

Costly Transparency*

Justin Fox[†] and Richard Van Weelden[‡]

July 9, 2010

Abstract

We consider whether a career-minded expert would make better decisions if the principal could observe the consequences of the expert's action. The previous literature has found that this "transparency of consequence" can only improve the efficacy of the expert's decision making. We show, however, that this conclusion is very sensitive to the specified cost structure. While learning the consequences of the expert's action makes the expert more likely to choose the action most likely to correspond to the true state of the world, when costs are asymmetric, this may be associated with a decrease in the principal's expected welfare. In addition, for a range of parameters, if the principal benefits from learning the consequences of the expert's action, her utility is higher if she observes only the consequences and not the action taken. As such, the optimal transparency regime will involve either the principal observing only the expert's action or only the consequences of the expert's action; it will but never be optimal to observe both. We illustrate these results with examples from finance and public policymaking.

Key Words: transparency; reputation; herding; principal-agent; welfare

*We would like to thank Dirk Bergemann, Stu Jordan, Ekaterina Malinova and Craig Volden for helpful comments.

[†]Assistant Professor of Political Science, Yale University. 77 Prospect St., New Haven, CT 06520. Phone: (203) 606-2608 Email: justin.fox@yale.edu

[‡]Corresponding Author: Max Weber Fellow, Department of Economics, European University Institute. Via delle Fontanelle 10, I-50014 San Domenico, Italy. Phone: +39-0346-8354-093 Email: richard.vanweelden@gmail.com

1 Introduction

In many situations individuals delegate authority to an informed expert: investors give money to a fund manager in the hopes that it will be invested wisely; voters elect politicians to make decisions for them; judges are charged with interpreting the law and upholding the constitution on behalf of the people. Of course, in situations where the principal delegates to an expert, there is always the concern that the expert will not act in the principal's interest. In particular, we consider a situation where the expert's objective is to promote his own career rather than to promote the principal's welfare. In such a situation there is the concern that the expert might ignore or distort socially valuable information in order to safeguard his reputation. It is usually argued that these concerns are lessened, however, if the principal can observe the outcome of the chosen policy: learning the outcome of the chosen policy creates an incentive for the expert to choose the policy most likely to give the principal a high payoff.

This does not necessarily mean, however, that the principal benefits from observing the consequences of the policy the expert selects; the principal may well prefer the expert to select a policy which is less likely to succeed given that some mistakes are more costly than others. There are many situations where the ex-ante less likely outcome is associated with a greater potential cost: stock market crashes are relatively rare events but can lead to large losses or even bankruptcy; the defendant in a trial is more likely to be guilty but the costs of convicting the innocent are greater than the costs of acquitting the guilty; a Senator who believed that Iraq more likely than not had weapons of mass destruction could still believe the war was misguided if the costs of invading without finding weapons were greater than the costs of not invading if Iraq had weapons. In such a setting, it might be socially optimal for an expert to go against the prior even if, conditional on his private information, that alternative is more likely to match the state of the world. In such a setting we show that observing the outcome of the chosen policy will *increase* the incentive to herd in a socially harmful manner.

Consider a setting where an expert may be high-ability, in which case he observes a perfectly accurate signal of the state, or low-ability, in which case he observes a noisy but informative signal of the state. Now consider the decision of a low-type expert who observes a signal that goes against the prior. If only the high-type expert were to ever go against the prior, and the state of the world is not learned, the low-type expert can mimic the behavior of the high-type expert by choosing the ex-ante less-likely alternative. So, in equilibrium, some low-type experts will follow their signal. Now, suppose the state of the world will be learned with certainty prior to assigning the expert's reputation. The low-type expert can no longer mimic the high-type expert by going against the prior – if the state matches the prior the expert will reveal himself to be low-ability. So, if the prior is high enough, the expert will ignore his information and stick with the action which corresponds

to the prior – even if it would be socially beneficial for the expert to follow his signal.

One extremely natural interpretation of our model is of a fund-manager. First, the manager’s reputation for competence is supremely important so it makes sense to model the manager’s objective as maximizing his reputation for competence. Second, a market crash is a relatively rare event but carries with it extremely large costs for investors. Third, there is extremely fast feedback about the outcomes resulting from the manager’s decisions. Finally, there is much concern in the financial literature about herd behavior (e.g. Scharfstein and Stein 1990). Our results then raise the possibility that there is so much herding precisely *because* there is so much information available about the quality of the manager’s choice: a manager who observes a signal that a crash is possible but not likely would be reluctant to act on such information for fear of being proven wrong. So, the career concerns of the manager lead him to expose the investors to excessive risk of a catastrophic loss – a distortion which is heightened when the manager’s performance is more easily evaluated.

Another situation where our model might apply is to a legislator who must decide whether to approve an executive’s proposal. Consider, for example, a U.S. Senator asked to vote on whether to authorize the Bush administration to use force against Iraq. Suppose the Senator (or more accurately the constituents the Senator represents) feel that war is justified if weapons of mass destruction exist but not justified if they don’t. In the lead-up to the war the prior that weapons of mass destruction existed was high but so were the costs of invading a country without weapons. If the Senator had weak information that weapons of mass destruction did not exist, our results indicate that such a Senator would be more likely to oppose the war if it would not be learned whether the weapons existed.¹

As the above examples indicate, we show that, in environments with a sufficiently unbalanced prior, learning the consequences of an expert’s action will only increase the incentive for experts to herd on the prior. As such, efforts to improve transparency, and thereby increase the speed with which the outcome of the chosen policy is revealed, though frequently beneficial, can, in some cases, re-enforce the distortions. Consequently, increasing the frequency of disclosure for financial institutions may have the counter-productive effect of increasing the manager’s incentive to herd. Similarly, increasing the effectiveness of the media, which makes it more likely that the consequences of a politician’s actions will be learned before the next election, can make politicians less likely to act pre-emptively to avert a disaster.²

In most of the agency literature, the more information the principal has, the better the agent’s behavior will be. However, the literature provides some examples where the principal may prefer

¹This assumes that the objective of the Senator was to signal competence rather than to signal ideology or toughness. It also assumes that the Senator’s vote would not affect the probability of learning the state.

²See Ashworth and Shotts (2010) for a model of an informative media and its effect on the governmental decision-making.

less information.³ Of particular relevance for our work, Prat (2005) argues that observing the action a career-minded expert takes may lead to socially harmful distortions.⁴ In the established results in the literature (e.g. Canes-Wrone et al. 2001, Maskin and Tirole 2004, Prat 2005, Gentzkow and Shapiro 2006), however, observing the consequences of the policy the expert selects can only benefit the principal. In particular, Prat (2005) describes “the main contribution of [Prat’s] paper is to show that, while transparency on consequences is beneficial, transparency on action can have detrimental effects.” (p.863). When costs are sufficiently asymmetric, however, we find the opposite: the principal is made better off observing the action the expert selects (“transparency of action”) but worse off observing the consequences of that action (“transparency of consequence”). That is, while transparency of consequence is potentially harmful, transparency of action is beneficial.

Whether “transparency of consequence” or “transparency of action” is beneficial will depend on the cost structure. In equilibrium, the behavior of a career-minded expert will not be affected by the distribution of costs; in contrast, the first-best decision rule will depend on the costs associated with different outcomes. As such, statements about the welfare implications of different transparency regimes will be sensitive to the specification of costs.

We show that, when the prior on the state of the world is sufficiently unbalanced, and the expert has private information about his own competence, it will be possible, in general, to order the likelihood of the low-type expert following his signal across the three different transparency regimes: the expert is most likely to herd on the prior when only the consequence is observed, and least likely to when only the action is observed. Interestingly, this means that, except when the expert’s decision corresponds to the first-best outcome under full transparency, some type of non-transparency will always result in an increase in the principal’s welfare; which type of non-transparency will depend on the costs of different kinds of mistakes. We show that, if the high-type expert’s signal is not perfectly accurate, the expert’s decision under full transparency will generically not result in the first best outcome, so some form of non-transparency will almost always be beneficial.

Section 2 describes the model, Section 3 presents our results, and Section 4 concludes. The proofs of our results are left to the Appendix.

³See Prat (2005) for a discussion. For example, the agent might not work as hard if more is known about his type (Holmstrom 1999) or if his action is more accurately observed (Dewatripont et al. 1999).

⁴In Prat’s setting, as in ours, the expert’s objective is to appear competent. Fox (2007) establishes similar results when the objective of the decision maker is to signal congruence rather than competence.

2 Model

A privately informed expert must make a decision on behalf of a principal.⁵ We assume that there are two states of the world, $\omega \in \{0, 1\}$, with the prior probability that the state is 1,

$$\Pr(\omega = 1) = \pi \geq \frac{1}{2}.$$

There are two possible policies which could be selected, $p \in \{0, 1\}$. The utility to the principal depends on whether the policy matches the state

$$u(p, \omega) = \begin{cases} 0 & \text{if } p = \omega, \\ -1 & \text{if } p = 0, \omega = 1, \\ -\kappa & \text{if } p = 1, \omega = 0. \end{cases} \quad (1)$$

Normalizing the payoff when the policy matches the state to be 0 in either state will be without loss of generality, as what matters is the difference between the payoff when the policy matches the state and when it does not. Notice that we are not assuming that the two types of errors are equally costly: if κ is large then selecting policy 1 in state 0 is more costly than selecting policy 0 in state 1. In what follows we will focus on the case where $\kappa \geq 1$, so the ex-ante less-likely outcome is also the one with the greater potential cost.⁶

The expert observes a private signal, $s \in \{0, 1\}$, of the state of the world. Experts are heterogeneous with regard to their ability: the expert can be either high ability or low ability, $t \in \{l, h\}$. We assume the expert is the high type with probability $\alpha \in (0, 1)$. The expert knows his own type but the principal knows only the distribution of expert types. A high-type expert observes a signal that is accurate with probability q_h , whereas a low-type expert observes a signal that is accurate with probability q_l . That is, $\Pr(s = 1|\omega = 1, t = h) = \Pr(s = 0|\omega = 0, t = h) = q_h$ and $\Pr(s = 1|\omega = 1, t = l) = \Pr(s = 0|\omega = 0, t = l) = q_l$. We assume $\frac{1}{2} < q_l < q_h \leq 1$.

The expert is the only active player in our model. His strategy maps his type and his signal of the state into a probability of selecting each action. Formally, the expert's strategy is represented by

$$\sigma(t, s) \in [0, 1],$$

the probability with which the expert selects action 0 for each type $t \in \{l, h\}$ and signal of the state, $s \in \{0, 1\}$. The principal takes no action but updates her belief about the expert's type based on the information she observes. Finally, we assume that the expert's concern is to appear competent, so he seeks to maximize the probability that the principal places on him being the high type. The key assumption underlying this payoff structure is that long term contracting between

⁵We use male pronouns for the expert and female pronouns for the principal.

⁶See Fox and Van Weelden (2010) for a model of executive-legislative interaction and oversight which uses this cost structure.

the principal and the expert is not possible.⁷ We compare the efficiency – defined as the principal’s expected utility in the current period – of different information structures.⁸

Following Prat (2005), we describe learning the outcome of the chosen policy before assigning reputation as “transparency on consequence” and observing the action taken by the expert as “transparency on action”. As Prat (2005) considers the effect of transparency over actions, our analysis will focus more heavily on the role of transparency over consequences.

Under the first information structure, the principal observes both the action taken by the expert and the outcome of the selected policy. We refer to this as *Full Transparency* (FT). If the state of the world is learned before reputations are assigned the expert’s reputation is given by

$$\lambda(p, \omega) \equiv Pr(t = h|p, \omega), (p, \omega) \in \{0, 1\}^2.$$

Under the information structure in which the principal does not observe the outcome of the chosen policy (equivalently does not observe the state of the world) we define

$$\lambda(p) \equiv Pr(t = h|p), p \in \{0, 1\}.$$

We refer to this as the *Non-Transparent Consequences* (NC) information structure. Notice that with both *Full Transparency* and *Non-Transparent Consequences* the probability of learning the state is independent of the chosen policy. Our results then will not be driven by a correlation between the action taken and the probability of learning the state, something which has been explored in the previous literature (e.g. Canes-Wrones et al. 2001, Levy 2005).

Finally, we consider the case where the principal does not observe the chosen policy but does observe whether or not the chosen policy matches the state. We can then define the updated reputation of the expert whether or not the policy matches the state,

$$\bar{\lambda}(i) \equiv Pr(t = h|\mathbb{I}(p, \omega) = i), i \in \{0, 1\}.$$

We refer to this as *Non-Transparent Action* (NA).⁹

⁷Clearly, if the principal were able write contracts with the expert which conditioned his payoff on the observed action or consequences, greater transparency can never be harmful to the principal.

⁸This means that we are ignoring the possible benefits to the principal from learning about the expert’s type. One setting in which the principal would derive no benefit from learning about the expert’s type is a perfectly competitive market without the possibility of long term contracting. As all experts would then be paid according to their expected value, all the gains/losses from learning about the expert’s competence would accrue to the expert. These assumptions are standard in the career concerns literature pioneered by Holmstrom (1982, 1999).

⁹As costs are asymmetric, we assume that the principal observes whether the policy matches the state but not her utility; consequently, the principal cannot infer the action chosen from the utility she receives. If the principal were to observe her utility, when costs are not symmetric, the principal must be able to infer the action taken either when the policy matches the state or when the policy and the state do not match. Further, when the principal would be able to back out the state would depend on the utility levels assigned for different actions. In our setup the assumption that $u(\omega, \omega) = 0$ in both states is without loss of generality. Note also that the information possessed by the principal, and the principal’s updating, is exactly the same in our model as in Prat (2003, 2005).

Under all three information structures, in equilibrium, the principal assigns her beliefs based on Bayesian updating and the expert chooses his strategy to maximize his expected reputation. Before proceeding to the results, we discuss what behavior would maximize the principal's welfare. In order to make the problem interesting, we restrict attention to parameters where delegating to an expert is potentially beneficial to the principal. We first assume that

$$\kappa \leq \frac{\pi}{1 - \pi},$$

which means that, if the principal did not have access to an expert, she would go with her prior and select policy 1. Similarly, the principal's welfare is higher if any expert, regardless of his type, who observes signal $s = 1$ chooses policy $p = 1$ rather than $p = 0$. While this assumption will not be necessary for any of our results which follow it does guarantee that the principal's welfare under the informative equilibrium we consider will be higher than in an equilibrium in which the expert always chooses policy $p = 0$. Second, we assume that

$$q_h > \pi,$$

which ensures that the principal would like the high-type expert to follow his signal even if it goes against the prior.

As $\pi \geq \frac{1}{2}$, the probability that the state is 1, after a low-type expert observes a signal 0, may still be greater than $\frac{1}{2}$. However, if κ is reasonably large, this does not mean that it is socially optimal for such an expert to ignore his signal. We can calculate

$$Pr(\omega = 1|t = l, s = 0) = \frac{\pi(1 - q_l)}{\pi(1 - q_l) + q_l(1 - \pi)}.$$

Note that the probability that the state is 1, conditional on a low-type expert observing a signal of 0, is greater than $\frac{1}{2}$ when $\pi > q_l$. Further, the expected utility of the principal if $p = 1$ is chosen is $-\kappa Pr(\omega = 0|t = l, s = 0)$, whereas the expected utility from selecting $p = 0$ is $-Pr(\omega = 1|t = l, s = 0)$. Re-arranging the above expression we have that, when κ is sufficiently high, the principal would benefit if a low-type expert selects policy $p = 0$ after observing signal $s = 0$.

Remark 1 *If $\kappa > \frac{\pi(1 - q_l)}{q_l(1 - \pi)}$ then it is socially optimal for low-type experts to select $p = 0$ after observing $s = 0$; if $\kappa < \frac{\pi(1 - q_l)}{q_l(1 - \pi)}$ then it is socially optimal for low-type experts to select $p = 1$ after observing $s = 0$.*

Notice that, when $\kappa > 1$, it is possible for the principal to prefer the low-type expert to take action $p = 0$, even when $\pi > q_l$. As the main contribution of this paper – showing that transparency over consequences can be socially harmful – will require that $\pi > q_l$, in our analysis we will assume

$\pi \in (q_l, q_h)$.¹⁰ Under this range of priors, with reasonable restrictions on the expert’s behavior, the equilibrium will be unique.

3 Results

3.1 Preliminaries

We now turn to the analysis of equilibrium behavior. As the expert’s payoff does not depend on the policy selected there will be many equilibria of this game. As is standard in the literature,¹¹ we restrict attention to informative equilibria and ignore “mirror” or “perverse” equilibria in which the high-type expert chooses the policy opposite to his signal. We will say that an equilibrium is *non-pooling* if both policies, $p = 1$ and $p = 0$, are selected by the expert with positive probability along the equilibrium path. Further, to rule out babbling and perverse equilibria, we focus on equilibria in which the expert plays a *monotone* strategy.

Definition 1 *The expert’s strategy, σ , is monotone if, for any type and signal combination $(t, s) \in \{0, 1\}^2$, if $\sigma(t, s) < 1$, then for all (t', s') such that*

$$Pr[\omega = 0|(t', s')] < Pr[\omega = 0|(t, s)],$$

we have $\sigma(t', s') = 0$. Similarly, if $\sigma(t, s) > 0$, then for all (t', s') such that

$$Pr[\omega = 0|(t', s')] > Pr[\omega = 0|(t, s)],$$

we have $\sigma(t', s') = 1$.

We will refer to an equilibrium as monotone if the expert is playing a monotone strategy. Monotone equilibria are then those in which the experts who have the strongest information that the state is 1 (respectively 0) are the ones most likely to propose policy $p = 1$ ($p = 0$). In particular, it assumes that if an expert chooses a specific policy, any expert who places a higher probability on that policy matching the state of the world must always choose that policy.

Our first result is that there will always exist a unique non-pooling, monotone Perfect Bayesian Equilibrium for any parameters under any of the three information systems we consider.

¹⁰The case where $\pi < q_l$ has been more heavily studied in the literature (e.g. Canes-Wrone et al. 2001). In that case, the low-type expert’s signal is more accurate than the prior, and learning the state of the world will increase the expert’s incentive to follow his signal, even if it goes against the prior. It would still be possible for learning the state of the world to be socially harmful when $\pi < q_l$, but only when $\kappa < 1$.

¹¹For example, see Prat (2005) and Levy (2004, 2007).

Proposition 1 *Suppose $q_l < \pi < q_h$. Under all three information structures, $j \in \{FT, NA, NC\}$, there exists a non-pooling, monotone Perfect Bayesian Equilibrium. This equilibrium involves the high-type expert following his signal and the low-type expert selecting policy $p = 1$ if he observes $s = 1$ and policy $p = 0$ with probability $\sigma_0^j \in [0, 1]$ if he observes $s = 0$. Further, this equilibrium is unique, up to the beliefs at off-path information sets.*

As the high-type expert will always follow his signal we can then characterize the expert's strategy in the non-pooling, monotone, equilibrium in each environment by the number σ_0^j , the probability with which a low-type expert selects policy $p = 0$ after observing signal $s = 0$ under information structure $j \in \{FT, NA, NC\}$. We now proceed to compare the equilibria across the different information structures.

3.2 Results ($q_h = 1$)

We first consider the case where the high-type's signal is perfectly accurate. That is, we assume that $q_h = 1$ and $q_l = q$. This simplifies the algebra as the principal would know with certainty that any "mistake" must have been made by the low-type expert. The main result of this section is that the low-type expert is most willing to act on his private information when information about the consequences of his action is suppressed. We further show that the low-type expert is least likely to act on his private information when the principal cannot observe the action taken.

Proposition 2 *Define $\pi^* \equiv \frac{q}{q+\alpha(1-q)} \in (q, 1)$. Then,*

1. *for $\pi \in [\pi^*, 1)$,*

$$0 = \sigma_0^{NA} = \sigma_0^{FT} < \sigma_0^{NC}.$$

2. *there exists $\pi_* \in [q, \pi^*)$ such that for $\pi \in (\pi_*, \pi^*)$,*

$$0 = \sigma_0^{NA} < \sigma_0^{FT} < \sigma_0^{NC}.$$

We first discuss the relationship between (*NC*) and (*FT*) when $\pi \geq \pi^*$. Consider first the case where the principal doesn't observe the consequences of the policy the expert selects (*NC*). If only the high-type expert were to ever select $p = 0$ then selecting 0 would reveal the expert to be high-ability with certainty. Hence, in equilibrium, the low-type expert must choose policy $p = 0$ with positive probability after observing $s = 0$.

Now consider the effect of revealing the consequences of the expert's decision (*FT*). If the consequences of the selected policy are not observed, the low-type expert who observes $s = 0$ can mimic the high-type expert who observes signal $s = 0$; if the principal observes the consequences of the chosen policy, however, this is not possible. If the low-type expert received an incorrect signal this will be revealed. Hence, if, conditional on the low-type expert's signal, the probability that

the state is 1 is sufficiently high, the expert would prefer to select policy $p = 1$ rather than increase the probability of revealing himself to have observed an incorrect signal by choosing policy $p = 0$.

As a high-type expert never observes an incorrect signal, the expert's expected reputation, from either policy choice, is equal to his reputation if proven correct multiplied by the probability that his choice will match the state. Now, as the principal would infer that the expert is the high-type with certainty if he is proven correct after choosing policy 0 but not after choosing policy 1, in order to support an equilibrium where all low-type experts choose policy $p = 1$, the low-type expert must believe that the state of the world is sufficiently more likely to be $\omega = 1$ than $\omega = 0$ after observing signal $s = 0$. That is, we will have that the low-type expert will always choose policy $p = 1$ if

$$\pi \geq \frac{q}{q + \alpha(1 - q)} \equiv \pi^*,$$

which is a stronger condition than $\pi > q$. Notice also that the smaller α is, the higher π must be for the low-type to always choose policy $p = 1$. This is because selecting $p = 0$ also reveals that the expert is more likely to be the high type, so, conditional on being correct, the expert's reputation is enhanced even more. As α gets larger, this "selection effect" becomes less important, and the condition approaches simply $\pi > q$.

Whether or not the principal benefits from hiding the consequences of the expert's decision depends on the costs: if κ is large, the principal would prefer low-type experts to select $p = 0$ when they observe $s = 0$, so the welfare of the principal is higher if she does not observe the consequences of the selected policy. Combining part (1) of the above proposition with Remark 1 we get the following corollary.

Corollary 1 *If*

$$\kappa > \frac{\pi(1 - q)}{q(1 - \pi)} \geq \frac{1}{\alpha},$$

then the utility of the principal is higher with Non-Transparent Consequences than under Full Transparency.

Notice that the above condition provides a sufficient, but not a necessary condition for "transparency over consequences" to be harmful. We have then established that, when κ and π are sufficiently large, the principal's welfare is higher if she does not observe the outcome from the selected policy. Note that in order for this corollary to apply we must have, $\kappa \geq \frac{1}{\alpha} > 1$, which is most easily satisfied when α is large – that is, when there are many high-type experts in the population. Finally note that, in order for the principal to be harmed by learning the outcome of the policy, two things must be true: first, the prior on the state being 1 must be sufficiently high that a low-type expert believes that state is more likely even if he observes a signal of 0; second, the cost to the ex-ante less likely outcome must be high enough that it is in the principal's interest for the low-type expert to follow his signal anyways.

Finally, note that while the low-type expert must choose policy $p = 0$ with positive probability under *Full Transparency* when $\pi < \pi^*$, a simple continuity argument establishes that, for a non-degenerate interval (π_*, π^*) he will do so with greater probability under *Non-Transparent Consequences*.

Having explored the implications of suppressing information about the consequences of the expert’s action, we now examine what happens when the expert’s action is not revealed, but the consequences are (*NA*). This is the “lack of transparency” considered in Prat (2005). In particular, Prat (2005) shows that the expert’s chosen policy is more likely to match the state of the world if the principal only observes the outcome of the policy, as opposed to the specific policy the expert selects.¹² In our setting, this means that the low-type expert will choose policy $p = 1$ regardless of his signal for a wider range of priors on the state (part (2) of Proposition 2). As the expert’s choice of policy cannot signal anything about his competence (since the principal doesn’t observe it), the expert will then choose the policy most likely to match the state of the world; when $\pi > q$ this means that the low-type will select $p = 1$ regardless of his signal. Hence there is a wider range of parameters under which the low-type expert will always choose policy $p = 1$ regardless of his signal when the action the expert takes is hidden.

We now consider whether this lack of transparency over actions is beneficial in our setting. In fact, we see that the result goes in the opposite direction of Prat’s when κ is large: transparency of action is beneficial, but transparency of consequence is not. Prat’s positive results about transparency over actions still hold in our setting – lack of transparency of action makes it more likely the expert’s policy choice matches the state – but, with asymmetric costs, this can be associated with a decrease in the principal’s expected welfare. When costs are symmetric ($\kappa \approx 1$) non-transparency of action can only be beneficial.

We now present the general results when the expert’s signal is not perfectly accurate. We will show that, when the prior is sufficiently unbalanced, full transparency is (generically) not the optimal information structure. When the costs are sufficiently asymmetric, the principal would prefer not to observe the consequences of the expert’s action; when the costs are symmetric, the principal would prefer to observe the consequences of the action but not the action itself.

3.3 Results ($q_h < 1$)

We now consider the equilibrium behavior when the high type’s signal is not perfectly accurate. While assuming that the high type’s signal is perfectly accurate simplifies the algebra it also means that the principal will believe that the expert is the low type with certainty if she ever learns that

¹²In Prat (2005), the expert has no private information about his own ability, though his results continue to hold in the case where the expert does. This is shown in the unpublished working paper version, Prat (2003) which corresponds to the case we consider.

the expert's signal did not match the state. Hence, for a wide range of parameters, if the principal observes the consequences of the expert's action, the low-type expert will always stick with the prior out of fear of being proven wrong. In contrast, if the high-type expert's signal is not perfect, the expert who received an incorrect signal may still be the high type; consequently, if only high-type experts were to take a certain action, then the principal would infer that they were the high-type regardless of the consequences of the action. So, when the action is observed, the low-type expert must also be willing to go against the prior with some probability in equilibrium. Under *Non-Transparent Action* the principal cannot distinguish which type of mistake was made, so the expert will choose the policy most likely to match the state. Our next result shows that, if the high-type expert's signal is sufficiently accurate, but not perfectly accurate, then we can rank the probability that the low-type expert will go against the prior under each of the three information structures.

Proposition 3 *For all $\pi > \pi_*$, there exists a $\bar{q}_h(q_l, \pi) < 1$ such that, for all $q_h \in (\bar{q}_h(q_l, \pi), 1)$,*

$$0 = \sigma_0^{NA} < \sigma_0^{FT} < \sigma_0^{NC}.$$

So we see that the low-type expert is most likely to act on his signal of the state when the consequences of his action are unobserved, and least likely when only the consequences are observed. If $\kappa \neq \frac{\pi(1-q_l)}{q_l(1-\pi)}$, then the principal will not be indifferent over the low-type expert's decision after observing $s = 0$. As such an expert randomizes with a non-degenerate probability under the *Full Transparency* information structure, some form of non-transparency will lead to a strict increase in the principal's welfare. Whether *Non-Transparent Consequences* or *Non-Transparent Action* is preferred will depend on the cost κ .

Corollary 2 *Suppose $\pi > \pi_*$ and $q_h \in (\bar{q}_h(q_l, \pi), 1)$. Then,*

1. *if $\kappa > \frac{\pi(1-q_l)}{q_l(1-\pi)}$, then the principal's welfare is highest with Non-Transparent Consequences.*
2. *if $\kappa < \frac{\pi(1-q_l)}{q_l(1-\pi)}$, then the principal's welfare is highest with Non-Transparent Action.*
3. *Full Transparency is only welfare maximizing if $\kappa = \frac{\pi(1-q_l)}{q_l(1-\pi)}$, in which case all three information structures generate the same payoff to the principal.*

So we see that, generically, when the prior is sufficiently unbalanced and the high-type expert's signal is sufficiently accurate, there is some form of non-transparency which strictly increases the principal's welfare. Which type of non-transparency is beneficial will depend on the cost structure. Further, when one form of non-transparency is beneficial the other form will be harmful. Thus, *Full Transparency* will always be the second most preferred information structure for any parameters. Hence, if some of the relevant parameters (π, q_l, κ) are unknown, *Full Transparency* may well be the most preferred information structure ex-ante. Further, it should be noted that if the principal

benefits from learning about the expert’s type – which she does not in this model – this could provide additional benefits to the *Full Transparency* information structure. Finally, as previously noted, if the principal could contract with the expert based on her observations, then greater transparency can only benefit the principal.

3.4 Discussion

Before concluding, we pause to discuss the implications of our results. While the main focus of our analysis has been on the welfare implications of the different information regimes, we have also provided positive predictions about the behavior of career-minded experts. While the previous literature has identified environments in which improved information about the consequences of an expert’s decision will mitigate the distortions induced by career concerns (e.g. Canes-Wrone et al. 2001, Prat 2005, Gentzkow and Shapiro 2006), we have shown that when the prior is sufficiently unbalanced, experts will be *more* willing to go against the prior when the consequences of their actions are *less* transparent. Consequently, in many circumstances, we should expect individuals to be more willing to challenge the conventional wisdom when it will be more difficult to evaluate the consequences of their choices. That is, the incentive to herd can be heightened with increased transparency as experts are afraid to be proven wrong.¹³

4 Conclusions

We considered a setting in which the costs of one type of mistake are greater than the other. We showed that a career-minded expert will not adjust his behavior appropriately to deal with such costs, and showed that learning the consequences of the policy the expert selects may only increase the distortions. While learning the state increases the incentives for the expert to select the policy more likely to be “correct,” this may not be the policy which maximizes ex-ante welfare. So we have that, when the cost structure is sufficiently asymmetric, observing the outcome of the policy the expert selects makes the principal *worse* off.

As there are many situations in which costs are not symmetric, our results stand as an important caveat to the standard results in the literature (Prat 2005; Canes-Wrone et al. 2001; Maskin and Tirole 2004). In the career concerns literature, the expert’s objective, and hence equilibrium behavior, does not depend on the distribution of costs. The first-best decision rule, in contrast, will be closely tied to the cost structure. As such, statements about welfare will always be very sensitive to specified cost structure in the model.

¹³Robert Rodriguez, CEO of First Pacific Advisors, a hedge fund which divested its portfolio of subprime mortgages well before the financial crises, argues that when it comes to risky but widely held assets managers feel compelled to “be fully invested for fear of underperforming” (Rodriguez, 2009).

We have shown that hiding the consequences of the action increases the incentives to choose the action less likely to match the state, whereas, hiding the action taken will decrease this incentive. As the outcome with full transparency will generically fall short of the first best, some form of non-transparency will always be beneficial; which type of non-transparency is desired, however, will depend on the specific cost structure.

5 References

- Ashworth, Scott and Kenneth Shotts (2010). "Does Informative Media Commentary Reduce Politicians' Incentive to Pander?" *Journal of Public Economics*, forthcoming.
- Canes-Wrone, Brandice, Michael Herron, and Kenneth Shotts (2001). "Leadership and Pandering: A Theory of Executive Policymaking." *American Journal of Political Science*, 45, 3, 532-550.
- Dewatripont, Mathius, Ian Jewitt and Jean Tirole (1999). "The Economics of Career Concerns, Part I: Comparing Information Structures." *Review of Economic Studies*, 66, 1, 183-198.
- Fox, Justin (2007). "Government Transparency and Policymaking." *Public Choice*, 131, 23-44.
- Fox, Justin and Richard Van Weelden (2010). "Partisanship and the Effectiveness of Oversight." *Journal of Public Economics*, forthcoming.
- Gentzkow, Matthew and Jesse Shapiro (2006). "Media Bias and Reputation." *Journal of Political Economy*, 114, 2, 280-316.
- Holmstrom, Bengt (1982). "Moral Hazard in Teams." *Bell Journal of Economics*, 13, 2, 324-340.
- Holmstrom, Bengt (1999). "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies*, 66,1, 169-182.
- Levy, Gilat (2004). "Anti-herding and Strategic Consultation." *European Economic Review*, 48, 3, 503-525.
- Levy, Gilat (2005). "Careerist Judges and the Appeals Process." *Rand Journal of Economics*, 36, 2, 275-297.
- Levy, Gilat (2007). "Decision Making in Committees: Transparency, Reputation and Voting Rules." *American Economic Review*, 97, 1, 150-168.
- Maskin, Eric and Jean Tirole (2004). "The Politician and the Judge." *American Economic Review*, 94, 4, 1034-1054.

Prat, Andrea (2003). “The Wrong Kind of Transparency.” Center for Economic Policy Research. CEPR Discussion Paper No. 3859.

Prat, Andrea (2005). “The Wrong Kind of Transparency.” *American Economic Review*, 95, 3, 862-877.

Rodriquez, Robert L. (2009). “Reflections and Outrage”. Speech to the Morningstar conference on May 29, 2009.

Scharfstein, David and Jeremy Stein (1990). “Herd Behavior and Investment.” *American Economic Review*, 80, 3, 465-479.

6 Appendix

In this section we provide the proofs of Propositions 1-3. We begin with the proof of Proposition 1, the existence and uniqueness of a non-pooling, monotone equilibrium. As this result encapsulates statements about three different information structures we prove each one as a separate Lemma.

Lemma 1 *Suppose $\pi \in (q_l, q_h)$ and the principal only observes the action the expert takes but not the consequences ($j = NC$). Then there exists a unique non-pooling, monotone Perfect Bayesian Equilibrium. In this equilibrium the high-type expert always chooses the policy which matches his signal of the state, $p = s$. The low-type expert chooses policy $p = 1$ after observing $s = 1$ and randomizes with a non-degenerate probability after observing $s = 0$.*

Proof. We begin by noting that in any non-pooling, monotone Perfect Bayesian Equilibrium the high-type expert must always choose $p = s$. This follows immediately from the following argument: in a non-pooling equilibrium both actions must be chosen with positive probability; by monotonicity, if any expert ever chooses a given action, p , then the high-type expert who observed signal p must always choose that action; if the high-type were to randomize after a certain signal, all other experts would have to choose the other action with probability 1; if only the high-type expert were to choose a certain action, the principal would assign reputation $\lambda = 1$ after observing that action and all other experts would have an incentive to deviate and choose that action too.

Given that the high-type expert will always choose policy $p = s$, we can then restrict our analysis to the low-type expert’s behavior given that the high-type always follows his signal of the state. We can represent any expert strategy by the double (σ_1, σ_0) , the probability of the low-type expert choosing 0 after observing signal $s = 1$ and $s = 0$ respectively.

Our next task is to show that we must have $\sigma_1 = 0$. We do this by contradiction: suppose $\sigma_1 > 0$. Then, by monotonicity, we have $\sigma_0 = 1$. Now we can calculate the reputations of the

expert after proposing each policy for each $\sigma_1 \in (0, 1]$. Note that by Bayes' rule we can calculate,

$$\begin{aligned}\lambda(1|\sigma_1, \sigma_0 = 1) &= \frac{Pr(t = h, p = 1)}{Pr(t = h, p = 1) + Pr(t = l, p = 1)} \\ &= \frac{\alpha[\pi q_h + (1 - \pi)(1 - q_h)]}{\alpha[\pi q_h + (1 - \pi)(1 - q_h)] + (1 - \alpha)(1 - \sigma_1)[\pi q_l + (1 - \pi)(1 - q_l)]},\end{aligned}$$

and

$$\begin{aligned}\lambda(0|\sigma_1, \sigma_0 = 1) &= \frac{Pr(t = h, p = 0)}{Pr(t = h, p = 0) + Pr(t = l, p = 0)} \\ &= \frac{\alpha[(1 - \pi)q_h + \pi(1 - q_h)]}{1 - \alpha[\pi q_h + (1 - \pi)(1 - q_h)] - (1 - \alpha)(1 - \sigma_1)[\pi q_l + (1 - \pi)(1 - q_l)]}.\end{aligned}$$

Note that $\lambda(0|\sigma_1, \sigma_0 = 1)$ is decreasing in σ_1 and $\lambda(1|\sigma_1, \sigma_0 = 1)$ is increasing in σ_1 . Further, after simplifying the algebra we can see that,

$$\lambda(0|\sigma_1 = 0, \sigma_0 = 1) = \frac{\alpha[(1 - \pi)q_h + \pi(1 - q_h)]}{[(1 - \pi)q_h + \pi(1 - q_h)] + (1 - \alpha)(q_h - q_l)(2\pi - 1)} < \alpha,$$

and

$$\lambda(1|\sigma_1 = 0, \sigma_0 = 1) = \frac{\alpha[\pi q_h + (1 - \pi)(1 - q_h)]}{[\pi q_h + (1 - \pi)(1 - q_h)] - (1 - \alpha)(q_h - q_l)(2\pi - 1)} > \alpha,$$

as $\pi > q_l > \frac{1}{2}$. So we have

$$\lambda(0|\sigma_1, \sigma_0 = 1) < \lambda(1|\sigma_1, \sigma_0 = 1),$$

whenever $\sigma_1 > 0$. Hence, the expert's reputation would be strictly higher from choosing policy $p = 1$, which means the above strategy cannot be part of an equilibrium. Hence in any equilibrium we must have $\sigma_1 = 0$.

Now we must show that there exists a unique equilibrium with $\sigma_1 = 1$ and $\sigma_0 \in [0, 1]$. We can then define $\lambda(p|\sigma_0)$ to be the reputation of the expert after each policy is chosen, given that σ_0 is the probability a low-type expert chooses $p = 0$ after observing $s = 0$. Now, given π , q_h , and q_l , we can define the reputational difference between choosing policy $p = 0$ and $p = 1$ for any probability of randomization σ_0 ,

$$\phi(\sigma_0) = \lambda(0|\sigma_0) - \lambda(1|\sigma_0). \quad (2)$$

Notice that we have an equilibrium iff $\phi(\sigma_0) = 0$. Note also that we have already established that $\phi(\sigma_0 = 1) < 0$.

By Bayes' rule we have

$$\begin{aligned}\lambda(0|\sigma_0) &= \frac{Pr(t = h, p = 0)}{Pr(t = h, p = 0) + Pr(t = l, p = 0)} \\ &= \frac{\alpha[(1 - \pi)q_h + \pi(1 - q_h)]}{\alpha[(1 - \pi)q_h + \pi(1 - q_h)] + (1 - \alpha)((1 - \pi)q_l + \pi(1 - q_l))\sigma_0}.\end{aligned}$$

Notice then that $\lambda(0|\sigma_0)$ is clearly decreasing in σ_0 with

$$\lambda(0|0) = 1 > \alpha.$$

Similarly we have that

$$\lambda(1|\sigma_0) = \frac{Pr(t = h, p = 1)}{Pr(t = h, p = 1) + Pr(t = l, p = 1)} = \frac{Pr(t = h, p = 1)}{1 - [Pr(t = h, p = 0) + Pr(t = l, p = 0)]}.$$

As the denominator is decreasing in σ_0 , we have that $\lambda(1|\sigma_0)$ is increasing in σ_0 . In addition,

$$\lambda(1|0) = \frac{\alpha - Pr(p = 0|\sigma_0 = 0)\lambda(0|0)}{1 - Pr(p = 0|\sigma_0 = 0)} < \alpha,$$

and we can conclude that $\phi(\sigma_0 = 0) > 0$. So we can conclude that $\phi(\sigma_0)$ is a decreasing function of σ_0 with $\phi(0) > 0$ and $\phi(1) < 0$. So we have established that there exists a unique solution to $\phi(\sigma_0) = 0$, and that $\sigma_0 \in (0, 1)$.

Hence, we conclude there exists a unique non-pooling, monotone Perfect Bayesian Equilibrium. This involves the high-type following his signal, and the low-type choosing policy $p = 1$ after observing signal $s = 1$ and choosing policy $p = 0$ with probability $\sigma_0^{NC} \in (0, 1)$ after observing signal $s = 0$. ■

Lemma 2 *Suppose $\pi \in (q_l, q_h)$ and the principal observes both the action the expert took and the consequences of that action ($j = FT$). Then there exists a non-pooling, monotone Perfect Bayesian Equilibrium which is unique up to the beliefs at off path information sets. In this equilibrium the high-type expert always chooses the policy which matches his signal of the state ($p = s$) and the low-type expert chooses policy $p = 1$ when $s = 1$ and chooses policy $p = 0$ with probability $\sigma_0 \in [0, 1]$ when $s = 0$.*

Proof. We begin by noting that in any non-pooling, monotone Perfect Bayesian Equilibrium the high-type expert must always choose $p = s$. To see this, first note that in a non-pooling equilibrium both actions must be chosen with positive probability; by monotonicity, if any expert choose as a given action p , then the high-type expert who observed signal p must choose that action; if the high-type were to randomize after a certain signal, all other experts would have to choose the other action with probability 1. Now there are two cases to consider: $q_h = 1$ and $q_h < 1$. When $q_h < 1$, if only the high-type expert were to choose a certain action, the principal must assign reputation $\lambda = 1$ after observing that action, regardless of the consequences, and all other experts would have an incentive to deviate and choose that action too. When $q_h = 1$, then the principal must assign reputation 1 after observing that policy and that $p = \omega$: hence, the high-type expert would strictly prefer to propose policy $p = s = \omega$ in that state of the world and it would not be an equilibrium to randomize.

Given that any non-pooling, monotone PBE must involve the high-type always choosing policy $p = s$, the expert's strategy can then be represented by the double (σ_1, σ_0) as in the Non-Transparent Consequences case. Recall also that monotonicity implies that either $\sigma_1 = 0$ or $\sigma_0 = 1$. We now show that we must have $\sigma_1 = 0$.

To see that we must have $\sigma_1 = 0$, suppose $\sigma_1 > 0$. Then we know $\sigma_0 = 1$ and so can calculate the reputations for the expert after each combination (p, ω) :

$$\begin{aligned}\lambda(1, 1|\sigma_1, \sigma_0 = 1) &= \frac{\alpha q_h}{\alpha q_h + (1 - \alpha)(1 - \sigma_1)q_l}, \\ \lambda(1, 0|\sigma_1, \sigma_0 = 1) &= \frac{\alpha(1 - q_h)}{\alpha(1 - q_h) + (1 - \alpha)(1 - \sigma_1)(1 - q_l)}, \\ \lambda(0, 1|\sigma_1, \sigma_0 = 1) &= \frac{\alpha(1 - q_h)}{\alpha(1 - q_h) + (1 - \alpha)(1 - q_l) + (1 - \alpha)q_l\sigma_1}, \\ \lambda(0, 0|\sigma_1, \sigma_0 = 1) &= \frac{\alpha q_h}{\alpha q_h + (1 - \alpha)q_l + (1 - \alpha)(1 - q_l)\sigma_1}.\end{aligned}$$

Note that when $q_h = 1$ and $\sigma_1 = 1$ then $(p, \omega) = (1, 0)$ is off the equilibrium path. As we are showing that an equilibrium cannot exist because the expert would have an incentive to deviate to action 1, and the action $p = 1$ is least attractive if the belief after $(1, 0)$ is 0, we can without loss of generality assume that the belief is 0. In the above equations, when $q_h = 1$, $\lambda(1, 0|\sigma_1, \sigma_0 = 1)$ equals 0 for all σ_1 including as σ_1 is taken to 1.

From the above equations we can note the following properties for any $\sigma_1 > 0$: $\lambda(1, 1) > \lambda(1, 0)$ and $\lambda(0, 0) > \lambda(0, 1)$ so being proven correct is beneficial. In addition we have $\lambda(1, 1) > \lambda(0, 0)$ and $\lambda(1, 0) \geq \lambda(0, 1)$.

Finally, as

$$Pr(\omega = 1|t = l, s = 1) > \pi > \frac{1}{2},$$

we can conclude that

$$E[\lambda(1, \omega|\sigma_1, \sigma_0 = 1)|t = l, s = 1] > E[\lambda(0, \omega|\sigma_1, \sigma_0 = 1)|t = l, s = 1].$$

Hence, the low-type expert would have a strict incentive to choose policy $p = 1$ rather than $p = 0$ after observing $s = 1$ for the specified strategies. Hence we cannot have a non-pooling, monotone equilibrium with $\sigma_1 > 0$.

We now turn to showing that there exists a unique PBE in which the high-type expert always choose policy $p = s$ and the low-type expert's strategy is represented by $\sigma_1 = 0$ and $\sigma_0 \in [0, 1]$. We begin by calculating, via Bayes's rule, the updated reputation for each policy-outcome pair while holding $\sigma_1 = 0$,

$$\begin{aligned}\lambda(1, 1|\sigma_0) &= \frac{\alpha q_h}{\alpha q_h + (1 - \alpha)q_l + (1 - \alpha)(1 - q_l)(1 - \sigma_0)}, \\ \lambda(1, 0|\sigma_0) &= \frac{\alpha(1 - q_h)}{\alpha(1 - q_h) + (1 - \alpha)(1 - q_l) + (1 - \alpha)q_l(1 - \sigma_0)}, \\ \lambda(0, 1|\sigma_0) &= \frac{\alpha(1 - q_h)}{\alpha(1 - q_h) + (1 - \alpha)\sigma_0(1 - q_l)}, \\ \lambda(0, 0|\sigma_0) &= \frac{\alpha q_h}{\alpha q_h + (1 - \alpha)\sigma_0 q_l}.\end{aligned}$$

Note that when $q_h = 1$ and $\sigma_0 = 0$ the information set $(p, \omega) = (0, 1)$ is off the equilibrium path. As the statement of the theorem says nothing about the uniqueness of off-path beliefs, and because $\lambda(0, 1) = 0$ is the belief most conducive to supporting an equilibrium with $\sigma_0 = 0$, there is no loss of generality in setting the beliefs at that information set equal to 0.

Next we show that, if we can verify that the low-type expert is optimizing after observing $s = 0$ then it will follow that all other experts are optimizing. To see this, note that, as in the previous case, being proven correct is beneficial: $\lambda(0, 0) > \lambda(0, 1)$ and $\lambda(1, 1) > \lambda(1, 0)$ for all σ_0 . Hence, if the low-type expert is willing to randomize with probability $\sigma_0 \in (0, 1)$ after observing $s = 0$, all other experts who observed $s = 1$ have a strict incentive to choose $p = 1$ and a high-type expert who observed $s = 0$ has a strict incentive to choose $p = 0$. If $\sigma_0 = 1$ then the low-type must have a (weak) incentive to choose policy $p = 0$ after observing $s = 0$, so the high-type has a strict incentive to choose $p = 0$ after observing $s = 0$; we have already established above that experts who observe $s = 1$ have a strict incentive to choose $p = 1$. Finally, if $\sigma_0 = 0$ then we established in the first paragraph of this proof that the high-type who observed $s = 0$ must choose policy $p = 0$; as the low-type who observed $s = 0$ has a (weak) incentive to choose $p = 1$ all experts who observe $s = 1$ must have a strict incentive to choose policy $p = 1$. Hence, we have shown that, if the low-type expert is optimizing under the specified strategy, all other experts will have a strict incentive to follow the specified strategies for all $\sigma_0 \in [0, 1]$. So all that remains to show is that there exists a unique $\sigma_0 \in [0, 1]$ such that it is optimal for the low-type expert upon observing $s = 0$ to choose $p = 0$ with probability σ_0 and $p = 1$ with probability $1 - \sigma_0$.

Notice that $\lambda(1, 1)$ is increasing, $\lambda(1, 0)$ is weakly increasing, $\lambda(0, 1)$ is weakly decreasing and $\lambda(0, 0)$ is decreasing in σ_0 . So we have that

$$\psi(\sigma_0) \equiv E[\lambda(0, \omega|\sigma_0) - \lambda(1, \omega|\sigma_0)|t = l, s = 0], \quad (3)$$

is decreasing in σ_0 for all q_h, q_l and π . Hence we can conclude that:

- if $\psi(0) \leq 0$ then we have an equilibrium if and only if $\sigma_0 = 0$.
- if $\psi(0) > 0$ and $\psi(1) < 0$ then there exists a unique $\sigma^* \in (0, 1)$ such that we have an equilibrium if and only if $\sigma_0 = \sigma^*$.
- if $\psi(1) \geq 0$ then we have an equilibrium if and only if $\sigma_0 = 1$.

So we can conclude that there exists a PBE with $\sigma_0 \in [0, 1]$ and further that this equilibrium is the unique monotone, non-pooling PBE up to the belief at off-path information sets (if such information sets exist). ■

Lemma 3 *Suppose $\pi \in (q_l, q_h)$ and the principal does not observe the action the expert took but does observe the consequences ($j = NA$). Then there exists a unique non-pooling, monotone Perfect Bayesian Equilibrium. In this equilibrium the high-type expert always chooses $p = s$ and the low-type expert always chooses $p = 1$ regardless of his signal of the state.*

Proof. We first show that in any non-pooling, monotone PBE we must have

$$\bar{\lambda}(1) > \bar{\lambda}(0).$$

To see this, we begin by noting that both actions $p = 0$ and $p = 1$ must be taken with positive probability. Now, by monotonicity, at least one action $p' \in \{0, 1\}$ must never be taken when the expert observes signal $s = 1 - p'$. Hence the expert who observes $s = p'$ must choose action p' with probability σ_h if they are the high-type, and σ_l if they are the low-type. Note that by monotonicity we must have either $\sigma_h = 1$ or $\sigma_l = 0$. Define $\pi' = Pr(\omega = p') \in \{\pi, 1 - \pi\}$. Now we can calculate $\bar{\lambda}(1)$ from the equation,

$$\bar{\lambda}(1) = \frac{Pr(t = h, p = \omega)}{Pr(p = \omega)}.$$

Note that, as $q_h > q_l$ and $\sigma_h \geq \sigma_l$,

$$\begin{aligned} Pr(p = \omega) &= Pr(p = \omega = p') + Pr(p = \omega = 1 - p') \\ &= \pi'[\alpha q_h \sigma_h + (1 - \alpha) q_l \sigma_l] + (1 - \pi')[\alpha q_h + (1 - \alpha) q_l + \alpha(1 - q_h)(1 - \sigma_h) + (1 - \alpha)(1 - q_l)(1 - \sigma_l)] \\ &< \pi'(q_h \sigma_h) + (1 - \pi')(q_h + (1 - \sigma_h)(1 - q_h)), \end{aligned}$$

and

$$Pr(t = h, p = \omega) = \alpha[\pi'(q_h \sigma_h) + (1 - \pi')(q_h + (1 - \sigma_h)(1 - q_h))] > \alpha Pr(p = \omega).$$

Therefore, we can conclude that $\bar{\lambda}(1) > \alpha$. Now, as

$$\alpha = Pr(p = \omega) \bar{\lambda}(1) + (1 - Pr(p = \omega)) \bar{\lambda}(0),$$

we can conclude that $\bar{\lambda}(1) > \bar{\lambda}(0)$ in any non-pooling, monotone equilibrium.

Finally, note that the expert will be optimizing, and hence we will have an equilibrium, if and only if, conditional of the expert's private information he always chooses the policy which is most likely to match the state. Now as $\pi \in (q_l, q_h)$, we have that,

$$P(\omega = 0 | s, t) = \begin{cases} \frac{q_h(1-\pi)}{q_h(1-\pi) + \pi(1-q_h)} > \frac{1}{2} & \text{if } s = 0, t = h, \\ \frac{q_l(1-\pi)}{q_l(1-\pi) + \pi(1-q_l)} < \frac{1}{2} & \text{if } s = 0, t = l, \\ < 1 - \pi < \frac{1}{2} & \text{if } s = 1, t \in \{l, h\}. \end{cases}$$

So we can conclude that in the unique non-pooling, monotone, Perfect Bayesian Equilibrium the high-type expert will choose policy $p = s$ and the low-type expert will choose policy $p = 1$ regardless of his signal. ■

Proof of Proposition 1. This result is immediate combining Lemmas 1-3. ■

We now turn our attention to Proposition 2, the proof of which we divide into two lemmas.

Lemma 4 *When $\pi \geq \frac{q}{q+\alpha(1-q)} \equiv \pi^*$ in the unique non-pooling, monotone Perfect Bayesian Equilibrium the high-type expert follows his signal of the state, the low-type proposes policy $p = 1$ whenever he observes signal $s = 1$, and:*

1. *if the principal observes only the action the expert has taken ($j = NC$) then the low-type expert who observes $s = 0$ chooses policy $p = 0$ with probability $\sigma_0^{NC} \in (0, 1)$.*
2. *if the principal observes the consequences of the expert's action, then whether the principal observes the action itself or not, $j \in \{FT, NA\}$, the low-type expert who observes $s = 0$ never chooses policy $p = 0$, so $\sigma_0^j = 0$.*

Proof. *Proof of Part 1:* This follows immediately from Lemma 1.

Proof of Part 2: There are two separate claims here as it encompasses both the case where the expert's action is observed and when the expert's action is hidden. Recall, however, that by Lemma 3 this claim holds for when only the consequences are observed ($j = NA$), so we need only prove this result with Full Transparency ($j = FT$). Notice that, in order to prove the result, we need only verify that the specified strategies constitute an equilibrium since, by Proposition 1, equilibrium behavior must be unique.

We begin by calculating the reputation $\lambda(p, \omega)$ for any policy-outcome pair given the above expert strategy. Note that the policy-outcome pair $(0, 1)$ will be off-path given the specified expert behavior; to make the equations as easy to satisfy as possible take $\lambda(0, 1) = 0$, though there may be other beliefs which support this as an equilibrium. Next we note that, as only the high-type expert ever chooses $p = 0$, and only the low-type ever chooses $p = 1$ when the state is 0, the principal must assign beliefs

$$\lambda(1, 0) = 0,$$

and

$$\lambda(0, 0) = 1.$$

Note also that by substituting in $q_h = 1$ and $\sigma_0 = 0$ into the equation for $\lambda(1, 1|\sigma_0)$ we get

$$\lambda(1, 1) = \alpha.$$

Finally, recalling that $\pi \geq \pi^*$, the probability that the state is 1 is

$$Pr(\omega = 1|t = l, s = 0) = \frac{\pi(1 - q)}{\pi(1 - q) + q(1 - \pi)} \geq \frac{1}{1 + \alpha}.$$

Consequently, the expected reputation from choosing $p = 1$, is at least as high as the expected reputation from choosing $p = 0$:

$$\begin{aligned} E(\lambda(1, \omega)|t = l, s = 0) &= \alpha Pr(\omega = 1|t = l, s = 0) \\ &\geq Pr(\omega = 0|t = l, s = 0) = E(\lambda(0, \omega)|t = l, s = 0). \end{aligned}$$

So we can conclude that it is a Perfect Bayesian Equilibrium for the low-type expert to choose policy $p = 1$ after observing signal $s = 0$. ■

Lemma 5 *There exists a $\pi_* < \pi^*$, such that for $\pi \in (\pi_*, \pi^*)$ in the unique non-pooling, monotone Perfect Bayesian Equilibrium the high-type expert follows his signal of the state, the low-type chooses policy $p = 1$ whenever he observes signal $s = 1$, and:*

1. *if the principal observes the expert's action, the low-type expert who observes signal $s = 0$ will choose policy $p = 0$ with probability $\sigma_0^j \in (0, 1)$ for $j \in \{NC, FT\}$. Further,*

$$0 < \sigma_0^{FT} < \sigma^{NC}.$$

2. *if the principal does not observe the action the expert took (NA), then the low-type expert who observes $s = 0$ never chooses policy $p = 0$ so $\sigma_0^{NA} = 0$.*

Proof. *Proof of Part 1:* We begin by considering the equations that determine σ_0^{FT} for each prior π . We found in Lemma 4 that the unique monotone, non-pooling equilibrium involves $\sigma_0^{FT} = 0$ when $\pi \geq \pi^*$, but that an equilibrium cannot be supported with $\sigma_0^{FT} = 0$ when $\pi < \pi^*$. So now we consider $\pi < \pi^*$, where we must then have $\sigma_0^{FT} > 0$. Note that when $\sigma_0^{FT} > 0$ there are no off-path information sets, so all beliefs can be derived by Bayes's rule. Substituting in $q_h = 1$ into the updated reputations we calculated in the proof of Lemma 2:

$$\begin{aligned} \lambda(1, 1|\sigma_0) &= \frac{\alpha}{\alpha + (1 - \alpha)q + (1 - \alpha)(1 - q)(1 - \sigma_0)}, \\ \lambda(1, 0|\sigma_0) &= 0, \\ \lambda(0, 1|\sigma_0) &= 0, \\ \lambda(0, 0|\sigma_0) &= \frac{\alpha}{\alpha + (1 - \alpha)\sigma_0 q}. \end{aligned}$$

Notice that and that when $\sigma_0 = 1$,

$$\lambda(1, 1|1) = \frac{\alpha}{\alpha + (1 - \alpha)q} = \lambda(0, 0|1),$$

and so

$$\begin{aligned} E[\lambda(1, \omega|\sigma_0 = 1)|t = l, s = 0] &= Pr(\omega = 1|t = l, s = 0) \frac{\alpha}{\alpha + (1 - \alpha)q} \\ &> Pr(\omega = 0|t = l, s = 0) \frac{\alpha}{\alpha + (1 - \alpha)q} = E[\lambda(0, \omega|\sigma_0 = 1)|t = l, s = 0], \end{aligned}$$

whenever $\pi > q$. So when $\pi > q$ in any equilibrium we must have $\sigma_0 < 1$. In equilibrium, because the low type must randomize, we must equate the expected reputations from the two actions.

Now recall the definition of $\psi(\sigma_0)$ from equation (3). As $\psi(\sigma_0)$ is decreasing we can implicitly define σ_0^{FT} to be the unique solution to

$$\psi(\pi, \sigma_0) = E[\lambda(0, \omega|\sigma_0) - \lambda(1, \omega|\sigma_0)|t = l, s = 0] = 0.$$

We write ψ as a function of π to make explicit that the equation which determines the probability of randomization depends on π . Note that ψ depends on π only through $Pr(\omega = 0|t = l, s = 0)$, and since $Pr(\omega = 0|t = l, s = 0)$ is continuously decreasing in π , so too is $\psi(\pi, \sigma_0)$. Further, as ψ is continuously differentiable and decreasing in σ_0 , by the implicit function theorem, we can write the solution to the above equation, $\sigma_0^{FT}(\pi)$, as a continuously differentiable function of π for $\pi \in (q, \pi^*)$.

Now recall the definition of $\phi(\sigma_0)$ from equation (2). For each π we can solve for the equilibrium probability that the low-type expert chooses policy $p = 0$ when only the action is observed ($j = NC$), σ_0^{NC} , by setting

$$\phi(\pi, \sigma_0) = \lambda(0|\pi, \sigma_0) - \lambda(1|\pi, \sigma_0) = 0.$$

Further, recall that

$$\lambda(0|\pi, \sigma_0) = \frac{(1 - \pi)\alpha}{(1 - \pi)\alpha + (1 - \alpha)((1 - \pi)q + \pi(1 - q))\sigma_0},$$

$$\lambda(1|\pi, \sigma_0) = \frac{\pi\alpha}{1 - [(1 - \pi)\alpha + (1 - \alpha)((1 - \pi)q + \pi(1 - q))\sigma_0]}.$$

We can then see immediately that $\phi(\pi, \sigma_0)$ is a continuous in π and σ_0 and decreasing in σ_0 . In addition, as $\sigma_0^{FT}(\pi^*) = 0$,

$$\phi(\pi^*, \sigma_0^{FT}(\pi^*)) > 0.$$

Due to the above inequality and the continuity of ϕ we can then define,

$$\pi_* = \inf\{\pi \in (q, \pi^*) : \forall \pi' > \pi, \phi(\pi', \sigma_0^{FT}(\pi')) > 0\} < \pi^*.$$

Consequently, for all $\pi \in (\pi_*, \pi^*)$, $\phi(\pi, \sigma_0^{FT}(\pi)) > 0$ and so

$$\sigma_0^{NC}(\pi) > \sigma_0^{FT}(\pi).$$

Proof of Part 2: This follows immediately from Lemma 3. ■

Proof of Proposition 2. Follows immediately from Lemmas 4 and 5. ■

We now turn our attention to the proof of Proposition 3.

Proof of Proposition 3. Recall that, by Lemma 3, $\sigma_0^{NA} = 0$ for all $\pi > \pi_* \geq q_l$ for all $q_h > \pi$. Therefore, it is sufficient to show that for all $\pi > \pi_*$ there exists a $\bar{q}_h(q_l, \pi)$ such that for all $q_h \in (\bar{q}_h(q_l, \pi), 1)$,

$$0 < \sigma_0^{FT} < \sigma_0^{NC}.$$

First note that it is immediate that we must have $\sigma_0^{FT} > 0$. If $\sigma_0^{FT} = 0$, then only high-type experts would ever choose policy $p = 0$ and the principal would then infer that any expert who chose policy $p = 0$ is the high-type with certainty, regardless of the policy consequences. All experts would then have a strict incentive to choose policy $p = 0$.

We now establish the second inequality, that $\sigma_0^{FT} < \sigma_0^{NC}$ when $\pi > \pi_*$ and q_h is sufficiently high. We begin by recalling the definition of the function ϕ from equation (2) and allow q_h to vary while holding π and q_l fixed,

$$\phi(q_h, \sigma_0) = \lambda(0|q_h, \sigma_0) - \lambda(1|q_h, \sigma_0).$$

Recall also that ϕ is continuously decreasing in σ_0 for all parameters. Next we note that $\phi(q_h, \sigma_0)$ is continuously differentiable in both its arguments, q_h and σ_0 , and strictly decreasing in q_h . To see this recall that

$$\begin{aligned} \lambda(0|q_h, \sigma_0) &= \frac{\alpha[(1-\pi)q_h + \pi(1-q_h)]}{\alpha[(1-\pi)q_h + \pi(1-q_h)] + (1-\alpha)((1-\pi)q_l + \pi(1-q_l))\sigma_0} \\ &= \frac{\alpha[\pi + (1-2\pi)q_h]}{\alpha[\pi + (1-2\pi)q_h] + (1-\alpha)((1-\pi)q_l + \pi(1-q_l))\sigma_0}. \end{aligned}$$

So we can see immediately that, as $\pi > \frac{1}{2}$ and $(1-\alpha)((1-\pi)q_l + \pi(1-q_l))\sigma_0 \geq 0$,

$$\frac{\partial \lambda(0|q_h, \sigma_0)}{\partial q_h} \leq 0.$$

Further, recall that

$$\begin{aligned} \lambda(1|q_h, \sigma_0) &= \frac{Pr(t=h, p=1)}{1 - [Pr(t=h, p=0) + Pr(t=l, p=0)]} \\ &= \frac{\alpha[\pi q_h + (1-\pi)(1-q_h)]}{1 - \alpha[\pi + (1-2\pi)q_h] - (1-\alpha)((1-\pi)q_l + \pi(1-q_l))\sigma_0} \\ &= \frac{\alpha[(2\pi-1)q_h + (1-\pi)]}{\alpha[(2\pi-1)q_h + (1-\pi)] + (1-\alpha)(1 - ((1-\pi)q_l + \pi(1-q_l))\sigma_0)}. \end{aligned}$$

So, as $\pi > \frac{1}{2}$ and $(1-\alpha)(1 - ((1-\pi)q_l + \pi(1-q_l))\sigma_0) > 0$, we can conclude that,

$$\frac{\partial \lambda(1|q_h, \sigma_0)}{\partial q_h} > 0.$$

Hence, we can conclude that

$$\frac{\partial \phi(q_h, \sigma_0)}{\partial q_h} < 0.$$

Hence, for each π and q_l , since $\phi(q_h, \sigma_0)$ is continuously differentiable and decreasing in both its arguments, by the implicit function theorem, we can implicitly define the continuous function $\sigma_0^{NC}(q_h)$ as the solution to

$$\phi(q_h, \sigma_0) = 0,$$

for each q_h . Now, recalling the definition of ψ from equation (3) we can define, for each π and q_l ,

$$\bar{q}_h(q_l, \pi) = \inf\{q_h \in (\pi, 1] : \forall q'_h > q_h, \psi(q'_h, \sigma_0^{NC}(q'_h)) < 0\}.$$

Note that as ψ is continuous and

$$\psi(q_h = 1, \sigma_0^{NC}(1)) < 0,$$

we have that

$$\bar{q}_h(\pi, q_l) < 1.$$

As we have now established that for all $q_h \in (\bar{q}_h, 1)$,

$$\psi(q_h, \sigma_0^{NC}(q_h)) < 0,$$

and given that ψ is decreasing in σ_0 with $\psi(q_h, \sigma_0^{FT}(q_h)) = 0$, we see that

$$\sigma_0^{FT}(q_h) < \sigma_0^{NC}(q_h).$$

Hence, we can conclude that for all $q_h \in (\bar{q}_h, 1)$,

$$0 = \sigma_0^{NA} < \sigma_0^{FT} < \sigma_0^{NC},$$

as claimed. ■