



Institute for Empirical Research in Economics
University of Zurich

Working Paper Series
ISSN 1424-0459

Working Paper No. 496

Biased Experts, Costly Lies, and Binary Decisions

Roland Hodler, Simon Loertscher and Dominic Rohner

July 2010

Biased Experts, Costly Lies, and Binary Decisions*

Roland Hodler[†], Simon Loertscher[‡] and Dominic Rohner[§]

July 2, 2010

Abstract

Decision makers lacking crucial specialist know-how often consult with better informed but biased experts. In our model the decision maker's choice problem is binary and her preferred option depends on the state of the world unknown to her. The expert observes the state and sends a report to the decision maker. His bias is such that he prefers the same decision for all states. Lying about the state leads to a cost that increases in the size of the lie. As a function of the size of the expert's bias and the decision maker's prior about the underlying state, three kinds of equilibrium behavior occur. In each case equilibrium consists of separating and pooling segments, and the decision maker takes the expert's preferred decision for some states for which she would not take this decision had she observed the state herself. The model has a variety of applications and extends to situations in which the decision maker may be naive and take the report by its face value, and to situations with multiple experts and uncertainty about the size of the expert's bias.

JEL classification: C72, D72, D82

Keywords: Experts, Policy Advice, Information Distortion, Costly Signalling

*We would like to thank Andrew Daughety, Winand Emons, Navin Kartik, Michael Raith, Jennifer Reinganum, Klaus M. Schmidt, and Jean Tirole for helpful discussions. This working paper supersedes a note that circulated under the title "False Alarm? Terror Alerts and Reelection." Dominic Rohner gratefully acknowledges financial support from the Swiss National Science Foundation (SNF grant no. 100014-122636).

[†]Study Center Gerzensee; and Department of Economics, University of Melbourne. Email: roland.hodler@szgerzensee.ch

[‡]Department of Economics, University of Melbourne. Email: simonl@unimelb.edu.au.

[§]Department of Economics, University of Zurich. Email: drohner@iew.uzh.ch.

1 Introduction

Many decision problems are binary in nature and characterized by the uncertainty that the decision maker faces about crucial decision-relevant facts. Examples include a policy maker's decision whether or not to realize a given infrastructure project or to support a struggling industry, a Board's decision whether or not to replace a company's CEO, the voters' decisions whether or not to reelect an incumbent government, or a judge's decision whether or not to convict a defendant. To reduce uncertainty decision makers often consult experts who are better informed about the underlying facts. However, experts may themselves have a preference over the decision, and this preference may not be well-aligned with the decision maker's preference. Examples include industry experts who are interested in receiving public support,¹ CEOs and incumbent governments who want to remain in power and who know more about their performance and competence than Boards and voters, respectively, and (expert) witnesses in trials who have private information about the level of fault of the defendant, but may be biased towards a specific outcome of the trial.

The question of how much information the decision maker can extract from the expert is of obvious interest and relevance for public policy and the economics of organization. To answer this and related questions, we study a model that can be sketched as follows. The underlying state of the world is unknown to the decision maker, but known to the expert. The decision maker's preferences are such that she likes to realize a given project if the state is high, and to reject it if the state is low.² The expert sends a – possibly biased – report to the decision maker about the state of the world. His preferences are such that he gets a benefit, which is independent of the state and referred to as his bias, if and only if the decision is to realize the project. Further, the expert incurs a cost of lying which increases in the size of his lie. This cost may be due to reputational concerns or psychological constraints, or may arise because misreporting is unlawful and may be punished. Both his bias and the cost of lying are common knowledge, and so are the preferences of the decision maker.

¹Defence spending illustrates this: Parliaments lack the technical know-how to assess the needs of the army, and therefore have to seek the advice of experts from the military-industrial complex who may benefit from procurement contracts.

²Throughout the paper we refer to the decision maker as “she” and to the expert as “he”.

Equilibrium behavior depends on the expert's bias and the decision maker's prior beliefs about the state of the world. In particular, three cases need to be distinguished: (i) moderately biased expert, (ii) strongly biased expert when the decision maker's prior is in favor of the expert's preferred decision, and (iii) strongly biased expert when her prior is unfavorable. In each case, equilibrium exhibits both separating and pooling segments. Hence some information is always revealed. In any equilibrium the expert manages to obtain his preferred decision for some states in which the decision maker would decide differently if she could observe the state directly. There is also an asymmetry for strongly biased experts: If the decision maker's prior is favorable, the expert manages to conceal information that would not be in his favor, so that the project is realized for all states, while for an unfavorable prior the expert succeeds in conveying some relevant information that is in his favor.

In terms of comparative statics, increasing lying costs cause moderately biased experts to misrepresent their information less, but have no effect on the behavior of strongly biased experts in the presence of a favorable prior (unless the cost increase is so large that the expert becomes moderately biased). If the expert is strongly biased and the decision maker's prior unfavorable, the probability of project realization *increases* with the costs of lying (unless the cost increase is so large that the expert becomes moderately biased). The reason is that in this case equilibrium play by the decision maker involves mixed strategies for the pooling observation. To keep the expert indifferent, costlier lying has to be rewarded with more frequent project realization.

As shown in an extension, results do not change when there is some chance that the decision maker is not fully rational and takes the expert's report by its face value. The model also extends naturally to situations with multiple experts with biases of different sizes and to situations with uncertainty about the size of the expert's bias. Moreover, we briefly discuss what happens when the expert has the right to remain silent.

The canonical model to study this sort of strategic interactions between an expert (or sender) and a decision maker (or receiver) is Crawford and Sobel (1982). We depart from their framework in two important ways. First, we assume that it is costly for the expert to misreport the state. Second, the choice variable of the decision maker is binary rather than

continuous. This second difference implies that there is a conflict of interest between the expert and the decision maker for low, but not for high states. While in Crawford and Sobel’s model the expert and the decision maker perpetually disagree about the optimal policy (even under complete information), their disagreement is only partial in our setup.³

Lying costs were suggested by Crawford and Sobel (1982, p. 1450) and have recently been introduced into their framework by Kartik et al. (2007) and Kartik (2009). The model of Kartik (2009) is closest to ours. He introduces lying costs, but maintains Crawford and Sobel’s assumptions that the decision maker’s choice variable is continuous and that the expert’s preferred action increases in the state of the world. Our analysis thus complements his with the key modification that the decision maker’s choice variable is discrete and that the expert’s preferences are state independent and, consequently, aligned with the decision maker’s preferences for some but not all states.⁴

Krishna and Morgan (2001) modify the framework of Crawford and Sobel by allowing for two experts who move sequentially.⁵ We discuss how our results differ from theirs in the extension in which we allow for multiple experts.

Our paper further relates to a number of contributions that do not directly build on Crawford and Sobel (1982): Morris (2001) analyzes political correctness by an expert who does not want to appear biased. His model is in some sense converse to ours, as there are two states of the world while the decision maker’s policy space is continuous (and there are no lying costs). Wrasai and Swank (2007) study expert behavior in a model in which the expert’s message space is binary and the decision maker is uncertain about the expert’s bias.⁶

Our framework permits various applications. First, some “experts” may have a vital self-

³Melumad and Shibano (1991) show for the linear-quadratic specification of the Crawford and Sobel (1982) model with a somewhat more general bias that ex ante communication needs not yield Pareto improvements for the sender and receiver. It follows from the extension presented in Section 4.4 that the same is true in our model when the expert is strongly biased and the prior favorable.

⁴Edmond (2010) studies a model with a continuum of decision makers who play a coordination game. When the ‘lying’ costs in his model are linear, his equilibrium structure is similar to ours when the expert’s bias is moderate.

⁵Cummins and Nyman (2005) provide a related model of competing experts. There are two experts and two states, and the experts simultaneously choose binary policies. The decision maker is uncertain about the state but updates her beliefs after observing the policies. Cummins and Nyman show that in equilibrium no information may be transmitted from the experts to the decision makers.

⁶Questions that are similar to ours are also analyzed in the literature on informational lobbying and bureaucracy (see, e.g., Swank et al., 1999, and van Winden, 2004).

interest that a “project” is continued rather than canceled. For example, CEOs sometimes have incentives to embellish the financial situation of their company or their track-record to induce the Board to retain them or to influence a shareholder vote.⁷ In this respect, our model complements the existing literature which mainly focuses on CEO’s influencing continuous outcome variables, such as the market price of the firm (Fischer and Verrecchia, 2000), his compensation (Goldman and Slezak, 2006), or the range of possible projects (Adams and Ferreira, 2007). For the same reasons, product managers may have incentives to exploit their information advantage and send biased reports on the profitability of some products to prevent capital reallocation. Similarly, universal banks face conflicts of interest when their private banking or brokerage section recommends an Initial Public Offerings that their investment bank section has recently taken public, or when they promote funds or Structured Investment Products that are managed by the same bank.⁸ Analogously, country directors in international development agencies may have incentives to embellish the reforms in the countries they are assigned to in order to increase their budget. In a similar vein, scientific experts advising the World Health Organization on whether or not to recommend that countries stockpile certain drugs may have incentives to inflate their reports about the severity of the disease if they have financial ties to the pharmaceutical firms producing these drugs.⁹

Second, the model applies to incumbent government behavior. For example, it predicts that an incumbent government who is known to be tough on terrorism may inflate its report on the threat from future terrorist attacks to improve its reelection prospects. This prediction is consistent with the allegations that former President George W. Bush may have misused terror alerts to obtain an electoral edge in the campaign against John Kerry in 2004. Indeed, by far the longest lasting rise to “high risk” in the Homeland Security Advisory System during

⁷A recent example is the merger of Bank of America with Merrill Lynch. It has been argued that Bank of America’s CEO Ken Lewis, who was in favor of the deal, had detailed knowledge that Merrill’s last quarter losses were going to be over \$12 billion, but did not share this crucial information with Bank of America shareholders when they voted on accepting the merger on December 5, 2008 (cf. *The Atlantic Monthly*, “An Offer He Couldn’t Refuse”, September 2009, p. 64/5). According to this interpretation, Lewis traded off substantial lying costs (the risk of being persecuted for “securities fraud”) with large rents from office.

⁸Empirical evidence suggests that the bank’s advice is biased in these situations and that investors are harmed by the information asymmetry (Michaely and Womack, 1999; Ber et al., 2001).

⁹The H1N1 influenza (“swine flu”) is a case in point. Cohen and Carter (2010, p. 1274) document that “scientists advising the World Health Organization on planning for an influenza pandemic had done paid work for pharmaceutical firms that stood to gain from the guidance they were preparing.”

the first term of George W. Bush was issued on August 1, 2004, three days after John Kerry accepted his presidential nomination at the Democratic National Convention, and lasted until November 10, 2004, eight days after President Bush's reelection. Many observers argued that the Bush administration may have been motivated by strategic political considerations, as voters preferred Bush over Kerry for the fight against terrorism.¹⁰

The application of our model to elections complements Banks (1990) and Callander and Wilkie (2007). While they analyze how two symmetric candidates make costly lies about the policies they will implement when elected, our model permits a focus on elections that are asymmetric in that an incumbent with an informational advantage about the state of the world runs for reelection. This application also relates our model to Rogoff and Sibert (1988) and Hodler et al. (2010), where an incumbent with private information about his competence or the state of the world may choose socially inefficient policies to improve his reelection prospects.

Third, our framework applies to courts, where (expert) witnesses provide – potentially biased – evidence that influences the binary decision of a judge or jury between “guilty” or “innocent”. Indeed, “witnesses often have a material interest in the court’s judgment. The plaintiff and defendant, for example, are interested in the stakes in the dispute, and an expert has an interest in future employment as a witness” (Cooter and Emons, 2003: p. 259). Biased (expert) witnesses thus trade off rents from their preferred verdict with lying costs. For example, as pointed out by Posner (1999), economic experts often have vested interests, as they are normally hired by one of the process parties. At the same time, they also face substantial lying costs, such as perjury penalties or the loss of their academic reputation. Our model complements the existing literature on biased witnesses in courts that either features binary state spaces (e.g., Cooter and Emons, 2003) or does not include a binary decision space (e.g., Sanchirico, 2001; Emons and Fluet, 2009).

The remainder of this paper is organized as follows: Section 2 sets out the basic model,

¹⁰Conservative Fox News was one of these observers, writing that “the advantage of incumbency were in full display” when the Bush administration issued this high terror alert (Fox News, “Dems Question Timing of Terror Alert”, August 4, 2004; see also Financial Times, “Bush administration defends latest warning on terrorist threat: Americans’ reaction to the government action has been sceptical”, August 4, 2004). Statements made by Homeland Security Secretary Tom Ridge after his resignation in February 2005 also point out potential threat exaggerations (USA Today, “Ridge reveals clashes on alerts”, May 10, 2005).

which is then solved in Section 3. Section 4 presents our extensions, and Section 5 concludes. Appendix A contains lengthy proofs.

2 The Model

There are two strategic players, the decision maker M and the expert E. M takes the binary decision whether or not to realize a certain project, but she is uncertain about the costs and benefits of this project. E is better informed about these costs and benefits, but he personally benefits if the project is realized.

We assume that the social costs and benefits depend on the state of the world s , where s is randomly drawn from the distribution $F(s)$ with density $f(s) > 0$ and support $[a, b]$. While $F(s)$ is common knowledge, only E observes s .¹¹ The timing is as follows: First, E observes s and sends a report $x(s) \in [a, b]$. Second, M observes the report x (but not s) and decides to realize the project with probability $v(x) \in [0, 1]$.¹²

In the state of the world s , M's utility is $u(1, s)$ if the project is realized, and $u(0, s)$ otherwise. Her net utility from realizing the project is $W(s) \equiv u(1, s) - u(0, s)$. We assume that $W(s)$ is continuous and satisfies $W'(s) > 0$ and $W(a) < 0 < W(b)$. Consequently, there exists a unique threshold $\tilde{s} \in (a, b)$ satisfying $W(\tilde{s}) = 0$. If she knew s , M would thus realize the project if $s > \tilde{s}$, but cancel it if $s < \tilde{s}$. But since she does not know s , she updates her beliefs $\mu(s|x)$ about s after observing x using Bayes' rule (where possible) before deciding whether or not to realize the project.

E is biased and gets a rent $\Phi > 0$ if and only if the project is realized. He thus has an incentive to misreport what he knows about the state of the world. That is, he might want to report $x(s) \neq s$. However, it is costly for him to distort his report x . In particular, he has to bear the distortion costs $c = c(d)$, where $d \equiv |x - s|$ and $c_d(d) > 0$.¹³ For simplicity, we

¹¹Assuming that E observes the state s directly eases the exposition. All our results go through even if E does not observe s , but only a noisy signal of it, provided the state and the signal are affiliated random variables.

¹²We allow for M to use mixed strategies. There would be no equilibrium in the model with the strongly biased expert who faces an unfavorable prior if M could not mix between her binary actions.

¹³Costly signal distortions are represented by similar cost functions in, e.g., Lacker and Weinberg (1989), Maggi and Rodriguez-Clare (1995), Crocker and Morgan (1998), Callander and Wilkie (2007), Kartik et al. (2007), Kartik (2009), and Emons and Fluet (2009). This simple cost function eases the exposition, but

assume $c(0) = 0$. These costs could represent mental or moral costs of lying. However, we prefer to think of them as reputational costs or expected punishment, i.e., as costs that arise because in the future M (or some third party) may learn s and punish E , e.g., by reclaiming some of E 's remuneration, by not employing E in the future, or even by inflicting an additional penalty. In this scenario, the expected costs c may increase in the distortion d for various reasons: First, M may later observe s , and the penalty may directly increase in d . Second, the penalty may be independent of d , but the probability that E is punished may increase in d .

Therefore, when distorting the report x by d , E 's utility is $\Phi - c(d)$ if the project is realized, and $-c(d)$ otherwise. We distinguish two cases depending on whether or not the rent Φ exceeds the costs $c(b-a)$ resulting from the largest possible distortion. If $\Phi \geq c(b-a)$, E is willing to send any report $x(s)$ for any state s if this guarantees that the project is realized rather than canceled for sure. If $\Phi < c(b-a)$, we therefore say that E is strongly biased. On the other hand, if $\Phi < c(b-a)$, E would not be willing to misreport that bluntly. For example, he would never report $x(a) = b$ or $x(b) = a$. In this case, we say that E is only moderately biased. Notice that it depends on both the rent Φ and the cost function $c(d)$ whether E 's bias is strong or moderate. We will look at the cases of moderately and strongly biased experts in turn.

The solution concept is perfect Bayesian equilibrium, and we focus on equilibria that satisfy the Intuitive Criterion (Cho and Kreps, 1987). We use the abbreviation PBE to stand for perfect Bayesian equilibria that satisfy the Intuitive Criterion. To be precise, denote by $x(s)$ and $v(x(s))$ the action a perfect Bayesian equilibrium prescribes E respectively M to take after observing s respectively x . Let $U_E(v, x | s)$ be the expected payoff to E , conditional on s , when he plays x and M plays v . Thus, $U_E(v(x(s)), x(s) | s)$ is E 's expected equilibrium payoff given s . Then:

Definition 1 *A PBE is a perfect Bayesian equilibrium in which M 's belief after observing some \hat{x} which E does not play in equilibrium, i.e., $x(s) \neq \hat{x}$ for all $s \in [a, b]$, is $\mu(s|\hat{x}) = 0$ for*

our results would also hold for a more general cost function $c = c(x, s)$ satisfying $c_x(\cdot) > 0$ and $c_s(\cdot) < 0$ if $x > s$, and $c_x(\cdot) < 0$ and $c_s(\cdot) > 0$ if $x < s$.

all s such that $U_E(1, \hat{x} | s) < U_E(v(x(s)), x(s) | s)$.

There will typically exist a continuum of PBE with the same equilibrium structure. We follow Grossman and Perry (1986) and further narrow our focus on Perfect Sequential Equilibria (PSE). We adopt the definition of PSE from Hörner and Sahuguet (2007) and adjust it to our model:¹⁴

Definition 2 *A PSE is a PBE in which M's relative likelihood assessment after observing some \hat{x} which E does not play in equilibrium satisfies $\frac{\mu(s_1|\hat{x})}{\mu(s_0|\hat{x})} = \frac{f(s_1)}{f(s_0)}$ if $U_E(1, \hat{x} | s_i) \geq U_E(v(x(s_i)), x(s_i) | s_i)$ for $i = 0, 1$ and $\frac{\mu(s_1|\hat{x})}{\mu(s_0|\hat{x})} \geq \frac{f(s_1)}{f(s_0)}$ if $U_E(1, \hat{x} | s_1) > U_E(v(x(s_1)), x(s_1) | s_1)$ and $U_E(1, \hat{x} | s_0) \geq U_E(v(x(s_0)), x(s_0) | s_0)$.*

Notice that the PSE concept agrees with the Intuitive Criterion to put zero probability on states at which E could not possibly benefit from a deviation, but it adds the requirement that after receiving an off-equilibrium report \hat{x} , M should put the same probability on the possibility that E has deviated at any state s at which the deviation \hat{x} could potentially be profitable for him.

3 Equilibrium Analysis

We now derive the equilibrium behavior, starting by describing some general properties that E's report $x(s)$ must satisfy in any perfect Bayesian equilibrium, independent of the strength of his bias. First, E never sends a distorted report $x \neq s$ that leads to $v(x) = 0$, i.e., to a cancelation of the project with certainty. Second, E never distorts his report x so much that $c(d) > \Phi$.¹⁵ Hence, he always sends a report $x \in [s - \bar{d}, s + \bar{d}]$, where $\bar{d} \equiv c^{-1}(\Phi) > 0$. He also never distorts his report x more than necessary to ensure the project's realization, i.e., $v(x) = 1$. It follows from these observations:

¹⁴The refinement introduced by Grossmann and Perry (1986) is slightly more involved than our definition of PSE. However, because of the binary nature of M's decision, what we define as PSE selects in our model exactly the same equilibria as Grossmann and Perry's refinement.

¹⁵To simplify the exposition assume that if E is indifferent between various reports x , he chooses the one that maximizes $v(x)$.

Lemma 1 *If M plays $v(x) = 0$ for all $x \in [s - \bar{d}, s + \bar{d}]$, then E reports $x(s) = s$. If M plays $v(x) = 1$ for some $x \in [s - \bar{d}, s + \bar{d}]$, but $v(x) = 0$ for all other $x \in [s - \bar{d}, s + \bar{d}]$, then E sends the report $x(s)$ with the lowest distortion d necessary to ensure $v(x(s)) = 1$.*

3.1 Moderate Bias

We now confine our attention to the case in which expert E is moderately biased, i.e.,

$$\Phi < c(b - a), \quad (1)$$

which is equivalent to $\bar{d} < b - a$. M knows that a report x is inconsistent with any state $s \notin [x - \bar{d}, x + \bar{d}]$ and, consequently, that $s > \tilde{s}$ if $x > \tilde{x} + \bar{d} \equiv x_H$, and $s < \tilde{s}$ if $x < \tilde{x} - \bar{d} \equiv x_L$, where $\tilde{x} \equiv \tilde{s}$. To simplify the exposition, we assume $a \leq s_L$ and $b \geq s_H$, where $s_H \equiv x_H$ and $s_L \equiv x_L$, in the remainder of this section.¹⁶ It follows:

Lemma 2 *M plays $v(x) = 1$ for $x > x_H$ and $v(x) = 0$ for $x < x_L$.*

We know from Lemma 2 that M plays either $v(x) = 0$ or $v(x) = 1$ for any report $x \notin [x_L, x_H]$. Assume for the moment that $v(x) \in \{0, 1\}$ also holds for any $x \in [x_L, x_H]$. (Proposition 1 will confirm that with a moderately biased E this holds in any PBE.) Further, denote by x^* the lowest x for which M plays $v(x) = 1$, which satisfies $x^* \in [x_L, x_H]$, and let $s^* \equiv x^*$. Lemmas 1 and 2 then imply:

Corollary 1 *E reports $x(s) = x^*$ for $s \in [s^* - \bar{d}, s^*]$ and $x(s) = s$ for $s < s^* - \bar{d}$.*

This corollary describes E's best reply for all $s \leq s^*$ to M's playing $v(x) = 0$ for $x < x^*$ and $v(x^*) = 1$. In particular, E's strategy is separating for low states $s < s^* - \bar{d}$, but involves pooling for intermediate states $s \in [s^* - \bar{d}, s^*]$.

Given Corollary 1, M is uncertain about the true state s when observing x^* . Assume for the moment that E plays $x(s) > x^*$ for all $s > s^*$. (Again Proposition 1 will confirm

¹⁶These assumptions play no role in Section 3.2. For our results in Section 3.1 to hold, it would be sufficient to assume $a \leq s' - \bar{d}$ and $b \geq s'$, where s' is defined below and where $s' - \bar{d} > s_L$ and $s' < s_H$. If $a > s' - \bar{d}$, the equilibrium would be similar to the one in the first part of Section 3.2 despite the expert being only moderately biased; and if $b < s'$, the equilibrium would be similar to the one in the second part of Section 3.2.

that this holds in any PBE with a moderately biased E.) E then knows that the report x^* is consistent with the continuum of states $s \in [s^* - \bar{d}, s^*]$. Her updated belief, using Bayes' rule, is thus $\mu(s|x^*) = \frac{f(s)}{F(s^*) - F(s^* - \bar{d})}$ for $s \in [s^* - \bar{d}, s^*]$ and $\mu(s|x^*) = 0$ for any other s , where $F(y) \equiv \int_a^y f(s)ds$. Consequently, her expected net utility from realizing the project is

$$E_\mu(W(s)|x^*) = \frac{1}{F(s^*) - F(s^* - \bar{d})} \int_{s^* - \bar{d}}^{s^*} W(s)f(s)ds, \quad (2)$$

which increases in s^* as $W_s(s) > 0$. Since $x^* = s^*$, $E_\mu(W(s)|x^*)$ must increase in x^* as well. Define s' by $\int_{s' - \bar{d}}^{s'} W(s)f(s)ds = 0$ and let $x' \equiv s'$. Note that $s' \in (\tilde{s}, s_H)$ and, consequently, $s' - \bar{d} \in (s_L, \tilde{s})$. The definition of x' and the fact that $E_\mu(W(s)|x^*)$ increases in x^* imply that realizing the project when observing x^* , and inferring that $s \in [s^* - \bar{d}, s^*]$, maximizes M's expected utility if and only if $x^* \geq x'$. This leads to our first main result:

Proposition 1 *There exists a continuum of PBE with the same structure, and no other PBE. Each PBE is characterized by a $x^* \in [x', x_H]$ such that E sends a distorted report $x(s) = x^*$ if $s \in [s^* - \bar{d}, s^*]$, and an undistorted report $x(s) = s$ otherwise; M plays $v(x) = 1$ if $x \geq x^*$, and $v(x) = 0$ for any other x on the equilibrium path. On the equilibrium path, $v(x(s)) = 1$ for $s \geq s^* - \bar{d}$, and $v(x(s)) = 0$ otherwise. For all x off the equilibrium path, $v(x)$ must be sufficiently small (e.g., $v(x) = 0$).*

The unique PSE corresponds to the PBE with $x^ = x'$ and $v(x) = 0$ for all x off the equilibrium path.*

Figure 1 illustrates the PSE described in Proposition 1. Thick lines depict E's equilibrium reporting strategy $x(s)$. M's equilibrium strategy $v(x)$ is shown along the vertical axis. The figure illustrates why E has no incentive to deviate at any s . If $s < s' - \bar{d}$, he cannot change M's decision to his benefit with a distortion $d \leq \bar{d}$; and if $s \geq s' - \bar{d}$, he cannot ensure the project's realization with any lower d . M can easily infer s for $x < x' - \bar{d}$ and $x > x'$. Her decision is then the same as she would take if she could observe s herself. After observing x' M is by definition of x' as well off when realizing the project as when canceling it (while in any other PBE she would be strictly better off realizing the project after making the pooling observation $x^* > x'$). For off equilibrium observations $x \in [x' - \bar{d}, x')$, the PSE concept

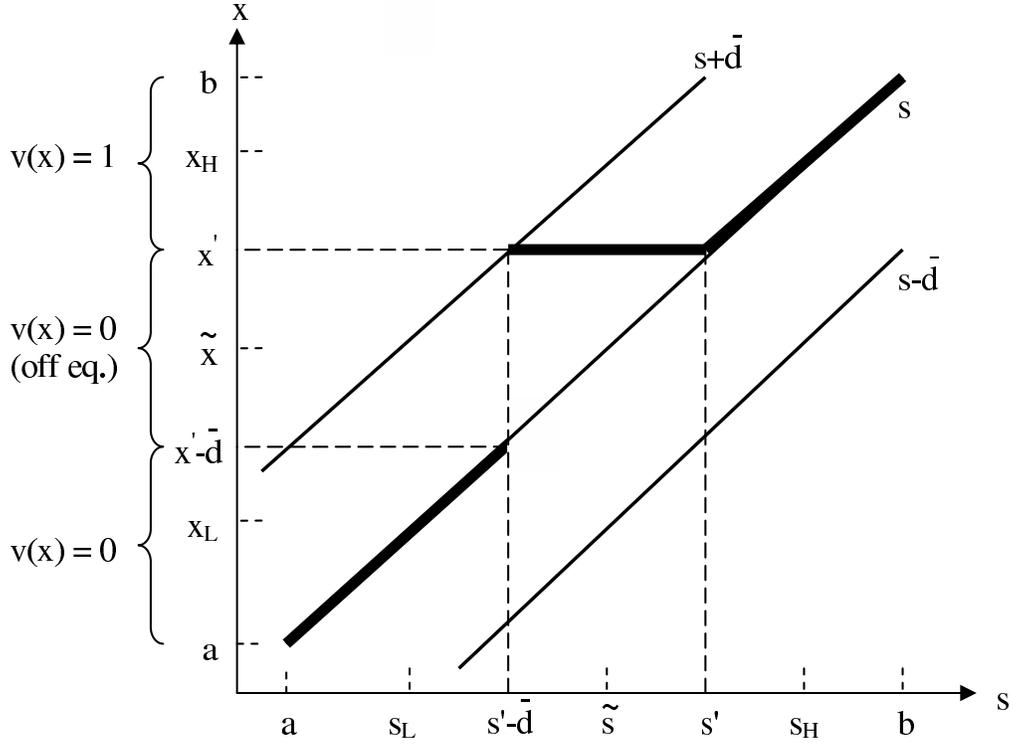


Figure 1: The PBE Structure with Moderate Bias

requires M to cancel the project with probability one, as such deviations could potentially benefit E at a continuum of states starting below $s' - \bar{d}$ and ending below s' .

Proposition 1 implies that in any PBE the project is realized for all states $s \geq s^* - \bar{d}$ while it would only be realized for the states $s \geq \tilde{s}$ if M could observe state s herself. Since $s^* - \bar{d} < \tilde{s}$ if $s^* < s_H$, the project is realized for more states than with symmetric information in the whole continuum of PBE, except at the endpoint when $s^* = s_H$ and $x^* = x_H$. With ex ante probability $F(\tilde{s}) - F(s^* - \bar{d}) \geq 0$, the project is thus only realized due to E's possibility to distort his report $x(s)$ while it would not be realized if the report were sent by an equally well informed, but unbiased expert, or, equivalently, if M could directly observe the state s . When observing x^* and realizing the project, M is aware that she would be better off canceling the project with probability $\frac{F(\tilde{s}) - F(s^* - \bar{d})}{F(s^*) - F(s^* - \bar{d})} \geq 0$. But she does not do so because the state s may have been high indeed, as E also reports $x(s) = x^*$ for $s \in [\tilde{s}, s^*]$, and because he manipulates her beliefs about s in such a way that she is weakly better off realizing the project. Proposition 1 further shows that the unique PSE is the PBE which is best for E, as

it leads to the realization of the project for all $s \geq s' - \bar{d}$. When observing $x = x'$, M realizes the project in this unique PSE with probability one even though she is indifferent and we did not impose any tie-breaking restriction on her strategy.

We next study how changes in E's bias and his lying costs affects our results. A higher rent Φ leads to an increase in \bar{d} but a decrease in $s' - \bar{d}$, such that the ex ante probability that the project is realized increases (assuming that condition (1) still holds). Higher lying costs $c(d)$ have the opposite effect: they lower \bar{d} and the ex ante probability that the project is realized. A tougher punishment for misreporting state s , which increases the costs $c(d)$, therefore makes E worse off, but M better off. These comparative static results are in line with much of the economic literature on crime and punishment following Becker (1968).

3.2 Strong Bias

We now focus on the case in which expert E is strongly biased, i.e.,

$$\Phi \geq c(b - a) \tag{3}$$

or, equivalently, $\bar{d} \geq b - a$. In this case, E is willing to make any distortion necessary to get a project realized for sure that would be cancelled for sure if he reported truthfully. A priori, it may thus seem as if decision maker M should not pay any attention to E's report x . In this case, E would however truthfully report $x(s) = s$ for all s . But if E reports $x(s) = s$ for all s , M would want to pay close attention to x .

We start by describing some general properties that any PBE must satisfy when E is strongly biased. If M plays $v(x) = 1$ for some $x \in [a, b]$, then E would never send a report $x(s)$ for any s that leads to $v(x(s)) = 0$. Therefore:

Lemma 3 *There is no PBE in which E's reporting strategy $x(s)$ is fully separating.*

Consequently, there must again exist a pooling range, i.e., a range of states s for which E sends the same report $x(s)$. However, in contrast to the case when E is only moderately biased, the equilibrium structure depends now crucially on how M would decide in the absence of any information about s . We know that M would realize the project in this case if and

only if

$$\int_a^b W(s)f(s)ds \geq 0. \quad (4)$$

We say that M's prior is favorable (for the project's realization and, hence, for E) if condition (4) holds, and unfavorable otherwise. Below we derive the PBE and PSE when E is strongly biased and M's prior favorable or unfavorable, respectively.

Strong Bias and Favorable Prior

We define the threshold s'' by

$$\int_a^{s''} W(s)f(s)ds = 0, \quad (5)$$

and $x'' \equiv s''$. M is indifferent between realizing or canceling the project when getting a report x that E sends for all $s \leq s''$, but strictly prefers to realize the project when getting a report x that E sends for all $s \leq s^{**}$, where s^{**} is an arbitrary state greater than s'' . Observe that $s'' > \tilde{s}$. Thus, M also realizes the project when getting a report x that E only sends for some $s > s''$.

Proposition 2 *There exists a continuum of PBE with the same structure, and no other PBE. Each PBE is characterized by a $x^{**} \geq x''$ such that E sends a distorted report $x(s) = x^{**}$ for all $s < s^{**}$, and $x(s) = s$ for all $s \geq s^{**}$; and M plays $v(s) = 1$ for all $x \geq x^{**}$. On the equilibrium path, $v(x(s)) = 1$ for all s . For all x off the equilibrium path, $v(x)$ is sufficiently small.*

*The unique PSE corresponds to the PBE with $x^{**} = x''$ and $v(x) = 0$ for all x off the equilibrium path.*

According to Proposition 2, the project is always realized if E is strongly biased and M's prior is favorable. The reason is that E again misreports states around \tilde{s} (in states $s \in [a, s^{**})$ to be precise) in such a way that M (weakly) prefers to realize the project when observing the report x^{**} even though she is aware that an unbiased expert would advice against the project's realization with probability $\frac{F(\tilde{s})}{F(s^{**})}$.

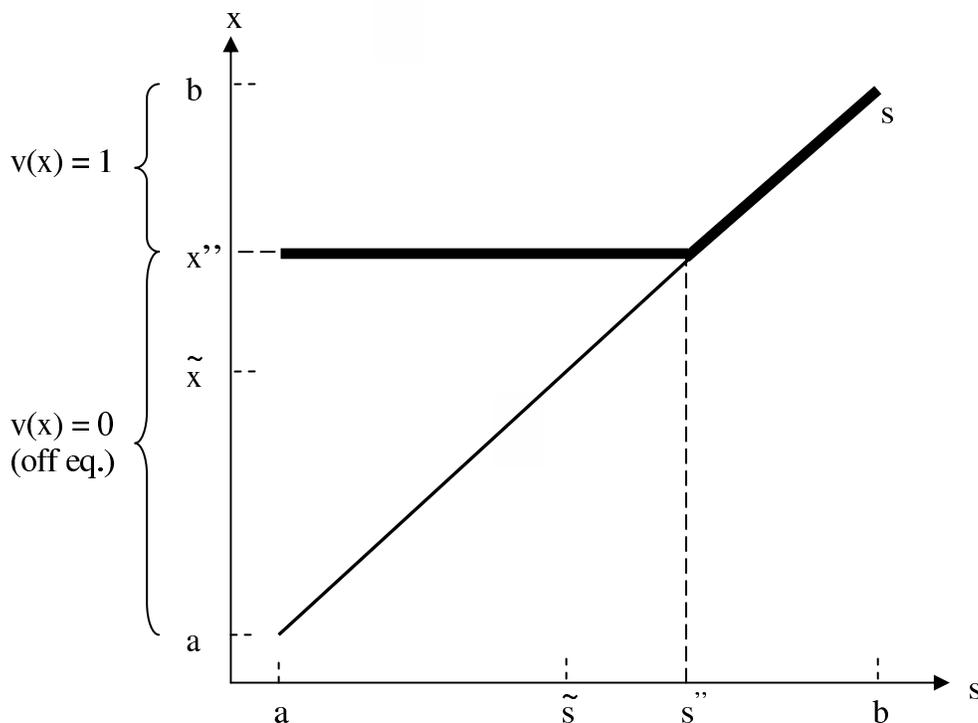


Figure 2: PSE with a Strong Bias and a Favorable Prior

Figure 2 illustrates the unique PSE. This is the PBE in which E distorts his reports $x(s)$ least. That is, E cannot only ensure the project's realization, but he can do so without distorting his report x by more than $x'' - a$ even though he would be willing to distort his report much more. In fact, he distorts his report $x(s)$ so little that M is again indifferent between realizing or canceling the project when observing the report x'' which E sends for a whole range of states s . M's equilibrium strategy is however to realize the project when observing x'' . The unique PSE is thus again the PBE that is best for E.

It is interesting to note that changes in the rent Φ and the distortion costs $c(d)$ have no effect on the equilibrium behavior as long as condition (3) still holds. Assuming that costs $c(d)$ are deadweight loss, M's utility is not affected by small changes in the costs, while E is strictly better off with lower costs. Welfare maximization would thus require reducing these costs, e.g., by reducing punishment, at least if they cannot be made sufficiently high that condition (3) no longer holds.

As discussed in the introduction, our model differs from Crawford and Sobel's (1982) in two respects: First, the sender E's talk is not cheap but costly if he lies. Second, the receiver M's decision is binary rather than continuous. Letting $\Phi \rightarrow \infty$, lying costs $c(d)$ become negligible relative to Φ , and our model approaches a model with cheap talk and a binary decision of the receiver. If M's prior is favorable, it follows from the above that E's equilibrium reporting is the same for any $\Phi \geq c(b-a)$ and that M always realizes the project even in the cheap talk limit of our game.

Strong Bias and Unfavorable Prior

We now analyze the case in which E is strongly biased and M's prior is unfavorable, i.e., in which condition (3) holds but condition (4) does not. For that purpose, we define the threshold s''' by

$$\int_{s'''}^b W(s)f(s)ds = 0, \quad (6)$$

and $x''' \equiv s'''$. M is indifferent between realizing or canceling the project when getting a report x that E sends for all $s \geq s'''$. Because $s''' < \tilde{s}$ holds, M cancels the project when getting a report x that E only sends for some $s < s'''$.

Proposition 3 *There exists a continuum of PBE with the same structure, and no other PBE. Each PBE is characterized by a $x^{***} \in [x''', b]$ such that E sends a distorted report $x(s) = x^{***}$ for all $s \geq s'''$, and $x(s) = s$ for all $s < s'''$; and M plays $v(x) = \theta \equiv \frac{c(x^{***}-s''')}{\Phi} < 1$ for $x = x^{***}$, and $v(x) = 0$ for $x < x'''$. On the equilibrium path $v(x(s)) = 0$ for $s < s'''$ and $v(x(s)) = \theta$ otherwise. For all x off the equilibrium path, $v(x)$ is sufficiently small.*

*There exists a PSE that corresponds to the PBE with $x^{***} = b$ and $v(x) = 0$ for all x off the equilibrium path.*

Proposition 3 implies that the project will be realized with strictly positive probability even if M's prior is unfavorable and despite the fact that M knows that E is willing to distort his report x severely. Note that because M mixes in equilibrium for all states $s \geq s'''$ and because $s''' < \tilde{s}$, it is ambiguous whether the project is realized with a larger ex ante probability than if M could observe s herself. Since θ is increasing in x^{***} , the PSE with

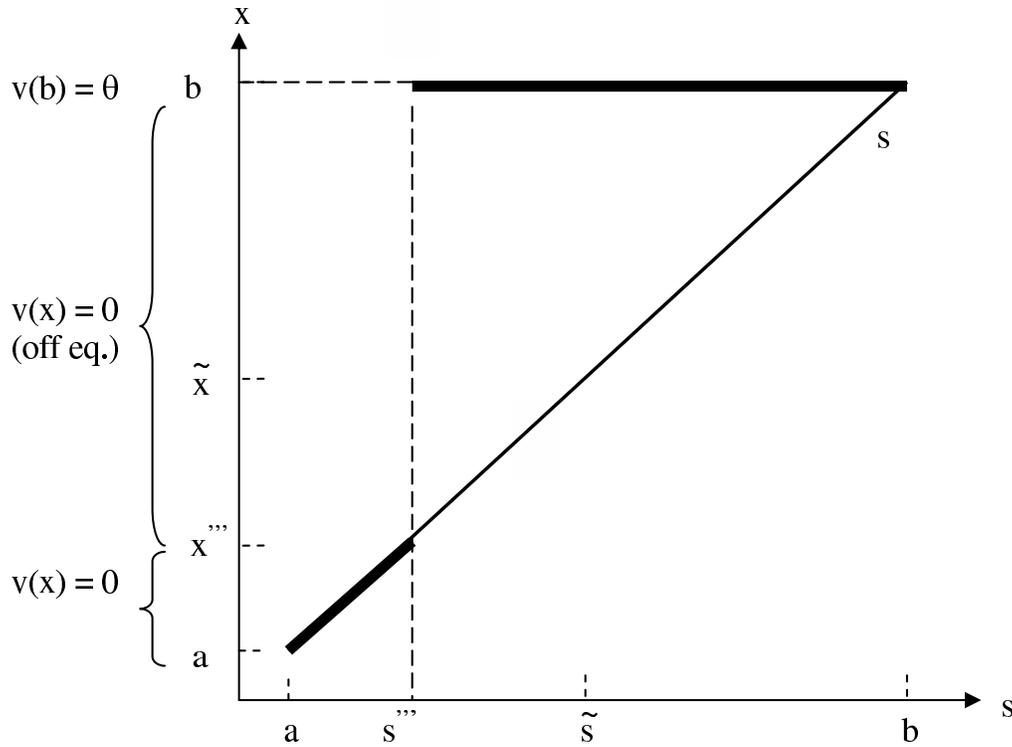


Figure 3: PSE with a Strong Bias and an Unfavorable Prior

$x^{***} = b$ is the PBE associated with the highest probability θ that the project is realized when the state is $s \geq s'''$.¹⁷ Figure 3 provides an illustration.

It is interesting to contrast Propositions 2 and 3. If E is strongly biased and the prior is favorable, his equilibrium reporting strategy $x(s)$ is such that the project is always realized for any state s ; but if he is strongly biased and the prior is unfavorable, his equilibrium reporting is such that the project is not always canceled. Therefore, E cannot only manipulate M's decision to his own advantage if he is known to be moderately biased, but even if he is known to be strongly biased and willing to severely misreport the true state of the world.

Observe also that E's equilibrium behavior is independent of his rent Φ and his lying costs $c(d)$ if he is strongly biased and if M's prior is unfavorable. The probability θ that M realizes the project when observing x^{***} , i.e., when $s \geq s'''$, however decreases in Φ and increases in

¹⁷Other PSE may exist. It follows from Proposition 3 that these PSE would have a very similar structure, but involve pooling at lower x . After the proof of Proposition 3 in Appendix A we briefly explain why we cannot determine in general whether or not other PSE exist in this model with strong bias and an unfavorable prior.

$c(x^{***} - s''')$. A tougher punishment for misreporting states thus harms E by increasing his expected lying costs $c(d)$, but it also has the consequence of increasing the ex ante probability that the project is realized.

Finally, consider again the cheap talk limit of our game. It follows from above that $\theta \rightarrow 0$ as $\Phi \rightarrow \infty$. Hence, as talk becomes arbitrary cheap relative to Φ , no decision relevant information can be transmitted in equilibrium and M does not realize the project, i.e., she acts as if E had not even sent a report.

4 Extensions

In this section, we discuss four extensions of our model.¹⁸ While confirming the general pattern discussed in the previous section, these extensions also provide genuinely new insights.

4.1 Naive Decision Maker

An interesting question is to what extent the results derived so far are robust when the decision maker may not be perfectly rational in the sense that she may not fully understand the model in which she acts. This question is relevant as there are many instances in which experts cannot be sure whether the decision maker is savvy and, despite being incompletely informed, understands the model, or whether she is naive and will take his report by its face value. To capture this we extend our model and assume that M can be of two types. With probability $\gamma \in [0, 1)$ she is naive and interprets any report x as truthful, therefore realizing the project if and only if $x \geq \tilde{x}$.¹⁹ With probability $1 - \gamma$, M is the sophisticated type studied hitherto. These types and the probability γ are common knowledge. To guarantee a single crossing condition, we now also assume $c(d)$ to be weakly convex.

This extended model leads to the following insights: First, the results derived in Section 3 are robust insofar as in each case (moderate bias, strong bias and favorable prior, and strong bias and unfavorable prior) there exists a unique threshold $\hat{\gamma} \in (0, 1)$ such that for $\gamma \leq \hat{\gamma}$ the PSE described in Propositions 1-3 are still PSE.

¹⁸Appendix B (not intended for publication) contains the corresponding derivations, propositions and proofs.

¹⁹With $\gamma = 1$, E would simply report $x(s) = \tilde{x}$ for $s \in (\tilde{s} - \bar{d}, \tilde{s}]$ and $x(s) = s$ for all other s .

Second, for $\gamma \geq \hat{\gamma}$ there exists a PSE with two pooling ranges. In relatively low states s , E sends the pooling report \tilde{x} , which is aimed at the naive type of M and hence gets the project realized with probability γ . In higher states s , E sends the pooling report $x_2 > \tilde{x}$, which makes the sophisticated type of M indifferent. For all other s (if such s exist), E sends a truthful report. When E is moderately biased, or when he is strongly biased and M's prior favorable, the project is always realized upon observations $x \geq x_2$.²⁰

Third, the possibility that she may be thought of as naive can improve the welfare of the sophisticated type of M. In particular, when E's bias is strong and M's prior is favorable, the sophisticated type of M is strictly better off when $\gamma > \hat{\gamma}$ than when $\gamma < \hat{\gamma}$. The main reason is that the pooling report \tilde{x} reveals very low states z when $\gamma > \hat{\gamma}$, so that she can make an informed decision in these states, whereas she always realizes the project when $\gamma < \hat{\gamma}$.

4.2 Multiple Like Biased Experts

In some relevant instances an expert enjoys a natural monopoly, e.g., if he is the informed head of an organization. In other important applications, however, there are naturally multiple experts who may differ with respect to the strength of their biases. In the jargon of Krishna and Morgan (2001), these experts are thus “like biased”. For example, when a decision maker decides on a big public infrastructure project, typically several industry experts are consulted who may all benefit from the “project” being realized, but to different extents.

To account for that, consider $N \geq 2$ experts who may differ with respect to their biases $\Phi_i > 0$, where $i = 1, \dots, N$. Without loss of generality we label experts in increasing order of their biases, i.e., $\Phi_1 \leq \dots \leq \Phi_N$. These biases are common knowledge, and all experts have the same costs of lying $c(d)$ and observe state s .

When all experts submit their reports simultaneously there exists a PSE in which each expert reports as if he were the only expert and his bias were Φ_1 , and M decides as she would if there were only one expert whose bias was Φ_1 . There also exists a fully separating PSE in which each expert reports truthfully simply because if every one else is saying the truth no

²⁰In these two cases, the two PSE exist under mutually exclusive conditions. With a strong bias and an unfavorable prior, the sophisticated type of M may play a mixed strategy upon making the pooling observation x_2 .

expert can gain by unilaterally deviating. There are no other PSE.

Now consider as in Krishna and Morgan (2001) a setting where experts report sequentially and in an arbitrary order, and where every expert can observe the reports of his predecessors. In this case the first equilibrium of the simultaneous reporting case still persists, but the fully separating PSE ceases to exist, and there are no other PSE.

These findings are in line with the results of the basic model. The implication that the decision maker fares as well with N experts as with a single expert if the single expert is the one with the smallest bias is consistent with Krishna and Morgan's (2001) finding that the most informative equilibrium with two like biased experts is never more informative than the most informative equilibrium with a single expert. However, there is a subtle, but important difference between our result and theirs: In Krishna and Morgan's cheap talk model the decision maker may be worse off with two experts than with a single expert, whereas having more experts can never make her worse off in our model. Quite to the contrary, ever more information gets transmitted in equilibrium as the number of experts rises if having more experts means that the minimal bias of all experts decreases.

4.3 Uncertain Size of Expert's Bias

Decision makers often face some uncertainty about how biased the expert is. Therefore consider a single expert E whose bias Φ is not known by M . For simplicity, assume that Φ can be of two types only. It is common knowledge that Φ takes the value Φ_1 with probability α and the value Φ_2 with probability $(1 - \alpha)$ where $0 < \Phi_1 < \Phi_2$. Call an expert with bias Φ_1 (Φ_2) weakly (strongly) biased.

The general structure of the PSE of the basic model remains the same under uncertainty about the bias in the following sense: If both (or all) types are strongly biased and the prior is favorable, the PSE strategies of all types are identical to the equilibrium strategy of the single, strongly biased expert of the basic model. Also for strong bias and unfavorable prior and for moderate bias the structure of the basic model with pooling and separating segments remains the same. The only difference is that now the thresholds at the beginning and end of the pooling segments are different. In particular, type 1 gets the project passed less frequently

than if M knew his bias, while type 2 gets the project passed more frequently. Consequently, the expert tends to benefit if his bias is underestimated, but to be worse off if his bias is overestimated.

4.4 The Right to Remain Silent

There exist both situations in which experts are obliged to send a report, and situations in which they have the choice to withhold a report, i.e., to say “no comment”. While we have so far focused on the former situations, we now briefly discuss the latter. Consider again a single expert with known bias. Clearly, E has incentives to avoid lying costs. Thus, whenever M’s prior is favorable (both for moderate and for strong bias), there is an equilibrium where E never sends a report and M realizes the project for sure. There are also other equilibria in which E sends a truthful report $x(s) = s$ for some but not all s . In all of these equilibria, though, the project is realized for all s .

If the bias is moderate and the prior unfavorable, E is indifferent between sending $x(s) = s$ and sending no report whenever $s < s' - \bar{d}$ (where $v(x) = 0$ in both cases), while for $s \geq s' - \bar{d}$ he sends the same report and M selects the same actions as in Proposition 1. If E’s bias is strong and the prior unfavorable, E is indifferent between sending $x(s) = s$ and sending no report whenever $s < s'''$ (where $v(x) = 0$ in both cases), while for $s \geq s'''$ he sends the same report and M selects the same actions as in Proposition 3. Hence, in contrast to the case with a favorable prior, it cannot occur that E remains silent for all s when M’s prior is unfavorable.

5 Concluding Remarks

We have studied the behavior of biased experts in a setting with a continuous state space and a binary choice of the decision maker. Further, in our model lying entails some cost that is increasing in the extent of distortion. This framework is suitable to analyze various real-world problems such as expert advice and testimony related to public policies or to court decisions, CEO and manager behavior given binary Board or investor decisions, as well as incumbent government behavior before reelections.

We show that three kinds of biased experts can be distinguished according to equilibrium play: Moderately biased, strongly biased with favorable prior of the decision maker, and strongly biased with unfavorable prior of the decision maker. Interestingly, in equilibrium in all three cases the strategy of the expert includes separating and pooling segments. Further, in all three cases the fact that the expert has private information about the state of the world and is able to send a report to the decision maker increases the set of parameter values for which the decision maker may realize the project. In extensions we show that these results also hold when the decision maker may be naive, when there are multiple like biased experts, when the decision maker faces some uncertainty about the size of a single expert's bias, and when the expert has the possibility to remain silent.

References

- [1] Adams, Renée, and Daniel Ferreira, 2007, A Theory of Friendly Boards, *Journal of Finance* LXII: 217-250.
- [2] Banks, Jeffrey S., 1990, A Model of Electoral Competition with Incomplete Information, *Journal of Economic Theory* 50: 309-325.
- [3] Becker, Gary, 1968, Crime and Punishment: An Economic Approach, *Journal of Political Economy* 76: 169-217.
- [4] Ber, Hedva, Yishay Yafeh, and Oved Yosha, 2001, Conflict of interest in universal banking: Bank lending, stock underwriting, and fund management, *Journal of Monetary Economics* 47: 189-218.
- [5] Callander, Steven, and Simon Wilkie, 2007, Lies, Damned Lies and Political Campaigns, *Games and Economic Behavior* 60: 262-286.
- [6] Cho, In-Koo, and David M. Kreps, 1987, Signaling Games and Stable Equilibria, *Quarterly Journal of Economics* 102: 179-221.

- [7] Cooter, Robert, and Winand Emons, 2003, Truth-Revealing Mechanisms for Courts, *Journal of Institutional and Theoretical Economics* 159: 259-279.
- [8] Cohen, Deborah, and Philip Carter, 2010, Conflicts of Interest: WHO and the pandemic flu “conspiracies”, *British Medical Journal* 340: 1274-1279.
- [9] Cummins, Jason G., and Ingmar Nyman, 2005, The Dark Side of Competitive Pressure, *RAND Journal of Economics* 36: 361-377.
- [10] Crawford, Vincent, and Joel Sobel, 1982, Strategic Information Transmission, *Econometrica* 50: 1431-1451.
- [11] Crocker, Keith J., and John Morgan, 1998, Is Honesty the Best Policy? Curtailing Insurance Fraud through Optimal Incentive Contracts, *Journal of Political Economy* 106: 335-375.
- [12] Edmond, Chris, 2010, Information Manipulation, Coordination, and Regime Change, *Working Paper, University of Melbourne*.
- [13] Emons, Winand, and Claude Fluet, 2009, Accuracy Versus Falsification Costs: The Optimal Amount of Evidence under Different Procedures, *Journal of Law, Economics and Organization* 25: 134-156.
- [14] Fischer, Paul, and Robert Verrecchia, 2000, Reporting Bias, *Accounting Review* 75: 229-245.
- [15] Goldman, Eitan, and Steve Slezak, 2006, An equilibrium model of incentive contracts in the presence of information manipulation, *Journal of Financial Economics* 80: 603-626.
- [16] Grossman, Sanford J., and Motty Perry, 1986, Perfect Sequential Equilibrium, *Journal of Economic Theory* 39, 97-119.
- [17] Hodler, Roland, Simon Loertscher, and Dominic Rohner, 2010, Inefficient Policies and Incumbency Advantage, *Journal of Public Economics*, forthcoming.
- [18] Hörner, Johannes, and Nicolas Sahuguet, 2007, Costly Signalling in Auctions, *Review of Economic Studies* 74: 173-206.

- [19] Kartik, Navin, 2009, Strategic Communication with Lying Costs, *Review of Economic Studies* 76: 1359-1395.
- [20] Kartik, Navin, Marco Ottaviani, and Francesco Squintani, 2007, Credulity, lies and costly talk, *Journal of Economic Theory* 134: 93-116.
- [21] Krishna, Vijay, and John Morgan, 2001, A Model of Expertise, *Quarterly Journal of Economics* 116: 747-775.
- [22] Lacker, Jeffrey M., and John A. Weinberg, 1989, Optimal Contracts under Costly State Falsification, *Journal of Political Economy* 97: 1345-1363.
- [23] Maggi, Giovanni, and Andres Rodriguez-Clare, 1995, Costly Distortion of Information in Agency Problems, *RAND Journal of Economics* 26: 675-689.
- [24] Melumad, Nahum D., and Toshiyuki Shibano, 1991, Communication in settings with no transfers, *RAND Journal of Economics* 22: 173-198.
- [25] Michaely, Roni, and Kent Womack, 1999, Conflict of Interest and the Credibility of Underwriter Analyst Recommendations, *Review of Financial Studies* 12: 653-686.
- [26] Morris, Stephen, 2001, Political Correctness, *Journal of Political Economy* 109: 231-65.
- [27] Rogoff, Kenneth, and Anne Sibert, 1988, Elections and Macroeconomic Policy Cycles, *Review of Economic Studies* 55: 1-16.
- [28] Posner, Richard, 1999, The Law and Economics of the Economic Expert Witness, *Journal of Economic Perspectives* 13: 91-99.
- [29] Sanchirico, Chris, 2001, Relying on the Information of Interested - and Potentially Dishonest - Parties, *American Law and Economics Review* 3: 320-357.
- [30] Swank, Otto, Wilko Letterie, and Hendrik van Dalen, 1999, A Theory of Policy Advice, *Journal of Law, Economics and Organization* 15: 602-614.

- [31] van Winden, Frans, 2004, Interest Group Behavior and Influence, in Charles Rowley and Friedrich Schneider (eds.), *Encyclopedia of Public Choice*, Boston: Kluwer Academic Publishers, pp. 118-128.
- [32] Wrasai, Phongthorn, and Otto Swank, 2007, Policy makers, advisers, and reputation, *Journal of Economic Behavior and Organization* 62: 579-590.

Appendix

A Proofs

Proof of Proposition 1: We proceed as follows: First we show that the strategy profile and the associated beliefs constitute PBE. Second we prove that no other PBE exist. Third we show that any PBE with $x^* > x'$ is not a PSE. Fourth we show that the PBE with $x^* = x'$ is a PSE.

It follows from Lemma 1 that E's strategy is optimal given M's strategy. M's strategy is optimal if and only if her beliefs are such that $E_\mu(W(s)|x) \leq 0$ for $x < x^*$ and $E_\mu(W(s)|x) \geq 0$ for $x \geq x^*$. It remains to be shown that M's beliefs on equilibrium are updated according to Bayes' rule and consistent with E's strategy, and that her beliefs off equilibrium satisfy the Intuitive Criterion. Since M's beliefs must be $\mu(s|x) = 1$ for $s = x$ if $x < x^* - \bar{d}$ or $x > x^*$, it follows that $E_\mu(W(s)|x) < 0$ for $x < x^* - \bar{d} \leq \tilde{x}$ and $E_\mu(W(s)|x) > 0$ for $x > x^* > \tilde{x}$. Lemma 1, equation (2) and the definition of x' further imply that M's beliefs are such that $E_\mu(W(s)|x^*) \geq 0$ since $x^* \geq x'$. Thus, given E's strategy, M has correct and consistent beliefs on equilibrium. M's beliefs off equilibrium must be such that $E_\mu(W(s)|x) \leq 0$ for $x \in [x^* - \bar{d}, x^*)$. Note that if $s = s^* - \bar{d} \leq \tilde{s}$, E would be better off by playing any $\hat{x} \in [x^* - \bar{d}, x^*)$ when followed by $v(\hat{x}) = 1$ than by playing his equilibrium strategy $x(s) = x^*$. Off equilibrium beliefs $\mu(s|\hat{x})$ that imply $E_\mu(W(s)|\hat{x}) \leq 0$ are thus consistent with the Intuitive Criterion.

We now prove that no other PBE exist in which $v(x) \in \{0, 1\}$ for all x . First, note that there can be no PBE with the structure described in Corollary 1 and with $x^* < x'$. To see this, suppose to the contrary that $x^* < x'$. Corollary 1 then requires M to play $v(x^*) = 1$, but the definition of x' implies $E_\mu(W(s)|x^*) < 0$. Therefore, M would want to deviate and play

$v(x^*) = 0$. Second, note that there can be no PBE in which $v(\hat{x}) = 0$ for any $\hat{x} > x^*$ given $x^* \geq x'$. It follows from the definition of x' and Lemma 1 that if $x(s) = \hat{x}$ is played for some s , then $E_\mu(W(s)|\hat{x}) > 0$ must hold and, hence, $v(\hat{x}) = 1$. If E does not play $x(s) = \hat{x} > x^*$ for any s , then M's off-equilibrium beliefs must satisfy $E_\mu(W(s)|\hat{x}) \leq 0$ to sustain a PBE in which $v(\hat{x}) = 0$. However E would be strictly worse off by playing $x(s) = \hat{x}$ (rather than his equilibrium strategy) for all $s < s^*$ and, hence, for all $s \leq \tilde{s}$, independently of $v(\hat{x})$, but better off by playing $x(s) = \hat{x}$ for some $s > s^*$, e.g., $s = \hat{s} \equiv \hat{x}$, if M played $v(\hat{x}) = 1$ (rather than $v(\hat{x}) = 0$). The Intuitive Criterion thus requires that when observing $\hat{x} > x^*$, M should put zero probability on the possibility that E has deviated at any $s \leq \tilde{s}$. Hence, her off-equilibrium beliefs must be such that $E_\mu(W(s)|\hat{x}) > 0$. Thus, there cannot exist any PBE in which M plays $v(\hat{x}) = 0$ for some $\hat{x} > x^*$. As a consequence, E plays $x(s) = s$ for all $s \geq s^*$ in any PBE.

We next prove that no PBE exists in which $v(x) \in (0, 1)$ for some x on the equilibrium path, i.e., in which M plays a non-degenerate mixed strategy for some x on equilibrium. M would only be willing to play $v(\hat{x}) \in (0, 1)$ for some \hat{x} if $E_\mu(W(s)|\hat{x}) = 0$. If $x(s) = s$ for all s and $v(x(\tilde{s})) < 1$, then E would deviate at $s = \tilde{s}$ to play $x(\tilde{s} + \epsilon)$, where ϵ is arbitrarily small. Hence, whenever $v(\hat{x}) \in (0, 1)$ for some \hat{x} , then it must hold that $x(s) = \hat{x}$ for a range of states $s \in [s_a, s_b]$ satisfying $E_\mu(W(s)|\hat{x}) = 0$, which requires $s_a < \tilde{s} < s_b$. This cannot be part of a PBE if $x(s) = s$ for $s \in (s_b, s_c]$ for any $s_c \in (s_b, b]$ because E would want to deviate if $s = s_b$ (or if s were slightly below s_b) in order to ensure $v(x(s)) = 1$. Hence, it only remains to be shown that there cannot be a PBE with $x(s) = \hat{x}$ for $s \in [s_a, s_b]$ and $x(s) = \hat{\hat{x}} > \hat{x}$ for $s \in (s_b, s_c]$. [For this to be a PBE, it would have to hold that $v(\hat{\hat{x}}) = 1$ (since $s_b > \tilde{s}$) and that $\Phi - c(|\hat{\hat{x}} - s_b|) = v(\hat{x})\Phi - c(|\hat{x} - s_b|)$]. There cannot exist such a PBE with $\hat{\hat{x}} < x_c \equiv s_c$ because E could only benefit from deviating and playing some $x \in (\hat{\hat{x}}, x_c]$ for some $s > \hat{\hat{s}} > \tilde{s}$, where $\hat{\hat{s}} \equiv \hat{\hat{x}}$. The Intuitive Criterion would thus require M to believe $\mu(s|x) = 0$ for all $s \leq \tilde{s}$, and E would thus want to deviate at $s \in (\hat{\hat{s}}, s_c]$. There also cannot exist such a PBE with $\hat{\hat{x}} \geq x_c$ because E could only benefit from deviating and playing some $x = \hat{\hat{x}} - \epsilon$, where ϵ is positive but can be arbitrarily small, if $s \in [s_b - \tilde{\epsilon}, \hat{\hat{s}} - \epsilon/2]$, where $\tilde{\epsilon}$ is small as well. The Intuitive Criterion would thus require M to believe $\mu(s|x) = 0$ for all $s \leq \tilde{s}$, and E would thus

want to deviate at $s \in [s_b - \tilde{\epsilon}, \hat{s} - \epsilon/2]$. This completes the proof that no equilibrium exists in which M plays a non-degenerate mixed strategy for some x on equilibrium.

We now prove that any PBE with $x^* > x'$ is not a PSE. Consider a deviation $\hat{x} = x^* - \epsilon$, where ϵ is positive but can be arbitrarily small. We first need to identify the states in which E could have potentially benefitted from the deviation \hat{x} . The interval $(s^* - \epsilon - \bar{d}, s^* - \epsilon/2)$ contains all the states for which E would be strictly better off by the deviation \hat{x} when followed by $v(\hat{x}) = 1$ than by playing his equilibrium strategy. Second, Bayesian updating by M under the conditions in Definitions 1 and 2 leads to

$$E_\mu^{PSE}(W(s)|\hat{x}) \equiv \frac{1}{F(s^* - \epsilon/2) - F(s^* - \epsilon - \bar{d})} \int_{s^* - \epsilon - \bar{d}}^{s^* - \epsilon/2} W(s)f(s)ds. \quad (7)$$

Letting $\epsilon \rightarrow 0$, we have $E_\mu^{PSE}(W(s)|\hat{x}) \rightarrow E_\mu(W(s)|x^*)$. Now, because $E_\mu(W(s)|x^*) > 0$ for $x^* > x'$, M optimally plays $v(\hat{x}) = 1$ in the continuation game after observing a deviation \hat{x} close enough to x^* . But given that M plays $v(\hat{x}) = 1$, E has indeed an incentive to deviate at all states $s \in (s^* - \epsilon - \bar{d}, s^* - \epsilon/2)$. Consequently, for all PBE with $x^* > x'$, there exist states s and deviations \hat{x} that are profitable for E when followed by a PBE in the continuation game in which M updates her beliefs according to Definitions 1 and 2. Thus, these PBE are not PSE.

We are left to show that the PBE with $x^* = x'$ is a PSE. Applying the same arguments as above, it is straightforward to see that given a deviation $\hat{x} = x' - \epsilon$ and updating on behalf of M according to definitions 1 and 2 leads to $E_\mu^{PSE}(W(s)|\hat{x}) \equiv \frac{1}{F(s' - \epsilon/2) - F(s' - \epsilon - \bar{d})} \int_{s' - \epsilon - \bar{d}}^{s' - \epsilon/2} W(s)f(s)ds < 0$, where the inequality follows from the definition of s' . Hence, it will be sequentially rational for M to play $v(\hat{x}) = 0$. Given this, the deviation will not be profitable for E at any state s . Last, for any strictly smaller deviation \hat{x} , $E_\mu^{PSE}(W(s)|\hat{x})$ must be negative as well (because $s^* - \epsilon - \bar{d}$ and $s^* - \epsilon$, obviously, decrease in ϵ). Hence, the PBE with $x^* = x'$ is a PSE. ■

Proof of Proposition 2: We first show that the strategy profile and the associated beliefs constitute PBE. We then prove that no other PBE exist.

It is straightforward to see that E's strategy is optimal given M's equilibrium strategy,

and that M's strategy is optimal given her beliefs $\mu(s|x)$ are such that $E_\mu(W(s)|x) \geq 0$ if $x \geq x^{**}$, and $E_\mu(W(s)|x) \leq 0$ if $x < x^{**}$. It follows from the definition of s'' and $s'' > \tilde{s}$ that $E_\mu(W(s)|x) \geq 0$ for all $x \geq x^{**} \geq x''$. M's beliefs off equilibrium must be such that $E_\mu(W(s)|x) \leq 0$ for $x < x^{**}$. Note that if $s = a$, E would be better off by playing any $\hat{x} < x^{**}$ when followed by $v(\hat{x}) = 1$ than by playing his equilibrium strategy $x(a) = x^*$. Thus, there are off equilibrium beliefs $\mu(s|\hat{x})$ such that $E_\mu(W(s)|\hat{x}) \leq 0$ which are consistent with the Intuitive Criterion.

We start proving that no other PBE exist by noting that in any PBE $v(x(s)) = 1$ must hold for some s . Otherwise, it would have to hold for all $s > \tilde{s}$ that $x(s)$ is part of a pooling range x_p and that M's beliefs $\mu(s|x)$ are such that she plays $v(x_p) < 1$. But given the favorable prior, these two requirements cannot jointly hold. Hence, $v(x(s)) = 1$ for some s . Together with Lemma 3, this implies that $v(x(s)) = 0$, where both $v(\cdot)$ and $x(s)$ are equilibrium strategies, cannot hold for any s . There must thus be a pooling range with $x(s) = \hat{x}$ for all $s \in [a, s^{**}]$ where $s^{**} \geq s''$.

We next show that there are no other PBE with $v(x(s)) = 1$ for all s . Given that $x(s) = \hat{x}$ is optimal for E if and only if $s \in [a, s^{**}]$, it must hold that $x(s) > \hat{x}$ for $s > s^{**}$. Further, we can use the same strategy as in the proof of Proposition 1 to prove that there are no PBE in which M plays $v(x) < 1$ for some $x > x^{**}$. Hence, it must hold that $\hat{x} = x^{**}$ and that $x(s) = s$ for all $s > s^{**}$.

We also have to show that there are no PBE in which $v(x(s)) \in (0, 1)$ for $s \in [a, s'']$. (We know that there cannot exist a PBE with $v(x(s)) < 1$ for any $s > s''$.) We can again do this using the same strategy as in the proof of Proposition 1.

We can also use the same strategy as in the proof of Proposition 1 to show that any PBE with $x^{**} > x''$ is not a PSE, and that the PBE with $x^{**} = x''$ and $v(x) = 0$ for all x off the equilibrium path is a PSE. ■

Proof of Proposition 3: First we show that the strategy profile and the associated beliefs constitute PBE. Then we prove that no other PBE exist.

Given M's equilibrium strategy, with θ ensuring that E is indifferent between $x(s) = x^{***}$ and $x(s) = s'''$ if $s = s'''$, E's strategy is optimal. M's strategy is optimal if her beliefs $\mu(s|x)$

are such that $E_\mu(W(s)|x^{***}) = 0$ and $E_\mu(W(s)|x) \leq 0$ for $x \neq x^{***}$. It follows from the definition of s''' and $s''' < \tilde{s}$ that $E_\mu(W(s)|x^{***}) = 0$ and that $E_\mu(W(s)|x) < 0$ for $x < s'''$. Note that if $s = a$, E would be better off playing any $\hat{x} \in [x''', x^{***})$ or any $\hat{x} > x^{***}$ when followed by $v(\hat{x}) = 1$ than playing his equilibrium strategy $x(a) = a$. The Intuitive Criterion thus allows for off equilibrium beliefs $\mu(s|\hat{x})$ that imply $E_\mu(W(s)|\hat{x}) \leq 0$.

We next prove that it must hold in any PBE that $v(x(s)) \in (0, 1)$ for $s \geq s'''$ and $v(x(s)) = 0$ otherwise. By an argument symmetric to the one used in the proof of Proposition 2, one can show that in any PBE $v(x(s)) = 0$ for some s must hold. Together with Lemma 3, this implies that $v(x(s)) = 1$ cannot hold for any s in any PBE. Therefore, there must be a pooling range with $x(s) = x^{***}$ for all $s \in [\hat{s}, b]$ with $\hat{s} \leq s'''$. Lemma 1 further implies $x(s) = s$ for all s for which $v(x(s)) = 0$. Therefore, it must hold that $\hat{s} = s'''$ such that $v(x^{***}) \in (0, 1)$ is possible. Since $s''' < \tilde{s}$, Lemma 1 further implies that $x(s) = s$ and, hence, $v(x(s)) = 0$ for all $s < s'''$. In any PBE of this structure, $v(x^{***}) = \theta$ must satisfy $\theta\Phi - c(x^{***} - s''') = 0 \Leftrightarrow \theta = c(x^{***} - s''')/\Phi < 1$.

For the PBE with $x^{***} = b$ to be a PSE, it must hold that M's off equilibrium beliefs $\mu(s|x)$ are such that $E_\mu(W(s)|x) < 0$ for any $x \in [x''', b)$. Note that E could potentially benefit from playing any $\hat{x} \in [x''', b)$ for all s except possibly some s close to b . Hence, PSE requires M's off equilibrium beliefs $\mu(s|\hat{x})$ indeed to be such that $E_\mu(W(s)|\hat{x}) < 0$. ■

Potentially Non-Unique PSE with Unfavorable Prior: We briefly discuss why the existence of other PSE cannot be ruled out when the prior is unfavorable and the bias is strong. For PBE with $x^{***} < b$ to be PSE, it must hold that all (off equilibrium) $\hat{x} \in (x^{***}, b]$ could potentially benefit E for a set of s such that $E_\mu(W(s)|\hat{x}) \leq 0$. This holds for $\hat{x} = x^{***} + \epsilon$ when $\epsilon > 0$ is arbitrarily small, as \hat{x} could then potentially benefit E for all s . Higher \hat{x} , e.g., $\hat{x} = b$, could certainly benefit E if s is close to b or if $s < s'''$ (since $\Phi > c(a - b)$), but possibly not if s is slightly above s''' . Hence, it is unclear whether or not some of the PBE with $x^{***} < b$ are PSE.

B Extensions (not for publication)

B.1 Naive Decision Maker

M is naive with probability $\gamma \in [0, 1)$, and sophisticated with probability $1 - \gamma$. If she is naive, she interprets any report x as truthful and realizes the project if and only if $x \geq \tilde{x}$. All of this is commonly known by E and by M when she is sophisticated. We further assume $c''(d) \geq 0$.

Proposition B.1 *i) In each of the three cases considered in Propositions 1-3, there exists a unique threshold $\hat{\gamma} \in (0, 1)$ such that for $\gamma \leq \hat{\gamma}$ the PSE described in these Propositions are still PSE. There cannot exist any other PSE with only one pooling range.*

ii) In each of these three cases, there exists a PSE with two pooling ranges when $\gamma \geq \hat{\gamma}$. E reports $x(s) = \tilde{x}$ for $s \in [s_0, s_1)$, $x(s) = x_2 > \tilde{x}$ for $s \in [s_1, s_2]$, and $x(s) = s$ for all other s (if such s exist), where $a \leq s_0 < s_1 < s_2 \leq b$ and $s_2 = x_2$. When \tilde{x} is reported, a naive M is indifferent and the project is thus realized with probability γ . When x_2 is reported, a sophisticated M is indifferent and the project is realized with probability 1 if E is moderately biased, or if E is strongly biased and M's prior favorable, but a sophisticated M may play a mixed strategy if E is strongly biased and her prior unfavorable.

Sketch of proof: i) To prove existence of the PSE characterized in Propositions 1-3, we need to show that E does not want to deviate at any s . In case of moderate bias, we need to ensure that E does not want to deviate and play $x(s) = \max\{\tilde{x}, s\}$ at any $s < s'$. It follows from $c''(d) \geq 0$ that E does not have an incentive to deviate at any $s \neq s' - \bar{d}$ if he has no incentive to deviate at $s = s' - \bar{d}$. E has no incentive to deviate at $s = s' - \bar{d}$ if $\Phi - C(\bar{d}) \geq \gamma\Phi - C(\tilde{x} - (s' - \bar{d})) \Leftrightarrow \gamma \leq \hat{\gamma} \equiv \frac{\Phi + C(\tilde{x} - (s' - \bar{d})) - C(\bar{d})}{\Phi} \in (0, 1)$. Hence the PSE characterized in Proposition 1 still exists if and only if $\gamma \leq \hat{\gamma}$.

The proofs for the cases of strong bias are analogous. The value of the threshold $\hat{\gamma}$ below which the PSE characterized in Propositions 2 and 3 still exist is however different, because of the different pooling ranges and the mixed strategy played by E when making the pooling observation in Proposition 3.

To prove that there cannot exist any other PSE with only one pooling range, note first that there cannot exist PSE with a unique pooling report other than x' or \tilde{x} , as there would

always be some s where E could benefit from deviating and playing either $x(s) = s$, x' or \tilde{x} instead. Second, consider a situation with a unique pooling report \tilde{x} . In such a situation E would play $x(s) = \tilde{x}$ for s somewhat below \tilde{s} and $x(s) = s$ for $s > \tilde{s}$. The best response of a sophisticated M is to play $v(x) = 0$ for $x \leq \tilde{x}$ and $v(x) = 1$ for $x > \tilde{x}$. But the best response of E is then to play some x slightly above \tilde{x} for a whole range of $s \leq \tilde{s}$. Hence there cannot exist a PSE with a pooling range where \tilde{x} is reported, but no second pooling range.

ii) In a PSE with two pooling ranges the thresholds s_0 , s_1 and s_2 must satisfy three conditions. One is for a sophisticated M: $\int_{s_1}^{s_2} W(s)f(s)ds = 0$. It implicitly defines s_2 as a function of s_1 , $s_2(s_1)$, which satisfies $ds_2(s_1)/ds_1 < 0$ and $s_2(\tilde{s}) = \tilde{s}$. Since $s_1 < s_2$, it follows that $s_1 < \tilde{s} < s_2$. The other two conditions are for E: First, $\gamma\Phi - C(\tilde{x} - s_0) = 0$ if $s_0 > a$, and $\gamma\Phi - C(\tilde{x} - s_0) \geq 0$ if $s_0 = a$. Second, $\Phi - C(x_2 - s_1) = \gamma\Phi - C(\tilde{x} - s_1)$ if $s_2 < b$, and $[\gamma + (1 - \gamma)\theta]\Phi - C(x_2 - s_1) = \gamma\Phi - C(\tilde{x} - s_1)$ if $s_2 = b$, where θ denotes the probability that the sophisticated M realizes the project when observing x_2 . The first of these two conditions determines s_0 ; the second one s_1 and $s_2(s_1)$. Note that $s_0 < s_1$ and that these two conditions can hold simultaneously if only $\gamma \geq \hat{\gamma}$.

Moreover, it must hold that E does not have an incentive to deviate and to play $x(s) = s$ for $s \in (\tilde{s}, s_2)$. It follows from $c''(d) \geq 0$ and the second condition for E that he is better off with $x(s) = s_2$ than with $x(s) = s$ at any $s \in (s_1, s_2)$. ■

B.2 Multiple Like Biased Experts

Consider $N \geq 2$ experts with biases $\Phi_i > 0$, where $i = 1, \dots, N$, $\Phi_1 \leq \dots \leq \Phi_N$. Biases are common knowledge, and each expert i obtains Φ_i if and only if M realizes the project. All experts face $c(d)$ and observe s . We first treat simultaneous and then sequential reporting.

Preliminaries M now forms her beliefs about state s and about the expected net utility $W(s)$ of the project after observing a set of N reports $\{x_1, \dots, x_N\}$. We thus have to specify what M believes when observing a set of reports that she should not observe in equilibrium. Following the logic of Nash equilibrium, we thereby focus on unilateral deviations. Given $N \geq 3$, it is easy for M to figure out which expert has unilaterally deviated. When the

information from the other $N - 1$ experts is sufficient to deduce s with certainty, M has no need to think about the motives of the deviating expert. But when the information from the other $N - 1$ experts is not sufficient to derive s with certainty, e.g., because they are all sending some pooling report \hat{x} , M takes the information that is possibly provided by the deviating expert into account as suggested by PSE. Given $N = 2$ and expert i has sent a report x_i that he should never send in equilibrium, M assumes that this expert has deviated unilaterally. If both reports are part of the experts' equilibrium strategies, i.e., if there are $s, s' \in [a, b]$ such that $x_1(s) = x_1$ and $x_2(s') = x_2$ but with, say, $s < s'$, we assume that M believes with probability one that the expert whose report is associated with the higher state (x_2 in the example) has unilaterally deviated. This is motivated by the incentives to over-report. When updating her beliefs about s , M decides in the same way for $N = 2$ as when $N \geq 3$ whether or not to take the report of the deviating expert into account.

Simultaneous Reporting Under simultaneous reporting the results are as follows:²¹

Proposition B.2 *When the multiple like biased experts report simultaneously, there exists a PSE in which each expert i reports as he would in the (or a) PSE if he were the only expert and his bias were Φ_1 , and on equilibrium M decides as she would if there were only one expert whose bias was Φ_1 .*

There also exists a fully separating PSE in which each expert i reports truthfully $x_i(s) = s$ for all s . There are no other PSE.

Proof: The first part of the statement follows directly from the results of Section 3 that on equilibrium M 's beliefs are updated using Bayes' rule given the experts' strategies, and that given these beliefs her strategy is optimal. It remains to show that no expert i wants to deviate unilaterally: No i wants to deviate and report any higher x_i because for any s this would increase distortion costs without affecting M 's decision. No i wants to deviate and report any lower \hat{x}_i if $x_i(s) = s$ is prescribed because this would also increase distortion costs without affecting M 's decision. No i wants to deviate and report any lower x_i if $x_i(s) \neq s$ is prescribed, i.e., if all experts send the pooling report x^p . The reason is that M is indifferent

²¹As seen earlier, there may be multiple PSE with a single strongly biased expert who faces an unfavorable prior. Propositions B.2 and B.3 do not preclude such multiplicity.

when knowing that $s \in [s^p - \bar{d}_1, s^p]$ (after observing x^p), where $s^p \equiv x^p$, but that the logic of PSE requires M to prefer not to realize the project when observing some $x < x^p$, as this indicates that $s < s^p$.

It is also a PSE that $x_i(s) = s$ for all i and s because any unilateral deviation only increases distortion costs without affecting M 's decision. There cannot be another fully separating PSE, as each single i could reduce distortion costs without negatively affecting M 's decision whenever $x_i(s) < s$, but also whenever $x_i(s) > s$ and either $s < \tilde{s}$ or $x_i(s) > x_i(\tilde{s})$. The proof that there are no further PSE with pooling segments follows from the observation that no expert i would ever play a pooling strategy $x_i(s)$ for some range of s for which another expert j plays a separating strategy $x_j(s)$, as this would only lead to higher distortion costs without affecting M 's decision. Hence, if there is a pooling segment, it must be the same for all experts. The remainder of the proof follows along the lines of the proofs with a single expert. ■

Sequential Reporting Now assume that the experts report sequentially and that every expert can observe the reports of his predecessors.²² The experts may report in an arbitrary order with expert 1 moving first, last or anywhere in between. The PSE in which all experts behave as if they were the only expert and their bias were the smallest bias Φ_1 still exists with sequential reporting. Each expert i knows that expert 1 never lies more than \bar{d}_1 , and that M learns the true s whenever one (or more) expert(s) report truthfully. Hence, no expert i is willing to lie any more than \bar{d}_1 , no matter whether he reports before or after expert 1.

There exists no fully separating PSE. The expert who reports first would have an incentive to deviate and to report some $x \geq \tilde{x}$ whenever $s \in (\tilde{s} - \bar{d}_1, \tilde{s})$, and the other experts would then have an incentive to follow suit and send the same report as the first expert. To summarize:

Proposition B.3 *When multiple like biased experts report sequentially, there exists a PSE in which each expert i reports as he would in the (or a) PSE if he were the only expert and his bias were Φ_1 , and on equilibrium M decides as she would if there were only one expert with bias Φ_1 . There are no other PSE.*

²²Without this observability assumption we would be back in the game with simultaneous moves.

B.3 Uncertain Size of Expert's Bias

Consider a single expert E with bias Φ unknown to M. It is common knowledge that Φ takes value Φ_1 with probability α and value Φ_2 with probability $(1 - \alpha)$ with $0 < \Phi_1 < \Phi_2$. E with bias Φ_i is called type i , with $i = 1, 2$. The timing is such that E first learns his type i and the state s and then sends a report x . We restrict our attention to cases in which both types are relatively similar and belong to the same kind of expert, i.e., either Φ_2 is small enough to still qualify as a moderately biased expert in a single expert setup with no uncertainty about E's type, or Φ_1 is big enough to qualify as a strongly biased expert in such a setup.

Moderate Biases Denote by $\bar{d}_i \equiv c^{-1}(\Phi_i)$ the maximal lie type i is willing to tell. Assume both types are moderately biased. Let s'_i be the cutoff state above which type i gets his preferred decision when reporting truthfully in the single expert setup without uncertainty about the expert type. It holds true that $\bar{s} < s'_1 < s'_2$, where \bar{s} is such that $W(\bar{s}) = 0$. Denote further by x'_i the pooling report sent by i in the PSE of the single non-stochastic expert type setup. The two types are sufficiently similar so that $s'_2 - \bar{d}_1 < \bar{s}$. Observe that both types playing their equilibrium strategy of the single expert setup with non-stochastic type is not part of an equilibrium here because type 2 would have an incentive to send report x'_1 for all $s \in [s'_1 - \bar{d}_2, s'_1]$ and $x(s) = s$ otherwise. Define

$$EU_M(\alpha, s_0) \equiv \alpha \int_{s_0 - \bar{d}_1}^{s_0} \frac{W(s)f(s)ds}{F(s_0) - F(s_0 - \bar{d}_1)} + (1 - \alpha) \int_{s_0 - \bar{d}_2}^{s_0} \frac{W(s)f(s)ds}{F(s_0) - F(s_0 - \bar{d}_2)}, \quad (8)$$

and note that EU_M increases in s_0 . Let $s'(\alpha)$ be such that $EU_M(\alpha, s'(\alpha)) = 0$.²³

Proposition B.4 *With uncertainty about the expert's moderate bias, there exists a unique PSE. In this equilibrium, type 1 plays $x_1(s) = s'(\alpha)$ for all $s \in [s'(\alpha) - \bar{d}_1, s'(\alpha)]$ and $x_1(s) = s$ otherwise. Type 2 plays $x_2(s) = s'(\alpha)$ for all $s \in [s'(\alpha) - \bar{d}_2, s'(\alpha)]$ and $x_2(s) = s$ otherwise. M plays $v(x) = 1$ if $x \geq s'(\alpha)$ and $v(x) = 0$ otherwise.*

Proof: Proofs of existence and uniqueness closely follow along the lines of the proofs of the

²³Observe that at $s_0 = s'(\alpha)$ the first term in (8) is positive and the second is negative. Therefore, $\partial EU_M(\alpha, s'(\alpha))/\partial \alpha > 0$ and the implicit function theorem can be invoked to establish $ds'(\alpha)/d\alpha < 0$. Moreover, $s'(1) = s'_1$ and $s'(0) = s'_2$.

model with a single expert with a known bias analyzed in Section 3. ■

The above logic extends straightforwardly to a setup with K types, labeled in increasing order, as long as $s'_K - \bar{d}_1 < \tilde{s}$. Proposition B.4 highlights that type 1 (2) gets the project passed less (more) frequently than if M knew his type, as $s'_1 < s'(\alpha) < s'_2$ for any $\alpha \in (0, 1)$.

Strong Biases If both (or all) types are strongly biased, the PSE strategies of all types are identical to the equilibrium strategy of the single, strongly biased expert of Section 3 if the prior is favorable. Indeed, the single expert's equilibrium strategy does not depend on the exact size of his bias since the point where he starts sending separating (and truthful) reports is entirely determined by the decision maker's payoffs.

In contrast, when both types of biases are strong but the prior is unfavorable, a PSE is slightly more involved, yet still very similar to the one in the model with one expert of known bias. The thresholds s'''_1 and s'''_2 , above which the expert sends the pooling report $x_1(s) = b$, are in this case type specific. These thresholds satisfy $s'''_1 > s'''_2$ and

$$EU_M(\alpha, s'''_1, s'''_2) \equiv \alpha \int_{s'''_1}^b \frac{W(s)f(s)ds}{1 - F(s'''_1)} + (1 - \alpha) \int_{s'''_2}^b \frac{W(s)f(s)ds}{1 - F(s'''_2)} = 0. \quad (9)$$

The corresponding PSE can be sketched as follows: Let $\theta \in (0, 1)$ be the probability that M realizes the project when observing $x = b$ and assume that type i sends $x = b$ for $s \geq s'''_i$ and $x = s$ otherwise. Then $\theta\Phi_1 = c(b - s'''_1)$ and $\theta\Phi_2 = c(b - s'''_2)$ must hold, implying $c(b - s'''_2)/\Phi_2 = c(b - s'''_1)/\Phi_1$. Thus, for any given s'''_1 there is a unique s'''_2 such that this equality holds and satisfies $s'''_2 = b - c^{-1}\left(\frac{\Phi_2}{\Phi_1}c(b - s'''_1)\right) \equiv \delta(s'''_1)$, where $\delta(s_1)$ increases continuously in s_1 and Φ_1 , decreases in Φ_2 and satisfies $\delta(s_1) < s_1$ for all $s_1 < b$ and $\delta(b) = b$. Let $s'''(\alpha)$ be such that $EU_M(\alpha, s'''(\alpha), \delta(s'''(\alpha))) = 0$. To ensure that such a $s'''(\alpha)$ satisfying $s'''(\alpha) \leq b$ and $\delta(s'''(\alpha)) \geq a$ exists, we assume $EU_M(\alpha, \delta^{-1}(a), a) \leq 0$. Then, there exists a PSE in which type 1 plays $x_1(s) = s$ for all $s < s'''(\alpha)$ and $x_1(s) = b$ otherwise, type 2 plays $x_2(s) = s$ for all $s < \delta(s'''(\alpha))$ and $x_2(s) = b$ otherwise, and M plays $v(x) = 0$ for $x < b$ and $v(b) = \theta$, where $\theta = c(b - s'''(\alpha))/\Phi_1$. If $EU_M(\alpha, \delta^{-1}(a), a) > 0$, the PSE is similar except that type 2 plays $x(s) = b$ for all s .

Thus, in equilibrium type 1 (2) therefore sends the pooling report $x_1(s) = b$ again less (more) often than if his bias were known.