
A Model of Expertise

Author(s): Vijay Krishna and John Morgan

Source: *The Quarterly Journal of Economics*, Vol. 116, No. 2 (May, 2001), pp. 747-775

Published by: The MIT Press

Stable URL: <http://www.jstor.org/stable/2696478>

Accessed: 27/04/2010 19:44

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=mitpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The MIT Press is collaborating with JSTOR to digitize, preserve and extend access to *The Quarterly Journal of Economics*.

A MODEL OF EXPERTISE*

VIJAY KRISHNA AND JOHN MORGAN

We study a model in which perfectly informed experts offer advice to a decision maker whose actions affect the welfare of all. Experts are biased and thus may wish to pull the decision maker in different directions and to different degrees. When the decision maker consults only a single expert, the expert withholds substantial information from the decision maker. We ask whether this situation is improved by having the decision maker sequentially consult two experts. We first show that there is no perfect Bayesian equilibrium in which full revelation occurs. When both experts are biased in the same direction, it is never beneficial to consult both. In contrast, when experts are biased in opposite directions, it is always beneficial to consult both. Indeed, in this case full revelation may be induced in an extended debate by introducing the possibility of rebuttal.

I. INTRODUCTION

The power to make decisions rarely resides in the hands of those with the necessary specialized knowledge. Instead, decision makers often solicit experts for advice. Thus, a division of labor has arisen between those who have the relevant expertise and those who make use of it. The diverse range of problems confronted by decision makers, such as corporate CEOs or political leaders, almost precludes the possibility that they themselves are experts in all relevant fields, and hence, the need for outside experts naturally arises. CEOs routinely seek the advice of marketing specialists, investment bankers, and management consultants. Political leaders rely on a bevy of economic and military advisers. Investors seek tips from stockbrokers and financial advisers.

These and numerous other situations share some common features.

First, the experts dispensing advice are by no means disinterested. Experts may attempt to influence the decision maker in ways that are not necessarily in the latter's best interests. Investment banks stand to gain from new issues and corporate mergers, decisions about which they regularly offer advice. The political future of economic and military advisers may be affected by the decisions on which they give counsel. Stockbrokers are obviously interested in the investment decisions of their clients.

* This research was supported by a grant from the National Science Foundation (SBR 9618648). We are grateful to Gene Grossman, George Mailath, Tomas Sjöström, Joel Sobel, and the referees for sharing their expertise with us.

Second, decision makers are often bombarded with advice from numerous experts, with possibly different agendas. Moreover, experts may strategically tailor their advice to counter that offered by other, rival, experts. For instance, hawks may choose more extreme positions on an issue if they know that doves are also being consulted, and vice versa. Thus, the decision maker faces the daunting task of sifting through the mass of conflicting opinions and deciding as to the best course of action. Indeed, this ability is routinely touted as the mark of a good leader.

In determining the size and composition of her "cabinet" of advisers, the decision maker must carefully consider the following questions: is it possible to extract all information relevant to the decision from a cabinet? Is it better to actively consult a number of advisers or only a single, well-chosen, adviser? Is an advocacy system, where the decision maker appoints experts with opposing viewpoints, helpful in deciding on the correct action? How do experts with extreme views affect the advice offered to the decision maker? Does it help to have an extended debate in which experts can make counterarguments? These questions form the central focus of our paper.

To address these questions, we use a simple model of the interplay among a single decision maker and *two* perfectly informed, but interested experts. The experts offer advice to the decision maker in order to influence the decision in a way that serves their own, possibly differing, objectives. We ask how a decision maker should integrate the opinions of experts when faced with this situation. We begin with a model where each expert speaks only once. Speeches are made sequentially and publicly.

In our model, an expert's preferences are parameterized by his inherent bias relative to the decision maker. The experts may differ both in terms of how biased they are and in which direction. They may have opposing biases: one expert may wish to pull the decision maker to the left and the other to the right. Alternatively, they may have like biases: both wish to pull in the same direction but possibly to differing degrees. The absolute value of the bias parameter indicates how "loyal" an expert is to the decision maker.

Like biases. When experts have like biases, we find that the decision maker derives *no* benefit relative to consulting only one expert, in the sense that no equilibrium with multiple experts is

superior to the most informative equilibrium with a single expert. Moreover, *ex ante* all parties, including the less loyal expert, would agree that the best course of action is for the decision maker to consult only the more loyal expert. This implies that even with two identically informed experts all equilibria result in some information loss.¹

Opposing biases. The situation changes dramatically when experts have opposing biases. Now, the decision maker always derives some benefit from consulting both experts relative to consulting only one. However, this conclusion holds only if at least one of the experts is not an “extremist.” If both of the experts are extremists, no information is revealed in any equilibrium—either when they are consulted separately or in combination.²

Rebuttal. We then explore what happens when experts engage in an extended back-and-forth debate. When experts have opposing biases, the introduction of a rebuttal stage in the debate can result in *full revelation* and first-best outcomes. Extended debate does not, however, lead to full revelation when experts have like biases.

Related Work. Our basic model is closely related to the model of Crawford and Sobel [1982] of strategic information transmission between two parties, one of whom has information useful for the other. We depart from Crawford and Sobel, hereinafter referred to as CS, in that we allow for *multiple* sources of information. The additional strategic considerations that arise with multiple experts lead to technical complications not present with a single expert. Differences between the single expert model of CS and our model are highlighted in later sections.

Also closely related are papers by Gilligan and Krehbiel [1989] and Krishna and Morgan [2001]. Gilligan and Krehbiel study a model where a committee consisting of two “experts” with opposing biases communicate to the decision-making legislature by simultaneously submitting bills to the floor of the House. They argue that the restrictive “closed rule,” which does not permit the amendment of bills submitted to the floor, is informationally

1. Our results in the case of like biases concern monotonic equilibria where the action taken is a monotonic function of the state. While nonmonotonic equilibria exist, we provide sufficient conditions for all equilibria to be monotonic.

2. Our results in the case of opposing biases hold generally for all equilibria, not just monotonic equilibria.

superior to the “open rule” under which bills are freely amendable. This model differs from that in the present paper in that experts’ advice is offered simultaneously and considers only the case of opposing biases. Krishna and Morgan have reexamined the Gilligan and Krehbiel model and have shown that full informational efficiency is attainable under the open rule but not under the closed. As we show below, full efficiency cannot be attained in our model of sequential communication. Thus, in contrast to previous work, our paper highlights how the informativeness of multiple experts differs depending on whether they are of like or opposing biases, as well as on the timing of the debate.

Austen-Smith [1993] also studies a model with two experts. His model differs from ours in two regards. First, the state space is binary, as is the signal space for each expert. Second, the experts are imperfectly informed about the state. Together, these imply that full revelation (the expert reveals all that he knows) is possible even with a *single* expert, a property that is impossible with a richer signal space. In his model, it is sometimes the case that full revelation is possible with a single expert, but not when two experts are consulted simultaneously. This is exactly the opposite of the results we obtain. Also, contrary to our findings, Austen-Smith [1993] finds that sequential consultation is superior to simultaneous consultation. In our model, full revelation is obtainable in the simultaneous model, whereas it is impossible in the sequential model. Finally, while Austen-Smith [1993] finds little difference between like and opposing biases, we find that these situations are in stark contrast.

Dewatripont and Tirole [1999] examine reward schemes to induce information gathering by “advocates.” Advocates differ from experts in that they are not interested in the decision per se but care about it only to the extent that their reward may depend on the decision. Dewatripont and Tirole [1999] find that informational benefits are maximized by making each of the advocates responsible for a distinct area and compensating them accordingly. In contrast, informational benefits in our model derive solely from differences in the interests and ideologies of the experts.

Sobel [1985] and Morris [2001] study how reputational considerations affect a single expert’s advice when the bias of the expert is uncertain. Morgan and Stocken [1998] consider this problem in a static CS-like setting and focus on information transmission by equity analysts. Ottaviani and Sorensen [1997]

study reputational issues with multiple experts. In their model the experts are not directly affected by the decisions but care only about making recommendations that are validated *ex post*.

Banerjee and Somanathan [2001] and Friedman [1998] examine information transmission in a setting in which there is a continuum of potential experts with differing prior beliefs. Austen-Smith [1990] examines the effect of debate on voting outcomes. In his model “expert” legislators themselves vote on legislation, so the separation between the experts and the decision maker is absent. The effects of combining information provided by experts with opposing incentives has also been examined by Shin [1994] in the context of persuasion games (see Milgrom and Roberts [1986]). Finally, although we are not concerned with implementation *per se*, the problem of how a decision maker should extract information from interested experts is related to Baliga, Corchon, and Sjöström [1997], who study abstract implementation problems where the planner is unable to commit to a mechanism.

As is well-known, models with cheap talk suffer from a plethora of equilibria and efforts to identify some as salient have led to the development of a substantial literature on refinements in this context [Matthews, Okuno-Fujiwara, and Postlewaite 1991; Farrell 1993]. Farrell and Rabin [1996] present a concise survey. The models we consider also have multiple equilibria; however, for the most part, our focus is on the “most informative” equilibrium.

II. PRELIMINARIES

In this section we sketch a simple model of decision making with multiple experts. Rather than modeling any of the examples mentioned in the introduction explicitly, we consider a stylized representation of the interaction among a decision maker and experts that applies to a broad range of institutional settings. Precisely, we extend the model of CS to a setting with multiple experts.

Consider a *decision maker* who takes an action $y \in \mathbf{R}$, the utility from which depends on some underlying state of nature $\theta \in [0,1]$ which is distributed according to the density function $f(\cdot)$. The decision maker has no information about θ , but there are two *experts* each of whom observes θ .

The two experts then offer “advice” to the decision maker by sending messages $m_1 \in [0,1]$ and $m_2 \in [0,1]$, respectively. The

messages are sent *sequentially* and *publicly*. First, expert 1 offers his advice, which is heard by both the decision maker and expert 2. Expert 2 then offers his advice, and the decision maker takes an action. The decision maker is free to interpret the messages however she likes as well as to choose any action.

In Section VI we allow for the possibility of extending the length of the "debate." In particular, we study a situation where following the initial round of messages, m_1 and m_2 , each expert offers a "rebuttal" message, r_1 and r_2 , respectively. As in the initial round, these messages are also sent sequentially and publicly.

The utility functions of all agents are of the form $U(y, \theta, b_i)$, where b_i is a parameter which differs across agents. For the decision maker, agent 0, b_0 is normalized to be 0. We write $U(y, \theta) \equiv U(y, \theta, 0)$. For the experts, agents 1 and 2, $b_i \neq 0$. We suppose that U is twice continuously differentiable, $U_{11} < 0$, $U_{12} > 0$, $U_{13} > 0$. Since $U_{13} > 0$, b_i is a measure of how *biased* expert i is, relative to the decision maker. For each i , $U(y, \theta, b_i)$ attains a maximum at some y . Since $U_{11} < 0$, the maximizing action is unique. The biases of the two experts and the decision maker are commonly known.

These assumptions are satisfied by "quadratic loss functions," in which case,

$$(1) \quad U(y, \theta, b_i) = -(y - (\theta + b_i))^2.$$

When combined with the assumption that θ is uniformly distributed on $[0, 1]$, this is referred to as the "uniform-quadratic" case, first introduced by CS.

The multiple experts problem is divided into two cases. If both $b_1, b_2 > 0$, then the experts are said to have *like biases*. If $b_i > 0 > b_j$, then the experts are said to have *opposing biases*.³

Define $y^*(\theta) = \arg \max_y U(y, \theta)$ to be the *ideal* action for the decision maker when the state is θ . Similarly, $y^*(\theta, b_i) = \arg \max_y U(y, \theta, b_i)$ is the ideal action for expert i . Since $U_{13} > 0$, $b_i > 0$ implies that $y^*(\theta, b_i) > y^*(\theta)$, and since such an expert always prefers a higher action than is ideal for the decision maker, we refer to him as being *right-biased*. Similarly, if $b_i < 0$, then $y^*(\theta, b_i) < y^*(\theta)$, and we refer to such an expert as being *left-biased*. With quadratic loss functions, $y^*(\theta, b_i) = \theta + b_i$.

3. The case where both $b_1, b_2 < 0$ is qualitatively no different from the case where both $b_1, b_2 > 0$.

A word of caution is in order. Our results fall into two categories. Some concern the structure of equilibria of the multiple experts game and are derived under the assumptions given above. Others concern welfare comparisons among equilibria and, require the same assumption as made in CS, Assumption M [p. 1444 of CS], in order to derive unambiguous welfare results (specifically, Proposition 2). This assumption, while not so transparent, is satisfied by the uniform-quadratic case.

III. EQUILIBRIUM WITH EXPERTS

Single Expert. In the single expert game studied by CS, a strategy for the expert, μ , specifies the message $m = \mu(\theta)$ that he sends in state θ . A strategy for the decision maker, y , specifies the action $y(m)$ that she takes following any message m by the expert.

CS show that every Bayesian equilibrium of the single expert game has the following structure. There are a finite number of equilibrium actions y_1, y_2, \dots, y_N . The state space is partitioned into N intervals $[0, a_1), [a_1, a_2), \dots, [a_{n-1}, a_n), \dots, [a_{N-1}, 1]$ with action y_n resulting in any state $\theta \in [a_{n-1}, a_n)$. The equilibrium actions are monotonically increasing in the state, that is, $y_{n-1} < y_n$. Finally, at every “break point” a_n the following “no arbitrage” condition is satisfied:

$$(2) \quad U(y_n, a_n, b) = U(y_{n+1}, a_n, b).$$

In other words, in state a_n the expert is indifferent between actions y_n and y_{n+1} . Since $U_{12} > 0$, for all $\theta < a_n$, the expert prefers y_n to y_{n+1} and for all $\theta > a_n$, the reverse is true. Thus, (2) serves as an incentive (or self-selection) constraint.

Multiple Experts. In the multiple experts game a pure strategy for expert 1, μ_1 , specifies the message $m_1 = \mu_1(\theta)$ that he sends in state θ . A pure strategy for expert 2, μ_2 , specifies the message $m_2 = \mu_2(\theta, m_1)$ that he sends in state θ after hearing message m_1 from expert 1. A strategy for the decision maker, y , specifies the action $y(m_1, m_2)$ that she takes following messages m_1 and m_2 . Let $P(\cdot | m_1, m_2)$ denote the posterior beliefs on θ held by the decision maker after hearing messages m_1 and m_2 .

A (pure strategy) *perfect Bayesian equilibrium* (PBE) entails: (1) for all messages m_1 and m_2 , $y(m_1, m_2)$ maximizes the decision maker’s expected utility given her beliefs $P(\cdot | m_1, m_2)$; (2) the

beliefs $P(\cdot | m_1, m_2)$ are formed using the experts' strategies μ_1 and μ_2 by applying Bayes' rule wherever possible; (3) given the decision maker's strategy y , for all θ and m_1 , $\mu_2(\theta, m_1)$ maximizes expert 2's utility; and (4) given the strategies y and μ_2 , for all θ , $\mu_1(\theta)$ maximizes expert 1's expected utility.⁴

Given a PBE, we denote by Y the *outcome function* that associates with every state the resulting action. Formally, for each θ , $Y(\theta) = y(\mu_1(\theta), \mu_2(\theta, \mu_1(\theta)))$. Denote by $Y^{-1}(y) = \{\theta : Y(\theta) = y\}$. Every Y determines an *equilibrium partition* of the state space, $\mathcal{P} = \{Y^{-1}(y) : y \text{ is an equilibrium action}\}$. The partition \mathcal{P} is then a measure of the informational content of the equilibrium.

An action y is said to be *rationalizable* if there are some beliefs that the decision maker could hold for which y is a best response. Clearly, an action y is rationalizable if and only if $y^*(0) \leq y \leq y^*(1)$.

A PBE always exists. In particular, there is always a "babbling" equilibrium in which all messages from both experts are completely ignored by the decision maker. Obviously, information loss is most severe in such an equilibrium. Typically, there are also other, more informative, equilibria.

Example 1. Consider the uniform-quadratic case with $b_1 = 1/40$ and $b_2 = 1/9$, so that the experts have *like* biases and expert 1 is less biased than is expert 2. A PBE for this game is depicted in Figure I, where the states $a_1 = 1/180$, $a_2 = 22/180$, $a_3 = 61/180$ and the actions $y_1 = 1/360$, $y_2 = 23/360$, $y_3 = 83/360$, $y_4 = 241/360$.

In Figure I the outcome function Y is the step function depicted by the dark lines. The dotted lines depict the experts' ideal actions $y^*(\theta, b_i) = \theta + b_i$. The resulting information partition is $\{[0, a_1], [a_1, a_2], [a_2, a_3], [a_3, 1]\}$. The action y_1 is optimal for the decision maker given that $\theta \in [0, a_1]$, y_2 is optimal given $\theta \in [a_1, a_2)$, etc.

To see that this is an equilibrium configuration, notice that in state a_2 expert 1 is exactly indifferent between actions y_2 and y_3 (In Figure I this indifference is indicated by the vertical double-pointed arrow centered on $a_2 + b_1$.) Expert 1 prefers y_2 to y_3 in all states $\theta < a_2$ and prefers y_3 to y_2 in all states $\theta > a_2$. Thus, given the decision maker's strategy, he is willing to distinguish between states $\theta < a_2$

4. The formal definition of a PBE requires only that the various optimality conditions hold for *almost every* state and pair of messages. This would not affect any of our results.




FIGURE I
A PBE with Like Biases

and states $\theta > a_2$. Similarly, in state a_3 expert 2 is indifferent between y_3 and y_4 and is willing to distinguish between states $\theta < a_3$ and states $\theta > a_3$. Thus, in states a_2 and a_3 the CS “no arbitrage” condition (2) holds for either expert 1 or expert 2.

In state a_1 , however, neither expert is indifferent between y_1 and y_2 . Indeed, expert 1 prefers y_1 to y_2 in state a_1 . (Expert 1’s ideal action is closer to y_1 than y_2 .) Expert 2, on the other hand, prefers y_2 to y_1 . The equilibrium calls for expert 1 to “suggest” action y_1 and for expert 2 to “agree.” Expert 2 has the option of “disagreeing” with expert 1 and inducing action y_3 . Expert 2 is indifferent between y_1 and y_3 in state a_1 and so prefers y_3 to y_1 if $\theta > a_1$. Thus, even though in states just above a_1 , expert 1 prefers y_1 to y_2 , if he were to suggest y_1 , expert 2 would disagree, resulting in y_3 . Since y_2 is preferred to y_3 by expert 1 in these states, expert 1 will not deviate. Here we see how the strategic interaction of the two experts creates the possibility of “disciplining” the experts in a manner not possible for the single expert case.⁵

5. A detailed specification of the equilibrium strategies and beliefs for this and all other examples in the paper may be obtained from the authors.

III.1. Full Revelation

A natural question is whether the increase in discipline possible with the addition of multiple experts can lead to full revelation. Full revelation means that for all θ , $Y(\theta) = y^*(\theta)$. Recall that with only a single expert, CS show that full revelation is *not* a *Bayesian equilibrium (BE)*.

With multiple experts, however, full revelation can occur in a BE. Suppose that experts have like biases and the decision maker holds the beliefs $P(\theta = \min\{m_1, m_2\} | m_1, m_2) = 1$. The associated strategy of the decision maker is then $y(m_1, m_2) = y^*(\min\{m_1, m_2\})$. Let expert 1 follow the strategy $\mu_1(\theta) = \theta$ and expert 2 follow the strategy: $\mu_2(\theta, m_1) = \theta$. In state θ , both experts send messages $m_1 = m_2 = \theta$, and the action is $y^*(\theta)$ which is preferred by both experts to any action $y < y^*(\theta)$. Reporting an $m_i < \theta$ will only decrease i 's utility, whereas reporting an $m_i > \theta$ will have no effect.⁶

The equilibrium constructed above, however, involves non-optimizing behavior on the part of expert 2 off the equilibrium path. Specifically, in state $\theta \in [0, 1)$ if expert 1 were to choose a message $m_1 = \theta + \epsilon$, for $\epsilon > 0$ small enough, it is no longer optimal for expert 2 to play $\mu_2(\theta, m_1) = \theta$. Indeed, he is better off also deviating to $m_2 = m_1$. Thus, the full revelation BE constructed above is not a PBE. This is true in general.

PROPOSITION 1. There does not exist a fully revealing PBE.

Proof. See Appendix. ■

IV. EXPERTS WITH LIKE BIASES

In this section we focus on two questions: first, what is the information content of advice offered by a given panel of like biased experts; and second, how should a decision maker determine the composition of such a panel.

IV.1. Choosing a Cabinet

To illustrate the key underlying issues, it is useful to study the example given earlier in greater detail.

6. When messages are sent simultaneously and experts have like biases, the above construction is a fully revealing PBE. With simultaneous messages and opposing biases, Krishna and Morgan [2001] show that full revelation is a PBE, but the construction is somewhat more involved.

Example 2. Once again consider the uniform-quadratic case where $b_1 = 1/40$ and $b_2 = 1/9$. If only expert 1 is consulted, the most informative equilibrium results in the partition \mathcal{P}_1 :

+	-	+	-	-	+	-	-	-	-	-	-	+
		<i>e1</i>			<i>e1</i>							
0		$\frac{1}{10}$			$\frac{3}{10}$					$\frac{6}{10}$		1

This means that if the true state θ lies in the interval $[0,1/10)$, the expert sends a message suggesting the action $y_1 = 1/20$. Similarly, if $\theta \in [1/10,3/10)$, he suggests $y_2 = 4/20$, and so on. The decision maker's expected utility is -0.0083 .

We will refer to points such as $1/10$, $3/10$, and $6/10$ as "break points." In the figure each break point is labeled with the name of the expert, e_1 or e_2 , whose message distinguishes states below the break point from states above.

Similarly, if the decision maker solicited only expert 2 for advice, the most informative equilibrium partition is \mathcal{P}_2 :

+	-	-	-	-	+	-	-	-	-	-	-	-	+
					<i>e2</i>								
0					$\frac{5}{18}$								1

resulting in an expected utility of -0.0332 to the decision maker. Notice that expert 2 withholds more information than does expert 1, in the sense that the variance of the true state, given the equilibrium partition, is higher with expert 2 than expert 1. Intuitively, since expert 2 wishes the decision maker to choose a larger value of y than does expert 1, he withholds more information than does expert 1.

If there were no further strategic considerations, that is, neither expert knew of the other's existence, the decision maker could *combine* the reports of the two experts to obtain the partition $\mathcal{P}_1 \wedge \mathcal{P}_2$:

+	-	+	-	-	+	+	-	-	-	+	-	-	-	-	+
		<i>e1</i>			<i>e2</i>	<i>e1</i>				<i>e1</i>					
0		$\frac{1}{10}$			$\frac{5}{18}$	$\frac{3}{10}$				$\frac{6}{10}$					1

which is the coarsest common refinement (*join*) of \mathcal{P}_1 and \mathcal{P}_2 . The decision maker's expected utility is now -0.0081 . Thus, it seems

plausible that the addition of another expert, even an expert more biased than expert 1, might be helpful in overcoming the problem of strategic information withholding.

Of course, this ignores strategic interaction among the experts. Indeed, the specification above is not a PBE in the multiple experts game. What is the structure of PBEs when the experts are aware of each other's presence? One such equilibrium was described in Example 1 with the following information partition:

	e1	e1	--	e2	--	--	--	--	--	--	--	+
+	+	+	--	+	--	--	--	--	--	--	--	+
0	$\frac{1}{180}$	$\frac{22}{180}$		$\frac{61}{180}$								1

Recall that at the point $\theta = 1/180$ neither expert is indifferent between y_1 and y_2 . The decision maker's expected utility in this equilibrium is $-.0250$.

A PBE with the property that expert 1 is indifferent at the first point of discontinuity results in the equilibrium partition \mathcal{Q} :

	e1	e1	--	e2	--	--	--	--	--	--	--	+
+	+	+	--	+	--	--	--	--	--	--	--	+
0	$\frac{1}{72}$	$\frac{23}{180}$		$\frac{41}{120}$								1

where expert 1 is also indifferent at the second break point and expert 2 is indifferent at the third. This results in expected utility of $-.0247$ to the decision maker and is better than the equilibrium of Example 1. Intuitively, by shifting the first break point to the right, informativeness is improved since all of the other break points shift to the right as well. Hence the resulting partition is "closer" to the most informative partition with three break points; that is, where the break points are equally spaced. In Example 1 since expert 1 is not indifferent at the first break point, such a rightward shift is possible. Indeed, one can show that \mathcal{Q} is the most informative PBE in which there are three break points and both experts' messages are relevant. It is not possible to have a fourth interior break point.

Comparing \mathcal{Q} to \mathcal{P}_1 , we see that the more biased expert 2 distorts the third break point to the left, from $6/10$ to $41/120$. This reduces the information content in the rightmost interval. Moreover, the leftward shift by expert 2 shifts all of the other break points to the left; thus, it also distorts expert 1's break points to

the left, from $3/10$ down to $23/180$ and from $1/10$ to $1/72$. The aggregate effect of these distortions is to reduce the expected utility of the decision maker and that of both experts.

To summarize, whenever there is a break point at which expert 1 is not indifferent, informativeness is enhanced by shifting this break point to the right until expert 1 becomes indifferent. As a result of this shift, there is a complementary, and still more informative, shift to the right of all of the other break points. Thus, any PBE containing a break point where expert 2 is indifferent is detrimental. Put differently, *when experts have like biases, the decision maker can do no better than to consult only the more loyal expert.*

In coming to this conclusion, we relied essentially on two properties of any PBE. First, that a rightward shift of one break point led to a rightward shift of all others and this improved informativeness—this is Assumption M of CS. Second, we relied on the fact that equilibrium partitions consisted only of intervals; that is, the equilibrium outcome function $Y(\cdot)$ is monotonic in θ . We now present the main result of this section, which shows that as long as these two properties are satisfied, the conclusion from the example generalizes.⁷

PROPOSITION 2. When experts have like biases, there is *never* a monotonic PBE with both experts that is informationally superior to the most informative PBE with a single expert.

Proof. See Appendix. ■

Despite the fact that the messages of one expert can be used to discipline the other, in the case of monotonic equilibria with like biases, this disciplining only has the effect of reducing informativeness. Thus, the advice of one of the experts is, at best, redundant. The redundancy result holds regardless of whether the biases of the two experts are close to one another or far apart.

Proposition 2 applies to monotonic PBEs. We now turn to a study of such PBEs in the like biased case and provide some sufficient conditions for all PBEs to be monotonic.

7. Although our analysis in the like bias case is confined to monotonic equilibria, we know of no instance where admitting a nonmonotonic equilibrium reverses the welfare comparisons.

IV.2. Monotonic Equilibria

A PBE is said to be *monotonic* if the corresponding outcome function $Y(\cdot)$ is nondecreasing. Recall that in the case of a single expert we know from CS that all equilibria are monotonic. While the PBE constructed in Example 1 has this property, this is not true in general when there are multiple experts.⁸

Monotonic equilibria must satisfy the following conditions, which are in the nature of incentive constraints.

LEMMA 1. Suppose that Y is monotonic. If Y has a discontinuity at θ and

$$\lim_{\epsilon \downarrow 0} Y(\theta - \epsilon) = y^- < y^+ = \lim_{\epsilon \downarrow 0} Y(\theta + \epsilon)$$

then

$$(3) \quad U(y^-, \theta, \min\{b_1, b_2\}) \geq U(y^+, \theta, \min\{b_1, b_2\}), \text{ and}$$

$$(4) \quad U(y^-, \theta, \max\{b_1, b_2\}) \leq U(y^+, \theta, \max\{b_1, b_2\}).$$

Proof. See Appendix. ■

We next show that when the experts have like biases there can be at most a finite number of equilibrium actions played in any monotonic PBE.

LEMMA 2. Suppose that experts have like biases and Y is monotonic. Then there are a finite number of equilibrium actions.

Proof. See Appendix. ■

The intuition for Lemma 2 is that if two equilibrium actions are too close to one another, then there will be some state where the lower action is called for, while both experts prefer the higher action. But then the first expert can deviate and send a message inducing the higher action, confident that expert 2 will follow his lead.

In the uniform-quadratic case with like biases, if the less loyal expert speaks first, then all PBEs are monotonic. More generally, when $b_1 \geq b_2 > 0$, all PBEs are monotonic as long as the following conditions are satisfied. First, utility functions are *symmetric*, that is, for all θ and k , $U(y^*(\theta, b_i) + k, \theta, b_i) = U(y^*(\theta, b_i) - k, \theta, b_i)$. Second, the utility functions satisfy *nondecreasing divergence*; that is, $y^*(\theta, b_i) - y^*(\theta)$ is nondecreasing in θ . Clearly, utility functions of the quadratic loss function variety

8. An example of a nonmonotonic equilibrium is available from the authors.

satisfy both of these conditions. Notice that the monotonicity of PBEs follows from properties of the utility functions alone and does not rely on any assumptions on the distribution of states.

LEMMA 3. Suppose that utility functions are symmetric and satisfy nondecreasing divergence. If $b_1 \geq b_2 > 0$, then all PBEs are monotonic.

Proof. A proof is omitted but is available from the authors. ■

V. EXPERTS WITH OPPOSING BIASES

A cabinet composed of two experts with like biases is no more effective than simply consulting the more loyal expert alone. In this section we examine whether it is helpful to choose a cabinet where the experts have opposing biases. We begin with an example to show that the extreme conclusion of Proposition 2 no longer holds.

Example 3. Consider the uniform-quadratic case with $b_1 = 1/12$ and $b_2 = -1/12$. With expert 1 alone, the most informative partition \mathcal{P}_1 is

$$\begin{array}{cccccccc}
 & & & e1 & & & & \\
 \vdots & - & - & + & - & - & - & \vdots \\
 0 & & & \frac{1}{3} & & & & 1
 \end{array}$$

Likewise, when expert 2 alone is consulted, \mathcal{P}_2 is

$$\begin{array}{cccccccc}
 & & & & e2 & & & \\
 \vdots & - & - & - & - & - & + & - & - & \vdots \\
 0 & & & & & & \frac{2}{3} & & & 1
 \end{array}$$

In both cases, the expected utility of the decision maker is -0.028 .

Consider a partition equilibrium such that expert 1 reports a size 2 partition \mathcal{Q}_1 with a break at $2/9$ and expert 2 also a size 2 partition \mathcal{Q}_2 with at break at $7/9$. The information available to the decision maker $\mathcal{Q} = \mathcal{Q}_1 \wedge \mathcal{Q}_2$ is then a size 3 partition:

		e1						e2		
+	-	+	-	-	-	-	-	+	-	+
0		$\frac{2}{9}$						$\frac{7}{9}$		1

There are strategies that support this as a PBE and have the property that, when the decision maker hears “inconsistent” messages, he believes one or the other of the experts. The expected utility of the decision maker in this equilibrium is -0.016 .

The example demonstrates that *when experts have opposing biases, the decision maker can benefit by consulting both experts.*

To see why consulting experts with opposing biases is beneficial, it is useful to consider single expert games over truncated portions of the unit interval. First, consider a single expert game where the state space is $[0, 7/9]$. In this case, if only expert 1 is consulted, then he can credibly convey whether or not the state lies below $\theta = 2/9$. Notice, however, that if the state space were the whole unit interval, expert 1 could not *himself* credibly break the interval at $2/9$. With two experts, since expert 2 is truncating the state space at $7/9$, it is as if expert 1 is playing a truncated game, and can credibly break the state at $2/9$. In determining the second break point, expert 2 faces a similar truncated game where the state space is $[2/9, 1]$. By a symmetric argument, expert 2 is able to credibly convey whether or not the state lies above $7/9$. Notice, however, that this reasoning would not work if expert 2's bias were in the same direction as expert 1. In that case, a break point at $7/9$ would not be credible for expert 2 in the truncated game.

By alternating break points between expert 1 and expert 2, the incentives of expert 1 to exaggerate the state upward are mitigated by making relatively longer intervals to the right of his break points. By the same token, the incentives of expert 2 to exaggerate downward are mitigated by making relatively longer intervals to the left of his break points. In the example, the interval $[2/9, 7/9]$ prevents expert 1 from exaggerating upward when $\theta = 2/9$ and expert 2 from exaggerating downward at $\theta = 7/9$.

When experts have opposing biases, partition equilibria which are Pareto superior to consulting a single expert are not atypical. In addition, consulting experts with opposing biases also creates the possibility of “semi-revealing” equilibria. These are equilibria where the decision maker learns the true state over a

portion of the state space. Semi-revealing equilibria are always informationally superior to consulting a single expert. Their construction, however, requires that at least one of the experts not be an “extremist.”

Extremists and Moderates. An expert with bias $b_i > 0$ holds extreme views in state θ if $U(y^*(\theta), \theta, b_i) \leq U(y^*(1), \theta, b_i)$. Similarly, an expert with bias $b_j < 0$ holds extreme views in θ if $U(y^*(0), \theta, b_j) \geq U(y^*(\theta), \theta, b_j)$. If a right-biased (left-biased) expert holds extreme views in θ , then all actions that are higher (lower) than $y^*(\theta)$ are attractive to the expert.

An expert who holds extreme views in every state is said to be an *extremist*. While an extremist will reveal no information if consulted alone, it is not the case that all experts who reveal no information are extremists. In the uniform-quadratic case, an expert is an extremist if $|b_i| \geq 1/2$. An expert with bias $1/2 > |b_i| \geq 1/4$ will reveal no information when consulted alone, but is not an extremist.

An expert who is not an extremist is a *moderate*. Notice, however, that every right-biased moderate has extreme views once the state is large enough. Define $\alpha(b_i) < 1$ to be a state such that for all $\theta > \alpha(b_i)$ a moderate expert i with bias $b_i > 0$ holds extreme views in θ and observe that

$$U(y^*(\alpha(b_i)), \alpha(b_i), b_i) = U(y^*(1), \alpha(b_i), b_i).$$

Similarly, for a left-biased moderate define $\alpha(b_j) > 0$ to be such that for all $\theta < \alpha(b_j)$ expert j with bias $b_j < 0$ holds extreme views in θ . Notice that for both left- and right-biased experts, $\alpha(b)$ is a decreasing function of b . For quadratic loss functions, $\alpha(b_i) = 1 - 2b_i$ if $b_i > 0$ and $\alpha(b_j) = -2b_j$ if $b_j < 0$. As we shall see, states where at least one of the experts does not hold extreme views are conducive to full revelation.

V.1. Semi-Revealing PBE

Consider the case of quadratic loss functions with $b_1 < 0 < b_2 < 1/2$.

Figure II depicts the outcome function Y associated with a semi-revealing PBE. In this equilibrium the state is revealed when it is below $1 - 2b_2$ and not otherwise. In equilibrium, for all $\theta \leq 1 - 2b_2$, expert 1 sends the “true” message $m_1 = \theta$ and expert 2 “agrees” by sending $m_2 = m_1$. If $m_1 < \theta$, expert 2 sends $m_2 = \max(m_1 + 2b_2, \theta + b_2)$. If $\theta < m_1 < 1 - 2b_2$, expert 2




FIGURE II
A PBE with Opposing Biases

sends $m_2 = \min(m_1, \theta + b_2)$. For all $\theta > 1 - 2b_2$, expert 1 suggests $m_1 = 1 - b_2$ and expert 2 agrees.

Following $m_1 \leq 1 - 2b_2$ and any m_2 , the decision maker tentatively believes that expert 1 is telling the truth. Expert 2's recommendation is then deemed to be "self-serving" if under the hypothesis that 1 is telling the truth, the adoption of expert 2's recommendation strictly benefits 2 relative to 1's recommendation; that is, if and only if $U(m_2, m_1, b_2) > U(m_1, m_1, b_2)$. Expert 2's recommendation is adopted if and only if it is not deemed to be self-serving. Otherwise, 1's recommendation is adopted. Following $m_1 > 1 - 2b_2$ and any m_2 , the decision maker chooses $y = 1 - b_2$.

It is clear that if expert 1 tells the truth, the use of the self-serving criterion guarantees that expert 2 can do no better than to also tell the truth. However, if expert 1 chooses to lie and "suggest" a lower action $m_1 < \theta$, expert 2 would counter this by recommending $m_2 = \max(m_1 + 2b_2, \theta + b_2)$ and this would not

be deemed self-serving. Thus, any attempt by expert 1 to deviate by suggesting a lower action will fail since expert 2 will recommend an even *higher* action that will be adopted. Any attempt by 1 to deviate by suggesting a higher action $m_1 > \theta$ will in fact lead to the action $\min(m_1, \theta + b_2) > \theta$.

For $\theta > 1 - 2b_2$, the decision maker cannot credibly apply the self-serving criterion. To see this, suppose that $\theta > 1 - 2b_2$ and $m_1 = \theta - \epsilon$ for some small ϵ . Expert 2 can do no better than to recommend $\max(m_1 + 2b_2, \theta + b_2)$. But since this is greater than $y^*(1)$, the decision maker cannot rationalize this choice even if it is not self-serving. In other words, using expert 2 to discipline expert 1 in a state θ is only possible when expert 2 does not hold extreme views in that state.⁹

Finally, observe that the strategies of the semi-revealing equilibrium depend only on b_2 and are valid for all $b_1 < 0$ as long as $0 < b_2 < 1/2$. The exact value of b_1 plays no role in the construction.

While the construction above was for the uniform-quadratic case, it can be readily extended. For general utility functions and distributions of the state of nature, a construction analogous to the one given above is a semi-revealing PBE where all states at which expert 2 does not hold extreme views are revealed.

V.2. Choosing a Cabinet

We now show that the semi-revealing equilibrium constructed above is informationally superior to the most informative equilibrium with a single expert of bias b_2 as long as he is a moderate.

Recall that in a semi-revealing equilibrium, all states where expert 2 does not hold extreme views are revealed. In contrast, when expert 2 alone is consulted, CS's condition (2) must hold at every break point. In particular, at the last break point a_{N-1} , $U(y_{N-1}, a_{N-1}, b_2) = U(y_N, a_{N-1}, b_2)$.

We claim that $a_{N-1} < \alpha(b_2)$. If $a_{N-1} \geq \alpha(b_2)$ then, since $y_{N-1} < y^*(a_{N-1})$, we have that

$$U(y_{N-1}, a_{N-1}, b_2) < U(y^*(a_{N-1}), a_{N-1}, b_2) \leq U(y^*(1), a_{N-1}, b_2).$$

This implies that for (2) to hold at a_{N-1} , we must have $y_N >$

9. A referee pointed out that if the state space were the whole real line, the same construction would result in full revelation since, in that case, there would be no largest rationalizable action.

$y^*(1)$, but this is impossible because then y_N is not a rationalizable choice for the decision maker.

Since all states below a_{N-1} are revealed when two experts are consulted, we have shown the following:

PROPOSITION 3. When experts have opposing biases and at least one is a moderate, there is *always* a PBE with both experts that is informationally superior to the most informative PBE with a single expert.

Proposition 3 shows that whenever the more loyal expert is willing to reveal some information on his own, the addition of a second expert with opposing bias, regardless of how extreme, creates the possibility of an equilibrium that is strictly preferred by the decision maker and both of the experts.

V.3. Extremists and the "Crossfire Effect"

In adversarial proceedings, it is common to have individuals with extreme views offering opinions, often as expert witnesses. As we showed above, when one of the experts is a moderate, the addition of a second expert, regardless of his bias, is helpful. We now turn to the case where both experts are extremists.

PROPOSITION 4 (CROSSFIRE EFFECT). If both experts are extremists, no information is transmitted in any PBE.

Proof. See Appendix. ■

The Crossfire Effect severely limits the information that can be garnered from opposing extremists.¹⁰ This result highlights the essential role played by the "disagreement" action in constructing the semi-revealing equilibrium. When experts are extremists, all "disagreement" actions that are able to discipline the experts exceed the highest rationalizable action for the decision maker (namely, $y^*(1)$). Thus, there are no beliefs that the decision maker could hold that would lead expert 2 to anticipate such extreme actions being taken. The inability of the decision maker to commit to a disagreement action dramatically reduces the informativeness of the equilibrium.

Finally, we show when an extremist is paired with a moder-

10. The television talk show *Crossfire* regularly pits an avowed right-wing extremist against an avowed left-wing extremist. The debate is singularly uninformative.

ate, a semi-revealing equilibrium arises only if the moderate is consulted second.

Example 4. Consider the uniform-quadratic case when $b_i \leq -1/2$, $1/4 < b_j < 1/2$. When $i = 1$ and $j = 2$, the semi-revealing PBE constructed earlier is preferred by all to consulting either expert singly.

Now suppose that the order of polling is reversed so that $j = 1$ and $i = 2$. In this case, the most informative PBE involves babbling. In the case where expert 2 is an extremist, it can be directly argued that any equilibrium must be monotonic and involve only a finite number of actions. One can then show that one of the experts must be indifferent at points of discontinuity. Then, since neither expert will reveal any information when polled alone, it follows that there can be no points of discontinuity where one of the experts is indifferent.

Thus, we have shown that both the composition of the cabinet as well as the order of polling can have a profound impact on the information revealed in equilibrium.

VI. REBUTTAL

We now amend the basic model slightly to allow for *rebuttal*. Specifically, there is an extended debate in which, as in the previous sections, expert 1 sends a message m_1 and expert 2, having heard m_1 , sends a message m_2 . In the rebuttal stage, first expert 1 is allowed to rebut m_2 by sending a second message r_1 , and then finally, expert 2 is allowed to rebut m_1 and r_1 by sending a second message r_2 .

Opposing Biases. We show that when experts have opposing biases, it is possible to obtain full revelation as an equilibrium in the game with extended debate. This requires that there be no state in which both experts hold extreme views. This is a fairly weak condition. For example, in the uniform-quadratic case, it is satisfied provided that $|b_1| + |b_2| \leq 1/2$. Indeed, it is sufficient in that case that each expert be willing to convey *some* information when consulted alone.

PROPOSITION 5. Suppose that experts have opposing biases and there is no state in which both hold extreme views. Then there exists a fully revealing PBE of the game with rebuttal.

Without loss of generality, suppose that $b_1 < 0 < b_2$. Again, we illustrate the construction for quadratic loss functions.

Recall that in the semi-revealing equilibrium of the previous section, all states in which expert 2 does not hold extreme views are revealed, that is, states in the interval $[0, 1 - 2b_2)$. We can construct a similar equilibrium when the order of moves is reversed so that all states in which expert 1 does not hold extreme views are revealed, that is, states in the interval $(-2b_1, 1]$. If, in every state θ , at least one expert does not hold extreme views, a fully revealing PBE can be constructed by "patching" the two semi-revealing equilibria together.

In equilibrium, for all $\theta \leq 1 - 2b_2$, expert 1 sends $m_1 = \theta$, and expert 2 agrees by sending $m_2 = m_1$. In stage 3, expert 1 then "passes." For all $\theta > 1 - 2b_2$, expert 1 passes in the first stage, expert 2 then tells the truth $m_2 = \theta$, and expert 1 agrees in stage 3 by sending $r_1 = m_2$. All rebuttal messages r_2 by expert 2 are ignored.

Following $m_1 \leq 1 - 2b_2$ and any m_2 , the decision maker applies the self-serving criterion in deciding whether or not to adopt 2's recommendation, as in the semi-revealing equilibrium. However, if expert 1 "passes" in stage 1 by recommending $m_1 > 1 - 2b_2$, then following $m_2 > -2b_1$ and any r_1 , the decision maker applies the self-serving criterion in deciding whether or not to adopt 1's stage 3 recommendation, again as in the semi-revealing equilibrium. If expert 1 passes in stage 1 and $m_2 \leq -2b_1$, the decision maker simply adopts m_2 .

Clearly, the decision maker is following an optimal strategy. If $\theta < 1 - 2b_2$, the argument that no deviation $m_1 < 1 - 2b_2$ is profitable for 1 is the same as that for a semi-revealing equilibrium. If $\theta > -2b_1$ and expert 1 passes, the argument that both experts are optimizing is also the same as that for a semi-revealing equilibrium. If $\theta \leq -2b_1$ and expert 1 passes, then by sending the message $m_2 = \min(\theta + b_2, -2b_1)$ expert 2 can guarantee that his recommendation will be adopted, and indeed this is optimal. Since this is higher than θ , this deviation leaves expert 1 worse off.

Thus, there is a fully revealing PBE once the possibility of rebuttal is admitted. Again, while the construction above was for quadratic loss functions, it extends in a straightforward fashion to general utility functions.¹¹

11. The construction above shows, in fact, that it is enough that only expert 1 is allowed to rebut.

Like Biases. Extended debate and the possibility of rebuttal does *not* lead to full revelation when the experts have like biases.

PROPOSITION 6. Suppose that experts have like biases. Then there does not exist a fully revealing PBE of the game with rebuttal.

Proof. See Appendix. ■

The fact that extending the length of the debate does not lead to full revelation in the like bias case does not rely on there being only two rounds of debate. Indeed, the proof highlights the fact that one could add arbitrarily many rounds of (sequential) debate without affecting this result. A determination of the most informative equilibrium with more than one round of debate is beyond the scope of this paper.

VII. OTHER EXTENSIONS

Our assumption that both experts are perfectly informed about θ and that biases are commonly known ensures that any improvement in information from combining the advice of the experts arises solely from the strategic interaction. In practice, the information of experts is neither perfect nor identical. Hence, in addition to the strategic motives highlighted in this paper, information aggregation motives also influence the composition of cabinets. Thus, our model should be thought of as only a partial description of the problem of choosing a cabinet. Incorporating both motives would obviously enhance the realism of the model but would obscure the circumstances in which strategic interaction among the experts is helpful.

An alternative model is one in which experts send messages simultaneously. As shown in Section III, in such a model, the most informative PBE with like biases is full revelation. Thus, the introduction of a second expert has a dramatic effect on information transmission. We showed that this PBE does not survive when we consider a sequential model. It also does not survive if experts' information is noisy, although a characterization of the most informative PBE with noisy information remains an open question.

Finally, our analysis only concerns itself with cabinets consisting of two experts. Obviously, the sequential framework we adopt is not particularly conducive to exercises where more and

more experts are added. Nonetheless, we believe that the basic intuition that satisfying the incentive constraints of the most loyal agent leads to the most informative equilibrium in the like bias case will carry over into the n agent case. In the case of opposite bias, again it is the most loyal agent who determines the length of the revealing interval in our construction of a semi-revealing equilibrium. Thus, we expect that our construction would continue to be an equilibrium provided that the most loyal expert does not speak first. Whether this can be improved upon by combining the information of more experts also remains an open question.

APPENDIX: PROOFS

Proof of Proposition 1. Suppose not. Then there is a PBE in which the state is fully revealed, and thus for all θ , the equilibrium action $Y(\theta) = y^*(\theta)$. We first consider the case of opposing biases.

Case 1: Opposing biases. First, consider the subcase where $b_1 < 0 < b_2$.

Let $\bar{\theta} < 1$ be such that $y^*(\bar{\theta}, b_2) > y^*(1)$. Such a $\bar{\theta}$ exists since $b_2 > 0$.

Let $\theta \in (\bar{\theta}, 1)$. Since $b_1 < 0$, we have that $y^*(\theta, b_1) < y^*(\theta)$. Choose a $\theta' > \theta$ close enough to θ so that $y^*(\theta', b_1) < y^*(\theta)$. Suppose that m_1 and m_2 are the equilibrium messages in state θ . Since the equilibrium is fully revealing, $y(m_1, m_2) = y^*(\theta)$.

Let $m'_2 = \mu_2(\theta', m_1)$ be expert 2's best response to the message m_1 in state θ' . Then by definition, $U(y(m_1, m'_2), \theta', b_2) \geq U(y(m_1, m_2), \theta', b_2)$ and since $y(m_1, m_2) = y^*(\theta) < y^*(\theta, b_2) < y^*(\theta', b_2)$, $U_1(y(m_1, m_2), \theta', b_2) > 0$ and so $y(m_1, m'_2) \geq y^*(\theta)$.

Next observe that $y(m_1, m'_2) \geq y^*(\theta')$. Suppose that $y(m_1, m'_2) < y^*(\theta')$. Then by sending the message m_1 in state θ' expert 1 can induce the action $y(m_1, m'_2)$, and since $y^*(\theta', b_1) < y^*(\theta) \leq y(m_1, m'_2) < y^*(\theta')$, this is a profitable deviation for 1. This is a contradiction, and so $y(m_1, m'_2) \geq y^*(\theta') > y^*(\theta)$.

By the definition of a PBE, it must be the case that the out-of-equilibrium action $y(m_1, m'_2) \leq y^*(1) < y^*(\bar{\theta}, b_2)$.

Thus, we have deduced that $y^*(\theta) < y(m_1, m'_2) < y^*(\theta, b_2)$. Since $b_2 > 0$, $U(y^*(\theta), \theta, b_2) < U(y(m_1, m'_2), \theta, b_2)$. But this contradicts the assumption that $y^*(\theta)$ is an equilibrium action in state θ . Thus, full revelation cannot be an equilibrium.

The subcase where $b_2 < 0 < b_1$ is treated similarly.

Case 2: Like Biases. The proof for the case of like biases is analogous. ■

Proof of Lemma 1. In order to economize on notation, in what follows, we will denote $\theta - \epsilon$ by θ^- and $\theta + \epsilon$ by θ^+ .

Case 1. $b_1 \leq b_2$.

To establish (3), suppose the contrary; that is, suppose that $U(y^-, \theta, b_1) < U(y^+, \theta, b_1)$. Then by continuity, for all $\epsilon > 0$ small enough,

$$(5) \quad U(Y(\theta^-), \theta^-, b_1) < U(Y(\theta^+), \theta^-, b_1).$$

Now suppose that in state θ^- , expert 1 were to send the message $m_1^+ = \mu_1(\theta^+)$, and let m_2 be expert 2's best response to this off-equilibrium message in state θ^- so that

$$U(y(m_1^+, m_2), \theta^-, b_2) \geq U(y(m_1^+, m_2^+), \theta^-, b_2).$$

This implies that $y(m_1^+, m_2) \leq y(m_1^+, m_2^+)$ since otherwise we would have that $U(y(m_1^+, m_2), \theta^+, b_2) > U(y(m_1^+, m_2^+), \theta^+, b_2)$ contradicting the fact that $Y(\theta^+) = y(m_1^+, m_2^+)$ is the equilibrium action in state θ^+ .

But now since $y(m_1^+, m_2) \leq y(m_1^+, m_2^+)$ and expert 2 weakly prefers the former in state θ^- , the fact that $b_1 \leq b_2$ implies that expert 1 also weakly prefers the former. Thus, $U(y(m_1^+, m_2), \theta^-, b_1) \geq U(Y(\theta^+), \theta^-, b_1)$ and hence by (5)

$$U(y(m_1^+, m_2), \theta^-, b_1) > U(Y(\theta^-), \theta^-, b_1).$$

Thus, by sending the message m_1^+ in state θ^- expert 1 can induce an action that he prefers to the equilibrium action. This is a contradiction, and thus (3) holds.

To establish (4), again suppose the contrary; that is, $U(y^-, \theta, b_2) > U(y^+, \theta, b_2)$. Then since $b_1 \leq b_2$, $U(y^-, \theta, b_1) > U(y^+, \theta, b_1)$.

Then by continuity, for small enough $\epsilon > 0$,

$$U(Y(\theta^-), \theta^+, b_1) > U(Y(\theta^+), \theta^+, b_1)$$

and

$$U(Y(\theta^-), \theta^+, b_2) > U(Y(\theta^+), \theta^+, b_2).$$

Hence if, in state θ^+ , expert 1 were to send the message $m_1^- = \mu_1(\theta^-)$, expert 2 will induce an action $y(m_1^-, m_2)$ that is strictly

lower than $Y(\theta^+)$. This is a profitable deviation for 1 and hence a contradiction. Thus, (4) holds.

Case 2. $b_1 > b_2$.

The proof for this case is similar. If either (3) or (4) does not hold, then expert 1 has a profitable deviation. ■

Proof of Lemma 2. Let $\epsilon = \min_j \min_\theta [y^*(\theta, b_j) - y^*(\theta)] > 0$.

Suppose that $\theta' < \theta''$ are two states such that $Y(\theta') \equiv y' < y'' \equiv Y(\theta'')$. Then there exist m'_1, m'_2 satisfying $m'_1 = \mu_1(\theta')$, $m'_2 = \mu_2(\theta', m'_1)$ and $y(m'_1, m'_2) = y'$ and similarly for the double primes. We will argue that $y'' - y' \geq \epsilon$.

Suppose that $y'' - y' < \epsilon$.

Since for all θ , $Y(\theta) \in [y^*(0), y^*(1)]$, there exist σ', σ'' such that $y^*(\sigma') = y'$ and $y^*(\sigma'') = y''$. Clearly $\sigma' < \sigma''$.

CLAIM. $\sigma' \in Y^{-1}(y')$ and $\sigma'' \in Y^{-1}(y'')$.

Proof of Claim. Let $\underline{\theta} = \min Y^{-1}(y')$ and $\bar{\theta} = \max Y^{-1}(y')$. Then $y^*(\underline{\theta}) \leq y' \leq y^*(\bar{\theta})$. If $y' < y^*(\underline{\theta})$, then $U(y', \underline{\theta}) < U(y^*(\underline{\theta}), \underline{\theta})$, and since $U_{12} > 0$, for all $t \in [\underline{\theta}, \bar{\theta}]$, $U(y', t) < U(y^*(\underline{\theta}), t)$. If $y' > y^*(\bar{\theta})$, a similar argument holds.

Now since $y^*(\cdot)$ is increasing, $\underline{\theta} \leq \sigma' \leq \bar{\theta}$ and $Y(\cdot)$ is monotonic, $\sigma' \in Y^{-1}(y')$.

This establishes the claim. □

Now since $U_1(y', \sigma') = 0$, $U_{13} > 0$ implies that for $j = 1, 2$, $U_1(y', \sigma', b_j) > 0$ and since $y'' - y' < \epsilon$, $U_1(y'', \sigma', b_j) > 0$ also. Similarly, since $U_1(y'', \sigma'') = 0$, $U_{13} > 0$ implies that $U_1(y'', \sigma'', b_j) > 0$ and since $y' < y''$, $U_1(y', \sigma'', b_j) > 0$ also.

Now let $z' \neq y''$ be such that

$$U(y'', \sigma', b_2) = U(z', \sigma', b_2).$$

Likewise, let $z'' \neq y''$ be such that

$$U(y'', \sigma'', b_2) = U(z'', \sigma'', b_2).$$

Since $U_1(y'', \sigma'', b_2) > 0$ and $U_{11} < 0$, $U_1(z'', \sigma'', b_2) < 0$, and so $z'' > y''$. Next since $U_{12} > 0$, $U(y'', \sigma'', b_2) < U(z', \sigma'', b_2)$ and so $z'' > z'$.

Now in state σ' , if expert 1 sent the message m''_1 in lieu of m'_1 , then we claim that expert 2 could do no better than sending message m''_2 resulting in action y'' . This is because all actions in the interval (y'', z'') cannot be induced by expert 2 following m''_1 that is, there does not exist an m_2 such that $y(m''_1, m_2) \in (y'', z'')$.

If there were such a message, then y'' would not be the equilibrium action in state σ'' . Thus, following m''_1 , no action greater than y'' is preferred by expert 2 to y'' . Thus, if expert 1 sends the message m''_1 in state σ' , expert 2 will respond by sending the message m''_2 , thereby resulting in action y'' . This deviation is then profitable for expert 1. ■

Proof of Proposition 2. We give the proof for the uniform-quadratic case. However, one can show that the argument generalizes in a straightforward fashion when Assumption M is satisfied.

Suppose that a_1, a_2, \dots, a_{N-1} are points where the function Y is discontinuous. Let $a_0 = 0$ and $a_N = 1$. Define $b = \min\{b_1, b_2\}$. Lemma 1 implies that these points satisfy the system of inequalities: for $n = 1, 2, \dots, N - 1$,

$$(a_n + b) - \frac{a_{n-1} + a_n}{2} \leq \frac{a_n + a_{n+1}}{2} - (a_n + b),$$

which results in the following recursive system of inequalities:

$$a_n \leq \frac{n}{n + 1} a_{n+1} - 2nb.$$

Now let $\bar{a}_1, \bar{a}_2, \dots, \bar{a}_{N-1}$ be the solution to the corresponding system of equations. Then clearly we have that $a_1 \leq \bar{a}_1, a_2 \leq \bar{a}_2, \dots, a_{N-1} \leq \bar{a}_{N-1}$. We can now directly apply Theorem 4 of CS. This implies that the single expert equilibrium is informationally superior. ■

Proof of Proposition 4. Suppose, without loss of generality, that $b_1 < 0 < b_2$. We argue that Y is constant.

Consider two states, σ and τ , such that $\sigma < \tau$.

Suppose that $Y(\sigma) > Y(\tau)$. Then there exists a $\sigma' \leq \sigma$ such that $Y(\sigma') = Y(\sigma)$ and $Y(\sigma') \geq y^*(\sigma')$ since otherwise, there are no beliefs that the decision maker could hold that would rationalize the choice of $Y(\sigma)$. Notice that since expert 1 is an extremist, then $U(Y(\sigma'), \sigma', b_1) \leq U(y^*(\sigma'), \sigma', b_1) \leq U(y^*(0), \sigma', b_1)$. In state σ' if expert 1 were to send the message $\mu_1(\tau)$, expert 2 will induce some action $z \leq Y(\tau)$. To see this, notice that if expert 2 chose to induce an action $z > Y(\tau)$ in state σ' , then expert 2 would also prefer z to $Y(\tau)$ in state τ . Since for all $z \leq Y(\tau)$, $U(z, \sigma', b_1) > U(Y(\sigma'), \sigma', b_1)$, this is a profitable deviation for 1.

Suppose that $Y(\sigma) < Y(\tau)$. Then there exists a $\tau' \leq \tau$ such

that $Y(\tau') = Y(\tau)$ and $Y(\tau') \geq y^*(\tau')$. Since expert 1 is an extremist, she prefers any action $z < Y(\tau')$ to $Y(\tau')$ in state τ' . There also exists a state $\sigma'' \geq \sigma$ such that $Y(\sigma'') = Y(\sigma)$ and $Y(\sigma'') \leq y^*(\sigma'')$. Clearly, following $\mu_1(\sigma'')$, the highest inducible action by expert 2 is $Y(\sigma'')$ since in state σ'' , all higher actions are preferred to $Y(\sigma'')$. Thus, following the message $\mu_1(\sigma'')$, expert 2 will induce some action $z \leq Y(\sigma'')$. Since expert 1 is an extremist, all such actions are preferred to $Y(\tau')$; hence this is a profitable deviation.

We have shown that Y is constant and so involves only babbling. ■

Proof of Proposition 6. Suppose that there is a fully revealing equilibrium. Then in each state θ , the equilibrium action $Y(\theta)$ is $y^*(\theta)$.

Let θ' and θ'' be two states close together such that $\theta' < \theta''$. Suppose that (m'_1, m'_2, r'_1, r'_2) and $(m''_1, m''_2, r''_1, r''_2)$ are the equilibrium messages in the two states, respectively.

First, notice that following the messages m''_1 , m''_2 , and r''_1 , if $z > y^*(\theta'')$ is an action that expert 2 can induce via his rebuttal message r_2 , then we must have that $U(z, \theta'', b_2) \leq U(y^*(\theta''), \theta'', b_2)$. Since $U_{12} > 0$, this implies that $U(z, \theta', b_2) < U(y^*(\theta''), \theta', b_2)$. Since $b_2 > 0$, this means that if m''_1 , m''_2 , and r''_1 are sent in state θ' , then expert 2 cannot do better than to induce $y^*(\theta'')$ by sending r''_2 .

Second, notice that following the messages m''_1 and m''_2 , if $z > y^*(\theta'')$ is an action that expert 1 can induce (assuming that expert 2 plays his equilibrium strategy) via his rebuttal message r_1 , then we must have that $U(z, \theta'', b_1) \leq U(y^*(\theta''), \theta'', b_1)$. As before, this implies that $U(z, \theta', b_1) < U(y^*(\theta''), \theta', b_1)$. Since $b_1 > 0$, this means that if m''_1 and m''_2 are sent in state θ' , then expert 1 cannot do better than to induce $y^*(\theta'')$ by sending r''_1 .

Finally, notice that a similar argument shows that following m''_1 in state θ' , expert 2 can do no better than to send m''_2 .

Thus, we have shown that if expert 1 were to send the message m''_1 in state θ' , the resulting action would be $y^*(\theta'') > y^*(\theta')$. For θ' and θ'' close to each other, this is a profitable deviation for expert 1. ■

REFERENCES

- Austen-Smith, David, "Information Transmission in Debate," *American Journal of Political Science*, XXXIV (1990), 124–152.
- , "Interested Experts and Policy Advice: Multiple Referrals under Open Rule," *Games and Economic Behavior*, V (1993), 1–43.
- Baliga, Sandeep, Luis Corchon, and Tomas Sjöström, "The Theory of Implementation When the Planner Is a Player," *Journal of Economic Theory*, LXXVII (1997), 15–33.
- Banerjee, Abhijit, and Rohini Somanathan, "A Simple Model of Voice," *Quarterly Journal of Economics*, CXVI (2001), 189–228.
- Crawford, Vincent, and Joel Sobel, "Strategic Information Transmission," *Econometrica*, L (1982), 1431–1451.
- Dewatripont, Mathias, and Jean Tirole, "Advocates," *Journal of Political Economy*, CVII (1999), 1–39.
- Farrell, Joseph, "Meaning and Credibility in Cheap Talk Games," *Games and Economic Behavior*, V (1993), 514–531.
- Farrell, Joseph, and Matthew Rabin, "Cheap Talk," *Journal of Economic Perspectives*, X (1996), 103–118.
- Friedman, Eric, "Public Debate among Experts," mimeo, Northwestern University, 1998.
- Gilligan, Thomas, and Keith Krehbiel, "Asymmetric Information and Legislative Rules with a Heterogeneous Committee," *American Journal of Political Science*, XXXIII (1989), 459–490.
- Krishna, Vijay, and John Morgan, "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review* (2001), forthcoming.
- Matthews, Steve, Masahiro Okuno-Fujiwara, and Andrew Postlewaite, "Refining Cheap Talk Equilibria," *Journal of Economic Theory*, LV (1991), 247–273.
- Milgrom, Paul, and John Roberts, "Relying on the Information of Interested Parties," *RAND Journal of Economics*, XVII (1986), 350–391.
- Morgan, John, and Philip Stocken, "An Analysis of Stock Recommendations," mimeo, Princeton University, 1998.
- Morris, Stephen, "Political Correctness," *Journal of Political Economy* (2001), forthcoming.
- Ottaviani, Marco, and Peter Sorensen, "Information Aggregation in Debate," mimeo, University College, London, 1997.
- Shin, Hyun, "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, LXIV (1994), 253–264.
- Sobel, Joel, "A Theory of Credibility," *Review of Economic Studies*, LII (1985), 557–573.