

The Lure of Authority: Motivation and Incentive Effects of Power

Ernst Fehr and Holger Herz and Tom Wilkening*

November 22, 2012

Abstract

Authority and power permeate political, social, and economic life, but empirical knowledge about the motivational origins and consequences of authority is limited. We study the motivation and incentive effects of authority experimentally in an authority-delegation game. Individuals often retain authority even when its delegation is in their material interest — suggesting that authority has non-pecuniary consequences for utility. Authority also leads to over-provision of effort by the controlling parties, while a large percentage of subordinates under-provide effort despite pecuniary incentives to the contrary. Authority thus has important motivational consequences that exacerbate the inefficiencies arising from suboptimal delegation choices.

Keywords: Organizational Behavior, Incentives, Experiments and Contracts.

JEL Classification Codes: C92, D83, D23

*Ernst Fehr: Department of Economics, Zurich University, Bluemlisalpstrasse 10, CH-8006 Zurich. E-mail: ernst.fehr@econ.uzh.ch. Holger Herz: Department of Economics, Zurich University, Bluemlisalpstrasse 10, CH-8006 Zurich. E-mail: holger.herz@econ.uzh.ch. Tom Wilkening: Department of Economics, The University of Melbourne, Economics and Commerce Building, Victoria, 3010, Australia. E-mail: Tom.Wilkening@unimelb.edu.au. We are grateful to Sam Bowles, Wouter Dessen, Jon Elster, Bob Gibbons, Sergei Izmalkov, Bentley MacLeod, Nikos Nikiforakis, Klaus Schmidt, Chris Smith, seminar participants at the European University Institute in Florence, the London School of Economics, Bocconi, Harvard, MIT, NYU, Zurich, participants at the ESA International Meetings 2008 at Caltech, Participants at the 24th annual congress of the European Economic Association, the 2nd Australian Workshop for Experimental Economics, the 2nd Behavioral Economics Conference at Cornell and three anonymous referees for helpful suggestions and feedback. We are particularly grateful to a fourth referee whose careful reading of our work greatly improved the paper. We acknowledge the generous support of the Swiss National Science Foundation (Grant No. 100014_130127/1 on “The Social Dynamics of Normative Behavior”). The project was conducted while Tom Wilkening was at MIT and received human subjects approval from MIT’s institutional review board (COUHES No 0703002170). The project also received the approval of the IRB at the faculty of economics, business administration, and computer science of the University of Zurich.

Authority and power play an important role in human societies. Influential scholars from various social science disciplines — such as Marx (1867), Russell (1938), Parsons (1963), Dahl (1957), and Weber (1978) — have contributed to our understanding of the origins, characteristics, and potential consequences of these forces.

Despite some notable early exceptions (Simon (1951); Zeuthen (1968); Harsanyi (1978); Bowles & Gintis (1988)), the study of authority and power has not been a major focus in economics. More recently, however, organizational economists have taken interest in the incentive effects of decision rights by studying situations where one party has the contractual right to make decisions that influence another party’s payoffs and potential choices (Grossman & Hart (1986); Hart & Moore (1990); Aghion & Tirole (1997); Baker, Gibbons & Murphy (1999); Dessein (2002); Aghion, Dewatripont & Rey (2004)). The granting of decision rights can mitigate inefficiencies by shielding the controlling party from potential holdup and expropriation.

There is, however, very little *empirical* work in economics that examines the behavioral consequences of authority and power or their motivational origins. This paper explores these forces using a laboratory experiment where we study how individuals manage and respond to authority in a hierarchical relationship. We propose a new “authority-delegation game” based on a model developed in Aghion & Tirole (1997). A principal and an agent must select one of a large number of potential projects for implementation. One party, initially the principal, has the right to decide which project to implement. The other party, initially the agent, can only make a project recommendation but lacks direct power to determine the project. We follow Aghion & Tirole (1997) by defining authority as the right to determine the project.

Payoffs to the principal and agent for implementing a project are unknown *ex ante*, and both parties can provide effort which directly controls the probability with which they will be informed about the value of each project. One of the projects is best for the principal, while a different project is best for the agent. After the parties have exerted effort, four states are possible: both parties are informed, only the principal is informed, only the agent is informed, or neither party is informed. Before the parties provide effort, the principal can delegate authority to the agent and become the subordinate party. Delegation of authority means that the agent becomes the controlling party and has the right to choose the project.

Both the controlling party (i.e. the party with the decision right) and the subordinate party have pecuniary incentives to exert effort in this setup because both parties earn only a low outside option in the absence of any information. The subordinate’s incentives are lower, however, because if both parties are informed, the controlling party may overrule the subordinate’s proposal and choose the project which is best for him or her. Delegation

therefore increases the agent's effort because he can now implement his preferred project in cases where he is informed. However, delegation also reduces the principal's control over project choice. When the principal's return from the agent's preferred project is high, the cost of losing authority is small. A rational principal who maximizes her expected payoff should thus delegate authority in this case. When a principal's return from the agent's preferred project is low, however, the cost of losing authority is high and a rational principal who maximizes expected payoff should retain authority.

Our first main result is that the principals show a proclivity for retaining authority in situations in which they could improve their expected income by delegating it, i.e. in situations in which their return from the agent's preferred project is relatively high. However, the principals only delegate in roughly 40 percent of these cases. Pessimistic expectations about the agent's effort in case of delegation cannot explain this reluctance. On the contrary, the principals have quite reasonable beliefs about the agent's effort, meaning that it would be profitable to delegate in the clear majority of cases based on these beliefs. Nevertheless, principals prefer retaining authority.

These findings suggest that the principals might view authority not just as an instrument that helps them increase their earnings, but that the allocation of decision rights has non-pecuniary consequences that inhibit the delegation of authority.¹ In our experiment, the fact that the principals are willing to sacrifice some of their earnings to keep authority suggests a preference for the decision right.

Why are the principals willing to forgo money in order to keep the decision right? Our empirical data indicate that a disutility for being overruled appears to be an important driver behind their reluctance to delegate. A principal is overruled if (i) she delegates authority and (ii) both the principal and the agent are informed about project values so that (iii) the agent can disregard the principal's information and choose his preferred project. Note that the principal's pecuniary payoff from the agent's preferred project is the same regardless of whether the principal is informed about the project values or not. Conditional on effort, an expected utility maximizing principal who is the subordinate of an informed agent should thus be indifferent between the case where she is informed and overruled by the agent and the case where she remains uninformed. It follows that her behavior after being informed, overruled, and receiving the payoff from the agent's preferred project should be identical to her behavior after being uninformed and receiving the same payoff.

If, however, a principal experiences a non-pecuniary disutility from being overruled, her

¹Psychologists have postulated a preference for power (Mulder (1975)) or a preference for agency, autonomy and self determination (Rotter (1966); deCharms (1968); Deci (1981)). To our knowledge, however, no evidence yet exists that demonstrates a willingness to pay (i.e. a preference) for power, agency, or autonomy.

behavior after these two outcomes may differ: the principal may be less willing to delegate in the next period if overruled. This is exactly what we find in our data. Principals who are overruled are significantly less likely to delegate in the next period relative to those who are uninformed, even if it is in their pecuniary interest to delegate. Moreover, we observe significantly higher delegation rates in a control treatment where delegation is profitable and the principal cannot be overruled after delegation (because she always remains uninformed) relative to a treatment in which the profitability of delegation is larger but the principal can be overruled. Thus, a disutility from being overruled appears to be an important non-pecuniary factor behind the reluctance to delegate.

Our second main finding is that the controlling party substantially over-provides effort relative to the Nash equilibrium and relative to his or her best response to the subordinate party's anticipated effort. This over-provision is persistent, with no convergence to the Nash Equilibrium over time. Being in the position of the controlling party thus generates additional motivation for effort provision.

Our third main finding is that many subordinates substantially under-provide effort relative to the Nash equilibrium. In fact, a substantial minority of the subordinate parties (between 30 and 50 percent across various treatments) chooses a zero effort level even though — due to the very small cost of low positive effort levels — zero effort is almost never an optimal choice. This result suggests that the lack of authority has a demotivating effect on a substantial minority of the subordinate parties.

It turns out that many aspects of our data can be captured by the notion that subjects want to avoid *ex-post* regret about their choices — a hypothesis that was introduced by Loomes & Sugden (1982). Regret aversion is a form of reference dependent utility and is based on the idea that subjects derive disutility from regret that arises by comparing their actual *ex-post* outcomes with those the subject could have had by choosing a different action.

A distaste for *ex-post* regret can explain our under-delegation result in the following manner:² the principals exert strictly less effort throughout the experiment as a subordinate (i.e. after delegation) than they did as the controlling party. Thus, a principal who delegates and is informed in the role of a subordinate would also have been informed as the controlling party.³ As being overruled leads to a project with a lower value to the principal, a principal

²For further details see appendix A where we apply a formal model of regret aversion to our experiment.

³Effort in the experiment is identical to the probability of becoming informed about the value of each project. Each party's effort is a number in the set $\{0, 5, \dots, 100\}$; this effort is compared to a random number equally distributed between 1 and 100. If the random number is below the effort of the party, the party was informed about the value of all projects; if the random number is above the effort of the party, the party remained uninformed. Thus, if the principal's effort as a controlling party is higher than her effort as a subordinate party, then being informed as a subordinate party implies that the principal also would have been informed as a controlling party.

who delegates and is informed and overruled may regret her delegation choice *ex post*, because if she had kept her decision right she could have chosen her preferred project. Anticipating such regret, an individual may maintain control in order to mitigate the potential for regret. Moreover, a principal who has delegated can minimize feelings of regret by minimizing the probability of being informed, i.e., by choosing a zero effort. Thus, the desire to avoid regret can explain both the propensity for under-delegation and the low effort choices of the subordinates.⁴ Finally, regret may analogously generate the over-provision of effort by the controlling parties if they experience regret when they remain uninformed because the right to choose the project is of little value in this case. We discuss the evidence in favor of regret aversion in more detail in section 3.2.2, section 3.5, and appendix A.

Our paper is related to the experimental literature on the consequences of delegation on punishment choices (Bartling & Fischbacher (2012), Coffman (2011)). While these papers study the assignment of punishment in response to the allocation choices of either a principal or a delegate, our paper studies the willingness of the principal to delegate and the willingness of a principal and an agent to invest effort in response to the assignment of decision rights. Our results on effort provision is also related to the literature on the hidden costs of control (Fehr & Rockenbach (2003), Fehr & List (2004), Falk & Kosfeld (2006), Charness et al. (2012)). While this literature shows that the exercise of control reduces an individuals' positive reciprocity towards the principal, our paper shows that lack of control has demotivating consequences on subordinates that induce them to act against their material self interest.

We believe that our results have potentially important implications across many domains. In relation to the property rights literature (Grossman & Hart (1986); Hart & Moore (1990)), (re)allocating property rights across firms may be difficult if decision rights have non-pecuniary consequences for utility because organization members with decision rights may oppose their re-allocation, even if they would benefit economically from it. The under-delegation of authority not only reduces the principals' earnings, but also causes the agent to lose money in some of our treatments. Thus the distortion in the allocation of control rights can lead to organizational structures that reduce the value of the organization as a whole. The identification of motivational obstacles to delegation adds an important component to the theoretical work by Baker, Gibbons and Murphy (1999), Sliwka (2001), and Bester & Kräbmer (2008), which predicts limits to delegation in environments with limited commitment, dynamic incentives, or limited liability.

⁴Note that regret can affect not only principals, but also agents in the subordinate role. The agents' efforts might be wasted even if they are informed, since in cases where the principal is also informed, the agent will be overruled. In these cases, agents may regret positive effort levels *ex post*. Therefore, regret averse agents in the subordinate role may reduce effort relative to an agent that maximizes his or her expected earnings.

A reluctance to delegate decision rights may also play a role in corporate finance, in the political sphere, and in the design of optimal institutions for regulating relations between firms. Models of empire-building investment (Jensen (1986), Hart & Moore (1995)), which have been used extensively in the literature to understand the trade-offs between financial instruments may, in part, be founded on non-pecuniary motives to retain authority. In view of incumbents' advantages for re-election (Gelman & King (1990)), these motives also strengthen the case for term limits because politicians may otherwise try to keep their political power positions beyond what is good for the polity. In addition, they may provide a rationale for models in the spirit of Niskanen (1971) which assume that bureaucrats seek to maximize their discretionary budget.

The motivational consequences of authority for effort provision may be equally important. The motivation enhancing effect for the controlling parties and the detrimental effect on the motivation of a large minority of the subordinates suggest that the incentive effects of authority are larger than the standard model predicts: a reallocation of authority may cause a marked increase in effort by the new controlling party and a large reduction in effort by the party previously in control. The noteworthy gap between the controlling and the subordinate parties' efforts also implies that when contracts are incomplete, the efficiency losses due to authority are likely to be higher than the standard model predicts. Furthermore, our finding that a lack of authority only seems to demotivate a minority of people suggests that putting the right people into positions that lack authority is important. The development of tools for detecting this type of employees may thus be important in minimizing the cost associated with the (re)allocation of authority.

Despite the systematic deviations from the predictions of the Aghion & Tirole (1997) model, we believe that their model is very useful for the study of authority because the main comparative static predictions of the model are nicely met and the precise numerical predictions of the model enabled us to detect the motivational forces we described above. The model is thus incomplete in terms of the underlying motivational forces, but the (incomplete) model is remarkably robust in terms of the comparative static predictions. It remains to be seen whether this robustness is a general feature of the broader organizational economics literature where communication (Dessein (2002), Rantakari (2008)), monetary incentives (Athey & Roberts (2001)), and dynamic learning (Aghion, Dewatripont & Rey (2004)) are possible. However, even if the robustness of the comparative static predictions of the Aghion and Tirole model extends to the broader organizational economics literature, we believe that this literature should take the behavioral forces observed in our paper into account because — as we show here — they may have important consequences.

The remainder of the paper is structured as follows. We present a simplified version of

the model of Aghion & Tirole (1997) in section 1 and derive its theoretical predictions. Section 2 details our experimental design and hypotheses. Section 3 reports the main results of our experiment and is separated into three parts. Section 3.1 summarizes the data and provides an overview of the major results. Section 3.2 explores possible reasons why principals might choose to keep control rights. The third part, consisting of sections 3.3-3.5, examines the reasons for the controlling parties' over-provision of effort and why subordinate parties might want to under-provide effort relative to the risk neutral Nash equilibrium. Section 4 concludes.

1 Theoretical Motivations

The basis of our experimental design is a model of authority developed in Aghion & Tirole (1997). We consider a world in which a principal (she) and an agent (he) are organized in a hierarchical structure and must decide to implement one or zero projects out of a set of $n \geq 3$ potential projects. With each project $k \in \{1, \dots, n\}$, there is an associated non-contractible gain of P_k for the principal and a private benefit A_k for the agent. If no project is implemented, the profit and private value are both equal to a known outside value of P_0 and A_0 respectively.

For ease of exposition, we index the principal's preferred project by 1 and the agent's preferred project by 2. The principal's preferred project yields known profit P_1 to the principal and A_1 to the agent where $P_1 > P_0$ and $A_1 > A_0$. Likewise, the agent's preferred project yields known benefit P_2 to the principal and A_2 to the agent with $A_2 > A_0$ and $P_2 > P_0$. As their name suggests, the principal's preferred project yields a strictly higher value to the principal than the agent's preferred project ($P_1 > P_2$). Likewise, the agent's preferred project yields strictly higher value to the agent than the principal's preferred project ($A_2 > A_1$).

While the potential values of projects are known, all projects look identical *ex ante* and information must be collected in order to differentiate between them. The principal and agent acquire information in a binary form. At private cost $g_A(e)$, the agent learns his payoffs to all candidate projects with probability e . With probability $1 - e$, the agent learns nothing and cannot differentiate between the projects. Similarly, at private cost $g_P(E)$, the principal becomes perfectly informed about the payoffs of all projects with probability E and learns nothing with probability $1 - E$. Effort choices are made simultaneously and privately. We concentrate on the case where $g_A(e)$ and $g_P(E)$ are quadratic, $g'_A(0) = g'_P(0) = 0$, $P_1 - g'_P(1) < 0$, and $A_2 - g'_A(1) < 0$. These assumptions ensure that the reaction functions

are linear and that a unique interior solution exists for both authority allocations.⁵

We consider a four stage game which relates decision rights, incentive conflict, and effort. In the first stage, the principal decides whether to keep decision rights or to delegate them to the agent. In the second stage, both parties privately and simultaneously gather information about the n projects' payoffs. In the third stage, the subordinate recommends a project to the controlling party. Finally, the controlling party implements a project or the outside option on the basis of his information and the information communicated by the subordinate.

We assume that the principal and agent are risk neutral. For a given effort level and implemented project k , the principal's utility is $P_k - g_P(E)$. The agent's utility is $A_k - g_A(e)$. As outcomes and effort choices are non-contractible, performance or outcome-contingent payments are ruled out and the introduction of wages is necessary only to satisfy the agent's participation constraint, which, to avoid further notation, we assume to be satisfied.

Information in the model is *soft* so that information passed between parties cannot be verified. As such, if one party is informed and the other party is uninformed, the informed party can limit the amount of information given to the other party. As there is always an incentive conflict between the parties and outcomes are non-contractible, there is always an incentive to restrict information to the preferred project of the informed individual. It follows that communication between parties is reduced to a recommendation for a single project choice.

1.1 Analysis and Theoretical Implications

We denote the party that has authority as the controlling party while the party without authority is called the subordinate. For each party, the expected value for selecting a project at random is less than their respective outside option. Thus, under the assumption of risk neutrality or risk aversion, the subordinate prefers to recommend the outside option rather than a random project. Similarly, an uninformed controlling party never chooses unilaterally to undertake a project other than the outside option.

Given that $A_2 > A_1 > A_0$, $P_1 > P_2 > P_0$, and information is soft, the subordinate always has an incentive to recommend his or her preferred project to the controlling party. The controlling party has an incentive to follow this recommendation if uninformed and to overrule the project and implement his or her preferred project if informed. It follows that

⁵In the experiment, we also use a discrete effort space to reduce cognitive burden. All predictions in section 2.2 are relative to the restricted effort space.

if the principal keeps control, the utilities of a risk-neutral principal and agent are

$$EV_P = E\hat{P}_1 + (1 - E)e\hat{P}_2 + P_0 - g_P(E), \quad (1)$$

$$EV_A = E\hat{A}_1 + (1 - E)e\hat{A}_2 + A_0 - g_A(e), \quad (2)$$

where

$$\hat{P}_i = P_i - P_0, \text{ for } i \in \{1, 2\}, \quad (3)$$

$$\hat{A}_i = A_i - A_0, \text{ for } i \in \{1, 2\}. \quad (4)$$

If the agent receives control, the utility of the principal and agent are

$$EV_P^d = (1 - e)E\hat{P}_1 + e\hat{P}_2 + P_0 - g_P(E), \quad (5)$$

$$EV_A^d = (1 - e)E\hat{A}_1 + e\hat{A}_2 + A_0 - g_A(e), \quad (6)$$

where the superscript d denotes the expected payoffs in the delegation case.

From Equations 1 and 2, the reaction functions if the principal keeps control are the solutions to the following first order conditions:

$$\hat{P}_1 - e\hat{P}_2 = g'_P(E), \quad (7)$$

$$(1 - E)\hat{A}_2 = g'_A(e). \quad (8)$$

Equation 7 describes the principal's reaction function which we denote by $r_P(e)$. Equation 8 describes the agent's reaction function denoted by $r_A(E)$. Note that both $r_P(e)$ and $r_A(E)$ are downward sloping in (E, e) -space, implying that the principal's and agent's effort level are strategic substitutes. Thus, an increase in the agent's effort induces the principal to reduce her effort and vice versa. By the additional assumptions placed on $g_P(E)$ and $g_A(e)$ above, the reaction functions are also linear and there exists an interior intersection of reaction functions, (e^{NE}, E^{NE}) , which constitutes the Nash equilibrium of this subgame.

If the agent receives control, the reaction curves of the principal and agent are the solutions to:

$$(1 - e)\hat{P}_1 = g'_P(E), \quad (9)$$

$$\hat{A}_2 - E\hat{A}_1 = g'_A(e), \quad (10)$$

and denoted by $r_P^d(e^d)$ and $r_A^d(E^d)$. As in the case when the principal keeps control, the reaction functions are downward sloping in (E^d, e^d) space. Our uniqueness criteria assumed above

again ensure the existence of an interior intersection of reaction functions, $(e^{d^{NE}}, E^{d^{NE}})$, which constitutes the Nash equilibrium of this subgame.

A careful examination of the reaction functions if the principal keeps control and if the agent receives control reveals that the principal decreases her effort when giving up control while the agent increases his effort. Delegation thus has two effects on the principal's payoff: 1) a cost saving effect since delegation reduces the equilibrium effort of the principal and increases the agent's equilibrium effort, and 2) a project selection effect which decreases the probability that the principal's preferred project is undertaken. As these effects are, in general, of opposite sign, the overall incentive for delegation depends on the specifics of the cost function and the degree of interest alignment. In our experiment, we chose cost functions and parameters such that the magnitude of \hat{P}_2 determines whether delegation or retention is optimal for the principal. Full details of the experimental design and its parameterizations are discussed in more detail in the next section.

2 The Experiment

2.1 The Authority Game

At the center of our experimental design is a computerized authority-delegation game with the following features. In each of ten periods, a principal is matched with an agent and shown a set of 36 cards on her computer screen representing potential projects.⁶ One of these cards has a small positive payoff for both players and is placed face up representing the outside option. The remaining thirty five cards are shuffled face down so that the location of each project is unknown. One of these cards is red and represents the principal's preferred project. Following the theory section, we refer to this card as project 1. A second card is blue and represents the agent's preferred project. We refer to this card as project 2. The remaining thirty-three cards are white and result in zero payoff for both parties. The task of each principal-agent pair is to select a card which will be used for payment. The payoffs ensure that individuals prefer to implement the outside option relative to picking a card at random.

Play of the game is done in six stages which are illustrated in Figure 1 and discussed here. Initially principals are given the decision right which corresponds to being able to select a card at the end of the game. In the first stage of the game, each principal is asked whether he wishes to keep this right or to transfer the right to the agent. Giving the right to the

⁶Subjects are randomly assigned the role of a principal or of an agent and remain in this role throughout the experiment. In the instructions, they were referred to as participant A and participant B.

agent is binding and irreversible.

In the second stage, subjects choose their effort levels simultaneously and in private.⁷ Both subjects select their effort in increments of 5 from $\{0, 5, \dots, 95, 100\}$. This effort corresponds to the probability that the subject learns the location of all projects. Effort has an associated cost generated via a quadratic cost function which is constant across treatments and player types:

$$g_P(E) = 25 \left(\frac{E}{100} \right)^2, \quad g_A(e) = 25 \left(\frac{e}{100} \right)^2. \quad (11)$$

Subjects are presented information on the cost of effort in a table where each possible effort and its associated cost is displayed. In all but one session, agents' effort levels are recorded via the strategy method where an effort level is elicited both for the case where principals keep decision rights and the case where these rights are delegated.⁸ Thus, agents choose their effort levels before they know whether the principal has delegated authority to them.

In the third stage, we elicit beliefs of both subjects. Principals and agents are asked their beliefs about the effort of the other party both in the case where decision rights are kept and where they are delegated. For principals this is done in two steps. Beliefs are first elicited for the chosen authority assignment followed by beliefs for the counterfactual. For agents, beliefs for both potential authority assignments are elicited simultaneously. To prevent hedging, no incentives are used in the elicitation of beliefs. In the fourth stage, agents are informed about whether principals kept or transferred decision rights. Then, given a subject's effort for the principal's assignment of authority, a random process determines whether that subject learns the payoffs of all projects or whether he stays uninformed. The effort of the other subject is not revealed nor is information indicating the success or failure of the other subject's effort. All information gained at this stage is private.

In the fifth stage, the subordinate is given the ability to recommend a project to the controlling party. This is accomplished by visibly marking a single project on the computer screen, which can include the outside option. The recommendation is shown to the controlling party, but the payoffs associated with the recommended project are kept hidden in the case where the controlling party remains uninformed.

⁷In the experiment we refer to effort as "search intensity".

⁸We test whether the strategy method influences our results by comparing the results of three sessions of the HIGH treatment where the strategy method was run ($N = 70$) to the session of the HIGH treatment where a standard elicitation method was used ($N = 30$). We find no difference across treatments. P-Values of a Kolmogorov-Smirnov test, which tests whether the distribution of agent effort is identical in treatments with and without strategy method, are 0.79 for effort with decision rights and 0.81 for effort without decision rights. Delegation frequencies differ by 1.6 percent. This difference is also not significant ($p=0.71$ in a Fisher's exact test). The data from the treatments using the strategy method are therefore pooled with the data from treatments using the standard elicitation method in the subsequent analysis.

In the final stage, after seeing the recommendation of the subordinate and the information from his own effort, the controlling party selects a project. Payment for the round is based on the selected project and the costs of effort of each subject.

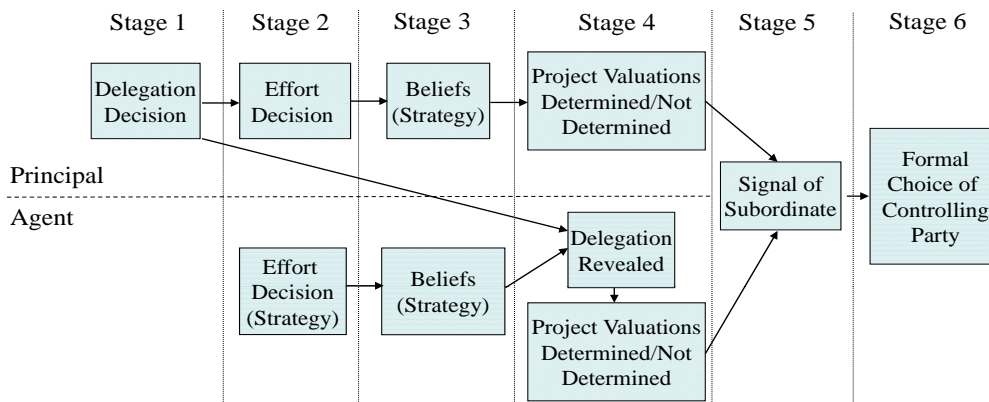


Figure 1: Experimental procedures in the authority game

2.2 Experimental Design and Hypotheses

The experimental design involves four treatments implemented in a between-subjects design. Treatments vary in the amount that principals and agents are paid for the selection of the project preferred by the other party (P_2 and A_1). By changing the payoff given to the other party, the level of incentive conflict in the environment is changed, which, as indicated by the first order conditions in section 1.1, leads to differences in predicted delegation and effort levels.

Table 1 summarizes the value of projects across the four treatments. In each treatment, each party earns 40 points for the selection of their preferred project and a smaller amount for the other party's preferred project. Treatments are divided into two groups — symmetric and asymmetric — where symmetry refers to the relative values of P_2 and A_1 . In the symmetric treatments (LOW and HIGH) the payoffs from the other party's preferred project are the same for the principal and agent. In the low alignment treatment (LOW), the payoffs from the other party's preferred project are small (20) leading to a high degree of incentive conflict. In the high alignment treatment (HIGH), the payoffs from the other party's preferred project are large (35) leading to less incentive conflict. In the asymmetric treatments (PLOW and PHIGH), the payoffs from the other party's preferred project are large for one of the two parties (35) and small for the other (20). As a naming convention, we use PHIGH to denote the case where the principal's value is high under the agent's preferred project. The PLOW treatment is the case where the principal's value is low under the agent's preferred project.

Table 1: Overview of Project Payoffs

	Project 1		Project 2		Outside Option	Other Projects
	Principal	Agent	Principal	Agent		
PLOW	40	35	20	40	10	0
LOW	40	20	20	40	10	0
HIGH	40	35	35	40	10	0
PHIGH	40	20	35	40	10	0

Table 2 shows the predicted Nash equilibrium effort levels and expected profits for each treatment under the case where authority is kept and transferred. As in the model developed in Section 1, E represents the effort level of the principal while e represents the effort level of the agent. As can be seen in table 2, the LOW treatment has a high degree of incentive conflict and authority should be kept by the principal, because the principals expected profit if she keeps control, EV_P , is 20.1 while the expected payoff if she delegates control, EV_P^d , is only 17.3. In the HIGH treatment, incentive conflict is reduced and the principal should delegate authority ($EV_P = 23.3$ vs. $EV_P^d = 24.0$)⁹.

Table 2: Predicted effort levels and expected profits

	Principal has control				Agent has control				Delegation?
	E^{NE}	e^{NE}	EV_P	EV_A	E^{dNE}	e^{dNE}	EV_P^d	EV_A^d	
PLOW	55	25	20.1	25.6	35	45	17.2	23.3	No
LOW	55	25	20.1	17.3	25	55	17.3	20.1	No
HIGH	45	35	23.3	24.0	35	45	24.0	23.3	Yes
PHIGH	45	35	23.3	17.2	25	55	25.6	20.1	Yes

E^{NE} and E^{dNE} denote Nash equilibrium predictions for the principal depending on the control allocation. e^{NE} and e^{dNE} denote Nash equilibrium predictions for the agent depending on the control allocation. EV_P , and EV_P^d denote expected equilibrium profits for the principal depending on the control allocation. EV_A and EV_A^d denote expected equilibrium profits for the agent depending on the control allocation.

In the asymmetric treatments the rewards to delegation are either exacerbated or further diminished relative to the symmetric treatments. Of the four treatments, principals are predicted to have the highest expected value from delegation in the PHIGH treatment ($EV_P^d = 25.6$) and the lowest expected value from delegation in the PLOW treatment ($EV_P^d = 17.2$).

In addition to the delegation predictions, the different interest alignments also lead to different predictions with regard to equilibrium effort levels. All point predictions are given

⁹As will be shown below, while the difference in *theoretical* expected value is small, the empirical difference turns out to be large.

in Table 2. Note that the delegation decisions predicted by the Nash equilibrium are always in the set of welfare maximizing delegation choices. In the PLOW treatment, aggregate expected earnings $EV_p + EV_A$ are highest if the principal keeps authority, while in the PHIGH treatment aggregate expected earnings are highest if the principal delegates authority. In the symmetric treatments, LOW and HIGH, the delegation decision has no effect on the overall welfare if subjects choose Nash equilibrium effort levels.

In the experiment described above, the delegation decision of the principal and the joint effort decisions of the principal and agent generate one of many potential compound lotteries. If an individual's preferences are reference dependent (as, e.g., stipulated in regret theory), the individual's preferred action profile may depart from the equilibrium action profile which assumes players are expected value maximizers. In order to control for such heterogeneity in preferences, we ran a lottery task and used choices from this lottery task as a proxy for the degree to which a subject's preferences exhibit reference dependence. In the lottery task, each subject is presented with the opportunity to participate in six different lotteries, each having the following form:

Win CHF 6 with probability $\frac{1}{2}$, lose CHF X with probability $\frac{1}{2}$. If subjects reject the lottery they receive CHF 0.

The six lotteries varied in the amount X , that could be lost, where X took on the values $X \in \{2, 3, 4, 5, 6, 7\}$. One of the six gambles was randomly selected and paid. As these lotteries are binary, any reference dependent utility function with a reference point between the lowest and the highest outcome can lead to a rejection of gambles with $X \leq 6$. In particular, individuals who are regret averse and compare their outcome to the action which is optimal *ex post* will reject actuarially fair gambles. Thus, the amount X at which a subject starts rejecting the lottery can therefore be taken as an indicator of the degree to which a subject's preference exhibits reference dependence, such as, e.g., regret aversion. For example, a subject that rejects all lotteries with a potential loss of $X > 3$ is classified as exhibiting more regret aversion than a subject that only rejects all lotteries with a potential loss of $X > 5$.¹⁰

In principle, the rejection of actuarially fair gambles in the lottery choice task may also reflect a subject's loss or risk aversion. However, in Section 3.5 and the appendices D and E we show that loss and risk aversion have little explanatory power with regard to effort

¹⁰143 out of 150 subjects who participated in the lottery task and played the authority-delegation game in one of our main treatments have a unique switching point. We use the accepted gamble with the largest potential loss as the independent variable when using the lottery task and do not exclude subjects in the analysis. Excluding subjects with multiple switching points does not significantly alter any of our results.

choices while regret aversion can rationalize both the controlling parties' over-provision of effort and the preference for extremely low (i.e. zero) effort levels among the subordinates.¹¹

Thus, if regret aversion is a motive behind subjects' rejection of lotteries in the lottery task then the propensity to reject lotteries should also be a predictor of subjects' effort level as a controlling party *and* the tendency to provide zero effort as a subordinate party. Likewise, if regret aversion is a motive behind both the propensity to reject lotteries and the reluctance to delegate then we should observe a correlation between the two phenomena. The lottery task may thus provide further evidence regarding the motivational forces behind effort and delegation choices.

2.3 Procedures

Typically, between 20 and 30 subjects participated in each experimental session which consisted of three parts.¹² In part one, subjects played 7 periods of a single player version of the authority game. This single player game is identical to the authority game except that there is no second party. Subjects choose an effort and receive information probabilistically based on their effort. Each individual must then select a project based solely on his own information. The selected project does not affect the payoff of a second party nor does a second party recommend a project. This single player variant gives subjects a chance to get familiar with the effort cost schedule and the computer program.

In part two, the subjects are divided into matching groups of 10 subjects consisting of 5 principals and 5 agents. Subjects play 10 periods of the main authority game in one of the four treatments. Subjects are informed that in a new period they would be matched with another randomly chosen partner.

In part three, subjects are asked to take a short questionnaire in which demographic information is recorded. Instructions for the experiment include a control quiz and a verbal summary of the authority game.

Our subject pool consisted primarily of students at Zurich University and the Federal Institute of Technology in Zurich.¹³ The first series of experiments took place in May and June 2007 with a second series of experiments conducted in May and October 2008. Further

¹¹Moreover, it has been shown by Rabin (2000) that the rejection of lotteries for $X \leq 6$ cannot be reconciled with the assumption that utility is a (reference-independent) strictly concave function of total wealth. Risk aversion based on concave utility of wealth at such low stake levels would imply unreasonable levels of risk aversion at higher stakes, which makes risk averse behavior in this task incompatible with expected utility theory.

¹²More details on individual sessions is provided in appendix B.

¹³Subjects were drawn from a database of volunteers using ORSEE (Greiner (2004)).

control experiments were conducted in May 2009 and April 2011.¹⁴ In total, 504 subjects participated in the experiment, divided into 17 sessions. Experiments were computerized using the software z-tree (Fischbacher (2007)). Payment was given for each period of the main authority game, for the last five periods of the single player game, and for one randomly chosen gamble from the lottery task. On average, an experimental session of the main treatments lasted 1 hour with an average payment of 33.5 CHF (\$35.00).¹⁵

3 Experimental Results

3.1 The Main Facts

Our experimental design generates predictions with regard to delegation, effort and project choices. With regard to project choices and project recommendations the theory does very well:

Result 1 *Controlling parties who are informed about the project valuations almost always choose their preferred projects, implying that they overrule the subordinates' recommendations. Informed subordinates almost always recommend their preferred project and uninformed controlling parties almost always implement this recommendation.*

Result 1 is supported by the following numbers. Principals (Agents) in the role of controlling parties who were informed implemented their preferred project in 100 percent (97.3 percent) of the cases. Principals (agents) in the role of the informed subordinate party recommended their preferred project in 92.6 percent (92.5 percent) of the cases. Finally, principals (agents) in the role of the uninformed controlling party followed the subordinate party's recommendation in 94.1 percent (96.5 percent) of the cases. If the subordinate parties were not informed they typically recommended the outside option (principals: 95.3 percent; agents: 97.0 percent)

Result 1 indicates that the controlling parties used the decision right in their favor. As predicted by theory, this generates a disincentive for subordinates' effort provision but it also makes it reasonable for the principals to delegate authority if their payoff loss at the agents' preferred project is low. Therefore, we next turn to the principals' delegation choices. Recall that in case of Nash equilibrium effort choices by the principal and agent, the principal has an incentive to delegate authority in the HIGH and PHIGH treatments and to keep authority in the LOW and PLOW treatments. Empirically, we find in our experiment:

¹⁴See also appendix B.

¹⁵The 25 and 50 period control treatments took longer. Additional information on these treatments is given in section 3.2.

Result 2 (a) When the principals' interests are misaligned with the agent (*LOW* and *PLOW*) such that the principals are predicted to keep authority, delegation decisions are close to the equilibrium predictions. (b) When the principals' interests are strongly aligned with the agent (*HIGH* and *PHIGH* treatments) such that principals should delegate, we observe strong under-delegation of authority relative to the equilibrium predictions.

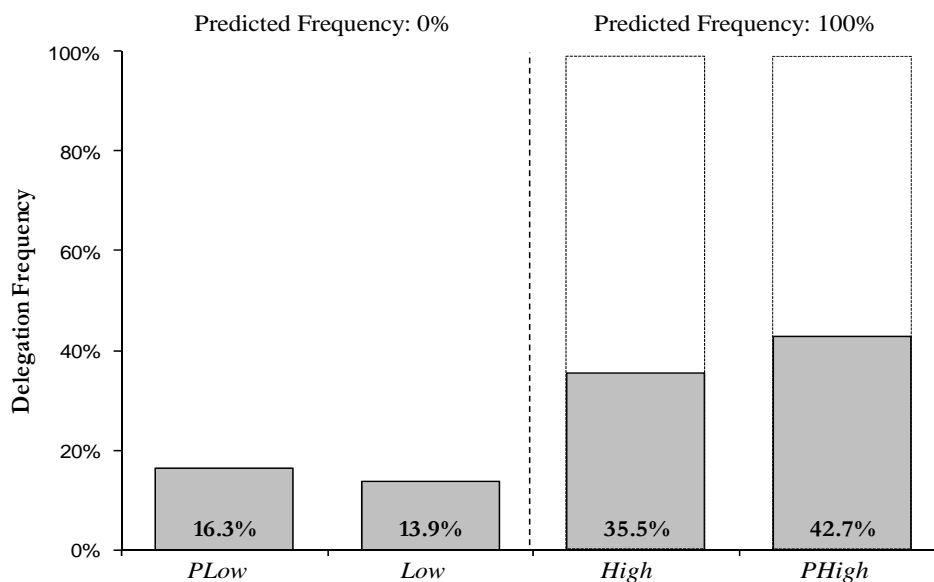


Figure 2: Delegation frequencies by alignment

Figure 2 shows the frequency of delegation for each treatment graphically. As can be seen on the left hand side of the figure, delegation rates in the *PLOW* and *LOW* treatments are 16.3 percent and 13.9 percent. While these levels are above the predicted level of zero, deviations from the prediction appear to be due to infrequent experimentation rather than heterogeneity in delegation strategies. There is little persistence in the strategy of delegation, with 67.4 percent of individuals who delegated authority in one period switching to keeping authority in the next. The frequency of delegation for most individuals is also low, with 39.4 percent of individuals choosing to never delegate and 89.4 percent of individuals delegating in three periods or less.

Average delegation rates in the *HIGH* and *PHIGH* treatment are 35.5 percent and 42.7 percent, far below the predicted rate of 100 percent. These low delegation rates are also rather stable over time. In the *HIGH* treatment the overall delegation rate is 33.5 percent in the first five periods and 37.5 percent in periods 6-10. In the *PHIGH* treatment the overall delegation rate is 36.7 percent in the first five periods and stabilizes around 48.7 percent

from period 6 onwards.¹⁶

In contrast to the LOW and PLOW treatments, the under-delegation of authority in the HIGH and PHIGH treatments appears to be due to heterogeneity in delegation strategies across individuals. Less than 20 percent of individuals delegate eight or more times in the experiment, and individuals who delegate in one period are more likely to delegate in the next period suggesting some persistence in the delegation strategy. However, even in the PHIGH treatment in which delegation incentives are highest according to the Nash prediction, 40 percent of individuals have a delegation frequency of zero suggesting that under-delegation is rather pervasive.

One possible reason for the observed under-delegation might be that actual effort provision if the principal keeps control compared to the case in which the agent receives control makes it more profitable to keep authority. Table 3, which shows the realized profits of principals who kept and delegated authority, shows that this is not the case. In the HIGH and PHIGH treatments, realized profits for the principal are lower than predicted if she keeps control and higher than predicted if the agent receives control. Principals who delegate have on average 30.4 percent greater earnings in the HIGH treatment and 44.5 percent greater earnings in the PHIGH treatment.

Table 3: Realized profits and predicted equilibrium profits for principals

	Principal has control			Observations	Agent has control			Observations
	Actual		Predicted		Actual		Predicted	
PLOW	18.3	*	<i>20.1</i>	251	17.6		<i>17.2</i>	49
LOW	19.0		<i>20.1</i>	310	15.0	**	<i>17.3</i>	50
HIGH	19.1	***	<i>23.3</i>	316	24.9		<i>24.0</i>	174
PHIGH	18.4	***	<i>23.3</i>	172	26.6		<i>25.6</i>	128

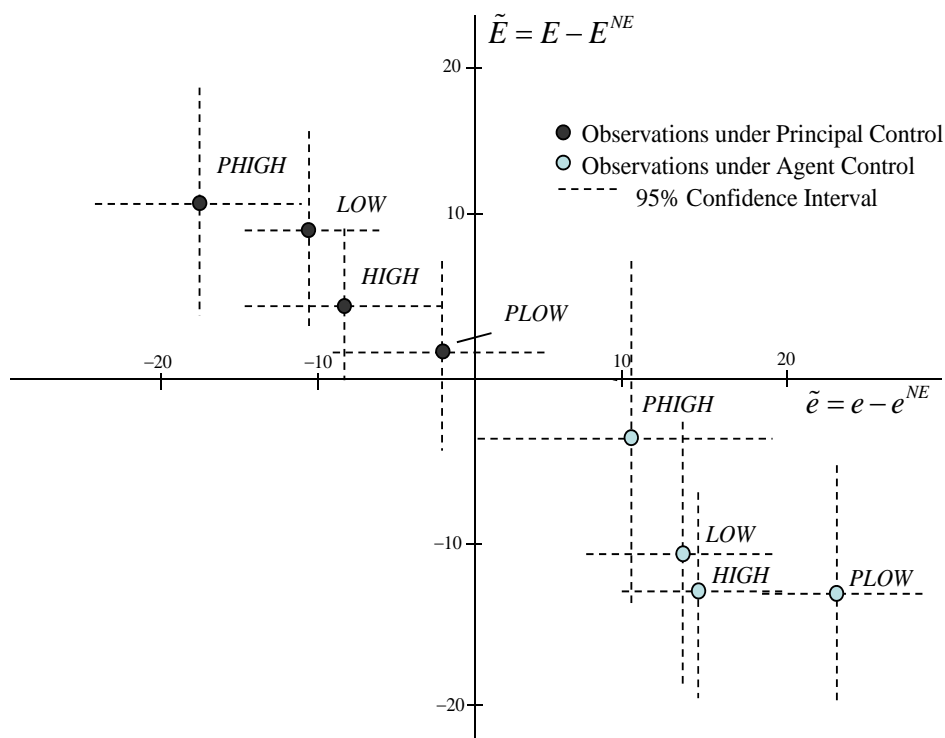
Significance levels calculated by regressing earnings on a constant and testing whether the constant is equal to the prediction. Errors clustered at the individual level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

The second main hypothesis of the experiment is about effort provision. In theory, an incentive conflict leads the controlling party to put in more effort than would be optimal in the case of contractible effort and causes the subordinate to put in less. Relative to this Nash equilibrium benchmark, we observe:

¹⁶While the difference in delegation rates between the first and second half of the experiment is insignificant for the HIGH treatment, the difference of 12 percentage points in the PHIGH treatment is significant ($p < 0.01$ in a probit regression of delegation on a dummy for periods 6-10). We also ran a 50 period session to test for potential long-run learning effects. While delegation increased over the first 20 periods, under-delegation was still pervasive. See section 3.2.1 for details.

Result 3 *Controlling parties provide an excess of effort relative to the Nash equilibrium. Subordinates under-provide effort relative to the Nash equilibrium.*

Figure 3 plots the average deviation of effort levels from the predicted equilibrium values by the principal and agent with both means and 95 percent confidence intervals calculated from individual average efforts. It can be seen that, when authority is kept, the principal over-provides and the agent under-provides relative to the prediction. This phenomenon is reversed, again in all treatments, when authority is delegated, and these deviations are significant for the majority of treatments.¹⁷ In the low treatment, for example, the principal over-provides effort by roughly 10 units relative to the prediction if she keeps control while the agent under-provides effort by about 10 units. This deviation pattern is reversed when the agent is the controlling party.



95% Confidence Intervals and mean effort calculated at the individual level.

Figure 3: Deviations from equilibrium effort predictions. The vertical axis shows the difference (\tilde{E}) between the principal's observed effort (E) and the Nash Equilibrium effort (E^{NE}). The horizontal axis shows the difference (\tilde{e}) between the agent's observed effort (e) and the Nash Equilibrium effort (e^{NE}).

¹⁷We report results from a non-parametric Wilcoxon Signed-Rank test in Table C.1 of appendix C.

The deviations in effort levels from the equilibrium prediction cause inefficiencies that are reflected in the low actual payoff levels of the principals and the agents relative to the predicted payoff levels. Table 3 shows that the principals earn less than predicted in 5 out of 8 cases. In particular, if control is kept, which occurs most frequently in all treatments, the principals always earn less than predicted. For the agents the income loss relative to the prediction is even more extreme (see Table 4): In all 8 cases they earn on average less than predicted.

Table 4: Realized profits and predicted equilibrium profits for agents

	Principal has control			Observations	Agent has control			Observations
	Actual		Predicted		Actual		Predicted	
PLOW	23.0	***	25.6	251	18.8	**	23.3	49
LOW	16.1	**	17.3	310	17.9		20.1	50
HIGH	21.0	***	24.0	316	20.1	***	23.3	174
PHIGH	15.9	**	17.2	172	18.1	**	20.1	128

Significance levels calculated by regressing earnings on a constant and testing whether the constant is equal to the prediction. Errors clustered at the individual level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

The combined effect of under-delegation and deviations in effort provision has particularly strong pecuniary consequences in the PHIGH treatment. Unlike the other treatments, delegation in this treatment leads both the principal and agent to be better off in expectations relative to held control. Principals who delegate receive 45 percent more income compared to those who hold on to decision rights. Likewise agents who are left as the subordinate receive 13.8 percent less profit than those who are delegated to. Taken together, in PHIGH the welfare loss of keeping authority amounts to 30 percent in terms of expected income.¹⁸ The loss in aggregate payoff due to the deviations from the Nash equilibrium are not restricted to PHIGH, however. Table C.2 in appendix C shows that both the principals and the agents earn less than in the Nash equilibrium in each of the 4 treatments. We summarize these findings in the following result:

Result 4 *In each treatment, the deviation in effort provision and delegation leads to monetary losses for both parties relative to the Nash equilibrium. Monetary losses are most acute in the PHIGH treatment where delegation would lead to higher average earnings for both parties.*

¹⁸The implicit assumption in this calculation is that if the principals who kept control rights were to delegate instead, they would exert effort similar to those who delegated. If the principal were to exert less effort, the overall monetary loss would be slightly smaller. In the extreme case where we assume principals exert zero effort in the counterfactual, the monetary loss would amount to 27 percent.

3.2 Exploring the principals' reluctance to delegate

3.2.1 Pecuniary Motivations

A natural initial hypothesis for the observed under-delegation of authority in the HIGH and PHIGH treatments is that individuals believe that they are monetarily better off retaining authority. To see whether this hypothesis has merit, we consider the following counterfactual: Suppose that a principal who did not delegate would elect to delegate instead. Given her beliefs about the agent's actions if she keeps control and if she delegates control, what would be her gain or loss in expected earnings?

As the effort of the principal was elicited only in the case of her chosen authority allocation a comparison of the principal's expected earnings for the cases of delegation and non-delegation requires assumptions about her effort in the counterfactual authority allocation. As we have the principals' beliefs about the agents' effort from both the delegation case and non-delegation case, a natural approach is to use the principal's best reply effort as a proxy for effort. If, for example, the principal kept authority we can compute the principal's best reply effort for the case in which the principal had delegated authority. Using this effort proxy and the principal's belief about the agent's effort enables us to compute the principal's expected profit for the counterfactual case of delegation.¹⁹

As a comparison value, we next compute the expected profits of the principal for the case of retained authority, taking the principal's actual effort and his beliefs about the agent's effort into account.²⁰ Subtracting the expected profit from retained authority from the expected profit from delegation yields our first measure for the expected gains from delegation.

Figure 4 shows the cumulative density function of the gains from delegating under the assumption that the principal would have played a best reply in case he had delegated. As can be seen in this graph by looking at the mass to the right of the zero line, 68 percent of observations in the HIGH treatment and 92 percent of observations from the PHIGH

¹⁹Under the assumption that the principal best replies to his beliefs the expected earnings for the counterfactual case of delegation is given by:

$$EV_P^d(E^d = r_P^d(\hat{e}^d), \hat{e}^d) = \hat{e}^d \hat{P}_2 + (1 - \hat{e}^d) r_P^d(\hat{e}^d) \hat{P}_1 + P_0 - g_P(r_P^d(\hat{e}^d)),$$

where \hat{e}^d is the principal's belief about the agent's effort under delegation, P_0 is the principal's payout under the outside option, \hat{P}_2 is the principal's payment under the agents preferred project net of P_0 , \hat{P}_1 is the principal's payment under the principal's preferred project net of P_0 , and $r_P^d(\hat{e}^d)$ is the best response function constructed in Equation 9.

²⁰This comparison value is given by

$$EV_P(E, \hat{e}) = E \hat{P}_1 + (1 - E) \hat{e} \hat{P}_2 + P_0 - g_P(E).$$

treatment are from individuals who would have been better off if they had delegated. The retention of authority in the PHIGH treatment is especially noteworthy since both the principal and the agent would be made better off through delegation. Thus, in this treatment, the under-delegation is suboptimal not only from the principals' perspective, but also from the perspective of the organization as a whole.

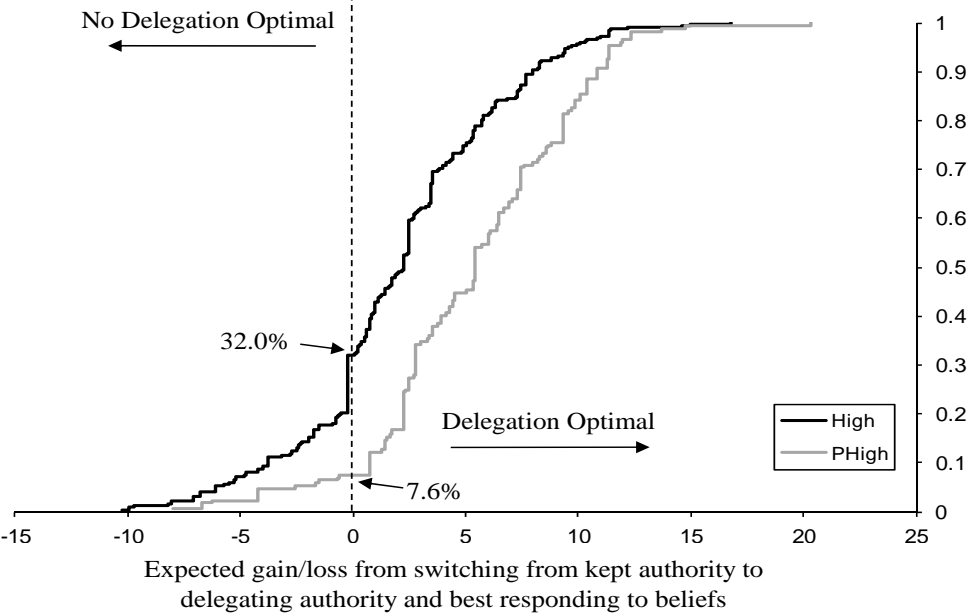


Figure 4: Cumulative density function of expected gain from delegation if principal best replies to beliefs

One might worry that using the best response to beliefs as a proxy for effort might overstate the expected return to delegation. Perhaps, some individuals may not perfectly best respond to their beliefs.

As a conservative secondary measure for the expected gains from delegation, we next consider the case where the principal provides zero effort after delegation. This criterion is selected for three reasons. First, an individual who puts in zero effort has no potential losses and minimal exposure to risk. Relative to the actual strategies typically employed by principals, the zero effort criterion should thus be an attractive strategy for principals who are extremely risk or loss averse. Second, besides very high effort choices which are observed very infrequently, zero effort minimizes the expected value of delegation giving us the lowest reasonable expected value of delegation. Finally, zero effort is in fact the modal strategy taken after delegation suggesting it is a relevant benchmark for analysis.

In Figure 5 we depict the cumulative density function for the expected gains from del-

egation under the assumption that the principal would have chosen zero effort if he had delegated. We find, that 46.8 percent of observations in the HIGH treatment and 75 percent of observations in the PHIGH treatment are from individuals who would have been better off in case of delegation. This result is remarkable because even if we assume that principals choose highly suboptimal effort levels after delegation, it would have often been better for them (given their beliefs) to delegate authority.

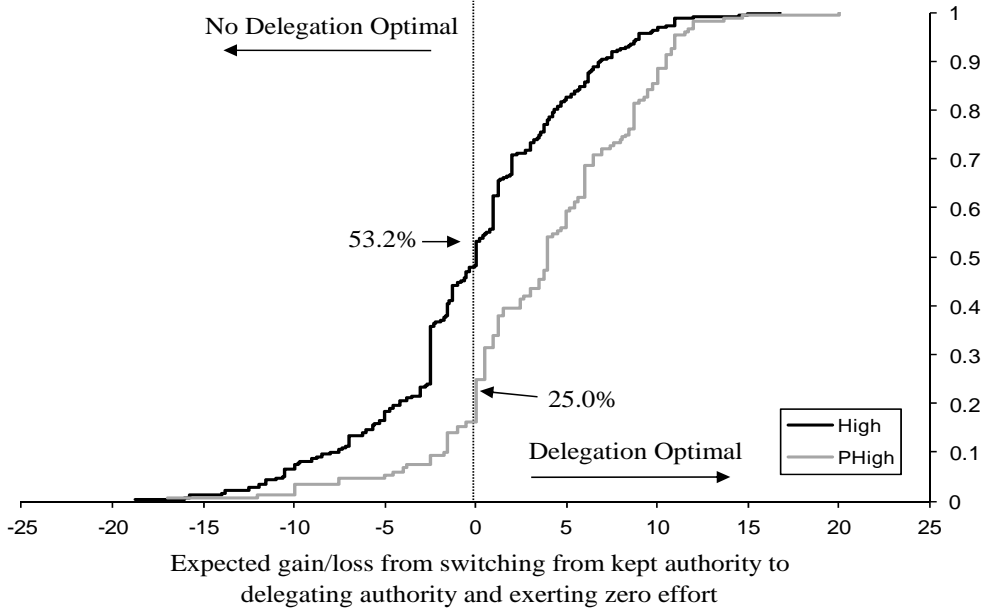


Figure 5: Cumulative density function of expected gain from delegation if principal chooses zero effort after delegation

One might also be concerned that the under-delegation effect is due to having a limited number of periods in which to learn the optimal delegation strategy. Perhaps, under-delegation is a consequence of incomplete learning in our 10 period experiments.

In order to study this hypothesis, we ran an additional treatment (PHIGH50) in which we increased the number of periods to 50 and increased the equilibrium returns to delegation from 10 percent (in the PHIGH treatment) to 17 percent. As with the original PHIGH treatment, we use an asymmetric design in which the agent’s payment for his preferred project is much higher than his payment under the principal’s preferred project, and where both the principal and agent would benefit highly from delegation at the equilibrium effort levels.²¹

²¹In the PHIGH50 treatment, the payment for the principal and agent under the principal’s preferred project were 45 and 20 respectively. Under the agent’s preferred project, the payments were 40 (for the principal) and 45 (for the agent). Equilibrium payoffs were 31 points for the principal and 23.1 points for the agent if the agent had control, and 26.5 points for the principal and 18.1 points for the agent if the principal

To further facilitate learning we also simplified the design by making the subordinate’s recommendation automatic, i.e. the controlling party knew with certainty that an informed subordinate would recommend his preferred project while an uninformed subordinate would recommend the outside option.

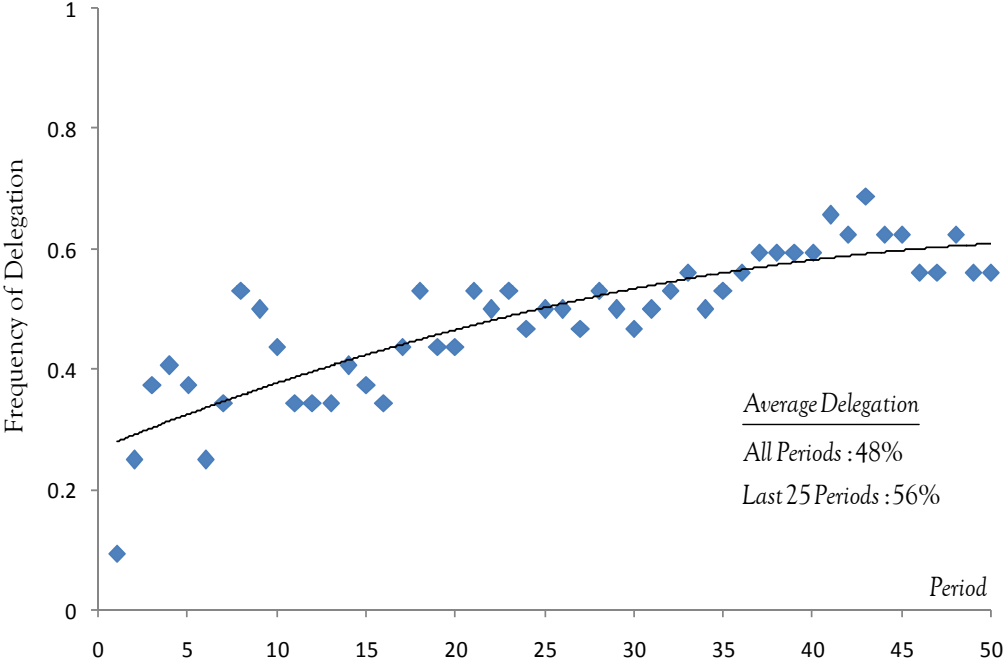


Figure 6: Average delegation frequencies of principals in the PHIGH50 treatment. Period averages shown by the diamonds, quadratic time trend shown by the solid line. In total, 32 principals participated in these 50 period experiments.

Figure 6 shows the time path of delegation decisions in this treatment. As can be seen, while there is an increase in delegation in the first 25 periods, delegation rates remain fairly constant in the remaining 25 periods. Of the 32 principals in the treatment, 43.8 percent never switched their delegation decision after period 25. Subjects are also on average fairly persistent in their delegation choice. On the basis of a median split, we find that those who predominantly keep authority delegate in only 11 percent of the cases in the second half of this treatment. Those who predominantly delegate authority keep it in only 12.5 percent of the cases in the second half of this treatment. Overall, the delegation rate is 8 percent higher in the second half of this treatment compared to the second half of the PHIGH treatment, but this is what one would expect because of the increased incentives to delegate in the PHIGH50 treatment. The overall delegation rates in PHIGH50 are, however, had control. We randomly paid 20 of the 50 periods at the end of the experiment. 64 subjects participated in two sessions of this treatment, and the average session time of this experiment was 2.5 hours with an average payment of 72 CHF (\$75).

still well below the level predicted by theory, with an average delegation rate of 56.5 percent in the last 25 periods of the experiment. As in the PHIGH treatment, this large amount of under-delegation occurred despite the fact that both the principals and the agents were substantially better off in cases where decision rights were delegated relative to cases where decision rights were held.

3.2.2 Non-pecuniary motivations

The results above suggest that many principals had strong pecuniary incentives to delegate. They further suggest that a large portion of under-delegation is not due to incomplete learning. Why then do we observe this strong reluctance to delegate?

One non-pecuniary force behind the principals' choices appears to be a disutility from being overruled. Recall that the principal's return from the agent's preferred project is the same regardless of whether the principal is informed or uninformed. Thus, for a principal who has delegated and faces an agent who selects his preferred project, the pecuniary value of being informed and uninformed is the same. Conditional on effort, an expected utility maximizing principal in the subordinate position should be indifferent between the case where she is informed and overruled by the agent and where she remains uninformed. Thus, her response to being overruled and receiving the payoff from the agent's preferred project should be identical to her response to being uninformed and receiving the payoff of the agent's preferred project. However, if a principal experiences a non-pecuniary disutility from being overruled, her response to these two outcomes may differ: In the next period, the principal may correctly anticipate the potential disutility from being overruled and, therefore, she may be less willing to delegate.

To examine spillovers across periods, we take the principals who delegate in period $t - 1$ and regress the probability of delegating in period t on the information the principal and agent received in the previous period. To account for potential differences in effort and beliefs, we also condition on the effort of the principal in the previous period and the beliefs the principal has about the agent's action in the current period.

Table 5 reports the marginal effects of this regression for various subsets of our data. The omitted category in all regressions is the case in which both the principal and the agent are uninformed. As can be seen in column (1), the principal is more likely to delegate in period t if either she was informed in $t - 1$ or the agent was informed in $t - 1$. However, when both the principal and agent were informed in $t - 1$ and the principal was overruled by the agent, her delegation probability falls to the same level that also prevails if both were uninformed in the previous period, i.e. when the outside option was implemented. This suggests that

the principal reacts negatively to being overruled — a behavior that cannot be explained if the principal maximizes the expected utility from her monetary payoffs only. The behavioral response to overruling is robust to specifications which include period dummies and belief data (column (2)), and specifications which use only the HIGH and PHIGH treatments where delegation is expected to take place (column (3)).

The importance of the overruling effect is also corroborated by data from questions in the *ex-post* survey administered after the PHIGH50 treatments in April 2011. We asked participants in the role of the principal to evaluate two delegation scenarios in both of which they delegated and the agent’s project is implemented, and which differ only in whether they were informed or not. On a 7-point likert scale, the scenario in which the principal was uninformed had an average evaluation of 5.97 while the scenario in which the principal was informed and thus overruled had a lower average evaluation of 4.59. This difference is highly significant both in a paired t-test ($p < 0.01$) and in a non-parametric sign-rank test ($p < 0.01$).²² Hence, even though the monetary outcome is identical across both scenarios, principals on average assigned significantly lower value to the scenario in which they were overruled.

Table 5: Delegation conditional on previous experience

	(1)	(2)	(3)
principal informed in $t - 1$	0.170 (0.120)	0.194* (0.110)	0.191 (0.129)
agent informed in $t - 1$	0.208*** (0.062)	0.189*** (0.061)	0.122** (0.059)
overruled in $t - 1$	-0.371** (0.153)	-0.410*** (0.149)	-0.323** (0.159)
effort in $t - 1$	-0.004** (0.002)	-0.005 (0.003)	-0.002 (0.002)
Treatment Controls?	Yes	Yes	Yes
Belief & Time Controls?	No	Yes	Yes
Pseudo R^2	0.170	0.249	0.199
Observations	360	360	271

Marginal Effects from a probit regression are reported in the table. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses, clustered by individual. Sample is restricted to principals who delegated in $t - 1$. Neither principal nor agent informed is the omitted category. Regressions (1) and (2) contain the data from all treatments. Regression (3) contains only data from the HIGH and PHIGH treatments, where delegation is predicted. Belief controls are beliefs of principals under both authority allocations. Time controls are period fixed effects.

²²The two survey questions were randomly ordered and asked on separate screens of the exit survey.

We performed another test of the hypothesis that the non-pecuniary disutility of being overruled causes a reluctance to delegate by conducting two additional 25-period treatments which had nearly identical equilibrium returns to delegation but varied in the extent to which overruling was possible. The first of these treatments restricted the subordinate’s effort to zero and did not allow them to make recommendations. This treatment was symmetric with a payment of 40 if the own preferred project was chosen and 30 if the other party’s preferred project was chosen. We refer to this treatment as HIGH NOREC. There are relatively high (15.8 percent) equilibrium returns to delegation since the commitment not to exert effort as the subordinate increases the effort exerted by the controlling party. We compare this data to a 25 period version of the PHIGH50 treatment (PHIGH25) which had similar equilibrium returns to delegation.²³

Recall that the difference between the HIGH NOREC and the PHIGH25 treatment is that in the former the principals cannot be overruled. Therefore, if the principals derive disutility from being overruled they should be more willing to delegate in HIGH NOREC. We find indeed that the average delegation rate of the HIGH NOREC treatment was higher (67.1 percent) than in the PHIGH25 treatment (41.3 percent) - a difference that is statistically significant ($p = 0.011$).²⁴ The higher delegation rate in HIGH NOREC occurred despite the fact that the *empirical* return on delegation is 21 percentage points *higher* in the PHIGH25 treatment than in HIGH NOREC. For this reason, the higher delegation rates in HIGH NOREC provide additional support for the hypothesis that being overruled is associated with disutility.

What is the source of this disutility? A plausible answer to this question is provided by the notion of regret aversion (which we formalize in appendix A). To examine the role of regret aversion in our setting it is useful to recall that the principals chose considerably higher effort in the role of the controlling party compared to when they were the subordinate party. Therefore, if the principals in the subordinate role became informed about which project was best for them they knew that they would also have been informed if they had kept the decision right. In other words, they could have chosen their best project if they had kept their decision right. It is thus quite plausible that the principals regretted their delegation decision when they were informed and the agent’s preferred project was implemented. If this regret is psychologically aversive the delegation option becomes less valuable. This account of under-delegation in terms of regret aversion can explain the following four facts:

²³The 25 period experiments lasted on average 1.75 hours, 32 subjects participated in the PHIGH25 treatment and 28 subjects participated in the HIGH NOREC treatment. Pooling earnings across the treatments, subjects earned 44 CHF (\$47) on average.

²⁴This p-value is calculated using a probit regression controlling for period fixed effects. Standard errors are clustered by the individual.

(i) the under-delegation of authority in the HIGH and the PHIGH treatment (ii) the sharp reduction in delegation rates after being overruled; (iii) the subjects' preference (in the post-experimental survey) for being not informed about their best project when the agent's preferred project is implemented anyway; and (iv) the much higher delegation rate when the experimental design rules out *ex-post* regret by preventing the principal from providing effort after delegation.

Our explanation of under-delegation in terms of regret aversion would receive further support if an individual measure of regret aversion would predict individual differences in the reluctance to delegate. Subjects' behavior in the lottery task described in Section 2.2 may be interpreted as such an individual difference measure. If the subjects' reluctance to accept the lotteries is at least partly driven by subjects' regret aversion, the propensity to reject lotteries should also predict the reluctance to delegate (in the treatments where overruling can occur). This is in fact what we observe. In appendix C (Table C3) we report probit regressions of principals' delegation choices with treatment conditions, period fixed effects, principals' beliefs about agents' effort and the number of rejected lotteries as explanatory variables. Individuals in the HIGH and PHIGH treatments are 12 percent less likely to delegate for each gamble in the lottery task that they decline ($p = .021$). If we combine the observations in the HIGH and PHIGH treatment, and split the sample at median lottery acceptance, the difference in delegation frequency between the group which accepts more gambles and the group which accepts fewer gambles is 20 percent. These findings further support the view that regret aversion appears to be an important motive behind the reluctance to delegate authority. The next subsections will show that the same motive may also partly explain the controlling and the subordinate parties' effort choices.

3.3 Exploring the controlling parties' over-provision of effort

We saw in Figure 3 that the provision of effort by the controlling party exceeds the Nash equilibrium prediction across all treatments while the effort of the subordinate is below the Nash prediction. These deviations are persistent, with no convergence to the Nash equilibrium over time.²⁵

Persistent deviations from the Nash equilibrium might be due to one of two sources. First, for a given belief about the other party's effort, an individual may respond to those beliefs differently than the best reply. For example, if the controlling party systematically over-provides effort relative to the best reply, his or her effort is likely to be higher than the

²⁵Looking at agents' effort in the first 5 periods and the last 5 periods of the four main treatments, average effort declines by 2.4 points as a subordinate ($p = 0.014$ if effort is regressed on a dummy for periods 6-10) and decreases by .2 points as controlling party ($p = 0.76$).

Nash equilibrium effort. Likewise, if the subordinate party under-provides effort relative to the best reply, then the effort is likely to be lower than the Nash equilibrium effort.

Second, beliefs about the other party’s effort provision may deviate from those predicted in the Nash equilibrium. Because of strategic substitutability, a controlling party whose beliefs about subordinate effort are below those predicted by the Nash equilibrium will increase his or her effort relative to the Nash equilibrium. Likewise a subordinate party whose beliefs are above the Nash equilibrium will decrease effort in substitution. In this subsection we examine both the best reply channel and the belief channel as potential sources of the controlling parties’ over-provision of effort.

We first look at systematic deviations from the best response function by constructing the theoretical best response for the controlling party if control is kept and if control is delegated under the assumption of risk neutrality:

$$r_P(\hat{e}) = \frac{100\hat{P}_1 - \hat{e}\hat{P}_2}{50}, \quad r_A^d(\hat{E}^d) = \frac{100\hat{A}_2 - \hat{E}^d\hat{A}_1}{50}. \quad (12)$$

By comparing these best responses with the actual response of controlling parties to their beliefs, we can examine systematic deviations from the best response function.

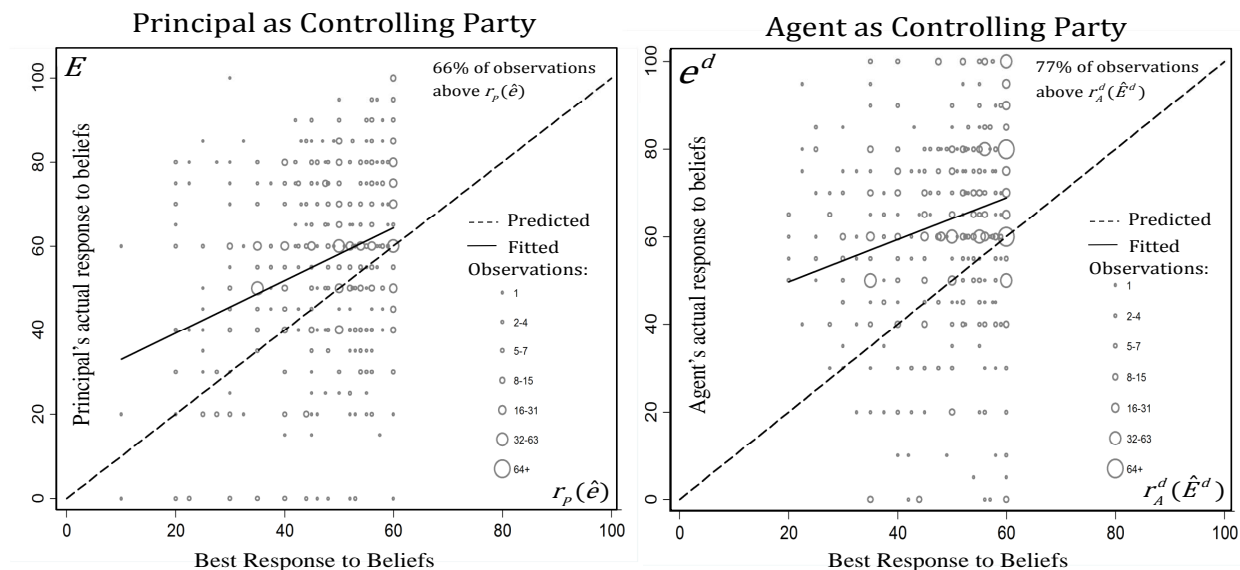


Figure 7: Controlling party: Actual effort vs. best response to beliefs (Combined data from all main treatments)

Figure 7 shows this comparison pooled for all treatments. The dashed 45° line represents those cases where the actual response to beliefs coincides with the best response to these beliefs. Points above the 45° line represent observations in which the controlling party

over-provides relative to the best response while points below the 45° line represent an under-provision of effort.

The solid line in Figure 7 shows the empirical relationship between the actual response to beliefs about the subordinates’ effort and the best response. The positive slope of this line indicates that the best response has some (qualitative) predictive power. However, the overwhelming feature in the data is the systematic over-provision of effort by the controlling party relative to the best response. Counting all observations strictly above the 45° line, 66 percent of observations for principals and 77 percent of observations for agents provide more effort than is predicted by a best response to beliefs. The magnitude of this over-provision is typically large, with 48 percent of observations 15 points or more above prediction.

Table 6: Comparison of effort provision of the controlling party to the best response to beliefs

	Principal has control			Agent has control		
	actual effort		best response effort	actual effort		best response effort
PLOW	55.7		53.9	68.1	***	49.1
LOW	66.1	***	54.5	68.3	***	55.8
HIGH	48.2	***	42.1	58.7	***	45.3
PHIGH	58.2	**	45.9	65.1	**	56.2

Significance levels calculated using a Signed-Rank test with beliefs and effort averaged by individual prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

While Figure 7 shows the combined data across all treatments, the pattern doesn’t vary qualitatively across treatments. Table 6 shows the average effort of the controlling party and the corresponding average of the best response to beliefs. As can be seen, effort provision of the controlling party is above the average best response prediction in all treatments and for both authority allocations, and in 7 of these 8 cases the difference is significant.²⁶

Based on this data, we conclude:

Result 5 *Controlling parties over-provide effort relative to their best response to beliefs about the subordinate’s effort.*

²⁶Technically, the effort provisions of individuals within a matching group may be correlated due to shared histories. However, as the information concerning the actions of others in the matching group is limited, we expect the effect of heterogeneous learning to be limited. As an additional control, we ran matching group-clustered versions of each sign-rank and rank-sum test included in this paper to check whether matching group-level effects are driving our results. As expected, the p-values of these tests are slightly higher but have a minor effect on the significance levels reported throughout the paper. See Datta & Satten (2008) and Datta & Satten (2005) for details of the two tests.

Result 5 suggests that having authority appears to have a motivational effect on the effort provision of the controlling parties. We next turn to beliefs. Since the effort of the two parties are strategic substitutes, deviations from the Nash Equilibrium prediction may partially be explained by pessimistic beliefs of controlling parties.

Table 7 compares actual beliefs to the Nash Equilibrium beliefs for all treatments and both authority allocations. As can be seen by comparing the first two columns, the principal’s belief about agent effort if control is kept is comparable to the Nash Equilibrium prediction. In fact, in three out of four cases (i.e. in PLOW, LOW and HIGH) the principals’ effort expectation is above e^{NE} , but the deviation is not significant. Thus, pessimistic beliefs of the principal cannot contribute to the over-provision of effort in these cases. The situation is somewhat different if authority was delegated. Here, the controlling party (the agents) expected in all four treatments that the subordinate party will under-provide effort relative to the Nash equilibrium. Thus, beliefs of the agents do account for some of the over-provision of effort relative to the Nash Equilibrium prediction.²⁷

Table 7: Comparison of actual beliefs of the controlling party to the Nash prediction

	Principal has control		Agent has control	
	Nash prediction	actual belief	Nash prediction	actual belief
PLOW	25	30.4	35	*** 21.8
LOW	25	27.5	25	* 20.9
HIGH	35	35.8	35	* 29.4
PHIGH	35	* 28.2	25	** 19.0

Significance levels calculated using a Signed-Rank test with beliefs and effort averaged by individual prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

3.4 Exploring the subordinate parties’ under-provision of effort

We next examine possible reasons for deviations from the Nash equilibrium on the part of the subordinates. We will again examine the role of beliefs as well as the role of systematic deviations from the best response function as potential sources of the observed deviation of effort relative to the Nash equilibrium.

Figure 8 shows the relationship between the theoretical best response and the actual response function for the subordinates. As before, the 45° line represents the predicted best response function of the subordinate in response to beliefs about the effort of the controlling party while the solid line shows the empirical best response behavior from a simple linear

²⁷Table C.4 in appendix C shows that beliefs of controlling parties about subordinate effort do, on average, exceed actual subordinate effort.

regression. Points above the 45° line represent observations in which the subordinate overprovides effort relative to the best response while points below the 45° line represent an under-provision of effort.

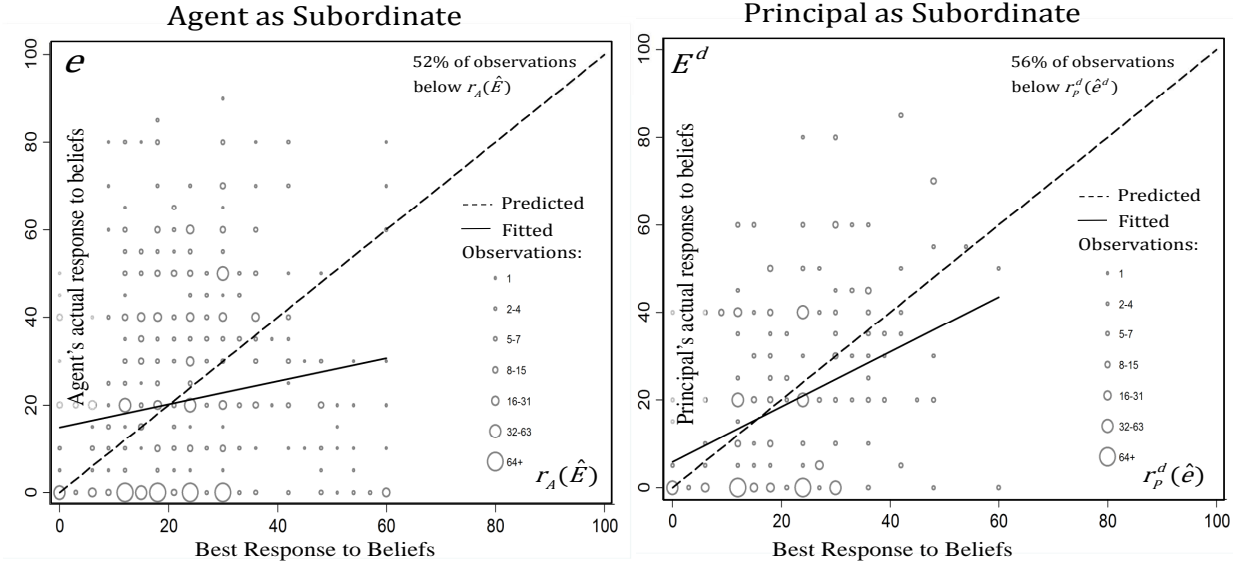


Figure 8: Subordinates: actual effort vs. best response to beliefs (Combined data from all main treatments)

As can be seen in the left hand panel of the figure, the actual response function is positively sloped but relatively flat, suggesting a relatively weak effort response to beliefs. Unlike the controlling parties' efforts, which were clustered above the best response correspondence, effort provisions by subordinates are heterogeneous: 52 percent of individual choices are below the best response to beliefs for agents, while 56 percent of individual choices are below the best response to beliefs for principals. In addition, a large number of individual choices are considerably below the best response.

A particularly salient fact in Figure 8 is that a large number of subordinates put in zero effort, i.e., lack of control appears to have a strong demotivational effect for a large minority. Recall that the cost of effort is convex with the cost of increasing effort from 0 to 5 equalling $g_P(5) - g_P(0) = .06$ points. Since incremental effort is nearly costless, zero effort is predicted only in cases where the subordinate believes in an effort of 100 by the controlling party, which almost never occurred.

The heterogeneous behavior of subordinates also appears to be a robust phenomenon at the treatment level. As shown in Table 8, which reports the average effort of the subordinate, the average theoretical best response to beliefs, and the proportion of individuals who provide zero effort for each treatment and for principals and agents separately. As can be seen by

comparing the first two columns of each treatment, there is little difference between the actual effort and the theoretical best response to beliefs at the mean. The similarity in these two averages reflects the heterogeneous nature of subordinate effort provision where both under and over-provision of effort is observed.

Table 8: Comparison of effort of subordinates to their best response to beliefs

	Agents in subordinate role			Principals in subordinate role		
	actual effort	best response to beliefs	percent zero	actual effort	best response to beliefs	percent zero
PLOW	22.8	21.1	39.0	16.5	24.2	36.7
LOW	14.3	** 19.8	49.4	16.2	18.9	54.0
HIGH	26.5	24.6	28.5	19.6	26.3	36.8
PHIGH	17.3	18.4	50.3	20.7	22.6	36.7

Significance levels calculated using a Signed-Rank test with beliefs and effort averaged by individual prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Looking at the third column of each row, it is apparent that zero-effort is a modal strategy for the subordinates in all treatments. Zero effort is observed 25 percent of the time in all eight cases and in three of these cases, zero effort accounts for for roughly 50 percent of the observations.²⁸ The high frequency of observed zeroes is due in large part to a subset of individuals who always exert zero effort in the subordinate role.²⁹

Based on these observations, we conclude:

Result 6 *The subordinates effort behavior is heterogeneous. While on average effort provision is close to the theoretical best response, there is a large group of subordinates who provide zero effort, far below the optimal best response. For this group, authority appears to have a strong demotivational effect. In addition, there is a smaller group of subordinates who systematically over-provide effort.*

²⁸One possible explanation for zero effort is that individuals who exert zero effort simply don't understand the environment. In studying the effort that agents exert as the controlling party, however, this explanation is unlikely. Remember that we collect effort choices of agents in both roles, the subordinate and the controlling party. We can therefore directly compare whether those subjects who exert zero effort as subordinates are different from those who do not when in the controlling party role. A regression of controlling party effort on a dummy that takes on the value 1 if subordinate effort is zero, controlling for treatment and clustering standard errors at the individual level, reveals that those agents who chose zero effort as subordinates on average exert 3 units of additional effort (this difference is not significant ($p=0.35$)). In fact, as will be discussed later in section 3.5, a positive difference is to be expected if regret aversion directly affects effort choices. This suggests that a lack of understanding is not the driver of zero effort choices.

²⁹Between 18 and 33 percent of agents exert zero effort in nine or ten periods. Between 8 and 50 percent of principals who delegated at least once exerted zero effort after each delegation.

Turning to beliefs, Table 9 shows the beliefs of the subordinate compared to the Nash Equilibrium. As can be seen, agents and principals have optimistic beliefs relative to the Nash equilibrium. As optimistic beliefs are expected to lead to a decrease in effort, beliefs may be contributing to the under-provision of effort by the agent and the principals. However, as we noted in Figure 8, the actual response function is much flatter than would be predicted by the best response. Whereas theory would predict that effort increases by 6 points when beliefs fall by 10 points, the actual response to beliefs is significantly smaller. For agents, a 10 point reduction in beliefs about the controlling parties' effort only leads to a 1.6 point increase in effort.³⁰

Table 9: Comparison of actual beliefs of subordinates to the Nash prediction

	Agents in subordinate role		Principals in subordinate role			
	Nash prediction	actual belief	Nash prediction	actual belief		
PLOW	55	***	64.8	45	***	59.6
LOW	55	***	66.9	55	***	68.4
HIGH	45	***	59.0	45	***	56.2
PHIGH	45	**	69.3	55	***	62.3

Significance levels calculated using a Signed-Rank test with beliefs and effort averaged by individual prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

3.5 The motivational and demotivational forces of authority: non-pecuniary explanations

Thus far we have seen that for both the controlling party and the subordinate, deviations from best response behavior account for much of the observed departure from the Nash predictions. A significant proportion of controlling parties provide effort which exceeds the best response function leading to effort levels higher than predicted. Similarly, a significant proportion of subordinates provide zero effort despite the almost negligible cost of providing incremental effort. Having established these observational facts, the question remains which behavioral force shows promise in rationalizing our data.

As we saw in the delegation section, many aspects of our data support the interpretation that regret aversion is an important force behind the under-delegation of authority. Might the same behavioral force also have promise to rationalize the observed deviations in effort

³⁰For principals, effort increases by 3.8 points when beliefs fall by 10 points. Both coefficients are significantly smaller than the theoretically expected 6 point increase ($p < 0.01$) in a simple regression of effort on beliefs. Table C.4 in appendix C reports beliefs about controlling party effort in comparison to actual controlling party effort choices.

choices?

To begin our analysis, we take the same overruling effect which was found to be of importance in delegation and ask to what extent it can explain the effort patterns of a subordinate. Recall that as a subordinate there is the potential of being overruled. If such overruling generates non-pecuniary losses — as predicted by regret aversion — a subordinate who anticipates this disutility may reduce effort thereby reducing the probability with which overruling occurs (see appendix A for more details). For cases where this non-pecuniary loss is particularly strong, effort provision in the subordinate role may be driven to zero.

The explanation that zero effort is a response to anticipatory regret would receive support if our individual measure of regret aversion would correlate with individual differences in zero effort. Our interpretation is supported in precisely this way: In a probit regression, the probability that a subordinate exerts zero effort increases by 5.9 percent per additional gamble rejected ($p=0.030$).³¹

Regret in the domain of effort choices may also result in an over-provision of effort by the controlling party if there is a non-pecuniary disutility for being unsuccessful in implementing one’s own preferred project. Filiz-Ozbay & Ozbay (2007) have shown that individuals who lose in first price sealed bid auctions to bids which are below their true valuation behave as if they experience “loser regret”, i.e., these individuals behave as if they suffer a disutility from losing when they could have won, and earned profit, by making a higher bid. In our setting, the controlling party may hence regret his or her effort choice in cases where it remains uninformed and thus could have improved the project selection by putting in more effort. An individual in the role of the controlling party who anticipates such regret optimally raises his or her effort (see appendix A for more details). Again, we can test this conjecture by using regression analysis and looking at the correlation between over-provision of effort and our individual measure of regret aversion. In a probit regression, we find that the probability of over-provision increases by 4.5 percent per additional gamble rejected ($p=0.058$).³² Thus, taken together, regret aversion may be a driving force behind all three major experimental patterns - the under-delegation of authority, the high frequency of zero effort choices among subordinates and the over-provision of effort by the controlling parties.³³

³¹A dummy variable is created that takes on the value 1 in case of zero effort provision in the role of the subordinate. Data from the four main treatments is included in this regression. The probit regression also contains controls for the treatment, being in the role of the agent, interactions of role and treatment, controls for beliefs, and period dummies. Standard errors are clustered at the individual level.

³²A dummy variable is created that takes on the value 1 if a subject over-provided effort in the role of the controlling party on average. Data from the four main treatments is included in this regression. The regression also contains controls for the treatment, being in the role of the agent, and interactions of role and treatment. Each individual in our data set who chose effort in the controlling party role at least once is an observation.

³³As was stated in result 6, we also observe a minority of subordinates who actually over-provide effort

While regret aversion shows promise in rationalizing our data, it is reasonable to ask whether alternative hypotheses show similar promise. Could it be, for example, that the agents' effort choices are influenced by reciprocity or other forms of social preferences? Or, may risk or loss aversion generate a similar pattern to regret aversion?

A common reason for deviations from standard economic predictions is the existence of social preferences. In our setting, if agents view the delegation of authority as a kind act they may over-provide effort because of reciprocal motivations. Likewise, if they view a lack of delegation as an unkind act they may under-provide effort relative to their best response. Thus, positive and negative reciprocity may, in principle, explain the agents' effort pattern. We tested for the impact of reciprocity motives by conducting an additional treatment in the HIGH condition in which the delegation decision was decided exogenously by the computer. In this HIGH RAND treatment, a virtual coin is flipped each period which determines whether control rights are kept by the principal or whether the principal is forced to delegate them. Since the agents know that the principals are forced to make a choice it is impossible to attribute kind or unkind intentions to the principal. If positive or negative reciprocity play a role, the agents' effort choices in the HIGH RAND condition will deviate from their choices in the HIGH condition. However, neither as a controlling party (Kolmogorov Smirnov test, $p = .20$) nor in the position of the subordinate party (Kolmogorov Smirnov test, $p = .99$) do the agents' effort choices differ in the two conditions, implying that reciprocity is unlikely to explain their effort pattern.

In all of our treatments, the controlling party over-provides effort relative to her best response which directly increases the expected earnings of the subordinate. Thus, altruism on the part of the controlling parties could explain this pattern of effort. To control for this possibility we implemented an additional control treatment with the following features. Only one of the two subjects was given the ability to provide effort and to choose the project, but both parties were paid based on the controlling party's project choice. Thus, in this treatment the passive party never receives the decision right and never makes an effort choice but only collects her payoffs. We compare this treatment with the single player game (described at the beginning of Section 2.3) which is identical to the above control treatment except that no passive recipient exists. Thus in the additional control treatment social preferences can affect the active subject's effort while in the single player game social preferences cannot

relative to their best response. In appendix A we show that heterogeneity in subordinate effort is also consistent with a model with both loser regret and overrule regret. Recall that the role of controlling party is influenced only by loser regret, and therefore individuals are predicted to exert effort above the Nash Equilibrium in this role. As the subordinate, an individual is exposed both to the potential of being overruled and the potential of loser regret. These forces go in opposite directions, and the response in effort will therefore depend on an individual's inclination towards both types of regret.

play a role. It turns out that the effort choices of the active party and the single player are indistinguishable (Kolmogorov Smirnov test, $p = 0.44$), indicating that social preferences do not significantly affect effort.

Another potential reason for deviations from theoretical predictions is that the assumption of expected value maximization may be violated due to loss aversion. In the appendix D, we show that loss aversion with a reference point at the outside option cannot explain the over-provision of effort by the controlling parties. The intuitive reason for this claim is as follows. For loss averse individuals, an increase in effort above the risk neutral optimum increases the magnitude of a potential loss which reduces utility. This follows from the fact that an increase in effort causes a sure increase in costs but as long as the possibility of success is below 1 the controlling parties' *ex post* payoff from unsuccessful search may not cover the effort cost. Thus, for reasonable amounts of loss aversion, optimal effort is decreasing in an individual's degree of loss aversion. For unrealistically extreme levels of loss aversion, an individual may prefer to guarantee a payoff rather than playing any lottery. For controlling parties with such extreme levels of loss aversion, providing maximal effort (which guarantees a payoff of 15) may be preferable to providing low effort and hoping for success by the subordinate. In these cases, loss aversion would predict a maximal effort level of 100.

Looking at both cases in combination, loss aversion cannot explain effort levels which are above the best response function but below an effort level of 100. As these are the observations which need to be rationalized in order to explain the over-provision of effort by the controlling parties, loss aversion cannot explain our effort results.

In appendix E, we show that similar to loss aversion, neither risk aversion nor risk loving preferences can account for over-provision of effort by the controlling party. Moreover, neither risk nor loss aversion can explain the subordinates' choice of zero effort levels because effort costs are negligible at low effort levels. For example, assuming risk aversion, beliefs about controlling party effort need to be extremely high to rationalize a subordinate's effort choice of zero. Using a CRRA utility specification of the following form, $U(x) = \frac{x^{1-\sigma}}{1-\sigma}$, an effort of 0 is only predicted if the belief in controlling party effort is 100 (up to $\sigma = 8$). Hence, only for counterfactually high beliefs (or unrealistically extreme levels of risk aversion) is zero effort expected to occur.

4 Conclusion

Authority and power permeate political, social and economic interactions. It is therefore important to understand their motivation and incentive effects. In this paper we tackle this question by using a novel experimental design. We find a strong behavioral bias among

principals to retain authority against their pecuniary interests and often to the disadvantage of both the principal and the agent. We demonstrate that under-delegation cannot be attributed to principals' beliefs nor incomplete learning, and that the individual and aggregate income losses of this delegation bias are substantial. Our results suggest that authority has non-pecuniary consequences that inhibit the reallocation of authority.

Our results also show that authority has effects on the motivation to provide effort that are not captured by the theoretical model. The fundamental trade-off between incentives and control, as modeled by Aghion & Tirole (1997), indeed exists; relative to the first best the subordinate provides too little effort, and the controlling party provides too much. Further, the comparative statics between treatments are well met. However, the inefficiency generated by the incentive conflict is much greater than predicted by theory. The controlling parties provide significantly more effort and the subordinate parties provide significantly less relative to the Nash Equilibrium prediction. For controlling parties and a large fraction of subordinates, this is also true relative to the best response to beliefs.

A deeper look at our data suggests that a distaste for being overruled is a substantial determinant of the desire to retain control. Principals who are overruled after delegation and earn the return from the agent's preferred project are less likely to delegate in the future relative to those who are uninformed and earn the very same return. This difference in delegation behavior — driven by informational states and not by pecuniary payoffs — suggests that individuals are suffering a non-pecuniary disutility from being overruled. One potential source of such non-pecuniary disutility is regret, a theory which can also help to explain the high frequency of zero effort among subordinates and the over-provision of effort by the controlling parties.

Given the importance of authority and power in the functioning of economic and political organizations we believe that the motivational biases revealed by our data should receive more attention. In addition, further explorations into the motives behind delegation and control are suggested by our data. Although a significant part of underdelegation can be explained by principals' regret aversion we also observed a nonnegligible underdelegation in the treatment HIGH NOREC where the principals' recommendation could not be overruled after delegation. This raises the question whether some subjects intrinsically prefer to be the controlling rather than the subordinate party. In Bartling, Fehr & Herz (2012) we show that this is indeed the case. Moreover, it is well possible that delegation of authority is affected by the mechanism by which authority has been initially assigned. For example, if the possession of authority is the result of prior superior performance, the principals might even be more reluctant to delegate compared to the random assignment of authority. This additional underdelegation may result from the perception of superior competence, from

overconfidence, or from the status gains associated with a position of authority. In addition, research on the cultural determinants of the sources of under-delegation and the motivational effects of decision rights may be interesting because societies seem to be quite heterogeneous with regard to how they view and legitimize hierarchical structures.

References

- Aghion, Philippe, and Jean Tirole.** 1997. “Formal and Real Authority in Organizations.” *Journal of Political Economy*, 105(1): 1–29.
- Aghion, Philippe, Mathias Dewatripont, and Patrick Rey.** 2004. “Transferable Control.” *Journal of the European Economic Association*, 2(1): 115–138.
- Athey, Susan, and Donald J. Roberts.** 2001. “Organizational Design: Decision Rights and Incentive Contracts.” *The American Economic Review*, 91(2): 200–205.
- Baker, George, Robert Gibbons, and Kevin J. Murphy.** 1999. “Informal Authority in Organizations.” *Journal of Law, Economics, and Organization*, 15(1): 56–73.
- Bartling, Björn, and Urs Fischbacher.** 2012. “Shifting the Blame: On Delegation and Responsibility.” *Review of Economic Studies*, 79: 67–87.
- Bartling, Björn, Ernst Fehr, and Holger Herz.** 2012. “The Intrinsic Value of Authority.” *Working Paper, Department of Economics, University of Zurich*.
- Bester, Helmut, and Daniel Krähmer.** 2008. “Delegation and Incentives.” *RAND Journal of Economics*, 39(3): 664–682.
- Bowles, Samuel, and Herbert Gintis.** 1988. “Contested Exchange: Political Economy and Modern Economic Theory.” *American Economic Review*, 78: 145–150.
- Charness, Gary, Ramon Cobo-Reyes, Natalia Jimenez, Juan A. Lacomba, and Francisco Lagos.** 2012. “The Hidden Advantages of Delegation: Pareto-Improvements in a Gift-Exchange Game.” *American Economic Review*, 102(5): 2358–2379.
- Coffman, Lucas C.** 2011. “Intermediation Reduces Punishment (and Reward).” *American Economic Journal: Microeconomics*, 3: 1–30.
- Dahl, Robert.** 1957. “The Concept of Power.” *Behavioral Science*, 2: 201–215.
- Datta, Somnath, and Glen Satten.** 2005. “Rank-Sum Tests for Clustered Data.” *Journal of the American Statistical Association*, 471(1): 908–915.
- Datta, Somnath, and Glen Satten.** 2008. “A Signed-Rank Test for Clustered Data.” *Biometrics*, 64(2): 501–507.

- deCharms, R.** 1968. *Personal Causation: The Internal Affective Determinants of Behavior*. New York:Academic Press.
- Deci, Edward L.** 1981. *The Psychology of Self-Determination*. Lexington, MA.:Health.
- Dessein, Wouter.** 2002. "Authority and Communication in Organization." *Review of Economic Studies*, 69(1): 811–838.
- Falk, Armin, and Michael Kosfeld.** 2006. "The Hidden Cost of Control." *The American Economic Review*, 96(1): 1611–1630.
- Fehr, Ernst, and Bettina Rockenbach.** 2003. "Detrimental Effects of Sanctions on Human Altruism." *Nature*, 422(13): 137–140.
- Fehr, Ernst, and John A. List.** 2004. "The Hidden Costs and Returns of Incentives-Trust and Trustworthiness Among CEOs." *Journal of the European Economic Association*, 2(5): 743–771.
- Filiz-Ozbay, Emel, and Erkut Y. Ozbay.** 2007. "Auctions with Anticipated Regret: Theory and Experiment." *The American Economic Review*, 97(4): pp. 1407–1418.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10(2): 171–178.
- Gelman, Andrew, and Gary King.** 1990. "Estimating Incumbency Advantage Without Bias." *American Journal of Political Science*, 34: 1142–1164.
- Greiner, Ben.** 2004. "The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics." University of Cologne, Department of Economics Working Paper Series in Economics 10.
- Grossman, Sanford J, and Oliver D Hart.** 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy*, 94(4): 691–719.
- Harsanyi, John C.** 1978. "The Dimension and Measurement of Social Power." In *Power in Economics*. , ed. K. W. Rothschild. Harmondsworth:Penguin Books.
- Hart, Oliver, and John Moore.** 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy*, 98(6): 1119–58.
- Hart, Oliver, and John Moore.** 1995. "Debt and Seniority: An Analysis of the Role of Hard Claims in Constraining Management." *American Economic Review*, 85(3): 567–85.
- Jensen, Michael C.** 1986. "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers." *American Economic Review*, 76.

- Koszegi, Botond, and Matthew Rabin.** 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics*, 121(4): 1133–65.
- Loomes, Graham, and Robert Sugden.** 1982. "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty." *The Economic Journal*, 92(368): 805–824.
- Marx, Karl.** 1867. *Das Kapital*.
- Mulder, Mauk.** 1975. *The Daily Power Game*. Leiden:Martinus Nijhoff.
- Niskanen, William A.** 1971. *Bureaucracy and Representative Government*. Chicago:Aldine-Atherton.
- Parsons, Talcott.** 1963. "On the Concept of Political Power." *Proceedings of the American Philosophical Society*, 232–262.
- Rabin, Matthew.** 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica*, 68(5): 1281–92.
- Rantakari, Heikki.** 2008. "Governing Adaptation." *Review of Economic Studies*, 75(4): 1257–1285.
- Rotter, Julian B.** 1966. "Generalized Expectancies for Internal Versus External Control of Reinforcement." *Psychological Monographs: General and Applied*, 80: 1–28.
- Russell, Bertrand.** 1938. *Power: A New Social Analysis*. London:Allen and Unwin.
- Simon, Herbert.** 1951. "A Formal Theory of the Employment Relation." *Econometrica*, 19(3): 293–305.
- Sliwka, Dirk.** 2001. "On the Costs and Benefits of Delegation in Organizations." *Journal of Institutional and Theoretical Economics (JITE)*, 157(4): 568–.
- Weber, Max.** 1978. *Economy and Society, translated by G. Roth and C. Wittich*. Berkeley and Los Angeles, CA:University of California Press.
- Zeuthen, Frederik.** 1968. "Economic Welfare." In *Problems of Monopoly and Economic Welfare*. 104–151. London, UK:Routledge and Kegan Pail Ltd. (original publication, 1930).

Appendix A: Regret Theory

This appendix examines the extent to which regret theory can rationalize our data. It is divided into three parts. In part one, we consider the effort stage of the experiment and concentrate on the decision problem of an agent. We consider two different sources of regret:

loser regret and overrule regret. An agent experiences *loser regret* if he remains uninformed but could have achieved a higher payoff had he chosen a higher effort and been informed. An agent experiences *overrule regret* if he is in the role of the subordinate, the controlling party is informed, and the agent’s recommendation is disregarded or his effort is wasted. We show that loser regret and overrule regret can rationalize important aspects of the agents’ behavior. In particular, loser regret induces agents in the position of the controlling party to overexert effort relative to the best reply of those without loser regret. Overrule regret, by contrast, induces agents in the subordinate position to reduce effort relative to agents who have no overrule regret.

In part two, we show that these two regret forces also suffice to explain effort choices as well as under-delegation of authority by the principals. In part three, we extend the analysis of the principal and include a third form of regret that only principals can experience: *delegation regret*. In contrast to the agents, regret experienced by principals can also stem from their delegation decision, i.e., having delegated or not having delegated the decision right. We show that including regret that stems from delegation can further decrease a principal’s utility from delegating and has effects on effort similar to the regret forces studied in the first two parts.

Part 1: Regret and Effort Decisions of the Agent: In the auction literature, it has recently been proposed by Filiz-Ozbay & Ozbay (2007) that a reason for overbidding in the first price sealed bid auction is that individuals experience “loser regret.” Loser regret occurs when an individual bids below their valuation and is beaten by a higher bid that is below their true valuation. In these cases, individuals experience regret because they would have preferred to bid higher *ex post* than is optimal to bid *ex ante*. An individual who anticipates such regret optimally increases their bid relative to the risk neutral Nash equilibrium in order to reduce the potential states for which regret occurs.

In our experiment, individuals may similarly experience loser regret in cases where they remain uninformed and thus cannot implement their preferred project. These will be cases in which an individual’s effort is below the number drawn by the random number generator that guides success and failure of an individual’s effort.

As the likelihood of being informed, and therefore the likelihood of regret, is based on an agent’s effort relative to a number drawn by nature, we require a formal way of expressing these draws. Let x_A be the realization of the random number generator (uniform between 0 and 1) for the agent, where the agent is informed if his effort is above or equal to the realization of x_A and uninformed if it is below x_A . Likewise, let x_P be the realization of the number generator for the principal, with the principal being informed when his effort is

above or equal to x_P and uninformed if his effort is below x_P .³⁴

In developing a formal model, we follow Loomes & Sugden (1982) and assume that loser regret enters utility linearly. Individuals experience loser regret when they remain uninformed and a project with a lower payoff is implemented compared to the payoff that would have resulted from the *ex-post* optimal effort decision of the individual. The magnitude of regret is related to the difference between the actual payoff and the payoff from this optimal effort decision.

We begin by considering an agent who has received decision rights and is now the controlling party. Given an implemented project k , an exerted cost of effort e^d , and a draw from the number generator x_A , the agent experiences loser regret equal to

$$\lambda_{LR} \max\{[A_2 - g(x_A)] - [A_k - g(e^d)], 0\} \quad (13)$$

any time his preferred project is not implemented, where the parameter $\lambda_{LR} \geq 0$ is the agent's degree of loser regret. Note that the max function explicitly rules out rejoicing and that the utility from the improved project choice, $A_2 - A_k$ must exceed the additional cost of effort $g(x_A) - g(e^d)$ in order for loser regret to be positive.

Let \hat{E}^d be the agent's belief about the principal's effort in the role of the subordinate. Based on these beliefs and the realizations of x_P and x_A , the outcome space can be partitioned into four distinct "cells" which differ in the extent to which regret influences utility. These cells are shown in figure (9).

In cell (i), we assume that an agent does not experience regret since he is informed and therefore is able to implement his preferred project. In the remaining three cells, however, the agent's effort is below the threshold for success ($e^d < x_A$) and the agent is uninformed. In these cases, the agent may regret his insufficient level of effort if the gain from being informed through improved project selection exceeds the incremental cost of attaining this information.

In cell (ii), the principal is uninformed and project 0 is recommended. If such a state is realized, the agent always prefers to be informed and regrets his insufficient effort level. In cells (iii) and (iv), the principal is informed and recommends project 1. Cell (iii) contains states in which the increased returns due to improved project choice exceed the cost of raising effort from e^d to x_A , i.e., states where $\hat{A}_2 - \hat{A}_1 \geq g_A(x_A) - g_A(e^d)$, and therefore the agent experiences loser regret. Cell (iv) contains states in which the additional cost of being

³⁴Individuals in the experiment were informed of their own draw from the number generator in each period. They were uninformed about the other party's draw. For loser regret and overrule regret considered in the first part of the appendix, only the information state of the other party matters for regret, not their actual draw.

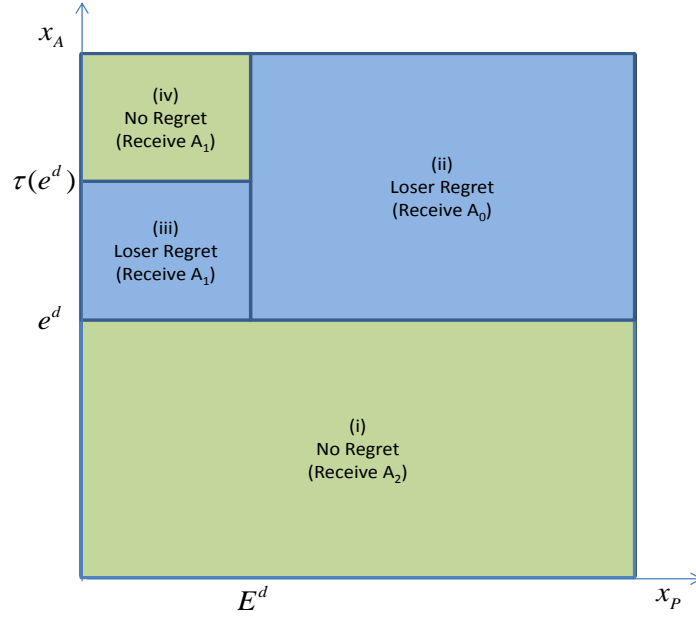


Figure 9: **Agent as Controlling Party:** As a controlling party, the outcome space can be partitioned into four cells that differ with regard to the regret experienced by the agent: In cell (i), the agent does not experience regret because he is informed and can implement his preferred project. In cells (ii) and (iii), however, a regret averse agent will experience loser regret since he could have achieved a higher payoff had he increased his effort. In cell (ii), the principal is also uninformed and the agent always experiences loser regret, since it would have been profitable to increase effort. In cells (iii) and (iv), the principal is informed and the principal's preferred project is implemented. In cell (iii), the return from having the agent's preferred project implemented, $\hat{A}_2 - \hat{A}_1$, exceeds the additional cost of acquiring the necessary information, given by $g_A(x_A) - g_A(e^d)$. Therefore, the agent experiences loser regret. In cell (iv), the additional effort cost exceeds the increased project return, and the principal does not experience regret. The cutoff between cells (iii) and (iv) is given by $\tau(e^d) \equiv \min\{g_A^{-1}(\hat{A}_2 - \hat{A}_1 + g_A(e^d)), 1\}$, where the min function is included to bound $\tau(e^d)$ in cases in which $\hat{A}_2 - \hat{A}_1$ exceeds the potential increase in effort costs.

informed exceeds its value, and therefore the agent does not experience loser regret. The threshold between cells (iii) and (iv) is given by $\tau(e^d)$, where $\tau(e^d) \equiv \min\{g_A^{-1}(\hat{A}_2 - \hat{A}_1 + g_A(e^d)), 1\}$.³⁵

Considering all four cells for the computation of utility, an agent in the role of the controlling party has utility $u_A^d(e^d|x_A, x_P, \hat{E}^d) =$

$$= \begin{cases} A_2 - g_A(e^d) & \text{if } x_A \leq e^d, \\ A_0 - g_A(e^d) - \lambda_{LR}[\hat{A}_2 - g_A(x_A) + g_A(e^d)] & \text{if } x_A > e^d \ \& \ x_P > \hat{E}^d, \\ A_1 - g_A(e^d) - \lambda_{LR}[\hat{A}_2 - \hat{A}_1 - g_A(x_A) + g_A(e^d)] & \text{if } \tau(e^d) \geq x_A > e^d \ \& \ x_P \leq \hat{E}^d, \\ A_1 - g_A(e^d) & \text{if } x_A > \tau(e^d) > e^d \ \& \ x_P \leq \hat{E}^d. \end{cases} \quad (14)$$

Intuitively, individuals who are in the role of the controlling party experience loser regret only in cases where they exert less effort than the amount needed to be informed. Thus, individuals who anticipate loser regret are likely to increase their effort relative to that of the best response of a standard expected value maximizer. The following proposition formalizes this intuition:

Proposition 1 *In the effort stage of the authority-delegation game, an agent who anticipates loser regret and who has received control from a delegating principal will over exert effort relative to the best response of an individual who maximizes expected value.*

Proof. An agent who has anticipatory regret maximizes the expected value of $u_A^d(e^d|x_A, x_P, \hat{E}^d)$ over all realizations of x_A and x_P . Taking into consideration the cases in which regret will occur, this is equivalent to maximizing:

$$\begin{aligned} \max_{e^d} \quad & e^d \hat{A}_2 + (1 - e^d) \hat{E}^d \hat{A}_1 - g_A(e^d) \\ & - \lambda_{LR}(1 - \hat{E}^d)(1 - e^d)[\hat{A}_2 - \mathbb{E}_{x_A}(g_A(x_A)|x_A > e^d) + g_A(e^d)] \\ & - \lambda_{LR} \hat{E}^d (\tau(e^d) - e^d)[\hat{A}_2 - \hat{A}_1 - \mathbb{E}_{x_A}(g_A(x_A)|x_A \in (e^d, \tau(e^d))) + g_A(e^d)] \end{aligned}$$

As the derivative of $\tau(e^d)$ is discontinuous at 1, the first order condition is solved separately for $\tau(e^d) < 1$ and $\tau(e^d) = 1$. In the case of $\tau(e^d) < 1$, two intermediate results are useful for constructing the first order condition. First note that:

$$- \frac{d}{de^d}(1 - e^d) \mathbb{E}_{x_A}(g_A(x_A)|x_A > e^d) = - \int_{e^d}^1 g_A(z) dz = g_A(e^d) \quad (15)$$

³⁵Since x_A is bounded above at 1, $\tau(e^d)$ is also bounded above at 1. If $\tau(e^d) = 1$ the agent always experiences loser regret. Note that $\hat{A}_2 - g_A(1) + g_A(0) > 0$ for the parameters chosen so that the agent always regrets not implementing his best project when the outside option is implemented.

by Leibniz's rule. Further,

$$-\frac{d}{de^d}(\tau(e^d) - e^d)\mathbb{E}_{x_A}(g_A(x_A)|x_A \in (e^d, \tau(e^d))) = -\tau'(e^d)[\hat{A}_2 - \hat{A}_1 + g_A(e^d)] + g_A(e^d) \quad (16)$$

by Leibniz's rule and the fact that $g_A(\tau(e^d)) = \hat{A}_2 - \hat{A}_1 + g_A(e^d)$. Using these intermediate calculations, the first order condition of this equation can be expressed as the following implicit function:

$$\begin{aligned} \hat{A}_2 - \hat{E}^d \hat{A}_1 - g'_A(e^d) + \lambda_{LR}(1 - \hat{E}^d)[\hat{A}_2 - (1 - e^d)g'_A(e^d)] + \\ + \lambda_{LR}\hat{E}^d[\hat{A}_2 - \hat{A}_1 - (\tau(e^d) - e^d)g'_A(e^d)] = 0. \end{aligned} \quad (17)$$

Effort is strictly above the best response of a standard expected value maximizer without regret if for a positive λ_{LR} the last two terms are positive when evaluated at or below the standard best response. Note that at the standard best response correspondence, $g'_A(e^d) = \hat{A}_2 - \hat{E}^d \hat{A}_1$, we can substitute in for $g'_A(e^d)$ to test this restriction. Looking at the last two terms with $g'_A(e^d)$ replaced with $\hat{A}_2 - \hat{E}^d \hat{A}_1$ yields

$$\lambda_{LR}(1 - \hat{E}^d)[e^d \hat{A}_2 + (1 - e^d)\hat{E}^d \hat{A}_1] \quad (18)$$

for the second to last term and

$$\lambda_{LR}\hat{E}^d[\hat{A}_2 - \hat{A}_1 - (\tau(e^d) - e^d)[\hat{A}_2 - \hat{E}^d \hat{A}_1]] \quad (19)$$

for the last term. Subtracting $\hat{E}^d(1 - \hat{E}^d)\hat{A}_1$ from expression (18) and adding it to expression (19), these two sub-equations can be further rewritten as

$$\lambda_{LR}(1 - \hat{E}^d)e^d[\hat{A}_2 - \hat{E}^d \hat{A}_1] \quad (20)$$

and

$$\lambda_{LR}\hat{E}^d[\hat{A}_2 - \hat{A}_1 + (1 - \hat{E}^d)\hat{A}_1 - (\tau(e^d) - e^d)[\hat{A}_2 - \hat{E}^d \hat{A}_1]]. \quad (21)$$

Expression (20) is clearly positive since e^d and E^d take intermediate values between zero and one along the best response function and $\hat{A}_2 > \hat{A}_1$. Expression (21) is decreasing in $\tau(e^d)$ and thus is (weakly) larger than

$$\lambda_{LR}\hat{E}^d[\hat{A}_2 - \hat{A}_1 + (1 - \hat{E}^d)\hat{A}_1 - (1 - e^d)[\hat{A}_2 - \hat{E}^d \hat{A}_1]] = \lambda_{LR}\hat{E}^d e^d[\hat{A}_2 - \hat{E}^d \hat{A}_1], \quad (22)$$

which is also strictly positive. Thus, expression (21) is positive. As both terms are positive,

it follows that an individual who experiences loser regret will exert more effort than an individual who maximizes expected value for any given belief about the other parties effort.

■

Proposition 1 shows that controlling agents who experience loser regret tend to overprovide effort relative to an expected value maximizer which rationalizes an important aspect of our data.³⁶

Turning to the subordinate role, the agent’s optimization problem and the potential sources of regret change considerably. In particular, as a subordinate the agent can experience regret whenever the controlling party is informed since subordinate effort is wasted in these cases and it would have been optimal *ex post* to free ride on the informed principal. We refer to this form of regret as *overrule regret*. Regretting wasted effort is likely to be particularly salient when the agent is successful and the information generated from his effort is wasted. To account for the particular salience of this event, we assume that the agents who are overruled not only experience regret due to their wasted effort, but also in proportion to the foregone payoffs lost due to their information being ignored.³⁷

In figure (10) we again partition the state space into cells that differ with regard to the regret experienced by the agent. We continue to assume that no regret is experienced if the agent receives the payoff from his own preferred project. This is the case in cell (i), since the agent is informed and the controlling party is not. In cell (ii), the agent experiences loser regret since both parties remain uninformed and project 0 is implemented. In such states, the agent regrets his insufficient effort level since he could have improved project selection had he chosen $e = x_A$.

³⁶Note that in principle, individuals may also experience “winner regret” in which an individual regrets over exertion relative to the level of effort needed to be informed. This force could be added to our model and would not change the main propositions as long as anticipated loser regret is not outweighed by anticipated winner regret. Note that if winner regret would be stronger than loser regret one cannot explain the overprovision of effort by the controlling parties. This suggests that winner regret is weaker than loser regret in our setting. This is precisely the result reported in Filiz-Ozbay & Ozbay (2007), who find strong evidence of loser regret but conclude that “winner regret” is either a weaker force or unanticipated by subjects. Therefore, to keep our model simple and parsimonious, we have excluded winner regret from the analysis.

³⁷Formally, the agent experiences overrule regret equal to $\lambda_{OR}g_A(e)$ if the principal is informed and the agent is uninformed, and overrule regret equal to $\lambda_{OR}[\hat{A}_2 - \hat{A}_1 + g_A(e)]$ if both parties are informed and the agent is overruled. We also considered specifications in which (i) wasted effort and the disutility of being overruled had different coefficients and (ii) where one of the two forces was excluded. As both forces move in the same direction, both forms of regret lead to a reduction in effort relative to the standard best response, and therefore there is no qualitative differences across these models. However, regret proportional to the foregone payoff ($\hat{A}_2 - \hat{A}_1$) due to the overruled recommendation is necessary to predict zero effort choices by subordinates. If wasted effort is the only source of overrule regret, the marginal increase in anticipatory regret at an effort level of zero is zero, implying that positive effort is always predicted. This is not the case when the foregone payoff ($\hat{A}_2 - \hat{A}_1$) matters for overrule regret.

Cells (iii) and (iv) are cases in which the controlling party is informed. As the agent in the subordinate role is not in control of final project selection, subordinate effort in these cases is effectively wasted, and the agent experiences overrule regret. In cell (iii), the agent remained uninformed and therefore regrets having wasted his effort. In cell (iv), the agent was informed himself and experiences overrule regret not only from wasted effort, but also from having his recommendation ignored by the principal.

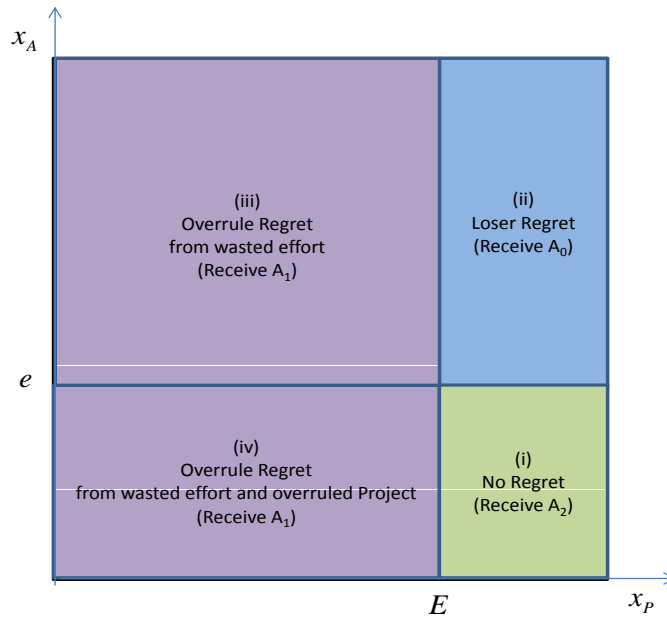


Figure 10: **Agent as Subordinate:** For agents as subordinates, the state space can be partitioned into four cells, which differ in the extent to which the agent experiences regret: In cell (i), the agent experiences no regret because his preferred project is implemented. In cell (ii), the agent experiences loser regret. Here both parties remain uninformed, implying that the agent could have improved the outcome by raising his own effort to $e = x_A$. The agent experiences overrule regret from wasted effort whenever the controlling party is informed, which is the case in cells (iii) and (iv). In cell (iv), overrule regret is particularly strong because the agent is also informed, but the agent’s recommendation is ignored.

As with loser regret, we model “overrule regret” in a linear fashion. Let e be the effort of the agent in the role of the subordinate and let \hat{E} be the agent’s belief about the principal’s effort in the role of the controlling party. The utility of an agent in the role of the subordinate

is then given by:

$$u_A(e|x_A, x_P, \hat{E}) = \begin{cases} A_2 - g_A(e) & \text{if } x_A \leq e \ \& \ x_P > \hat{E} \\ A_0 - g_A(e) - \lambda_{LR}[\hat{A}_2 - g_A(x_A) + g_A(e)] & \text{if } x_A > e \ \& \ x_P > \hat{E} \\ A_1 - g_A(e) - \lambda_{OR}g_A(e) & \text{if } x_A > e \ \& \ x_P \leq \hat{E} \\ A_1 - g_A(e) - \lambda_{OR}[\hat{A}_2 - \hat{A}_1 + g_A(e)] & \text{if } x_A \leq e \ \& \ x_P \leq \hat{E} \end{cases} \quad (23)$$

where $\lambda_{OR} \geq 0$ is the agent's degree of overrule regret.

Just as loser regret can increase effort relative to the best response, individuals who anticipate overrule regret will decrease effort in order to reduce the possibility of being overruled. Depending on whether an individual is more sensitive to loser regret or overrule regret, effort in the subordinate role can be either higher or lower than the standard best response. Effort may also be zero if individuals experience a significant amount of overrule regret and the degree of loser regret isn't too strong.

Proposition 2 *In the effort stage of the authority-delegation game, an agent who is in the role of the subordinate may experience either loser regret or overrule regret depending on the realized state. Individuals who anticipate a disutility of being overruled will decrease effort relative to those who do not. Individuals who anticipate loser regret will increase effort relative to those who do not. As these forces move in different directions, heterogeneity in anticipatory regret may lead to observed effort choices both above and below the best response.*

Proof. As before, an agent who has anticipatory regret maximizes the expected value of $u_A(e^d|x_A, x_P, \hat{E})$ over all realizations of x_A and x_P . After some simplifications, the agent maximizes:

$$\begin{aligned} \max_e \quad & \hat{E}[\hat{A}_1 - e\lambda_{OR}(\hat{A}_2 - \hat{A}_1)] - g_A(e) + [1 - \hat{E}]\hat{A}_2[e - \lambda_{LR}(1 - e)] + \\ & + [1 - \hat{E}](1 - e)\lambda_{LR}\mathbb{E}_{x_A}(g_A(x_A)|x_A > e) - [1 - \hat{E}](1 - e)\lambda_{LR}g_A(e) - \hat{E}\lambda_{OR}g_A(e) \end{aligned} \quad (24)$$

Taking the first order condition yields the following implicit function:

$$(1 + \lambda_{LR})[1 - \hat{E}]\hat{A}_2 - \hat{E}\lambda_{OR}(\hat{A}_2 - \hat{A}_1) = g'_A(e)[1 + (1 - \hat{E})(1 - e)\lambda_{LR} + \hat{E}\lambda_{OR}] \quad (25)$$

As $[1 - \hat{E}][1 - e] < 1 - \hat{E}$, effort is again higher when λ_{LR} is positive. However, since the left hand side is decreasing in λ_{OR} while the right hand side is increasing, overrule regret leads to a decrease in effort relative to an expected value maximizer. As overrule regret and loser regret go in opposite directions, effort choices as a subordinate should be heterogeneous

depending on the magnitude of these forces in individuals' utility functions. ■

Part 2: Overrule Regret and the Delegation Decision of the Principal: Having considered how regret affects the effort decision of an agent in the controlling party and the subordinate role, we next turn to the effort and delegation decisions of the principal. Just as with the agent, a principal can experience loser regret in cases where her best project is not implemented and overrule regret in cases where she is in the subordinate role and the agent is informed. Analogous to the agent, these forces increase the principal's effort as a controlling party and can lead her to under or over-exert effort after delegation, i.e. when she is the subordinate.

Propositions (1) and (2) and analogous results for the principal thus show that loser regret and overrule regret can rationalize the effort patterns observed in our experimental data. Controlling parties with loser regret will over provide effort relative to the risk neutral best reply of an individual without loser regret. Rational subordinates who anticipate the increased effort of the controlling party will update their beliefs upward (as observed in the data) and have an incentive to reduce their effort relative to the risk neutral Nash equilibrium. In addition, regret averse subordinates with strong enough overrule regret will have an incentive to further decrease their effort below the risk neutral best reply because this reduces overrule regret. If anticipated, this decrease in subordinate effort will further increase effort of the controlling party. Taken together, equilibrium effort provision is expected to be larger for controlling parties and smaller for subordinates if regret aversion exists compared to the risk neutral Nash equilibrium without regret.

It turns out that these same forces can also result in under delegation by the principal. As overrule regret has a negative utility that arises only in the case of delegation, overrule regret decreases the utility of delegation relative to the utility of keeping control. Thus overrule regret can lead to under delegation relative to a standard expected utility maximizer.³⁸

Proposition 3 *Overrule regret decreases the utility of delegation and has no effect on the utility of keeping control. Thus individuals who experience overrule regret may keep control rights even in cases in which expected value comparisons predict delegation.*

Proof. This proposition follows from a direct comparison of the utility for a principal holding control and delegating. ■

Part 3: Regret Due to the Delegation Choice: While we can capture all the main deviations observed in our data with loser regret and overrule regret, a formulation using only

³⁸The effects of loser regret on delegation are more subtle and may go in either direction depending on the efforts chosen by the principal and agent.

these two forces ignores the fact that the principal's decision problem and the agent's decision problem differ in the delegation stage. In order to understand how regret over the delegation decision might affect the principal's decisions, this section extends the model to include regret that might occur due to the principal's delegation choice.

In modeling regret over the delegation choice, we take a direct extension of the baseline model where a principal compares the outcome of his selected delegation and effort decision pair with the decision pair which would maximize his payoff *ex post* given information about the state of nature and beliefs about agent behavior. To ensure consistency, we hold the beliefs about the effort of the agent in the subgame which was not entered constant.³⁹ We also rule out the analogue of winner regret by assuming that the minimum effort that an individual believes she will exert in the counterfactual where she kept control rights is equal to the amount of effort actually exerted after delegation (i.e., we restrict the counterfactual E to be greater or equal to E^d).

We begin by studying the effort decision of a principal who keeps control. In cases where the principal's preferred project is not implemented, a principal has two possible ways in which she might alter her actions to improve her final payoff. First, if she continues to maintain control, the principal can increase effort to $E = x_P$, thus ensuring her preferred project is implemented. Second, if in case of delegation the agent is informed, the principal could instead delegate control to the agent. In this alternative case, the principal's optimal subordinate effort is zero since the informed agent will anyway implement his preferred project, regardless of the principal's recommendation. As a naming convention, we define *delegation regret* as regret which occurs in states where the principal would prefer to change her delegation decision.

In figure (11) we again partition the state space into cells that differ with regard to the regret experienced by the principal. In cell (i), we continue to assume that no regret is experienced if the principal is informed and can choose her preferred project. If the principal remains uninformed, however, she experiences either loser regret or delegation regret. In cell (ii), the agent remains uninformed even if the principal delegates decision rights. In this case, the principal can improve her payoff by increasing her own effort, and therefore she experiences loser regret. In cells (iii) and (v), the principal remains uninformed and the agent is informed, such that project 2 is chosen. In these cases, the principal could have improved her payoff by either keeping the decision right and increasing her own effort to x_P , or by delegating the decision right to the informed agent and choosing zero effort herself. This will depend on the profitability of these alternative strategies, i.e., on whether $P_1 - g_P(x_P)$

³⁹For example, if the principal keeps control, she does not update her beliefs about the effort the agent would have put in she had delegated regardless of the effort observed from the agent in the subordinate role.

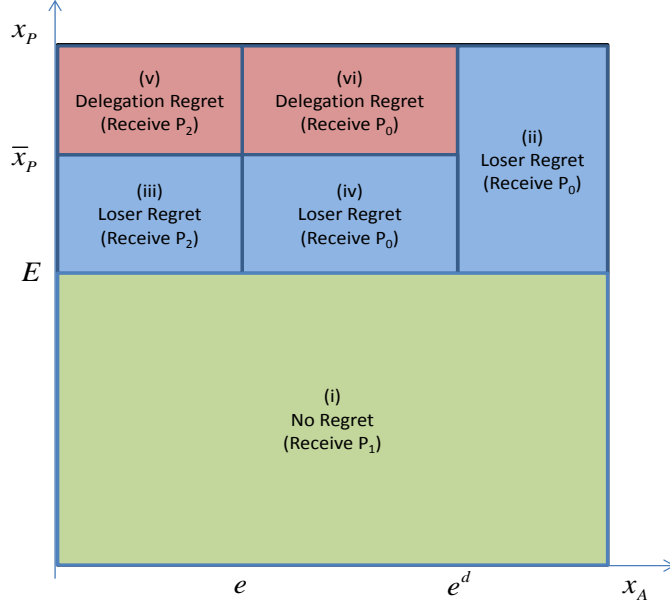


Figure 11: **Principal as Controlling Party:** For principals in the role of the controlling party, the state space can be partitioned into six cells, which differ with regard to the regret experienced by the principal. In cell (i), the principal experiences no regret because her preferred project is implemented. In cell (ii), both the principal and the agent remain uninformed in both subgames. In the remaining cells, the principal is uninformed but the agent is either informed, as in cells (iii) and (v), or would have been informed if delegation had taken place, as in cells (iv) and (vi). In these cells, the principal either regrets his effort choice and experiences loser regret or regrets his delegation choice and experiences delegation regret. This depends on whether it is *ex-post* optimal to keep the decision right and increase effort to $E = x_P$, as in cells (iii) and (iv), or whether it is *ex-post* optimal to delegate the decision right to the agent and choose $E^d = 0$, as in cells (v) and (vi). The cutoff between cells with loser regret and delegation regret if the agent would be informed after delegation depend on whether $P_1 - g_P(x_P)$ is greater or less than $P_2 - g_P(0)$ and is defined by $\bar{x}_P \equiv g_p^{-1}(P_1 - P_2)$.

is greater or less than $P_2 - g_P(0)$. The threshold between these strategies is defined by \bar{x}_P , where $\bar{x}_P \equiv g_P^{-1}(P_1 - P_2)$. In cell (iii), $P_1 - g_P(x_P) \geq P_2 - g_P(0)$, such that the principal prefers to keep the decision right and to increase effort. Therefore, she experiences loser regret. In cell (v), $P_1 - g_P(x_P) < P_2 - g_P(0)$, such that the principal prefers to delegate the decision right and to choose zero effort. Therefore, she experiences delegation regret.

Cells (iv) and (vi) differ from cells (iii) and (v) in that the agent is uninformed as a subordinate and therefore project 0 is implemented. However, if the principal delegates the decision right the agent is informed, and therefore project 2 is implemented. Whether the principal prefers to increase her own effort or to delegate the decision right to the agent and to choose zero effort herself will therefore again depend on which of these two strategies is more profitable, i.e., whether $P_1 - g_P(x_P)$ is greater or less than $P_2 - g_P(0)$. In cell (iv), the principal prefers to raise her effort and hence she experiences loser regret. In cell (vi), the principal prefers to delegate the decision right and to choose zero effort, and hence she experiences delegation regret.⁴⁰

Combining all cells, the utility of a principal in the role of the controlling party is given by $u_P(E|x_A, x_P, \hat{e}, E^d) =$

$$\left\{ \begin{array}{ll} P_1 - g_P(E) & \text{if } x_P \leq E \\ P_0 - g_P(E) - \lambda_{LR}[\hat{P}_1 - g_P(x_P) + g_P(E)] & \text{if } x_P > E \ \¬\ x_A > \hat{e}^d \\ P_0 - g_P(E) - \lambda_{LR}[\hat{P}_1 - g_P(x_P) + g_P(E)] & \text{if } \bar{x}_P \geq x_P > E \ \¬\ \hat{e} < x_A \leq \hat{e}^d \\ P_2 - g_P(E) - \lambda_{LR}[\hat{P}_1 - \hat{P}_2 - g_P(x_P) + g_P(E)] & \text{if } \bar{x}_P \geq x_P > E \ \¬\ x_A \leq \hat{e} \\ P_0 - g_P(E) - \lambda_D[\hat{P}_2 + g_P(E)] & \text{if } x_P > \bar{x}_P \ \¬\ x_P > E \ \¬\ \hat{e} < x_A \leq \hat{e}^d \\ P_2 - g_P(E) - \lambda_D[g_P(E)] & \text{if } x_P > \bar{x}_P \ \¬\ x_P > E \ \¬\ x_A \leq \hat{e} \end{array} \right. \quad (26)$$

where λ_D is the principal's degree of delegation regret.

As can be seen by comparing equations (14) and (26) as well as figures (9) and (11), the principal's utility is similar to that of the agent except that for those realizations of x_P and x_A for which the principal would have preferred to delegate rather than to have increased effort, loser regret is substituted by delegation regret. Note that both forms of regret can be reduced by increasing E and therefore affect the controlling party's effort decision in similar ways. The following remark summarizes the effects of delegation regret on effort:⁴¹

⁴⁰Note that $g_P^{-1}(P_1 - P_2) \leq g_P^{-1}(P_1 - P_2 + g_P(E))$, i. e., $\bar{x}_P \leq \tau(E)$. Hence, unlike the agent, the principal will always experience either delegation regret or loser regret in case she remains uninformed. If $\bar{x}_P > 1$, the principal never experiences delegation regret and always regrets not having invested more effort.

⁴¹The proof for this remark follows directly from the first order condition of the principal's decision problem and is omitted.

Remark 1 *In the effort stage of the authority-delegation game, a principal who anticipates delegation regret and has held decision rights will over exert effort relative to the best response of an individual who maximizes expected value.*

Remark (1) shows that delegation regret has a positive effect on the principal's effort as a controlling party. As with loser regret, the principal attempts to avoid states where he is uninformed in order to reduce the likelihood of regretting his delegation decision.

Finally, we can turn attention to the case of a principal who delegated control. In this subgame, a principal can potentially experience all three forms of regret: loser regret, overrule regret and delegation regret. In figure (12) we again partition the state space into cells that differ with regard to the regret experienced by the principal.

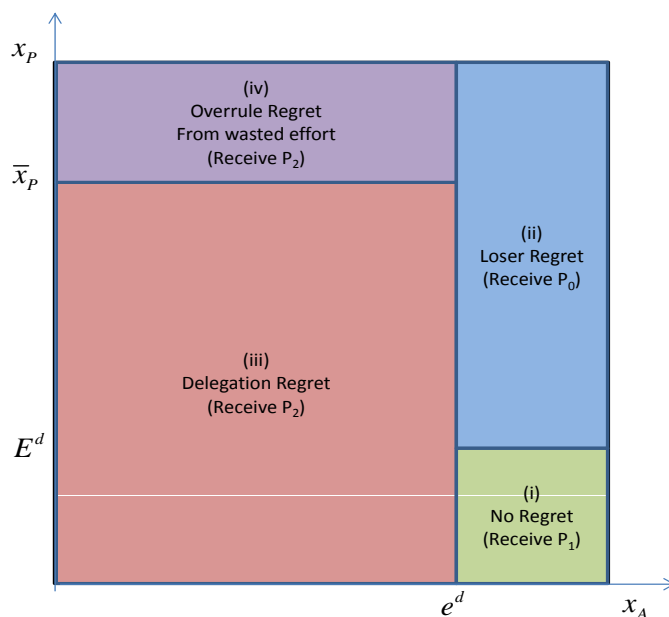


Figure 12: **Principal as Subordinate:** For principals in the role of the subordinate party, the state space can be partitioned into four cells, which differ with regard to the regret experienced by the principal. In cell (i), a principal experiences no regret because only she is informed and therefore her preferred project is implemented. In cell (ii), both parties remain uninformed and therefore the principal experiences loser regret from not having chosen $E^d = x_P$. In cells (iii) and (iv), the agent is informed and implements his preferred project. In these cells, the principal either regrets delegating or regrets his effort choice. Which regret force is felt depends on whether it is *ex-post* optimal to have kept the decision right and exerted $E = \max\{E^d, x_P\}$, as in cell (iii), or to have delegated the decision right and exerted $E^d = 0$, as in cell (iv). The cutoff between cells with delegation regret and overrule regret depend on whether $P_1 - g_P(x_P)$ is greater or less than $P_2 - g_P(0)$ and is defined by $\bar{x}_P \equiv g_p^{-1}(P_1 - P_2)$.

In cell (i), the principal can implement her preferred project and therefore she does not

experience regret. In cell (ii), both the principal and the agent are uninformed. As the principal could have been informed by increasing effort, she experiences loser regret. In cells (iii) and (iv) the agent is informed and thus the agent implements his preferred project. As the principal receives the agent's preferred project, she experiences either delegation regret or overrule regret depending on her *ex-post* optimal strategy. If $x_P \leq \bar{x}_P = g_p^{-1}(P_1 - P_2)$, the principal's *ex-post* optimal action is to keep the decision right and choose $E = \max\{E^d, x_P\}$, where the "max" comes from the assumption that the principal never expects to exert less effort with held control rights than after delegation. If $x_P > \bar{x}_P$, however, the principal's *ex-post* optimal action is to continue to delegate and choose zero effort. Therefore, in cell (iii), the principal experiences delegation regret and in cell (iv) the principal experiences overrule regret.

Combing the cells into a single utility function, the utility of a principal in the role of the subordinate is given by $U_P^d(E^d|x_A, x_P, \hat{e}^d, E) =$

$$= \begin{cases} P_1 - g_P(E^d) & \text{if } x_P \leq E^d \ \& \! \! \! \& \! \! \! x_A > \hat{e}^d \\ P_0 - g_P(E^d) - \lambda_{LR}[\hat{P}_1 - g_P(x_P) + g_P(E^d)] & \text{if } x_P > E^d \ \& \! \! \! \& \! \! \! x_A > \hat{e}^d \\ P_2 - g_P(E^d) - \lambda_D[\hat{P}_1 - \hat{P}_2 - g_P(\max\{E^d, x_P\}) + g_P(E^d)] & \text{if } x_P \leq \bar{x}_P \ \& \! \! \! \& \! \! \! x_A \leq \hat{e}^d \\ P_2 - g_P(E^d) - \lambda_{OR}g_P(E^d) & \text{if } x_P > \bar{x}_P \ \& \! \! \! \& \! \! \! x_A \leq \hat{e}^d \end{cases} \quad (27)$$

Comparing equations (27) and (23) as well as figures (10) and (12), it can again be seen that overrule regret is substituted by delegation regret whenever delegation regret is of larger magnitude than overrule regret. This implies that the effort choice of the principal as a subordinate is increasing or decreasing relative to the standard best response, depending on the strength of loser regret and the combined strength of overrule and delegation regret.

Remark 2 *Principals in the subordinate role either experience delegation regret or overrule regret when the controlling party is informed. Anticipation of both forms of regret will decrease effort relative to those who do not. Individuals who anticipate loser regret will increase effort relative to those who do not. As these forces move in different directions, heterogeneity in anticipatory regret may lead to observed effort choices both above and below the best response.*

We now turn attention to the effects of delegation regret on the delegation decision. Delegation regret further reduces the utility of delegation since the principal will ex-post experience delegation regret in a multitude of states. Delegation regret may also reduce the utility in case of kept control, since the principal may also regret not having delegated

ex-post. However, as we explain below, it seems plausible that in our experiment regret after delegation played a more important role.

A principal who delegates observes x_P ex-post and thus knows with certainty whether she would have had a better outcome had she kept decision rights. She also directly experiences her recommendation being overruled, which may be particularly salient. By contrast, the principal is never informed of x_A . If a principal keeps control and experiences the agent recommending the outside option (which indicates that the agent is not informed) she does not know whether the agent would have been informed if she had delegated. It therefore seems reasonable to assume that the experience of delegation regret after the principal kept control is much less salient than the delegation regret experienced after the principal delegated and was informed. If this was the case, delegation regret is likely to have reduced the incentive to delegate.⁴²

⁴²One way to account for these saliency differences in the model might be to allow for different degrees of delegation regret λ_D in the delegation and in the no-delegation subgames. To avoid further notation, however, we abstracted from this differentiation in this appendix.

Appendix B: Session Overview

Table B.1: Session Overview

Date	Treatment	Subjects	Matching Groups	Periods
Main Treatments				
May 2008	PLOW	30	3	10
May 2008	PLOW	30	3	10
May 2007 ¹	LOW	12	1	10
May 2007	LOW	30	3	10
May 2008	LOW	30	3	10
May 2007 ¹	HIGH	10	1	10
May 2007	HIGH	30	3	10
June 2007	HIGH	28	3	10
Oct 2008 ²	HIGH	30	3	10
June 2007	PHIGH	30	3	10
May 2008	PHIGH	30	3	10
Control Treatments				
Oct 2008	HIGH RAND	30	3	10
May 2009	HIGH RAND	30	3	10
May 2009	HIGH RAND	30	3	10
April 2011	PHIGH50	32	2	50
April 2011	PHIGH50	32	2	50
April 2011	HIGH NOREC	28	2	25
April 2011	PHIGH25	32	2	25

¹ This session was split into two matching groups with different treatments.

² This session did not use the strategy method for eliciting agent effort.

Appendix C: Additional Tables

Table C.1: Average effort levels vs. Nash predictions across treatments

	Controlling Party						Subordinate					
	Principal			Agent			Agent			Principal		
	E	E^{NE}	e^d	e^d	e^{dNE}	e^{dNE}	e	e^{NE}	E^d	E^d	E^{dNE}	E^{dNE}
PLOW	55.7		55	68.1	***	45	22.8		25	16.5	***	35
LOW	66.1	***	55	68.3	***	55	14.3	***	25	16.2	**	25
HIGH	48.2	*	45	58.7	***	45	26.5	***	35	19.6	***	35
PHIGH	58.2	***	45	65.1	**	55	17.3	***	35	20.7		25

Significance Levels for Wilcoxon Signed-Rank Tests against Nash predictions with data averaged by individual prior to estimation. Significance Levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Table C.2: Overall profit of principals and agents by treatment

	Principals		Agents	
	Actual ^a	Predicted ^b	Actual ^a	Predicted ^b
PLOW	18.23	20.1	22.35	25.6
LOW	18.40	20.1	16.32	17.3
HIGH	21.13	24.0	20.69	23.3
PHIGH	21.89	25.6	16.83	20.1

^aActual earnings in treatment.

^bPredicted earnings with Nash equilibrium effort and delegation.

Table C.3: Delegation decisions by principals

	(1)	(2)	(3 ^a)	(4 ^b)
PLOW	0.035 (0.068)	0.061 (0.073)	0.106 (0.080)	
HIGH	0.245*** (0.061)	0.310*** (0.097)	0.462*** (0.153)	
PHIGH	0.326*** (0.085)	0.356*** (0.118)	0.503*** (0.160)	0.003 (0.144)
Belief if subordinate		-0.003** (0.001)	-0.002 (0.001)	-0.005 (0.003)
Belief if controlling party		0.003*** (0.001)	0.006*** (0.002)	0.012*** (0.004)
# of Lotteries Declined			-0.062** (0.026)	-0.120** (0.052)
Period Dummies?	Yes	Yes	Yes	Yes
Pseudo. R^2	.062	.112	.179	.176
Observations	1450	1450	750	300

Marginal effects of a probit regression. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard error in parentheses, clustered by individual. ^a Regret aversion measures are available only for sessions conducted in 2008-2011. ^b Column (4) includes data only from the HIGH and PHIGH treatments for which we have regret aversion measures, and HIGH is the omitted category.

Table C.4: Average effort levels vs. average beliefs across treatments

	Controlling Party Effort						Subordinate Effort					
	Principal			Agent			Agent		Principal			
	has control			has control			has control		has control			
	E	**	\hat{E}	e^d	**	\hat{e}^d	e	\hat{e}	E^d	\hat{E}^d		
PLOW	55.7	**	64.8	68.1	**	59.6	22.8	30.4	16.5	21.8		
LOW	66.1		66.9	68.3		68.4	14.3	***	27.5	16.2	*	20.9
HIGH	48.2	***	59.0	58.7		56.2	26.5	**	35.8	19.6	*	29.4
PHIGH	58.2	**	69.3	65.1		62.3	17.3	**	28.2	20.7		19.0

Significance levels calculated using a Wilcoxon Rank-Sum test with beliefs and effort averaged by individual prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$. E is the principals' average effort with control. \hat{E} is the agents' average belief about principals' effort with control. e^d is the agents' average effort with control. \hat{e}^d is the principals' average belief about agents' effort with control. e is the agents' average effort in the subordinate role. \hat{e} is the principals' average belief about agents' effort in the subordinate role. E^d is the principals' average effort in the subordinate role. \hat{E}^d is the agents' average belief about principals' effort in the subordinate role.

Appendix D: Loss Aversion and Effort

In discussing the effort provision of a loss averse individual, we made the intuitive argument that loss aversion cannot explain the observed effort choices of the controlling party. This appendix shows that a controlling party who is loss averse will never choose effort which is above 60 but below 100. To simplify the equations, we follow the theory section and express all effort choices in decimal form (i.e., an effort of 60 is expressed as .6).

Following Koszegi & Rabin (2006), we assume that subjects have a utility function of the following form:

$$v(x) = \begin{cases} x - R & \text{if } x \geq R \\ (1 + \lambda)(x - R) & \text{if } x < R \end{cases}, \quad (28)$$

where $\lambda \geq 0$ denotes the degree of loss aversion and R denotes the reference point. A natural reference point is $R = 10$, the value of project P_0 in each experiment. Recall that if subjects provide zero effort, they can always ensure a payoff of $P_0 = 10$ by choosing the known outside option. Also recall that \hat{e} is the belief of the principal about the effort of the agent when she is the controlling party. We begin by proving the following:

Lemma 1 *Let $E^*(\lambda, \hat{e})$ be a local maximum of the principal's utility maximization problem when she is the controlling party with loss aversion λ and beliefs \hat{e} . Then $E^*(\lambda, \hat{e})$ is decreasing in loss aversion if $E^*(0, \hat{e}) < .65$.*

Proof. If $E < 0.65$, the cost of effort is below 10. Given the parameters in the authority game, this implies that losses relative to the reference point can only occur in the case that both the controlling party and the subordinate remain uninformed. We use this fact to circumvent non-differentiability around the reference point by restricting analysis to this region. The optimization problem of the principal when she is the controlling party is

$$\begin{aligned} \max_E U(E) = & E(P_1 - R - g_P(E)) + (1 - E)\hat{e}(P_2 - R - g_P(E)) \\ & - (1 + \lambda)(1 - E)(1 - \hat{e})(P_0 - R - g_P(E)). \end{aligned} \quad (29)$$

By assumption $R = P_0$, which implies that the corresponding first order condition is:

$$\begin{aligned} U'(E) = & (\hat{P}_1 - g_P(E)) - E g'_P(E) - \hat{e}(\hat{P}_2 - g_P(E)) - g'_P(E)\hat{e}(1 - E) - \\ & (1 + \lambda)(1 - \hat{e})[(g_P(E)) - g'_P(E)(1 - E)] = 0. \end{aligned} \quad (30)$$

Rearranging this equation and replacing $g'_P(E)$ and $g_P(E)$ and \hat{P}_1 with their values which

were constant across treatments yields:

$$U'(E) = -50E + 30 - \hat{e}\hat{P}_2 + 50\lambda(1 - \hat{e})E \left[\frac{3}{2}E - 1 \right] = 0. \quad (31)$$

Writing 31 as an implicit function, the FOC is satisfied when:

$$E = \frac{30 - \hat{e}\hat{P}_2}{50} + \lambda(1 - \hat{e})E \left[\frac{3}{2}E - 1 \right]. \quad (32)$$

The last term is negative for $E \in [0, \frac{2}{3}]$ and $\lambda > 0$. Thus, effort is decreasing in λ for all $E^*(0, \hat{e}) < .65$ (our initial condition for the considered case). ■

We now prove our main result:

Proposition 4 *Effort of a loss averse individual will never be above 60 but below 100.*

Proof. Equation 31 can be rewritten as follows:

$$U'(E) = 75\lambda(1 - \hat{e})E^2 - 50[1 + \lambda(1 - \hat{e})]E + 30 - \hat{e}\hat{P}_2 = 0. \quad (33)$$

Note that this equation is quadratic and thus has two roots. Taking the second derivative of U with respect to E we have:

$$U''(E) = 150\lambda(1 - \hat{e})E - 50[1 + \lambda(1 - \hat{e})]. \quad (34)$$

Thus, there is a unique inflection point at $E = \frac{1}{3} \frac{1 + \lambda(1 - \hat{e})}{\lambda(1 - \hat{e})}$. The second derivative is negative to the left of this reflection point and positive to the right of this inflection point.

Solving the quadratic equation, E is a local maxima/minima at:

$$\frac{50[1 + \lambda(1 - \hat{e})] \pm \sqrt{Z(\lambda)}}{150\lambda(1 - \hat{e})}, \quad (35)$$

where $Z(\lambda) = 2500[1 + \lambda(1 - \hat{e})]^2 - 300\lambda(1 - \hat{e})[30 - \hat{e}\hat{P}_2]$. Also note that $Z(\lambda)$ is always greater than 0 so both roots exist. Comparing this to the inflection point, the left root is the local maximum. Next, using L'Hôpital's rule,

$$E^*(0, \hat{e}) = \lim_{\lambda \rightarrow 0} \frac{50[1 + \lambda(1 - \hat{e})] - \sqrt{Z(\lambda)}}{150\lambda(1 - \hat{e})} = \frac{[30 - \hat{e}\hat{P}_2]}{50} \leq .6 \quad (36)$$

By lemma 1, it follows that this unique local maximum is decreasing in loss aversion. As

the unique local maximum is always below 60 and $E \in [0, 100]$, it follows that the global maxima are either below 60 or at the boundaries of $E = 0$ and $E = 100$. ■

Appendix E: Risk Aversion and Effort

In discussing the effort provision of a risk averse individual, we made an informal argument as to why risk aversion and risk lovingness cannot account for the effort provisions of the controlling party. This appendix provides numeric support for this argument for the case of CRRA utility. To simplify the equations, we follow the theory section and express all effort choices in decimal form (i.e., an effort of 60 is expressed as .6).

Recall that a controlling principal with belief \hat{e} about the effort of the subordinate and a concave utility function has an expected utility of

$$U(E) = Eu(P_1 + w - g_P(E)) + \hat{e}(1 - E)u(P_2 + w - g_P(E)) + (1 - \hat{e})(1 - E)u(P_0 + w - g_P(E)) \quad (37)$$

where w is wealth, $P_1 = 40$, $P_2 \in \{35, 20\}$, $P_0 = 10$, $g_P(E) = 25E^2$, and $\hat{e} \in \{0, .05, \dots, 1\}$. As can be seen by studying the arguments on the right hand side of this equation, increasing effort has two effects. First, an increase in effort increases the probability of winning the highest valued gamble which strictly increases utility. Second, increasing effort decreases the utility earned for each of the three possible outcomes. As this second effect necessarily depends on the marginal utility of three separate points, it is easy to construct cases in which locally, effort is increasing in risk aversion. Such local non-monotonicity makes analytic analysis both tedious and unenlightening, particularly for extremely concave utility or those which do not satisfy decreasing relative risk aversion.

As the decision problem of the controlling party is inherently discrete, we take a more direct approach to determining the potential effect of risk aversion on effort. Starting with common parameterized risk aversion utility functions such as CRRA and CARA, we find the risk aversion parameters which maximize effort and then compare these effort levels to the risk neutral baseline.

As with loss aversion, there is potential that an extremely risk averse controlling party will choose an effort of 100 and ensure themselves P_1 . As a first step of the analysis, we start by finding the lowest σ for which an individual with a CRRA utility will choose an effort of

1. Let

$$E(\sigma, \hat{e}) = \arg \max_E Eu(P_1+w-g_P(E))+\hat{e}(1-E)u(P_2+w-g_P(E))+(1-\hat{e})(1-E)u(P_0+w-g_P(E)) \quad (38)$$

be the optimal effort of an individual with CRRA utility of the form $u(x) = \frac{x^{1-\sigma}}{1-\sigma}$ where $w \geq 16$ so that utility is always well defined. Next, define σ_1 to be the smallest risk aversion parameter such that $E(\sigma_1, \hat{e}) = 1$. It can be shown analytically that $E(\sigma, \hat{e}) = 1$ for all $\sigma > \sigma_1$ and thus that σ_1 is a sufficient statistic for the parameter space where full effort is predicted.

Our interest in risk aversion lies in being able to predict effort levels above the risk neutral prediction but below an effort of 1. The next step of our analysis is to look at the maximum possible effort which can be predicted for all $\sigma \in [-\infty, \sigma_1)$. Let

$$\sigma^*(\hat{e}) = \arg \max_{\sigma \in [-\infty, \sigma_1)} E(\sigma, \hat{e}) \quad (39)$$

and define $E(\sigma^*(\hat{e}), \hat{e})$ as the effort level which corresponds to $\sigma^*(\hat{e})$. For all initial beliefs, we find $E(\sigma^*(\hat{e}), \hat{e})$ and compare this to $E(0, \hat{e})$, the effort predicted when an individual is risk neutral.

Table E.1: Maximum effort predicted by risk aversion

\hat{e}	<i>Low Treatment</i>				\hat{e}	<i>High Treatment</i>			
	$\sigma^*(\hat{e})$	σ_1	$E(\sigma^*, \hat{e})$	$E(0, \hat{e})$		$\sigma^*(\hat{e})$	σ_1	$E(\sigma^*, \hat{e})$	$E(0, \hat{e})$
0	-0.7 - 0.6	1.2	60	60	0	-0.7 - 0.6	1.2	60	60
10	-0.3 - 0.3	1.4	60	60	10	-.9 - 0.7	1.6	55	55
20	-1.3 - 0.7	1.6	55	55	20	-1.1 - 0.7	2.2	50	50
30	-0.9 - 0.5	2.0	55	55	30	-1.2 - 0.8	2.9	45	45
40	-2 - 0.9	2.5	50	50	40	-1.4 - 0.9	3.8	40	40
50	-1.8 - 0.6	3.2	50	50	50	-1.5 - 1.3	5.0	35	35

Table E.1 reports $\sigma^*(\hat{e})$, σ_1 , as well as $E(\sigma^*(\hat{e}), \hat{e})$ and $E(0, \hat{e})$ for initial beliefs \hat{e} in intervals of 10. As can be seen, $\sigma^*(\hat{e}) < 0$ for all initial beliefs revealing that an individual who is slightly risk loving will provide the highest effort. As can be seen in the last two columns of the table, however, the increase in effort for these individuals is not large enough to alter the effort predictions.

As we typically are most interested in small amounts of risk aversion, it is useful to also look at σ in the domain of $[0, \sigma_1)$. For all wealth and beliefs, it is the case that effort is

maximal in this domain when $\sigma = 0$.

Just as with loss aversion, effort provision under risk aversion has a difficult time explaining effort levels above the risk neutral prediction. For all $w \geq 16$, all beliefs \hat{e} , and using both CRRA and CARA utility, it is never the case that $E(\sigma^*(\hat{e}), \hat{e}) - E(0, \hat{e}) > 5$. As 50 percent of our data lies 15 points above the risk neutral prediction, we cannot rationalize the over-provision of effort by the controlling party with risk preferences.