuring Beliefs in an Experimental Lost Wallet Game," *Games and Economic Behavior*, 2000, Vol. 30, 163-82.

[11] Dufwenberg, M., M. Servátka, and R. Vadovič, "Handshakes," University of Canterbury working paper, 2007.

[12] Falk, A., E. Fehr, and U. Fischbacher, "On the Nature of Fair Behavior," *Economic Inquiry*, 2003, Vol. 41(1), 20-26. 9

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

[14] Fehr, E. and S. Gächter, "Fairness and Retaliation: The Economics of Reciprocity," *Journal of Economic Perspectives*, 2000, Vol. 14, 159-191.