Information Acquisition in Expert Games

Tobias Böhm*

10.07.2007

Abstract

We analyze the incentives for information acquisition by an privately informed expert who is concerned about being regarded as well informed. We show that deviations from efficient decision making may increase if the agents received more signals. We discuss how these inefficiencies can be mitigated by organizational structure.

JEL Classification: D82, D83

Keywords: Reputation, Cheap Talk, Advice

*Department of Economics, University of Munich, Ludwigstr. 28 VG, D-80539 Munich, Germany, tobias.boehm@lrz.uni-muenchen.de. I would like to thank Florian Englmaier, Nadine Riedel, and seminar participants at the University of Munich for helpful discussions and comments.

1 Introduction

In many situations of economic relevance a principal has to either delegate a decision to an agent or rely on information the agent transmits, since the principal lacks sufficient information or expertise to perform the task independently. Examples include such diverse settings as delegated portfolio management or advisers of a political decision maker (as in Morris (2001)). If the agent's information can be verified full information disclosure can often be sustained as an equilibrium outcome (see e.g. Bolton and Dewatripont (2005) on disclosure of private certifiable information.). The same holds true if the principal can design a mechanism to elicit information.

If these condition are not satisfied, however, inefficiencies might arise due to divergence of preferences as in standard cheap talk games[1] or due to the agent's desire to appear well informed. These situations have been analyzed in so called expert games: here an agent (the expert) is hired in order to make a decision on behalf of the principal.[2] He bases his decision on his private information whose accuracy is determined by his type. In contrast to a classical cheap talk game the agent does not care about the decision per se, but only about the decision's impact on the principal's assessment of his type. These kind of preferences can be rationalized by means of future reemployment considerations, for example. In the examples mentioned above such incentives arguably play an important role: politicians or their advisers care about reelection which is partially determined by the electorate's assessment of their ability.[3] The same is probably true for other state officials.[4] As a further example, the behavior of fund managers is driven by such career concerns, a point empirically confirmed by Chevalier and Ellison (1999) and theoretically elaborated by Dasgupta and Prat (2005, 2006).

This paper examines the consequences of information acquisition by experts. To analyze this, we extend an expert model to two periods where in each period the agent acquires additional

[1]The first paper analyzing the outcome in cheap talk games is Crawford and Sobel (1982). A more recent reference is Battaglini (2002); Krishna and Morgan (2007) gives a very comprehensive overview.

[2]It could also be that the agent is just supposed to making a recommendation to the principal who then makes the decision. In the framework of this paper this distinction is immaterial.

[3]See for example Ottaviani and Sorensen (2001) or Majumdar and Mukand (2004).

[4]Levy (2005) builds a model of the incentives of judges where the behavior of the judges in driven by such career concerns.

information. The agent can act at different points in time and the principal can observe the timing decision of the agent. There are a lot of situations where this setting seems realistic: fund managers can decide whether to invest early or late in certain stocks, for example. The set up allows us to shed light on questions like: Do agents have an incentive to accumulate information? If so, which agents? Does the principal necessarily benefit from better informed agents? Or might there be an incentive for the principal to restrict information acquisition and force agents to act early?

It turns out that more information on the agent's side does not unambiguously benefit the principal. Of course, more information improves the quality of decision making which benefits both the principal and the agent. However, in some situations the behavior of the agents is further distorted through more information acquisition. If the agent has gathered more than one signal an additional effect comes into play which is best understood if one considers efficient decision making first. Efficiency dictates that the agent chooses those actions, which from an ex ante point of view are less likely to be optimal, only if his information in favor of these actions is sufficiently strong. In case this is not true, which will happen especially if the expert has received contradictory information, he should opt for a "standard" action. However, since the better the agent the more correlated is his information over time, predominantly bad experts will end up with conflicting signals.[5] Hence, under efficient decision making, the standard actions carry a reputational discount as they are selected mainly by unable experts. In equilibrium this effect might lead to inefficiencies as some agents will be tempted to choose the reputationally more valuable action even if their own information is not precise enough. As we will illustrate in the model, this wedge between different actions only emerges if agents have accumulated more than one piece of information. Therefore it might be beneficial for the principal to restrict the agents and force them to make their decision early.[6]

The paper can thus also be seen as a contribution to the literature on optimal delegation. Here the principal does not restrict *which* actions are available to the agent (as for example in Alonso and Matouschek (2007)), but at *which point in time* the agent is supposed to act. There is also a completely different rationale behind the principal's desire to restrict the agents. The reason does not lie in an imperfect alignment of preferences which induces the

[5]The better the agent the higher the probability that the signal is correct. Hence better agents are more likely to receive identical signals, as in case of contradictory signals at least one signal must have been wrong.

[6]One way to achieve this may be work overload of the agents. See section 4 for a more detailed discussion.

agents to choose the wrong action from the principal's point of view. In fact the principal and the agent share the same objective as the agents want to choose the correct action. Restricting the agents may nevertheless be valuable, since the accumulation of more information might aggravate distortions due to strategic behavior. In the framework of an expert game, this paper is the first one which shows that the principal might indeed be hurt by agents holding better information.

1.1 Related Literature

The paper is related to several strands of the literature a few of which have been mentioned already. There is a link to the literature on optimal delegation and, obviously, on cheap talk. As the agent cares about the principal's assessment of his ability, this paper fits into the literature on career concerns, which has been pioneered by Holmström (1999) and generalized by Dewatripont, Jewitt, and Tirole (1999a,b). There an agent's ability increases expected output and the agent exerts unobservable effort to improve some performance measure. Here, in contrast, the agent's type determines the precision of the information she receives. This information is private but the principal can observe the action chosen by the agent.

The literature on experts started with a famous paper by Scharfstein and Stein (1990) who analyzed sequential decision making by multiple experts.[7] This allows them to study incentives to herd on previous agent's actions. The herding result depends crucially on the agent's objective function: Effinger and Polborn (2001) show that anti-herding can be an equilibrium outcome if the agents care sufficiently about their relative reputation.

The issue of the timing of decision making has not received attention in the literature on experts, but was analyzed in the context of statistical herding models. Statistical herding occurs whenever an agent disregards his own signal because the actions of her predecessors at least partially reveal their information. If this information is stronger than the agent's own information it is optimal to follow the predecessors independently of one's own signals. See Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) for early contributions in this realm. Following these papers there are a few contributions where agents can endogenously determine when to act but can not influence the amount of information they receive.

[7]See Ottaviani and Sorensen (2000) for a generalization.

Gul and Russell (1995) establish that endogenous sequencing leads to delay and a clustering of agent's decision. As actions are public and reveal (part of) the private information agents delay their actions in order learn other agent's private information. As soon as the first agent has moved (this will be the agent with the most extreme signal realization) others follow immediately (clustering of actions). In Zhang (1997) the model is extended to two dimensional uncertainty: not only are the signal realizations private information but also the accuracy of the information. The equilibrium of this model also exhibits initial delay; once the first agent (here the one with the most precise signal) has moved, again all others follow immediately.
While we share with these papers the focus on the timing of decision making, I will only consider settings with a single agent who has to make only one decision, so herding is not an issue. Moreover, in contrast to the statistical herding literature agents in this chapter care about their reputation. In what follows I will give a short overview over the reputational expert literature, which can be divided into two parts depending whether the agent knows his own type or not.
The problem if the agent is ignorant about his own type and receives only one signal, has been extensively studied in Ottaviani and Sorensen (2006a,b). They show that quite generally agents distort their behavior in order to signal competence. Prat (2005) shows that the problem of inefficient signalling can be alleviated if the principal can only observe whether the agent has chosen the correct action or not, but not which action in particular.
If the agent knows his type new distortions can arise as shown for example by Trueman (1994), Prendergast and Stole (1996) Avery and Chevalier (1999), and Levy (2004). Good agents, knowing their type will find it optimal to follow their signal more often and contradict the prior. As the bad agent has an incentive to appear well informed, he will try to mimic the behavior of good agents. Hence, it will be optimal for him to contradict the prior even if his information is not precise enough to outweigh the prior.
The only paper in this literature which allows for an endogenous information structure is Levy (2004). Here the agent can resort to an external consultant and gather additional information. She shows that the agent has an incentive to ignore or even excessively contradict the consultant in order to gain reputation. The result is thus complementary to ours as it too studies the behavioral distortions arising if more information about the optimal course of action is available. However, there are two important differences which make the mechanism at work quite distinct from the one at work in our model. First, the quality of additional information is independent of the expert and second, the consultant's recommendation is

publicly observable. By contradicting the consultant, the agent can signal that his information is superior.

The paper is organized as follows. In the next section I present the model. The next two parts of the chapter present the analysis when agents have to move in the first or second period respectively. Section 5 concludes. All proofs are relegated to the appendix.

2 The Model

I consider a model with two time periods $t = 1, 2$ and two players, an agent A and an evaluator E. In every period the agent can choose an action $d_t \in \{a, b\}$ which I assume to be irreversible.

Which action is best depends on the realization of a state of the world $x \in \{a, b\}$. The prior probability of $x = a$ being the true state of the world is denoted by $q \geq \frac{1}{2}$. In addition, the agent receives a signal $s_t \in \{a, b\}$ about x in every period immediately before she can make the decision. The precision of the signal depends on the agents type $\theta \in \Theta := \{\underline{\theta}, 1\}$, $\underline{\theta} \in [\frac{1}{2}, 1]$ in the following way:

$$\Pr(s_t = a | x = a, \theta) = \Pr(s_t = b | x = b, \theta) = \theta.$$

Hence, we assume that the agent's type denotes the probability with which a correct signal is obtained. The quality of the signal therefore increases in the agent's type with the case of $\underline{\theta} = \frac{1}{2}$ corresponding to the situation where the bad agent receives pure noise. Note that the good type $\theta = 1$ gets a perfect signal in every period.[8] I assume that conditional on the true state x the signals are drawn independently in every period. The true state is revealed to all players after the agent has made his decision.

Importantly, I assume that the agent's type is only known to her. The prior probability of $\theta = 1$ is given by p.

An important ingredient of the model is that the agent's payoff u_A depends positively on the evaluators assessment of his type.[9] I assume that the agent's payoff is given by her expected

[8]For the model to be interesting it is important that the information of the good type is sufficiently precise. In particular, it must be optimal for the principal to follow the agent's advice if he knew that the agent was good. As long as this is satisfied, the exact precision of the good agent's information does not qualitatively affect the equilibrium.

[9]This can be, e.g. due to reemployment or promotion decisions the evaluator has to make in the future.

type conditional on all information the evaluator has. The evaluator observes the decision of the agent, the point in time when the decision was made and the true state of the world.

$$u_A = \mathbb{E}(\theta|d_1, d_2, x, p).$$

The evaluator is a passive player whose only task it is to assess the quality of the agent.[10] Throughout the paper I will consider Perfect Bayesian Equilibria (PBE). In such an equilibrium the agent's strategy must be optimal given the evaluator's beliefs. The evaluator forms beliefs according to Bayes' Rule using all of his information whenever this is possible. More formally the agents strategy consists of functions $d_1(s_1|\theta) : \{a,b\} \rightarrow \Delta(\{a,b\})$ and $d_2(s_1, s_2|\theta) : \{a,b\}^2 \rightarrow \Delta(\{a,b\})$.[11] The evaluator uses an updating function $\mu(d_1, d_2, x, p)$: $\{a,b\} \times \{a,b\} \times \{a,b\} \times [0,1] \rightarrow [0,1]$ which denotes the posterior probability of facing a good agent. Using the updating function the evaluator can compute $\mathbb{E}(\theta|d_1, d_2, x, p) = \mu(d_1, d_2, x, p) + (1 - \mu(d_1, d_2, x, p))\underline{\theta}$.

In what comes I will restrict attention to informative equilibria where the agent conditions his actions on his information.[12]. Moreover I will ignore all "mirror" equilibria which take some equilibrium and just flip every action from a to b and vice versa. [13]

The timing of the game is as follows:

1. The agent learns his type θ.
2. The agent receives signal s_1 and chooses d_1.
3. The agent receives s_2 and chooses d_2.
4. Evaluator observes d_1, d_2 and x and updates about agent's type.
5. Payoffs realized.

[10]This is common in the literature on career concerns. One could interpret the evaluator as consisting of possible future employers of the agent, who, after having observed the agent's performance, are willing to offer a wage equal to the experts expected reputation for his services.

[11]Given some set A, $\Delta(A)$ denotes the set of all probability distributions over A.

[12]As common in all cheap talk games there also exists an "babbling" equilibrium in which the agent's decisions does not convey any information about his type and the evaluators belief is independent of any of the agent's actions

[13]This is standard, see e.g. Levy (2004)

3 Analysis with a Single Period

As a benchmark case and in order to gain some intuition into the workings of the model, consider first a situation where the agent has to make his decision in $t = 1$.[14] To this end define $\overline{V}_1^i := \mathbb{E}(\theta | d_1 = i, x = i)$, $i = a, b$ as the reputational payoff for the agent if he chooses $d = i$ in the first period and $x = i$. Analogously we define $\underline{V}_1^i = \mathbb{E}(\theta | d_1 = i, x \neq i)$, $i = a, b$ as the agent's reputation if he chooses the wrong action. As there is no incentive for the good type $\theta = 1$ to contradict his signal, we have $\underline{V}_1^i = \underline{\theta}$.

Assume first that the bad agent decides to follow his signal as well. Note that if the evaluator correctly anticipates the behavior of the agent this implies $\overline{V}_1^a = \overline{V}_1^b$ as

$$\begin{aligned}\overline{V}_1^a &= \mathbb{E}(\theta | d_1 = a, x = a) = \Pr(\theta = 1 | d = a, x = a) \cdot 1 + \Pr(\theta = \underline{\theta} | d = a, x = a) \cdot \underline{\theta} \\ &= \frac{qp}{qp + q(1-p)\underline{\theta}} \cdot 1 + \frac{q(1-p)\underline{\theta}}{qp + q(1-p)\underline{\theta}} \cdot \underline{\theta} = \frac{p + (1-p)\underline{\theta}^2}{p + (1-p)\underline{\theta}},\end{aligned}$$

and

$$\begin{aligned}\overline{V}_1^b &= \mathbb{E}(\theta | d_1 = b, x = b) = \Pr(\theta = 1 | d = b, x = b) \cdot 1 + \Pr(\theta = \underline{\theta} | d = b, x = b) \cdot \underline{\theta} \\ &= \frac{(1-q)p}{(1-q)p + (1-q)(1-p)\underline{\theta}} \cdot 1 + \frac{(1-q)(1-p)\underline{\theta}}{(1-q)p + (1-q)(1-p)\underline{\theta}} \cdot \underline{\theta} = \frac{p + (1-p)\underline{\theta}^2}{p + (1-p)\underline{\theta}}.\end{aligned}$$

As the agent's payoff depends only on making the correct decision the bad agent has always an incentive to follow his signal if $s_1 = a$. Formally we see that given $s_1 = a$ the agent prefers $d_1 = a$ over $d_1 = b$ as

$$\Pr(x = a | s_1 = a, \underline{\theta}) \overline{V}_1^a + \Pr(x = b | s_1 = a, \underline{\theta}) \underline{V}_1^a \geq \Pr(x = b | s_1 = a, \underline{\theta}) \overline{V}_1^b + \Pr(x = a | s_1 = a, \underline{\theta}) \underline{V}_1^b,$$

because $\Pr(x = a | s_1 = a, \underline{\theta}) \geq q \geq \frac{1}{2}$.

Following $s_1 = b$ in turn is only optimal if

$$\Pr(x = b | s_1 = b, \underline{\theta}) \overline{V}_1^b + \Pr(x = a | s_1 = b, \underline{\theta}) \underline{V}_1^b \geq \Pr(x = a | s_1 = b, \underline{\theta}) \overline{V}_1^a + \Pr(x = b | s_1 = b, \underline{\theta}) \underline{V}_1^a.$$

Note that this condition is only satisfied if $\Pr(x = b | s_1 = b, \underline{\theta}) \geq \Pr(x = a | s_1 = b, \underline{\theta})$, i.e. only if $\underline{\theta} \geq q$.[15] The intuition for this is straightforward. If $\underline{\theta}$ is sufficiently small then a signal in favor of state b does not outweigh the prior, i.e. even after having observed $s_1 = b$ the

[14] The results in this section can in a slightly different form already be found in Trueman (1994) and in Avery and Chevalier (1999).

[15] This implies that agents will always follow their signal if $q = \frac{1}{2}$. Notice that this is also efficient.

bad agent considers state a more likely to be true. Since her utility depends only on making the correct decision, bad agents have an incentive to contradict their signal. If the prior probability of state a being true is high enough, this effect may be so strong that the bad agent does not choose $d_1 = b$ anymore. Note that in this case observing $d = b$ and $x = a$ is an out of equilibrium event. In what follows I assume that the evaluator holds the belief that he faces a bad agent for sure whenever he observes the inappropriate decision. The equilibrium of the game can be summarized as follows:

PROPOSITION 3.1 *In equilibrium the good type will always follow his signal. The behavior of the bad type depends on the parameters.*

1. *If $\underline{\theta} \geq q$ the bad agent will also follow his signal.*
2. *Let $\beta_1 = Pr(d = b|s_1 = b)$. If $p \leq \frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}$ the bad agent will choose $d_1 = a$ whenever $s_1 = a$ and $d_1 = b$ with probability $\widehat{\beta}_1$ if $s_1 = b$ where $\widehat{\beta}_1$ is the unique solution to*

$$\frac{p + (1-p)\underline{\theta}\widehat{\beta}_1}{1 - (1-p)(1-\underline{\theta})\widehat{\beta}_1} = \frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}.$$

3. *If $p > \frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}$ the bad agent will always choose $d_1 = a$.*

PROOF: See the appendix.

Part one of the proposition asserts that the bad agent will always follow his signal if it is efficient to do so. The last part stipulates that the bad agent will abstain from $d_1 = b$ if, even after having observed a signal in favor of state b, she still puts sufficiently large probability on state a being true. In this case the likelihood ratio $\frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}$ on the right hand side becomes sufficiently small. This will hold true whenever the difference between q and $\underline{\theta}$ is large, i.e. if either there is a strong prior in favor of state a and/or the information of the bad agent is very noisy. In intermediate cases (part 2 of the proposition) the bad agent randomizes between both actions given she received signal b. To understand the equilibrium condition note that one can write the left hand side as

$$\frac{p + (1-p)\underline{\theta}\widehat{\beta}_1}{p + (1-p)[\underline{\theta} + (1-\underline{\theta})(1-\widehat{\beta}_1)]}$$

which exactly gives the probability that $d = b$ will be chosen conditional on $x = b$ relative to the probability of $d = a$ conditional on $x = a$. This expression is a measure of the relative reputational payoffs attached to the different actions. If $\widehat{\beta}_1$ declines, the bad agent switches away from action b and hence, the expected type upon observing $d = a$ decreases. As a direct consequence the reputation earned upon having selected $d = a$ correctly decreases as well, and so does the left hand side of the equilibrium condition. When will it be optimal to shy away from the unexpected action b? Only when the right hand side of the equilibrium condition goes down as well, hence, if state a is ex ante more likely and if the information quality of the agent deteriorates (i.e. $\underline{\theta}$ declines).

We can therefore conclude that the randomization decision trades off two effects: on the one hand, as the bad agent chooses $d = b$ less often, a higher reputation can be gained by correctly selecting action b compared to action a. On the other hand however, given that the information of the bad type is not precise enough to outweigh the prior, $d = b$ is correct with a smaller probability. It is clear then that the randomization probability decreases if the latter effect becomes larger, i.e. if the prior rises or $\underline{\theta}$ decreases.

Note that the agent behaves inefficiently in this intermediate case. Although his information suggests state a being the most probable, the agent chooses $d_1 = b$ with some probability.

4 Second Period Decision Making

I now turn to the analysis when the agent must make his decision in the second period. Define $\overline{V}_2^i$ ($\underline{V}_2^i$) analogously as the reputation of the agent after having correctly (incorrectly) chosen action i in the second period.

The behavior of the good type is not affected. She will still choose the action prescribed by her signals. The bad agent in contrast can now end up with three different posteriors. Either she has received two identical signals ($s_1 = s_2 = a$ or $s_1 = s_2 = b$) or two contradictory signals ($s_1 \neq s_2$).[16] In the first case the agent believes with a higher probability that the signals indicate the true state of the world compared to a situation where she received only one signal. In the latter case the signals exactly offset each other and the agent puts probability q on state a being true. The following figure illustrates the different posterior assessments of the agent dependent whether she received only one signal (as in the previous section) or two signals.

As one moves from the left to the right, state b is considered to be more likely. The lower

[16]Notice that the good agent never receives two different signals.

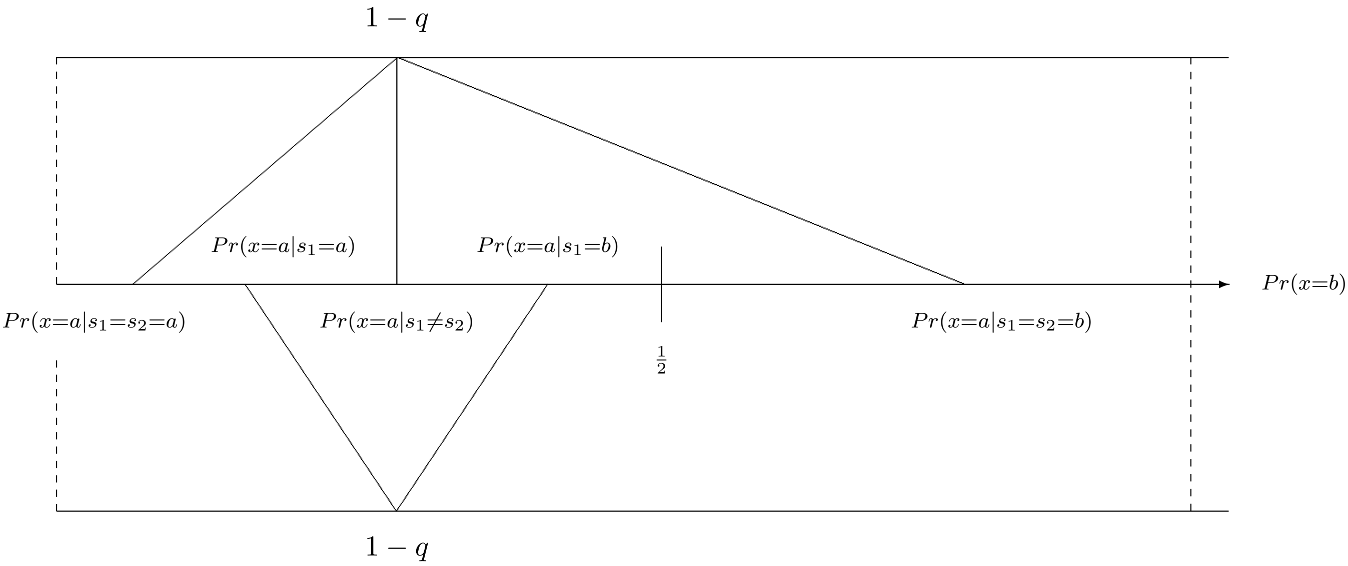

Figure 1: Posteriors of the bad agent

half of the figure illustrates how the bad agent's assessment of state b being true changes conditional on the signal she receives. Starting from the prior $1-q$ the agent puts even less probability on state b being true if he obtains $s_1 = a$. In the figure we implicitly assumed that $\underline{\theta} < q$ so even after having observed a signal in favor of b the bad agent assigns a probability smaller than $\frac{1}{2}$ of state b being true.

The upper side of the figure illustrates the agent's posterior if he has received two signals. In case of two conflicting signals the posterior equals the prior. Identical signals, however, pull the posterior more strongly to the indicated state. In the figure the bad agent's signal is good enough, such that two signals in favor of state b more than offset the prior.

It turns out that it is of crucial importance which decision is made given the two signals are not equal, as this influences the reputation which is attached to an agent successfully choosing action a or b (still assuming that the evaluator correctly anticipates the behavior of the agents.). In particular, the action which is chosen after $s_1 \neq s_2$ carries a lower reputational value. As contradictory signals are only received by bad agents, the action chosen after $s_1 \neq s_2$ is made by a bad agent with a higher probability. This already indicates that it cannot be an equilibrium outcome that $d_2(s_1 \neq s_2) = b$, since in this case $d = b$ would carry a reputational discount and would be correct with smaller probability compared to $d = a$. This intuition can be formally confirmed.

Lemma 4.1 *If $s_1 \neq s_2$ the bad agent will play $d_2 = a$ with strictly positive probability.*

Proof: See the appendix.

How large this probability is in equilibrium depends on q. Assume that $d_2(s_1 \neq s_2) = a$ with probability one. In this case the decision $d_2 = a$ carries a reputational discount while $d_2 = b$ gives a higher reputation to the agent. Choosing $d_2 = b$ correctly is a better signal about the agent's type since the evaluator knows that the agent must have received two correct signals. In contrast, upon observing that decision a was selected correctly, the evaluator can only infer that the agent has received at least one correct signal. The former is much more probable for good relative to bad agents, hence $\overline{V}_2^b > \overline{V}_2^a$ (see the appendix for a formal proof). However, for this to be an equilibrium outcome, an agent with two conflicting signals must find it optimal to select action a. He will do so if

$$q\overline{V}_2^a + (1-q)\underline{\theta} \geq (1-q)\overline{V}_2^b + q\underline{\theta}.$$

Given the wedge between $\overline{V}_2^b$ and $\overline{V}_2^a$ this condition is satisfied only if q is larger than some threshold value $\widehat{q}$.[17]

Case 1: $q \geq \widehat{q}$ Assume first that this is the case. It is clear that if $d_2(s_1 \neq s_2) = a$ is optimal than $d_2(s_1 = s_2 = a) = a$ will be optimal as well. But similar to the the one period case the bad agent has an incentive to contradict his signals in case of $s_1 = s_2 = b$ if his information is sufficiently bad. Formally, it holds true that

$$\begin{aligned} \Pr(x = b|s_1 = s_2 = b, \underline{\theta})\overline{V}_2^b &+ \Pr(x = a|s_1 = s_2 = b, \underline{\theta})\underline{\theta} \\ &< \\ \Pr(x = b|s_1 = s_2 = b, \underline{\theta})\underline{\theta} &+ \Pr(x = a|s_1 = s_2 = b, \underline{\theta})\overline{V}_2^a. \end{aligned}$$

given that $\underline{\theta}$ is low enough.[18] The structure of the equilibrium will resemble the one in the one period case as the agent will contradict signals in favor of b with at least some probability. The equilibrium is summarized in the following proposition where I assume again that E believes to face a bad agent with probability one in case of an inappropriate decision.

Proposition 4.1 *Suppose $q \geq \widehat{q}$ and define $\kappa := \frac{p+(1-p)\underline{\theta}^2}{p+(1-p)\underline{\theta}(2-\underline{\theta})}$. In equilibrium the good agent will always follow his signal.*

[17] As $\overline{V}_2^a$ is always larger than $\underline{\theta}$ the existence of $\widehat{q}$ is guaranteed.

[18] In the extreme case where $\underline{\theta} = \frac{1}{2}$ the agent does not receive any information and so the posterior equals the prior. We know then from the definition of $\widehat{p}$ that the bad type will not find it optimal to follow his signals if $s_1 = s_2 = b$.

1. *The bad agent always chooses* $d_2 = a$ *if* $p > \frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2}$.

2. *Suppose* $p \leq \frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2} \leq \kappa$. *Then the bad agent will set* $d_2 = a$ *if either* $s_1 = s_2 = a$ *or* $s_1 \neq s_2$. *If* $s_1 = s_2 = b$ *the agent chooses* $d = b$ *with probability* $\widehat{\beta}_2 > \widehat{\beta}_1$ *implicitly defined by*
$$\frac{p + (1-p)\underline{\theta}^2\widehat{\beta}_2}{1 - (1-p)(1-\underline{\theta})^2\widehat{\beta}_2} = \frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2}.$$

3. *Suppose* $\frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2} > \kappa$. *Then the bad agent follows his signal whenever* $s_1 = s_2$ *and chooses* $d_2 = a$ *in case of* $s_1 \neq s_2$.

Proof: See the appendix.

The intuition for this equilibrium is similar to proposition 3.1. Note first that the likelihood ratio $\frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2}$ again measures how probable it is that state b is true relative to state a, given that two signals in favor of b have been observed.

The bad agent will again disregard his signal and always choose $d_2 = a$ if the evidence in favor of state b is weak. This happens for large values of q and small values of $\underline{\theta}$. If the information of the agent becomes better she will randomize between decision a and b if she receives two signals in favor of state b. The equilibrium condition again equates the likelihood ratio of both states being true with the relative reputational payoffs attached to the different actions. Only if the posterior probability of state b is sufficiently high the agent follows his information.

It is important to note that conditional on $s_1 = s_2 = b$ the agent chooses $d_2 = b$ more often if he has to decide in the second period compared to the decision in the first period. There are two reasons for this. First, given $s_1 = s_2$ the evidence in favor of state b being true is now stronger. However there is a second effect which makes the agent more aggressive. As already noted above, the fact that $d_2 = b$ is chosen only with two confirmatory signals while for $d_2 = a$ at least one signal in favor of state a must have been received, leads to a reputational "premium" of decision b relative to decision a even if the agent does not randomize.

Hence letting the agent decide in the second period may lead to strategic distortions in the behavior of the agents. An interesting fact is that this distortion can more than offset the benefits of better information which can be accumulated.

Proposition 4.2 *If $\underline{\theta}$ is low enough (and smaller than q) the probability of a wrong decision is higher if the agent acts in the second period only.*

It should be rather obvious that $\underline{\theta} < q$ is necessary for inefficient decision making to increase. If the agent does not randomize in any of the two equilibria under consideration, decision making will be better in the second period as more information is utilized.
However, if the bad agent randomizes, additional distortions can arise. As an illustration consider a parameter constellation such that the agent randomizes between a and b if he has obtained one, respectively two signals in favor of state b. We know from the proposition that $\widehat{\beta}_2 > \widehat{\beta}_1$. Consider now the limit case where $\underline{\theta} = \frac{1}{2}$. In this case the bad agent receives only noise while the information quality of the good agent does not improve through a second perfect signal either. The probability of a correct decision declines strictly when a second signal is acquired, as the bad agent should efficiently choose $d = a$ regardless of his information (remember that $q \geq \widehat{q} > \frac{1}{2}$). As $\underline{\theta}$ rises the decision made by the bad agent improves since $d_2 = b$ if and only if the agent has received two signals in favor of it. Hence the evidence of state b being true becomes stronger. For $\underline{\theta}$ high enough this effect of better information offsets the stronger inclination to choose $d = b$ if agents acquire two signals.
Hence, if the evaluator or a principal is interested in correct decision making he may benefit from forcing the agent to act early and forgo useful information. An interesting question is how the principal can achieve this. One possible way might be work overload of the agents which prevents them from gathering additional information. The role of work overload in mitigating agency problems has been already identified in previous work. In Aghion and J.Tirole (1997) work overload on the principal's side serves as a commitment device not to interfere with the agent's decisions. Although not optimal ex post, this commitment gives better incentives to the agents ex ante. Laux (2001) focuses on a different mechanism. He notes that bundling several tasks and allocating them to a single agent might reduce agency cost stemming from limited liability. The rationale for overload offered here is quite different. In contrast to Aghion and J.Tirole (1997), here it may be optimal to overburden the *agent*, not his principal in order to achieve better decision making. Moreover, in Aghion and J.Tirole (1997) the *probability* that the agent makes a valuable decision increases in the principal's work load, but the decision per se remains the same. Here it is that work overload has a direct impact on which decision the agent makes.

Case 2: $q < \widehat{q}$ In the analysis above the "inferior" action $d = a$ was sustainable as in the case of $s_1 \neq s_2$ the agent was compensated for the lower reputation with a higher success probability. However, if q is sufficiently close to $\frac{1}{2}$, playing $d_2(s_1 \neq s_2) = a$ with probability one is no longer optimal as the reputational wedge between the different actions is too large to be offset by the different success probabilities. Formally,

$$q\overline{V}_2^a + (1-q)\underline{\theta} \geq (1-q)\overline{V}_2^b + q\underline{\theta}$$

is violated for small values of q. Hence in the only equilibrium the agent will now randomize between both actions if she has received two contradictory signals. As the agent is now indifferent between both actions after having received no information, she will strictly prefer to follow her signal in case of $s_1 = s_2$. The equilibrium is formally described in the following proposition.

PROPOSITION 4.3 *Suppose $q < \widehat{q}$. Both agents will follow their signal if $s_1 = s_2$. In case of contradictory signals the agent chooses $d = b$ with probability $\widehat{\beta}_3$, where $\widehat{\beta}_3 < \frac{1}{2}$ is implicitly defined by the unique solution to*

$$\frac{p + (1-p)\underline{\theta}(1 - (1-\underline{\theta})(1-2\widehat{\beta}_3))}{p + (1-p)\underline{\theta}(1 + (1-\underline{\theta})(1-2\widehat{\beta}_3))} = \frac{1-q}{q}.$$

PROOF: See the appendix.

If ex ante both states of the world are considered to be almost equally likely, there does not exist an "inferior" action anymore which is chosen without new information. Still as long as $q \geq \frac{1}{2}$, action a will carry a lower reputational value than action b.[19] This is only possible if the bad agent chooses $d_2 = a$ with higher probability, hence $\widehat{\beta}_3 < \frac{1}{2}$.

Next, we will examine the consequences of additional information on the principal's well being, in particular, on the probability of a wrongful decision. It turns out that the principal always benefits from better information if the prior on the two states is balanced enough (i.e. $q \leq \widehat{q}$).

PROPOSITION 4.4 *Consider the case of $q < \widehat{q}$. Then under second period decision making the probability of success is always higher.*

[19] Again the agent must be compensated for the lower success probability if he chooses $d_2(s_1 \neq s_2) = b$ with a higher reputation in case of success.

PROOF: See the appendix.

If the agents postpone their decision they will make better informed decisions which benefits the principal directly. Additionally, if $q < \widehat{q}$ the distortions arising from strategic behavior are also smaller. Two forces drive this result. Given that the agent must decide in the first period, she holds weaker information in favor of state b when selecting $d = b$. In the second period the agent has either acquired two contradictory signals or two signals in favor of state b but the randomization probability is such that on average, the evidence in favor of $x = b$ is stronger. Additionally, as shown in the appendix the agent chooses the inefficient action b with strictly lower probability.

5 Conclusion

This paper examined the consequences of information acquisition in an expert setting, where an agent (the expert) who is primarily concerned with his reputation, has to make a recommendation to a principal. It was shown that while better informed agents make correct recommendations more often, more information also has a potential downside. As bad agents try to mimic good ones, their behavior suffers from excessive "experimentation", i.e. they select ex ante less likely actions to often from an efficiency point of view. This distortion can be aggravated if agents hold better information, since an additional reputational wedge is driven between different actions. This additional wedge can hurt the principal if he is concerned with correct decision making. But a different dimension of interest might be the selection of agents. If the principal's primary focus is on selecting able agents and sorting out bad ones, another rationale for restricting agent's access to information arises. Superior competence can be better assessed if the error probability of bad agents increases. So especially in the beginning of their career, when arguably selection is more important than correct decision making, agent's might be overburdened with work.[20] More generally speaking, one could think of the organizational structure as a whole being designed such that career concerns of agents are optimally exploited. Koch and Peyrache (2005) is a very interesting first step in that direction.

Although this paper attempted to advance our understanding of expert models by consider-

[20]The issue of selection in expert settings has received little attention so far. A notable exception is Prat (2005).

ing a rather obvious extension, a host of further extensions are still unexplored. To name a few, the role of multiple experts is still poorly understood in that setting. As shown by Dewatripont and Tirole (1999), forcing agents to advocate a certain standpoint increases effort. In a pure cheap talk setting, Krishna and Morgan (2001) explore the role of consulting multiple agents. Under certain conditions, if the agents are not too strongly biased in opposite directions, this can improve on information transmission.

To make full use of multiple experts, it might be necessary to augment reputational incentives with explicit incentive schemes. Zwiebel (1995) already noted that relative reputational concerns are of major interest.[21] However, those relative concerns can also be created with contracts specifying some form of relative performance evaluation. Hopefully those issues will be examined in more detail in the near future.

[21] See Effinger and Polborn (2001).

6 Appendix

Proof of proposition 3.1

Part one of the proposition is already shown in the text. For the other parts, first define the reputational values of choosing $d = a$ and $d = b$ if a bad agent sets $d(s_1 = b) = b$ with probability β_1 and plays $d = a$ otherwise.

$$\begin{aligned}\overline{V}_1^a(\beta_1) &= \Pr(\theta = 1 | d = a, x = a) \cdot 1 + \Pr(\theta = \underline{\theta} | d = a, x = a) \cdot \underline{\theta} = \\ &\frac{p}{p + (1-p)(\underline{\theta} + (1-\underline{\theta})(1-\beta_1))} \cdot 1 + \frac{(1-p)(\underline{\theta} + (1-\underline{\theta})(1-\beta_1))}{p + (1-p)(\underline{\theta} + (1-\underline{\theta})(1-\beta_1))} \cdot \underline{\theta}\end{aligned}$$

$$\begin{aligned}\overline{V}_1^b(\beta_1) &= \Pr(\theta = 1 | d = b, x = b) \cdot 1 + \Pr(\theta = \underline{\theta} | d = b, x = b) \cdot \underline{\theta} = \\ &\frac{p}{p + (1-p)\underline{\theta}\beta_1} \cdot 1 + \frac{(1-p)\underline{\theta}\beta_1}{p + (1-p)\underline{\theta}\beta_1} \cdot \underline{\theta}\end{aligned}$$

To understand this expression note that good types (which occur with probability p) always implement the correct policy. Bad agents, in turn, choose action a correctly if they receive the correct signal (which happens with probability $\underline{\theta}$) and if they obtain a wrong signal but decide to contradict it (which happens with probability $(1-\underline{\theta})(1-\beta_1)$).

If the agent randomizes between both actions the following equality must hold

$$\Pr(x = a | s_1 = b, \underline{\theta})\overline{V}_1^a(\beta_1) + \Pr(x = b | s_1 = b, \underline{\theta}) \cdot \underline{\theta} = \Pr(x = b | s_1 = b, \underline{\theta}) \cdot \overline{V}_1^b(\beta_1) + \Pr(x = a | s_1 = b, \underline{\theta}) \cdot \underline{\theta},$$

which is true if

$$q(1-\underline{\theta})\left[\overline{V}_1^a(\beta_1) - \underline{\theta}\right] = (1-q)\underline{\theta}\left[\overline{V}_1^b(\beta_1) - \underline{\theta}\right]$$

$$\Longleftrightarrow$$

$$q(1-\underline{\theta})\left[\frac{p(1-\underline{\theta})}{p + (1-p)(\underline{\theta} + (1-\underline{\theta})(1-\beta_1))}\right] = (1-q)\underline{\theta}\left[\frac{p(1-\underline{\theta})}{p + (1-p)\underline{\theta}\beta_1}\right]$$

$$\Longleftrightarrow$$

$$\frac{p + (1-p)\underline{\theta}\beta_1}{p + (1-p)(\underline{\theta} + (1-\underline{\theta})(1-\beta_1))} = \frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}$$

$$\Longleftrightarrow$$

$$\frac{p + (1-p)\underline{\theta}\beta_1}{1 - (1-p)(1-\underline{\theta})\beta_1} = \frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}.$$

The right hand side is smaller than one if $\underline{\theta} < q$. The left hand side is equal to p if $\beta_1 = 0$ and goes to one as $\beta_1 \to 1$. Moreover, the left hand side is monotonically increasing in β_1.

Hence by the intermediate value theorem, if $\frac{(1-q)\theta}{q(1-\underline{\theta})} \in [p, 1]$ there exist a unique $\beta_1 \in [0,1]$ which solves the condition in proposition 3.1.
If $\frac{(1-q)\theta}{q(1-\underline{\theta})} < p$ then it can be easily seen that the payoff from choosing $d = a$ always exceeds the payoff from $d = b$. ||

Proof that $\overline{V}_2^b > \overline{V}_2^a$ if $d_2(s_1 \neq s_2) = a$

$$\overline{V}_2^a = \frac{p + (1-p)\underline{\theta}(\underline{\theta}^2 + 2\underline{\theta}(1-\underline{\theta}))}{p + (1-p)(\underline{\theta}^2 + 2\underline{\theta}(1-\underline{\theta}))} = \frac{\mathbb{E}(\theta^3) + (1-p)2\underline{\theta}^2(1-\underline{\theta})}{\mathbb{E}(\theta^2) + (1-p)2\underline{\theta}(1-\underline{\theta})}.$$

$$\overline{V}_2^b = \frac{p + (1-p)\underline{\theta}^3}{p + (1-p)\underline{\theta}^2} = \frac{\mathbb{E}(\theta^3)}{\mathbb{E}(\theta^2)}.$$

The good type will again always choose action a if this is appropriate. Bad agents will do so if they receive the correct signal twice (which occurs with probability $\underline{\theta}^2$) or in case of contradictory signals (which happens with probability $2\underline{\theta}(1-\underline{\theta})$). $\overline{V}_2^b > \overline{V}_2^a$ if

$$\mathbb{E}(\theta^3) \cdot \mathbb{E}(\theta^2) + \mathbb{E}(\theta^3)(1-p)2\underline{\theta}(1-\underline{\theta}) > \mathbb{E}(\theta^3) \cdot \mathbb{E}(\theta^2) + \mathbb{E}(\theta^2)(1-p)2\underline{\theta}^2(1-\underline{\theta})$$

$$\Longleftrightarrow$$

$$\underline{\theta}\mathbb{E}(\theta^2) < \mathbb{E}(\theta^3),$$

which is always satisfied. ||

Proof of Lemma 4.1

We have just shown that the action which is chosen only with two confirmatory signals bears a higher reputation in case of success. Assume this action would be a, i.e. $d_2(s_1 \neq s_2) = b$. This would imply $\overline{V}_2^b < \overline{V}_2^a$. In case of two contradictory signals, the bad agent must have an incentive to choose b. But

$$(1-q)\overline{V}_2^b + q\underline{\theta} \geq q\overline{V}_2^a + (1-q)\underline{\theta}$$

can only be satisfied if $\overline{V}_2^b > \overline{V}_2^a$, a contradiction. ||

Proof of proposition 4.1

Let $q \geq \widehat{q}$ which implies $d(s_1 \neq s_2) = a$. Define $\beta_2 = \Pr(d_2 = b | s_1 = s_2 = b, \underline{\theta})$.

$$\begin{aligned} \overline{V}_2^b(\beta_2) &= \Pr(\theta = 1 | d_2 = b, x = b) \cdot 1 + \Pr(\theta = \underline{\theta} | d_2 = b, x = b) \cdot \underline{\theta} = \\ &= \frac{p}{p + (1-p)\underline{\theta}^2\beta_2} \cdot 1 + \frac{(1-p)\underline{\theta}^2\beta_2}{p + (1-p)\underline{\theta}^2\beta_2} \cdot \underline{\theta}. \end{aligned}$$

$$\overline{V}_2^a(\beta_2) = \Pr(\theta = 1|d_2 = a, x = a) \cdot 1 + \Pr(\theta = \underline{\theta}|d_2 = a, x = a) \cdot \underline{\theta} =$$

$$\frac{p}{p + (1-p)[\underline{\theta}^2 + 2\underline{\theta}(1-\underline{\theta}) + (1-\beta_2)(1-\underline{\theta})^2]} \cdot 1 + \frac{(1-p)[\underline{\theta}^2 + 2\underline{\theta}(1-\underline{\theta}) + (1-\beta_2)(1-\underline{\theta})^2]}{p + (1-p)[\underline{\theta}^2 + 2\underline{\theta}(1-\underline{\theta}) + (1-\beta_2)(1-\underline{\theta})^2]} \cdot \underline{\theta}.$$

To understand these expressions note that the denominator of $\overline{V}_2^a$ gives the probability that action a is chosen correctly in equilibrium. The good type will always choose correctly, while the bad type sets $d_2 = a$ if either he receives two signals in favor of a (this happens with probability $\underline{\theta}^2$) or two mixed signals (probability $2\underline{\theta}(1-\underline{\theta})$). She will also choose $d = a$ with probability $1 - \beta_2$ if she receives two signals in favor of state b (probability $(1-\underline{\theta})^2$). The other expressions can be interpreted analogously.

For randomization to be an equilibrium outcome we must have

$$\Pr(x = a|s_1 = s_2 = b, \underline{\theta})\overline{V}_2^a(\beta_2) + \Pr(x = b|s_1 = s_2 = b, \underline{\theta}) \cdot \underline{\theta} =$$

$$\Pr(x = b|s_1 = s_2 = b, \underline{\theta}) \cdot \overline{V}_2^b(\beta_2) + \Pr(x = a|s_1 = s_2 = b, \underline{\theta}) \cdot \underline{\theta},$$

$$\Longleftrightarrow$$

$$q(1-\underline{\theta})^2 \left[\overline{V}_2^a(\beta_2) - \underline{\theta}\right] = (1-q)\underline{\theta}^2 \left[\overline{V}_2^b(\beta_2) - \underline{\theta}\right]$$

$$\Longleftrightarrow$$

$$q(1-\underline{\theta})^2 \left[\frac{p(1-\underline{\theta})}{p + (1-p)[\underline{\theta}^2 + 2\underline{\theta}(1-\underline{\theta}) + (1-\beta_2)(1-\underline{\theta})^2]}\right] = (1-q)\underline{\theta}^2 \left[\frac{p(1-\underline{\theta})}{p + (1-p)\underline{\theta}^2\beta_2}\right]$$

$$\Longleftrightarrow$$

$$\frac{p + (1-p)\underline{\theta}^2\beta_2}{p + (1-p)[\beta_2(2\underline{\theta} - \underline{\theta}^2) + (1-\beta_2)]} = \frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2}$$

$$\Longleftrightarrow$$

$$\frac{p + (1-p)\underline{\theta}^2\beta_2}{p + (1-p)[1 - \beta_2(1-\underline{\theta})^2]} = \frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2},$$

which is the condition stated in the proposition which implicitly defines the randomization probability (existence can be shown completely analogous to proposition 3.1).

Note that the left hand side is monotonically increasing in β_2 and takes on values between p for $\beta_2 = 0$ and $\frac{p+(1-p)\underline{\theta}^2}{p+(1-p)\underline{\theta}(2-\underline{\theta})} := \kappa$ for $\beta_2 = 1$. If the right hand side is larger than κ, $\beta_2 = 1$ is optimal, i.e. the agent always follows his signal. If the right hand side falls short of p the agent will optimally set $\beta_2 = 0$. This proves the optimality of the strategy stated in the proposition.

What remains to be shown is that $\beta_2 > \beta_1$. As $\frac{(1-q)\underline{\theta}^2}{q(1-\underline{\theta})^2} \geq \frac{(1-q)\underline{\theta}}{q(1-\underline{\theta})}$ it is sufficient to show that

$$\forall \beta : \frac{p + (1-p)\underline{\theta}^2\beta}{p + (1-p)[1 - \beta(1-\underline{\theta})^2]} \leq \frac{p + (1-p)\underline{\theta}\beta}{1 - (1-p)(1-\underline{\theta})\beta}.$$

Figure 2 illustrates this.

We have already seen that for $\beta = 0$ both terms are equal to p. Note that one can write the

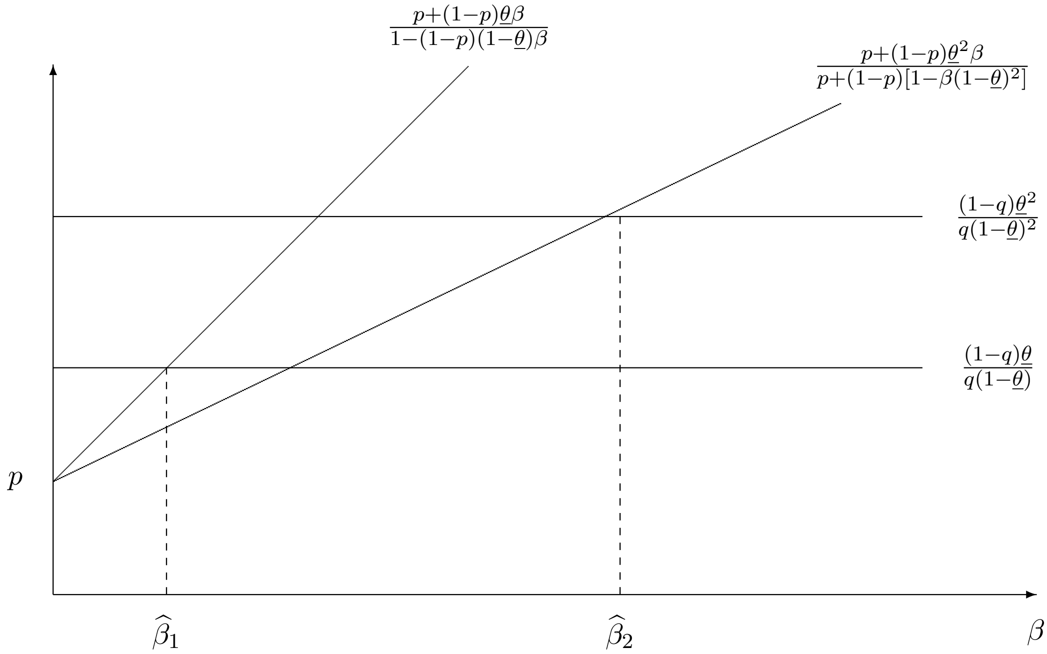

Figure 2: Illustration of the equilibrium

term on the left hand side also as

$$\frac{p + (1-p)\underline{\theta}^2\beta}{1-(1-p)\beta(1-\underline{\theta})^2}.$$

As the numerator is smaller and the denominator is larger compared to the respective terms on the right hand side it is now obvious to see that the condition is satisfied, hence $\beta_2 > \beta_1$. ||

Proof of proposition 4.2

The probability of a wrong decision in the first period is given by

$$q[(1-\underline{\theta})\beta_1] + (1-q)[(1-\underline{\theta}) + \underline{\theta}(1-\beta_1)] = q[(1-\underline{\theta})\beta_1] + (1-q)[1-\underline{\theta}\beta_1].$$

Whenever the true state is a then the wrong decision is only made if the agent receives an incorrect signal (which happens with probability $1-\underline{\theta}$) and chooses action b (which happens with probability β_1). If $x = b$ the agent chooses $d_1 = a$ if she obtained the wrong signal (which happens with probability $(1-\underline{\theta})$), but also with probability $(1-\beta_1)$ if she received the correct signal.

In an completely analogous way we derive the probability of a wrong decision in the second period as

$$q[(1-\underline{\theta})^2\beta_2] + (1-q)[1-\beta_2\underline{\theta}^2].$$

Therefore, the probability of a mistake in the second period exceeds the one in the first period if

$$q(1-\underline{\theta})[\beta_1 - (1-\underline{\theta})\beta_2] < (1-q)\underline{\theta}[\beta_1 - \beta_2\underline{\theta}].$$

From proposition 3.1 and 4.1 we can derive $\beta_1(q) = \frac{(1-q)\underline{\theta}-pq(1-\underline{\theta})}{\underline{\theta}(1-\underline{\theta})(1-p)}$ and $\beta_2(q) = \frac{(1-q)\underline{\theta}^2-pq(1-\underline{\theta})^2}{\underline{\theta}^2(1-\underline{\theta})^2(1-p)}$. Inserting , one obtains

$$q\frac{pq(2\underline{\theta}-1)(1-\underline{\theta})}{\underline{\theta}^2(1-p} > (1-q)\frac{(1-q)(2\underline{\theta}-1)\underline{\theta}}{(1-\underline{\theta})^2(1-p)},$$

which is satisfied whenever

$$\frac{q^2}{\underline{\theta}^2}p(1-\underline{\theta}) > \frac{(1-q)^2}{(1-\underline{\theta})^2}\underline{\theta}.$$

One can directly see that $\underline{\theta} < q$ is necessary for this condition to be satisfied. If $\underline{\theta}$ approaches $\frac{1}{2}$ the condition becomes $q^2p > (1-q)^2$ which is satisfied if p and q are large enough. $\|$

PROOF OF PROPOSITION 4.3

Suppose $q \leq \widehat{q}$ and define $\beta_3 := \Pr(d = b|s_1 \neq s_2)$. The reputational payoffs of the different actions are given by[22]

$$\begin{aligned}
\overline{V}_2^a(\beta_3) &= \Pr(\theta = 1|d_2 = a, x = a)\cdot 1 + \Pr(\theta = \underline{\theta}|d_2 = a, x = a)\cdot\underline{\theta} = \\
&= \frac{p}{p + (1-p)[\underline{\theta}^2 + (1-\beta_3)2\underline{\theta}(1-\underline{\theta})]}\cdot 1 + \frac{(1-p)[\underline{\theta}^2 + (1-\beta_3)2\underline{\theta}(1-\underline{\theta})]}{p + (1-p)[\underline{\theta}^2 + (1-\beta_3)2\underline{\theta}(1-\underline{\theta})]}\cdot\underline{\theta} \\
&= \frac{p + (1-p)[\underline{\theta} + \underline{\theta}(1-\underline{\theta})(1-2\beta_3)]\underline{\theta}}{p + (1-p)[\underline{\theta} + \underline{\theta}(1-\underline{\theta})(1-2\beta_3)]}
\end{aligned}$$

$$\begin{aligned}
\overline{V}_2^b(\beta_3) &= \Pr(\theta = 1|d_2 = b, x = b)\cdot 1 + \Pr(\theta = \underline{\theta}|d_2 = b, x = b)\cdot\underline{\theta} = \\
&= \frac{p}{p + (1-p)[\underline{\theta}^2 + \beta_3 2\underline{\theta}(1-\underline{\theta})]}\cdot 1 + \frac{(1-p)[\underline{\theta}^2 + \beta_3 2\underline{\theta}(1-\underline{\theta})]}{p + (1-p)[\underline{\theta}^2 + \beta_3 2\underline{\theta}(1-\underline{\theta})]}\cdot\underline{\theta} \\
&= \frac{p + (1-p)[\underline{\theta} - \underline{\theta}(1-\underline{\theta})(1-2\beta_3)]\underline{\theta}}{p + (1-p)[\underline{\theta} - \underline{\theta}(1-\underline{\theta})(1-2\beta_3)]}
\end{aligned}$$

Indifference implies that

$$q\overline{V}_2^a(\beta_3) + (1-q)\underline{\theta} = (1-q)\overline{V}_2^b(\beta_3) + q\underline{\theta}$$

[22] Notice that we make use of the fact that if the agent randomizes in case of mixed signals she will have an incentive to follow her signals if $s_1 = s_2$.

$$\Longleftrightarrow$$

$$q[\overline{V}_2^a(\beta_3) - \underline{\theta}] = (1-q)[\overline{V}_2^b(\beta_3) - \underline{\theta}]$$

$$\Longleftrightarrow$$

$$q\left[\frac{p(1-\underline{\theta})}{p+(1-p)[\underline{\theta}+\underline{\theta}(1-\underline{\theta})(1-2\beta_3)]}\right] = (1-q)\left[\frac{p(1-\underline{\theta})}{p+(1-p)[\underline{\theta}-\underline{\theta}(1-\underline{\theta})(1-2\beta_3)]}\right],$$

from which the proposition follows immediately. ‖

<u>PROOF OF PROPOSITION 4.4</u>

The probability of a mistake in the second period can be written as

$$(1-\underline{\theta})^2 + 2\underline{\theta}(1-\underline{\theta})[\beta_3(q)q + (1-\beta_3(q))(1-q)],$$

where from proposition 4.3 we obtain $\beta_3(q) = \frac{1}{2} - \frac{(2q-1)\mathbb{E}(\theta)}{2(1-p)(1-\underline{\theta})\underline{\theta}}$. The first term denotes the probability of receiving two wrong signals, while the second term gives the probability of getting two contradictory signals and choosing the wrong action.

Consider first the case where the bad agent does not randomize in the first period, i.e. $\underline{\theta} \geq q$. The error probability in the first period is then simply given by $1-\underline{\theta}$. If the error probability in the second period was higher it must be true that

$$1-\underline{\theta} < (1-\underline{\theta})^2 + 2\underline{\theta}(1-\underline{\theta})[\beta_3(q)q + (1-\beta_3(q))(1-q)]$$

$$\Longleftrightarrow$$

$$\frac{1}{2} < [\beta_3(q)q + (1-\beta_3(q))(1-q)]$$

$$\Longleftrightarrow$$

$$-\frac{(2q-1)q\mathbb{E}(\theta)}{2(1-p)(1-\underline{\theta})\underline{\theta}} + \frac{(2q-1)(1-q)\mathbb{E}(\theta)}{2(1-p)(1-\underline{\theta})\underline{\theta}} > 0,$$

which can never be the case given that $(1-q) < q$.

Let us now turn to the case where the agent randomizes in the first period. We now obtain the following expression for the error probability in the first period:

$$q(1-\underline{\theta})\beta_1(q) + (1-q)[(1-\underline{\theta}) + \underline{\theta}(1-\beta_1(q))] = (1-\underline{\theta}) - (\underline{\theta}-q)(1-\beta_1(q)).$$

When the agent receives a wrong signal if the true state is a, she chooses the wrong action with probability $\beta_1(q)$, while she will always make the wrong decision upon receiving the wrong signal in state b. Additionally, if the state is b the bad agent wrongly contradicts a correct signal with probability $(1 - \beta_1(q))$.

Inserting $\beta_1(q) = \frac{(1-q)\underline{\theta}-pq(1-\underline{\theta})}{\underline{\theta}(1-\underline{\theta})(1-p)}$, one obtains

$$(1 - \underline{\theta}) - (q - \underline{\theta})^2 \frac{\mathbb{E}(\theta)}{\underline{\theta}(1 - \underline{\theta})(1 - p)}$$

as expression for the error probability in period 1.

The likelihood of a wrongful decision is lower in period 1 if

$$(1 - \underline{\theta}) - (q - \underline{\theta})^2 \frac{\mathbb{E}(\theta)}{\underline{\theta}(1 - \underline{\theta})(1 - p)} \leq 2\underline{\theta}(1 - \underline{\theta}) \left[\frac{1}{2} - \frac{(2q - 1)^2 \mathbb{E}(\theta)}{2(1 - p)(1 - \underline{\theta})\underline{\theta}} \right]$$

$$\Longleftrightarrow$$

$$(1 - \underline{\theta})^2 \leq [(q - \underline{\theta})^2 - (2q - 1)^2] \frac{\mathbb{E}(\theta)}{\underline{\theta}(1 - \underline{\theta})(1 - p)},$$

which can never be satisfied since $(q - \underline{\theta})^2 < (2q - 1)^2$, so the right hand side is always smaller than zero. $\|$

References

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

Alonso, R., and N. Matouschek (2007): "Relational Delegation," *Review of Economic Studies*, p. forthcoming.

Avery, C. N., and J. A. Chevalier (1999): "Herding over the career," *Economics Letters*, 63(3), 327–333.

Banerjee, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

Battaglini, M. (2002): "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70(4), 1379–1401.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, 100(5), 992–1026.

Bolton, P., and M. Dewatripont (2005): *Contract Theory*. MIT Press.

Chevalier, J., and G. Ellison (1999): "Career Concerns Of Mutual Fund Managers," *The Quarterly Journal of Economics*, 114(2), 389–432.

Crawford, V. P., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–51.

Dasgupta, A., and A. Prat (2005): "Asset Price Dynamics When Traders Care About Reputation," CEPR Discussion Papers 5372, C.E.P.R. Discussion Papers.

——— (2006): "Financial equilibrium with career concerns," *Theoretical Economics*, 1(1), 67–93.

Dewatripont, M., I. Jewitt, and J. Tirole (1999a): "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies*, 66(1), 183–98.

——— (1999b): "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies," *Review of Economic Studies*, 66(1), 199–217.

DEWATRIPONT, M., AND J. TIROLE (1999): "Advocates," *Journal of Political Economy*, 107(1), 1–39.

EFFINGER, M. R., AND M. K. POLBORN (2001): "Herding and anti-herding: A model of reputational differentiation," *European Economic Review*, 45(3), 385–403.

GUL, F., AND L. RUSSELL (1995): "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy*, 103(5), 1039–1066.

HOLMSTRÖM, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182.

KOCH, A. K., AND E. PEYRACHE (2005): "Tournaments, Individualized Contracts and Career Concerns," IZA Discussion Papers 1841, Institute for the Study of Labor (IZA).

KRISHNA, V., AND J. MORGAN (2001): "A Model Of Expertise," *The Quarterly Journal of Economics*, 116(2), 747–775.

——— (2007): "Cheap Talk," in *New Palgrave Dictionary of Economics*. Palgrave-Macmillan.

LAUX, C. (2001): "Limited-Liability and Incentive Contracting with Multiple Projects," *RAND Journal of Economics*, 32(3), 514–26.

LEVY, G. (2004): "Anti-Herding and Strategic Consultation," *European Economic Review*, 48(3), 503–525.

——— (2005): "Careerist judges and the appeals process," *RAND Journal of Economics*, 36(2), 275–297.

MAJUMDAR, S., AND S. W. MUKAND (2004): "Policy Gambles," *American Economic Review*, 94(4), 1207–1222.

MORRIS, S. (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

OTTAVIANI, M., AND P. N. SORENSEN (2000): "Herd Behavior and Investment: Comment," *American Economic Review*, 90(3), 695–704.

——— (2001): "Information aggregation in debate: who should speak first?," *Journal of Public Economics*, 81(3), 393–421.

——— (2006a): "Professional advice," *Journal of Economic Theory*, 126(1), 120–142.

——— (2006b): "Reputational Cheap Talk," *RAND Journal of Economics*, 37(1), 155–175.

Prat, A. (2005): "The Wrong Kind of Transparency," *American Economic Review*, 95(3), 862–877.

Prendergast, C., and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104(6), 1105–34.

Scharfstein, D. S., and J. C. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80(3), 465–79.

Trueman, B. (1994): "Analyst Forecasts and Herding Behavior," *The Review of Financial Studies*, 7(1), 97–124.

Zhang, J. (1997): "Strategic Delay and the Onset of Investment Cascades," *RAND Journal of Economics*, 28(1), 188–205.

Zwiebel, J. (1995): "Corporate conservatism and relative compensation," *Journal of Political Economy*, 103, 1–25.