Department of Economics

Working Paper Series

The Lure Authority:
Motivation and Incentive Effects of Power

Ernst Fehr, Holger Herz &
Tom Wilkening

Dec 2010

Research Paper Number 1115

ISSN: 0819-2642
ISBN: 978 0 7340 4468 6

Department of Economics
The University of Melbourne
Parkville VIC 3010
www.economics.unimelb.edu.au

The Lure of Authority: Motivation and Incentive Effects of Power

Ernst Fehr and Holger Herz and Tom Wilkening*

July 22, 2010

Abstract

Authority and power permeate political, social, and economic life but there is limited empirical knowledge about the motivational origins and consequences of authority. We experimentally study the motivation and incentive effects of authority in an authority-delegation game. Individuals exhibit a strong tendency to retain authority even when its delegation is in their material interest — suggesting that they value authority per se. Moreover, this tendency to hold on to authority strongly increases with individuals' degree of loss aversion, suggesting an endowment effect with regard to authority. Authority also leads to a substantial over provision of effort by the controlling party, while a large percentage of subordinates under provide effort despite pecuniary incentives to the contrary. Thus, authority 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: Institute for Empirical Research in Economics, Zurich University, Bluemlisalpstrasse 10, CH-8006 Zurich. E-mail: efehr@iew.uzh.ch. Holger Herz: Institute for Empirical Research in Economics, Zurich University, Bluemlisalpstrasse 10, CH-8006 Zurich. E-mail: hherz@iew.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, Jon Elster, Sergei Izmalkov, Nikos Nikiforakis, Chris Smith, Participants at the ESA International Meetings 2008 at Caltech, Participants at the 24th annual congress of the European Economic Association, the Experimental Workshop at Zurich University, the 2nd Australian Workshop for Experimental Economics and the 2nd Behavioral Economics Conference at Cornell for helpful suggestions and feedback.

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 and 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 in which one party has the contractual right to make decisions that influence the payoffs and potential choices of another (Grossman and Hart (1986); Hart and Moore (1990); Aghion and Tirole (1997); Baker, Gibbons and Murphy (1999); Dessein (2002); Aghion, Dewatripont and Rey (2004)). By granting decision rights, inefficiencies can be eliminated 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 in which we study how individuals manage and respond to authority in a hierarchical relationship. We propose a new "authority-delegation game" based on the model of authority developed in Aghion and 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 authority which implies that she 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.

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 learn the value of each project. One of the projects is best for the principal, while a different project is best for the agent. Relative to first best, this conflict of interest leads to under provision of effort by the agent since his suggestions may be overruled when both parties are 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. Delegation 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 his expected payoff should thus delegate authority in these cases. When a principal's return from the agent's preferred project is low, however, the cost of losing authority is high and a rational, expected payoff maximizing principal should retain authority.

This setting neatly captures the notion of power as defined by Max Weber.[1] According to Weber an actor has power in a social relationship if he "is in the position to carry out his own will despite resistance" (Weber (1978, p. 53)). In the setting of Aghion and Tirole (1997), the controlling party — which is the agent in the case of delegation, and the principal in the case of retention — has power over the subordinate party because the controlling party can overrule the subordinate party's project preference. In the experiment, we find that the controlling parties use their power extensively. If they are informed about the project valuations they almost always overrule the subordinate's recommendation.

Our first main result is that the principals show a strong preference for retaining authority in situations in which they could substantially improve their expected income by delegating authority to the agents. More specifically, rational principals who maximize their expected income should delegate authority in those treatments where the principal's return from the agent's preferred project is relatively high. However they only do so 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 to keep authority.

These findings suggest that the principals might view authority not just as an instrument that helps them increase their earnings, but that they value the decision right (i.e. authority) per se. Several psychologists have postulated a preference for power (McClelland (1975), Mulder (1975), Poppe (2003)) or a preference for agency, autonomy and self determination (Rotter (1966); deCharms (1968); Deci (1981); van Dijk and Poppe (2006)). To our knowledge, however, no evidence yet exists that proves a willingness to pay (i.e. a preference) for power, agency, or autonomy.[2] In our experiment, the fact that the principals are willing

[1] We view authority as a form of power in which the controlling party has the right to make decisions which directly affect the payoffs of another individual who has no means of successfully resisting the controlling party's decision. Aghion and Tirole used the term "authority" to describe the hierarchical relationship between the controlling and the subordinate party, but the relation between the two parties can also be characterized as a power relationship in Max Weber's sense.

[2] A preference for power can explain why principals hesitate to delegate authority. Also, being subject to the choice of a controlling party may be viewed as a constraint on autonomy; if the principal values autonomy positively, he has a reason to avoid delegation. Interestingly, Adam Smith stipulated the existence of a preference for "domination and authority" in his lectures on jurisprudence (Smith 1978, p. 186): "yet the love of domination and authority and the pleasure men take in having every<thing> done by their express orders, rather than to condescend to bargain and treat with those whom they look upon as their inferiors and are inclined to use in a haughty way; this love of domination and tyrannizing, I say, will make

to sacrifice some of their earnings to keep authority indicates a willingness to pay for the decision right.

If it is indeed the case that decision rights are valued per se, then we might also observe an endowment effect in decision rights. There is evidence (Knetsch (1989), Kahneman, Knetsch and Thaler (1990)) that subjects have a higher valuation for goods if they are randomly endowed with them compared to a situation in which they have to buy them. One important potential reason for the tendency to assign higher valuations to owned goods is loss aversion (Knetsch (1995)); the tendency of losses to loom larger than gains. In our experiment, the subjects were randomly assigned the role of the principal, i.e. principals were randomly endowed with decision rights. Thus, if there is an endowment effect in decision rights, then those principals who are more loss averse should show a greater reluctance to delegate authority. We indeed find that loss aversion has a large and significant effect on subjects' reluctance to delegate. Subjects with a degree of loss aversion above the median are 20 percentage points less likely to delegate than those with below-median loss aversion. Moreover, subjects assigned to the principal's role expressed a much stronger preference for the role of the controlling party in an exit survey after the experiment than subjects who were assigned to the role of the agent. Thus, despite the fact that the principals earned less money during the experiment when they were the controlling party, the large majority of the principals expressed a preference for being the controlling party.

Our second main finding is that the controlling party substantially over provides effort relative to the Nash equilibrium and relative to his best response to the subordinate party's anticipated effort. We show that loss aversion cannot explain this behavior. Furthermore, neither deviations from risk neutrality nor social preferences can explain the controlling party's excess effort. It thus seems that the mere fact of having authority has strong motivational effects on effort choices.

Our third main finding is that, relative to the Nash equilibrium, subordinates substantially under provide effort. In fact, a substantial minority of the subordinate parties (between 30 percent and 50 percent across various conditions) choose 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 strong demotivating effect on a substantial minority of the subordinate parties.

Our paper is related to the experimental literature on the consequences of delegation on punishment choices (Bartling and Fischbacher (2008), Coffman (2010)), the willingness to behave selfishly (Weber, Hamman and Loewenstein (2010a)), on public goods provision

it impossible for slaves in a free country ever to recover their liberty". We owe this reference to John Elster.

(Weber, Hamman and Woon (2010b)) and on the hidden costs of control (e.g. Falk and Kosfeld (2006), Charness, Cobo-Reyes, Jimenez, Lacomba and Lagos (2009)). However, none of these papers examines and identifies the reluctance to delegate in the context of the optimal allocation of decision rights.

We believe that our results have potentially important implications across many domains. In relation to the property rights literature (Grossman and Hart (1986); Hart and Moore (1990)), if people value decision rights per se, it may be difficult to (re)allocate authority in organizational hierarchies to the benefit of the organization because even if organization members with authority would benefit economically from delegation, they may oppose it. In one of our treatments, the under delegation of authority not only reduces the principals' earnings, but also causes the agent to lose money. 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 and Krahmer (2008) which predicts limits to delegation in environments with limited commitment, dynamic incentives, or limited liability.

A reluctance to delegate power may also play a role in both corporate finance and the political sphere. Models of empire-building investment (Jensen (1986), Hart and Moore (1995)) which have been used extensively in the literature to understand the trade-offs between financial instruments may, in part, be founded on a desire for power. Similarly, the taste for power may provide a rationale for term limits because otherwise politicians may try to keep their political power positions beyond what is good for the polity. In addition, the desire for power may also 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 much larger than the standard model predicts; a reallocation of authority causes much larger effort increases for the new controlling parties and may cause a much larger effort reduction for the previously controlling party. The noteworthy gap between the controlling and the subordinate parties' efforts also implies that the efficiency costs of authority are likely to be higher than predicted by the standard model, as, in the presence of strictly convex (and identical) effort cost functions, the first best effort allocation requires effort to be identical across parties. Additionally, the result that a lack of authority only seems to demotivate a minority of people strongly suggests that putting the right people into positions that lack authority is important.[3] The development of tools for detecting this type of employee may thus be important in minimizing the cost associated with the (re)allocation of authority.

Our results on effort behavior are related to the psychological literature on the consequences of power (Anderson and Berdahl (2002); Anderson, Gruenfeld and Keltner (2003); Anderson and Galinsky (2006)). These studies put forward the hypothesis that power induces approach-related behaviors while lack of power causes inhibition-related behaviors. According to this view, approach-related behaviors focus on potential gains in risky situations, while inhibition-related behaviors focus on the downside risk. In our setting, this hypothesis implies that the controlling party is strongly focused on the large payoff that accrues if the party can choose the preferred project. The controlling party will thus tend to provide a high effort in order to make sure that the preferred project can be identified. In contrast, the subordinate party focuses on the worst case, given by a high effort and a complete lack of information about which projects yield a positive return.[4] By reducing the effort to very low levels, the subordinate can improve the payoff in this worst case, which may provide a reason for many subordinates' low effort levels. Our effort pattern is thus compatible with this approach/inhibition hypothesis, but we also believe that more research is needed to identify the motivational forces behind the effort choices more precisely.[5]

Despite the systematic deviations from the predictions of the Aghion and 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 (Rantakari (2008), Dessein (2002)), monetary incentives (Athey and Roberts (2001)), and dynamic learning (Aghion et al. (2004)) are possible.

[3]We find strong evidence that the vast majority of subordinates consists of two types. Most subordinates display either persistent underprovision or persistent overprovision of effort relative to the best reply.

[4]In the experiment, the party either knew the valuations of all available projects or of none. In the worst case neither the principal nor the agent knows the project valuations.

[5]The subordinates' low effort levels may, for example, also be a consequence of their distaste for the constraints on the project choice. As in Falk and Kosfeld (2006) they may react to a constraint on their action space with a reduction in effort. There is, however, a decisive difference between their result and ours. While providing the lowest feasible effort level was always in the agent's material self-interest in Falk and Kosfeld (2006), the subordinate hurts himself if he chooses a zero effort level in our setting. The principal's constraint on the agent's action space thus makes agents more selfish in Falk and Kosfeld. In our paper, the lack of authority induces the agents to make choices against their material self interest. This behavior resembles the "discouraged worker effect" which describes workers who have ceased to search for a job. Perhaps, the fact that subordinates can be overruled generates a kind of "discouraged subordinates' effect".

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 into account the behavioral forces observed in our paper 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 and 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 Nash equilibrium. Section 4 concludes.

1 The Model

The basis of our experimental design is a model of authority developed in Aghion and 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 project out of a set of $n \geq 3$ potential projects. With each project $k \in \{1, \ldots, n\}$, 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 would suggest, 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 her 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. To impose a unique interior solution in the current general form, we assume that both g_A and g_P are strictly convex, satisfy $g_i(0) = 0$ and $g_i'(0) = 0$ for $i = A, P$ and that $A_2 - g_A(1) \leq A_0$, and $P_1 - g_P(1) \leq P_0$. These last assumptions imply that the principal and agent would rather accept the outside option with certainty than guarantee themselves their preferred project.[6]

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, the principals utility is $P_k - g_P(E)$. The agents utility is $A_k - g_A(e)$. As the benefit to the principal is non-contractible, 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.

Formal authority is defined as having the right to make the final decision. We analyze two cases, a P-Formal authority structure in which the principal maintains decision rights and an A-Formal authority structure in which the principal delegates decision rights to the agent. In the P-Formal case, a principal may always overrule the agent. She does so if she is informed and if the agent's recommendation is not the principal's preferred project. Otherwise, she (optimally) rubber-stamps the agent's proposal any time she is not informed since $P_2 > P_0$.[7] Under A-Formal authority, the principal delegates decision rights to the agent, giving him the irrevocable right to make the project choice.

[6] In the experiment, we depart from this assumption and ensure uniqueness by imposing linearity in the reaction functions.

[7] Aghion and Tirole (1997) refer to this case as the agent having real authority.

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 under both authority structures 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 under P-Formal authority, the utilities of a risk-neutral principal and agent are

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

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

where

$$\hat{P}_i = P_i - P_0, \;\; \textit{for } i \in \{1,2\}, \tag{3}$$

$$\hat{A}_i = A_i - A_0, \;\; \textit{for } i \in \{1,2\}. \tag{4}$$

Under A-Formal authority, the utility of the principal and agent are

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

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

where the d in the A-Formal utility functions stands for the mnemonic *delegation*.

From Equations 1 and 2, the reaction functions under P-Formal authority are the solutions to:

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

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

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. In order to ensure a unique and stable intersection of the reaction function in (E, e)-space, we assume that the absolute value of the agent's reaction function (given by $\frac{g''_A}{\hat{A}_2}$) is larger than the absolute value of the principal's reaction function (given by $\frac{\hat{P}_2}{g''_P}$).

Looking at the reaction functions under P-Formal authority with $\hat{A}_2$ and $\hat{P}_1$ constant, the principal's effort is decreasing in $\hat{P}_2$. The more value that the principal receives at the agents preferred project, the less incentive he has to get informed himself to overrule the agent. This leads to an increase in effort by the agent since it is more likely that the project he suggests will be implemented. Notice also that neither the principals' nor the agents' best response function depends on $\hat{A}_1$, the agents valuation under the principals best project.

The same logic applies under A-Formal authority, where the reaction curves of the principal and agent are the solutions to:

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

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

and denoted by $r^d_P(e^d)$ and $r^d_A(E^d)$. As in the case of P-Formal authority, the reaction functions are downward sloping in (E^d, e^d) space. Under similar stability and uniqueness criteria, there exists an interior intersection of reaction functions, $(e^{d^{NE}}, E^{d^{NE}})$, which constitutes the Nash equilibrium. Given the symmetry in the best response functions, it is unsurprising that under A-Formal authority, the agent's effort is decreasing in $\hat{A}_1$, the principal's effort is increasing in $\hat{A}_1$, and $\hat{P}_2$ does not affect the equilibrium effort choices.

A careful examination of the reaction functions under P-Formal authority and A-Formal authority reveals that the principal decreases her effort moving from P-Formal authority to A-Formal authority while the agent increases his effort. Delegation by the principal thus has two effects in our model: 1) a cost saving effect since delegation reduces the equilibrium effort of the principal, 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.[8] 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. These cards ensure that individuals prefer to implement the outside option relative to implementing a project at random. The task of each principal-agent pair is to select a card which will be used for payment.

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 agent is binding and irreversible.

In the second stage, subjects choose their effort levels simultaneously and in private.[9] Both subjects select their effort in increments of 5 from $\{0, 5, \ldots, 95, 100\}$. 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:

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

Subjects are presented information on the cost of effort in a table where each possible effort and its associated cost is displayed. 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.[10] Thus, agents choose their effort levels for both

[8] Subjects are randomly assigned the role of a principal or of an agent and remain in this role throughout the experiment.

[9] In the experiment we refer to effort as "search intensity".

[10] We test whether the strategy method influences our results by comparing results to treatments that use a standard elicitation method. We find no difference across treatments. P-Values of a Kolmogorov-Smirnov test whether the distribution of agent effort is identical in treatments with and without strategy method are

authority structures 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 structure followed by beliefs for the counterfactual. For agents, beliefs for both authority structures are elicited simultaneously. To prevent hedging, no incentives are used in the elicitation of beliefs.[11] In the fourth stage, agents are informed about whether principals kept or transferred decision rights. Then, given a subject's effort in the chosen authority structure, 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.

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.

2.2 Experimental Design and Hypotheses

The experimental design involves four treatments implemented in a between-subject design. Treatments vary in the amount that principals and agents are paid for the selection of the project preferred by the other party ($\hat{P}_2$ and $\hat{A}_1$). By changing the payoff given to the other party, the level of incentive conflict in the environment is changed, which, as shown 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

0.70 for effort with decision rights and 0.72 for effort without decision rights. Delegation frequencies differ by 1.6 percent. This difference is also not significant (p=0.62 in a Fisher's exact test)

[11]See Blanco, Engelmann, Koch and Normann (2008) for a discussion of hedging

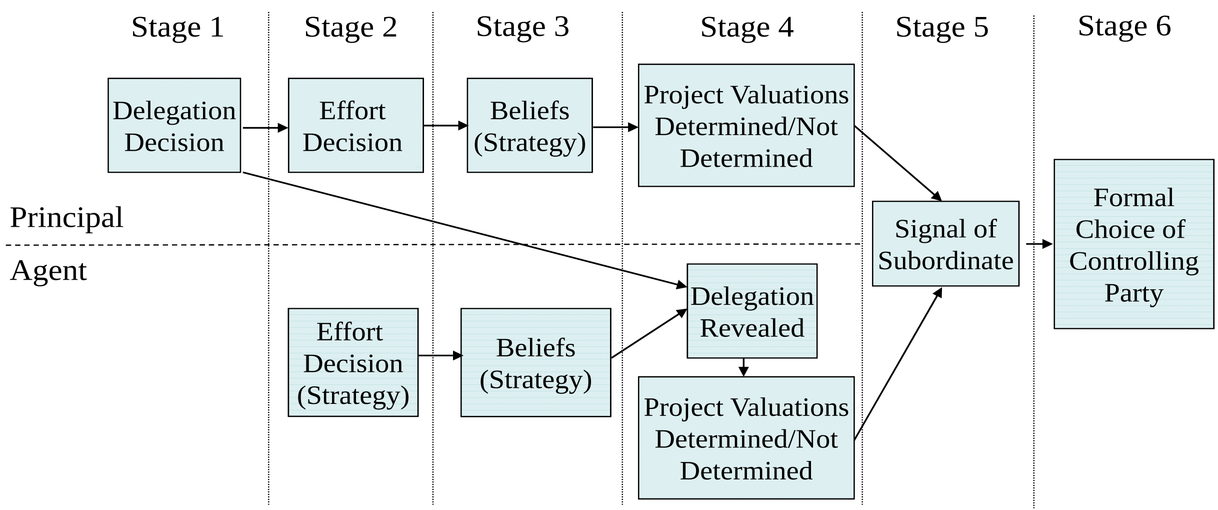

Figure 1: Experimental procedures in the authority game

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**	**Other**
	Principal	Agent	Principal	Agent	**Option**	**Projects**
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 (P-Formal authority) and transferred (A-Formal authority). 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 this table, the LOW treatment has a high degree of incentive conflict and authority should be kept by the principal, because the principals expected profit under P-Formal, EV_P, is 20.1 while the expected payoff under A-Formal, 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. $EV_P^d = 24.0$).

Table 2: Predicted effort levels and expected profits

	P-Formal				**A-Formal**				**Predicted**
	E^{NE}	e^{NE}	EV_P	EV_A	$E^{d^{NE}}$	$e^{d^{NE}}$	EV_P^d	EV_A^d	**Delegation**
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 $E^{d^{NE}}$ denote Nash equilibrium predictions for the principal under P- and A-formal authority.
e^{NE} and $e^{d^{NE}}$ denote Nash equilibrium predictions for the agent under P- and A-formal authority.
EV_P, and EV_P^d denote expected equilibrium profits for the principal under P- and A-formal authority.
EV_A and EV_A^d denote expected equilibrium profits for the agent under P- and A-formal authority.

In the asymmetric treatments the rewards to delegation are either exacerbated or further diminished relative to the symmetric treatments. Of the four treatments, principal's 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 estimate predictions are given in Table 2. Note that the delegation decisions predicted by the Nash equilibrium are always in the set of welfare maximizing delegation choices. In PLOW aggregate expected earning $EV_p + EV_A$ are highest if the principal keeps control, while in PHIGH overall welfare is 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 principals are endowed with the right to choose the project. Thus, if principals value authority per se, loss aversion can play a role here because more loss averse principals may be more reluctant to give up authority than less loss averse principals. In order to better understand the potential individual heterogeneity in delegation decisions, we therefore measured subjects loss aversion with a lottery choice task. Each subject was 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. These lotteries enable us to construct individual measures of loss aversion. The amount X at which a subjects starts rejecting the lottery is an indicator of a subjects' loss aversion. For example, a subject that rejects all lotteries with a potential loss of $X > 3$ is classified as more loss averse than a subject that only rejects all lotteries with a potential loss of $X > 5$.[12]

In principle, one might think that the rejection of these lotteries is also compatible with risk aversion arising from diminishing marginal utility of lifetime income. Matthew Rabin's calibration theorem (Rabin (2000)) rules out this interpretation, however. Rabin shows that a theory of risk averse behavior based on the assumption of diminishing marginal utility of *lifetime* income implies that people essentially must be risk neutral for low-stake gambles like our lotteries. Intuitively, this follows from the fact that risk-averse behavior for low-stake gambles implies ridiculously high levels of risk aversion for slightly higher, but still moderate, stake levels. Yet, such unreasonably high levels of risk aversion can be safely ruled out. For example, we show in the appendix that if one assumes that the rejection of the lottery with $X = 4$ is driven by diminishing marginal utility of lifetime income, then the subject will also reject a lottery where one can lose \$30 with probability $\frac{1}{2}$ and win *any* price with probability $\frac{1}{2}$. Thus, there is no finite prize that induces this subject to accept a 50-percent chance of losing \$30. Similar results are implied by rejection of lotteries with other potential losses X.

2.3 Procedures

30 subjects participate in each experimental session which consists of three parts.[13] In part one, subjects play 7 periods of a single player version of the authority game. This single player game is identical to the authority game except 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 their 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, subjects play 10 periods of the main authority game in one of the four treatments. The 30 subjects are divided into 3 matching groups of 10 subjects consisting of 5 principals and 5 agents. Subjects are informed that in a new period they would be

[12] 74 out of 75 subjects who participated in the lottery task and played the authority-delegation game as a principal have a unique switching point.

[13] One session consisted of only 20 subjects.

matched with another randomly chosen partner.

In part three, subjects are asked to take a short questionnaire in which gender and other demographics 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.[14] The first series of experiments took place in May and June 2007 with a second series of experiments conducted in May and October 2008 and a third series in May 2009. In total, 380 subjects participated in the experiment, divided into 13 sessions. Experiments were computerized using the software z-tree (Fischbacher (2007)). Payment was given for each period of the main authority game and for the last five periods of the single player game. On average, an experimental session lasted one hour with an average payment of 33.5 CHF ($33.00 at the time of the experiment).

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 *If the controlling party is informed about the project valuations it almost always chooses its preferred project, implying that it overrules the subordinates recommendations. Informed subordinates almost always recommend their preferred project and uninformed principals almost always implement this recommendation.*

Result 1 is supported by the following numbers. Controlling parties in the role of the principal (agent) who were informed implemented their preferred project in 100 percent (98.3 percent) of the cases. Principals (agents) in the role of the informed subordinate party recommended their preferred project in 93.4 percent (94.3 percent) of the cases. Finally, principals (agents) in the role of the uninformed controlling party followed the subordinate party's recommendation in 94.2 percent (96.5 percent) of the cases. The high credibility of the subordinate's advice stems from the fact that they rarely misled the controlling party. If the subordinate parties were not informed they typically recommended the outside option (principals: 95.2 percent ; agents: 97.3 percent)

[14]Subjects were drawn from a database of volunteers using ORSEE (Greiner (2004)).

Result 1 indicates that the controlling parties used the decision right in their favor. 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 decision rights in the HIGH and PHIGH treatments and to keep authority in the LOW and PLOW treatments. Empirically, we find in our experiment:

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.*

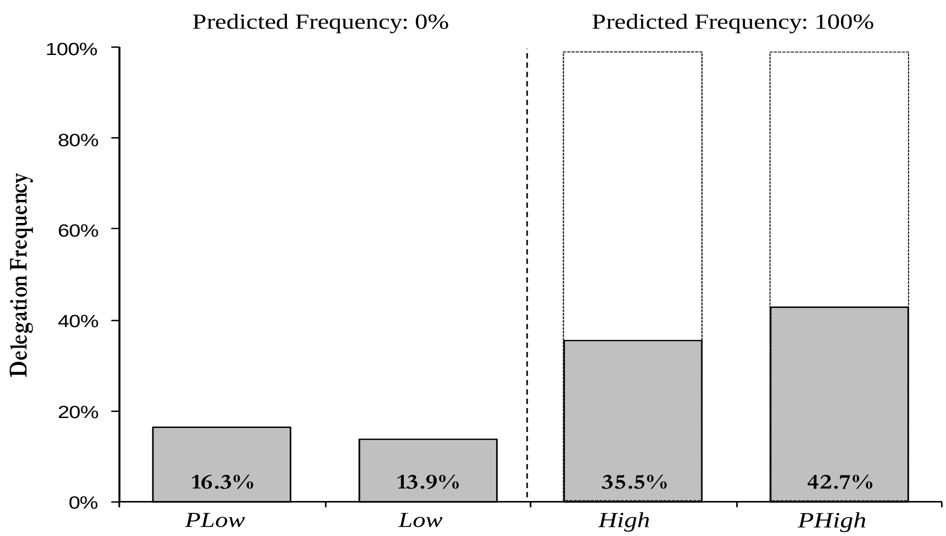

Figure 2: Delegation frequencies by alignment

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 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.

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. From period 2 onwards the delegation rate is roughly constant at 37.2 percent in this treatment. In the PHIGH treatment the overall delegation rate is 36.7 percent in the first five periods and stabilizes at roughly 48.7 percent from period 6 onwards.

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 seven 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, 30 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 in P-Formal and A-Formal authority makes it more profitable to keep authority. Table 3, which shows the realized profits of principal's 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 under P-Formal authority and higher than predicted under A-Formal authority. Principals who delegate have on average 30.3 percent greater earnings in the HIGH treatment and 45.5 percent greater earnings in the PHIGH treatment.

Table 3: Realized profits and predicted equilibrium profits for principals

	P-Formal		Number of	**A-Formal**		Number of
	Actual	**Predicted**	Observations	**Actual**	**Predicted**	Observations
PLOW	18.4 **	*20.1*	251	17.6	*17.2*	49
LOW	19.0	*20.1*	310	15.0 ***	*17.3*	50
HIGH	19.1 ***	*23.3*	316	24.9	*24.0*	174
PHIGH	18.4 ***	*23.3*	172	26.6	*25.6*	128

Significance levels calculated by regressing earnings on a constant and testing whether the constant is equal to the prediction. 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:

Result 3 *Controlling parties provide an excess of effort relative to the Nash equilibrium. Subordinates strongly 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.[15] In the low treatment, for example, the principal over provides effort by roughly 10 units relative to the prediction under P-Formal authority while the agent under provides effort by about 10 units. This deviation pattern is reversed under A-Formal authority when the agent is the controlling party.

Recall that in the Nash equilibrium the controlling party provides too much, and the subordinate party too little effort relative to the welfare maximizing effort levels. The fact that the controlling party over provides and subordinate party under provides effort relative to the equilibrium means that the actual effort allocation across parties is even more inefficient than in the predicted equilibrium. This inefficiency is 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, under P-Formal authority, 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.

As predicted by the model, the total welfare of the principal and agent are higher under delegation in both the HIGH and PHIGH treatments. Thus, the inefficiency of effort allocation is further exacerbated by the inefficiency in delegation. This effect is particularly acute in the PHIGH treatment in which delegation of authority would have made both parties better off. Comparing the average realized profits of principals and agents to the equilibrium profits *under the equilibrium authority structure* shows, that on average, earnings are between 8.5 percent and 14.4 percent lower than predicted for the principal and between 5.7 percent and 16.2 percent lower for the agent. [16] We can therefore conclude:

[15] We report results from a non-parametric Wilcoxon Signed-Rank test in Table A.1 of Appendix A.

[16] see Table A.2 of appendix A.

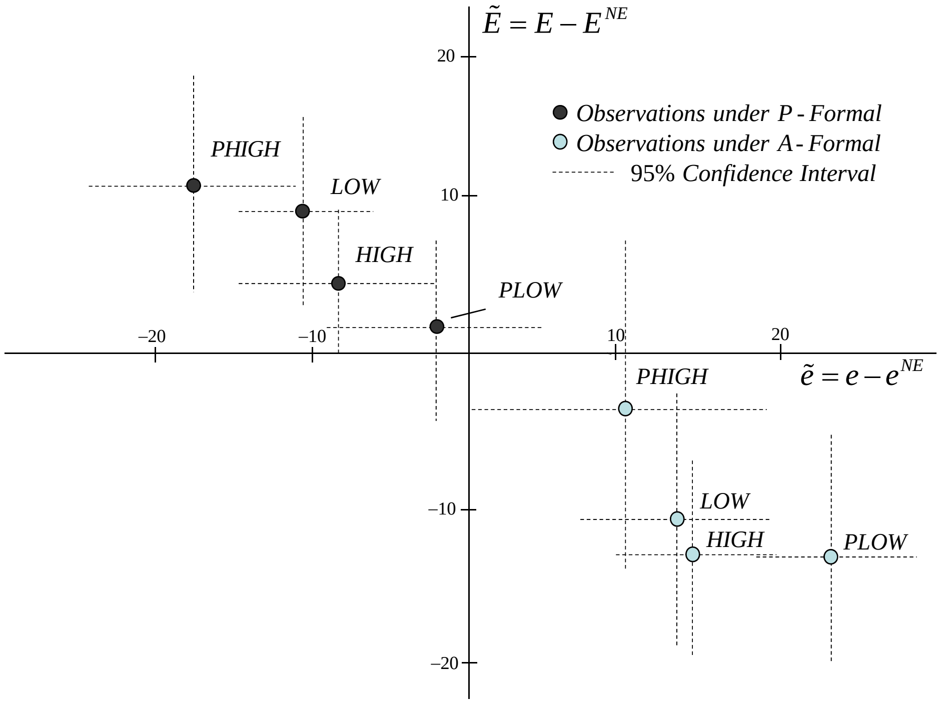

95% Confidence Intervals and mean effort calculated at the individual level.

Figure 3: Deviations from equilibrium effort predictions

Result 4 *The deviation in effort provision and delegation leads to welfare losses by both parties in all treatments. Welfare losses are most acute in the PHIGH treatment where delegation would make both parties better off.*

3.2 Exploring the principals' reluctance to delegate

A natural initial hypothesis for the observed under delegation of authority in the HIGH and PHIGH treatments is that individuals believe that they are 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 his beliefs about the agents actions under both P-Formal and A-Formal authority, what would be his gain or loss in expected earnings?

As the effort of the principal was elicited only in the case of her chosen authority structure a comparison of the principal's expected earnings for the case of delegation and non-delegation requires assumptions about her effort in the counterfactual authority structure. As we have beliefs data from both the delegation case and non-delegation case, a natural

Table 4: Realized profits and predicted equilibrium profits for agents

	P-Formal		Number of	A-Formal		Number of
	Actual	**Predicted**	Observations	**Actual**	**Predicted**	Observations
PLOW	23.0 ***	*25.6*	251	18.8 **	*23.3*	49
LOW	16.1 ***	*17.3*	310	17.9	*20.1*	50
HIGH	21.1 ***	*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. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

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.[17]

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.[18] 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 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,

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

$$EV_P^d(E^d = r_P^d(\hat{e}^d), \hat{e}^d) = \hat{e}^d \hat{P}_2 + (1 - \hat{e}^d) r_P^d(\hat{e}^d) \hat{P}_1 + P_0 - g_P(r_P^d(\hat{e}^d)), \tag{12}$$

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.

[18]This comparison value is given by

$$EV_P(E, \hat{e}) = E\hat{P}_1 + (1 - E)\hat{e}\hat{P}_2 + P_0 - g_P(E). \tag{13}$$

the under delegation is suboptimal not only from the principals' perspective, but also from the perspective of the organization as a whole.

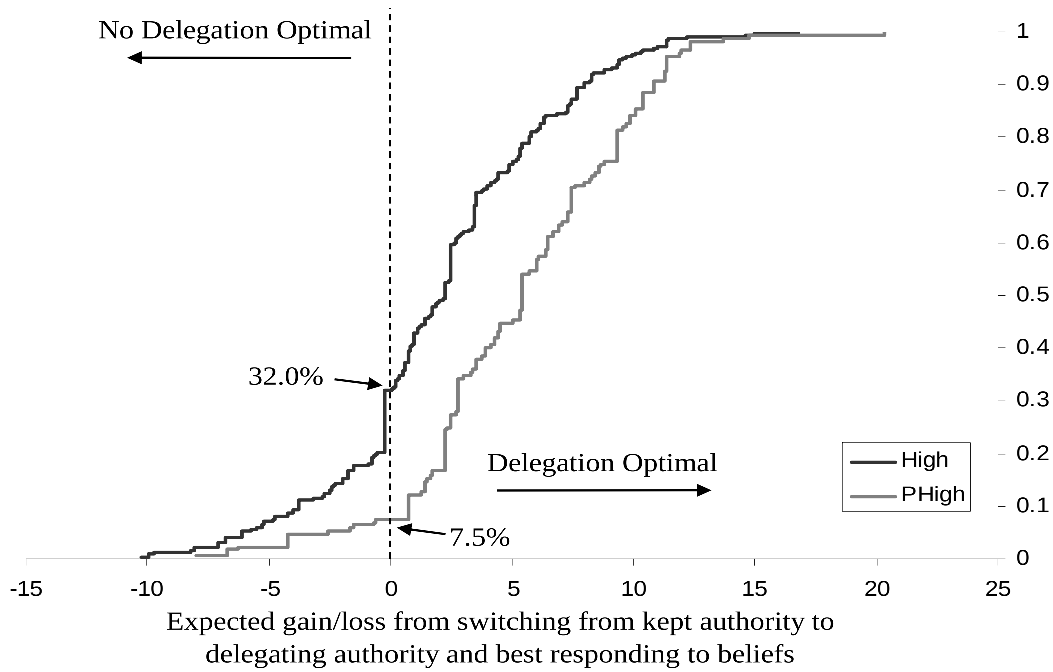

Figure 4: Cumulative density function of expected gain from delegation if principal best replies to beliefs

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 under A-Formal authority even with extreme risk and loss aversion. Second, besides very high effort choices which are observed very infrequently, zero effort minimizes the expected value for delegation giving us the lowest reasonable expected value for 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 delegation 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.

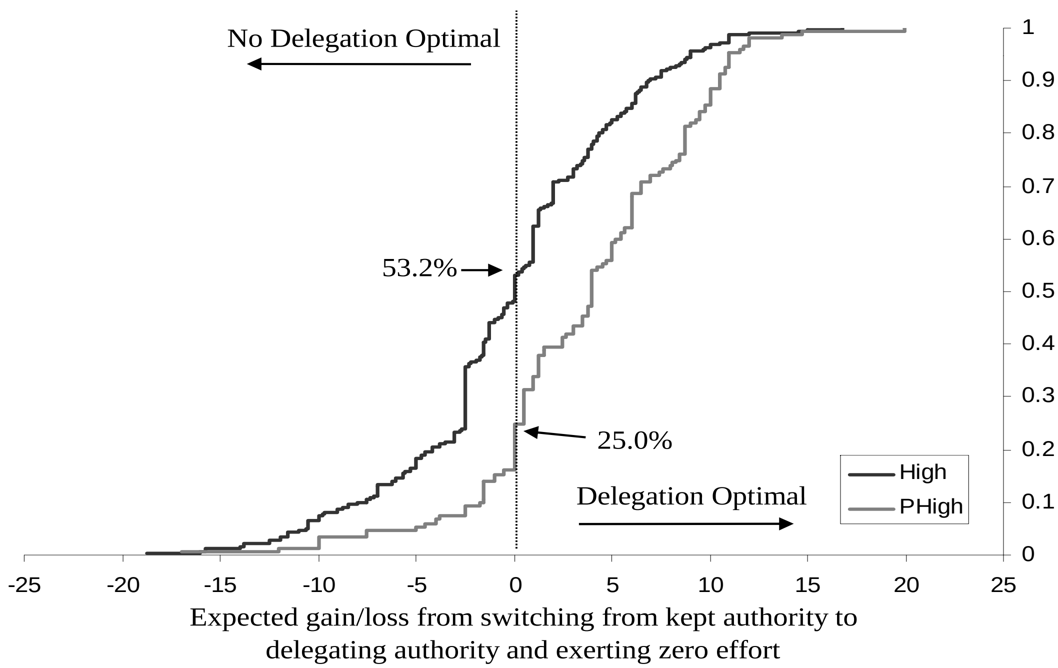

Figure 5: Cumulative density function of expected gain from delegation and if principal chooses zero effort after delegation

The results above suggest that the principals had little or no pecuniary reason to retain authority in the HIGH and PHIGH treatment. In fact, a large share of the principal's had strong pecuniary incentives to delegate. Why then do we observe this strong reluctance to delegate?

One potential explanation is that people receive direct value from holding decision rights, independent of any monetary consequences or associated risk. Assigning value to decision rights changes the trade-off between delegation and retention of authority. If subjects value decision rights per se, they may perceive the delegation of control rights as a loss. Loss aversion could then amplify the likelihood to keep decision rights. This hypothesis would be supported by evidence showing that individuals who are more loss averse are more likely to keep control.

To study this hypothesis, we turn to regression analysis. We first run the following regression

$$Delegation_{i,t} = \alpha_0 + \Sigma\alpha_t + \beta_{PLOW} I_{PLOW} + \beta_{HIGH} I_{HIGH} + \beta_{PHIGH} I_{PHIGH} + \epsilon_{i,t}, \quad (14)$$

where I_{PLOW}, I_{HIGH}, and I_{PHIGH} are treatment indicator variables and the LOW treatment is the omitted category. As the delegation decision is discrete, we use a probit specification with data clustered at the individual level. All results reported are robust to a linear specification of the regression model, a panel probit model, or a poisson regression which uses the absolute delegation frequency of principals as the dependent variable. We report the probit results from regression 14, and extensions of it, as these have the best controls for possible learning over time.[19]

Column (1) of Table 5 shows the marginal effects of the probit regression. The delegation prediction from section 2.2 predicts that $I_{HIGH} = I_{PHIGH} = 1$ and $I_{PLOW} = 0$. While the HIGH and PHIGH treatments are significantly different from the LOW treatment baseline, their magnitudes are significantly smaller than the predictions. Column (2) extends this regression to include beliefs and gender. As can be seen in the coefficients on the beliefs variables, a principal who believes the agent will put in more effort under A-Formal authority and less effort under P-Formal authority is more likely to delegate. Gender effects in delegation appear to be small and not significant.

Columns (3) and (4) present our main results with regard to loss aversion. As seen in Column (3), which includes observations from all four treatments, loss aversion strongly correlates with held authority. At the margin, subjects across all treatments are 6.9 percent more likely to keep authority for each additional gamble rejected in the lottery treatment. This effect is diluted, however, as individuals in the LOW and PLOW treatment are not expected to delegate. When the observations are restricted to only the HIGH and PHIGH treatments, as in Column (4), the effect of loss aversion is magnified, with a 12.6 percent decrease in probability of delegating for each gamble declined. As the variance of loss aversion is high (the standard deviation is 1.41 lotteries), loss aversion appears to constitute a major force in the under delegation phenomenon. In fact, combining observations in the HIGH and PHIGH treatment, the difference in delegation frequency between the group with high loss aversion (above the median) and low loss aversion (below the median) is 20 percent.

If loss aversion is a major driving force in the delegation decision, we should also see its fingerprint in other parts of our data. Kahneman and Tversky (1991) have shown the-

[19] As noted in the previous section, there is a small time trend in the HIGH and PHIGH treatments. The specification shown uses period fixed effects and period fixed effects interacted with the HIGH and PHIGH treatments as controls. Omitting these controls or replacing them with a linear time trend does not affect the stated results.

Table 5: Delegation decisions by principals

	(1)	(2)	(3[a])	(4[b])
PLOW	0.034	0.061	0.100	
	(0.068)	(0.073)	(0.079)	
HIGH	0.246***	0.216**	0.503***	
	(0.062)	(0.090)	(0.133)	
PHIGH	0.328***	0.252**	0.535***	–0.003
	(0.086)	(0.117)	(0.148)	(0.145)
Female (d)		0.061	0.049	0.056
		(0.048)	(0.068)	(0.152)
Beliefs P-Formal		–0.002**	–0.002	–0.005
		(0.001)	(0.001)	(0.003)
Beliefs A-Formal		0.003***	0.006***	0.012***
		(0.001)	(0.002)	(0.004)
Loss Aversion			–0.069**	–0.126**
			(0.027)	(0.055)
Period Dummies?	Yes	Yes	Yes	Yes
Pseudo. R^2	.073	.116	.181	.178
Observations	1450	1450	750	300

Significance levels: * p<0.1, ** p<0.05, *** p<0.01. Robust standard errors in parentheses, clustered by individual. Probit regression with marginal effects shown.

[a] Loss aversion measures are available only for sessions conducted in 2008.

[b] Column (4) includes data only from the HIGH and PHIGH treatments for which we have a loss aversion measure. HIGH is the omitted category.

oretically that there is a positive relationship between an individual's loss aversion in risky choices and the individual's proneness to the endowment effect. The endowment effect means that individuals value a good more highly because of the mere fact that they possess the good. Thus, even if they are randomly endowed with the good, they value it more. The strong negative impact of loss aversion on the delegation decision therefore suggests that there may also be an endowment effect with regard to authority, i.e. the principals (who have been randomly endowed with authority) may display a higher preference for authority than the agents. In a survey conducted at the end of the sessions with the HIGH treatment run in 2008, we asked individuals whether they (a) preferred to have authority, (b) were indifferent to the authority allocation, or (c) preferred not having authority. Table 6 shows the proportion of principals and agents who prefer having authority. As can be seen in the first column, 73 percent of Principals prefer to be the controlling party while only 20 percent prefer to be in the role of the subordinate. By contrast, agents are much more likely to prefer the subordinate role and are much less likely to prefer being the controlling party. The principals' strong preference for authority is quite remarkable because in case of delegation the principals earn on average 24.87 while if they keep authority they only earn 19.07 on average. Therefore, this strong preference for authority cannot be explained by the pecuniary experiences of the principals during the experiment. If anything, the pecuniary experiences in the experiment should have taught them to prefer the subordinate role. Thus taken together our data strongly supports an endowment effect in authority.

Table 6: **Principals' and agents' preferences for authority in the exit survey**

	Type	
Preferred Role	Principal	Agent
Controlling Party	0.73	0.20
Indifferent	0.07	0.33
Subordinate	0.20	0.47
Observations	15	15

Fisher's exact test of the null hypothesis
that Principals and Agents have the same preferences
yields a p-value of .018

We summarize the findings in this subsection by

Result 5 *On the basis of the principals' own beliefs about the agent's effort choices the principals' would often have been better off by delegating authority in the HIGH and PHIGH treatment. In addition, loss aversion of the principal appears to be an important determinant of the reluctance to delegate suggesting an endowment effect in decision rights.*

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 apparent convergence to the Nash equilibrium over time.

Persistent deviations from the Nash equilibrium might be due to one of two sources. First, for a given belief about the other parties 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 their best reply, its effort is likely to be higher than the Nash equilibrium effort. Likewise, if the subordinate party under provides effort relative to its best reply, then the effort of the subordinate is likely to be below that of the Nash equilibrium.

Second, beliefs about the other parties 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 its 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 in both P-Formal and A-Formal authority under the assumption of risk neutrality:

$$r_P(\hat{e}) = \frac{100\hat{P}_1 - \hat{e}\hat{P}_2}{50}, \qquad r_A^d(\hat{E}^d) = \frac{100\hat{A}_2 - \hat{E}^d\hat{A}_1}{50}. \tag{15}$$

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

Figure 6 shows this comparison for the HIGH treatment. The dashed 45° line represents those cases when the actual effort response to beliefs coincides with the best response to these beliefs. Points above the 45° line represents observations in which the controlling party over provides relative to the best response while points below the 45° line represents an under provision of effort.

The solid line in Figure 6 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, 67 percent of observations for principals and 76 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 over 50 percent of observations 15 points or more above prediction.

While Figure 6 shows the best response data only for the HIGH treatment, it is a fair representation of the data as a whole. Table 7 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 authority structures, and in 7 of these 8 cases the difference is significant. Based on this data, we conclude:

Result 6 *Controlling parties over provide effort relative to their best response to beliefs*

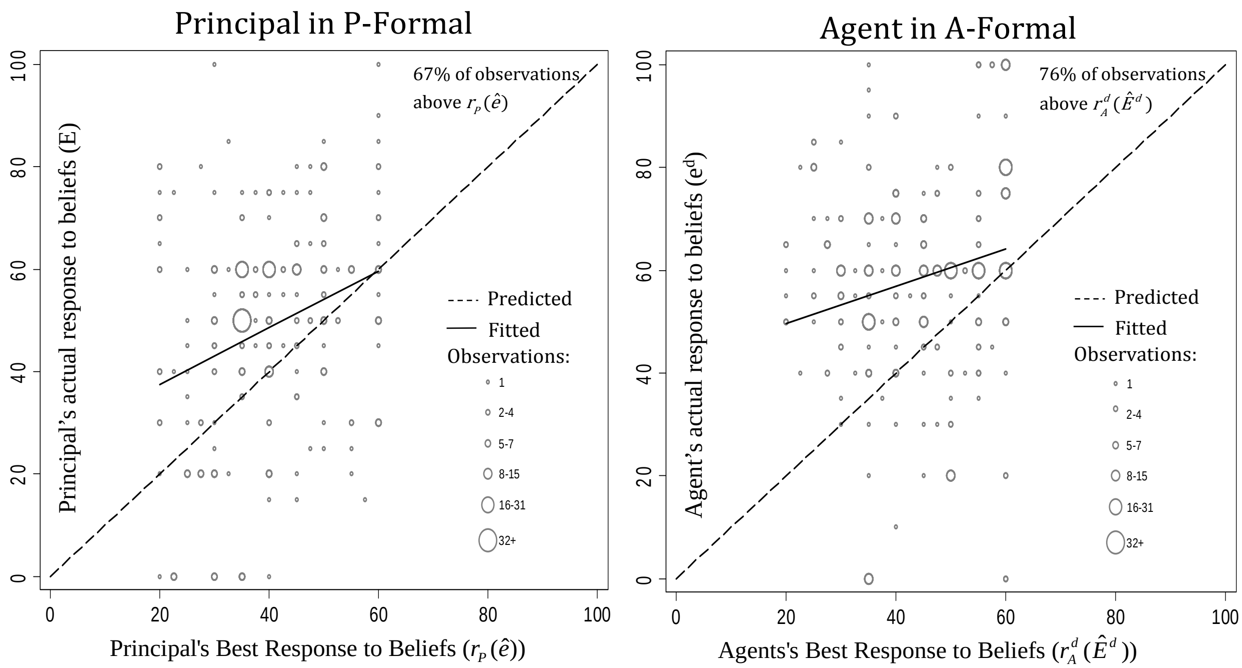

Figure 6: Controlling party: Actual effort vs. best response to beliefs (HIGH treatment)

Table 7: Comparison of effort provision of the controlling party to the best response to beliefs

	Principals in P-Formal			**Agents in A-Formal**		
	actual effort		best response effort	actual effort		best response effort
PLOW	55.7		53.7	68.1	***	49.1
LOW	66.1	***	54.6	68.3	***	55.8
HIGH	48.2	***	39.5	58.7	***	45.3
PHIGH	58.2	**	44.8	65.1	**	56.2

Significance levels calculated using a Mann-Whitney-Wilcoxon test with beliefs and effort averaged to the individual level prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

about the subordinate's effort.

Result 6 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 8 compares actual beliefs to the Nash Equilibrium beliefs for all treatments and authority structures. As can be seen by comparing the first two columns, the principal's belief about agent effort in P-Formal authority 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 for A-Formal authority. 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.

Table 8: Comparison of actual beliefs of the controlling party to the Nash equilibrium beliefs

	P-Formal			**A-Formal**		
	Nash prediction		actual belief	Nash prediction		actual belief
PLOW	25		30.7	35	***	21.8
LOW	25		26.8	25	*	20.9
HIGH	35		36.5	35	*	29.4
PHIGH	35	*	30.5	25	**	19.0

Significance levels calculated using a Mann-Whitney-Wilcoxon 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. In their case, a systematic under provision of effort relative to the best response may lead to a reduction of effort relative to the Nash equilibrium.

Figure 7 shows the relationship between the theoretical best response and the empirical response function in the case of 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 filled line shows the empirical best response behavior from a simple linear regression. Points above the 45° line represent observations in which the subordinate over provides effort relative to the best response while points below the 45° line represents an under provision of effort.

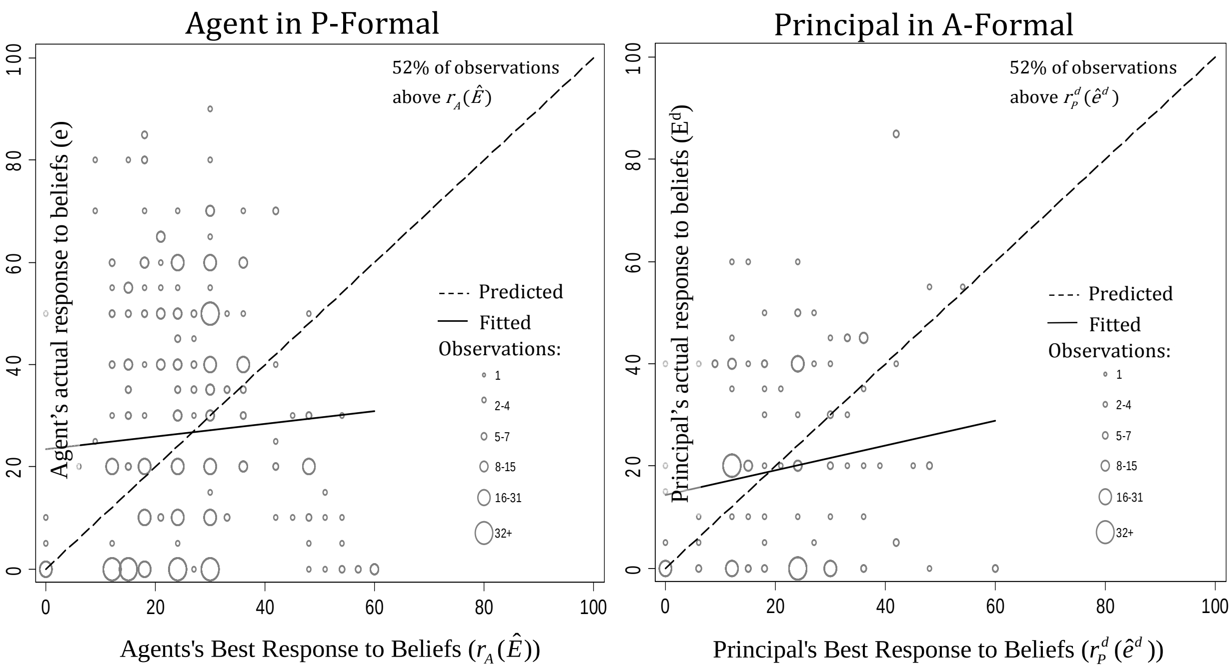

Figure 7: Subordinates: actual effort vs. best response to beliefs (HIGH treatment)

As can be seen in the left hand panel of the figure, the empirical response function is positive but 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 and roughly bimodal. 50 percent of individual choices are at or above the best response to beliefs for both agents in P-Formal authority and principals in A-Formal authority. In addition, a large number of individual choices is considerably below the best response. Interestingly, the heterogeneity of responses is not only present at the level of choices but also at the level of the individual. During the last five periods of the HIGH treatment 46 percent of the subordinates always underprovide effort and 39 percent of the individuals always overprovide effort relative to the best reply.

A particularly salient fact in Figure 7 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 for 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 response to the subordinate role appears to be a robust phenomenon across all four treatments. Across these treatments 58 percent of the subordinates always underprovide effort relative to the best reply during the last five periods and 24 percent always overprovide effort. Further evidence is shown in table 9, 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 authority structure. As can be seen by comparing the first two columns of each treatment and authority structure, 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 bimodal nature of subordinate effort provision where large outliers exist for both under and over provision of effort.

Looking at the third column of each row, however, the role of subordinate does appear to have a large demotivational effect in a large minority of individuals. In all conditions, at least 25 percent of individuals provide zero effort. In three of the eight conditions and authority structures, this group accounts for roughly 50 percent of observations. Based on these observations, we conclude:

Result 7 *The response to the subordinate role 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 large group, authority appears to have a strong demotivational effect. In addition, there is a smaller group of subordinates who systematically overprovide effort.*

Table 9: Comparison of effort of subordinates to their best response to beliefs

	Agents in P-Formal				**Principals in A-Formal**		
	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	18.4	36.7
LOW	14.3	*	19.8	49.4	16.2	20.0	54.0
HIGH	26.5		24.5	28.5	19.6	21.9	36.8
PHIGH	17.3		18.4	50.3	20.7	21.7	36.7

Significance levels calculated using a Mann-Whitney-Wilcoxon test with beliefs and effort averaged to the individual level prior to estimation. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Turning to beliefs, Table 10 shows the beliefs of the subordinate compared to the Nash Equilibrium beliefs. 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 7, the empirical response function is much flatter than would be predicted by the best response. Whereas theory would predict an increase in effort of 6 point per 10 point reduction in beliefs, the empirical response to beliefs is significantly smaller. For agents, a 10 point reduction in beliefs about the controlling parties' effort only leads to a 1.2 point increase in effort.

Moreover, beliefs about controlling party effort need to be extremely high[20] to rationalize a subordinate's effort choice of zero. Since such extreme beliefs are rarely observed, best replies to beliefs cannot explain the large fraction of zero effort choices.

Table 10: Comparison of actual beliefs of subordinates to the Nash equilibrium beliefs

	P-Formal			**A-Formal**		
	Nash prediction		actual belief	Nash prediction		actual belief
PLOW	55	***	64.8	45	***	63.9
LOW	55	***	66.9	55	***	68.8
HIGH	45	***	59.0	45	***	61.2
PHIGH	45	**	69.3	55	***	62.5

Significance levels calculated using a Mann-Whitney-Wilcoxon 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: Ruling out alternative explanations

Thus far we have seen that for both the principal and the agent deviations from best response behavior play an important role for departures from the Nash predictions. For all beliefs, a significant proportion of controlling parties provide effort which greatly exceeds the best response function leading to effort levels higher than predicted. Similarly, a significant proportion of subordinates provide zero effort despite the extremely small cost of low effort. This raises the question why subjects deviate from their best replies. Could it be, for example, that loss or risk aversion can account for the observed pattern of choices? Or, could the agents' effort choices be due to reciprocity or other forms of social preferences? Another possibility is that the deviations from best reply behavior are due to decision errors. In this section we show that all these hypotheses are very implausible explanations for the

[20]for example, 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 very extreme risk aversion, an effort of zero is predicted if the belief is 95, which is still very high and rarely observed.

dual deviation of controlling parties' and subordinates' effort. In our view, this means that decision rights (and the lack of them) may have motivational consequences that go beyond the motives that have so far been considered by economists.

A common reason for deviations from risk neutral predictions in laboratory experiments is the potential existence of loss aversion. As we saw in section 4.2, loss aversion seems to play a considerable role in the principal's delegation decision and encourages them to keep control. However, as we show in Appendix C, loss aversion 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. In rare cases where an individual has 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. Regression analysis supports this interpretation. If we regress the controlling parties' effort on our measure of loss aversion (excluding effort choices of 100) we find a very small and insignificant effect ($p = 0.35$). Thus, loss aversion cannot explain the controlling parties' overprovision of effort.

At the typical stakes available in laboratory experiments risk aversion is generally not a plausible explanation of behavior. If one rationalizes risk averse behavior in laboratory experiments with a utility function that is concave in wealth one inevitably predicts totally unreasonable levels of risk aversion at higher stake levels (Rabin (2000); see also Appendix B). Thus, there is a strong a priori reason why risk aversion is unlikely to be a reason behind the controlling parties' overprovision of effort. However, for the sake of the argument, let us assume that individuals are risk averse, i.e., they have a striktly concave utility function. An increase in risk aversion thus reduces the utility of the highest payment relative to the utility of the lowest payment. For the cost parameters chosen in this experiment, diminished marginal utility of money directly reduces the marginal utility for a successful search relative to its marginal cost. As search is now less attractive, effort of a risk averse individual is predicted to be strictly below the risk neutral optimum. We illustrate this argument more formally in Appendix D. Thus, just as with loss aversion, risk aversion cannot explain the over provision of effort by the controlling party.[21] Nor can risk and loss aversion explain the subordinates' choice of zero effort levels because effort costs are negligible at low effort levels.

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 principal, 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 = .382) nor in the position of the subordinate party (Kolmogorov Smirnov test, p = .449) 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 active 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

[21] It is also interesting that risk loving preferences cannot explain the over provision of effort by the controlling parties. A risk loving subject will increase effort beyond the level that maximizes the expected monetary payoff if the subject can "buy" additional risk with this behavior. Since the spread of the payoffs is fixed by our payoff parameters (see Table 1) the subjects can only vary risk by varying the effort. The variance of the payoff that a subject faces is maximal at $E = 50$ and it is zero at $E = 0$ and $E = 1$. Thus, below and above $E = 50$ the variance of payoffs is lower than at $E = 50$. Therefore, a risk loving subject generally will not sacrifice expected payoff for the sake of increasing effort above $E = 50$. However, the overprovision of effort by the controlling parties is exactly characterized by the fact that a large number of effort observations are substantially above 50.

the beginning of Section 3.3) 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 play a role. It turns out that the effort choices of the active party and the single player are indistinguishable (Kolmogorov Smirnov test, p = 0.43), indicating that social preferences do not affect effort.

Finally, we consider the possibility that decision errors might be responsible for the subjects' deviations from best response behavior. To examine this question we compute the quantal response equilibrium of the effort subgame. A basic idea of quantal response equilibrium (McKelvey and Palfrey (1995), Goeree and Holt (2001)) is (i) that subjects anticipate that their opponents make random errors and (ii) that they themselves play a best reply with errors to the other players' choices. In Appendix E we show that it is impossible to rationalize the overprovision of effort with such a model. This holds regardless of whether we assume small or large decision errors. In fact, a Kolmorov Smirnov test that compares the actual effort distribution with the distribution generated by the best fitting quantal response model clearly rejects the hypothesis that the two distributions are identical ($p < .001$).

4 Conclusion

Authority and power permeate political, social and economic interactions. It is therefore important to understand the motivation and incentive effects of these forces. 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 principal's beliefs, and that the individual and organizational welfare losses of this delegation bias are substantial. Our results suggest that individuals attribute a non-pecuniary value to the possession of authority. The positive valuation of authority per se is further endorsed by the finding that a principal's degree of loss aversion is highly predictive of his tendency to keep authority. Hence, there appear to be strong endowment effects with regard to decision rights.

Our results also show that authority has effects on motivation that are not captured by the theoretical model. The fundamental trade-off between incentives and control, as modeled by Aghion and Tirole (1997), indeed exists; relative to the first best the subordinate provides too little effort, and the controlling party provides too much. 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. These deviations from the best response do not appear to be generated by reciprocity or other forms of social preferences, risk attitudes or decision errors, but rather point towards motivational effects of authority itself.

The psychological literature on power has put forward the idea that humans value power per se. However, to our knowledge, this hypothesis has not been shown to hold in an environment that controls for the pecuniary, risk, and social preferences of individuals in the power relationship. That humans value power and authority per se is important for economic models of organizations as well as for governance considerations in the political sciences because the reluctance to delegate can cause considerable welfare losses for organizations and, perhaps, even society. Further empirical studies of the determinants and consequences of power motivations may thus yield important insights. We believe that our empirical approach may prove useful in this respect.

Our experimental design also allows us to identify the consequences of authority on effort provision. Given that risk aversion, loss aversion, social preferences, or decision errors are very unlikely explanations for the observed deviations from best response behavior, we conjecture that humans not only value authority for non-instrumental reasons, but that authority may also have direct consequences on motivation. This is in line with the approach-inhibition theory of power put forward by Anderson, Gruenfeld and Keltner (2003). According to this theory the possession of power induces approach-related behaviors that tend to focus on the potential gains in risky situations, while a lack of power induces inhibition-related behaviors that tend to focus on the downside risk. In terms of our experiment this means that the controlling parties focus primarily on the potential gains which can be reaped by a high effort while the subordinate parties primarily focus on the worst case, i.e., on the possibility that nobody knows the project valuations despite investments into effort. Reducing the effort to very low levels, the subordinate can improve the payoff in this worst case.

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 particular, the result that a lack of authority does not demotivate all people strongly suggests that putting the right people into positions that lack authority is important. The development of tools for detecting this type of employee may thus be important in minimizing the cost associated with the (re)allocation of authority. Further experiments are necessary to better understand the sources and consequences of under delegation and the effect of decision rights on motivation.

References

Aghion, Philippe and Jean Tirole, "Formal and Real Authority in Organizations," *Journal of Political Economy*, 1997, *105* (1), 1–29.

____ **, Mathias Dewatripont, and Patrick Rey**, "Transferable Control," *Journal of the European Economic Association*, 03 2004, *2* (1), 115–138.

Anderson, Cameron and Adam D. Galinsky, "Power, Optimism and Risk Taking," *European Journal of Social Psychology*, 2006, *36* (4), 511–536.

____ **and Jennifer L. Berdahl**, "The Experience of Power: Examining the Effects of Power on Approach and Inhibition Tendencies," *Journal of Personality and Social Psychology*, 2002, *83* (6), 1362–1377.

____ **, Deborah H. Gruenfeld, and Dacher Keltner**, "Power, Approach and Inhibition," *Psychological Review*, 2003, *110* (2), 265–284.

Athey, Susan and Donald J. Roberts, "Organizational Design: Decision Rights and Incentive Contracts," *The American Economic Review*, 2001, *91* (2), 200–205.

Baker, George, Robert Gibbons, and Kevin J. Murphy, "Informal Authority in Organizations," *Journal of Law, Economics, and Organization*, 1999, *15* (1), 56–73.

Bartling, Bjoern and Urs Fischbacher, "Shifting the Blame: On Delegation and Responsibility," IEW - Working Papers iewwp380, Institute for Empirical Research in Economics - IEW July 2008.

Bester, Helmut and Daniel Krahmer, "Delegation and Incentives," *RAND Journal of Economics*, 2008, *39* (3), 664–682.

Blanco, Mariana, Dirk Engelmann, Alexander K. Koch, and Hans-Theo Normann, "Belief Elicitation in Experiments: Is there a Hedging Problem?," IZA Discussion Papers 3517, Institute for the Study of Labor (IZA) May 2008.

Bowles, Samuel and Herbert Gintis, "Contested Exchange: Political Economy and Modern Economic Theory," *American Economic Review*, 1988, *78*, 145–150.

Charness, Gary, Ramon Cobo-Reyes, Natalia Jimenez, Juan A. Lacomba, and Francisco Lagos, "The Hidden Costs of Control: Comment," Mimeo 2009.

Coffman, Lucas C., "Intermediation Reduces Punishment (and Reward)," Mimeo 2010.

Dahl, Robert, “The Concept of Power,” *Behavioral Science*, 1957, *2*, 201–215.

deCharms, R., *Personal Causation: The internal affective determinants of behavior*, New York: Academic Press, 1968.

Deci, Edward L., *The Psychology of Self-Determination*, Lexington, MA.: Health, 1981.

Dessein, Wouter, “Authority and Communication in Organization,” *Review of Economic Studies*, 2002, *69* (1), 811–838.

Falk, Armin and Michael Kosfeld, “The Hidden Cost of Control,” *The American Economic Review*, 2006, *96* (1), 1611–1630.

Fischbacher, Urs, “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 2007, *10* (2), 171–178.

Goeree, Jacob K. and Charles A. Holt, “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions,” *The American Economic Review*, 2001, *91* (5), 1402–1422.

Greiner, Ben, “The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics,” Working Paper Series in Economics 10, University of Cologne, Department of Economics June 2004.

Grossman, Sanford J and Oliver D Hart, “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy*, August 1986, *94* (4), 691–719.

Harsanyi, John C., “The dimension and measurement of social power,” in K. W. Rothschild, ed., *Power in economics*, Harmondsworth: Penguin Books, 1978.

Hart, Oliver and John Moore, “Property Rights and the Nature of the Firm,” *Journal of Political Economy*, December 1990, *98* (6), 1119–58.

___ **and** ___ , “Debt and Seniority: An Analysis of the Role of Hard Claims in Constraining Management,” *American Economic Review*, June 1995, *85* (3), 567–85.

Jensen, Michael C., “Agency costs of free cash flow, corporate finance, and takeovers,” *American Economic Review*, 1986, *76.*

Kahneman, Daniel and Amos Tversky, “Loss Aversion in Riskless Choice: A Reference-Dependent Model.,” *The Quarterly Journal of Economics*, 1991, *106* (4), 1039–1061.

___ **, Jack L Knetsch, and Richard H Thaler**, “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy*, December 1990, *98* (6), 1325–48.

Knetsch, Jack L., "The Endowment Effect and Evidence of Nonreversible Indifference Curves," *American Economic Review*, 1989, *79.*

——, "Asymmetric Valuation of Gains and Losses and Preference Order Assumptions," *Economic Inquiry*, 1995, *33.*

Koszegi, Botond and Matthew Rabin, "A Model of Reference-Dependent Preferences.," *Quarterly Journal of Economics*, 2006, *121* (4), 1133–65.

Marx, Karl, *Das Kapital* 1867.

McClelland, David C., *Power: The Inner Experience*, Irvington, New York: Halsted Press, 1975.

McKelvey, Richard D. and Thomas R. Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, July 1995, *10* (1), 6–38.

Mulder, Mauk, *The daily power game*, Leiden: Martinus Nijhoff, 1975.

Niskanen, William A., *Bureaucracy and Representative Government*, Chicago: Aldine-Atherton, 1971.

Parsons, Talcott, "On the Concept of Political Power," *Proceedings of the American Philosophical Society*, June 1963, pp. 232–262.

Poppe, Matthijs, "Preferences for changing power positions and power distances: A social value orientation approach," *European Journal of Social Psychology*, 2003, *33.*

Rabin, Matthew, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 2000, *68* (5), 1281–92.

Rantakari, Heikki, "Governing Adaptation," *Review of Economic Studies*, October 2008, *75* (4), 1257–1285.

Rotter, Julian B., "Generalized Expectancies for Internal Versus External Control of Reinforcement," *Psychological Monographs: General and Applied,*, 1966, *80*, 1–28.

Russell, Bertrand, *Power: A New social Analysis* 1938.

Simon, Herbert, "A Formal Theory of the Employment Relation," *Econometrica*, 1951, *19* (3), 293–305.

Sliwka, Dirk, "On the Costs and Benefits of Delegation in Organizations," *Journal of Institutional and Theoretical Economics (JITE)*, December 2001, *157* (4), 568–.

Smith, Adam, *Lectures on Jurisprudence*, Oxford: Oxford University Press, 1978.

van Dijk, Marius and Matthijs Poppe, "Striving for personal power as a basis for social power dynamics," *European Journal of Social Psychology*, 2006, *36.*

Weber, Max, *Economy and Society, translated by G. Roth and C. Wittich* 1978.

Weber, Roberto A., John R. Hamman, and George Loewenstein, "Self-interest through delegation: An additional rationale for the principal-agent relationship," forthcoming, The American Economic Review 2010.

____, ____, and Jonathan Woon, "An Experimental Investigation of Delegation, Voting and the Provision of Public Goods," Mimeo 2010.

Zeuthen, Frederik, "Economic Welfare," in "Problems of Monopoly and Economic Welfare," London, UK: Routledge and Kegan Pail Ltd. (original publication, 1930), 1968, pp. 104–151.

Appendix A: Additional Tables

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

	Controlling Party						**Subordinate**					
	P-Formal			**A-Formal**			**P-Formal**			**A-Formal**		
	E		E^{NE}	e^d		e^{dNE}	e		e^{NE}	E^d		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 with data averaged to the individual level prior to estimation. Significance Levels: *** $p < .01$, ** $p < .05$, * $p < .1$,

Table A.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	20.69	*24.0*	21.13	*23.3*
PHIGH	21.89	*25.6*	16.83	*20.1*

[a]Actual earnings in treatment

[b]Predicted earnings with Nash equilibrium effort and delegation

Appendix B: Measuring Loss Aversion (not for publication)

In the main text of the paper, we interpret decisions made in the lottery task as being a result of loss aversion rather than risk aversion. This interpretation is based on Rabin's Calibration Theorem (Rabin (2000)) which shows that strictly concave utility of wealth is an implausible explanation for risk averse behavior over modest stakes. In this appendix we apply Rabin's calibration theorem to our lottery game. We show that if individuals have a globally concave utility function over wealth $w \in [0, \infty]$ and rejects gamble three of our lottery game — a coin flip in which the individual can either win CHF 6 or lose CHF 4 — then he or she will reject *any* coin flip in which she could lose CHF 30 no matter how large the positive prize that is associated with the coin flip. This is an implausibly high level of risk aversion while a reference dependent utility function that incorporates loss aversion can easily capture this behavior.

We proceed in four steps:

(i) We adopt the convention that, if indifferent, the individual rejects the coin flip. Rejecting the coin flip implies

$$0.5u(w+6) + 0.5u(w-4) \leq u(w)$$
$$u(w+6) - u(w) \leq u(w) - u(w-4)$$

It follows from concavity that $6[u(w+6) - u(w+5)] \leq u(w+6) - u(w)$ and $u(w) - u(w-4) \leq 4[u(w-3) - u(w-4)]$. Define $MU(x) = u(x) - u(x-1)$ as the marginal utility of the xth dollar. Putting the last three inequalities together, it follows that

$$MU(w+6) \leq \frac{2}{3} MU(w-4)$$

and, by concavity, that $MU(x+10) \leq \frac{2}{3} MU(x)$ for all $x > w - 4$.

(ii) We now derive an upper bound on $u(\infty)$. The concavity of $u()$ implies

$$u(w+10) \leq u(w) + 10MU(w)$$

Using the same logic,

$$\begin{aligned} u(w+20) &\leq u(w) + 10MU(w) + 10MU(w+10) \\ &\leq u(w) + 10MU(w)[1 + \frac{2}{3}] \\ u(w+30) &\leq u(w) + 10MU(w)[1 + \frac{2}{3} + \frac{2}{3}^2] \end{aligned}$$

and so on. Thus, we can develop a geometric series starting from w. Taking the limit, we obtain

$$u(\infty) \leq u(w) + 30MU(w)$$

(iii) Concavity implies $u(w-30) \leq u(w) - 30MU(w)$.

(iv) Using the results from step (ii) and (iii), we get an upper bound on the value of a coin flip where the individual would either lose CHF 30 or win an infinite amount:

$$0.5u(w-30) + 0.5u(\infty) \leq u(w)$$

This implies that the individual would reject the gamble. This concludes the proof.

Appendix C: Loss Aversion and Effort (not for publication)

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.

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

$$v(x) = \begin{cases} x - R & \text{if } x \geq R \\ (1+\lambda)(x-R) & \text{if } x < R \end{cases}, \tag{16}$$

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 maxima 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's and the subordinate's effort is unsuccessful. We use this fact to circumvent non-differentiability around the reference point by restricting analysis to this region. The optimization problem of the controlling party is

$$\max_E U(E) = E(P_1 - R - g(E)) + (1 - E)\hat{e}(P_2 - R - g(E)) \\ -(1 + \lambda)(1 - E)(1 - \hat{e})(P_0 - R - g(E)). \tag{17}$$

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

$$U'(E) = (\hat{P}_1 - g(E)) - Eg'(E) - \hat{e}(\hat{P}_2 - g(E)) - g'(E)\hat{e}(1 - E) - \\ (1 + \lambda)(1 - \hat{e})[(g(E)) - g'(E)(1 - E)] = 0. \tag{18}$$

Rearranging this equation and replacing $g'(E)$ and $g(E)$ and $\hat{P}_1$ with their values which were constant across treatments yields:

$$U'(E) = -50E + 30 - \hat{e}\hat{P}_2 + 50\lambda(1 - \hat{e})E\left[\frac{3}{2}E - 1\right] = 0. \tag{19}$$

Writing 19 as an implicit function, the FOC is satisfied when:

$$E = \frac{30 - \hat{e}\hat{P}_2}{50} + \lambda(1 - \hat{e})E\left[\frac{3}{2}E - 1\right]. \tag{20}$$

The last term is negative for $E \in \left[0, \frac{2}{3}\right]$ 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 1 *Effort of a loss averse individual will never be above 60 but below 100.*

Proof. Equation 19 can be rewritten as follows:

$$U'(E) = 75\lambda(1-\hat{e})E^2 - 50[1+\lambda(1-\hat{e})]E + 30 - \hat{e}\hat{P}_2 = 0. \tag{21}$$

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

$$U''(E) = 150\lambda(1-\hat{e})E - 50[1+\lambda(1-\hat{e})]. \tag{22}$$

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.

By the properties of quadratic functions, E is a local maxima/minima at:

$$\frac{50[1+\lambda(1-\hat{e})] \pm \sqrt{Z(\lambda)}}{150\lambda(1-\hat{e})}, \tag{23}$$

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,

$$E^*(0,\hat{e}) = \lim_{\lambda\to 0} \frac{50[1+\lambda(1-\hat{e})] - \sqrt{Z(\lambda)}}{150\lambda(1-\hat{e})} = \frac{[30-\hat{e}\hat{P}_2]}{50} \leq .6 \tag{24}$$

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 D: Risk Aversion and Effort (not for publication)

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.

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

$$U(E) = Eu(P_1 + w - g(E)) + \hat{e}(1 - E)u(P_2 + w - g(E)) \\ +(1 - \hat{e})(1 - E)u(P_0 + w - g(E)) \quad (25)$$

where w is wealth, $P_1 = 40$, $P_2 \in \{35, 20\}$, $P_0 = 10$, $g(E) = 25E^2$, and $\hat{e} \in \{0, .05, \ldots, 1\}$. 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

$$E(\sigma, \hat{e}) = arg \max_E Eu(P_1 + \hat{e}(1 - E)u(P_2 + w - g(E)) + (1 - \hat{e})(1 - E)u(P_0 + w - g(E)) \quad (26)$$

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. It follows that 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

$$\sigma^*(\hat{e}) = arg \max_{\sigma \in [-\infty, \sigma_1)} E(\sigma, \hat{e}) \quad (27)$$

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 A.3: Maximum effort predicted by risk aversion

	Low Treatment			
$\hat{e}$	$\sigma^*(\hat{e})$	σ_1	$E(\sigma^*, \hat{e})$	$E(0, \hat{e})$
0	$-0.7 - 0.6$	1.2	60	60
10	$-0.3 - 0.3$	1.4	60	60
20	$-1.3 - 0.7$	1.6	55	55
30	$-0.9 - 0.5$	2.0	55	55
40	$-2 - 0.9$	2.5	50	50
50	$-1.8 - 0.6$	3.2	50	50

	High Treatment			
$\hat{e}$	$\sigma^*(\hat{e})$	σ_1	$E(\sigma^*, \hat{e})$	$E(0, \hat{e})$
0	$-0.7 - 0.6$	1.2	60	60
10	$-.9 - 0.7$	1.6	55	55
20	$-1.1 - 0.7$	2.2	50	50
30	$-1.2 - 0.8$	2.9	45	45
40	$-1.4 - 0.9$	3.8	40	40
50	$-1.5 - 1.3$	5.0	35	35

Table A.3 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.

Appendix E: The impact of decision errors on effort choices (not for publication)

In this appendix we examine whether decision errors are a plausible explanation for subjects' deviations from best response effort levels. For this purpose, we construct the quantal response equilibrium (QRE) of the subgame in which principals and agents choose effort levels, and look at the noise parameter γ which best describes our data using a maximum likelihood criterion. The first step in this approach is to develop a logit choice model where the probability of selecting an effort level is determined by the ratio of the actual value of a choice relative to all potential alternatives. Let

$$p(E) = \frac{\exp(\gamma EV_P(E|p(e)))}{\sum_E \exp(\gamma EV_P(E|p(e)))} \tag{28}$$

be the probability that effort level E is chosen given the equilibrium choice distribution of the other party. The parameter γ controls the probability that an individual selects a suboptimal effort choice. As $\gamma \rightarrow \infty$ the QRE approaches the original unique Nash equilibria of the theoretical model. As $\gamma \rightarrow 0$, the choices of both parties goes toward the uniform distribution with equal probability weight on each outcome.

Across all treatments, the noise parameter γ which best describes our data is between 0.2 and .45, close to what would be predicted if all strategies were chosen randomly. The failure of the QRE to predict our results are best seen in Figure A.1, which plots QRE predicted effort against the empirical effort distribution for principals who kept control in the HIGH treatment. As can be seen by looking at the three QRE distributions plotted for various γ levels, the QRE model has a density function which is centered at the best response and which flattens as γ decreases. Note also, that since the QRE distributions are normally distributed, adding density to parts of the distribution which are away from the best response requires a reduction of the noise parameter γ.

Figure A.1 shows that the peak of the empirical distribution is to the right of the best response with a relatively tight distribution of effort levels. In order for the QRE model to add mass to the region which contains this empirical peak, γ must decrease. The low γ parameter is thus not informative of the data itself but merely reflects the fact that the QRE model predicts too low effort levels. As can be seen by comparing the QRE's predicted distribution against the empirical distribution, the two distributions are highly dissimilar. A similar argument can be made with regard to the subordinates' effort level. Since the mean of the distribution of effort predicted by the QRE is substantially above zero, the QRE model cannot explain the high frequency of zero effort choices. Therefore, taken together, QRE decision errors cannot explain the systematic deviations from best response effort levels.

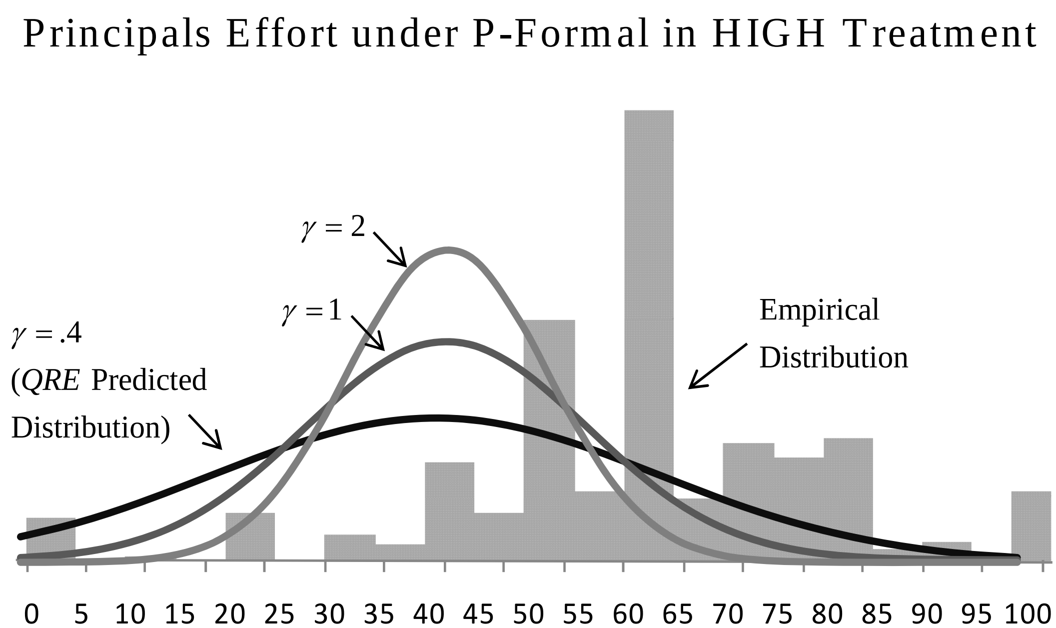

Figure A.1: Comparison of the QRE predicted effort distribution and the empirical effort distribution.

The distribution of effort predicted under QRE (shown by the smooth density lines) does not correctly predict the peak of the empirical effort distribution. In order to better fit the actual data (depicted by the histogram), the noise parameter γ must decrease. The best fitting QRE has a $\gamma = .4$ which is close to the uniform distribution. This distribution is a poor representation of the true distribution which has significantly lower variance and retains a bell curve shape. A Kolmogorov-Smirnov test rejects that they are the same distribution at the $p = .001$ level.