

Chapter 1

Introduction

A core principle of economics is that individuals respond to incentives. Agency theory endorses the view that tighter monitoring invariably leads to profitable increases in employees' effort and recommends employers to offer jobs characterized by limited effort discretion. However, since the work by Deci (1971), many social scientists reject this view and acknowledge that monitoring is likely to have two opposite effects on the employees' performance: A *disciplining* effect (monitoring limits shirking and therefore raises work effort) and a *crowding-out* effect (monitoring undermines intrinsic motivation and therefore reduces work effort).

The crowding-out effect is viewed as a major anomaly in economics (Frey and Jegen, 2001), and it has severe implications for the design and management of reward systems. Incentive strategies are needed to regulate the opportunistic employees but at the same time management strategies need to sustain the motivation of employees disinclined to shirk. In order to be able to design profitable employment contracts, we have to improve our understanding of the determinants of the crowding-out effect.

In three essays, I investigate the interaction between control and motivation in employment relationships. I consider a simple agency relationship where the agent chooses among productive activities after the principal determined the agent's choice set. To each productive activity is associated a cost for the agent and a benefit for the principal. The extent to which the principal restricts the agent's choice set corresponds to the degree of control implemented by the former in the work environment of the latter. Whether introducing control is beneficial to the principal depends on the relative size of the disciplining and the crowding-out effects and both effects are likely to vary in size with individual, situational and cultural conditions.

The three studies of my dissertation extend the economic literature on individual (Chapter 2), situational (Chapter 3), and cultural conditions (Chapter 4) under which crowding-out effects arise, and under which they are predicted to dominate disciplining effects.

Three studies on the determinants of control aversion

I will mainly use the term *control aversion* when referring to the crowding-out effect. Control averse individuals show a strictly higher effort if not controlled than if controlled. Control aversion can be interpreted as punishing the implementation of control, or as rewarding the absence of control. Accordingly, one could say that control averse individuals react negatively to control, or that control averse individuals react positively to freedom. To ease the discussions, I will mainly use the first phrasing. On a sample level,

control aversion implies that the principal suffers *hidden costs of control* which increase with the fraction of control averse agents and the relative size of their efforts if controlled and if not controlled.

Chapter 2

Chapter 2 reports on a robustness check of the surprising findings by Falk and Kosfeld (2006, henceforth F&K). The latter study implements a principal-agent game where the agent chooses to transfer part of his endowment to the principal who has the possibility to enforce a minimal transfer. F&K's experimental results show that the control strategy entails hidden costs caused by the existence of control-averse agents who choose a lower transfer if forced than otherwise and that these hidden costs outweigh the benefits of control.

In collaboration with Matteo Ploner and Anthony Ziegelmeyer I repeat the low and medium control treatment of the principal-agent game to check the robustness of F&K's surprising experimental results to small changes in subject pools. Each of the three repetitions employs a sample drawn from a standard subject pool of students and demographics vary across samples.

Our laboratory results largely confirm the existence of hidden costs of control but, contrary to the original study, hidden costs of control are usually not substantial enough to significantly undermine the effectiveness of economic incentives. Our participants were asked, at the end of the experimental session, to complete a questionnaire in which they had to state their work motivation in hypothetical scenarios. Our questionnaires are identical to the ones administered in F&K's questionnaire study. In contrast to the game play data, our questionnaire data are similar to those of the original questionnaire study. In an attempt to solve this puzzle, we report an extension where performance-contingent earnings are absent i.e. both principals and agents are paid according to a flat participation fee. We observe that hidden costs outweigh benefits of control which shows that control aversion is more prevalent under hypothetical than under real incentives. Still, in the low control treatment, we observe much weaker negative responses to control in our extension than in the original study. This observation, the fact that the original study uses real incentives, and preliminary findings on the relationship between demographics and the degree of control aversion lead us to conclude that F&K's experimental regularities originate from the characteristics of their participants.

Moving from Chapter 2 to Chapters 3 and 4

I believe that F&K's principal-agent game was a good starting point for my dissertation to study the relative importance of explicit and implicit incentives, and their interaction effects. Chapters 3 and 4 move away from F&K's design in four respects.

First, the two chapters rely on a repetitive trial environment meaning that participants play the game repeatedly. The fact that F&K's study and our three repetitions made use of a single-shot trial environment is likely to account for part of the observed variability in participants' behavior. A repetitive trial environment will allow participants to gain experience with the interactive situation and experienced behavior in the control game is likely to be a more reliable measure of control aversion. We implement a matching protocol which best preserves the nature of one-shot interactions.

Second, we consider a principal-agent game which is a straightforward extension of the game implemented by F&K. To diminish the confounding effect of opportunism on

control aversion, we define a payoff function with convex effort costs for the agent and concave benefits for the principal such that expressing the dislike of control bears low costs for selfish agents.

Third, we complement laboratory experiments with internet experiments employing heterogeneous populations. A substantial amount of my doctoral time has been dedicated to the development of a platform for conducting internet experiments. It has been developed for the purpose of this thesis and further cross-cultural experiments. This platform allows setting up synchronized experiments where all participants are present at the same time (drop-outs are taken care of) as well as asynchronous experiments where participants enter and leave the platform on many occasions not necessarily at the same time.

Fourth, game play data are complemented with survey data on a large spectrum of individual characteristics, subjective attitudes to control and information about the workplace environment.

Chapter 3

F&K's experimental work has informed us about some of the situational factors that influence control aversion (e.g. the level of control). Our Experiment 4 with hypothetical incentives has complemented this work by showing that our subjects are more control-averse in the considered principal-agent game under hypothetical than under real incentives. Chapter 3 is an attempt to improve our knowledge on the situational determinants of control aversion (in collaboration with Anthony Ziegelmeyer). In this thesis, Chapter 3 also serves as a transition from Chapter 2 to Chapter 4. We compare a laboratory and an internet implementation of an experimental principal-agent game where the principal can impose control at two different levels on the agent.

Bruno S. Frey and his coauthors have repeatedly argued that the more personal the relationship between employers and employees, the more likely monitoring reduces work effort and performance (Frey, 1993; Frey and Jegen, 2001). The argument distinguishes work situations where the intensity of the personal relationship is strong from work situations where relationships are impersonal. Agency relationships are arguably less personal in the internet than in the laboratory setting. In addition to providing an experimental test of Frey's argument, our study addresses several methodological issues related to the measurement of control aversion. The results of the two experiments inform us about the viability of the Internet as an alternative to the laboratory for the experimental investigation of control aversion in employment relationships.

Our experimental results can be summarized as follows. First, benefits of control outweigh hidden costs of control in both experiments. Second, we show that negative reactions to control are stronger on average in the laboratory experiment than in the internet experiment, and that this difference is stronger the more effort is enforced by the principal. This second result confirms Frey's argument since a smaller hierarchical distance induces stronger control aversion. Third, we show that negative reactions to control are of similar magnitude at the end of the two experiments which indicates that experienced agents react similarly to the implementation of control in both environments. Though hidden costs of control are quite stable in the laboratory experiment, they increase as the session progresses in the internet experiment. The observed dynamics of hidden costs of control in the internet experiment suggest that even in highly impersonal employment relationships the imposition of control on employees might backfire on employers. We

conclude that the Internet is a viable alternative to the laboratory for the experimental investigation of control aversion.

Chapter 4

The large variance in control aversion among student samples uncovered in Chapter 2 remains a puzzle. Though non-systematic, this evidence suggests that societal background could at least partially account for the behavioral heterogeneity. (How) does societal background affect control aversion?

Chapter 4 is an attempt to answer this question. We investigate systematically the impact of politico-economic systems on reactions to control by employing participants of different generations from East and West Germany. Reunified Germany offers a perfect opportunity to study the impact of politico-economic systems on control preferences as both parts of Germany are homogenous with respect to their basic culture but they have experienced different degrees of political control in the past. Our conjecture is that less individuals acquired control-averse preferences under the coercive regime of East Germany than under the liberal regime of West Germany, and we test this conjecture in a large-scale internet study all over Germany.

This chapter goes beyond a systematic documentation of East-West differences and suggests a mechanism for a causal link from politico-economic systems to control preferences. We present a model of cultural evolution in which the level of state coercion endogenously affects control-averse preferences, so that the distribution of preferences for control or autonomy in a society depends on state coercion.

The empirical part of Chapter 4 is organized in two experiments with four treatments. The core of all treatments is the repeated control game presented in Chapter 3, preceded by the survey on individual characteristics that could potentially be related to control preferences as presented in Chapter 3. In Experiment 1, we move away from established laboratory conditions with only stepwise methodological adjustments to finally conduct a large-scale experiment on the internet. It comprises a set of three treatments conducted with students. These students treatments are implemented in two East German and two West German locations which are matched based on a series of socio-economic attributes. Experiment 2 reflects one treatment conducted all over Germany with professionals, i.e. university graduates who are working.

Our most important results are as follows. Indeed, control aversion is largely stronger among West than among East Germans. The differences are significant and systematic for elder professionals but only in specific cases for students and never for younger professionals. The differences in control aversion are largely driven by lower fractions of control averse individuals in the East whereas the strength of control averse reactions is similar among East and West Germans. Our findings corroborate the idea that for elder generations of East Germans who have been socialized under a coercive regime as children and adolescents, the experience of a liberal regime in their adulthood does not make them abandon completely the preferences they acquired in early life. Differences in control aversion are vanishing in younger generations of East and West Germans who essentially grew up in the unified Germany, implying that the two parts of Germany are 'growing together'.

Chapter 2

Hidden Costs of Control: Three Repetitions and an Extension

(with Anthony Ziegelmeyer and Matteo Ploner)

2.1 Introduction

Standard principal-agent theory endorses the view that tighter monitoring invariably leads to profitable increases in agents' work effort and recommends principals to offer jobs characterized by limited effort discretion. However, since the work by Deci (1971), many social scientists reject this view and acknowledge that monitoring is likely to have two opposite effects on the agents' performance: A disciplining effect (monitoring limits shirking and therefore raises work effort) and a crowding-out effect (monitoring undermines intrinsic motivation and therefore reduces work effort). Whether monitoring agents is beneficial to the principal depends on the relative size of the two countervailing effects, and a worthwhile goal for economic research is to identify the conditions under which crowding-out effects are predicted to dominate disciplining effects.

Frey (1993) distinguishes between principal-agent relationships and he reasons that the disciplining effect is likely to dominate when the relationship is abstract, as in competitive markets, while the crowding-out effect is likely to dominate in more personalized relationships. This proposition is supported by the field evidence in Barkema (1995) and by a large body of evidence in psychology which indicates that distant principal-agent relationships are often characterized by a controlling leadership style whereas agents' autonomy is more prominent in close principal-agent relationships (Antonakis and Atwater, 2002; Stanton, 2000 and the references therein). Recent experimental evidence collected by economists provides additional support for Frey's proposition. For example, Dickinson and Villeval (2008) establish that a personal relationship between principal and agent is a major condition for a substantial crowding-out effect to emerge, and Bartling et al. (2012) show that offering discretion to an anonymous agent is not profitable on its own but that the combined strategy of offering discretion, paying high wages, and screening based on past efforts is profitable.

Contrary to this large body of evidence, Falk and Kosfeld (2006)—henceforth F&K—show experimentally that motivation crowding-out effects are substantial enough to significantly undermine the effectiveness of control in abstract relationships and in the absence of complementarities between effort discretion and other elements of the employment strategy. In their main treatments, F&K implement a principal-agent game where the principal either restricts the agent's choices by setting $x_{\min} = \underline{x}$ or leaves the agent's

choices unrestricted by setting $x_{\min} = 0$ and the agent then chooses a transfer $x \in \{x_{\min}, x_{\min} + 1, \dots, 119, 120\}$ resulting in a (monetary) payoff of $2x$ and $120 - x$ experimental currency units (ECUs) for the principal and the agent respectively. The minimum transfer requirement $\underline{x} \in \{5, 10, 20\}$ is a treatment variable which corresponds to the principal's degree of control. Experimental results show that control entails *hidden costs* caused by the existence of agents who choose a lower transfer if controlled than otherwise. Most strikingly, hidden costs outweigh the benefits of control in the low and medium control treatments (where \underline{x} equals 5 and 10 respectively) which implies that principals earn less when they control their agents than when they do not control them. If robust, F&K's experimental results have far-reaching implications for the design of employment contracts. Monitoring strategies would not only have to regulate the margin of employees who exploit their freedom to shirk but they should also sustain the motivation of the large fraction of employees disinclined to shirk.

In this chapter, we first report three *repetitions* of F&K's low and medium control treatments. Our repetitions largely reproduce the conditions of the original study with different subject pools (like F&K each of our experiments employs a standard subject pool of students). The primary purpose of our repetitions is to check the robustness of F&K's experimental results to small changes in subject pools. In the first repetition, subjects are students from the University of Jena (federal state of Thuringia, Germany) and the sample composition according to educational background is (almost) equally distributed among the primary fields of academic study. In the second repetition, subjects are students from the University of Trento (province of Trento, Italy) and business and economics majors are over-represented (about 60% of the sample). In the third repetition, subjects are students from the University of Jena and the sample composition according to educational background is identical to the one of the second repetition. Except in our second repetition, we confirm the existence of statistically significant hidden costs of control. More importantly, and in contrast to F&K's findings, hidden costs almost never statistically significantly outweigh benefits of control in our four repetitions.

Once game play data had been collected, subjects were asked to complete a questionnaire in which they had to state their work motivation in hypothetical scenarios. Our questionnaires are identical to the ones F&K administered in their questionnaire study. In contrast to the game play data, our questionnaire data are similar to F&K's questionnaire data. Like in the original questionnaire study, we observe that in each scenario stated work motivation is statistically significantly lower in the condition where the principal controls or uses explicit incentive devices than in the condition where the principal trusts the agent. Hence, we observe little revealed control aversion but stated work motivation is significantly reduced by the implementation of control in our first three experiments. F&K, on the other hand, observe that control has a similar influence on revealed and stated work motivation. In an attempt to solve this puzzle, we decided to investigate the influence of hypothetical incentives on the play of the principal-agent game. In the second part of this chapter, we report an extension of F&K's low and medium control treatments where performance-contingent earnings are absent i.e. both principals and agents are paid according to a flat participation fee. We observe that hidden costs statistically significantly outweigh benefits of control which suggests that the magnitude of hidden costs of control is larger under hypothetical than under real incentives.

We also investigate whether the magnitude of hidden costs of control varies according to subjects' demographics. This investigation is motivated by the following two observations on agents' behavior. First, the findings of our third repetition (conducted in Jena with 60% of business and economics majors) are more in line with F&K's findings than

the ones from our first repetition (conducted in Jena and majors balanced). This suggests that business and economics majors are more control-averse than other majors. Second, the findings of our third repetition are more in line with F&K's findings than the ones from our second repetition (conducted in Trento with 60% of business and economics majors). This suggests that Italian students in the province of Trento are less control-averse than German students in the state of Thuringia. In our general discussion, we test the reliability of these two suggestions with the help of regressions using experimental data from our second and third repetition as well as those from our extension. In these three experiments we collected information from subjects on basic demographic characteristics (academic major, age and gender) which enables us to compare the degree of control aversion across different sub-samples. Statistical results show that our Italian students with economics training are statistically significantly less control-averse than our German students with economics training but that the degree of control aversion of non-economists does not statistically significantly differ between the two regions. Additionally, none of the comparisons between sub-samples leads to the conclusion that economists are statistically significantly more control-averse than non-economists. Finally, in the absence of performance-contingent earnings, female non-economists are statistically significantly more control-averse than male non-economists. This gender effect is the only statistically significant one out of six comparisons.

The remainder of this chapter is organized as follows. Section 2.2 summarizes F&K's experimental findings in their main treatments. Sections 2.3 to 2.5 discuss our three repetitions. Section 2.6 analyzes our questionnaire data and Section 2.7 presents our extension. We provide a general discussion in Section 2.8.

2.2 Falk and Kosfeld's main treatments

In their main treatments F&K implement a principal-agent game to study the potential interaction between control and effort in a parsimonious way. An agent has an endowment of 120 ECUs and chooses a transfer of x ECUs to the principal. For every ECU that the agent gives up, the principal receives two ECUs. Before the agent decides how much to transfer, the principal decides whether or not to impose a compulsory transfer of \underline{x} ECUs. F&K consider a low ($\underline{x} = 5$), a medium ($\underline{x} = 10$), and a high ($\underline{x} = 20$) control treatment (referred to as C5, C10 and C20 respectively) where principals and agents interact only once and they make use of the strategy method to elicit agents' transfers. Concretely, each agent is asked to choose a pair of transfers (x^C, x^{NC}) where $x^C \in \{\underline{x}, \underline{x} + 1, \dots, 120\}$ is payoff-relevant in case the principal imposes a compulsory transfer and $x^{NC} \in \{0, 1, \dots, 120\}$ is payoff-relevant in case the principal does not impose a compulsory transfer. A total of 140, 144, and 134 subjects participated in the control treatments C5, C10, and C20 respectively.

For a given agent's pair of transfers (x^C, x^{NC}) , the decision to impose a compulsory transfer has a total effect on the principal's payoff which is given by $2(x^C - x^{NC})$. This total effect can be expressed as the sum of the direct and indirect effects of control which are given by $2 \max\{\underline{x} - x^{NC}, 0\}$ and $2(x^C - \max\{\underline{x}, x^{NC}\})$ respectively (see Schnedler and Vadovic, 2011). Hence, there is an indirect effect of control (or psychological reaction to control) if (i) $x^C \neq x^{NC}$ and (ii) $x^{NC} \leq \underline{x}$ implies $x^C > \underline{x}$. The indirect effect of control is positive if either $x^C > \underline{x} > x^{NC}$ or $x^C > x^{NC} \geq \underline{x}$, in which case there are hidden benefits of control. The indirect effect is negative, i.e., there are hidden costs of control, if $x^{NC} > x^C$. For a given distribution of pairs of transfers, statistically significant hidden costs of control are observed whenever negative psychological reactions to control "dominate" positive psychological reactions to control.

F&K's key findings can be summarized as follows: (i) Statistically significant hidden costs of control are observed in all control treatments (C5, C10 and C20); (ii) A clear majority of agents choose a pair of transfers such that $x^C - x^{NC} < 0$ in C5 and C10. Almost identical proportions of agents exhibit positive and negative reactions to control in C20; (iii) The average transfer by agents is higher when agents are not controlled than when they are ($\bar{x}^C - \bar{x}^{NC} < 0$). These differences are statistically significant in C5 and C10 but not in C20. In other words, in C5 and C10, principals earn less when they control their agents than when they do not control them, which implies that the hidden costs of control outweigh the benefits; and (iv) A clear majority of principals anticipates the adverse effect of control and choose to leave transfers unrestricted in C5 and C10. About half of the principals choose to leave transfers unrestricted in C20.¹

¹F&K ran three robustness treatments for which treatment C10 served as a basis of comparison. First, they checked whether their results were an artifact of the strategy method by applying the specific response method in treatment SR10 (agents decided only after knowing the principal's decision). Their results do not indicate any effect of using the strategy method versus the specific response method in the considered principal-agent relationship. Second, they established that control only hurts the agent's motivation when the principal has a choice whether to impose control or not. Indeed, in treatment EX10 which is the subgame of treatment C10 following an exogenously imposed control decision, meaning that the agent always has to choose a transfer of x ECUs in the set $\{10, 11, \dots, 120\}$, the negative effect of control vanishes. Third, F&K explored the validity of their results in a more general economic setting by implementing a gift-exchange treatment (GE10) where the principal not only determines the agent's minimum level of x , but also pays the agent a wage. They found that the reciprocal relationship between wages and effort is weaker in the presence of control in treatment GE10.

As F&K explain in their paper, profit-maximizing principals will force agents to transfer at least \underline{x} ECUs if all agents are profit-maximizers, and also if some agents have a concern for equity and/or efficiency. Even some applications based on psychological games rule out the possibility that, in equilibrium, the principal leaves transfers unrestricted and the agent responds positively to the principal's signal of kindness (Rabin, 1993; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006). Intention-based reciprocity might explain the crowding-out effect if the propensity for reciprocity varies in the population *and* propensities are private information (von Siemens, 2011). Signaling motives also have the potential to rationalize agents' negative reactions to control (Sliwka, 2007; Ellingsen and Johannesson, 2008).