Legitimacy of Control

Wendelin Schnedler*
Department of Economics
University of Heidelberg

Radovan Vadovic
Department of Economics
ITAM, Mexico City

November 8, 2010

Abstract

Control sometimes triggers negative responses. While there is empirical evidence for such negative reactions and theories that can explain them, it remains to be examined *when* they occur. We conjecture that these negative responses disappear if control is legitimate, i.e., if it averts anti-social behavior. Specifically, we predict that fewer individuals respond negatively to control if control prevents selfishness or theft. We confirm these predictions in an experiment.

JEL C9, M5, C7
Keywords: moral-hazard, intrinsic motivation, legitimacy, hidden costs of control

*Accepted at the Journal of Economics & Management Strategy. We like to thank to Gary Charness, Peter Dürsch, Martin Dufwenberg, Armin Falk, Price Fishbeck, Petra Hagemann, Bernd Irlenbusch, Michael Kosfeld, Jörg Oechssler, Matteo Ploner, Joop Sonnemans and seminar participants at the IZA/SOLE transatlantic meeting 2007, the ESA meeting 2007 in Rome for their helpful comments. Financial support by the Deutsche Forschungsgemeinschaft (SFB 504) and the Asociación Méxicana de Cultura is gratefully acknowledged.

What would, under ordinary circumstances, be justly condemned as persecution, may fall within the bounds of legitimate self-defence.

Thomas Babington Macaulay, *The History of England*, 1849.

1 Introduction

Whether an activity is regarded as persecution or as legitimate self-defense may depend on the circumstances. For example, in a well-performing, family-based business environment a managerial decision to restrict Internet access may not be considered legitimate. But in the anonymous environment of a large corporation, in which Internet misuse has occurred in the past, the same policy would likely be viewed as prudent and in everyone's best interest.[1] The difference lies in the justification of the policy. In the former case, control implies distrust and in the latter it is a justified prevention. Likewise, a bag search is legitimate if the opportunity to steal valuables is present but not otherwise.[2] In line with these examples, we argue that control is legitimate if it prevents anti-social behavior such as selfishness or theft.

Authorities (governments, firms, organizations, etc.) often appeal to the legitimacy of their behavior when they control.[3] But why do they care whether control is legitimate? A possible reason is that control is rarely complete and individuals may use the remaining latitude to punish the authority if they believe that control is not justified. For example, a worker who believes that monitoring his Internet use is not legitimate may withhold an idea that could improve productivity, exert less voluntary effort or stop cooperating with his supervisor. Apart from the obvious benefit of preventing unwanted behavior, control may thus entail *hidden costs.* In a recent experiment, Falk and Kosfeld (2006) provide indirect but striking evidence of hidden costs: they find in a principal-agent setting that it can be more profitable for the principal to refrain from controlling the agent.

Yet in reality, we observe many circumstances under which control is accepted and does not lead to adverse reactions by those who are restricted. This suggests that agents' responses to control and hence hidden costs vary with the context in which control is used. In this paper, we study the hypothesis that hidden costs are affected by the legitimacy of control. We assume that agents behave pro-socially if control is legitimate but act in their own interest, otherwise. Accordingly, we expect control to entail no hidden costs if and only if it is legitimate.

The idea that the same action may be legitimate under some circumstances but not under others, is not new to economists. Okun (1981) argues that consumers only regard a price increase as unfair if the producer's costs did not increase. Kahneman et al. (1986) find that it is indeed more acceptable for a company to raise prices and for an employer to cut wages if profits are threatened. Finally, in a recent series of studies, Charness and Levine observe that layoffs are more likely to be considered legitimate if product demand is low (Charness and Levine, 2000) and identify when sabotage is considered legitimate (Charness and Levine, 2004). These studies suggest that the legitimacy of the same action could be seen very differently under different circumstances.

Complementing these ideas, we employ the control of a laboratory experiment to vary legitimacy, i.e., whether control prevents anti-social behavior or not. We build on Falk and Kosfeld's findings by using their basic experimental design as our baseline treatment and complementing it with two new treatments. In the baseline treatment, a principal (she) can decide whether to control or not. Then, the agent (he) decides on effort. Control rules out effort below an exogenously given threshold level. However, it leaves the agent the freedom to choose any effort above this level. By exerting less effort when being controlled, the agent can inflict hidden costs of control if he thinks that control is not legitimate.

In our first treatment (ROBOT), we increase the legitimacy of control (relative to the baseline treatment) by augmenting the probability that the principal receives no effort. In half of the cases, the principal is matched with a selfish computerized agent who exerts the smallest possible effort. When deciding on control, the principal does not know whether she faces the selfish computerized agent or a human. This is clearly explained to the human agent in the instructions.

In our second treatment (ENDOWMENT), we affect the legitimacy of control (in relation to the baseline treatment) by changing the framing of both the principal's and agent's decision. The game in this treatment is formally equivalent to that in the baseline treatment. However, we move some of the initial endowment from the agent to the principal and relabel strategies. Control by the principal now prevents stealing.

Hidden costs of control are not formally defined by Falk and Kosfeld (2006) and the presence of such costs is inferred indirectly from the negative overall effect of control. Here, we suggest a definition for hidden costs. This allows us to directly estimate hidden costs and to examine how they change across treatments. In our baseline treatment, a sizable fraction of agents exerts less effort when being controlled; the principal incurs significant hidden costs. Unlike in Falk and Kosfeld (2006), however, these costs are not outweighed by the gains from control. More importantly, hidden costs are significantly lower in the ROBOT treatment, in which the probability that the principal receives no effort is larger. Similarly, when the endowment is shifted from agent to principal and the agent is now seen as taking rather than giving, hidden cost of control drop significantly and are no longer distinguishable from zero. These findings confirm that legitimacy of control is related to hidden costs.

The remainder of the paper is organized as follows. In the next section we describe the game and define hidden costs of control. In Section 3, we present our hypotheses and present our treatments. Section 4 presents our results. Section 5 discusses possible theoretical explanations for our findings and their link to framing; we also explain how our design and findings relate to studies that fail to replicate one of Falk and Kosfeld's results, namely that control reduces rather than increases principal's payoff (Hagemann, 2007; Ploner et al., 2010). Finally, Section 6 concludes.

2 Measuring Hidden Costs of Control

As a starting point, consider the game from the main treatment by Falk and Kosfeld (C10). In this game, there is a principal and an agent. The agent is endowed with $\pi^0_{\mathrm{A}} = 120$ points and the principal has no endowment, $\pi^0_{\mathrm{P}} = 0$. First, the principal decides whether to impose a minimum effort requirement $\underline{x} = 10$, i.e., she can control the agent ($y = \underline{x}$) or not ($y = 0$). Second, the agent decides how much effort, x, to exert, where x is an integer between y and the endowment of 120. The strategy of the agent is a pair $(x^{\mathrm{c}}, x^{\mathrm{nc}})$ that specifies an effort choice when the principal controls and when she does not control. Each unit of effort costs the agent x and increases the payoff of the principal by $2x$. Hence there are efficiency gains if the agent exerts effort. The monetary payoffs are:

$$\pi^1_{\mathrm{A}}(y,x) = \pi^0_{\mathrm{A}} - x \text{ for the agent and} \quad (1)$$
$$\pi^1_{\mathrm{P}}(y,x) = \pi^0_{\mathrm{P}} + 2x \text{ for the principal.} \quad (2)$$

We can distinguish two effects of control:

Definition 1 (Direct and indirect effect of control) *The direct effect of control forces the agent to supply at least the required benchmark, $\underline{x}$, it amounts to:* $2(\underline{x} - x^{nc})$ *if* $x^{nc} < \underline{x}$ *and zero else. The indirect effect of control reflects any other response by the agent and amounts to the difference between overall effect,* $2(x^c - x^{nc})$*, minus the direct effect. It can be more succinctly written as* $2(x^c - \max\{\underline{x}, x^{nc}\})$.

The indirect effect captures the consequences on the principal's payoff due to the agent's psychological reaction to control. A particularly interesting case is when this effect is negative.

Definition 2 (Hidden costs of control) *There are hidden costs of control if the indirect effect of control is negative.*

This definition attempts to capture Falk and Kosfeld's idea of 'hidden costs.' Based on this definition, it becomes possible to estimate hidden costs. Moreover, by separating the direct effect of control from the indirect effect and testing whether the latter is negative, we obtain a procedure to determine whether hidden costs are significant. In addition, the terminology helps to shed light on the inner structure of experimental findings. For example, Falk and Kosfeld's finding that control leads to lower effort than no control can be more specifically described as follows. First, there is an indirect effect of control. Second, this effect is negative, i.e., there are hidden costs; and third, the indirect effect outweighs the direct effect. Finally and most importantly, the formal definitions permit an analysis of the indirect effect even if it does not outweigh the direct effect.

3 Hypotheses and Treatments

Philosophers and social scientists have discussed legitimacy and the use and abuse of power for centuries.[4] Legitimacy is defined as a property of an institution, norm, or authority. Its key aspect is the process of validation, i.e., an agreement amongst the members of the society that the course of action or type of behavior is in line with their moral values and principles of justice—for a review on legitimacy from a psychological perspective see Tyler (2006); for an account of its development see Zelditch, Jr. (2001).

The concept of legitimacy is commonly discussed in the context of freedom of choice and violation of individual property rights. Höffe (1995, p. 36), for example, observes that control "reduces freedom of action, which is a cost or disadvantage to the affected party and which thus requires legitimation." He argues that mutual advantages can provide such a legitimation (Höffe, 1995, p. 40). From an economic perspective, norms such as property rights (which enable investments) or crop sharing (which provides insurance) are mutually advantageous (or welfare-enhancing). However, respecting others' property or sharing with them entails opportunity costs for the individual. Norms thus also need to deter individuals from violating property rights and not-sharing with others.

Our primary thesis results from three main ideas: (i) humans have a tendency to conform to norms that are legitimate (Hoffman, 1977; Suchman, 1996); (ii) anti-social behavior such as selfishness and theft are norm violations and not-legitimate—support for this claim comes from numerous streams of the literature;[5] and (iii) deterrence, punishment, or prevention of non-legitimate acts is in itself legitimate (Tyler, 2008).

These three ideas shape our hypotheses: controlling restricts someone's freedom and thus is not legitimate, unless it prevents anti-social behavior, in which case it becomes legitimate. This statement abstracts from many aspects that are relevant in order to determine whether a certain activity is in accordance with a norm in reality. For the purpose of structuring the data of our experiment, however, it provides a useful first approximation.

Definition 3 (Legitimacy of control) *Control is legitimate if it prevents anti-social behavior.*

Let us now turn to the deterrence mechanism that supports the norm. In infinitely repeated games, this can often be achieved with punishment strategies. A somewhat reduced form to ensure norm-adherence is that subjects act selfishly following non-legitimate actions and in some pro-social manner following legitimate actions; the precise form of pro-social behavior is irrelevant as long as it results in a payoff above the required minimum $\underline{x}$. We summarize this idea in the following basic hypothesis.

Basic Hypothesis 1 (Norm enforcing behavior) *Agents respond with pro-social behavior, $x > \underline{x}$, whenever control is legitimate and with selfish behavior, $x = \underline{x}$, otherwise.*

Similar reduced form hypotheses about subjects preferences that are sensitive to some norm have been evoked in various other contexts.[6] A direct consequence of the legitimacy definition and our basic hypothesis about norm enforcing behavior is that control entails no hidden costs if it prevents selfish behavior or theft. In the following, we suggest two variations of the game by Falk and Kosfeld (2006). In the first, control prevents selfish behavior, in the second theft.

We take the C10 treatment by Falk and Kosfeld as our BASELINE treatment. Our second treatment (ROBOT) varies the BASELINE in that the principal is no longer certain to be matched with a human subject. With a probability of one half, she is matched with a computerized selfish agent that always chooses the smallest possible effort, i.e., $x = 0$ under no control and $x = 10$ under control. Accordingly, the payoffs to the principal are 0 under no control and 20 under control. If the principal interacts with the computerized agent, the human agent is not affected and keeps his endowment. When deciding on control, the principal does not know whether she is matched with the human or with the computerized selfish agent.

The fact that the principal encounters a computerized agent in half of the cases and that this computerized agent only exerts minimal effort is publicly known. Had we chosen a second human agent rather than a computerized agent, the behavior of this agent would not have been publicly known and the beliefs of the first agent about the second agent's behavior would have been crucial (but unobservable to the experimenter). By computerizing the second agent, we thus increase experimental control and reduce unwanted variation.

In the BASELINE treatment, control is legitimate if it prevents selfishness. Some subjects may think that other subjects in the role of the agent are selfish and for them control is legitimate. Other subjects may have a different opinion and for those, control is not legitimate. In the ROBOT treatment, it is publicly known that control prevents selfishness and hence control is legitimate. Accordingly, we expect less punishment following control in the ROBOT than in the BASELINE treatment.

Hypothesis 1 *Control leads to lower hidden costs in the ROBOT than in the BASELINE treatment.*

Our third treatment (ENDOWMENT) differs from the BASELINE treatment in that the agent has the possibility to take from the principal. While in the BASELINE treatment effort ranges from zero to 120 and the minimum requirement is 10, effort in the ENDOWMENT treatment is between -10 and 110 and the minimum requirement is 0. If the agent neither takes nor gives, the principal has an endowment of $\pi^0_{\mathrm{P}} = 20$ and the agent of $\pi^0_{\mathrm{A}} = 110$. The game form of the ENDOWMENT treatment is identical to that of the BASELINE treatment. The only difference are the labels attached to the actions of principal and agent. Unlike in the BASELINE treatment, where control only violates the agent's freedom, it protects the principal's initial endowment in the ENDOWMENT treatment. Accordingly, control is legitimate in the ENDOWMENT treatment and agents should respond less adversely to it.

Hypothesis 2 *Control leads to lower hidden costs in the ENDOWMENT than in the BASELINE treatment.*

In order to test our hypotheses, we ran a total of 12 sessions: 4 for each treatment. We observe 32 principal-agent pairs in the BASELINE treatment, 33 in the ROBOT treatment and 36 in the ENDOWMENT treatment. All sessions were conducted in the experimental laboratory at the University of Mannheim in Fall of 2006. Subjects were primarily undergraduate students who were randomly recruited from a pool of approximately 1000 subjects using an e-mail recruitment system. Each subject only participated in one of the treatments. The software was written in Visual Basic 6 and the experiment lasted approximately 60 minutes (including time for reading the instructions and receiving payments).

After the subjects arrived at the laboratory, they were randomly and anonymously matched in pairs and seated at the computer terminals. They were handed instructions (included in the appendix) and given 15 to 20 minutes to study them. After everyone had finished reading, they were asked to complete a series of questions designed to verify their understanding of the experiment. Once all questions had been answered successfully, the experiment began. The game was not repeated: each subject only had to make one decision. Subjects in the role of the principal had to decide whether to control or not. For subjects in the role of the agent, we used the strategy method to elicit their decisions: the subjects had to decide on the effort level under control and no control.[7] The strategy method allows us to identify and estimate hidden costs. Moreover, we are able to learn agents' choices even if the behavior of principals varies little. At the end of the experiment, we paid each subject privately in cash. All payoffs were initially explained in points that were later converted using the rate that 1 point=10 cents. Subjects received a show-up fee of €4 and their total earnings were on average about €10 for the whole experiment.

4 Results

We first present the treatment effects concerning agents' behavior, which is followed by the discussion of the principals' payoffs and behavior.

4.1 Treatment Effects

For legitimacy to affect hidden costs of control, it is a necessary condition that control leads to hidden costs in the BASELINE treatment. Our first finding indicates that this is indeed the case.

Result 1 *In the BASELINE treatment, control entails hidden costs.*

The left panel in Figure 1 shows the distribution of effort choices under control and no control. Conditional on effort being larger than ten, the distribution under control is stochastically smaller. This indicates that a sizable fraction of subjects provides more effort in the absence of control. The right panel in the same figure depicts the distribution of the individual difference between the effort under control and without control for each subject. It

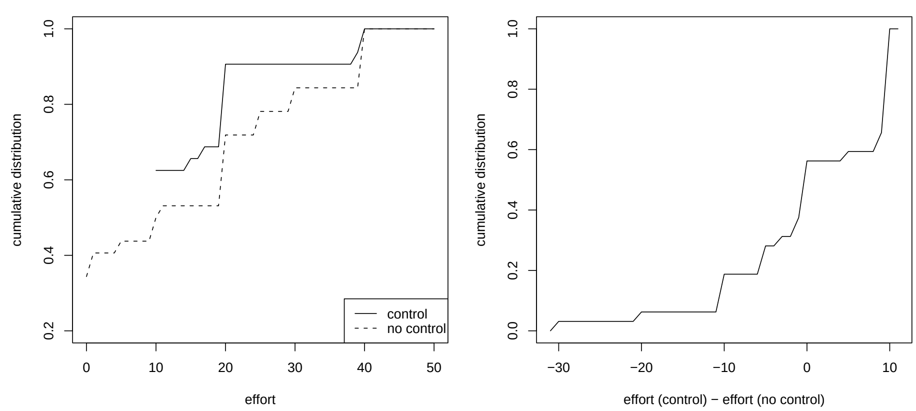

Figure 1: Cumulative distribution of effort (left panel) and of within-subject differences in effort choices for the BASELINE treatment

shows that around 37% of the subjects give less when they are controlled and 18% reduce their effort by ten points or more. We estimate the indirect effect of control on effort by replacing x^{c} and $\max\{x^{\mathrm{nc}}, \underline{x}\}$ with the respective sample means of the variables and find that it amounts to -6.88 points.

On average, principals thus incur hidden costs of about 7 points when controlling. Is this loss statistically significant? We test whether there are hidden costs by checking whether the median of the distribution of x^{c} is smaller than the median of $\max\{x^{\mathrm{nc}}, \underline{x}\}$.[8] In line with Falk and Kosfeld (2006), we find that hidden costs of control are significant (two-sided exact Wilcoxon signed rank test has a p-value below 0.001).

After having established the existence of costs of control, we now address our main hypotheses by examining how these costs change across treatments.

Result 2 *Hidden costs of control are significantly lower in the ROBOT than in the BASELINE treatment.*

In the ROBOT treatment one-third of the subjects voluntarily choose an effort above the minimum requirement even if they are not controlled. The left panel in Figure 2 illustrates that conditional on effort being larger then ten, agents once more choose higher effort in the absence of control. The right panel indicates that the differences between effort under control and in the absence of control are considerably larger in the ROBOT than in the BASELINE treatment. In the ROBOT treatment, only 9% of the subjects choose less effort under control while about 37% do so in the BASELINE treatment. Likewise, the share of subjects who 'punish' control by reducing effort by 10 or more points drops from 18% in the BASELINE to only 6% in the ROBOT treatment. The estimate for the indirect effect of control amounts to -2.06 points. The median subject, however, no longer punishes control and hidden costs are no longer significantly different from zero (the one-sided Wilcoxon signed rank test has a p-value of 0.78). Hidden costs in the ROBOT treatment are less than a third of those in the BASELINE treatment. We can verify whether this reduction is significant by testing whether the indirect effect in the ROBOT treatment is smaller than or equal to that in the BASELINE treatment.[9] Based on the two-sample one-sided test, this hypothesis is rejected (the Wilcoxon rank sum test has a p-value below 0.01). The hidden costs of control thus decrease when moving to an environment where control is not specifically aimed at the agent.

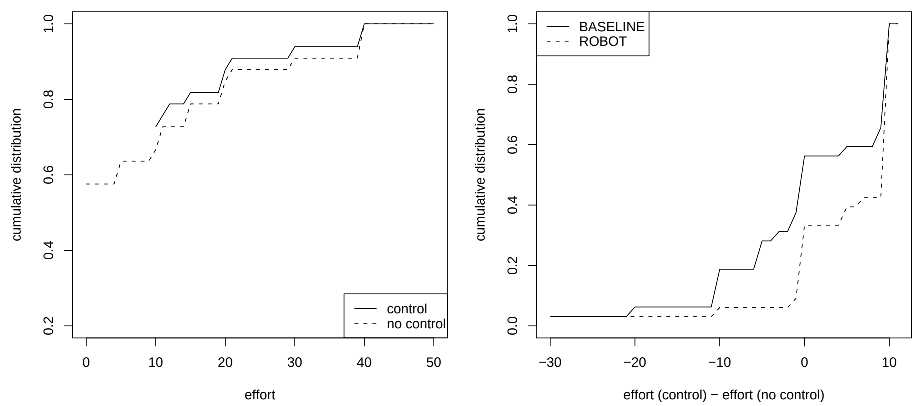

Figure 2: Cumulative distribution of effort for the ROBOT treatment (left panel) and of the differences in effort choices for each subject in the BASELINE and ROBOT treatment (right panel)

Result 3 *Hidden costs of control are significantly lower in the ENDOWMENT than in the BASELINE treatment.*

One-third of the subjects voluntarily give more than the required number of points. Subjects in the role of the agent, however, no longer choose a different effort under control and in the absence of control: the left panel in Figure 3 shows that the distribution of effort is not stochastically smaller when the agent is controlled. The share of subjects that punish the principal for controlling is lower than in the BASELINE treatment and their reduction in effort is smaller too: only 11% choose less effort when being controlled and only about 5% reduce their effort by more than 10—see right panel in Figure 3.

Recall that the indirect effect can be positive if agents choose to exert more effort when they are controlled. If the indirect effect is positive, then there are, of course, no hidden costs of control. This is indeed the case for the ENDOWMENT treatment: the estimate for the indirect effect amounts to 6.11 points. If we use the same testing approach as before, we find that the indirect effect in the ENDOWMENT treatment (and hence the hidden costs of control) is significantly lower than in the BASELINE treatment (Wilcoxon rank sum test has a p-value below 0.01).[10]

There seems to be an increased proportion of agents who only exert the minimum effort in our treatments relative to the BASELINE treatment. This

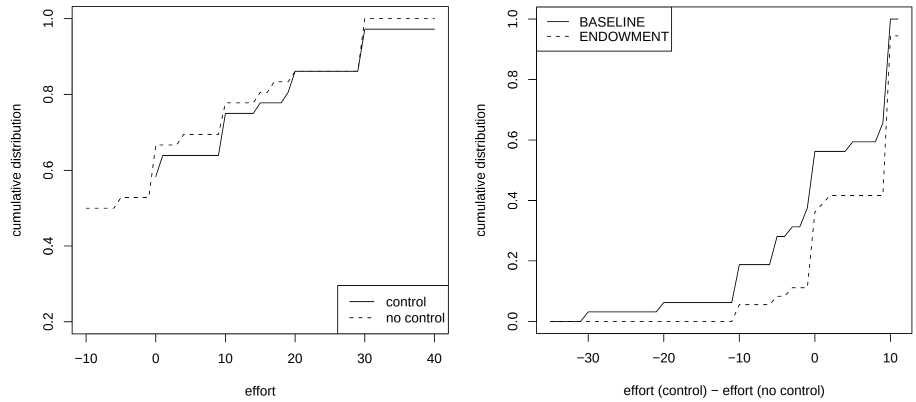

Figure 3: Cumulative distribution of effort (left panel) and of within-subject differences in effort choices for the ENDOWMENT treatment

unforeseen treatment difference is almost significant between BASELINE and ROBOT treatment (p-value of two-sided proportions test: 0.104) and not significant between BASELINE and ENDOWMENT treatment. Although the effect is not significant, we verify whether it affects our results on hidden costs. Our findings are robust: hidden costs remain significantly lower in our treatments relative to the baseline treatment even if we exclude all subjects that only exert minimum effort. The reduction of hidden costs is thus not driven by an increase in subjects who exert minimal effort.

4.2 Principal's Payoff and Behavior

The focus of our paper is agents' behavior, in particular, their response to legitimate and non-legitimate control. In this section, we report some additional interesting findings about the principal. Probably the most interesting concerns the overall effect of control.

Result 4 *In none of our treatments, principals are worse off when they control than when they do not control.*

Figure 4 and Table 1 indicate the average payoffs to principals from human subjects in all three treatments; visual inspection already suggests that the principal cannot gain from leaving the choice free. Also, we already know from our previous analysis that control entails no hidden costs in the ROBOT

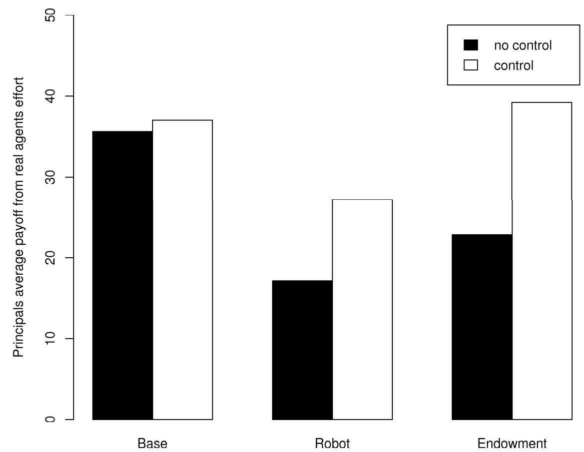

Figure 4: Principals' average payoffs from human subjects

and ENDOWMENT treatment, so it cannot harm. In the BASELINE treatment, the direct effect of control amounts to 8.31 points. Subtracting hidden costs yields an overall effect of $8.31 - 6.88 = 1.43$ and we cannot reject the hypothesis that this effect is zero (The two-sided Wilcoxon signed rank test has a p-value of 0.38).

	BASELINE		ROBOT		ENDOWMENT	
	control	no control	control	no control	control	no control
mean	37.03	35.63	27.21	17.15	39.22	22.83
median	20	22	20	0	20	5

Table 1: Principals' payoffs from human subjects

Later, we discuss how Result 4 relates to studies that aim to replicate Falk and Kosfeld's finding that 'no control' pays for the principal. In line with our idea of legitimacy, control is more beneficial in the ROBOT and ENDOWMENT treatment in comparison to the BASELINE treatment. Is this reflected in principals' behavior?

Result 5 *In the ROBOT treatment and in the ENDOWMENT treatment, principals control more often than in the BASELINE treatment.*

In the BASELINE treatment, the principal controls in 23 of 33 cases. These numbers contrast with the ROBOT treatment, where she controls in 32 out of 33 cases, and the ENDOWMENT treatment, where control happens in 32 out of 36 cases. Figure 5 depicts the shares of principals that control for the three treatments. The share in BASELINE is significantly lower than in ROBOT (p-value for Pearson's χ^2-test: < 0.01) and in ENDOWMENT (p-value for Pearson's χ^2-test: 0.046). This is consistent with the idea that principals are aware of the fact that controlling leads to lower costs when control is legitimate. In the ROBOT treatment, principals' behavior may of course be driven by the presence of the computerized agent. If this would be the only motive, however, then the difference between the shares of principals who control should be significantly larger in the ROBOT than in the ENDOWMENT treatment. This, however, is not the case.

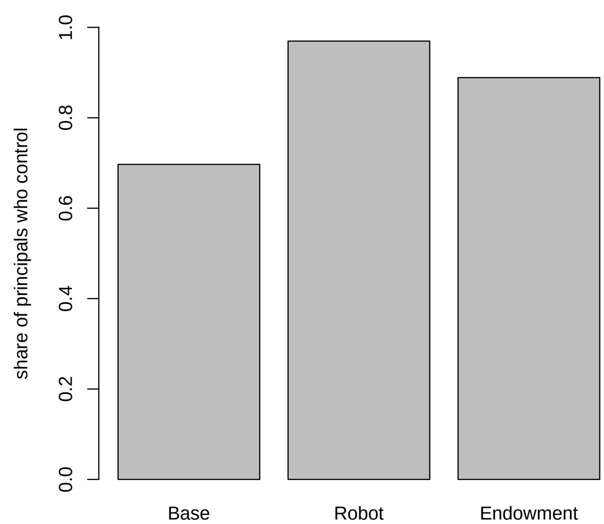

Figure 5: Control by principals across treatments

When interpreting this result, it is important to note that control is attractive in all three treatments: given the actual behavior of agents, control is a best response. So a significantly larger proportion of principals control in the new treatments *although* control already maximizes surplus in the BASELINE treatment.

5 Discussion and Related Literature

Here, we relate our findings to the literature. In particular, we discuss possible explanations for our treatment effects, the link between our results and framing, and to what extend we can replicate the original findings by Falk and Kosfeld (2006).

5.1 Theoretical Explanations

Falk and Kosfeld (2006) do not present an explicit theoretical model to explain their findings. In this section, we review some theories that have been proposed to explain why agents may respond with lower effort to an intervention.[11] We first explain that some of these theories do not apply to Falk and Kosfeld's experiment because they rely on features that are not present in this experiment. Then, we describe theories that can explain hidden costs in Falk and Kosfeld's experiment. Finally, we argue that these theories are not able to predict the observed differences between the BASELINE and ENDOWMENT treatment.

In their seminal paper to explain adverse reactions to control, Bénabou and Tirole (2003) assume that the principal is better informed about agents' effort costs, while Herold (2010) supposes asymmetric information about the production technology. Neither explanation applies to Falk and Kosfeld's design because the agent knows his costs and the production technology. Friebel and Schnedler (forthcoming) examine a team setting with complementarities in efforts. In their paper, control signals that some team-member is uncommitted and committed agents reduce their effort because they are afraid that it will not be matched. In contrast, agents' effort choices do not affect each other in any of our treatments. In Seabright (2009), effort itself is a signal, which indicates civic values and can no longer do so once incentives are in place. Unlike in his model, effort choices remain anonymous and hence cannot signal civic values. In Schnedler (forthcoming), interventions reduce effort because there is a stochastic link between effort and payment, whereas this link is deterministic in Falk and Kosfeld's and our experimental design.

While all these theories cannot be applied here, there are three prominent models that can explain hidden costs of control and apply to the experimental design proposed by Falk and Kosfeld (2006): the theory of guilt (Battigalli and Dufwenberg, 2007), the theory of conformism (Sliwka, 2007), and the theory of self-esteem (Ellingsen and Johannesson, 2008). In all three theories, the adverse response by the agent occurs because control is a signal about private information held by the principal. The theories differ in what information is revealed by the signal and why the agent cares about this information. Control may signal low expectations by the principal and the agent may care about these expectations because he is guilt-averse in the sense of Battigalli and Dufwenberg (2007). Alternatively, control may be a signal about the preferences of other agents' to an agent who is a conformists (Sliwka, 2007). Finally, control may signal the principal's selfishness to an agent who wants to impress the principal but only if the principal is not selfish (Ellingsen and Johannesson, 2008).

Since the game in our BASELINE treatment is identical to that in Falk and Kosfeld (2006), all three theories can explain why control leads to lower effort in this treatment. In the ROBOT treatment, the principal has a new reason to control: the selfish robot. As a consequence, control is less informative about the principal's expectations, her knowledge of other agents, or her selfishness. Lower hidden costs in the ROBOT treatment are thus consistent with these theories. The findings in our ENDOWMENT treatment, however, are not predicted by any of these theories. Recall that the game in the ENDOWMENT treatment is strategically equivalent to that in the BASELINE treatment. Whatever is signaled by control in the BASELINE treatment (e.g. low expectations, low trust, a large share of selfish subjects, or principal's preferences, etc.), is also signaled in the ENDOWMENT treatment. Accordingly, one would expect the same negative response to control in the BASELINE and ENDOWMENT treatment. However, we find that hidden costs are significantly lower in the ENDOWMENT treatment.

In order for control to be observed in equilibrium in the ENDOWMENT treatment but not in the BASELINE treatment, the game, which underpins both treatments, needs to have two equilibria: one in which control is not punished and another one in which it is punished. The observed behavior can then be explained by assuming that equilibrium selection is based on subjects' norms. More specifically, suppose that principal and agent agree on the conditions under which control is legitimate as well as on the fact that control is punished whenever it is not legitimate. Then, the differences in the ENDOWMENT and BASELINE treatment can be explained using our definition of legitimacy.

Applying the theory of guilt-aversion to our experiment, there are indeed multiple equilibria, some of which entail hidden costs, while others do not. These equilibria differ in their initial beliefs. If these beliefs are shaped by a norm, this can explain the treatment effects—for more details see Appendix B.

5.2 Replicating Falk and Kosfeld (2006)

There are two studies that are directly concerned with replicating Falk and Kosfeld's experiment (2006), in particular, their finding that trusting leads to a higher payoff for principals than control. Hagemann (2007) points out that this finding is sensitive to the wording of instructions: with a neutral frame, the principal is better off controlling. Ploner et al. (2010) conduct a replication study with over 228 subjects and three different subject pools. They come to the conclusion that Falk and Kosfeld's finding is driven by the subject pool.

Unlike Hagemann (2007) and Ploner et al. (2010), our experiment has not been designed as a replication study. In particular, our BASELINE treatment differs in various ways from the C10 treatment by Falk and Kosfeld (2006), which are possibly relevant. First, the experiment was conducted with a different subject pool. Second, we wrote the instructions to have the smallest possible difference across our treatments. As a consequence, they differ from those used by Falk and Kosfeld (2006). According to Hagemann (2007) and Ploner et al. (2010), both differences may be relevant. With these reservations in mind, the effect of control on the principal's payoff in our BASELINE treatment can be compared with that in Falk and Kosfeld's C10 treatment. Like Hagemann (2007) and Ploner et al. (2010) and unlike Falk and Kosfeld (2006), we cannot find that control harms the principal (see our Result 4).

Probably, the most important aspect of the experiment by Falk and Kosfeld (2006) is that they document the existence of hidden costs. They do so indirectly by observing that the overall effect of control is negative. The same overall effect is the central object of study in Hagemann (2007) and Ploner et al. (2010). Our paper, however, is not concerned with this overall effect but directly with the hidden costs of control and the question under which conditions they exist. We develop the necessary methodology to separate the direct (beneficial) effect from control from indirect psychological effects, which allows us to estimate hidden costs. Using this methodology, our findings support Falk and Kosfeld's observation that hidden costs of control exist (Result 1), although they do not outweigh the direct benefits in our BASELINE treatment.

5.3 Framing

Our experiment adds to the well established literature on framing in games. As early as 1960, Schelling has argued that appropriate framing of actions may help players coordinate their play on certain focal equilibria. The impact of framing on behavior was documented by Andreoni (1995) in the public goods and by Cooper et al. (1999) in the principal-agent setting. In their recent study, Dufwenberg et al. (2008) show that the framing of a decision problem may also affect behavior through higher order beliefs. Perhaps one of the nicest illustrations of framing effects is due to Liberman et al. (2004). In their experiment they found much more cooperation in the prisoner's dilemma game when it was presented as a cooperative 'community game' than when it was presented as a competitive 'Wall Street game.'

Closely related to our experiment is the study of behavior in the dictator game by List (2007). In all treatments, both the dictator and the recipient were given initial endowments, i.e., amounts equal to 10 and 5 respectively. In one of the treatments, the dictator could keep $10 > x > 5$ to himself and give the rest to the recipient. In two additional treatments, 'taking' was allowed. The dictators could allocate $11 > x > 4$ to themselves in one of the additional treatments and $15 > x > 0$ in the other. List finds that there is a significant spike (30-45%) at the recipient's initial endowment, $x = 10$. This, illustrates that the initial endowments might be inducing a feeling of entitlement or ownership rights that are respected by a large number of dictators. Our paper takes this finding one step further and claims that if the entitlements are successfully induced, then it becomes legitimate for the recipient to protect her initial endowment against a selfish dictator. Our data support this conjecture.

6 Conclusion

Control can have two types of consequences: the desired direct effect of preventing a certain behavior and the indirect effect of provoking counter-productive reactions. There are various theories that can explain such counter-productive reactions. We complement these theories with the conjecture that control is legitimate in contexts or situations in which it averts anti-social behavior. This leads to the predictions that counter-productive reactions are weaker when control averts anti-social behavior. We test these predictions in an experiment and find that they are borne out by the data. Our results highlight the fact that the behavioral consequences of control are specific to the context in which control is used. Falk and Kosfeld (2006) suggest that control may be an attractive option for the principal when it is particularly effective—in other words, when principals can enforce a high minimum level of effort. In this case, the direct, disciplinary effect of control dominates any other indirect, behavioral response to control. By looking at the two effects separately, our study uncovers another reason for principals to control: the detrimental behavioral response to control (the hidden costs) may vanish. Our experiment demonstrates that hidden costs are negligible if control is legitimate.

A Appendix: Instructions for the ROBOT treatment

General Instructions for Participants

You are now participating in an economic experiment. Please read the following instructions carefully. The instructions will provide you with all the information you require for participation in the experiment. Please ask for assistance if there is something that you do not understand. Your question will be answered at your workplace. There is a strict prohibition of communication during the experiment.

You will receive a show-up fee of 4 Euro at the beginning of the experiment. Over the course of the experiment you can earn points. This income will be converted into Euro at the end of the experiment. Please note that:

1 point = 10 Cent

The converted income and the show-up fee will be paid out in cash at the end of the experiment.

The Experiment

In this experiment, each participant A is associated with a participant B in a group of two. No participant knows with whom he is associated; all decisions are made anonymously.

At the beginning of the experiment, Participant A receives 120 points as an endowment and participant B 0 points.

Decision of Participant A

Participant A can decide how many points he wants to give to participant B. These points are deducted from A's endowment, doubled and added to B's endowment. Each point that A gives to B thus reduces the income of A by one point and increases that of B by two points.

After participant A has chosen, the incomes hence amount to:

Income of Participant A:	120 – chosen points
Income of Participant B:	0 + 2 * chosen points

The following examples illlustrate the computation of the income:

Example 1: A gives 0 points (chosen points: 0). Then, A's income is 120 and B's 0 points.
Example 2: A gives 20 points (chosen points: 20). The income is 100 for A and 40 for B.
Example 3: A gives 80 points (chosen points:80). The income is 40 for A and 160 for B.

Decision of Participant B

Participant B may require to be given at least 10 points or he can leave the choice of points completely free. There are hence two cases:

Case 1: Participant B requires to be given at least 10 points. Then participant A can chose any (integer) amount **between 10 and 120** points. In this case the income of B is at least 20 points.

Case 2: Participant B leaves the choice of points free. Then, participant A can chose any (integer) amount **between 0 and 120** points. In this case the income of B can amount to 0 points.

Decision of Computer Program

A random draw decides whether the decision of participant A has any consequences at all. With a probability of 50%, participant A is replaced by a computer program. In this case, participant A's income is equal to his endowment and the income of participant B results from the number of points chosen by the computer program. The program always chooses the *smallest possible* amount of points, i.e. 10 points if participant B requires at least 10 points (case 1) and 0 points if B leaves the choice free (case 2).

Suppose that the program replaces participant A. If B leaves the choice free (case 2), then the program will give B the lowest possible amount, i.e. 0 points. The income of B is then 0 points. If B requires to be given at least 10 points (case 1), the program gives 10 points. In this case, participant B gets at least 20 points.

When participant B decides whether to leave the choice free or not, he does not know whether he faces participant A or the program.

The Stages of the Experiment

1. Participant B decides whether to require at least 10 points or leave the choice free.

2. Participant A determines a number of points. If participant B requires at least 10 points (case 1), participant A can select a number of points between 10 and 120. If participant B leaves the choice free (case 2), participant A can select a number of points between 0 and 120.

3. A random move decides whether participant A or a program determines the income. There is a probability of 50% that A's number of points determines the size of the participants' income. There is a probability of 50% that A keeps his initial endowment and the program determine's B's income. If participant B requires at least 10 points (case 1), the program gives 10 points and B's income is 20 points. If participant B leaves the choice free (case 2), the program gives 0 points and B's income amounts to 0 points.

4. Participant A and participant B each learn the decision of the other, the result of the draw, and the size of the resulting incomes.

Then, the experiment is over.

B Appendix: Guilt-aversion and Norms

In this appendix, we use the framework of pro-social behavior based on guilt-aversion (Battigalli and Dufwenberg, 2007) to illustrate the effect of a legitimacy norm. We show that with a guilt-averse agent, the general game between principal and agent, which underpins all three treatments, has two equilibria: one in which control entails hidden costs and one in which it does not. Moreover, which of these two equilibria occurs depends on initial beliefs. If there is a norm in line with our definition of legitimacy, these beliefs are pinned down and lead to equilibrium behavior consistent with our hypotheses.

Suppose the principal expects the agent to supply a certain level of effort when she controls and a possibly different level when she does not control. In other words, the principal has some first-order belief about effort which depends on control. Further, suppose that the agent has a second-order belief:[12] he has an estimate of the principal's expectations given that the principal controls, μ^{c}, and given that the principal does not control, μ^{nc}.

The utility of a guilt-averse agent depends not only on monetary payoffs but also on his belief about what the principal expects him to do. Formally, we denote the guilt felt by the agent when he thinks the principal expects μ and he supplies x with $G(\mu, x) := \max\{0, 2(\mu - x)\}$. Let $\theta \geq 1/2$ describe the degree of guilt-aversion of the agent,[13] then the overall utility for the agent from monetary payoff and guilt is:

$$u_{\mathrm{A}}\left(y, (x^{\mathrm{nc}}, x^{\mathrm{c}}) \mid (\mu^{\mathrm{nc}}, \mu^{\mathrm{c}})\right) \;=\; \pi^{1}_{\mathrm{A}}\left(y, x^{i}\right) - \theta G(\mu^{i}, x^{i}), \tag{3}$$

where $i = \mathrm{c}$ if the principal controls and $i = \mathrm{nc}$ else. For simplicity, we assume that the principal only cares about her monetary payoffs.

Let us examine the following generalized game, which embeds the games played in all three treatments. Denote by x_{L} the lowest possible effort if the principal does not control. Suppose the principal is only matched with the agent with probability p, receives y with probability $1-p$, and does not know whether she is matched with the agent while deciding whether to control ($y = \underline{x}$) or not ($y = x_{\mathrm{L}}$). This game has two types of equilibrium.

Proposition 1 *There is a trust equilibrium in which the principal does not control ($y = x_L$) and the agent's effort matches the agent's beliefs ($x^{nc} = \mu^{nc}, x^c = \mu^c$) if and only if the agent's beliefs satisfy:*

$$\mu^{nc} - \mu^{c} \geq \frac{1-p}{p}(\underline{x} - x_L). \tag{4}$$

There is a control equilibrium in which the principal controls ($y = \underline{x}$) and the agent's effort matches the agent's beliefs ($x^{nc} = \mu^{nc}, x^c = \mu^c$) if and only if the agent's beliefs satisfy:

$$\mu^{nc} - \mu^c \leq \frac{1-p}{p}(\underline{x} - x_L). \tag{5}$$

Proof. In any equilibrium, the agent and the principal have correct beliefs. Since θ is sufficiently high ($\theta > 1/2$), the agent always meets the expectations of the principal and plays ($x^{nc} = \mu^{nc}$, $x^c = \mu^c$). Let us now turn to the principal. The payoff of the principal if she controls is: $2(x^{\mathrm{c}}p + \underline{x}(1-p))$. Under no control, it amounts to: $2(x^{\mathrm{nc}}p + x_{\mathrm{L}}(1-p))$. Accordingly, a deviation to control in a trust equilibrium is not profitable if $\mu^{\mathrm{nc}} - \mu^{\mathrm{c}} \geq \frac{1-p}{p}(\underline{x} - x_{\mathrm{L}})$. Analogously, a deviation to 'no control' in a control equilibrium is not profitable if $\mu^{\mathrm{nc}} - \mu^{\mathrm{c}} \leq \frac{1-p}{p}(\underline{x} - x_{\mathrm{L}})$. ■

In a trust equilibrium, the agent matches expectations, his effort under no control is larger than or equal to his effort under control, and control may entail hidden costs. The control equilibrium is consistent with low or no hidden costs of control because of (5). Which equilibrium arises depends on the initial beliefs of principal and agent; these beliefs may be determined by what is legitimate.

Next we argue that only the trust equilibrium can result in the BASELINE treatment, only the control equilibrium can result in the ROBOT treatment, while both are possible in the ENDOWMENT treatment. Legitimate actions are associated with high expectations and non-legitimate actions with low expectations. Recall that according to Definition 3 an action is legitimate, if it prevents theft or selfish behavior. In the BASELINE treatment, control by the principal reduces the initial endowment of the agent and hence it is legitimate *only* if it prevents theft or selfishness on the part of the agent. Since the principal has no initial endowment of his own, the agent cannot steal from her. However, selfish behavior (giving nothing: $x = 0$) is feasible under no control. Observe that controlling by the principal does not only prevent $x = 0$, but also $x \in \{1, ..., 10\}$, which are all legitimate acts of giving. Thus, in the BASELINE treatment whether controlling is legitimate depends on the common prior of selfish behavior ($x = 0$) in the absence of control. If this prior is low, controlling is not legitimate and hence the principal must have low expectations following control: $\mu^{\mathrm{c}} < \mu^{\mathrm{nc}}$. Given that $p = 1$, the inequality in Condition 4 is then strict and the BASELINE treatment results in a trust equilibrium with hidden costs of control.

In terms of legitimacy, the ROBOT treatment is identical to the BASELINE with the exception that there is now a common belief of at least 50% chance for the selfish action, $x = 0$, under no control. This means that both control and no control are legitimate. Accordingly, initial beliefs in the ROBOT treatment are identical for control and no control: $\mu^{\mathrm{c}} = \mu^{\mathrm{nc}}$. Given that $p = \frac{1}{2}$, the inequality in Condition 4 is strict. The ROBOT treatment thus results in a control equilibrium with no hidden costs.

In the ENDOWMENT treatment, the principal has 20 points of initial endowment of his own. Hence, under no control, any action by the agent that is lower than 10 ($x < 10$) reduces the principal's endowment and is therefore and act of theft. Definition 3 fits this situation perfectly, in which control prevents all acts of theft and *only* theft. This legitimizes control; control and no control are legitimate and therefore associated with high (and equal) expectations: $\mu^{\mathrm{c}} = \mu^{\mathrm{nc}}$. Both equilibria are consistent with these beliefs but whichever is played, there will be no hidden costs of control.

Notes

[1]See The New York Times' article "German Court Permits Limited Cybermonitoring" from February 2008 for a recent example involving the Constitutional Court in Germany.

[2]At a number of archaeological sites, e.g., the catacombs in Paris, check bags are routinely checked for stolen historical artifacts. This clearly does not discourage tourists from visiting these sites.

[3]As an example take the following quote from the Verizon Wireless Code of Business Conduct (2001):

> "Sometimes it is necessary to monitor employee personal communications or computer usage or to search employee workspaces for the protection of employees, company assets and other legitimate business reasons."

[4]The tradition dates back to Plato's Republic (390 BC) and Aristotle's Politics (335-323 BC) and includes political scientists, sociologists, social psychologists, historians and political economists throughout the centuries; major works have been dedicated to this topic — see e.g. Weber (1947), Habermas (1996) and Freeman (2000).

[5]See, e.g., Güth et al. (1982) and Gintis (2000) for reciprocity arguments, Hoffman et al. (1994) and Gächter and Riedl (2005) for entitlement formation, Bosman and Winden (2002) and Falk and Fischbacher (2002) for laboratory studies of theft and crime and Carlsmith et al. (2002) for a discussion of justified crime punishment.

[6]In Cox et al. (2008), the agent's preferences are sensitive to more or less generous choice sets; in Andreoni and Bernheim (2009) and Fehr and Schmidt (2000) the norm is the equal split; central to works of Rabin (1993) and Dufwenberg and Kirchsteiger (2004) is an appropriately defined kindness norm which is based on averaging the minimum and

maximum payoff following a choice; finally, Charness and Rabin (2002) and López-Pérez (2008) each use a norm that combines a human concern for both the equality and social efficiency.

[7]Falk and Kosfeld (2006) also use the strategy method. They validate this step by running an additional between-subject design and find that the method does not affect results. Notice that even if hidden costs were enhanced by possible demand effects induced by this method, this would have been the case in all our treatments. The fact that we do not find hidden costs in the ROBOT and ENDOWMENT treatment would then be in spite of this counteracting effect and would only reinforce our results.

[8]Falk and Kosfeld use the same procedure to test whether control has a 'behavioral impact.'

[9]Recall that a smaller indirect effect means larger hidden costs.

[10]The positive sign of the indirect effect is due to a single subject who chooses the maximal effort when being controlled and no effort otherwise. Excluding this agent from the data does not change the results: although the indirect effect becomes negative, hidden costs remain significantly lower in the ENDOWMENT treatment, and the Wilcoxon signed rank test (which is robust to outliers) still has a p-value below 0.01.

[11]For a survey of respective experiments see Bowles (2008).

[12]Higher order beliefs could be defined accordingly but are not needed for our analysis.

[13]If $\theta < 1/2$ marginal costs of exerting effort always exceed the marginal guilt and guilt-averse agents act like selfish agents.

References

Andreoni, James. "Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments." *The Quarterly Journal of Economics* 110, 1: (1995) 1–21.

Andreoni, James, and B. Douglas Bernheim. "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects." *Econometrica* 77, 5: (2009) 1607–1636.

Battigalli, Pierpaolo, and Martin Dufwenberg. "Guilt in Games." *American Economic Review* 97, 2: (2007) 170–176. Papers and Proceedings.

Bénabou, Roland, and Jean Tirole. "Intrinsic and Extrinsic Motivation." *Review of Economic Studies* 70: (2003) 489–520.

Bosman, Ronald, and Frans van Winden. "Emotional Hazard in a Power-to-Take Experiment." *The Economic Journal* 112, 476: (2002) 147–169.

Bowles, Samuel. “Policies Designed for Self-Interested Citizens May Undermine ‘the Moral Sentiments’: Evidence from Economic Experiments.” *Science* 320, 5883: (2008) 1605 — 1609.

Carlsmith, Kevin M., John M. Darley, and Paul H. Robinson. “Why Do We Punish? Deterrence and Just Deserts as Motives for Punishments.” *Journal of Personality and Social Psychology* 83: (2002) 284–299.

Charness, Gary, and David I. Levine. “When are Layoffs Acceptable? Evidence from a Quasi-Experiment.” *Industrial Labor Relations Review* 53, 3: (2000) 381–400.

———. “When is Employee Retailitation Acceptable at Work? Evidence from Quasi-Experiments.” *Industrial Relations* 49, 4: (2004) 499–523.

Charness, Gary, and Matthew Rabin. “Understanding Social Preferences with Simple Tests.” *The Quarterly Journal of Economics* 117, 3: (2002) 817–869.

Cooper, David J., John H. Kagel, Wei Lo, and Qing Liang Gu. “Gaming against Managers in Incentive Systems: Experimental Results with Chinese Students and Chinese Managers.” *The American Economic Review* 89, 4: (1999) 781–804.

Cox, James C., Daniel Friedman, and Vjollca Sadiraj. “Revealed Altruism.” *Econometrica* 76, 1: (2008) 31–69.

Dufwenberg, Martin, Simon Gächter, and Heike Hennig-Schmidt. “The Framing of Games and the Psychology of Play.” Technical Report 15/200, Bonn Graduate School of Economics, 2008.

Dufwenberg, Martin, and Georg Kirchsteiger. “A theory of sequential reciprocity.” *Games and Economic Behavior* 47, 2: (2004) 268–298.

Ellingsen, Tore, and Magnus Johannesson. “Pride and Prejudice: The Human Side of Incentive Theory.” *American Economic Review* 98, 3: (2008) 990–1008.

Falk, Armin, and Urs Fischbacher. “”Cimre ” in the Lab-Detecting Social Interaction.” *European Economic Review* 46, 4-5: (2002) 859–869.

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

Fehr, Ernst, and Klaus M. Schmidt. "Fairness, Incentives, and Contractual Choices." *European Economic Review* 44: (2000) 1057–1068.

Freeman, Samuel. "Deliberative Democracy: A Sympathetic Comment." *Philosophy and Public Affairs* 29: (2000) 371–418.

Friebel, Guido, and Wendelin Schnedler. "Team Governance: Empowerment or Hierarchical Control." *Journal of Economic Behavior and Organization*, forthcoming.

Gächter, Simon, and Arno Riedl. "Moral Property Rights in Bargaining with Infeasible Claims." *Management Science* 51, 2: (2005) 249—263.

Gintis, Herbert. "Strong Reciprocity and Human Sociality." *Journal of Theoretical Biology* 206, 2: (2000) 169 – 179.

Güth, Werner, Rolf Schmittberger, and Bernd Schwarze. "An experimental analysis of ultimatum bargaining." *Journal of Economic Behavior & Organization* 3, 4: (1982) 367 – 388.

Habermas, Jürgen. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy.* Massachusetts: MIT Press, 1996.

Hagemann, Petra. "What's in a Frame? Comment on: The Hidden Costs of Control.", 2007. University of Cologne.

Herold, Florian. "Contractual Incompleteness as a Signal of Trust." *Games and Economic Behavior* 68: (2010) 180–191.

Höffe, Otfried. *Political Justice: Foundations for a Critical Philosophy of Law and the State.* Cambridge, UK: Polity Press, 1995. Translated by Jeffrey C. Cohen.

Hoffman, Elizabeth, Kevin McCabe, Keith Shachat, and Vernon Smith. "Preferences, Property Rights, and Anonymity in Bargaining Games." *Games and Economic Behaviour* 7: (1994) 346–380.

Hoffman, Martin L. "Moral Internalization: Current Theory and Research." In *Advances in Experimental Social Psychology*, edited by Loeonard Berkowitz, New York: Academic Press, 1977, 86—135.

Kahneman, Daniel, Jack L Knetsch, and Richard Thaler. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review* 76, 4: (1986) 728–41.

Liberman, Varda, Steven Samuels, and Lee Ross. "The Name of the Game: Predictive Power of Reputations versus Situational Labels in Determining Prisoner's Dilemma Game Moves." *Personality and Social Psychology Bulletin* 30: (2004) 1175–1185.

List, John A. "On the Interpretation of Giving in Dictator Games." *Journal of Political Economy* 115, 3: (2007) 482–493.

López-Pérez, Raúl. "Aversion to norm-breaking: A model." *Games and Economic Behavior* 64, 1: (2008) 237–267.

Okun, Arthur M. *Prices and Quantities: A Macroeconomic Analysis*. Washington, D.C.: Brookings Insititute Press, 1981.

Ploner, Matteo, Katrin Schmelz, and Anthony Ziegelmeyer. "Hidden Costs of Control: Three Repetitions and an Extension." Technical Report 7, Max-Planck-Institute of Economics, Jena, 2010.

Rabin, Matthew. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83, 5: (1993) 1281–1302.

Schelling, Thomas C. *The Strategy of Conflict*. Cambridge, MA.: Harvard University Press, 1960.

Schnedler, Wendelin. "You Don't Always Get What You Pay For." *German Economic Review*, forthcoming.

Seabright, Paul B. "Continuous Preferences and Discontinuous Choices: How Altruists Respond to Incentives." *The B.E. Journal of Theoretical Economics* 9, 1, Contributions: (2009) Article 14. Contributions.

Sliwka, Dirk. "On the Hidden Costs of Incentive Schemes." *American Economic Review* 97, 3: (2007) 999–1012.

Suchman, Mark C. "Managing Legitimacy: Strategic and Institutional Approaches." *The Academy of Management Review* 20, 3: (1996) 571–610.

Tyler, Tom R. "Psychologicyl Perspectives on Legitimacy and Legitimation." *Annual Review of Psychology* 57: (2006) 375–400.

———. "Psychology and Institutional Design." *Review of Law & Economics* 4, 3: (2008) Article 6.

Weber, Max. *The Theory of Social and Economic Organization.* Glencoe: The Free Press and the Falcon's Bring Press, 1947. Translated by A.M. Henderson and Talcott Parsons.

Zelditch, Jr., Morris. "Theories of Legitimacy." In *The Psychology of Legitimacy*, edited by John T. Jost, and Brenda Major, Cambridge: Cambridge University Press, 2001, 33–53.