STRATEGIC CONSULTATION IN THE PRESENCE OF CAREER CONCERNS*

by

Gilat Levy
London School of Economics and Political Science

Contents:

Discussion Paper
No. TE/00/404
November 2000

The Suntory Centre
Suntory and Toyota International Centres
for Economics and Related Disciplines
London School of Economics and Political
Science
Houghton Street
London WC2A 2AE
Tel.: 020-7405 7686

---

* This paper is based on the first chapter of my dissertation from Princeton University. I thank Oriana Bandiera, Patrick Bolton, Allesandro Lizzeri, Meg Meyer, Marco Ottaviani, Ady Pauzner, Wolfgang Pesendorfer, Michele Piccione, Yossi Rapaport, Ronny Razin, Ran Spiegler, and especially my adviser Gene Grossman, for many valuable comments. I also thank seminar participants at Essex, Keele, LSE, Princeton, Madison-Wisconsin, Cornell, BU, Columbia and SITA 98. Financial support from the Sloan Dissertation Fellowship is acknowledged.

Abstract

In this paper I analyse the strategic interaction of decision makers and their advisers in a consultation process. I find that when agents are concerned about their reputation, consultation results in sub-optimal sharing of information; some decision makers may deliberately act unilaterally and not consult even when advice is costless. When they do consult, decision makers may excessively contradict their adviser's recommendation. Anticipating it, advisers may not report their information truthfully. These results are obtained without assuming either a tournament or a competition between decision makers and their advisers for wages or a future job.

Keywords: Reputation; consultation; relative performance evaluation.
JEL Nos.: D82, D83, G20

1 Introduction

Ever since the days of Ahithophel, consultation has been common in almost any organization. The use of advisers stems, of course, from the need of decision makers to gather information. The process of consulting, central as it is to the aggregation of information, has crucial impact on the efficiency of decision making. In the presence of career concerns, however, the process of consulting does not involve only the transmission of information. It also has bearing on the reputation of both the adviser and the decision maker. Seeking advice, giving advice and using advice can, and do, become strategic actions, inevitably compromising the efficiency of the decision making process as a whole.

This is a casual observation, which turns up in empirical studies of organizations. Blau (1954) interviewed employees of a federal regulation agency, who were encouraged to consult their experienced supervisor. However, the agents were reluctant to do so, since consulting their supervisor exposed them as incompetent or stupid. A similar phenomenon was observed by Allen (1977), in a study of engineers in private firms. The engineers reported that they prefer to rely on professional literature rather than on their colleagues' experience. They did acknowledge that reading the literature is a relatively inefficient method of gathering information; it takes more time to read through. But they avoided consulting a colleague, for this implied an admission of their inferiority. For a decision maker, choosing not to consult, or to disregard advice, has its benefits. Staw and Ross (1982) found that special praise is accorded to decision makers that stick to their views and ignore advice, a phenomenon they termed the 'hero' effect.

In this paper, I investigate the effect of career concerns, defined here as an agent's attempt to prove her competence to an evaluator, on the efficiency of internal consultation processes within organizations. The reaction of careerist managers to other agents' actions or messages is the subject of several theoretical and empirical studies, such as

Scharfstein and Stein (1990), Trueman (1994), and Chevalier and Ellison (1999). Similarly, Prendergast and Stole (1996) analyze how managers in a market react to common public prior. However, consultation within an organization is different from the transmission of information across firms. In particular, at market environment managers acquire information by observing the unsolicited actions of other agents. By contrast, within an organization managers can *choose* whether to solicit additional information from their colleagues. Even more significantly, internal communication between agents may be unobserved by an evaluator. Thus, actions by one agent may contain information about the ability of another agent, and can give rise to a strategic interaction between agents - an interaction that does not arise in markets.

Let us consider as an example the following situation. A manager in a venture capital firm is required to decide whether to invest in a new company. The manager collects the relevant information, such as the new company's business plan and background literature on relevant markets. Based on her interpretation of this data, she forms an opinion regarding the profitability of the investment. In a situation of uncertainty, however, she could ask for the advice of an assistant or a colleague. The adviser provides a recommendation based on his own independent research and interpretation. Now, if the manager is only motivated by outcome concerns, she should gather costless advice and follow it as long as it is more accurate than her own information. Assume, however, that the manager is not only interested in reaching the correct decision, but also in demonstrating her ability to an evaluator, be it her superiors or the firm's shareholders. This paper examines whether, in the presence of career concerns, the manager will still choose to seek advice, whether advisers will provide truthful information, and whether the advice will be used efficiently by the manager.

A simple model is introduced in order to answer these questions. A decision maker has a private signal about the state of the world, and the accuracy of this signal depends on her ability. An adviser also has private information whose accuracy depends on his ability. The manager can choose whether to consult, and eventually takes an action given the information she gathers. The adviser can choose which information to transmit. An evaluator observes the state of the world, whether the decision maker chose to consult, the direction (although not the accuracy) of the adviser's recommendation, and finally, the action taken. On the basis of these observations, the evaluator can assess the quality of the adviser and the decision maker's private information, which reflects their respective abilities. The decision maker and her adviser are concerned about the evaluator's assessment of their ability as well as about the outcome of the action taken. Thus, the decision maker's actions form a costly signal on her type. The adviser's message is a cheap talk message to two audiences, fully observed by the decision maker and partially by the evaluator.

The main results of the model are as follows. First, I consider the decision maker's choice whether to consult an adviser. I find an equilibrium in which able decision makers prefer to act alone, while lower quality decision makers consult an adviser. An able decision maker limits the amount of information that she gathers, even if it is costless, because it allows her to signal confidence in her private information. This equilibrium is sustained when decision makers are sufficiently concerned about the outcome. If outcome concerns are weak, even less able decision makers will prefer not to consult. This is so because they have no incentive to distinguish themselves from their more able counterparts.

Second, I consider the decision maker's use of the information provided by an adviser. In equilibrium, all types of decision makers contradict their advisers excessively and may take actions that run counter to their own information. This behavior may be termed 'anti-herding'.[1] This occurs because in equilibrium a *relative performance eval-*

[1]Herding is defined as the behaviour of decision makers who follow others' recommendations while ignoring their own information (Scharfstein and Stein (1990), Banerjee (1992)).

uation endogenously arises between the decision maker and her advisers. The decision maker has an incentive to signal her own *absolute* ability by maintaining to be better *relative* to the adviser, even though they are not in direct competition.[2] Contradicting advice is a possible signal only when decision makers know their type. Otherwise, decision makers can only signal their type by taking the correct decision, which leads to herding behavior (see Trueman (1994), Ottaviani and Sørensen (2000)).

Third, I allow for strategic behavior by the advisers. In equilibrium advisers motivated by career concerns do not provide truthful information. This also is a result of the relative performance evaluation that arises between decision makers and their advisers. When a decision maker signals that she is able, she also reveals as a by-product that her adviser is relatively less able. The adviser, who sends a cheap talk message, anticipates this endogenous conflict and is unable to transmit truthful information credibly.

To summarize, the main finding of this paper is that consultation in the presence of career concerns produces sub-optimal sharing of information. Decision makers tend to act alone, and even when they do consult, they excessively ignore advice. Moreover, advisers do not provide truthful information. The quality of advice is reduced and advice is misused. These results imply that a share-holder or a principal who tries to identify the most able managers to ensure better decision making in the future, faces an intertemporal trade-off, since current management becomes less effective. In the concluding section I discuss how this trade-off can determine the optimal size of a committee or a team of decision makers.

This paper proceeds as follows. In the next section I construct the model while assuming that the adviser reports his information truthfully. Section 3 allows for strate-

[2] Effinger and Polborn (1998) examine relative performance evaluation in this context, but they assume the advisers to be in direct competition for the same job. In my model, relative performance evaluation arises even though the decision maker and her advisers are not engaged in any competition or tournament.

gic advisers. Section 4 discusses an extension and section 5 concludes. All proofs are relegated to an appendix.

2 The model

Consider a decision maker D, who must take an action, a. For simplicity, suppose there are only two possible actions, $a \in \{l, h\}$. This could represent a low or a high level of investment. The appropriateness of each action depends on the state of the world ω, which is unknown and is realized only after the action a is taken. There are two possible states, which, for the sake of brevity, are also denoted by l or h. It is common knowledge that each state can occur with equal probability. Let action l be appropriate in state l and action h be appropriate in state h.

The decision maker receives a private signal $s \in \{l, h\}$ about the state of the world. The informativeness of the signal depends on the decision maker's ability. Let p represent the ability (or type) of D. Decision makers differ in their abilities. This is reflected in variation in the accuracy of their private information. In particular, the more able is D, the more likely it is that her signal is reliable. That is, the probability that $s = \omega$ rises with p; specifically, $\Pr(s = \omega \mid \omega) = p$. D knows p, which is drawn from a uniform distribution on $[.5, 1]$.

After the decision maker receives her signal but prior to taking an action, she can choose to consult an adviser in order to gather another opinion.[3] By an adviser, I mean an agent who has information about the state of the world, but no control rights over decision making. The adviser, denoted by A, can therefore represent one of the manager's assistants or a colleague in the research department of the firm. Like the decision maker, A receives a private signal $s^A \in \{l, h\}$. This signal reflects the way *he* interprets the state of the world. The signal is accurate with probability p^A, that is, $\Pr(s^A = \omega \mid \omega) = p^A$. Let

[3] I assume that advice is costless. Time and resource constraints are caputred by the assumption that the decision maker consults only one adviser.

p^A be drawn from a uniform distribution on $[.5, 1]$. The signals of the decision maker and the adviser are conditionally independent. In order to focus on the behavior of careerist decision makers, I assume in this section that the adviser transmits his information truthfully. The next section relaxes this assumption and allows for a strategic adviser. Thus, if the decision maker decides to consult, A transmits (s^A, p^A) privately to D. Finally, given the information that she had accumulated, the decision maker takes her action $a \in \{l, h\}$. Figure 1 presents the extensive form of the model.

Figure 1: The organizational form

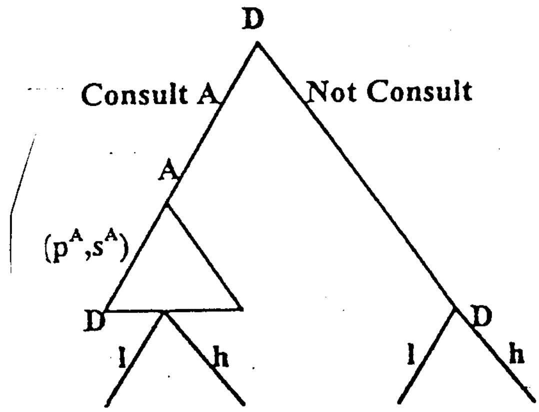

D gains utility both from making an appropriate decision and from advancing her career. For a meaningful description of career concerns, assume that there exists another player, namely an evaluator E. If E believes that D is an able decision maker, she can be rewarded with high wages or a promotion. We can think of E as representing a higher-level executive or a share-holder in the firm who can promote the middle-level manager D.

The decision maker tries to signal to E that she is able. E, on the other hand, tries to guess the decision maker's true type. E knows that p is distributed uniformly on $[.5, 1]$. I assume that E observes the action a chosen by D, the realized state of the world ω and whether the decision maker consulted or not.[4] I further assume that the

[4]There are, obviously, some consultations which a manager or a decision maker can conduct secretly. Here, however, we consider the effect of public consultations.

evaluator can, at least partially, observe the adviser's report. This is likely to occur if the advice is a public document, such as a memo issued by one of the manager's assistants. Assume, for simplicity, that the evaluator can observe the recommendation, s^A, but not p^A. In other words, E learns only whether the adviser was "for" or "against" D's action but not the adviser's level of confidence in his recommendation. The evaluator understands only the bottom line of a recommendation and not the lengthy explanation of its accuracy. This is justified by either communication costs, bounded rationality, or professional terminology that an evaluator is unfamiliar with.[5] Given these observations, E updates his beliefs about p rationally. Denote $E's$ posterior beliefs about p by π.

Another useful piece of notation is an indicator variable I, where $I = 1$ if the action taken by the decision maker is correct, i.e., $a = \omega$, and $I = 0$ otherwise. *The objective of D is to maximize $\pi + \theta I$.* The parameter θ, $\theta > 0$, represents how much D cares about taking the correct action in the current decision. θ might reflect her discount rate; that is, her preference for current income (e.g., if she gets a bonus for making the correct decision) relative to future rewards.

The players' strategies, and subsequently, the structure of the game, are as follows. Given (s, p), $D's$ strategy prescribes whether she consults A or not (recall that the decision maker does not observe the adviser's signal or type prior to consultation). Then, if she consults, her strategy determines which action a she takes given (s, p) and (s^A, p^A). If she does not consult, her strategy determines which action a she takes conditioned only on (s, p). The evaluator observes a, ω, whether D has consulted, and s^A if she does. He then uses a belief function π to update his beliefs about the type p of D. The decision maker and the evaluator are therefore the only players in the game.

I use the concept of a Perfect Bayesian Equilibrium to solve the model. PBE

[5] The assumption that s^A is observed by the evaluator, is an important assumption. I will discuss later ways to relax this assumption. On the other hand, the assumption that p^A is not observed is not crucial. I will discuss later how the results are slightly modified when p^A is observed.

means that $D's$ strategies, conditioned on her information, are optimal at any decision node given $E's$ beliefs and that the evaluator's belief function π is derived from $D's$ strategy using Bayes rule whenever possible. For the sake of interest, I will analyze only informative equilibria, i.e., equilibria in which the strategy of the decision maker is responsive to her signal. Moreover, I will ignore any "mirror" equilibrium, i.e., an equilibrium that takes an original equilibrium and switches each action from l to h and vice versa.

The equilibrium strategies can be described by cutoff points. Given a signal s, any equilibrium will define the set of types p who consult, the set of types p that follow their signal when they do not consult, and the set of types p that follow their signal given (s^A, p^A). Since the situation that I analyze is symmetric in the two states of the world, I examine equilibria in symmetric strategies. This means that if a decision maker of type p consults her adviser given s, she also consults her adviser given $s' \neq s$. Similarly, if a decision maker of type p does not consult and follows (contradicts) a signal s, she also follows (contradicts) s'. And so on.[6]

To analyze the equilibria of the model we will use backward induction, i.e., we will first analyze the decision maker's actions in the 'unilateral' subgame, i.e., the continuation game given that she had not consulted. We will then analyze her actions in the continuation game subsequent to a consultation (since the evaluator observes the decision whether to consult or not, these two continuation games are proper subgames and we can analyze their PBE). Finally, we will analyze how the decision maker chooses whether to consult or not.

[6]Formally, $D's$ consultation strategy is defined by $c : (s,p) \to \{0,1\}$ where 1 signifies a conusltation, $D's$ action strategy upon conusltation is defined by $a_c : (s,p,s^A,p^A) \to \{l,h\}$ and $D's$ action strategy given no consultation, is defined by $a_{nc} : (s,p) \to \{l,h\}$. Note that since the strategies can be described by cutoff points, mixed strategies are irrelevant. Then, for $s' \neq s$ and $a' \neq a$, symmetric strategies satisfy $c(s,p) = c(s',p)$, $a_{nc}(s,p) = a'_{nc}(s',p)$ and $a_c(s,p,s^A,p^A) = a'_c(s',p,s^{A\prime},p^A)$.

I will concentrate on checking whether the decision maker behaves efficiently, i.e., whether she maximizes the probability that the right decision is taken. An efficient strategy has D consulting A and taking the action that corresponds to the adviser's report, unless $s^A \neq s$ and $p^A \leq p$. That is, the decision maker should always consult and follow the adviser unless her own information is different and more reliable. An alternative social goal may be the evaluator's objective, which is to identify able decision makers. If the process of identifying able decision makers induces the latter to distort their actions, these two goals may be in a conflict. We will now check whether efficient behavior can be sustained as an equilibrium phenomenon or whether such a conflict arises.

2.1 The unilateral subgame

Consider the subgame in which the decision maker has to take an action unilaterally, i.e., she had not consulted her adviser. The only information she can condition her decision on is (s, p). By Bayes rule, $\Pr(\omega = s|s, p) = p \geq .5$. Thus, the efficient action in this subgame is to follow the signal s.

The evaluator, who observes that the decision maker had not consulted and plans to take an action unilaterally, believes that $D's$ type belongs to some set of types $U \subseteq [.5, 1]$. He can further refine his beliefs given the action a and the state ω. By the assumption that $D's$ strategies are symmetric, the only important information for the evaluator is whether $D's$ action is successful or not. The evaluator has to use only the indicator I in order to update his beliefs. The belief updating function π can be written as $\pi : I \rightarrow U$. Thus, $\pi(1)$ denotes the beliefs of E when D takes the correct decision and $\pi(0)$ denotes the beliefs of E on p when D fails to take the correct decision.

The decision maker has to decide, given her type p, whether to follow her signal or to contradict it. If she follows s, her expected utility is $p\pi(1) + (1 - p)\pi(0) + p\theta$ whereas by contradicting she obtains $(1 - p)\pi(1) + p\pi(0) + (1 - p)\theta$. She therefore follows her signal if

$$p\pi(1) + (1-p)\pi(0) + p\theta \geq (1-p)\pi(1) + p\pi(0) + (1-p)\theta$$

The first lemma asserts that in equilibrium this condition is satisfied.

Lemma 1. *When the decision maker does not consult there exists a unique equilibrium in which every type of D follows her signal.*

In equilibrium, when the decision maker's action is the correct one, it implies that her private information is relatively accurate. As a result, the Bayesian updating of the evaluator yields $\pi(1) > \pi(0)$. Taking the correct action becomes a signal that the decision maker is able.

Thus, despite the career concerns of D, the unique equilibrium in this subgame is efficient. The reason is that D does not possess a rich enough set of tools to allow for distortionary signaling. The only way with which she can prove her ability is by taking the correct action, an incentive which coincides with her outcome concerns.

2.2 The consultation subgame

I now analyze how the decision maker's choice of action responds to advice. More specifically, I check whether each of the decision maker's types puts proper weight on the advice that she receives, when considering which action to take. Recall that the efficient course of action is to follow the adviser unless $s^A \neq s$ and $p^A \leq p$.

The evaluator, who observes that the decision maker had consulted her adviser, believes that $D's$ type belongs to some set of types $C \subseteq [.5, 1]$. To further refine his beliefs, the evaluator has to consider how decision makers of different abilities use the information conveyed to them by their advisers. Formally, define another indicator, I^A, where $I^A = 1$ if the adviser's recommendation proves correct about the true state of the world, and $I^A = 0$ otherwise. By restricting $D's$ strategies to be symmetric, the updating function π depends only on whether the decision is correct and whether the adviser has agreed with it, i.e., $\pi : I \times I^A \to C$. For example, $\pi(1,1)$ are the beliefs of E if the decision maker took the correct action ($I = 1$) after the adviser had provided a correct recommendation ($I^A = 1$). Similarly, $\pi(1,0)$ means that D took the correct action despite faulty advice. And so on. When the decision maker's action and the adviser's recommendation coincide, we say that the decision maker *followed* the adviser's recommendation. When the action and the recommendation do not coincide, we say that the decision maker *contradicted* her adviser.

What is the equilibrium behavior of the decision maker in this subgame? Some intuition suggests the following. If the decision maker takes an action that is similar to the adviser's recommendation, she reveals that she might have relied on his recommendation and therefore that his information may have been more accurate than her own. The evaluator may therefore infer that the decision maker's information is less accurate than the average. If, however, the decision maker contradicts her adviser's recommendation, she suggests that her own information was at least as accurate as his. The evaluator may believe in this case that the decision maker has access to more accurate information than the average. Thus, a decision maker could signal her own absolute ability by claiming (by her actions) to be superior to her adviser.

This argument suggests that the decision maker may have an excessive incentive to contradict her adviser. Let us consider for which recommendations (s^A, p^A) of the adviser a decision maker with private information (s, p) would choose to contradict advice. There are two cases to consider: (i) the adviser's signal opposes that of the decision maker; (ii) the adviser's signal coincides with that of the decision maker.

Suppose that $s^A \neq s$. Then D will opt to contradict her adviser if the expected utility from doing so exceeds that from following his advice, i.e., if:

$$\Pr(\omega = s | s, p, s^A, p^A)(\pi(1,0) + \theta) + \Pr(\omega = s' | s, p, s^A, p^A)\pi(0,1) \geq$$
$$\Pr(\omega = s | s, p, s^A, p^A)\pi(0,0) + \Pr(\omega = s' | s, p, s^A, p^A)(\pi(1,1) + \theta)$$

where by Bayesian updating,

$$\Pr(\omega = s|s, p, s^A, p^A) = \frac{p(1-p^A)}{p(1-p^A) + (1-p)p^A}.$$

Rearranging terms, I find that D contradicts her adviser if

$$p^A \leq \frac{kp}{kp + (1-p)}, \tag{1}$$

where $k \equiv \frac{\pi(1,0)-\pi(0,0)+\theta}{\pi(1,1)-\pi(0,1)+\theta}$. Loosely speaking, k measures the relative gain of contradicting advice vis-a-vis following it. When k is positive, the RHS of equation (1) increases with p, implying that relatively able decision makers will more often contradict their advisers. These able types have the greatest incentive to contradict an adviser, because most of the adviser's recommendations are less accurate than their own. This is the reason why contradicting the adviser may become a signal of the decision maker's ability.

Now suppose that $s^A = s$. In this case, when the decision maker contradicts her adviser, she also acts counter to her own information. Here, all available evidence points to s as the appropriate action. Why then would the decision maker ignore this preponderance of evidence? The answer is that the reputation gains from contradicting the adviser may be higher than the gains from taking the seemingly correct action while appearing to follow the adviser's lead. Thus, D contradicts the adviser if

$$\Pr(\omega = s|s, p, s^A, p^A)\pi(0,1) + \Pr(\omega = s'|s, p, s^A, p^A)(\pi(1,0) + \theta) \geq$$
$$\Pr(\omega = s|s, p, s^A, p^A)(\pi(1,1) + \theta) + \Pr(\omega = s'|s, p, s^A, p^A)\pi(0,0)$$

I can rearrange terms again to show that D contradicts her adviser if

$$p^A \leq \frac{k(1-p)}{k(1-p) + p}. \tag{2}$$

When $k > 1$, some types of decision makers satisfy equation (2) and contradict the adviser even though their own information agrees with the adviser's recommendation. Since the RHS of equation (2) decreases with p, the *less* able is the decision maker, the more often she contradicts the adviser. The reason that the less able decision makers have the greatest incentive to contradict advice when $s = s^A$ is that the information these decision makers have access to is least decisive. In other words, these less able decision makers are not confident about what is the correct action to take, so they opt to contradict advice in order to mimic able decision makers who trust their own information and usually ignore advice.

Of course, the evaluator, being rational, can anticipate this behavior and may realize that a decision maker who contradicts advice is not necessarily an able one. Moreover, if he observes a decision maker who contradicts advice while taking the wrong action, it may be more probable that she is a less able type who has satisfied equation (2). This, in return, discourages all types of decision makers from contradicting any recommendation, and fixes a level of contradiction that can be sustained in equilibrium.

An equilibrium in the consultation subgame is defined by equations (1) and (2) and a fixed point of the equation $k = \frac{\pi_k(1,0)-\pi_k(0,0)+\theta}{\pi_k(1,1)-\pi_k(0,1)+\theta}$. These equations capture the requirements that D behaves optimally given the beliefs of the evaluator, and that the beliefs of the evaluator are derived from $D's$ strategy which is characterized only by k. We are now ready to characterize the equilibrium:

Proposition 1. *When the decision maker consults her adviser, any equilibrium is characterized by k>1, i.e., decision makers contradict their advisers excessively and sometimes act counter to their own information as well.*

The equilibrium strategy is illustrated in Figure 2a for the set of types $C = [.5, 1]$. It therefore describes the equilibrium when the decision maker always consults. For comparison, the efficient strategies are depicted in Figure 2b (the efficient strategies correspond to $k = 1$). In these figures, I depict the strategy for each type p given the signal s (by symmetry it is enough to describe it for one signal only). The $x-axis$ depicts the decision maker's ability $p \in [.5, 1]$. The $y-axis$ depicts the possible recommendations of the adviser so that as we go up the $y-axis$, the probability that $\omega = s$ increases; In the upper half of the figure, the adviser recommends $s^A = s$, for variable reliability levels p^A, which range from .5 to 1. In the lower half of the figure, the adviser recommends $s^A = s' \neq s$, i.e., he receives a signal which opposes the decision maker's signal but as we go up the $y-axis$, the reliability level of the adviser decreases from 1 to .5. The figures show that any type of the decision maker contradicts advice excessively compared to the efficient strategies but that the less able types distort their actions more often. Hence, the more able is the decision maker, the more likely it is that she takes the correct action.

Figure 2a: Equilibrium strategies

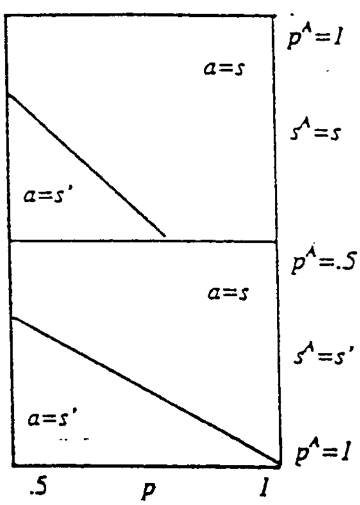

Figure 2b: The efficient strategies

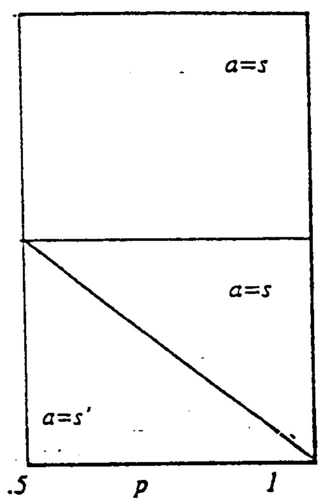

In the equilibrium of the consultation subgame, contradicting an adviser while taking the correct action is a signal for high ability. This is so because able decision makers are likely to receive more accurate signals than their advisers. On the other hand, taking the correct action enhances reputation regardless of the adviser's recommendation because it indicates that the decision maker has received, on average, a correct signal.[7] As

[7]The equilibrium feature that reputation is greater when the decision maker takes the correct action, implies that even if the decision maker has no outcome concerns and her actions are cheap talk messages on her type, i.e., when $\theta = 0$, informative actions are possible in equilibrium and Proposition 1 holds as well. The information held by the evaluator about the state of the world disciplines the decision maker and creates endogenous costs to different actions.

a result, the decision maker faces a trade-off; contradicting advice pays off if she manages to make the right decision, but is also risky. The risk increases with the adviser's accuracy. On the other hand, if she follows advice, she utilizes more information and is more likely to make the right decision. But, she also forgoes the opportunity to mimic able decision makers who tend to ignore advice.

The solution to this trade-off depends on the relative abilities of the decision maker and her adviser. Able decision makers follow any advice when p^A is sufficiently high and contradict advice when $s^A \neq s$, for relatively low p^A. On the other hand, the less able decision makers follows any advice if p^A is sufficiently high and contradict *any* advice if p^A is relatively low, since in this case they have a good chance to take and right action and ignore advice at the same time.

We therefore reach the following conclusions. Since contradicting advice signals high ability in equilibrium, all types of decision makers excessively contradict their advisers. Decision makers are tempted to show their ability by taking an action that opposes their advisers' recommendation. Contradicting advice can arise as a signal on ability only when decision makers know their type (as in Trueman (1994)). Otherwise, they cannot claim that it is better than someone else's type. In other words, when decision makers do not know their type, they can signal it only by taking the right decision. This incentive induces them to follow informative recommendations while ignoring their own information, which is termed herding behavior (see Scharfstein and Stein (1990), Ottaviani and Sørensen (2000)). Accordingly, I find that decision makers who know their type, may engage in 'anti-herding' behavior. When advice is relatively not reliable, less able decision makers contradict any recommendation while ignoring their own information.

In fact, figure 2a reveals that both herding and anti-herding may coexist in equilibrium. Since taking the correct decision still enhances reputation, the trade-off between the reputational benefits of taking the correct action and contradicting advice determines the level of equilibrium herding and anti-herding, as exercised by the less able decision makers.

Remark 1: What if the evaluator can observe p^A? this situation corresponds for example to the case in which the evaluator and the adviser are one agent, e.g., if the manager has to consult her supervisor, who is also responsible for her promotion. I can show that in the unique equilibrium, the decision maker excessively contradicts her adviser, but only when $s \neq s^A$. When $s = s^A$, as opposed to the result in Proposition 1, advice is always followed. To see the intuition for this altered result, note that when p^A *is* observed, it is only the less able types who contradict advice when $s = s^A$. Thus, when p^A is not observed, these types find it harder to pool themselves with the more able types who contradict advice only when $s \neq s^A$.

Remark 2: The public nature of recommendations, i.e., the ability of the evaluator to observe s^A, is a key assumption in this model. If E cannot observe s^A, the decision maker can signal her type only by taking the correct decision, which induces her to behave efficiently. I can weaken the observability assumption and assume instead that the recommendation is verifiable (i.e., 'hard' evidence which can be concealed but not misrepresented) and that it is revealed to the evaluator only if the decision maker wishes to. I can then show that there exist equilibria of the modified model with full revelation of recommendations: if the decision maker contradicts the adviser, she has an incentive to reveal the advice given to her because it indicates her superiority. This corresponds to a manager presenting her case in the share-holders or higher-level executives meeting, showing the memo of her assistant, with its bottom line: "we should not invest", while claiming: "I think that we should invest". A decision maker who does not reveal what she was told, is perceived as an advice-follower in this equilibrium. Thus, all relevant information can be extracted by the evaluator and our results are sustained.

The main result of the consultation subgame is that a *relative performance evaluation* competition emerges endogenously between the decision maker and her adviser. This competition creates an incentive for the decision maker to downplay the contribution of the adviser to the joint decision making process. This allows her to claim that her own information was sufficient and decisive. This endogenous competition will turn out to affect also the adviser's recommendations when we allow for a strategic adviser. But first, we complete the analysis of the model by proceeding to analyze the decision maker's choice whether to consult.

2.3 Choosing whether to consult

In the first stage of the model, the decision maker has to decide whether to consult her adviser or not, anticipating her equilibrium behavior and rewards in the subsequent stages. The previous analysis showed that when consulting, the decision maker distorts her actions in order to signal her type. We should therefore ask whether consulting an adviser is the efficient action at the first stage of the game, or maybe the probability of taking the right decision is higher in the unilateral subgame. The answer to this question is ambiguous, since the level of distortion depends on the value of k, which in turn depends on which types consult an adviser in equilibrium. The question of whether consulting is more efficient than acting on your own is endogenous in the model. However, for sufficiently high values of θ, I can verify (Lemma 3 in the appendix) that in any equilibrium the probability with which the decision maker takes the correct action is still higher in the consultation subgame, even though she abuses the adviser's information. Thus, if outcome concerns are strong enough, the efficient action is to consult the adviser. In order to compare equilibrium results with efficiency, I will restrict the analysis in this section to this range of values of θ.

An equilibrium in the model is defined by an equilibrium in each of the subgames as described in the previous analysis, calculated given some beliefs of E that a type p

that consults belongs to the set $C \subseteq [.5, 1]$ and a type p that acts unilaterally belongs to the set $U = [.5, 1] \backslash C$, and by the condition that a decision maker of type $p \in U$ prefers to act unilaterally and a decision maker of type $p \in C$ prefers to consult, anticipating her expected utility in each of the subgame's equilibrium.

The goal in this section is to determine whether the act of "going it alone" can emerge as an endogenous signal about the decision maker's ability. Why might the decision by D not to consult signal high ability? If D decides not to gather additional information, she reduces the likelihood of taking the appropriate action. But taking the correct action based only on her own information, is a more challenging test than doing so after collecting advice. The willingness of a decision maker to impose this test upon herself may show that she puts faith in her own private information.

Once D consults an adviser, she can no longer credibly claim that she could as easily have relied on her own information. However, as I showed previously, by challenging the adviser's suggestion, she may indicate that she had access to information at least as good as that of the adviser. Thus, D's decision whether to consult or not and her choice of how to use the advice she obtains can both serve as signaling devices. The question that arises is which form of signaling emerges in equilibrium. To answer this question, the next Proposition analyses equilibria in which some decision makers consult and some do not.

Proposition 2. *For sufficiently high values of θ, any equilibrium in which some types of decision makers consult advisers and others do not, is characterized by a cutoff point $\hat{p} \in (.5, 1)$. If $p \geq \hat{p}$, D does not consult and follows her own signal. If $p < \hat{p}$, D consults an adviser but excessively contradicts his advice.*

Figure 3: The equilibrium strategies

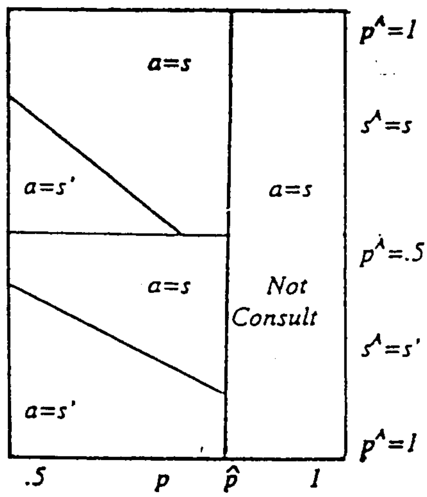

Both types of signals emerge in equilibrium; the decision maker distorts her actions by not consulting and by contradicting advice. However, these signals are also endogenously ranked. In particular, failing to consult an adviser is a signal that the decision maker is one of the most able ones, whereas contradicting advice signals moderate ability.

The reason for this ranking is as follows. If reputation for acting unilaterally is low, then none of the decision maker's types would follow this route, which is inferior to consulting both in terms of career and outcome concerns. However, if a decision maker who acts unilaterally is highly praised, then able decision makers are willing to acquire this reputation. This comes to them only at a small cost; the probability that they would fail to take the right decision is limited because they have precise enough private information. This insures the consistency of the evaluator's beliefs and establishes the act of "going it alone" as the most powerful signal on ability.

Contradicting an adviser provides another signal, which is, endogenously, ranked second. The decision maker consults an adviser because she does not feel confident enough to take the risk of making a decision without first checking the adviser's information. But, once the decision maker learns the adviser's recommendation, she sometimes can claim to have the better information. Thus, contradicting an adviser provides an *indirect* signal; when D consults an adviser, she does not know ahead of time whether she will end up following his recommendation or contradicting it.

As stated in the Proposition, the equilibrium is sustained only when outcome concerns are sufficiently meaningful. These concerns provide an incentive for the less able types to aggregate more information. Otherwise, these types would not be willing to distinguish themselves from the more able types and will rather mimic them by acting unilaterally.

Remark 3: In the equilibrium characterized in Proposition 2, the more able types, who act unilaterally, behave inefficiently. As a result, some of them take the right decision with a lower probability than their lesser able counterparts. To see that, note that type $\hat{p}$, who is indifferent between consulting and acting unilaterally, receives higher reputation when she does not consult than when she does. To counter that, she must gain utility in the form of θ and take the right decision with a higher probability when she consults. This means that in the neighborhood of $\hat{p}$, the types below $\hat{p}$ take the right action with a higher probability than the types above $\hat{p}$. This finding seems to be inconsistent with the evaluator's objective to isolate the more able types; it is not necessarily the case that these able types take better decisions. However, I believe that we should view the problem raised in this paper in an inter-temporal context. Career concerns probably exist when managers are young and inexperienced. As time goes by, more information is accumulated about these managers and their type is revealed. Their incentive to prove it becomes superfluous. Thus, an evaluator who tries to identify able decision makers knows that his objective may induce these able decision makers to take worse decisions than their lesser able colleagues in the short term. But, that it would prove beneficial in the long term, once career concerns dissolve. I will further discuss this inter-temporal trade-off, and ways to resolve it, in the concluding section.

Remark 4: The analysis reported in Proposition 2 had focused on the range of high values of θ. Although we cannot tell what is the efficient course of action when θ is low, i.e., whether it is consulting or not, we know (see claim 1 in the proof of Proposition 2) that the results are not reversed for low values of θ. Namely, it is not the case that consulting an adviser indicates the highest ability.

The analysis had also focused on equilibria in which some decision makers consult and some do not. Can other equilibria exist? An equilibrium in which all types of decision makers refrain from consulting can exist only for low values of θ. Otherwise, the least able types prefer to gather information to increase the probability that they make the correct decision, even at the cost of losing some reputation. On the other hand, an equilibrium in which all types consult exists for all values of θ. For example, if the evaluator believes that a decision maker who fails to consult is the least able type, then all types of D are forced to seek advice to avoid the harsh judgment (and loss of valuable information). The equilibrium strategies when all types consult are as depicted in figure 2a. However, for sufficiently high values of θ, this equilibrium does not survive the intuitive criterion. Able types would find it profitable to deviate and not consult because $E's$ beliefs must dictate that only able types would take the risk of doing so.

The main prediction of this section is therefore that decision makers do not necessarily aggregate the efficient amount of information. The decision-maker may limit the amount of information that she gathers, even if this information is costless. Casual observations suggest that sometimes managers and politicians are known to act on their own. The model explains why they choose to do so; these "authoritarian" decision makers are rewarded with high reputation for being capable.

The model may also suggest that in a more general setting, the number of consulted advisers can serve as a signal on managerial ability. It is reasonable to conjecture that the more able managers would consult fewer advisers, and by doing so, show off their reliance on their own private and accurate information. Thus, even if it would be possible in terms of resources to consult an unbounded number of advisers, decision makers would restrict the number of advisers they consult and thus under-utilize readily available information.

3 Careerist advisers

Advisers themselves may be motivated by career concerns. If the advisers are the employees of the research department of the firm or of a consulting firm, reputation is their main asset. If the adviser is the manager's assistant or colleague, he probably wishes, as the manager does, to be perceived as an able employee who is likely to be promoted.

I now relax the assumption of a sincere adviser and consider what kind of recommendations careerist advisers provide. Specifically, I examine how the strategic behavior of the decision maker affects the information transmitted by advisers, when the latter can withhold information from her.

Assume therefore that the adviser is interested in convincing the evaluator of his accurate private information. Let π^A denote the posterior belief of E on the reliability p^A of A. Thus, once an adviser is consulted, his (sequentially rational) strategy is set to maximize $\pi^A + \theta^A I$, where $\theta^A \geq 0$ describes the outcome concerns of the adviser. Since the adviser does not have to report his information sincerely, he uses a message function that describes which recommendation and reliability level he reports given his true signal and its accuracy. The adviser, not only can misrepresent his signal, but also can misrepresent the accuracy of his information.

In equilibrium, the evaluator's belief updating functions, π and $\pi^A : I \times I^A \rightarrow [.5, 1]$, are derived from the consultation and action strategies of the decision maker and from the message function of the adviser, given the appropriateness of the decision maker's action and whether it coincides with the adviser's recommendation or not. Note that if the decision maker does not consult, the beliefs of E on the type of A must remain the prior beliefs, i.e., that p^A is distributed uniformly on $[.5, 1]$. The reason is that when D chooses whether to consult, she has no information about $A's$ type and therefore cannot condition her actions on it.

The model with a strategic adviser becomes then a two-level message game. The first level, consists of the decision maker sending, by her actions, a message on her ability to the evaluator. The second level, consists of the adviser sending a cheap talk message about his ability, that is observed both by the evaluator (partially) and by the decision maker.[8] In the seminal paper about cheap talk games by Crawford and Sobel (1982), a sender transmits information to a receiver who has no information at all regarding the state of the world. The sender and the receiver may have conflicting preferences regarding the action that the receiver should take, given each truthful message of the sender. In my model, the second level message game between the adviser (sender) and the decision maker (receiver) is a cheap talk game generalized for the case of an informed receiver, since the decision maker in the model has information about the state of the world. Therefore, the sender and the receiver may have conflicting interests regarding the *strategy* the receiver should use (i.e., the action given each (s, p)) and not regarding some specific action.

Why would the adviser opt to misrepresent his information? Evidently, it must be to affect the actions that the decision maker takes given his recommendation. These actions shape the evaluator's belief about the adviser's type. The decision maker's actions may contain information about the reliability of the adviser, and as a result, the evaluator may attribute different beliefs to an adviser whose advice is followed compared to one who is contradicted.[9] The adviser may therefore possess preferences regarding the decision maker's strategy. He may prefer, for example, that a type (s, p) would contradict

[8]The second level message game takes place only if the decision maker consults her adviser.

[9]This is a consequence of the assumption that the adviser's message about his reliability is observed only by the decision maker.

his advice and that a type (s', p') would follow it. In equilibrium, although the adviser does not know (s, p), he can compute the expected action taken by the decision maker given any of his reports. If for some reports this expected action does not accord with his preferred action, the adviser may lie and manipulate the decision maker to take the action that best promotes his interests.

A priori, the model does not entail a conflict between decision makers and advisers. However, such a conflict is generated endogenously as described in the next Lemma.

Lemma 2. *If in equilibrium the adviser transmits truthful information, then $\pi^A(1,1) > \pi^A(0,1)$ and $\pi^A(0,0) > \pi^A(1,0)$, i.e., the adviser receives greater reputation when his advice is followed than when it is contradicted.*

Lemma 2 relies on the decision maker's behavior when recommendations are truthful, as described in Proposition 1. It establishes that when the decision maker contradicts advice and proves that she is able, she also shows as a by-product that she has consulted a relatively unable adviser. Since the decision maker has an excessive incentive to contradict, the adviser, to protect himself, must prevent the decision maker from signaling her ability at his expense. This incentive yields our next result.

Proposition 3. *There is no equilibrium with truthful information revelation.*

In equilibrium, the adviser does not necessarily fool the decision maker to the extent that she follows his advice more often than she should. However, his incentive to lie reduces the amount of information that he is able to credibly transmit.

It is clear, following Crawford and Sobel (1982), that if the adviser were allowed to send a message about his reliability to the evaluator, he could not have sent any credible messages, since poor-ability advisers have an incentive to fool the evaluator. The above result shows that the adviser cannot even send truthful information about his reliability secretly to a third party, namely the decision maker, whose action is observed by the evaluator.

I do not fully characterize the equilibria of the model with a strategic adviser but construct equilibria which maintain the nature of the previous results. For example, for sufficiently high values of θ, there exists an equilibrium analogous to the one described in Proposition 2; the most able decision makers do not consult while the less able decision makers consult. The next proposition describes another equilibrium, in which the decision maker always consults.

Proposition 4. *For all values of* θ *and* θ^A*, there exists an equilibrium which is characterized by a cutoff point* $\tilde{p}$*,* $\tilde{p} \in (.5, E_x(p^A))$*. If* $p \geq \tilde{p}$*, D consults but ignores advice and follows her own signal. If* $p < \tilde{p}$*, D consults and follows advice. The adviser does not send meaningful messages about his reliability but just reports his signal* s^A*.*

In this equilibrium, when the decision maker contradicts advice, it is evident that her type is above $\tilde{p}$ because all types below $\tilde{p}$ simply follow advice. Thus, contradicting an adviser still arises as a signal for high ability even in the presence of strategic advisers. In accordance with the previous results, the adviser is contradicted excessively; the efficient strategy is to contradict if $p \geq E_x(p^A)$, but the decision maker does so for some $p < E_x(p^A)$.

The adviser does not send any meaningful messages about his reliability in this equilibrium. Therefore, the actions of the decision maker do not reveal anything to the evaluator on the adviser's type beyond what the adviser's messages themselves reveal. Since the adviser is believed to report his signal s^A truthfully, he is judged only by the correctness of his prediction and receives higher reputation for the right guess. This induces him to report his signal truthfully. This incentive also coincides with his outcome concerns since when his recommendation is not ignored, it is followed, and the adviser rather lead the decision maker in the right direction.

To summarize this section, the main finding is that consultation in the presence of career concerns elicits low quality of information. Although the career concerns of the adviser are not conflicting with those of the decision maker, in the sense that they are not competing for the same job, his personal career concerns provide him with a sufficient incentive to conceal his information from the decision maker.

4 Extension: choosing an adviser

In many organizations, the decision maker or the manager may be able to choose among advisers of different qualities. The manager may know that one of her assistants specializes in some fields whereas in other fields the research department is more likely to provide a knowledgable recommendation.

Consider therefore the following variation of the model, in which instead of the first stage choice between consulting and not consulting, the decision maker chooses between consulting A and consulting B, which stands for a bad quality adviser. Let us consider the case of sincere advisers. Assume, for simplicity, that the bad adviser's information is of the lowest possible accuracy, so that consulting a bad adviser amounts to not consulting at all in terms of information value. Which type of adviser will the decision maker consult?

If the evaluator can observe the type of each adviser, the model reduces to the one that we analyzed. Consider therefore the case in which the evaluator cannot observe the type of the consulted adviser (the evaluator still partially observes the recommendation transmitted). It seems that there is no incentive to consult the bad adviser since the mere act of consulting him is not observed so it cannot serve as a signal. Moreover he does not provide valuable information. However, in a companion working paper (Levy 2000), I show that in any equilibrium, the bad adviser is consulted. The incentive to consult the bad adviser is an implication of the relative performance evaluation that is created in equilibrium upon consultation. Since the decision maker signals her ability by comparing herself with her adviser and in particular by contradicting his advice successfully, she finds it sometimes easier to be compared with an inferior challenger. By hiring a bad adviser, who is more likely to err, the decision maker can more comfortably use the signal of contradicting advice.

To conclude, the endogenous competition between decision makers and advisers induces decision makers to consult inefficient sources of information although efficient sources are available. We therefore identify another rationale for decision makers to restrict in advance the amount of available information.

5 Discussion and concluding remarks

In this paper, I have shown how consultation may be inefficient when agents have career concerns. Decision-makers tend to restrict in advance the amount of information that they gather by not consulting available advisers. Once they do ask for advice, their advisers conceal information from them and in any case, decision-makers tend to ignore advice too often. Career concerns result therefore in sub-optimal sharing of information. In some environments, the usual tools for correcting inefficient behavior cannot prevail. If the preference parameters (θ, θ^A) are unknown, contracts cannot restore efficiency.

In a dynamic environment, which is suppressed in this paper, it is reasonable to assume that career concerns diminish over time. The reason is that as time goes by, enough observations are accumulated about the decision-maker's type, which therefore becomes known. In light of this, I suspect that the results in this paper apply more often to inexperienced managers that have just finished their MBA studies, than to senior managers.

This reasoning reveals an inter-temporal trade-off that a principal or an evaluator faces. Once he identifies able managers, their career concerns diminish and better management is ensured. However, in the process of identifying these able decision makers, current management is relatively poor and risky decisions are taken. Could the evaluator mitigate this trade-off by, for example, using other methods of evaluation? If the evaluator updates his beliefs on the basis of successful outcomes only, while disregarding actions such as following or contradicting advice, then the decision maker is motivated to behave efficiently both through her outcome and career concerns. Thus, efficient decisions are taken and at the same time some information is revealed about the decision maker's type. However, in most markets and institutions, the evaluator is unable to commit to ignore available information that is relevant for his belief updating process. Ex-post, the evaluator prefers to use information about whether the decision maker had followed advice or not. The decision maker realizes this and distorts her actions accordingly. The model identifies therefore another case in which *too much information hurts the principal* (for other cases see Meyer and Vickers (1997), Cremer (1995)).

Finally, if commitment to ignore individual messages *is* possible, then the model points at tools that may better solve the inter-temporal trade-off. Specifically, the model suggests the advantage of treating the decision maker and her adviser as one unit. In other words, if agents would be promoted as a team, and the team's reputation would matter instead of the individual one, information would be shared optimally among the team members. The drawback of this procedure is that the evaluator or the principal would not be able to learn about individuals and would forgo the opportunity to break teams and create better ones. This trade-off, between good decisions in the future (which is enhanced if incompetent decision makers are replaced) and good decisions in the present (which is achieved if replacement is impossible and hence information is shared optimally), could determine the optimal size of a decision-making team or a committee.

References

Aghion, P. and J. Tirole (1996), "Formal and Real Authority in Organizations", *Journal of Political Economy*, vol. 105, 1-29.

Allen, T.J. (1977), *Managing the Flow of Technology: Technology Transfer and the*

Dissemination of Technological Information Within the R&D Organization, Cambridge, MA: MIT Press.

Austen-Smith, D. (1993), "Interested Experts and Policy Advice: Multiple Referrals under Open Rule", *Games and Economic Behavior*, vol. 5, 3-43.

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

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

Blau, P.M. (1954), "Patterns of Interaction Among a Group of Officials In a Government Agency", *Human Relations*, vol. 7, 337-348.

Calvert, R. (1985), "The Value of Biased Information: A Rational Choice Model of Political Advice", *Journal of Politics*, vol. 47, 530-555.

Chevalier, J. and G. Ellison (1999), "Career Concerns of Mutual Fund Managers", *Quarterly Journal of Economics*, vol. 114, 389-432.

Crawford, V. and J. Sobel (1982), "Strategic Information Transmission", *Econometrica,* vol. 50, 1431-1451.

Cremer, J. (1995), "Arm's Length Relationships", *Quarterly Journal of Economics*, vol. 110, 275-295.

Effinger, M.R. and M.K. Polborn (1998), "I Say Hello and You Say Good-bye", mimeo, University of Munich.

Farrell, J. and R. Gibbons (1989), "Cheap Talk with Two Audiences", *American Economic Review,* vol. 79, 1214-1223.

Feldman, M.S. and J.G. March (1981), "Information in Organizations as Signal and Symbol", *Administrative Science Quarterly*, vol. 26, 171-186.

Friebel, G. and M. Raith (1997), "Strategic Recruiting and the Chain of Command", mimeo, Universite Libre de Bruxelles.

Gibbons, R. and K.J. Murphy (1991), "Optimal Incentive Contracts in the Presence of Career Concerns", *Journal of Political Economy,* vol. 100, 468-505.

Holmström, B. (1982), "Managerial Incentives Problems: A Dynamic Perspective", in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School Econ.

Holmström, B. and J. Ricart i Costa (1986), "Managerial incentives and Capital Management", *Quarterly Journal of Economics,* vol. 101, 835-860.

Krishna, V. and J. Morgan (1998), "A Model of Expertise", mimeo, Pennsylvania State University and Princeton University.

Levy, G. (2000), "Strategic Consultation and Strategic Recruiting in the Presence of Career Concerns", http://econ.lse.ac.uk/staff/glevy/papers/glevy0900.pdf.

Meyer, M.A. and J. Vickers (1997), "Performance Comparisons and Dynamic Incentives", *Journal of Political Economy,* vol. 105, 547-581.

Milgrom, P. (1982), "Good news and Bad news: Representation Theorems and Applications", *Bell Journal of Economics*, vol. 17, 380-391.

Morris, S. (1998), "An Instrumental Theory of Political Correctness", mimeo, University of Pennsylvania and Northwestern University.

Okuno-Fujiwara, M., A. Postelwaite and K. Suzumura (1990), "Strategic Information Revelation", *Review of Economics Studies,* vol. 57, 25-47.

Ottaviani, M. and P. Sørensen (1999), "Professional Advice", mimeo, University College, London and University of Copenhagen.

Ottaviani, M. and P. Sørensen (2000), "Herd Behavior and Investment: A Comment", *American Economic Review,* vol. 90, 695-704.

Prendergast, C. (1993), "A Theory of 'Yes Men' ", *American Economic Review,* vol. 83, 757-770.

Prendergast, C. and L. Stole (1996), "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning", *Journal of Political Economy,* vol. 104, 1105-1134.

Scharfstein, D. and J. Stein (1990), "Herd Behavior and Investment", *American Economic Review,* vol. 80, 465-479.

Staw, B. and L. Ross (1987), "Behavior in Escalation Situations: Antecedents, Prototypes, and Solutions", in *Research in Organizations Behavior*, L.L.Cummings and B. Staw. Ed, vol. 9, 39-78.

Trueman, B. (1994), "Analyst Forecasts and Herding Behavior", *Review of Financial Studies,* vol. 7, 97-124.

Zweibel, J. (1996), "Corporate Conservatism and Relative Compensation", *Journal of Political Economy,* vol. 103, 1-25.

Appendix

Some preliminaries:

Let $g(p)$ and $\tilde{g}(p)$ denote continuous and twice differentiable density functions on $p \in X \subseteq \Re^1$. A sufficient condition for $\int_X pg(p)dp > \int_X p\tilde{g}(p)dp$ is the monotone likelihood ratio property (MLRP), i.e.,

$$\frac{g(p)}{g(p')} \geq \frac{\tilde{g}(p)}{\tilde{g}(p')}$$

for all $p \geq p'$, with a strict inequality for at least some $p > p'$ (see Milgrom (1982) for the proof that MLRP is a sufficient condition).

When D consults an adviser for $p \in X$, the updated beliefs of E are of the form $\pi(I, I^A) = \int_X pg(p|I, I^A)dp$, where $g(p|I, I^A)$ is the posterior density function of p, given the observed history (I, I^A). Specifically, by Bayesian updating,

$$g(p|I, I^A) = \frac{\Pr(I, I^A|p)f(p)}{\int_X \Pr(I, I^A|p)f(p)dp}$$

where $\Pr(I, I^A|p)$ is the probability that some history (I, I^A) is observed given p and $f(p)$ is the prior distribution over $p \in X$.

Let D consult A for some values of $p \in X$ and behave according to equations (1) and (2) in the text, i.e., when $s^A \neq s$, D contradicts advice if $p^A \leq \frac{kp}{kp+(1-p)}$ and if $s^A = s$, she contradicts advice whenever $p^A \leq \frac{k(1-p)}{p+k(1-p)}$. Then, $D's$ strategy depends only on the parameter k. Assume that E knows $D's$ strategy and updates his beliefs accordingly. $D's$ strategy is continuous in k, implying that $E's$ updated beliefs are continuous in k. Using the MLRP, I can show that:

(i) For all $k > 0$, $\pi(1,1) > \pi(0,0)$ and $\pi(1,0) > \pi(0,1)$.

(ii) For $0 < k \leq 1$, $\pi(1,0) > \pi(1,1)$ and $\pi(0,1) > \pi(0,0)$.

(iii) For $k \to \infty$, $\pi(1,1) > \pi(0,1)$.

To exemplify the nature of the proofs (which are not complicated but still tedious and therefore omitted), consider (ii). To see that $\pi(1,0) > \pi(1,1)$, by MLRP, we need to show that for $p > p'$,

$$\frac{g(p|1,0)}{g(p'|1,0)} \geq \frac{g(p|1,1)}{g(p'|1,1)} \to \frac{\Pr(1,0|p)f(p)}{\Pr(1,0|p')f(p')} \geq \frac{\Pr(1,1|p)f(p)}{\Pr(1,1|p')f(p')}$$

and therefore it suffices to show that $\frac{\Pr(1,0|p)}{\Pr(1,1|p)}$ increases with p.[10] These updated probabilities are simple expressions, such as $\Pr(1,0|p) = p\int_{\frac{kp}{kp+(1-p)}}^{1}(1-p^A)f(p^A)dp$ for $p > \frac{1}{1+k}$ and 0 otherwise when $0 < k \leq 1$.

Proof of Lemma 1:

Assume that equilibrium beliefs satisfy $\pi(1) > \pi(0)$. Then the decision maker must follow her signal in equilibrium because $p\pi(1) + (1-p)\pi(0) + p\theta > p\pi(0) + (1-p)\pi(1) + (1-p)\theta$. Then,

$$\pi(1) = \int_U p\frac{\Pr(1|p)f(p)}{\int_U \Pr(1|p)f(p)dp}dp, \text{ and } \pi(0) = \int_U p\frac{\Pr(0|p)f(p)}{\int_U \Pr(0|p)f(p)dp}dp$$

for $\Pr(1|p) = p$, $\Pr(0|p) = 1-p$ and $f(p)$ is uniform over U. It is easy to see, using the

[10] $\Pr(\cdot,\cdot|p)$ depends on k. Whenever no confusion is created I omit the index k.

MLRP, that indeed $\pi(1) > \pi(0)$ and therefore the equilibrium described in the lemma exists. Note that the expected utility of D, $p(\pi(1) - \pi(0) + \theta) + \pi(0)$ is linearly increasing in p.

Assume now that $\pi(1) < \pi(0)$. When θ is sufficiently high, the decision maker would still follow her signal for all p. Therefore, the beliefs are as stated above, inducing $\pi(1) > \pi(0)$, a contradiction. When θ is sufficiently low, a "mirror" equilibrium can exist, in which all types of D contradict their signal. These types of equilibria exist as long as the meaning of messages is not fixed and I disregard them. Finally, assume that $\pi(1) = \pi(0)$. For such beliefs to be rational in an informative equilibrium, it must be that the decision maker's strategy is not monotone in p, which is impossible for $\theta > 0$.■

Proof of Proposition 1:

An equilibrium is determined by a fixed point to the equation $k = h_\theta(k) \equiv \frac{\pi_k(1,0) - \pi_k(0,0) + \theta}{\pi_k(1,1) - \pi_k(0,1) + \theta}$, and equations (1) and (2) that describe $D's$ behavior in equilibrium given k. To prove the proposition, we have to show that in an informative equilibrium a fixed point exists and that its value k satisfies $k > 1$.

If $k \leq 0$, then all types of D must either contradict the adviser or follow the adviser which cannot constitute an informative equilibrium. Consider therefore the case of $k > 0$ where both the denominator and the nominator of $h_\theta(k)$ are positive (otherwise, it is a "mirror" equilibrium). By the preliminaries, $h_\theta(k)$ is continuous, $h_\theta(k) > k$ for all $k \leq 1$ and for $k \to \infty$, $h_\theta(k) < k$. Therefore, for all values of $\theta \geq 0$, a fixed point $k = h_\theta(k)$ exists and must admit $k > 1$.■

Lemma 3. *For sufficiently high values of θ, a decision maker of type p takes the right decision with a greater probability when she consults an adviser than when she relies only on her private information.*

Proof of Lemma 3:

Given some level of $k > 1$, a decision maker with a type $p \geq \frac{k}{1+k}$ contradicts A only when $s^A \neq s$ and $p^A \leq \frac{kp}{kp+(1-p)}$, and follows advice otherwise. Although the contradiction level is excessive, the probability of taking the right action is higher compared to acting unilaterally (which amounts to contradicting advice for all values of p^A when $s^A \neq s$, and following advice whenever $s^A = s$). We need therefore to consider all types $p < \frac{k}{1+k}$, who contradict advice also when $s^A = s$. Thus, the gain (relative to acting unilaterally) from following advice sometimes when $s^A \neq s$ may be outweighed by the loss from contradicting advice also when $s^A = s$. The greater is the value of k in equilibrium, the greater is this loss as equations (1) and (2) imply. According to these equations, when $k \to 1$ the decision maker's behavior upon consultation converges to the efficient behavior, thereby dominating acting unilaterally. When $k \to \infty$, the decision maker almost always contradicts advice, rendering consulting less efficient than acting unilaterally. Thus, there exists a cut-off value of k, $\tilde{k}$, such that if in equilibrium $k \leq \tilde{k}$, consultation is the efficient action relative to acting unilaterally. Finally, note that the value of k in equilibrium is bounded. That is, $k = h_\theta(k) < \frac{\theta+.5}{\theta-.5}$. Moreover, the equilibrium value of k decreases with θ, because whenever $k > 1$, $\frac{\partial h_\theta(k)}{\partial \theta} < 0$. Therefore, for all $\theta \geq \tilde{\theta}$, consulting an adviser is the efficient action where $\tilde{\theta}$ is defined by $\tilde{k} = \frac{\tilde{\theta}+.5}{\tilde{\theta}-.5}$.■

Proof of Proposition 2:

If the decision maker does not consult an adviser, she follows her signal, as in the proof of Lemma 1. Her expected utility $EU^N(p)$ is therefore a linearly increasing function in p. When she consults an adviser, it is easy to show that her expected utility, $EU_k^A(p)$, is an increasing convex function. Therefore, if there exists an equilibrium in which some players consult and some do not, there can be at most two types who are indifferent between consulting and not consulting, $\hat{p} \in [.5, 1)$ and $\check{p} \in (\hat{p}, 1]$ where D does not consult for $p \in [\hat{p}, \check{p}]$. Namely, $U = [\hat{p}, \check{p}]$ and $C = [.5, 1) \cup (\hat{p}, 1]$. The following conditions must hold, where at least one of the inequalities has to be satisfied with an equality:[11]

$$(i)\ EU_k^A(\check{p}) \geq EU^N(\check{p}), \text{with equality if } \check{p} < 1,$$

$$(ii)\ EU_k^A(\hat{p}) \leq EU^N(\hat{p}), \text{with equality if } \hat{p} > .5,$$

$$(iii)\ k = \frac{\pi(1,0) - \pi(0,0) + \theta}{\pi(1,1) - \pi(0,1) + \theta}$$

I will now prove three claims that establish the result.

Claim 1: For all θ, the least able types must consult in equilibrium.

Suppose to the contrary that there exists an equilibrium with $\hat{p} = .5$. To sustain an equilibrium in which some types do consult, it must be that $\check{p} < 1$ so that condition (i) must be satisfied with equality. However, the evaluator observes consultation. Thus, a decision maker that does not consult is believed to be in the open interval $(.5, \check{p})$ and a decision maker that consults is believed to be in the open interval $(\check{p}, 1)$. Since a decision maker that consults can always ignore advice and follow her own signal, it is a profitable deviation for any p to consult. Thus, $\hat{p} > .5$.□

Claim 2: For sufficiently high values of θ, the most able types do not consult.

We have to prove that $\check{p} = 1$. Assume that $\check{p} < 1$ but that $\check{p} \rightarrow 1$. Note that if θ is high enough, then $\hat{p} >> .5$ because types in the neighborhood of .5 rather consult and receive θ with a greater probability even if they are revealed in equilibrium. But, if $\check{p} \rightarrow 1$ and $\hat{p} >> .5$, then $\pi(\cdot) > \pi(\cdot, \cdot)$, i.e., the beliefs of the evaluator are higher if the decision maker does not consult than if she consults, disregarding the observed history. Thus, a preferred action by a type $\check{p} = 1$, who follows her signal and receives θ anyhow, is to act unilaterally, a contradiction. Assume instead that $\check{p} << 1$ and therefore condition (i) holds with equality. Since D prefers to consult for all $p \geq \check{p}$, then $\frac{\partial EU_k^A(p)}{\partial p}|_{p=\check{p}} > \frac{\partial EU^N(p)}{\partial p}|_{p=\check{p}}$ and $\frac{\partial EU_k^A(p)}{\partial p}|_{p=1} > \frac{\partial EU^N(p)}{\partial p}|_{p=1}$ by the convexity of $EU_k^A(p)$ and the linearity of $EU^N(p)$. I find that for all p, when θ is sufficiently high, condition (i)

[11]For brevity, I omit the equilibrium indices $k, \hat{p}$, and $\check{p}$, from the equilibrium beliefs and utility whenever no confusion is created.

and $\frac{\partial EU_k^A(p)}{\partial p}|_{p=1} > \frac{\partial EU^N(p)}{\partial p}|_{p=1}$ cannot hold simultaneously unless $\check{p} \to 1$, a contradiction (details are available upon request).□

Claim 3: For sufficiently high values of θ, there exists an equilibrium in which the decision maker does not consult when she is sufficiently able.

Let $p(\hat{p})$ satisfy $EU_k^A(p(\hat{p})) = EU^N(p(\hat{p}))$. The type $p(\hat{p})$ is the type that is indifferent between consulting or not if the evaluator believes that all types $p > \hat{p}$ do not consult and all types below $\hat{p}$ consult. By the continuity of the expected utility functions, $p(\hat{p})$ is a continuous function as well. Assume that $\hat{p} \to .5$, i.e., that almost all types of decision makers do not consult. When θ is high enough, a decision maker with type $p \to .5$ prefers to consult in order to increase the probability of receiving θ, implying that $p(\hat{p})_{\hat{p}\to.5} > \hat{p}$. If $\hat{p} \to 1$, the reputation from acting unilaterally is greater than that from consulting disregarding the history observed because only types above $\hat{p} \to 1$ do not consult. A decision maker with type $p \to 1$ prefers then to act unilaterally because she has the same probability of receiving θ in both cases. This implies $p(\hat{p})_{\hat{p}\to 1} < \hat{p}$. By the mean value theorem, there must exist a $\hat{p} \in (.5, 1)$ such that $p(\hat{p}) = \hat{p}$. □

By the preliminary steps, as in the proof of Proposition 1, for any $\hat{p}$ there exists $k > 1$ that satisfies $k = \frac{\pi_k(1,0)-\pi_k(0,0)+\theta}{\pi_k(1,1)-\pi_k(0,1)+\theta}$. This completes the proof.■

Proof of Lemma 2:

From the analysis of the game with a sincere adviser, we know that whenever the adviser is consulted in equilibrium, it is either the case that all types consult (i.e., $C = [.5, 1]$), or it is the case that some types do not consult. From the proof of Proposition 2, we know that whenever some types do not consult, it is either the most able types (i.e., $C = [\hat{p}, 1]$), or a combination of the most able and the least able ($C = [.5, \check{p}) \cup (\hat{p}, 1]$). These are the only sets of types of D that possibly consult A in equilibrium for any θ.

Whenever a type p of D consults A, she behaves according to Proposition 1, and contradicts advice excessively. We have to show that for all these specified C sets,

$\pi^A(1,0) < \pi^A(0,0)$ and $\pi^A(0,1) < \pi^A(1,1)$.

Consider $C = [.5, 1]$ and the Bayesian updating on p^A performed by E when the decision maker takes the right choice by contradicting her adviser who therefore makes the wrong recommendation:

$$g(p^A|1,0) = \begin{cases} \frac{2(1-p^A)\left(\int_{.5}^{\frac{k(1-p^A)}{k(1-p^A)+p^A}} 2(1-p)dp + \int_{.5}^{1} 2pdp\right)}{\Pr(1,0)} & \text{for } p^A \in [.5, \frac{k}{1+k}] \\ \frac{2(1-p^A)\int_{\frac{p^A}{k(1-p^A)+p^A}}^{1} 2(1-p)dp}{\Pr(1,0)} & \text{for } p^A \in [\frac{k}{1+k}, 1] \end{cases}$$

and when the decision maker is wrong as well, the updated density function over the types of the adviser is:

$$g(p^A|0,0) = \begin{cases} \frac{2(1-p^A)\int_{\frac{k(1-p^A)}{k(1-p^A)+p^A}}^{1} 2(1-p)dp}{\Pr(0,0)} & \text{for } p^A \in [.5, \frac{k}{1+k}] \\ \frac{2(1-p^A)\left(\int_{.5}^{\frac{p^A}{k(1-p^A)+p^A}} 2pdp + \int_{.5}^{1} 2(1-p)dp\right)}{\Pr(0,0)} & \text{for } p^A \in [\frac{k}{1+k}, 1] \end{cases}$$

Then, for all $p^A > \tilde{p}^A$,

$$\frac{g(p^A|0,0)}{g(\tilde{p}^A|0,0)} > \frac{g(p^A|1,0)}{g(\tilde{p}^A|1,0)} \Leftrightarrow \frac{\left(\frac{p^A}{k(1-p^A)+p^A}\right)^2}{\left(\frac{\tilde{p}^A}{k(1-\tilde{p}^A)+\tilde{p}^A}\right)^2} > \frac{1-\left(\frac{p^A}{k(1-p^A)+p^A}\right)^2}{1-\left(\frac{\tilde{p}^A}{k(1-\tilde{p}^A)+\tilde{p}^A}\right)^2}$$

which holds for all $p^A > \tilde{p}^A$. Thus, by MLRP, $\pi^A(1,0) < \pi^A(0,0)$. Analogous analysis holds for the case of $\pi^A(0,1) < \pi^A(1,1)$ and for the sets $C = [\hat{p}, 1]$ and $C = [.5, \check{p}) \cup (\hat{p}, 1]$.■

Proof of Proposition 3:

Assume that the adviser reports his information truthfully and consider a deviation from a truthful report (s^A, p^A) to a non-truthful report $(s^A, p^A + \varepsilon)$ for a small $\varepsilon > 0$. Since the decision maker contradicts advice excessively, when the reported reliability slightly increases, it increases the efficiency of decision making because the decision maker would follow advice more often when $p < p^A$. Moreover, since the decision maker is more likely to follow advice, the adviser is more likely to receive $\pi^A(1,1)$ instead of $\pi^A(0,1)$ and $\pi^A(0,0)$ instead of $\pi^A(1,0)$ and therefore to increase his expected utility. Thus, both his outcome concerns and his career concerns drive him to bias his report upwards, implying that a fully revealing equilibrium cannot exist.■

Proof of Proposition 4:

Suppose that the adviser behaves as the Proposition supposes, and that the evaluator believes that for all $p > \tilde{p}$, the decision maker ignores advice and otherwise she follows it. When $s^A = s$, expected utility is equal for both actions. When $s^A \neq s$, if the decision maker ignores advice her expected utility $E_x U_{\tilde{p}}(ignore|p)$ is:

$$\frac{p(1-E_x(p^A))}{p(1-E_x(p^A))+(1-p)E_x(p^A)}(\pi_{\tilde{p}}(1,0)+\theta)+\frac{(1-p)E_x(p^A)}{p(1-E_x(p^A))+(1-p)E_x(p^A)}\pi_{\tilde{p}}(0,1)$$

When she follows advice, her expected utility $E_x U_{\tilde{p}}(follow|p)$ is:

$$\frac{p(1-E_x(p^A))}{p(1-E_x(p^A))+(1-p)E_x(p^A)}\pi_{\tilde{p}}(0,0)+\frac{(1-p)E_x(p^A)}{p(1-E_x(p^A))+(1-p)E_x(p^A)}(\pi_{\tilde{p}}(1,1)+\theta)$$

If an equilibrium as described in the Proposition exists, then there exists $\tilde{p}$ for which $E_x U_{\tilde{p}}(\text{ignore}|\tilde{p})=E_x U_{\tilde{p}}(\text{follow}|\tilde{p})$, implying:

$$\frac{(1-\tilde{p})E_x(p^A)}{\tilde{p}(1-E_x(p^A))}=\frac{\pi_{\tilde{p}}(1,0)-\pi_{\tilde{p}}(0,0)+\theta}{\pi_{\tilde{p}}(1,1)-\pi_{\tilde{p}}(0,1)+\theta}$$

By MLRP, the RHS is greater than 1. The reason is that once a decision maker contradicts advice she is known to be above $\tilde{p}$ whereas whenever she follows advice, she may be above $\tilde{p}$ but she is also mixed with all types below $\tilde{p}$. Thus, if a fixed point exists, then it has to be that $\frac{(1-\tilde{p})E_x(p^A)}{\tilde{p}(1-E_x(p^A))} > 1 \rightarrow \tilde{p} < E_x(p^A)$.

To see that a fixed point exists, let $p(\tilde{p})$ satisfy $E_x U_{\tilde{p}}(\text{ignore}|p(\tilde{p}))=E_x U_{\tilde{p}}(\text{follow}|p(\tilde{p}))$. The type $p(\tilde{p})$ is the type that is indifferent between ignoring advice and following it if the evaluator believes that all types $p > \tilde{p}$ ignore advice and all types below $\tilde{p}$ follow advice. $p(\tilde{p})$ is a continuous function. Assume that $\tilde{p} \rightarrow .5$, i.e., that almost all types of decision makers ignore advice. Thus, the adviser's recommendation does not reveal any information about $D's$ types. Then, as in Lemma 1, the decision maker receives a greater reputation for taking the right decision, a reputation which is invariant to the adviser's message. This induces the less able types to follow advice, because they increase the probability of making the right decision, thereby increasing both their reputation and the probability of receiving θ, implying $p(\tilde{p})_{\tilde{p}\to .5} > \tilde{p}$. If $\tilde{p} \to 1$, then clearly $\pi_{\tilde{p}}(0,1) > \pi_{\tilde{p}}(1,1)$ because only the most able types contradict advice. Hence, a type $p \to 1$ that takes the right action even with no advice, would rather ignore advice to capitalise on the reputation gains, inducing $p(\tilde{p})_{\tilde{p}\to 1} < \tilde{p}$. By the mean value theorem, for all θ (including $\theta = 0$), there must exist a $\tilde{p} \in (.5, 1)$ such that $p(\tilde{p}) = \tilde{p}$. Finally, the incentives of the adviser to accord with the suggested equilibrium are explained in the text.■