Leadership, Coordination and Mission-Driven Management

Patrick Bolton*
Columbia University

Markus K. Brunnermeier†
Princeton University

Laura Veldkamp‡
New York University

November 26, 2007

- Preliminary and Incomplete Draft -

Abstract

What makes a good leader? A good leader coordinates his followers by making a precise mission statement, which credibly communicates the optimal course of action. In practice, leaders learn about the optimal action over time. Learning creates a time-consistency problem because the leader has an incentive to commit to a mission to achieve coordination and then adjust it when new information arrives. Overconfidence is a valuable attribute in such a setting, since it helps the leader stick with his prior belief. Even with a costly commitment technology available, overconfident leaders still facilitate better coordination and teamwork. The drawback of overconfidence is that it inhibits learning from the followers' actions. Hence, it is costly when followers have sufficiently valuable signals.

Keywords: Leadership, Coordination, Overconfidence
JEL Classification: D23, D7, D81

*Department of Economics and GSB, Columbia University, New York, NY 10027, e-mail: pb2208@columbia.edu, http://www0.gsb.columbia.edu/faculty/pbolton

†Department of Economics, Bendheim Center for Finance, Princeton University, Princeton, NJ 08544-1021, e-mail: markus@princeton.edu, http://www.princeton.edu/∼markus

‡Department of Economics, Stern School of Business, New York University, New York, NY 10012-1126, e-mail: lveldkam@stern.nyu.edu, http://pages.stern.nyu.edu/∼lveldkam

> "Consistency of word and deed on the leader's part is absolutely necessary if others are to commit themselves to the personal and business risks associated with new and unproven courses of action. The general manager who runs hot and cold will fail to encourage confidence in others. ... Nobody wants to go out on a limb and risk being abandoned at the first sound of cracking wood." Aguilar (1988)

1 Introduction

In this paper we consider a model of leadership in organizations. The role of leadership we focus on is that of helping coordinate the actions of the different members of an organization. The role of the leader is to give a sense of direction for the organization. The leader evaluates the environment in which the organization operates and determines the best strategy adapted to that environment. The leader's dilemma is that he would like to base the organization's focus (or mission) on all the relevant information about the environment available to him. But, since information about the environment only trickles in over time, the leader may then be led to revise the organization's direction as new information becomes available. His desire to modify the direction of the organization over time undermines his ability to coordinate the actions of the other members of the organization.

In other words, the essence of the leadership problem in our model is to reconcile the adaptation to a changing environment–which requires information acquisition and revision of the organization's strategy in response to new information–and coordination of the actions of the other members of the organization. Thus, the main question we are interested in here, is determining which attributes of a leader are most desirable in balancing the need for adaptation with coordination.

Our leadership problem can be captured in a simple setup involving five stages. In the first stage the leader observes a first signal of the environment (or state of nature) the organization is likely to be in. Based on that signal the leader can define a mission or overall strategy for the organization. In a second stage, the other members of the organization – the followers – decide how closely they want to stick to the leader's strategy. They may not be inclined to blindly follow the leader's proposed strategy because they also observe signals about the state of nature, and they may come up with different forecasts of what the ultimate direction for the organization will be. In a third stage the leader receives a second signal. This signal could be an aggregate of the signals of the followers or simply new information that becomes available. In a fourth stage the leader implements the organization's strategy given all the information he has available. Since the followers have already acted, there is no more incentive for the leader to coordinate them. The leader's only remaining goal is to adapt his strategy to all the information he has. But the members of the organization anticipate that the leader will revise the organization's strategy. This is what causes them to be insufficiently coordinated in the third stage. Thus, the leader is confronted with a time-consistency problem. In the fifth and last stage, once the strategy has been implemented, the organization's payoffs are realized. These will be higher the better adapted the strategy is to the environment and the better coordinated all the members' actions are.

In this setup, the followers in the organization may be concerned about two potential flaws in leadership: one is that the leader misdiagnoses the circumstances the organization finds itself in and chooses a mis-adapted mission for the group; the other is that the mission is poorly communicated and incoherently implemented with substantial coordination failures.

What makes a good leader in such a situation? We argue that a key attribute of a good leader is a form of overconfidence. An overconfident leader attaches too much weight to the signal he receives in the first stage. He is overconfident to the extent that he attaches an exaggerated information value on the first signal, or on the signals he processes himself. In other words, an overconfident leader trusts his own judgement more than that of other members of the organization. He then tends to define a strategy for the organization based disproportionately on his own best assessment of the environment the firm finds itself in. Given that he puts too little weight on other members' views he is also more likely to follow through the strategy he has intitially outlined for the organization. That is, when he receives new information at stage three he will not revise his assessment of the environment much and therefore will stick more or less to the strategy he defined in the first stage. The overconfident leader's "ability" to follow through a plan proposed in stage 1 clearly helps bring the followers around the plan and coordinate their actions. We show that this coordination benefit outweighs the potential misadaptation cost as long as the leader's overconfidence is not too extreme. There may be circumstances, however, where it may be preferable to have a leader who is a good 'listener' and is capable of formulating well-adapted missions (even if they get implemented somewhat incoherently).

While overconfidence helps a leader to commit to "staying the course" it also raises the risk for the the organization of pursuing the wrong strategy. One might wonder, therefore, whether there aren't better ways of achieving commitment, without at the same time putting too much weight on the signal obtained in the first stage. If a rational leader were able to commit to a strategy for the organization in the first stage by, say, staking his reputation on pursuing a clearly defined mission, in a manner similar to President George H.W. Bush's announcement, "read my lips: no new taxes", wouldn't that be a superior form of commitment? It turns out that even when such a commitment technology is available an overconfident leader could still outperform a rational one, since an overconfident leader makes a stronger mission statement and hence leads to more coordination among the followers and this more than compensates for the loss of flexibility.

2 Related Literature

There is a small but rapidly growing economics literature on leadership. Most of this literature, however, deals with different facets of leadership. One of the earliest contributions is by Rotemberg and Saloner (1993), who address the question of how a leader can motivate followers to exert effort and come up with proposals for improvements in the firm's operation. Followers value the fact that their proposals are taken into account and are adopted by the leader. They are therefore willing to exert (unobservable) costly effort to come up with proposals if they expect that there is a reasonable chance that they will be adopted. Rotemberg and Saloner consider two leadership styles. One is where the leader maximizes profits, and the other where managerial decision-making is more sensitive to the preferences of employees. They show that the latter approach can ultimately lead to higher profits, as it induces employees to exert more effort and thus brings about more improvements. In a subsequent related article, Rotemberg and Saloner (2000) also allow the leader to encourage employee effort by ruling out possible future courses of action, so that employees are better able to determine what kinds of initiatives will be favored. Ruling out certain activities amounts to defining the organization's focus. In this respect the leader's objective of delineating the scope of the organization in Rotemberg and Saloner (2000) is similar to the leader's objective of proposing a consistent plan in our setup. But, instead of considering overconfidence as a desireable quality, Rotemberg and Saloner emphasize the benefits of CEO bias or "vision".

Hermalin (1998) considers the role of leadership by example in a moral hazard in teams problem where organizational output depends on all members' efforts and where all members share the aggregate output. As is well known, in a team production problem individual team members may free-ride on other team members' efforts. Thus, the leader's problem is to motivate team members and help overcome free-riding. Hermalin assumes that the leader has private information about the return to effort and argues that the leader will then tend to overstate the return to effort so as to mitigate free-riding. He will be able to motivate other team members to put in effort by leading by example and exerting himself. Hermalin does not allow for leader overconfidence, but his notion of leading by example is related to our conception of leadership as giving a sense of direction to other organization members.

A handful of papers explore the role of overconfidence in leadership. In Van Den Steen (2005), managerial overconfidence helps attract and retain employees with similar beliefs. The resulting alignment of beliefs helps firms function more efficiently. In particular when similar followers and managers are paired, the manager is more likely to implement projects or ideas proposed by an employee (which provides private benefits to the employee). As in Rotemberg and Saloner, employees are then induced to put in more effort to identify new projects, which benefits the organization.

Goel and Thakor (2000) is a complementary explanation for overconfidence. It studies the process through which leaders get appointed and provides an explanation for why overconfident managers tend to be selected as leaders. They consider a model where managers with unknown ability compete for leadership. In their model managers make the best available project choices and the manager with the best project outcome is selected as leader. They show that overconfident managers tend to make riskier project choices and are therefore more likely to be selected as leader. Similarly, Gervais and Goldstein (2007) introduces overconfidence into a moral hazard in teams problem akin to Hermalin (1998). In their model an overconfident leader tends to work harder and thus induces all other team members to coordinate around a higher effort choice. Unlike in our model, however, they do not consider the time-consistency problem of the leader and how overconfidence can mitigate this problem.

The model of organizations that is most closely related to ours is that of Dessein and Santos (2006). As in our setup they also consider an organization's tradeoff between achieving greater coordination and greater adaptation. However, they do not allow for any role for leadership. In their model members of the organization coordinate through direct communication.

The remainder of our paper is organized as follows. Section 2 presents our model of coordination and adaptation for the organization and the role of leadership in an organization facing this tradeoff. Section 3 considers a slightly more general variant of our model, where the leader can obtain information from other members of the organization revealed by their actions. Section 4 concludes with a summary and directions for future research. Finally, an appendix contains the more involved proofs.

3 Coordination vs. adaptation

The tension between coordination and flexibility arises first from changes in the environment, which require adaptation, and second from the gradual arrival of information about the environment. To illustrate this problem we consider a setting where the leader receives an exogenous signal in each of two periods. Following the arrival of the first signal the leader must propose a strategy for the organization and get other members to coordinate their actions around it. But the leader may change his mind and reorient the strategy following the arrival of the second signal. While the ex-post reorientation helps bring about better adaptation, the anticipation of possible changes in strategy also make it harder to coordinate followers' actions. The reason is that the followers also observe a private signal about the environment and use this signal to forecast possible reorientations of the organization's strategy.

We show that leader overconfidence is a valuable attribute in such a situation (Section 3.1). The more overconfident the leader the less likely he is to change his mind and therefore the less likely is a possible reorientation of the organization's strategy. Remarkably, overconfidence remains a valuable attribute even when the leader can commit to a strategy by staking his reputation (Section 3.2). We assume for now that signals are exogenous. We explore endogenous signals, derived from the aggregate choice of followers, in Section 4.

3.1 Merits of overconfidence

Model setup. The organization we consider has one leader and a continuum of followers indexed by i. The leader differs from the followers in two ways: first he defines a mission statement for the organization based on his initial information S_1 before the followers make their moves. Second, after the followers choose their actions a_i and after receiving an additional signal S_2, the leader implements the strategy of the organization a_L.

Followers value three things:

1. taking an action that is close to (or consistent with) the organization's strategy;
2. belonging to a well-coordinated organization, and
3. belonging to an organization that is well-adapted to its environment θ.

Formally, represent these preferences with the following objective function for each follower:

$$\Pi_i = -(a_i - a_L)^2 - \int_j (a_j - \bar{a})^2 dj - (a_L - \theta)^2 \text{ for } i \in [0,1] \cup \{L\} \qquad (1)$$

One interpretation of this payoff function is that the followers gets a private payment for taking an action close to the ultimate policy choice of their organization/firm. In addition, all followers get a share of firm profits, which depends on the accuracy of the firm's stated goal, and on the degree of coordination among followers.

The leader's objective Π_L is the same as the followers and in this respect our model of the organization is essentially a team problem à la Marchak and Radner (1972). However, our model is different in two respects from a standard team problem. First, as the leader is inevitably well coordinated with himself, we always have $(a_i - a_L) = 0$, for $i = L$. Second, to the extent that a well coordinated action a_i by follower i benefits both him and all other members of the organization, there is a coordination externality among all members. And to the extent that the private and public values of coordination are misaligned there is an additional role for leader overconfidence in our model, namely to help internalize this coordination externality.

Neither the leader nor followers know the true environment of the organization, θ. Both begin with improper (flat) priors and update these priors based on the private signal each receives. The leader receives two signals about θ: S_1 and S_2. Because the first signal S_1 is observed before followers act, the leader can publicly announce it. This announcement, which we call the leader's mission statement, is credible because the leader has no incentive to manipulate the level of followers' expectations. Followers are all rational and know the true variance of all signals. But the leader may be overconfident in the sense that he may underestimate the variance of the first signal (or overestimate the precision of the first signal). The signal has a true distribution $S_1 \sim \mathcal{N}(\theta, 1)$. But the leader believes the signal to have a lower variance $\sigma_p^2 \leq 1$.

After followers choose their action a_i but before the leader chooses his action a_L, the second signal $S_2 \sim \mathcal{N}(\theta, \sigma_2^2)$ is observed. We assume for simplicity that the true and perceived precision of this signal are the same. The rationale for modeling overconfidence as higher precision of the first signal, rather than the second, is most clear in Section 4, when the second signal is generated by other agents' actions. In essence, overconfidence in our model means that a leader trusts his own information more than that acquired from others.

Finally, to form optimal actions, followers need to forecast the ultimate direction of the organization a_L. Followers form this forecast using their private signal $\theta_i \sim \mathcal{N}(\theta, \sigma_\theta^2)$ and the information in the leader's mission statement.

Definition 1 *A Perfect Bayesian Nash Equilibrium is given by*

(i) a strategy, or direction, for the organization a_L *that maximizes* $E\left[\Pi_L | S_1, S_2\right]$*;*

(ii) set of followers' actions $\{a_i\}_{i\in[0,1]}$ *that maximize* $E[\Pi_i|\theta_i, S_1]$.

Optimal actions. We solve the model by backwards induction. When the leader chooses the organization's strategy a_L, the actions of the agents $\{a_i\}_{i\in[0,1]}$ are already determined. Since the first term of his payoff function (1) is zero, the leader's payoff in the final stage of the game reduces to $-E[(a_L-\theta)^2]$. The leader's optimal choice of strategy ex post then is to set a_L as close to the true state as possible: $a_L = E[\theta|S_1, S_2]$. According to Bayes' law, this expectation is

$$a_L = \lambda S_1 + (1-\lambda)S_2, \tag{2}$$

where the weight on the first signal is

$$\lambda := \sigma_p^{-2}/(\sigma_p^{-2} + \sigma_2^{-2}).$$

A rational leader (with $\sigma_p^2 = 1$) would set the weight λ equal to the relative precision of the first signal and the second: $1/(1+\sigma_2^{-2})$.

Each follower takes the actions of the others as given and cannot influence the average action because he is of measure zero. Therefore, his objective function (1) reduces to $E[-(a_i - a_L)^2|\theta_i, S_1]$. Consequently, his optimal action a_i is equal to his expectation of the leader's action given his own private signal θ_i: $a_i = E[a_L|\theta_i, S_1]$. Again by Bayes' law the follower's belief is

$$a_i = E[a_L|\theta_i, S_1] = \lambda S_1 + (1-\lambda)\left[\phi S_1 + (1-\phi)\theta_i\right]. \tag{3}$$

The term in square brackets is the follower's expectation of the leader's second signal $E[S_2|\theta_i, S_1]$. Since S_2 is an independent, unbiased signal about θ, $E[S_2|\theta_i, S_1] = E[\theta|\theta_i, S_1]$. The expectation of θ is a precision-weighted sum of S_1 and θ_i, where the weight on S_1 is

$$\phi := 1/(1+\sigma_\theta^{-2}).$$

Optimal overconfidence. We define the organization's payoff Π (without superscript) as the integral over all followers' payoffs plus the leader's payoff, assuming that the leader along with all followers is of zero measure.[1] The organization's ex-ante expected payoff therefore also has three components:

1. the variance of each follower's action around the leader's,

$$E[-(a_i - a_L)^2] = -(1-\lambda)^2(\phi + \sigma_2^2)$$

2. the dispersion of followers' actions around the mean,

$$\int_j -(a_j - \bar{a})^2 dj = -(1-\lambda)^2(1-\phi)^2\sigma_\theta^2$$

3. the distance of the leader's action from the true state,

$$E[-(a_L - \theta)^2] = -\lambda^2 - (1-\lambda)^2\sigma_2^2.$$

Summing the three terms and rearranging yields,

$$E\Pi = -(1-\lambda)^2(\phi + 2\sigma_2^2 + \phi(1-\phi)) - \lambda^2.$$

Recall that $\lambda := \sigma_p^{-2}/(\sigma_p^{-2} + \sigma_2^{-2})$ is a measure of the leader's overconfidence. The higher is λ the more overconfident is the leader. Therefore a simple way of determining the effects of leader overconfidence on the organization's overall welfare is to differentiate the ex ante objective with respect to λ.

[1]Note that all our qualitative results survive even if the leader has non-zero weight but the optimal level of overconfidence may vary.

The partial derivative of the organization ex-ante expected payoff with respect to λ is:

$$\frac{\partial E\Pi}{\partial \lambda} = 2(1-\lambda)(\phi + 2\sigma_2^2 + \phi(1-\phi)) - 2\lambda.$$

This is positive if

$$2\sigma_2^2 + \phi(2-\phi) > \frac{\lambda}{1-\lambda}$$

With a rational leader we have $\sigma_p^2 = 1$, and therefore $\frac{\lambda}{1-\lambda} = \sigma_2^2$. Thus, $\frac{\partial E\Pi}{\partial \lambda} > 0$ at $\frac{\lambda}{1-\lambda} = \sigma_2^2$, so that some degree of overconfidence is always optimal. On the other hand, for an extremely overconfident leader who fails to update at all, $\lambda \longrightarrow 1$, and the LHS of the inequality approaches infinity, so that $\frac{\partial E\Pi}{\partial \lambda} < 0$. As $\frac{\lambda}{1-\lambda}$ is continous for $\lambda \in (0,1)$, $\frac{\partial^2 E\Pi}{\partial \lambda^2} < 0$ and since the weight λ is strictly increasing in the perceived precision σ_p^{-2}, there exists an interior optimal level of overconfidence that maximizes the organization's expected payoff, which is given by

$$\sigma_p^{-2} = .2 + \phi(2-\phi)\sigma_2^{-2}. \tag{4}$$

The 2 in Equation (4) is due to the fact that there are two reasons why overconfidence increases the expected payoff of the organization: First, overconfidence reduces the distance of the followers' actions from the leader's action $(a_i - a_L)^2$. Second, weighting the later signal less reduces the error in the leader's action that comes from the noise in S_2. Of course, there is a corresponding increase in the weight on the first signal that increases the error in the leader's action that comes from noise in S_1. That effect is captured in the second term. The net effect of overconfidence is to increase the error in the leader's choice $(a_L - \theta)^2$.

We summarize this discussion in the proposition below.

Proposition 1 *The organization's ex-ante payoff is maximized with a leader's overconfidence level of* $\sigma_p^{-2} = 2 + \phi(2-\phi)\sigma_2^{-2} > 1$.

In particular, since the second term in Equation (4) is always positive, it is strictly beneficial for an organization to have an overconfident leader.

3.2 Strength of the mission statement

As the preceding analysis highlights, overconfidence of a leader provides a form of commitment to staying within the broad outlines of the mission defined by the leader. It ensures that the leader's strategy choice after learning his second signal does not deviate too much from the mission he set for the organization, which is centered on his belief conditional on the first signal S_1. If the leader's beliefs do not change much, his strategy choice will be similar to his mission statement. This commitment in turn facilitates coordination. However, to the extent that leader overconfidence also introduces a bias in the organization's adaptation to the environment, it would seem that a more direct solution to the leader's time-consistency problem–allowing a rational leader to commit to staying the course–would be preferable. We explore this question in this section by introducing such a commitment device into our model. Specifically, we add the possibility for the leader to stake his or the organization's reputation on carrying through a proposed mission. Should the leader choose to deviate from the proposed course of action then the organization will incur an additional cost that is increasing in the distance between the initial mission statement and the final strategy. The higher is this cost the stronger is the leader's mission statement.

An alternative interpretation of our commitment device is an incentive scheme for the leader, where the organization sets a punishment for deviating from the proposed mission (or a reward for carrying out a mission) that is increasing in the size of the deviation. It would seem that if the organization can incentivize a rational leader to optimally stay the course then there is no longer any role for leader overconfidence.

We shall argue, however, that overconfidence is still valuable. The reason is that a leader does not commit as much as is socially optimal because some of the benefit of commitment comes from internalizing coordination externalities. As long as the leader does not appropriate this entire benefit there will be too little commitment by a rational leader to staying the course.

In contrast, an overconfident leader will also make commitments to staying the course, but such a leader will commit even more than a rational leader to sticking to a mission and thereby helps close the wedge between his marginal value of commitment and the socially optimal value. There are three differences between overconfidence and commitment:

1. Commitment is a choice the leader makes, not an immutable type,

2. Commitment has payoff consequences, and

3. Commitment cost (strength of the statement) is a more flexible policy instrument. It could vary from project to project, while leader overconfidence is not malleable. At the same time, there are similarities between overconfidence and commitment: comparative statics for commitment cost are the same as for overconfidence along almost every dimension.

Model extension. We add one additional choice to the model: The leader can choose a cost that he and the organization will incur that is increasing in the distance between his mission statement and the chosen strategy. We call this cost the *strength of a mission statement* and interpret it as being a reputational cost borne by the leader (and the organization).

The leader's payoff now has a new last term that captures the cost of lost reputation.

$$\Pi_L = -(a_i - a_L)^2 - \int_j (a_j - \bar{a})^2 dj - (a_L - \theta)^2 - c(a_L - S_1)^2. \tag{5}$$

The commitment cost c determines how big the quadratic loss is from having a final strategy far away from the initial mission statement.

Optimal actions. Given this payoff, the first order condition for the leader's action now yield:

$$a_L = \frac{1}{1+c}\left((\lambda + c)S_1 + (1-\lambda)S_2\right),$$

where $\lambda = \sigma_p^{-2}/(\sigma_p^{-2} + \sigma_2^{-2})$. As before, each follower chooses his action to match its expectation of the organization's strategy: $a_i = E[a_L|\theta_i, S_1]$. But follower expectations now take a different form:

$$a_i = \frac{1}{1+c}\left\{\left[\lambda + c + (1-\lambda)\phi\right] S_1 + (1-\lambda)(1-\phi)\theta_i\right\}$$

Given that all members' actions vary with the reputation cost c it is natural to ask what payoff the organization could achieve if the reputation cost parameter c was chosen optimally. Alternatively, the optimal choice of c could also be interpreted as an optimal incentive scheme.

Thus consider the leader's choice of cost parameter c to maximize the his own ex-ante expected payoff:

$$\max_c E\Pi_L = -\left(\frac{1-\lambda}{1+c}\right)^2 \left[c\sigma_p^2 + (1+c)\sigma_2^2 + (1-\phi)^2\sigma_\theta^2\right] - \left(\frac{\lambda + c}{1+c}\right)^2 \sigma_p^2$$

Note that the expectation is taken given the leader's distorted beliefs about the precision of the initial signal, σ_p^{-2}. A stronger mission statement (higher c) shows up in the leader's expected utility in a way similar to a lower σ_p^2. Both increase the weight the leader puts on the first signal, relative to the second.

The cost that maximizes this payoff is

$$c_L^* = \max\left\{\frac{2\phi(1-\phi)}{\sigma_p^2 + \sigma_2^2} - 1, 0\right\} \tag{6}$$

The intuition here is that if followers have perfect information $\sigma_\theta^2 = 0$, there is no role for a commitment of the leader to stay the course, as followers are able to coordinate their actions independently of the leader. The -1 term in (6) arises because the leader expects to make some changes in the organization's strategy away from the initial mission statement and therefore wants to keep the cost of these changes small by choosing a low c. Note also that it is possible for the optimal reputation cost to be negative $c_L^* < 0$. One might interpret this as *commitment to reform.* In the results that follow, we consider choices of $c_L > 0$.

Lemma 1 *A more overconfident leader chooses a higher commitment cost.*

Proof. $\partial c_L^*/\partial \sigma_p^2 = -2(1-\phi)^2\sigma_\theta^2/(\sigma_p^2+\sigma_2^2)^2$. Since the numerator and denominator terms are both squares and the fraction is multiplied by -2, it must be negative. Since more overconfidence is defined as a lower σ_p, and $\partial c_L^*/\partial \sigma_p^2 < 0$, the commitment cost is increasing in overconfidence. ∎

Overconfident leaders choose higher commitment costs because they believe that the probability of taking an action far away from their mission statement is low. They systematically underestimate the cost they will pay.

Optimal strength of mission statement. The organization's payoff is the same as before, with the added reputation cost term $-c(a_L - S_1)^2$. The expected reputation cost is almost equal to the leader's expected benefit from completing the mission, with two differences:

1. the first payoff term is zero for the leader but non-zero for the organization:

$$E[-(a_i - a_L)^2] = -(1+c)^{-2}(1-\lambda)^2 E[(E[\theta|\theta_i, S_1] - S_2)^2] = -\left(\frac{1-\lambda}{1+c}\right)^2 (\phi + \sigma_2^2),$$

2. the followers are not overconfident like the leader and hence, correctly believe

that $Var[S_1] = 1$, not σ_p^2:

$$E\Pi = -\left(\frac{1}{1+c}\right)^2 \left[(1-\lambda)^2 \left(\phi(2-\phi) + 2\sigma_2^2 + c(\sigma_2^2+1)\right) + (\lambda+c)^2\right]$$

The next step is to determine the marginal payoff for the organization of having a more overconfident manager. An overconfident leader now affects the organization in two ways: through the weight λ put on the first signal and through the chosen commitment cost c. The socially optimal level of overconfidence sets the partial derivative $\frac{\partial E\Pi}{\partial \sigma_p^2}$ to zero. Substituting in for λ and for c_L delivers an equation in σ_p^2 that implicitly determines the optimal degree of overconfidence:

$$\left(\frac{\sigma_P^2}{2\phi^2(1-\phi)^2}\right)\left[-\sigma_2^2 - \phi(2-\phi) - \left(\frac{\sigma_2^2}{\sigma_P^2+\sigma_2^2}+1\right)\frac{\phi(1-\phi)(1+\sigma_2^2)}{\sigma_P^2+\sigma_2^2}\right] + \frac{1}{\phi(1-\phi)} = 0$$

Proposition 2 *Even with a commitment device which allows the leader to vary the strength of his mission statement, it is still optimal to choose an overconfident leader. However, the level of overconfidence is lower than when $c = 0$.*

The proof can be found in the Appendix.

Optimal signal-contingent reputation cost

The reader may wonder if the need for an overconfident leader despite the fact that the organization can incentivize the leader to optimally stay the course is due to the fact that the incentive scheme is imperfect. This intuition is correct. Specifically, if the organization can impose a reputation cost on the leader that is contingent on the realized signal S_2, then the organization can maximize their payoff with a rational leader. The organization can impose no reputation cost on the leader for taking the first best action and impose a very large negative reputation cost for taking any other action. In this manner, they can effectively take the choice of action away from the leader and make his preferences irrelevant.

This highlights that while overconfidence alleviates the time-consistency problem, it does not perfectly resolve it. However, whenever the signal S_2 is private information to the leader, or is not verifiable, it is not possible to implement such a signal-contingent scheme and to get the leader to verifyably commit to an action a_L conditional on the ralization of S_2.

4 Lead by being led

In this section, we replace the exogenous second signal with an endogenous signal: The leader observes the average action of the followers, plus some noise. In other words, the leader also learns about the environment from the followers through their choice of action, since they are based on their signals. In such a situation it is more important for the leader to let followers' base their actions on the signal they observe, as this brings about better adaptation.

Our main conclusion is that this moderates the benefit of overconfidence. A leader who is very overconfident dissuades his followers from acting based on their private information and suppresses information revelation. In other words, *overconfident managers make bad listeners* and learn little.

The second main result is that observing followers' actions creates a feedback effect that can generate multiple equilibria: If followers expect the leader to ignore the information from their actions, then the leader's initial announcement is a good estimate of his final action and there is no need for followers to take actions that depend on their private information. If followers do not act based on their private information, the leader rightly ignores the aggregate action because it contains no new information. On the other hand, if followers expect the leader to use the average action in forming policy, then they want to take actions that will be closer to that average action. In order to do that, followers use their private signals. Actions reflect their information.

Model setup. The payoff functions are as before (but with the commitment costs removed). Therefore, the leader's and followers' first order conditions are the same as in (2). Followers also form expectations over the state as before. However, now followers' actions aggregate into the second signal for the leader, which is the publicly observable organization output A:

$$A = \int_j a_j dj + e,$$

where the independent noise term, $e \sim \mathcal{N}(0, \sigma_e^2)$. As before, the leader uses the signal A along with S_1 to make a final inference about θ. Suppose that followers' equilibrium strategies take the form

$$a_i(\theta_i) = \beta\theta_i + (1-\beta)S_1, \tag{7}$$

then we can rewrite the aggregate output signal as

$$\hat{S}_2 := \tfrac{1}{\beta}\left[A - (1-\beta)S_1\right] = \theta + \tfrac{1}{\beta}e.$$

Note that this signal's precision is given by $\beta^2\sigma_e^{-2}$. Thus, the more followers rely on their private information (the higher is β), the more accurate this second signal becomes. Of course, if followers rely more on their private signals θ_i there is also less coordination among them. Thus, in this setting coordinated actions have both a positive payoff externality and a negative information externality because they supress information revelation to the leader.

Optimal actions. As in Section 3.1, the leader's optimal action is

$$a^L = E[\theta | S_1, \hat{S}_2] = \lambda S_1 + (1-\lambda)\hat{S}_2.$$

But now

$$\lambda = \frac{\sigma_p^{-2}}{\sigma_p^{-2} + \beta^2 \sigma_e^{-2}}, \tag{8}$$

where β is chosen by the followers and will depend on the leader's overconfidence level σ_p^{-2}.

As before, each follower's optimal action is their forecast of the leader's action:

$$a_i(\theta_i) = E[a^L|\theta_i, S_1] = \lambda S_1 + (1-\lambda)(\phi S_1 + (1-\phi)\theta_i).$$

Note that this is linear in S_1 and θ_i, which validates the conjecture in (7). Matching coefficients reveals that the weight followers place on their private signal is

$$\beta = (1-\lambda)(1-\phi). \tag{9}$$

Thus, the only difference in this new setting is that now λ depends on β and conversely β depends on λ. Therefore, to solve for the equilibrium actions we need to solve the fixed point problem given by the equations (8) and (9).

Substituting for λ in equation (9) delivers a third-order polynomial in β

$$\beta^3 - (1-\phi)\beta^2 + \sigma_p^{-2}\sigma_e^2\beta = 0.$$

This equation has three potential solutions: (i) a "dictatorial equilibrium" characterized by $\beta = 0$, $\lambda = 1$; and (ii) two "lead by being led equilibria" with

$$\beta = \frac{1}{2}\left[(1-\phi) \pm \sqrt{(1-\phi)^2 - 4\sigma_p^{-2}\sigma_e^2}\right].$$

Since we focus on stable equilibria we neglect the unstable equilibrium with the smaller quadratic root for β. Note that while the dictatorial equilibrium exists for any set of parameter values, the "lead by being led equilibrium" only exists for

$$(1-\phi)^2 > 4\sigma_e^2\sigma_p^{-2}. \tag{10}$$

Proposition 3 *When leaders learn from followers' actions, there are two stable (linear) equilibria:*
(i) A ***dictatorial equilibrium*** *where there is perfect coordination* $a_i = a^L = S_1$*, but information flow from followers to leaders is totally supressed.*
(ii) A ***"lead by being led equilibrium"*** *where coordination is reduced, but the organization is better adapted to the environment, as it relies on more information to determine its strategy:*

$$a_i = \beta\theta_i + (1-\beta)S_1 \qquad \textit{where} \qquad \beta = \frac{1}{2}\left[(1-\phi) + \sqrt{(1-\phi)^2 - 4\sigma_p^{-2}\sigma_e^2}\right],$$

and

$$a^L = E[\theta|S_1, \hat{S}_2] = \lambda S_1 + (1-\lambda)\hat{S}_2 \qquad \textit{where} \qquad \lambda = \frac{1}{2}\left[1 \pm \sqrt{1 - 4\sigma_e^2\sigma_p^{-2}(1-\phi)^{-2}}\right]$$

The logic of the multiple equilibria is the following: If followers expect leaders to learn no new information from their actions, then they expect the leader's action to be the same as his initial announcement ($a^L = S_1$). Since agents want to take actions close to the leader's action, they choose the same action $a_i = S_1$. But when agents all take the same actions, they reveal no new information. So, their expectation is self-confirming. In contrast, when followers expect the leader to learn new information, they try to forecast what he will learn, using their private signals. Because their actions are based on this forecast and on their private signals, aggregate output reveals information. So, the expectation that the leader will learn is also confirmed.

Optimal overconfidence. In the dictatorial equilibrium (where $\beta = 0$, $\lambda = 1$ solution), overconfidence has no effect on the organization's ex-ante expected payoff because it only works through the coefficients β and λ which, in this case, do not depend on the leader's overconfidence.

In the stable lead-by-being-led equilibrium, the organization's expected payoff is

$$E\Pi = -(1-\lambda)^2(2\beta^{-2}\sigma_e^2 + \phi(2-\phi)) - \lambda^2.$$

Taking the deriviative with respect to leader's overconfidence, yields

$$\frac{\partial E\Pi}{\partial \sigma_p^2} = -\left[2(1-\lambda)\phi(2-\phi) - 2\lambda\right]\frac{-\sigma_2^2}{(\sigma_2^2+\sigma_p^2)^2}, \tag{11}$$

Overconfidence is optimal if the partial derivative $\frac{\partial E\Pi}{\partial \sigma_p^2}$ is negative at $\sigma_p^2 = 1$.

Proposition 4 *In the lead-by-being-led equilibrium, leader overconfidence increases the organization's expected payoff if and only if*

$$\beta^{-2}\sigma_e^2 < \phi(2-\phi). \tag{12}$$

Otherwise, "underconfidence" increases the expected payoff.

When is the leader's overconfidence beneficial? These are situations where the leader is already extracting most of the relevant information about θ. If the signal the leader sees from the followers output is already very precise (low $\beta^{-2}\sigma_e^2$), then the benefit of better coordination ($\phi(2-\phi)$) matters more than the marginal loss of signal quality. When the leader learns little from followers' actions (σ_e^2/β^2 is large), overconfidence worsens this problem. Underconfidence allows the leader to observe more precise information and take a better-directed final action.

Setting (11) equal to zero gives the optimal degree of overconfidence, as long as the learning equilibrium exists (10) and the second order condition hold. The existence condition is likely to be satisfied if the noise in output, the degree of leader overconfidence, and the true precision of the leader's initial signal are low, and the precision of agents' private information is high. In sum, overconfidence is most valuable when there is lots of noise in output and the leader's true signal precision is high. In these times, the risks of overconfidence, suppression of followers' information revelation and possibly taking a bad final action, are minimized. Note finally that the effect of information quality of followers' signals on the value of overconfidence is ambiguous.

Our solution has another interesting feature. So far we have normalized the variance of S_1 to one. If one generalizes it to σ_1^2, then one can see that if the leader's first period information is sufficiently precise, it becomes impossible for him to learn from actions of followers and overconfidence has no value. It might be better in some cases to have a leader who knows less, so that followers place some weight on their private information (at least in one equilibrium they do), and the leader learns from their actions. A leader might optimally inject extra noise into his signal S_1 to make it less desirable to coordinate on and generate more information for himself.

One way of interpreting the multiplicity of equilibria in this setting is that the role of leadership in an organization must be adapted to the organization's *culture*. In a dictatorial organization, where followers are expected to just coordinate around the leader's mission statement it is best to have a rational, well-informed, and competent leader. In contrast, and somewhat counter-intuitively, in a democratic organization, where followers are expected to take a lot of *initiatives* and where the leader learns from the followers' actions, it may nevertheless be best to have a somewhat overconfident leader. This is especially valuable if the leader has significantly better information about the environment than other members of the organization.

5 Conclusion

We have proposed a model of leadership in organizations that captures a fundamental tension between adaptation to changing circumstances and coordination of followers on a given course of action. Specifically, the leader's problem is to steer the organization towards the best overall strategy or mission, while communicating a clear mission to other organization members that helps them coordinate and implement the organization's strategy. We have stripped down our model of leadership to five main phases. In a first phase the leader assesses the environment and defines a mission for the organization. In a second phase, the other members attempt to coordinate around the leader's stated mission. Followers face their own dilemma, as they are aware that the leader may change the organization's strategy in a subsequent stage in light of new information he gets about the environment. Therefore, they use their private information to forecast the change in strategy. Since private information is heterogeneous, forecasts and resulting actions are heterogeneous. This is the coordination problem that the leader is trying to minimize. In a third phase, the leader gets new information and updates his assessment of the state. In a fourth phase, the leader chooses a direction for the organization. Fifth and last, the state is revealed and leader's and followers' payoffs are realized.

The main message of the paper is that the tension between coordination and adaptation creates a time-consistency problem. This problem is ameliorated when leaders are overconfident. Overconfidence causes the leader to stick to his guns because he fails to update as much as he rationally should. Even when the leader can pledge a commitment cost, being overconfident is helpful for two reasons: First, it induces the leader to make a stronger commitment not to change the organization's direction. The stronger commitment achieves better coordination. Second, overconfidence results in lower commitment costs paid because the overconfident leader makes smaller changes in direction. The model also illustrates the dangers of overconfidence in situations where followers have valuable information. Overconfident leaders are less likely to learn what their followers know and may therefore lead their organization in the wrong direction.

A Technical Appendix

A.1 Results: Basic Model

The organization's ex-ante expected payoff has three components:

1. the variance of each follower's action around the leader's,

$$\begin{aligned} E[-(a_i - a_L)^2] &= E[-(\lambda S_1 + (1-\lambda)\left[\phi S_1 + (1-\phi)\theta_i\right] - \lambda S_1 - (1-\lambda) S_2)^2] \\ &= E[-((1-\lambda)\left[\phi S_1 + (1-\phi)\theta_i\right] - (1-\lambda) S_2)^2] \end{aligned}$$

   Since θ_i, S_1, S_2 each have independent signal noise, and the coefficients in the previous expression add up to zero, we can subtract the true θ from each one and then have independent, mean-zero variables that we can take expectations of separately. The first term $\phi S_1 + (1-\phi)\theta_i - \theta$ is the posterior belief error of the follower. It has precision that is the sum of the signal precisions $(1 + \sigma_\theta^{-2})$ and therefore has variance $1/1 + \sigma_\theta^{-2} = \phi$. In the second term, $S_2 - y$ has variance σ_2^2. Therefore, expected utility is

$$\begin{aligned} E[-(a_i - a_L)^2] &= -(1-\lambda)^2\phi - (1-\lambda)^2\sigma_2^2 \\ &= -(1-\lambda)^2(\phi + \sigma_2^2) \end{aligned}$$

   But in the leader's utility, he sets this term =0 because i=L.

2. the dispersion of followers' actions around the mean,

$$\begin{aligned} \int_j -(a_j - \bar{a})^2 dj &= -E[(\lambda S_1 + (1-\lambda)\left[\phi S_1 + (1-\phi)\theta_i\right] - \lambda S_1 - (1-\lambda)\left[\phi S_1 + (1-\phi)\theta\right])^2] \\ &= -E[(1-\lambda)(1-\phi)\theta_i + (1-\lambda)(1-\phi)\theta)^2] \\ &= -(1-\lambda)^2(1-\phi)^2\sigma_\theta^2 \end{aligned}$$

   Since the first signal cancels out here, this expectation is the same for the leader and the organization.

3. the distance of the leader's action from the true state,

$$\begin{aligned} E[-(a_L-\theta)^2] &= -E[(\lambda S_1+(1-\lambda)S_2-\theta)^2] \\ &= -\lambda^2 E[(S_1-\theta)^2]+(1-\lambda)^2E[(S_2-\theta)^2] \\ &= -\lambda^2-(1-\lambda)^2\sigma_2^2. \end{aligned}$$

The leader believes that this is

$$E^P[-(a_L-\theta)^2] = -\lambda^2\sigma_P^2-(1-\lambda)^2\sigma_2^2$$

Organization's expected utility Summing the three terms and rearranging yields,

$$\begin{aligned} E\Pi &= -(1-\lambda)^2\phi-(1-\lambda)^2\sigma_2^2-(1-\lambda)^2(1-\phi)^2\sigma_\theta^2-\lambda^2-(1-\lambda)^2\sigma_2^2. \\ E\Pi &= -(1-\lambda)^2(\phi+2\sigma_2^2+(1-\phi)^2\sigma_\theta^2)-\lambda^2. \end{aligned}$$

while the expression in the paper was $. = -\left[(1-\lambda)^2\left(2\sigma_2^2+((1-\phi)^2+\phi)\sigma_\theta^2\right)+\lambda^2\right].$

The partial derivative of the organization ex-ante expected payoff with respect to λ is:

$$\frac{\partial E\Pi}{\partial\lambda} = 2(1-\lambda)(\phi+2\sigma_2^2+(1-\phi)^2\sigma_\theta^2)-2\lambda.$$

This is positive if

$$\begin{aligned} 2(1-\lambda)(\phi+2\sigma_2^2+(1-\phi)^2\sigma_\theta^2) &> 2\lambda \\ \phi+2\sigma_2^2+(1-\phi)^2\sigma_\theta^2 &> \frac{\lambda}{1-\lambda}. \end{aligned}$$

With a rational leader we have $\sigma_p^2 = 1$, and therefore $\frac{\lambda}{1-\lambda} = \sigma_2^2$. Thus, $\frac{\partial E\Pi}{\partial\lambda} > 0$ at $\frac{\lambda}{1-\lambda} = \sigma_2^2$, so that some degree of overconfidence is always optimal. On the other hand, for an extremely overconfident leader who fails to update at all, $\lambda \longrightarrow 1$, and the LHS of the inequality approaches infinity, so that $\frac{\partial E\Pi}{\partial\lambda} < 0$. As $\frac{\lambda}{1-\lambda}$ is continous for $\lambda \in (0,1)$, $\frac{\partial^2 E\Pi}{\partial\lambda^2} < 0$ and since the weight λ is strictly increasing in the perceived precision σ_p^{-2}, there exists an interior optimal level of overconfidence that maximizes the organization's expected payoff, which is given by

$$\phi + 2\sigma_2^2 + (1-\phi)^2\sigma_\theta^2 = \frac{\sigma_2^2}{\sigma_p^2}. \tag{13}$$

$$\sigma_p^{-2} = .2 + \phi\sigma_2^{-2} + (1-\phi)^2\sigma_\theta^2\sigma_2^{-2} \tag{14}$$

$$\sigma_p^{-2} = .2 + \phi\sigma_2^{-2} + \phi(1-\phi)\sigma_2^{-2} \tag{15}$$

$$\sigma_p^{-2} = .2 + \phi(2-\phi)\sigma_2^{-2} \tag{16}$$

This proves proposition 1.

A.2 Results: Overconfidence and Commitment

The leader's utility now has a new last term that captures the cost of commitment.

$$\Pi_L = -\int_j (a_j - \bar{a})^2 dj - (a_L - \theta)^2 - c(a_L - S_1)^2. \tag{17}$$

The commitment cost c determines how big the quadratic loss is from having a final action far away from the initial announcement. Given this utility, the first order condition for the leader's action yields

$$\begin{aligned}
-2(a_L - E[\theta]) - 2c(a_L - S_1) &= 0 \\
(1+c)a_L - E[\theta] - cS_1 &= 0 \\
(1+c)a_L &= \lambda S_1 + (1-\lambda)S_2 + cS_1 \\
a_L &= \frac{1}{1+c}\left((\lambda + c)S_1 + (1-\lambda)S_2\right),
\end{aligned}$$

where $\lambda = \sigma_p^{-2}/(\sigma_p^{-2} + \sigma_2^{-2})$. As before, each follower chooses his action to match its expectation of the organization's strategy: $a_i = E[a_L|\theta_i, S_1]$. But follower expectations now take a different form:

$$\begin{aligned}
a_i &= \frac{1}{1+c}\left((\lambda + c)S_1 + (1-\lambda)E[S_2|S_1, \theta_i]\right) \\
&= \frac{1}{1+c}\left((\lambda + c)S_1 + (1-\lambda)(\phi S_1 + (1-\phi)\theta_i)\right) \\
&= \frac{1}{1+c}\left\{\left[\lambda + c + (1-\lambda)\phi\right]S_1 + (1-\lambda)(1-\phi)\theta_i\right\}
\end{aligned}$$

The organization's ex-ante expected payoff now has 4 components

1. the variance of each follower's action around the leader's,

$$\begin{aligned} E[-(a_i - a_L)^2] &= E[-(\frac{1}{1+c}((\lambda+c)S_1 + (1-\lambda)S_2) - \frac{1}{1+c}\{[\lambda+c+(1-\lambda)\phi]\,S_1 + (1-\lambda)(1-\phi)\theta_i\})^2] \\ &= -\left(\frac{1}{1+c}\right)^2 E[((1-\lambda)S_2 - (1-\lambda)(\phi S_1 - (1-\phi)\theta_i))^2] \end{aligned}$$

   Since θ_i, S_1, S_2 each have independent signal noise, and the coefficients in the previous expression add up to zero, we can subtract the true θ from each one and then have independent, mean-zero variables that we can take expectations of separately. The second term $\phi S_1 + (1-\phi)\theta_i - \theta$ is the posterior belief error of the follower. As before, it has precision that is the sum of the signal precisions and therefore has variance $1/1+\sigma_\theta^{-2} = \phi$. In the second term, $S_2 - y$ has variance σ_2^2. Therefore, expected utility is

$$E[-(a_i - a_L)^2] = -\left(\frac{1}{1+c}\right)^2 (1-\lambda)^2(\phi + \sigma_2^2)$$

   For the leader, this component of utility is zero.

2. the dispersion of followers' actions around the mean,

$$\begin{aligned} \int_j -(a_j - \bar{a})^2 dj &= -\left(\frac{1}{1+c}\right)^2 E[[\lambda+c+(1-\lambda)\phi]\,S_1 + (1-\lambda)(1-\phi)\theta_i \\ &\quad - [\lambda+c+(1-\lambda)\phi]\,S_1 + (1-\lambda)(1-\phi)\theta)^2] \\ &= -\left(\frac{1}{1+c}\right)^2 E[(1-\lambda)^2(1-\phi)^2(\theta_i - \theta)^2] \\ &= -\left(\frac{1}{1+c}\right)^2 (1-\lambda)^2(1-\phi)^2\sigma_\theta^2 \end{aligned}$$

   The leader and the organization have the same perceived utility from this component because the first signal drops out.

3. the distance of the leader's action from the true state,

$$\begin{aligned} E[-(a_L - \theta)^2] &= -E[(\frac{1}{1+c}((\lambda+c)S_1 + (1-\lambda)S_2) - \theta)^2] \\ &= -\left(\frac{\lambda+c}{1+c}\right)^2 E[(S_1-\theta)^2] + \left(\frac{1-\lambda}{1+c}\right)^2 E[(S_2-\theta)^2] \\ &= -\left(\frac{\lambda+c}{1+c}\right)^2 - \left(\frac{1-\lambda}{1+c}\right)^2 \sigma_2^2. \end{aligned}$$

The leader believes that this component is

$$E^P[-(a_L-\theta)^2] = -\left(\frac{\lambda+c}{1+c}\right)^2\sigma_P^2 - \left(\frac{1-\lambda}{1+c}\right)^2\sigma_2^2.$$

4. Finally, expected utility depends on the commitment cost incurred.

$$\begin{aligned}
-c(a_L-S_1)^2 &= -c\left(\frac{1}{1+c}\left((\lambda+c)S_1+(1-\lambda)S_2\right)-S_1\right)^2 \\
&= -c\left(\frac{1-\lambda}{1+c}\right)^2(S_2-S_1)^2 \\
-E[c(a_L-S_1)^2] &= -c\left(\frac{1-\lambda}{1+c}\right)^2(1+\sigma_2^2)
\end{aligned}$$

The leader believes that this will be

$$-E^P[c(a_L-S_1)^2] = c\left(\frac{1-\lambda}{1+c}\right)^2(\sigma_P^2+\sigma_2^2)$$

Organization's expected utility Summing the three terms and rearranging yields the expected utility under the objective probability measure,

$$\begin{aligned}
E\Pi &= -\left(\frac{1-\lambda}{1+c}\right)^2(\phi+\sigma_2^2) - \left(\frac{1-\lambda}{1+c}\right)^2(1-\phi)^2\sigma_\theta^2 - \left(\frac{\lambda+c}{1+c}\right)^2 - \left(\frac{1-\lambda}{1+c}\right)^2\sigma_2^2 - c\left(\frac{1-\lambda}{1+c}\right)^2(1+\sigma_2^2). \\
&= -\left(\frac{1-\lambda}{1+c}\right)^2\left[c+\phi+(2+c)\sigma_2^2+(1-\phi)^2\sigma_\theta^2\right] - \left(\frac{\lambda+c}{1+c}\right)^2
\end{aligned}$$

Leader's expected utility But for the leader who has distorted beliefs, $(S_1-\theta)^2=\sigma_p^2$. Because the leader understands that the followers do not believe the same that he believes, this does not change ϕ. It changes only the third and fourth terms. Furthermore, the leader gets no adverse utility consequences from his utility being far away from itself, so the first term is zero. That makes the leader's expected utility

$$E\Pi_L = -\left(\frac{1-\lambda}{1+c}\right)^2\left[c\sigma_p^2+(1+c)\sigma_2^2+(1-\phi)^2\sigma_\theta^2\right] - \left(\frac{\lambda+c}{1+c}\right)^2\sigma_p^2$$

The leader chooses his strength of commitment c to maximize this expected utility. The first order condition is

$$\begin{aligned}
2(1+c)^{-3}\left((1-\lambda)^2(c\sigma_p^2+(1+c)\sigma_2^2+(1-\phi)^2\sigma_\theta^2)+(\lambda+c)^2\,\sigma_p^2\right)-\left(\frac{1-\lambda}{1+c}\right)^2(\sigma_p^2+\sigma_2^2)-\frac{2(\lambda+c)}{(1+c)^2}\sigma_p^2 &= 0 \\
(1-\lambda)^2(2c\sigma_p^2+2(1+c)\sigma_2^2+2(1-\phi)^2\sigma_\theta^2-(1+c)(\sigma_p^2+\sigma_2^2))+(\lambda+c)^2\,2\sigma_p^2-2(\lambda+c)(1+c)\sigma_p^2 &= 0 \\
(1-\lambda)^2((c-1)\sigma_p^2+(1+c)\sigma_2^2+2(1-\phi)^2\sigma_\theta^2)-2(\lambda+c)(1-\lambda)\sigma_p^2 &= 0 \\
c((1-\lambda)(\sigma_p^2+\sigma_2^2)-2\sigma_p^2)+(1-\lambda)(-\sigma_p^2+\sigma_2^2+2(1-\phi)^2\sigma_\theta^2)-2\lambda\sigma_p^2 &= 0
\end{aligned}$$

Note that $(1-\phi)^2\sigma_\theta^2=\sigma_\theta^{-2}/(1+\sigma_\theta^{-2})^2=\phi(1-\phi)$.

$$c((1-\lambda)\sigma_2^2-(1+\lambda)2\sigma_p^2)=(1-\lambda)(-\sigma_2^2-2\phi(1-\phi))+(1+\lambda)\sigma_p^2$$

Note that the second order condition is

$$(1-\lambda)(\sigma_p^2+\sigma_2^2)-2\sigma_p^2<0$$

In other words, in order for the FOC to characterize and optimum, it needs to be that

$$\begin{aligned}
(1-\lambda)\sigma_2^2 &< (2-(1-\lambda))\sigma_p^2 \\
(1-\lambda)\sigma_2^2 &< (1+\lambda)\sigma_p^2
\end{aligned}$$

If the SOC holds, the utility-maximizing commitment is

$$\begin{aligned}
c &= \frac{2(1-\lambda)\phi(1-\phi)}{(1+\lambda)\,\sigma_p^2-(1-\lambda)\sigma_2^2}-1 \\
&= \frac{2\phi(1-\phi)}{\sigma_2^2+\sigma_p^2}-1
\end{aligned}$$

Lemma 1: The partial derivative of c with respect to σ_p^2 is

$$\frac{\partial c_L^*}{\partial\sigma_p^2}=\frac{-2\phi(1-\phi)}{(\sigma_2^2+\sigma_p^2)^2}$$

Since $\phi<1$, this expression is always negative. Since more overconfidence means lower σ_p^2, more overconfidence increases the optimal c for the leader to choose.

Optimal overconfidence for the organization The expected utility for the organization is

$$
\begin{aligned}
E\Pi &= -\left(\frac{1}{1+c}\right)^2 (1-\lambda)^2(\phi+\sigma_2^2) - \left(\frac{1}{1+c}\right)^2 (1-\lambda)^2(1-\phi)^2\sigma_\theta^2 - \left(\frac{\lambda+c}{1+c}\right)^2 - \left(\frac{1-\lambda}{1+c}\right)^2 \sigma_2^2 \\
&\quad -c\left(\frac{1-\lambda}{1+c}\right)^2 (1+\sigma_2^2) \\
&= -\left(\frac{1-\lambda}{1+c}\right)^2 \left[\phi+\sigma_2^2+(1-\phi)^2\sigma_\theta^2+\sigma_2^2+c(1+\sigma_2^2)\right] - \left(\frac{\lambda+c}{1+c}\right)^2 \\
&= -\left(\frac{1-\lambda}{1+c}\right)^2 \left[2\sigma_2^2+\phi(2-\phi)+c(1+\sigma_2^2)\right] - \left(\frac{\lambda+c}{1+c}\right)^2
\end{aligned}
$$

For the leader, expected utility differs because there is no misalignment cost and because the perceived variance of the first signal is lower.

$$
\begin{aligned}
E\Pi_L &= -\left(\frac{1-\lambda}{1+c}\right)^2 (1-\phi)^2\sigma_\theta^2 + -\left(\frac{\lambda+c}{1+c}\right)^2 \sigma_P^2 - \left(\frac{1-\lambda}{1+c}\right)^2 \sigma_2^2 - c\left(\frac{1-\lambda}{1+c}\right)^2 (\sigma_P^2+\sigma_2^2) \\
&= -\left(\frac{1-\lambda}{1+c}\right)^2 \left[\phi(1-\phi)+\sigma_2^2+c(\sigma_P^2+\sigma_2^2)\right] - \left(\frac{\lambda+c}{1+c}\right)^2 \sigma_P^2
\end{aligned}
$$

The leader chooses his commitment cost to maximize this. The FOC is

$$
+2\frac{(1-\lambda)^2}{(1+c)^3}\left[\phi(1-\phi)+\sigma_2^2+c(\sigma_P^2+\sigma_2^2)\right] - \left(\frac{1-\lambda}{1+c}\right)^2(\sigma_P^2+\sigma_2^2) - 2\left(\frac{\lambda+c}{1+c}\right)\left(\frac{1-\lambda}{(1+c)^2}\right)\sigma_P^2 = 0
$$

$$
\begin{aligned}
2(1-\lambda)\left[\phi(1-\phi)+\sigma_2^2+c(\sigma_P^2+\sigma_2^2)\right] &= (1-\lambda)(1+c)(\sigma_P^2+\sigma_2^2)+2(\lambda+c)\sigma_P^2 \\
(1-\lambda)\left[2\phi(1-\phi)+2\sigma_2^2+2c(\sigma_P^2+\sigma_2^2)-\sigma_P^2-\sigma_2^2-c(\sigma_P^2+\sigma_2^2)\right] &= 2(\lambda+c)\sigma_P^2 \\
(1-\lambda)\left[2\phi(1-\phi)+\sigma_2^2+c(\sigma_P^2+\sigma_2^2)-\sigma_P^2\right] &= 2(\lambda+c)\sigma_P^2 \\
(1-\lambda)\left[2\phi(1-\phi)+\sigma_2^2-\sigma_P^2\right]-2\lambda\sigma_P^2 &= c(2\sigma_P^2-(1-\lambda)(\sigma_P^2+\sigma_2^2)) \\
\frac{(1-\lambda)\left[2\phi(1-\phi)+\sigma_2^2\right]-(1+\lambda)\sigma_P^2}{((1+\lambda)\sigma_P^2-(1-\lambda)\sigma_2^2)} &= c \\
c &= \frac{2\phi(1-\phi)(1-\lambda)}{((1+\lambda)\sigma_P^2-(1-\lambda)\sigma_2^2)} - 1
\end{aligned}
$$

substitute in for lambda.

$$\begin{aligned} c &= \frac{2\phi(1-\phi)\sigma_2^{-2}}{(\sigma_2^{-2}+2\sigma_P^{-2})\sigma_P^2-\sigma_2^{-2}\sigma_2^2}-1 \\ &= \frac{2\phi(1-\phi)\sigma_2^{-2}}{\sigma_2^{-2}\sigma_P^2+2-1}-1 \\ &= \frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}-1 \end{aligned}$$

Note that more overconfidence raises the commitment cost the leader chooses, unambiguously.

The next step is to determine the organization's expected utility, as a function of overconfidence. Substitute in for c and lambda, noting that $\lambda = \sigma_2^2/(\sigma_P^2+\sigma_2^2)$.

$$\begin{aligned} E\Pi &= -\left(\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2\left[2\sigma_2^2+\phi(2-\phi)+\left(\frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}-1\right)(1+\sigma_2^2)\right]-\left(\frac{\sigma_2^2+2\phi(1-\phi)-(\sigma_P^2+\sigma_2^2)}{2\phi(1-\phi)}\right)^2 \\ &= -\left(\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2\left[\sigma_2^2+\phi(2-\phi)+\left(\frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}\right)(1+\sigma_2^2)-1\right]-\left(1-\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2 \\ &= -\left(\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2\left[\sigma_2^2+\phi(2-\phi)+\left(\frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}\right)(1+\sigma_2^2)-1\right]-1+\frac{\sigma_P^2}{\phi(1-\phi)}-\left(\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2 \\ &= -\left(\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2\left[\sigma_2^2+\phi(2-\phi)+\left(\frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}\right)(1+\sigma_2^2)\right]-1+\frac{\sigma_P^2}{\phi(1-\phi)} \end{aligned}$$

Then, take the partial derivative w.r.t. overconfidence to get its marginal utility.

$$\begin{aligned} \frac{\partial E\Pi}{\partial \sigma_p^2} &= -2\left(\frac{\sigma_P^2}{2^2\phi^2(1-\phi)^2}\right)\left[\sigma_2^2+\phi(2-\phi)+\left(\frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}\right)(1+\sigma_2^2)\right]-\left(\frac{\sigma_P^2}{2\phi(1-\phi)}\right)^2\left[\frac{-2\phi(1-\phi)}{(\sigma_P^2+\sigma_2^2)^2}(1+\sigma_2^2)\right] \\ &\quad +\frac{1}{\phi(1-\phi)} \\ &= \left(\frac{\sigma_P^2}{2\phi^2(1-\phi)^2}\right)\left[-\sigma_2^2-\phi(2-\phi)-\left(\frac{2\phi(1-\phi)}{\sigma_P^2+\sigma_2^2}\right)(1+\sigma_2^2)+\frac{\sigma_P^2\phi(1-\phi)}{(\sigma_P^2+\sigma_2^2)^2}(1+\sigma_2^2)\right]+\frac{1}{\phi(1-\phi)} \\ &= \left(\frac{\sigma_P^2}{2\phi^2(1-\phi)^2}\right)\left[-\sigma_2^2-\phi(2-\phi)-\left(\frac{-\sigma_P^2}{\sigma_P^2+\sigma_2^2}+2\right)\frac{\phi(1-\phi)(1+\sigma_2^2)}{\sigma_P^2+\sigma_2^2}\right]+\frac{1}{\phi(1-\phi)} \\ &= \left(\frac{\sigma_P^2}{2\phi^2(1-\phi)^2}\right)\left[-\sigma_2^2-\phi(2-\phi)-\left(\frac{\sigma_2^2}{\sigma_P^2+\sigma_2^2}+1\right)\frac{\phi(1-\phi)(1+\sigma_2^2)}{\sigma_P^2+\sigma_2^2}\right]+\frac{1}{\phi(1-\phi)} \end{aligned}$$

Finally, evaluate this at $\sigma_P^2 = 1$ to see if overconfidence is optimal.

$$\frac{\partial E\Pi}{\partial \sigma_p^2}\,|_{\sigma p=1} = \left(\frac{1}{2\phi^2(1-\phi)^2}\right)\left[-\sigma_2^2-\phi(2-\phi)-\left(\frac{\sigma_2^2}{1+\sigma_2^2}+1\right)\frac{\phi(1-\phi)(1+\sigma_2^2)}{1+\sigma_2^2}+2\phi(1-\phi)\right]$$

This is negative if, meaning that some degree of overconfidence is optimal if

$$\begin{aligned}
-\sigma_2^2-\phi(2-\phi)-\left(\frac{\sigma_2^2}{1+\sigma_2^2}+1\right)\frac{\phi(1-\phi)(1+\sigma_2^2)}{1+\sigma_2^2}+2\phi(1-\phi) &< 0\\
-\sigma_2^2-\phi^2-\phi(1-\phi)\left(\frac{\sigma_2^2}{1+\sigma_2^2}+1\right) &< 0\\
-\sigma_2^2-\phi+\phi(1-\phi)-\phi(1-\phi)\left(\frac{\sigma_2^2}{1+\sigma_2^2}+1\right) &< 0\\
-\sigma_2^2-\phi-\phi(1-\phi)\frac{\sigma_2^2}{1+\sigma_2^2} &< 0.
\end{aligned}$$

Result 1 *It is* Π*-maximizing for a manager with a commitment technology to be overconfident.*

Proof. If $\frac{\partial E\Pi}{\partial \sigma_p^2}<0$ when evaluated at $\sigma_p=1$, then it is welfare-increasing for the manager to believe that the first-period signal variance is lower than it actually is.

$$\left.\frac{\partial E\Pi}{\partial \sigma_p^2}\right|_{\sigma_p=1}=-\sigma_2^2-\phi-\phi(1-\phi)\frac{\sigma_2^2}{1+\sigma_2^2}$$

Since this expression is always negative, overconfidence ($\sigma_p<1$) is always welfare-maximizing. ■

Result 2 *The optimal level of overconfidence with commitment is lower than without it.*

Proof: We prove this by substituting in the optimal level of overconfidence with no commitment σ_p^* from (4) into the first-order condition in the environment with commitment. We show that the resulting first-order condition is negative. Since the second derivative is negative, a value of the first order condition lower than zero implies that the level of no-commitment overconfidence is higher than what is optimal in the commitment setting.

Substitute the optimal sigmap $\sigma_p^{-2}=.2+\phi(2-\phi)\sigma_2^{-2}$,which is equivalent to $\sigma_p^2=.1/(2+\phi(2-\phi)\sigma_2^{-2})$ in section 1 into the first-order condition above from section 2 delivers $FOC(\sigma_p^{nocomm})$. This is an expression in two variables σ_2^2 and ϕ. Simply plotting the function reveals that the inequality holds for all $\phi\in[0,1]$ and for all $\sigma_2^2\in[0,1000]$.

A.3 Results: Learning from Followers

Substituting for $\lambda = \sigma_p^{-2}(\sigma_p^{-2} + \beta^2\sigma_e^{-2})$ into $\beta = (1-\lambda)(1-\phi)$ yields

$$\begin{aligned} \beta &= \left(\frac{\beta^2\sigma_e^{-2}}{\sigma_p^{-2} + \beta^2\sigma_e^{-2}}\right)(1-\phi) \\ \beta\sigma_p^{-2} + \beta^3\sigma_e^{-2} &= \beta^2\sigma_e^{-2}(1-\phi) \end{aligned}$$

One solution to this equation is $\beta = 0$. The others can be found by dividing through by β and using the quadratic formula.

$$\beta^2\sigma_e^{-2} - \beta\sigma_e^{-2}(1-\phi) + \sigma_p^{-2} = 0$$

$$\begin{aligned} \beta &= \left[\sigma_e^{-2}(1-\phi) \pm \sqrt{\sigma_e^{-4}(1-\phi)^2 - 4\sigma_e^{-2}\sigma_p^{-2}}\right]/2\sigma_e^{-2} \\ & \frac{1}{2}\left[1-\phi \pm \sqrt{(1-\phi)^2 - 4\sigma_e^2\sigma_p^{-2}}\right] \end{aligned}$$

This means that

$$\lambda = \frac{\beta}{1-\phi} = \frac{1}{2}\left[1 \pm \sqrt{1 - 4\sigma_e^2\sigma_p^{-2}(1-\phi)^{-2}}\right]$$

The organization's expected payoff is the same as in section 1, except that σ_2^2 is replaced with $\beta^{-2}\sigma_e^2$.

$$E\Pi = -(1-\lambda)^2(2\beta^{-2}\sigma_e^2 + \phi(2-\phi)) - \lambda^2$$

Next, substitute $\beta^{-2} = (1-\lambda)^{-2}(1-\phi)^{-2}$ to get an expression with only λ and parameters.

$$E\Pi = -2(1-\phi)^{-2}\sigma_e^2 - (1-\lambda)^2\phi(2-\phi) - \lambda^2$$

Then, take a partial derivative of this payoff with respect to overconfidence in order to derive the optimal level of overconfidence.

$$\frac{\partial E\Pi}{\partial \sigma_p^2} = [2(1-\lambda)\phi(2-\phi) - 2\lambda]\,\frac{\partial \lambda}{\partial \sigma_p^2}$$

We know that more overconfidence always makes the leader weight his first signal more.

$$\frac{\partial \lambda}{\partial \sigma_p^2} = \frac{\partial}{\partial \sigma_p^2} \frac{\sigma_2^2}{\sigma_2^2 + \sigma_p^2} = \frac{-\sigma_2^2}{(\sigma_2^2 + \sigma_p^2)^2}$$

Thus, $\frac{\partial \lambda}{\partial \sigma_p^2} < 0$. The partial derivative is negative at $\sigma_p^2 = 1$, meaning that overconfidence is optimal if

$$\begin{aligned} (1-\lambda)\phi(2-\phi) &> \lambda \\ \phi(2-\phi) &> \frac{\lambda}{1-\lambda} \\ \phi(2-\phi) &> \frac{1}{\beta^2 \sigma_e^{-2}}. \end{aligned}$$

This proves proposition 4.

References

Aguilar, F. (1988): *General Managers in Action.* New York: Oxford University Press, p.71.

Dessein, W., and T. Santos (2006): "Adaptive Organizations," *Journal of Political Economy*, 114(5), 956–995.

Gervais, S., and I. Goldstein (2007): "The Positive Effects of Biased Self-Perceptions in Firms," *Review of Finance*, forthcoming.

Goel, A., and A. Thakor (2000): "Overconfidence, CEO Selection and Corporate Governance," Michigan Business School.

Hermalin, B. (1998): "Torward an Economic Theory of Leadership," *American Economic Review*, 88, 1188–1206.

Rotemberg, J., and G. Saloner (1993): "Leadership Styles and Incentives," *Management Science*, 39, 1299–1318.

——— (2000): "Visionaries, Managers and Strategic Direction," *Rand Journal of Economics*, 31, 693–716.

Van Den Steen, E. (2005): "Organizational Beliefs and Managerial Vision," *Journal of Law, Economics, and Organization*, 21(1), 256–283.