



The hidden benefits of abstaining from control

This is the peer reviewed version of the following article:

Original:

Gabriel, B., Halliday, S., Landini, F. (2018). The hidden benefits of abstaining from control. JOURNAL OF ECONOMIC BEHAVIOR & ORGANIZATION, 147, 1-12 [10.1016/j.jebo.2017.12.018].

Availability:

This version is available <http://hdl.handle.net/11365/1277354> since 2024-11-04T17:07:20Z

Published:

DOI:10.1016/j.jebo.2017.12.018

Terms of use:

Open Access

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license.

For all terms of use and more information see the publisher's website.


(Article begins on next page)

Accepted Manuscript

Title: The hidden benefits of abstaining from control

Authors: Gabriel Burdin, Simon Halliday, Fabio Landini

PII: S0167-2681(17)30368-2
DOI: <https://doi.org/10.1016/j.jebo.2017.12.018>
Reference: JEBO 4232



To appear in: *Journal of Economic Behavior & Organization*

Received date: 18-4-2016
Revised date: 12-12-2017
Accepted date: 21-12-2017

Please cite this article as: Burdin G, Halliday S, Landini F, The hidden benefits of abstaining from control, *Journal of Economic Behavior and Organization* (2010), <https://doi.org/10.1016/j.jebo.2017.12.018>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

[Accepted for publication: Journal of Economic Behavior and Organization]

The hidden benefits of abstaining from control

Gabriel Burdin
Leeds University Business School & IZA

Simon Halliday
Department of Economics, Smith College

Fabio Landini
University of Parma & ICRIOS - Bocconi University

Highlights

- Laboratory experiment examining the role played by negative reciprocity, positive reciprocity and autonomy in explaining agents' reactions to control in principal-agent interactions.
- Two treatments are compared: one in which control is exerted by the principal and the other in which it is exerted by a third party.
- When the principal decides not to control, the agent exerts greater effort relative to the case when the third party decides not to control.
- Agents reward principals who abstain from control for their trust, rather than punish controlling ones for their distrust.

Abstract

This paper studies the role of negative reciprocity, positive reciprocity and preferences for autonomy in explaining agents' reactions to control in experimental principal-agent games. While most of the social psychology literature emphasizes the role of autonomy, recent economic research has provided an alternative explanation based on reciprocity. To understand the behavioral mechanisms underlying such reactions, we conduct an experiment in which we compare two treatments: one in which control is exerted directly by the principal; and the other in which it is exerted by a third party enjoying no residual claimancy rights (third-party control). The results indicate that when either the principal or a third party decides to control the average level of effort that is selected by the agents is similar. What changes remarkably are the agents' reactions to the decision of the other participants not to control. When the principal decides not to control, then the agent exerts greater effort relative to the case when the third party decides not to control. Agents seem to reward principals who abstain from control for their trust, rather than punish controlling ones for their distrust.

JEL Codes: C72, C91, D23, M54

Keywords: control aversion, autonomy, principal-agent game, social preferences, trust, negative reciprocity, positive reciprocity.

1. Introduction

In contemporary societies, significant resources are devoted to control people's actions. For instance, a substantial fraction of the labor force is allocated to supervisory tasks in both developed and developing countries (Acemoglu and Newman, 2002; Jayadev and Bowles, 2006; Fafchamps and Söderbom, 2006). According to figures computed from the European Working Condition Survey (EWCS), more than half (57%) of non-supervisory employees lack procedural autonomy at work in at least one dimension (i.e. the ability to change or choose the order of tasks, the speed or rate of work and the method of work) and 42% perceive that their work rate depends on the direct control of their bosses.¹ Hence, understanding the precise behavioral mechanisms underlying people's reactions to control and their economic consequences are important concerns.

Traditionally, two main streams of literature have focused on people's reactions to control. On the one hand social psychologists have emphasized the role of individual orientations towards autonomy and control. According to Self-Determination Theory (SDT), human beings have a basic psychological need for autonomy (Deci and Ryan, 1985). That is, humans require "a form of freedom in which a party experiences himself to be the locus of causality for his own behavior" (Gagné and Deci, 2005, p. 333). This approach sees people's wellbeing as inseparable from their experience of personal and motivational autonomy (Chirkov et al, 2011) and considers the quest for autonomy as one of the main drivers of individual reactions to control.²

On the other hand, behavioral economists have focused primarily on intention-based social preferences, in particular reciprocity (Falk and Fischbacher, 2006; Dufwenberg and Kirchsteiger, 2004; Von Siemens, 2013). Nowadays, there is ample evidence that many agents behave in a reciprocal manner even when acting on reciprocal preferences is costly and yields no future rewards (see, for instance, Fehr and Gächter, 2000). On

¹ Own calculations from EWCS wave 2010.

² Several experiments conducted by psychologists in highly differentiated contexts have shown that environments supporting autonomy (control) to significantly increase (decrease) intrinsic motivations and prosocial behavior, and therefore that autonomy and control can severely affect task performance (see Gagné, 2003; Greene-Demers et al., 1997; Pelletier et al., 1998; Fabes et al., 1989; Kunda and Schwartz, 1983; Upton, 1974; Batson et al., 1978; Sobus, 1995). Recent experimental evidence also suggests that individuals intrinsically value decisional autonomy over their own and others' outcomes (Bartling et al, 2014; Owens et al, 2014). Moreover, greater procedural autonomy and lower monitoring intensity appear to correlate positively with greater job satisfaction (Bartling et al, 2013).

this basis Falk and Kosfeld (2006, henceforth, F&K) have provided a reciprocity-based explanation for individual reactions to control. In a principal-agent game they explore the phenomena of hidden costs of control and the idea that ‘control aversion’ may be one of the reasons why incentives sometimes degrade performance. They found a sizeable fraction of agents react negatively to control and that control is not profitable, i.e. principals earn more if they leave agents to decide freely than if they control. F&K explain their result in terms of negative reciprocity on the side of the agents who punish the controlling principals for their distrust. An alternative interpretation can be based on positive reciprocity according to which the agents reward principals who abstain from control for their trust.³

Although both the autonomy and the reciprocity explanations are plausible, there is still no clear evidence on their relative importance. F&K explored the agents’ emotional perception of control in their experiment and the most frequent answers among those agents who react negatively to control were distrust and lack of autonomy. However, the experimental design does not allow the authors to separate the explanatory role of these two motives. Moreover, subsequent replications of F&K’s study (e.g. Ziegelmeyer et al, 2012) confirmed the existence of a hidden costs of control, namely the fact that a large fraction of agents reacts negatively to control by lowering their level of effort, but they failed to replicate the finding that such costs outweigh the benefit of control.

In this paper we extend F&K’s experimental design to separate the role of autonomy and reciprocity in explaining how individuals react to control in a principal-agent relationship. We vary their experiment to permit it to include three parties: the principal, who benefits from the effort of the agent, the agent, and a third party who is given a show-up fee and chooses whether or not to exert control over the agent, but does not directly benefit from the agent’s actions (i.e. he does not have any claim over the residual). On this basis, we obtain two main results. First, we reject autonomy as a

³ Ellingsen and Johannesson (2008) and Sliwka (2007) provide an alternative explanation for F&K’s results in terms of image concern and conformism, respectively. While Ellingsen and Johannesson (2008) present an alternative way of modeling a reciprocity-based reaction to control, Sliwka (2007) exploits the tendency to conform to a generalized social norm to explain control aversion. Although in this paper we do not explicitly compare alternative type of social preferences, our results suggest that reciprocity-based explanations are more reliable than conformism-based ones and we thus stick to reciprocity as the main reference in our analysis.

motive to explain agent's behavior in F&K's stylized principal-agent interaction. When a third party instead of the principal exerts control, the fraction of agents that reduce performance dramatically decreases and control results in a greater expected profit for the principal. Moreover, regression analysis shows no correlation between the probability of performing worse when exposed to control and a psychological measure of autonomy orientation (General Causality Orientation Scale; GCOS). In contrast, we find a significantly negative correlation between this probability and individuals' control orientation as measured by GCOS. Overall, reactions to control appear not to be driven (at least in this very specific and highly stylized experimental setting) by agents' preference for autonomy.

The second result that we obtain is that, in line with Ziegelmeyer et al (2012), we fail to replicate F&K's finding on the negative effect of control for the principal's profit. Although we largely confirm that when the principal is exerting control there exist hidden costs, we find that the latter are not substantial enough to significantly undermine the benefits of control. In our setting, however, the comparison with the case in which control is exerted by a third party makes it possible to enrich the interpretation of such a result. When either the principal or a third party decides to control the average level of effort that is selected by the agents is similar. What changes remarkably is the agent's reaction to the decision of not controlling. When a third party makes such a decision, it results in a significant lower level of effort. On the contrary, when the principal makes the decision not to control, it results in a level of effort that is equal to the case in which control is exerted. In this sense the behavioral response that F&K interpret as a negative reaction to the principal's choice of controlling (i.e. negative reciprocity) seems more sensibly interpreted as a positive reaction to the principal's choice not to control when he had the option to do so (i.e. positive reciprocity with respect to the choice of abstaining from control). In other words, in our experiments trust seems to have hidden benefits that make it a suitable substitute for a principal's control. As a result, principals who do not exert control perform as well as principals who do control. The same is not true when the option to control is in the hands of a neutral third party, as in that case control is by far the most profitable choice.

The paper contributes to the growing experimental economics literature on authority and control in organizations (Falk and Kosfeld, 2006; Ellingsen and Johannesson, 2008;

Sliwka, 2007; Zieglmeyer et al, 2012; Fehr et al., 2013; Charness et al., 2012; Schnedler and Vadovic, 2011; Riener and Wiederhold, 2016). The study also adds to the literature on crowding out (in) effects of incentives on intrinsic motives (see Frey and Jegen, 2000; Bowles and Polanía, 2012). Finally, the paper contributes to the research agenda in organizational economics trying to improve the mapping of individual preferences and assessing the consequences of the mismatch between preferences and organization design (Ben Nér, 2013). By deepening the study of the behavioral motives underlying reactions to control, the results presented in this paper may have implications for key aspects of organizational design, such as the optimal level of employees' discretion and monitoring practices. Specifically, our results may provide a rationale for why principals should retain control without exerting it rather than delegating control to a neutral third party. In the presence of reciprocal types, the principal's decision not to control may trigger workers' good behavioral dispositions, which translate into a level of effort that compensates the costs of running into some degree of shirking. On the contrary, when the option to control is delegated to a neutral third party (e.g. consultancies, governments) such trust-based behavioral dispositions are difficult to induce and the performance of the organization may worsen.

The rest of the paper is structured as follows. In section 2, we present the experimental design, including the original F&K design and our third-party treatment. Section 3 discusses the behavioral predictions. Section 4 describes practical procedures related to the experiment. Section 5 presents the main results. Finally, in section 6 we conclude and discuss potential extensions.

2. Experimental Design

Principal-agent game

In order to test the extent to which reactions to control depend on both reciprocity and autonomy we rely on a simple laboratory experiment. The experiment is based on the two-stage principal-agent game used in F&K and replicated in Zieglmeyer et al (2012). The agent chooses a productive activity x , which is costly to him but beneficial for the principal. The monetary cost for the agent is $c(x) = x$, while the benefit for the principal is $2x$; i.e., the marginal cost of providing the productive activity is always smaller than

the marginal benefit. The agent has an initial endowment of 120 experimental currency units (ECUs), while the endowment of the principal is 0.

Before the agent decides on x , the principal determines the agent's choice set. The principal can either restrict the agent's choice set, in which case the agent can choose any integer value $x \in \{\underline{x}, \underline{x} + 1, \dots, 120\}$, or the principal can leave the choice set unrestricted to $x \in \{0, 1, \dots, 120\}$. Thus the principal can control the agent's decision environment, thereby guaranteeing a minimal payoff of $2\underline{x}$, or the principal can leave the decision completely up to the agent, trusting that the agent will not choose an x below \underline{x} .⁴

Treatments

We conjecture that the principal's choice to control has two main effects. First, as conjectured by F&K, it motivates reciprocity on the side of the agent. Second, as a consequence of a reduction in decisional autonomy, it crowds out the agent's intrinsic motivation to contribute. We call the first the reciprocity effect, and the second the autonomy effect. In order to separate these two effects we consider 2 distinct treatments: Treatment 1 (C10) and Treatment 2 (TP10). In C10, the principal chooses whether or not to control (replicating F&K's baseline treatment with $\underline{x} = 10$). In TP10 the decision to control is taken by a neutral third party (i.e. a subject outside the main principal-agent interaction) whose payoff is not affected by the agent's choice as the third party is only paid a show-up fee (\$5). The third party chooses whether or not to require $x \geq \underline{x} = 10$. Each agent makes their decision using the strategy method specifying the level of x in the condition when the principal exerts control and the level of x when the principal does not exert control, or, in the TP10 treatment, the level of x in the condition when the third party exerts control and the level of x when the third party does not exert control. Since in TP10 the principal is only a passive player, no reciprocity motive can explain the agent's behavior in this treatment.

⁴The design does not involve different treatments in which principals can make choices to control or not. Control is not a treatment but is a strategy (the principal can decide either to control or to abstain from control of the agent). We can identify agents' reactions under both conditions by using the strategy method.

The treatment TP10 is different from the treatment EX10 included in F&K's original design. In EX10 the principal and the agent play only the sub-game that follows the decision to control in treatment C10. Such a treatment is thus used to control for the effects associated with an exogenously given smaller size of the agent's choice set. By fixing the size of the choice set ex-ante, however, EX10 cannot control for the effect associated with an exogenous variation in the size of the choice set, i.e. an exogenous variation in decisional autonomy. This can be important as long as the initial size of the choice set is seen as a reference point against which the imposed variation is compared.⁵ Finally, the design is between subjects as in F&K.

Questionnaire study

In addition to the experiment we conduct a questionnaire study to help evaluating the subjects' motivations. In contrast with previous research on control aversion, we do not use F&K's standard questionnaire. Rather, we use a psychological questionnaire aimed at measuring the strength of individuals' considerations for choices considering the roles of impersonal, autonomous or controlling forces (Deci and Ryan, 1985). The questionnaire is called the General Causality Orientation Survey (GCOS) and it has been used and verified in a variety of circumstances to understand peoples' preferences for self-determination or autonomy. In the GCOS, subjects answer questions relating to their preferences for an autonomy orientation, impersonal orientation, or control orientation. As the study focuses on adults' decisions in an economic setting, we employ the original 12-vignette version of the GCOS (see complete questionnaire in Appendix B).

Deci and Ryan define each of the orientations in the following ways. A person's autonomy orientation involves, "a high degree of experienced choice in the initiation and regulation of one's own behavior" and people who rate highly on the autonomy orientation "seek out opportunities for self-determination and choice" (p. 111) or they are more likely to experience intrinsic motivation. With the control orientation, people "seek out, select or interpret events as controlling" with a person who is rated highly on the scale being motivated significantly by extrinsic benefits and rewards. Lastly, with

⁵ An extended literature on a different but related topic shows that contract arrangements can play an important role as reference point in buyer-seller interactions. See, for instance, Fehr et al., (2011).

the impersonal orientation people experience their behavior as “beyond their intentional control.” A person who rates highly on the impersonal orientation may view himself or herself as incompetent, or see their behavior as subject to the whims of impersonal forces.

3. Behavioral Predictions

To outline the behavioral predictions of our experiment let us consider the case of purely distributional preferences first. There is ample evidence that many people are endowed with social preferences such as concern for equity. These intrinsic motives, for example, have been used to explain the fact that in dictator games proposers often give positive amount of money (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002). On the basis of these motives we should expect a large fraction of agents in our principal-agent game to choose an activity x that is strictly larger than the minimum. Moreover, whenever the agent’s choice exceeds constraint $\underline{x}=10$ we should observe no significant difference between the level of action x under the control and no-control condition (i.e. constraint $x \geq \underline{x}$ is not binding). Finally, no significant difference in agent’s behavior should emerge between the two treatments C10 and TP10, as the only agent’s concern is the distribution of payoffs.

F&K show that in treatment C10 the observed agent’s behavior significantly deviates from these predictions. Although the majority of agents selects an activity x that is well above the minimum, a large fraction of them select a lower level of x when the principal decides to control than when the principal decides not to control. They call the latter the “hidden costs of control”. Moreover, they find that the average level of x under the control condition is significantly smaller than the average level of x under the no-control condition, i.e. “the hidden costs of control outweigh the benefits of control.” As argued above F&K’s interpretation is based on negative reciprocity. Agents that are intrinsically motivated to act in the principal’s interest perceive control as a signal that the principal does not expect them to perform well. Some agents might perceive this as a signal of distrust; others might take it as an indication of how much voluntary performance the principal expects. In both cases, the agent reacts by choosing a lower performance than he would have chosen if the principal had decided not to control.

With regard to TP10, where no reciprocity motives are at stake, F&K would predict no hidden costs of control as well as no significant difference in the average level of x selected under the two control conditions for those subjects for which control is not binding ($x \geq 10$). In addition, the average level of x selected in TP10 should be similar to the one observed in C10 when no control is exerted. In the latter condition, in fact, the agent does not punish the principal for his distrust and his decision depends on equity concerns only.

Alternatively, the behavior of the agent in treatment C10 can be interpreted in terms of positive reciprocity. According to this view, intrinsically motivated agents reward the principals for the choice of not controlling, rather than punishing them for the choice of controlling. With respect to C10 the behavioral predictions are the same as in the case of negative reciprocity: a) a large fraction of agents select larger level of x when the principal decides not to control than when the principal decides to control and b) the average level of action x is greater under the no-control condition than under the control condition. The difference is that in this case prediction a) can be reinterpreted as suggesting the existence of “hidden benefits of abstaining from control” rather than “hidden costs of control.” Moreover, behavioral predictions differ with regard to TP10. In the latter case, in fact, the presence of a neutral third party reduces the possibility to reciprocate trust with greater performance so that the level of activity x when the third party decides not to control should be lower than or equal to the level of activity x when the third party decides to control.⁶ At the same time the level of activity x under the control condition should be similar to the one observed under the control condition in C10, when the agent does not reward the principal for his trust and his decision depends on equity concerns only. Therefore, on the basis of positive reciprocity one should predict benefits of abstaining from control in C10: the level of x under the no-control condition in C10 is higher than or equal to the level of x under the no-control condition in TP10 and the level x under the control condition is the same in both C10 and TP10.

Finally, the agent’s behavior in our principal-agent game can be interpreted on the basis of Self-Determination Theory (SDT). According to SDT the lack of decisional

⁶ To be more precise, under no control in TP10 the average transfer should be lower than the average transfer under control in TP10 whenever the transfer induced by the distributional preference is lower than the minimum threshold.

autonomy is one of the main factors that crowd out intrinsic motivations in human interactions. As stated above in our setting intrinsic motives are reflected by the existence of equity concerns, which have been shown to play a relevant role in this type of games. At the same time the results of a questionnaire study reported by F&K suggest that “distrust” and “lack of autonomy” are the two most common feelings among controlled agents, which suggest that also autonomy concerns are important. On this basis SDT would predict that being exposed to control will crowd out the agent’s equity concerns, reducing the level of action x compared to the case of no control. This result should hold independently of the subject exerting control, namely the principal or a third party. Therefore the behavioral predictions could be summarized in the existence of hidden costs of control in both C10 and TP10.

4. Practical Procedures

As in F&K and Ziegelmeyer et al (2012), all experiments were facilitated with the use of z-Tree experimental economics software (Fischbacher, 2007). We used a modified version of the official English-language translations of the F&K instructions, with the minor modifications proposed by the Institutional Review Board of Smith College to make certain differences clear to home language English-speakers.

All sessions were conducted at the Cleve E. Willis Experimental Economics Laboratory at the University of Massachusetts, Amherst. Subjects were invited using the ORSEE recruitment system (Greiner, 2015). All subjects were students at the University of Massachusetts, Amherst. Subjects did not participate in more than one session. Most subjects had participated in at least one other economics experiment, but all were inexperienced in that they had not participated in an experiment of this type before. The subjects interacted only once and each session lasted 45 minutes on average (including time for private payment). Table A.1 summarizes the experimental conditions of the two experiments. Excluding the show-up fee (\$5), participants earned an average of \$10.47. Subjects in the C10 treatment earned \$13.93 on average, whereas subjects in the TP10 experiment earned \$9.11 on average (excluding the third party, average payoffs in TP10 are \$13.67).

At the start of each experimental session subjects arrived and randomly drew a cubicle number. Cubicles are separated from each other visually and physically. Subjects are prohibited from speaking. In C10, half of the subjects were assigned the role of principal and half of the subjects were assigned the role of agent. In TP10, one third of the subjects were assigned to the role of principal, one third to the role of agent, and one third to the role of third party. All subjects received a common set of instructions and all questions were answered privately.

As in F&K and Zieglmeyer et al (2012), the subjects' understanding of the players' choice sets and payoffs were assured by three control questions. Once all subjects had answered the control questions correctly (with opportunities to ask questions privately), the subjects played the principal-agent experiment (C10) or principal-agent-third party experiment (TP10) once. After they had played and before they received information about their payoffs, they filled out the General Causality Orientation Scale discussed in section 2 and a basic demographic survey. Responding to the questionnaire was not incentivized and subjects were told that their responses on the survey were not connected to their final payments. After completing the survey, a payment screen showed final earnings in the experiment. Once payment information was revealed, subjects were called to a cubicle to receive their final earnings privately (including the show-up fee). One ECU was equivalent to \$0.20.

5. Results

In this section, we present our findings about the replication of the F&K experiment (C10) and discuss the subject's answers to the questionnaires. We proceed to discuss the results from the third-party variation (TP10) and the answers to the questionnaires in that experiment.

We report results from two-sided statistical tests and we either reject or do not reject the relevant null hypotheses based on a 5 percent level of significance. Consequently, any reference to "significance" in this section should be read as referring to statistical rather than economic or substantive significance. We refer to the agents' choices as occurring in either the "control" or "no-control" setting, consistent with Zieglmeyer et al (2012).

The sample comprises 235 subjects: 76 subjects in the C10 treatment with 38 subjects playing the role of the principal and 38 the role of the agent; 159 subjects in the TP10 treatment with 53 subjects playing the role of the principal, 53 the agent and 53 the third party. Much of the analysis refers to transfers x (by the agent to the principal).

5.1 The replication (C10)

Result 1: We observe significant hidden costs of control (or benefits of abstaining from control) in C10.

First, consistent with F&K, we present the cumulative distributions of the players' transfers in Figure 1. The no-control distribution is shown by the solid line and the control distribution is shown by the dashed line. Were there no hidden costs of control (or benefits of abstaining from control) then the two distributions would coincide for all $x \geq \underline{x}$. On the contrary, the distributions differ. For each value of $x > \underline{x}$ there are more agents in the no-control condition who choose at least that value of x than in the control condition. For instance, more than 40% of agents choose $x > 20$ when they are not controlled. In contrast, less than 20% of agents choose $x > 20$ if controlled and, hence, many are forced to choose at least 10. A greater mass of x -choices is centered at $\underline{x} = 10$ if the principal restricts the agent's actions.

Second, examining the distributions in greater detail, we follow F&K and Ziegelmeyer et al (2012) by constructing a modified distribution for the no-control condition, such that all $x \leq \underline{x}$ in the no-control condition are set equal to \underline{x} . We reject the null hypothesis that the modified distribution from the no-control setting and the distribution from the control setting are the same (Wilcoxon signed-rank test for paired observations, $z = -3.385$, $p = 0.007$).

We can therefore confirm the results from F&K and from Ziegelmeyer et al, that there are significant costs of control in dyadic principal-agent relationships. But, as Ziegelmeyer et al argue, we should be particularly concerned about hidden costs of

control if they are economically substantial and large enough to undermine the use of incentives in relevant settings. That is, do the costs of control outweigh the benefits of control? Consistent with Ziegelmeyer et al (2012), but inconsistent with F&K, in our replication we find that the costs of control do not outweigh the benefits.




Fig.1. Cumulative distributions of transfers for C10 (n = 38)

Result 2: We find at best weak evidence for a negative effect of control in C10

Table 1 presents the agents' transfers as a function of the principals' decisions in the two treatments. The first row presents the average transfers for each of the control (column 1) and no-control (column 2) conditions in the treatments and the difference between the two (column 3). The second row for each treatment reports the standard deviation, followed by the 1st quartile, the median, and the 3rd quartile. For the difference between x^{NC} and x^C , the 95% bootstrap confidence interval is reported in the second row based on 10^5 replications.⁷

⁷ F&K do not compute bootstrap confidence intervals. However, we have computed them based on the data available as supplementary materials to their article. Our calculations provided the following information on their three treatments based on the value of $x^{NC} - x^C$ reported as follows: mean [bootstrap confidence interval]: C5 Treatment: 12.92 [8.47, 17.39]; C10 Treatment: 5.46 [2.75, 8.17], C20 Treatment: 1.31 [-3.06, 5.69]. These results are consistent with those reported in Ziegelmeyer et al, who also performed bootstrap confidence intervals with F&K's data.

In C10, the mean and median are higher in the no-control condition than in the control condition but the difference is not significant at the 5%. The 95% bootstrap confidence interval of the difference $x^{NC} - x^C$ includes zero suggesting that the hidden costs of control (benefits of abstaining from control) do not significantly outweigh the benefits of control.⁸ On the contrary, the 90% confidence interval excludes zero so there is at best weak evidence for a negative effect of control as in F&K.⁹

Table 1
Agents' Transfers as a function of the principal's decision

	Control Condition	No-control condition	$x^{NC} - x^C$
C10	17.47 (11.76, 10, 10, 20)	21.24 (19.03, 0, 18.5, 40)	3.76 [-.71, 8.24]
TP10	18 (11.62, 10, 10, 25)	12.87 (15.50, 0, 5, 30)	-5.13 [-6.53, -3.73]

Notes: Mean agents' transfers. Standard deviation, 1st quartile, median, and 3rd quartile reported in parentheses. The 95% bootstrap confidence interval of the difference $x^{NC} - x^C$

5.2 The third-party treatment (TP10)

Result 3: We do not observe hidden costs of control (benefits of abstaining from control) in TP10

In Figure 2, we show the cumulative distributions of transfers in the control (dashed line) and no-control (solid line) conditions. When there are no hidden costs of control (or benefits of abstaining from control), the values of the cumulative distributions ought to coincide for all $x \geq \underline{x}$. As can be seen in the CDFs, the two distributions approximately coincide for all $x \geq \underline{x}$.

We also compare modified distributions in TP10, where we replace all $x < 10$ by $x = 10$. We do not reject the null hypothesis that the modified distribution from the no-control

⁸ We should note, however, that though we cannot reject the null that the bootstrap 95% confidence interval includes zero, we also cannot reject that it includes 5.46 (the average from F&K's experiment).

⁹ The 90% confidence interval is [.016, 7.51].

setting and the distribution from the control setting are the same (Wilcoxon signed-rank test for paired observations, $z = 0.840$, $p = 0.4007$).




Fig. 2. Cumulative Distributions of Transfers from TP10, $n = 53$ Agents

Result 4: The total effect of control is beneficial in TP10

The properties of Table 1 were discussed in detail in Result 2. In TP10, the mean transfer is lower in the no-control condition than in the control condition. The x transferred by the agent is significantly lower in the no-control condition than in the control condition (Wilcoxon signed-rank $z = 5.030$, $p < 0.01$). The mean of the difference $x^{\text{NC}} - x^{\text{C}}$ is negative and the bootstrap confidence interval around the mean excludes zero. Referring to Table 2, the expected payoff to a principal when a third party exerts control is greater than the expected payoff to a principal when a third party does not exert control. Therefore, the total effect of control is beneficial.

5.3 Comparing C10 and TP10

Result 5: We observe hidden benefits of abstaining from control. In the no-control condition, agents' transfers are lower in TP10 than in C10; in the control condition, transfers are similar across the two treatments.

Differences in agents' reaction to control in TP10 compared to C10 suggest that intention-based social preferences rather than preferences for autonomy account for the observed behavior in this particular experimental setting. However, it is not clear which type of reciprocity model fits the data better. As discussed in section 3, agents may be rewarding the principals for the choice of not controlling (positive reciprocity) or punishing them for the choice of controlling (negative reciprocity).

We provide further evidence by comparing the two treatments under both the control and no-control condition. It is immediately evident from Table 1 that the average transfer in the control condition is similar across the two treatments (17.47 and 18). In contrast, the average transfer in the no-control condition is substantially higher in C10 than in TP10 (21.24 and 12.87). We then compare the distributions under the control and no-control conditions between the two treatments using Kruskal-Wallis tests. For the control condition, we do not reject the null hypothesis that the distributions are equivalent across the two treatments ($X^2 = 0.007$ (1 d.f.), $p = 0.9326$). For the no-control condition, we reject the null hypothesis that the distributions are equivalent ($X^2 = 4.048$ (1 d.f.), $p = 0.0442$).

The result above sheds light on the specific behavioral mechanism at work in the F&K experiment. The fact that effort is higher in C10 than in TP10 in the no-control condition suggests that the interpretation of the result in terms of "hidden benefits of abstaining from control" rather than in terms of "hidden cost of control" is more convincing. Interestingly, effort is not smaller in C10 than in TP10 in the control condition. Hence, the behavioral response that F&K interpret as a negative reaction to the principal's choice of controlling seems more sensibly interpreted as a positive reaction to the principal's choice of not to controlling.

As a consequence of using the strategy method to elicit the agents' choices, we can observe whether the players are heterogeneous in their types by gaining greater understanding of whether players react positively (i.e. they increase the transfer), neutrally (i.e. they do not vary the transfer) or negatively (i.e. they reduce the transfer) to the choice of abstaining from control (control abstinence).

Result 6: Players react to control abstinence heterogeneously in both treatments.

Table 2 summarizes the agent's responses to control abstinence in each of the experiments. In C10, our result where 42.10% of agents react positively to control abstinence are consistent with Ziegelmeyer et al who found a proportion of agents who lower the transfer when controlled in the range 40.62% to 45.45%.¹⁰ This result contrasts with F&K who found that a majority of agents (56.94%) reduced performance when controlled. The proportion of agents who react negatively to control abstinence (36.84%) is consistent with Ziegelmeyer et al (39.40% to 60.00% in various C10 experiments) rather than relative to F&K (25% in C10). In C10, the minority of subjects responds neutrally to control abstinence (21.05%). Agents who positively reciprocate control abstinence transfer approximately the same in C10 as they do in in F&K and in Ziegelmeyer et al, that is, they transfer roughly double in the no-control condition relative to the control condition. Agents who respond neutrally to control abstinence transfer (25.13) within the range of what they did in F&K (22.3) and Ziegelmeyer et al (range: 14.8 to 30.71).¹¹

Table 2
Agents' behavioral reactions to control abstinence

		Negative	Neutral	Positive	Expected Transfer to P
C10	Relative Share	36.84%	21.05%	42.11%	
	Mean Control Transfer	11.64 (4.40)	25.13 (19.16)	18.75 (9.40)	17.47
	Mean no-control Transfer	1.93 (5.09)	25.13 (19.16)	36.19 (10.23)	21.24
TP10	Relative Share	54.72%	41.51%	3.77%	
	Mean Control Transfer	10.76 (2.21)	27.36 (12.67)	20 (7.07)	18
	Mean no-control Transfer	1.14 (2.70)	27.36 (12.67)	23.5 (9.19)	12.87

Notes: n = 38 in C10 and n = 53 in TP10; Expected Transfer to P is calculated as a weighted average where the mean transfer in each condition is multiplied by the proportion of each type

¹⁰ Ziegelmeyer et al. follow F&K in interpreting the results of their experiment in terms of negative reciprocity. It follows that their classification considers negative, neutral and positive reactions to the choice of controlling rather than positive, neutral and negative reactions to the choice of not controlling.

¹¹ In the Online Appendix (Tables OA.1 and OA.2), we use regression analysis to examine the extent to which our results differ from those in F&K and Ziegelmeyer et al.

In TP10, on the other hand, few agents respond positively to control abstinence (3.77%), a large proportion responds neutrally to control abstinence (41.51%) and the majority of agents respond negatively to control abstinence (54.72%). Of those agents who respond negatively to control abstinence, 86% increased their payoffs to exactly 10 when controlled. That is, they experienced control as binding on their choices. These results are consistent with the preceding results examining the distributions of transfers and the difference between control and no-control transfers. We follow Ziegelmeyer et al (2012) in using Fisher's exact tests to compare the frequencies of type of reaction across treatments. Fisher's exact tests contrasting differences in frequencies between the treatments suggest the following: we reject the null hypothesis that the frequencies are equal for neutral responses to control abstinence ($p = 0.045$) and positive responses to control abstinence ($p < 0.001$), but we cannot reject the null hypothesis that the proportions are equal for negative responses to control abstinence ($p = 0.136$).¹²

Result 7: The proportions of principals and third parties who exert control are not significantly different

In C10, 63.15% of principals exert control. In TP10, 77.34% of third parties exert control (Fisher's exact test, $p = 0.163$). A 95% bootstrap confidence interval of the difference in the proportions contains zero ($-0.051 < p_{C10} - p_{TP10} < 0.335$). These results are consistent with Ziegelmeyer et al who found proportions of control ranging from 57% to 83% in their C10 treatments. Both our results and Ziegelmeyer et al's results suggest significantly higher proportions of control than F&K who found 29% of principals choosing control.

5.4 Agents' reactions to control abstinence and GCOS score

Finally, we study the correlation between agents' reactions to control abstinence and their corresponding scores in the GCOS. As mentioned, the questionnaire measures the strength of individuals' considerations for choices considering the roles of impersonal, autonomous or controlling forces (Deci and Ryan, 1985). The purpose of this exercise is to test whether the GCOS, widely used in social psychology, is able to capture in

¹² We also include Figure OA.1 in the Online Appendix, which contrasts the frequencies of each reaction type (negative, neutral, positive) by treatment.

reasonable ways the observed heterogeneity in individuals' reactions to control abstinence in an incentive-compatible experiment.

For robustness across the treatments, we confirm that the samples are not statistically significantly different with respect to the subjects' reported attitudes using the GCOS. The means and standard deviations for the subjects' reported preference for each scale in each of the treatments are reported in Table 3.¹³

Table 3
Summary of General Causality Orientation Scale Indexes for Agents Only

	C10	TP10	t-stat/ (Mann-Whitney z)
GCOS: Autonomy Scale	73.53 (5.99)	71.30 (7.138)	1.56554 (1.468)
GCOS: Impersonal Scale	43.32 (10.64)	44.19 (9.892)	-.4022743 (-0.443)
GCOS: Control Scale	59.66 (7.235)	58.45 (7.523)	.7656292 (0.782)
Observations	38	53	91

Notes: Mean values. Standard deviation in parentheses. 12-vignette version of the GCOS (Deci and Ryan, 1985). Complete questionnaire included in the Online Appendix as well as the summary statistics for the entire sample. See also Figure OA.2 in the Online Appendix.

The means in the scales are not statistically significantly different across treatments, as shown by the t-statistics of the difference between their values by treatment.¹⁴ The regressions report results for our data only as F&K or Ziegelmeyer et al did not gather the GCOS attitudes.

In Table 4 the first column represents an OLS regression with x^{NC} (the transfer in the no-control condition). Column 2 uses the transfer difference $x^{NC} - x^C$ as the dependent variable as in F & K and Ziegelmeyer et al. In specifications 3 through 5, the dependent variable was a dummy variable indicating whether a subject displayed a negative (neutral or positive) response to control abstinence.

The standardized control GCOS variable was statistically significant and negative in the x^{NC} and $x^{NC} - x^C$ regression (columns 1 and 2), the negative response to control abstinence logit regression (column 3) and the positive response to control abstinence

¹³ Table OA.3 in the Online Appendix reports GCOS Indexes pooling all subjects.

¹⁴ Regarding the internal consistency of each of the three subscales, the Cronbach's α non-standardized values were autonomy, 0.8469; impersonal, 0.7394; and control, 0.6218.

logit regression (column 4). A one standard deviation increase in the standardized control scale corresponds with a decrease in the probability a subject will respond positively to control abstinence by 6.4%, an increase in the probability the subject will respond negatively by 16.6%, a decrease in the transfer (x^{NC}) in the no-control condition, and an decrease in the difference between x^{NC} and x^C of 2.17 units of x (with the coefficient of -2.17, the difference increases in magnitude).

Table 4
Regressions from our Subject Pool Only (n = 91) with GCOS variables

VARIABLES	(1) x^{NC}	(2) $x^{NC} - x^C$	(3) Negative	(4) Neutral	(5) Positive
D: TP10 Treatment = 1	-8.235** (3.875)	-8.864*** (2.444)	0.193* (0.111)	0.216** (0.0961)	-0.403*** (0.0830)
Standardized Autonomy Scale	3.090 (1.885)	1.198 (1.363)	-0.105 (0.0676)	0.0800 (0.0615)	0.00241 (0.0320)
Standardized Impersonal Scale	-0.190 (2.167)	-0.611 (1.259)	-0.0465 (0.0644)	0.0343 (0.0533)	0.0149 (0.0294)
Standardized Control Scale	-4.573** (2.265)	-2.177* (1.303)	0.166** (0.0701)	-0.0675 (0.0514)	-0.0641** (0.0299)
Constant	20.76*** (3.136)	3.551 (2.324)			
Observations	91	91	91	91	91
R-squared	0.131	0.212			
Log Likelihood			-57.47	-54.16	-32.22

Notes: Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This result is consistent with the psychological interpretation given to the control orientation, which assesses the extent to which a person is oriented toward being controlled by rewards and the directives of others (Deci and Ryan, 1985). In line with the idea that agents' reactions to control in the experiment are not driven by individuals' preferences for self-determination, the standardized autonomy GCOS does not show significant correlation either with the probability that the subject will respond positively to control abstinence or with the difference between x^{NC} and x^C . This result, however, should be interpreted cautiously as large standard errors may simply be due to a lack of statistical power.

6. Discussion and conclusion

We provide further evidence on the behavioral motives underlying individual reactions to control in an experimental principal-agent game by introducing a third-party

treatment. First, in line with Ziegelmeyer et al (2012), we fail to replicate F&K's finding on the negative effect of control for the principal's profit. Although we largely confirm that a large fraction of agents exert lower effort under the control condition than under the no-control condition, we find that the proportion is not substantial enough to significantly undermine the benefits of control. Second, in the presence of a third party rather than a principal who exerts control, the fraction of agents that reduce performance dramatically decreases and control results in a greater expected profit for the principal. But, this should not be viewed as a form of delegation or the third party acting on behalf of (or at the orders of) the principal,¹⁵ rather it suggests that the agents respond reciprocally toward the principals in C10, but do not have that motive in TP10.

Indeed, we find that when either the principal or a third party decides to control the average level of effort that is selected by the agents is similar. The difference between the two treatments is entirely driven by the agent's reaction to the decision not to control. When a third party decides not to control, the decision results in a significant lower level of effort by the agent compared to the case in which the decision not to control is taken by the principal. This favors an interpretation of the results in terms of positive reciprocity: agents reciprocate not controlling principals with greater effort rather than punishing controlling ones with lower effort (see for instance, Levine 1998). Our interpretation has important implications for the design of organizations. Rather than delegating control to third parties (e.g. consultancies, governments), as suggested by F&K's interpretation based on negative reciprocity, a principal would find it beneficial to be in a position in which he could control, but then abstain from doing so.

One could argue that the treatment EX10 conducted by F&K is a proper baseline treatment whereas our TP10 is not and, therefore, that our interpretation in terms of positive reciprocity is not guaranteed.¹⁶ While it is true that the actions of the third party

¹⁵ The effect of delegation of responsibility is studied, for instance, by Bartling and Fischbacher (2011), Coffman (2011) and Oexl and Grossman (2013). With respect to control aversion it could be interesting to study the factors that affect the perceived responsibility of delegated control. This aspect, however, goes beyond the scope of the present paper and is thus left for future research.

¹⁶ We are grateful to an anonymous reviewer for this observation. As explained in section 2, in EX10 the restriction on the agent's choice is exogenously imposed and, hence, intentions on the part of principal are omitted from the setting. The median transfer observed by F&K in EX10 is not significantly different from the median transfer observed in their C10 without control, while the transfers when control is endogenous are significantly lower than when it is exogenous. This leads F&K to interpret hidden cost of control in terms of negative reciprocity

in TP10 can be interpreted in different ways and could activate other motives¹⁷, several factors change between EX10 and C10 which could drive the results in F&K. For instance, the fact that in EX10 the strategy space is bounded between 10 and 120 makes the comparison with the trust subgame of C10 difficult, even if one takes the different support of x into account as F&K and we do. If the value of the minimum transfer affects individual decisions, e.g. acting as a reference point, the two subgames are hardly comparable and little can be said about the type of reciprocity driving individual behaviors. Moreover, in contrast with C10, in EX10 the agent makes only one choice that does not involve the strategy method. This limited choice transforms the game into a dictator-setting, which may also affect the results. Finally, as argued above, EX10 cannot control for the effect associated with a variation in the size of the choice set, which instead plays a crucial role in both C10 and TP10. Overall, we believe that although F&K's interpretation of the difference between EX10 and C10 should be taken into account, our interpretation in terms of positive reciprocity is more convincing.

Our results can be used to differentiate between alternative explanations of individual reactions to control, such as Sliwka (2007) and Ellingsen and Johannesson (2008). In Sliwka (2007), reactions to control derive from signals about general social norms that are sent by the controlling party. The author considers a pool of agents with three behavioral types: selfish, inequity averse and conformist, i.e. agents that change behavior depending on the signal sent by the controlling party. The decision to control (not to control) is interpreted as a signal that a large fraction of agents is expected to be selfish (inequity averse), which leads conformist agents to behave according to the generalized social norms. In this framework, the effect of control does not stem from principal-specific characteristics and the control decision of the third party should equally affect the updating about the generalized social norm. The prediction is thus that control by a third party should have the same effect as control by a second party. Our results, however, contradict that prediction. Ellingsen and Johannesson (2008), on the contrary, explain individual reactions to control on the basis of positive reciprocity. In their model, reduced transfers are the consequence of changed image concerns that

¹⁷ For instance, an anonymous referee suggests that the third party's decision not to control may be interpreted by the agent as a suggestion that being opportunistic is perfectly appropriate. However, if present, this signalling effect should also hold when the principal decides to abstain from control, which is not the case either in our data or in F&K's data.

directly relate to the principal's control decision. By controlling, the principal signals that “he is not worth impressing”, while by not controlling he signals the opposite. The signaling effect is thus principal-specific and it should disappear under third-party control. This prediction is indeed consistent with our findings. Our results suggest that positive reciprocity-based explanations are more consistent with the observed behavior than explanations based on conformism.

That positive reciprocity, rather than preferences for autonomy or negative reciprocity, drives the behavior of agents in these interactions is the main contribution of this paper. In demonstrating this result, we contribute to a wider literature engaged with understanding the employment relation, hierarchy, coercion and the exercise of power (Fehr et al, 2013; Nikiforakis et al, 2014). Of course, the limited role played by preferences for autonomy in our very specific and highly stylized experimental setting should not be interpreted as a general claim about the irrelevance of this type of preferences, which have been proven to be very salient in other settings (see, for instance, Bartling et al, 2014). It is worth noting that the simplicity of the proposed experimental task (i.e., to report a natural number via a computer interface) may reduce the salience of self-determination motives among subjects. In addition, the SDT literature has studied autonomy in interaction with two other basic psychological needs: competence and relatedness (Deci and Ryan, 2000).¹⁸ Such complex interplay may be difficult to recreate in simple experimental games.

Future work should examine the extent to which preferences evolve over repeated principal-agent interactions and interactions in which the hierarchical relationship between subjects in the experiments may be made clearer either through framing or through changes in experimental design where the loci of control for the principal are more diverse. This may permit researchers to examine more unambiguously the extent to which autonomy and reciprocity may complement or substitute for each other and, therefore, the extent to which extrinsic benefits may crowd out or in the effort of agents.

Acknowledgements

¹⁸ Competence refers to a feeling of efficacy with respect to a particular goal or task. Relatedness is defined as a “sense of belongingness and connectedness to the persons, group, or culture disseminating a goal (see Deci and Ryan, 2000: p.64).

The experiments were funded by the Santa Fe Institute Cowan Fund. Additional support was provided by the Smith College Committee on Faculty Compensation and Development. We thank the editor, associate editor, and two anonymous referees for valuable comments that have improved the paper. We are also grateful to Avner Ben-Ner, Sam Bowles, Alessandro Innocenti, Natalia Montinari, Marco Piovesan, Antonio Nicolò, and participants at the SFI seminar, 2015 Canadian Economic Association meeting and 2017 Socrates workshop, particularly to David Freeman, Glenn Harrison, David Kingsley, Peter Matthews, Rajiv Sethi, John Spraggon, and Sylvie Thoron. The usual disclaimer applies.

References

- Acemoglu, D. and F. Newman, A., 2002. "The labor market and corporate structure," *European Economic Review*, 46(10), 1733-1756.
- Alchian, A. and Demsetz, H., 1972. "Production, Information Costs, and Economic Organization," *American Economic Review*, 62(5), 777-95.
- Bartling, B., Fehr E., Schmidt K., 2013. "Discretion, Productivity, and Work Satisfaction," *Journal of Institutional and Theoretical Economics*, vol. 169(1), 4-22.
- Bartling B., Fehr E. and Herz H., 2014. "The Intrinsic Value of Decision Rights," *Econometrica*, 82(6), 2005-2039.
- Bartling, B., and Fischbacher, U., 2012. "Shifting the blame: on delegation and responsibility". *Review of Economic Studies*, 79(1), 67-87.
- Batson, C. D., Coke, J. S., Jasnoski, M. L., and Hanson, M., 1978, "Buying kindness: effect of an extrinsic incentive for helping on perceived altruism", *Personality and Social Psychology Bulletin*, 4(1), 86-91.
- Ben-Ner, A., 2013, "Preferences and Organization Structure: Towards Behavioral Economics Micro-Foundations of Organizational Analysis", *Journal of Socio-Economics*, 46, 87-96.
- Bolton G. E. and Ockenfels A., 2000. "ERC: A Theory of Equity, Reciprocity, and Competition", *American Economic Review*, 90 (1), 166-193
- Bowles, S. and S. Polania, 2012, "Economic Incentives and Social Preferences: Substitutes or Complements?", *Journal of Economic Literature*, 50(2), 368-425.
- Charness G. and Rabin M., 2002. "Understanding Social Preferences with Simple Tests," *The Quarterly Journal of Economics*, 117(3), 817-869
- Charness, G., Cobo-Reyes, R., Jinéñez, N., Lacomba, J. A, Lagos, F., 2012, "The Hidden Advantage of Delegation: Pareto-improvements in a Gift-exchange Game", *American Economic Review*, 102(5), 2358-79.
- Chirkov, V. I., R. M. Ryan and K.M. Sheldon, 2010. "Human Autonomy in Cross-Cultural Context", *Perspectives on the Psychology of Agency, Freedom, and Well-Being. Cross-Cultural Advancements in Positive Psychology*, Vol 1. Springer.
- Coffman, L. C., 2011. "Intermediation reduces punishment (and reward)", *American Economic Journal: Microeconomics*, 3(4), 77-106.
- Deci, E. L. and R. M. Ryan, 1985. "Intrinsic Motivation and Self-Determination in Human Behavior", New York: Plenum Press.
- Deci, E. L. and R. M. Ryan, 2000. "Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions", *Contemporary Educational Psychology* 25, 54-67.

- Dufwenberg M. and G. Kirchsteiger, 2004. "A Theory of Sequential Reciprocity", *Games & Economic Behavior* 47(2), 268-98.
- Ellingsen, T. and Johannesson, M., 2008. "Pride and Prejudice: The Human Side of Incentive Theory", *American Economic Review*, 98(3), 990-1008.
- Fabes, R. A., Fultz, J., Eisenberg, N., May-Plumlee, T., and Christopher, F. S., 1989. "Effects of rewards on children's prosocial motivation: a socialization study", *Developmental Psychology*, 25(4), 509-515.
- Fafchamps, M. and Söderbom M., 2006. "Wages and Labor Management in African Manufacturing," *Journal of Human Resources*, 41(2), 356-379.
- Falk, A. and Fischbacher, U., 2006. "A theory of reciprocity." *Games and Economic Behavior* 54(2), 293-315.
- Falk, A. and M. Kosfeld, 2006. "The Hidden Costs of Control", *American Economic Review*, 96(5): 1611-1630.
- Fehr E. and Schmidt K. M., 1999. "A Theory of Fairness, Competition, and Cooperation," *The Quarterly Journal of Economics*, 114(3), 817-868
- Fehr, E. and Gächter S., 2000. "Fairness and Retaliation: The Economics of Reciprocity." *Journal of Economic Perspectives* 14(3): 159-181.
- Fehr, E., Hart O., and Zehnder C., 2011. "Contracts as Reference Points: Experimental Evidence." *American Economic Review*, 101(2): 493-525.
- Fehr, E., Herz H., and Wilkening T., 2013. "The Lure of Authority: Motivation and Incentive Effects of Power." *American Economic Review*, 103(4): 1325-59.
- Fischbacher, U., 2007. "z-Tree: Zurich toolbox for ready-made economic experiments", *Experimental Economics*, 10(2): 171-178
- Frey, B. S. and R. Jegen, 2000, "Motivation Crowding Theory: A Survey of Empirical Evidence", CESifo Working Paper No. 245.
- Gagné, M., 2003, "The role of autonomy support and autonomy orientation in the engagement of prosocial behavior", *Motivation and Emotion*, 27(3), 199-223.
- Gagné, M. and E.L. Deci, 2005. "Self-determination Theory and Work Motivation", *Journal of Organizational Behavior*, 26(4), 331-362.
- Greiner, B., 2015. "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE". *Journal of the Economic Science Association*, 1(1), 114-125.
- Greene-Demers, I., Pelletier, L. G., and Ménard, S., 1997. "The impact of behavioral difficulty on the saliency of the association between self-determined motivation and environmental behaviors", *Canadian Journal of Behavioural Science*, 29(3), 157-166.
- Jayadev, A. and Bowles, S., 2006. "Guard labor," *Journal of Development Economics*, 79(2), 328-348.
- Kunda, Z., and Schwartz, S. H., 1983. "Undermining intrinsic moral motivation: external reward and self-presentation", *Journal of Personality and Social Psychology*, 45, 763-771
- Levine, D. K., 1998. "Modeling Altruism and Spitefulness in Experiment," *Review of Economic Dynamics*, 1(3), 593-622.
- Oexl, R., and Grossman, Z. J., 2013. "Shifting the blame to a powerless intermediary". *Experimental Economics*, 16(3), 306-312.
- Owens Jr. D., Z. Grossman Jr., and R. Fackler Jr., 2014. "The Control Premium: A Preference for Payoff Autonomy," *American Economic Journal: Microeconomics*, 6(4), 138-61.
- Nikiforakis, N., Oechssler, Jörg and Shah, Anwar, 2014. "Hierarchy, coercion, and exploitation: An experimental analysis," *Journal of Economic Behavior and Organization*, 97, 155-168.
- Pelletier, L. G., Tuson, K. M., Greene-Demers, I., Noels, K., and Beaton, A. M., 1998, "Why are you doing things for the environment? The Motivation Toward the Environmental Scale (MTES)", *Journal of Applied Social Psychology*, 28(5), 437-468.

- Ryan, R.M. and Deci, E. L., 2000. "Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being", *American Psychologist*, 55(1), 68-78
- Riener, Gerhard and Wiederhold, Simon, 2016. "Team building and hidden costs of control", *Journal of Economic Behavior and Organization*, 123, 1-18.
- Schnedler, W. and Radovan Vadovic, 2011. "Legitimacy of Control," *Journal of Economics and Management Strategy*, 20(4), 985-1009.
- Sliwka, D., 2007. "Trust as a Signal of a Social Norm and the Hidden Costs of Incentive Schemes", *American Economic Review*, 97(3), 999-1012.
- Sobus, M. S., 1995. "Mandating community service: psychological implications of requiring prosocial behaviour", *Law and Psychology Review*, 19, 153-182.
- Upton, W. E. III., 1974. "Altruism, attribution and intrinsic motivation in the recruitment of blood donors", in *Selected readings in donor recruitment* (Vol. 2, pp. 7-38). Washington, DC: American National Red Cross.
- Von Siemens, F. 2013, "Intention-based reciprocity and the hidden costs of control." *Journal of Economic Behavior and Organization*, 92, 55-56
- Ziegelmeyer, A, Schmelz, K., and Ploner, M., 2012, "Hidden Costs of Control: Four Repetitions and an Extension", *Experimental Economics*, 15(2), 323-340.

Appendix A

Table A.1
Experimental Conditions

C10	Number of Sessions	4
	Number of Subjects	76
	Gender (% Female)	44%
	Average age	21.02 (2.34)
	Agents' Average Earnings	20.14 (2.9486)
	Principals' Average Earnings	7.73 (5.8972)
TP10	Number of Sessions	7
	Number of Subjects	159
	Gender (% Female)	54%
	Average age	20.75 (4.2)
	Agents' Average Earnings	20.66 (2.5930)
	Principals' Average Earnings	6.69 (5.1860)
	Third Party's Average Earnings	0

Notes: Earnings are stated in dollars net of the show-up fee with standard deviations in parentheses. Third parties were simply paid the show-up fee and therefore would have no payoff net of the show-up fee.