Sequential Communication: Ally or Rival First?

Robert Evans[1] and Sönje Reiche[2]

This version: October 2013

Abstract

We analyze a cheap talk model with three possible decisions and two experts, one biased and one unbiased. Sequential public communication is always optimal for the decision maker. The optimal ordering of speaking depends on the conflict indices of the two experts, where the conflict index of an expert is his probability that there is a conflict conditional on a conflict being possible. In some cases, it is optimal to have the unbiased expert speak first in order to silence him. The decision maker may be better off if the biased expert knows all of the unbiased expert's information. He may also prefer to replace the unbiased expert with a second biased one.

Keywords: Experts, Cheap Talk

[1]St. John's College and University of Cambridge, UK. robert.evans@econ.cam.ac.uk
[2]Toulouse School of Economics (CNRS, GREMAQ)

1 Introduction

There are many situations in which a principal has more than one potential expert, or adviser, whom she can call on, each of whom has partial information which is relevant to the decision which the principal must make. Some of these advisers may be more biased than others, relative to the principal's preferences. How should she optimally consult these advisers? In particular, if there are two advisers, one of whom is 'friendly', in the sense that he has similar preferences to her, whereas the other's preferences are only partially aligned with hers, under what circumstances is it better to ask the friendly adviser to speak first, publicly, and under what circumstances is it better to ask the less friendly one to speak first? When is it optimal to consult them in private? How do the answers to these questions depend on the nature of the private information: for example, on whether the friendly adviser is less well-informed than the less friendly one?

The answers to these questions are relevant to the literatures on strategic communication and organizational design, in particular the design of procedures in committees. They are also relevant to the theory of voting since the model can also be interpreted as a model of sequential voting over a small number of alternatives.

We examine these questions in the context of a simple model in which the principal must choose between three decisions. The friendly adviser has private information about one of these potential decisions while the less friendly one has private information about the other two. Choosing the optimal decision for the principal requires knowing both advisers' private information. The principal cannot commit to a mechanism, and so her decision will be her optimal one given what she has learned from the advisers. She cannot use side-payments.

In the class of simple communication procedures which we consider it is always optimal to have the two advisers speak publicly in sequence. It turns out that the choice of optimal order of announcements depends on two belief parameters of the advisers, which we call the conflict indices of the good (i.e., friendly) adviser and the biased (i.e., less friendly) adviser respectively. Whether there are conflicting interests between the two advisers depends on the state of the world; a particular adviser's conflict index is his probability that there is a conflict, conditional on his information state being such that a conflict is possible. The relation between the conflict indices and the optimal procedure is somewhat subtle. There are two reasons why it might be strictly optimal to have the good adviser speak first - (a) to encourage the other adviser to reveal information by revealing information himself, and (b) to encourage the other adviser to reveal information by saying nothing. These correspond to two regions of the parameter space. Similarly, there are two reasons why it might be strictly optimal to have the less friendly adviser speak first - (a) because the likelihood that there is a conflict is low enough that he is willing to reveal his information, but only if he speaks first, and (b) that he will never reveal his information, whether first or second, while the friendly adviser will not reveal his information unless he speaks second.

Suppose the biased adviser (B) speaks first. If his conflict index is low then he will reveal his information, but not if his conflict index is high. This is because he knows that the unbiased, or good, adviser (G) will subsequently reveal the conflict if it exists, to his own detriment. In the case in which G speaks first the opposite is true - he will reveal his information if his conflict index is high, but not otherwise. This is because, if it is not high, he will want to pretend that there is no potential conflict (when in fact there is), in order to encourage B then to reveal his information.

The implication of this is that if B's conflict index is low then it is optimal to have him speak first, since then both advisers will reveal their information. If B's conflict index is too high for this, but G's index is high, then it is optimal to have G speak first, in which case he will reveal his information and, if in fact he reveals that there is no possibility of a conflict, B will then reveal his information. In effect, the motive in this case for having G speak first is that he can then signal, when appropriate, that it is safe for the other adviser to be truthful.

The second motive for having G speak first applies in the case when G's conflict index is too low for him to be truthful when speaking first, and B's conflict index takes an intermediate value. The idea is that this serves to silence G. He reveals nothing; when B speaks he is then willing to speak truthfully because he knows that nothing will be learned from G. The second motive for having B speak first applies when his conflict index is too high for him to reveal anything at all (whether he is first or second) and G's has an intermediate value. Placing G second enables him to speak freely, which he would not if he were placed first because his conflict index is not high enough.

We also consider a variation of the model in which the biased adviser is fully-informed - he knows all of the good adviser's information, in addition to his own private information. In this case the only motive for having the good adviser speak first is to silence him. Whether the decision maker is worse off with a fully-informed than a partially-informed biased adviser depends on the conflict indices; in particular, if the biased adviser's conflict index is high and the good adviser's is not, then the decision maker prefers to have a fully-informed biased adviser because she can extract his information, whereas she cannot if he is partially-informed.

Finally, we ask whether the decision maker is always better off having an adviser whose interests are fully aligned with her own, rather than a second biased adviser. We show via an example that she may not be. In the example, a biased adviser would tell her the truth, when speaking first, whereas an unbiased adviser would not.

Related Literature

The two closest papers to ours are Ottaviani and Sorensen (2000) and Krishna and Morgan (2001)[3]. In Ottaviani and Sorensen's paper, advisers are unbiased but motivated by career concerns and are imperfectly informed about the state of the world. The main insight is that with several advisers who give advice sequentially, herding will occur once the belief becomes too concentrated, since each adviser will want to agree with the current belief to appear better informed. In our paper herding does not arise since advisers have different pieces of information and are not motivated by career concerns. Nevertheless, experts learn from each other and since

[3]Ottaviani and Sorensen (2006) uses a more general signal structure and utility functions. They provide a full characterization of all equilibria but do not analyze the problem of several experts, so there are no results on scheduling.

learning affects their incentives for truthful revelation the way advice is scheduled is important. Ottaviani and Sorensen show that less reliable advisers should speak first, because they would otherwise herd and their information would be lost. In our model, Sometimes B should speak first, either because his own conflict index is low (and so he reports his information truthfully), or because G's conflict index is too high (and he can only be made to report truthfully if he goes second), and sometimes G should go first (to signal to B that it is safe to talk). Ottaviani and Sorensen also show that it is not always optimal to have better informed advisers since this might increase herding. We have a similar result albeit for a different reason: sometimes advisers are only willing to reveal their information if they are sufficiently ignorant about aspects of the state of the world that conflicts with their own interests.

Krishna and Morgan (2001) consider sequential information revelation by two perfectly informed, biased experts in an extension of Crawford and Sobel's (1982) cheap talk model (CS from now on). In contrast, we consider a problem where the decision maker has to choose between different discrete alternatives and our experts are only partially informed. As explained above, with partially informed advisers who learn from each other scheduling advice is an important consideration, while this is not an issue if experts are perfectly informed. Also, in our paper one of the experts is unbiased. If this expert was in addition perfectly informed, the decision maker could easily obtain all information by listening only to him. In our paper the unbiased adviser does not know the state of the world fully and he lies because he wants to influence the biased expert's advice in a certain direction. Krishna and Morgan show that full information revelation is possible provided that the two experts are biased in opposite directions. We show a related result: If the second adviser is biased in the opposite direction to B he might be more willing to reveal a potential conflict between B and DM, so that, counter intuitively, more information is transmitted than with an unbiased adviser. Krishna and Morgan also consider what happens when experts can speak twice and show that in this case full information revelation can be induced. We neglect this possibility but it would constitute an important extension of our paper.

In a related paper Austen-Smith (1993) compares sequential to simultaneous or joint referral of legislation by an uninformed House to two heterogeneous and partially informed committees. Under the 'open rule' the referral process is equivalent to a cheap talk game in which the committees serve as informed experts, and thus can be modelled as a version of CS. Again, this paper differs from ours since it considers a single dimensional decision space and information is not specific to a particular adviser. It can be optimal to have the more extreme (more biased) committee give advice first. In this context, the more biased committee separates only if the less biased one follows by conditionally separating, whereas in our paper, when the more biased expert B moves first, he either separates only if he is ignorant about G's information, or he always pools, in which case G should speak second to allow him to separate.

A problem that mirrors the one in Austen-Smith (1993) and Krishna and Morgan (2001) is considered by Wolinsky (2002). Here, the experts have the same biased preferences but have different pieces of information regarding a single dimensional decision. Here, information transmission works better if the decision maker groups experts together in different groups and has each group report a 'joint signal, rather than communicating with each expert individually or with all experts joint in a single group. Because advisers share the same preferences, how they communicate with each other once grouped together, is immaterial. In a problem where experts are both biased in favour of a status quo action and career concerned Bourjade and Jullien (2011) show that the optimal schedule of advice critically depends on whether or not the market can identify individual advisers. Sequential advice is optimal when advisers cannot be distinguished and simultaneous advice dominates in the opposite case. Dewatripont and Tirole (1999) provide a rationale for the existence of advocates, i.e. advisers, who have a vested interest in the decision(s) they report on, like the advisers in our paper. They consider the situation of an uninformed principal who needs advice on two possible alternatives to be implemented in lieu of a status quo. Advisers incur a cost for information collection and can either find hard information or no information. Effort is not observable and so advisers are paid based on the decision taken (this can be interpreted as short-hand for advisers with career concerns). Dewatripont and Tirole show that in their simplest setting (information cannot be concealed) it is more costly to use one adviser who reports on both activities rather than two who are asked to investigate one alternative each, so the principal 'creates' partisan advisers. In an extension, where advisers can find evidence both in favor and against an alternative and can decide to conceal this evidence, results are less clear cut. Nevertheless, it is shown that it can be optimal for the principal to hire two advisers who become either advocates (bias their reports in favor of their alternative) or prosecutors (bias their report in favor of the status quo). Their paper can thus be considered a precursor to ours. It explains why biased advisers exist, while we take them as given and ask how best to talk to them.

Our paper is also related to the paper by Che, Dessein and Kartik (2013) on pandering. In their paper an adviser has private information on a collection of alternative actions but always favours an action in this set over a status quo action. In our paper, two advisers hold information on different possible actions and either have a bias towards their controlled set of actions (adviser B) or no bias (adviser G). In Che, Dessein and Kartik the adviser biases his advice in the direction of conditionally more attractive looking actions even though his preferences are identical to the decision makers on those actions. This is reminiscent of the result in our paper that G does not reveal his information fully although he has identical preferences to those of the decision maker. He does so not because it makes his revelation more 'credible' as in Che, Dessein and Kartik, but because he strategically wants to influence the other expert's advice.

2 The Model

A decision-maker (DM) must choose one of a number of options. She is not fully informed about the merits of each option but can take advice from one or both of two agents who are better-informed than she is. One of these agents, the 'good' adviser, or G, has the same preferences as DM while the other, the 'biased' adviser, or B, has preferences which are only partially aligned with DM's. The potential options, or 'projects', can be divided into those which fall into the sphere of B and those which fall into the sphere of G. B is informed about the payoffs generated, both to himself and to DM, of the former projects, but less so about the latter ones, and vice versa for G. B's preferences among the former projects coincide with those which DM (and so G) would have if she were fully informed, but B always wants one of 'his' projects to be chosen in preference to any of the projects in G's sphere. G, on the other hand, always wants the best project for DM to be chosen.

Should the decision-maker in this type of setting listen to both advisers before taking her decision? If so, should she listen to their advice in private or in public? Does it matter in what order she consults them, and if so, under what conditions is each order optimal? If the biased adviser were instead fully informed and so knew all of G's private information as well as the information about his 'own' projects, what difference would that make to DM's optimal method of consultation? Finally, is DM necessarily better off by virtue of one of the advisers having preferences fully aligned with her own, or would it be better in some circumstances to have two advisers, both of whom had different preferences from DM's?

In order to address these questions we consider the simplest model which corresponds to the situation described above. There are three possible decisions, b_1, b_2 and g, of which DM must choose one. b_1 and b_2 are in B's sphere, while g is in G's. The players' payoffs from decisions b_1 and b_2 depend on a random variable s_b which has two possible values, 1 and 2. The payoffs from g depend on a random variable s_g which can take two values, h and l. The probability that $s_b = 1$ is denoted by β and the probability that $s_g = h$ is denoted by γ. s_b and s_g are independent. We will sometimes refer to the B-state being 1 (or 2) and similarly to the G-state being h or l.

Initially we consider the information structure in which each adviser has information which is unobserved by the other: B observes the B-state, i.e., the realization of s_b, but not the G-state and vice versa for G. Therefore each has some relevant information, about his own sphere, but neither can be said to be better informed than the other. Later we will consider the structure in which B is better informed than G:

G observes the G-state while B observes both states. In either case, DM observes neither s_b nor s_g.

The state-dependent utility of DM, and also of G, for decision i ($i \in \{b_1, b_2, g\}$) is $v_G(i, s_b, s_g)$. The corresponding utility for B is $v_B(i, s_b, s_g)$. Since the payoffs of all parties from decisions b_1 and b_2 depend only on the B-state, and the payoffs from g only on the G-state, $v_G(i, s, h) = v_G(i, s, l)$ for $i \in \{b_1, b_2\}, s \in \{1, 2\}$ and $v_G(g, 1, s) = v_G(g, 2, s)$ for $s \in \{h, l\}$; similarly for v_B.

We make the following assumption about the form of the payoffs:

Assumption 1

$$v_G(b_2, 2, .) > v_G(b_1, 2, .) > v_G(g, ., h) > v_G(b_1, 1, .) > v_G(g, ., l), v_G(b_2, 1, .);$$

$$v_B(b_2, 2, .) > v_B(b_1, 2, .) > v_B(b_1, 1, .) > v_B(b_2, 1, .) > v_B(g, ., h) = v_B(g, ., l).$$

The following example illustrates the structure of payoffs.

	b_1	b_2	g
1	4,4,4	3,0,0	(0,0),(4.5,3.5),(4.5,3.5)
2	5,5,5	7,7,7	(0,0),(4.5,3.5),(4.5,3.5)

Rows correspond to B-states and columns correspond to decisions. The first entry in a cell (i, j) is B's payoff $v_B(j, i, .)$, the second is G's and the third DM's. In the case of decision g there are two payoffs for a given agent and B-state; the first is the payoff if the G-state is h and the second the corresponding payoff for G-state l. That is, we assume that $v_G(g, i, h) = 4.5$ and $v_G(g, i, l) = 3.5$ for $i \in \{1, 2\}$. The only effect of the G-state is on DM's and G's payoff from decision g.

By Assumption 1, b_2 is better than b_1 in state 2 and vice versa in state 1, for all players. Also 2 is the better state in the sense that the payoffs in state 2, for each of b_1 and b_2, are higher than in state 1. Regardless of the value of s_b or s_g, g is the worst decision for B. For DM, on the other hand, g may or may not be optimal, but the information of both advisers is necessary to decide which. DM's payoff from decision g is high if $s_g = h$ and otherwise low. In the former case g is the optimal decision for DM if $s_b = 1$ but is inferior to both b_2 and b_1 if $s_b = 2$; on the other hand, if $s_g = l$ then g is inferior to b_1 regardless of B's information.

It is convenient to normalize the payoffs as follows

Assumption 2 $\quad v_G(b_2, 1, .) = v_B(g, ., h) = v_B(g, ., l) = 0.$

DM is indifferent between b_1 and b_2 if her B-state probability $\beta = \tilde{\beta}$, where

$$\tilde{\beta} = \frac{v_G(b_2, 2, .) - v_G(b_1, 2, .)}{(v_G(b_2, 2, .) - v_G(b_2, 1, .)) - (v_G(b_1, 2, .) - v_G(b_1, 1, .))},$$

and, if DM knew that $s_g = h$ then she would be indifferent between b_1 and g if $\beta = \beta^*$, where

$$\beta^* = \frac{v_G(b_1, 2, .) - v_G(g, ., h)}{v_G(b_1, 2, .) - v_G(b_1, 1, .)}.$$

We make the following assumption, which is equivalent to assuming that each of the three decisions b_1, b_2 and g is optimal for some belief β when the G-state is high (h).

Assumption 3 $\quad \beta^* > \tilde{\beta}.$

This implies that, when G's type (i.e., the G-state) is high, g is optimal if B-state 1 is likely ($\beta > \beta^*$), b_1 is optimal for intermediate values of β ($\beta \in (\tilde{\beta}, \beta^*)$) and b_2 is optimal if B-state 2 is likely ($\beta < \tilde{\beta}$). In our example above, $\tilde{\beta} = \frac{1}{3}$ and $\beta^* = \frac{1}{2}$.

Each player is an expected utility maximizer. It is assumed not to be possible to make money transfers. How should DM structure the communication between the two advisers and herself? There are many communication structures which she could use. She could consult the advisers privately or publicly. The advisers could send messages simultaneously or sequentially. There could be several rounds of communication. She could consult only one of them before taking the decision.

We restrict attention to games with at most two rounds of communication. Specifically, there are eight games which DM can choose from. Firstly, there are two public sequential message games, the *G-leader game* and the *B-leader game*. The G-leader game has the following rules. First, B and G observe the realizations of, respectively, s_b and s_g. Then G chooses a message from a finite set M and publicly announces it. B, after hearing G's message, chooses a message from M and publicly announces it. DM then chooses a decision from the set $\{b_1, b_2, g\}$. The B-leader game is identical except that the order of messages is reversed - B sends the first message and G, after observing it, sends the second. Secondly, there are two sequential private message games, in which an adviser's message is not observed by the other adviser. There are two simultaneous-move message games, one private and one public. Finally, DM can exclude one or other adviser and have a single message. It is easy to see that the two simultaneous-move games are equivalent to each other, and also to the two private sequential games. Here we analyze the (public) G-leader and B-leader games. In the next section we show that one of these two games is always optimal in the class which we consider.

In the G-leader game a pure strategy for G is a function $m_G : \{h, l\} \to M$ mapping his information type set to messages; a mixed strategy is a function $\mu_G : \{h, l\} \to \Delta(M)$, where $\Delta(M)$ is the set of probability distributions over M. A pure strategy for B is a function $m_B : \{1, 2\} \times M \to M$ mapping pairs consisting of B's information type and G's message to messages; a mixed strategy is a function $\mu_B : \{1, 2\} \times M \to \Delta(M)$. A pure strategy for DM is a function $d_D : M^2 \to \{b_1, b_2, g\}$ and a mixed strategy is a function $\delta_D : M^2 \to \Delta(\{b_1, b_2, g\})$. The strategies in the B-leader game are defined in an analogous way.

We analyze perfect Bayesian equilibrium (PBE) of these games. A PBE of the G-leader game consists of (i) a strategy profile, (ii) for each possible message $m \in M$ sent by G, a belief for B $\gamma_B(m) \in \Delta(\{h, l\})$ and (iii) for each possible pair of messages $(m_g, m_b) \in M^2$, a belief for DM $\sigma_D(m_g, m_b) \in \Delta(\{h, l\} \times \{1, 2\})$; such that each strategy is optimal given the beliefs and the beliefs are derived from Bayes' Rule after positive-probability messages. Note that DM's belief in a PBE after two positive-probability messages must be uncorrelated and her belief about G must be the same as B's belief, i.e., $\gamma_B(m_G)$. A PBE for the B-leader game is defined similarly.

As in any cheap talk game, there exist uninformative (babbling) equilibria. We take the view that these are implausible and so we use a refinement which excludes them. Suppose, for example, in the B-leader game, after B has sent his message, DM (and G) attach strictly positive probability to B being type 1. Then G's information is valuable and the only plausible equilibrium continuation is the separating one, in which G reveals his type truthfully. We assume that this is the equilibrium continuation which is played. Equally, in the G-leader game, after G has spoken, if there is an informative continuation which Pareto-dominates the babbling, or pooling, continuation for DM and both types of B, given their common belief about G's type, then we assume that the pooling equilibrium is not played. We denote by *communication equilibrium*, or *c-equilibrium*, a PBE which satisfies this refinement.

DM chooses a game, or communication structure, in order to maximize her equilibrium expected payoff. We assume that DM's preferred equilibrium of the chosen game is played. Initially we restrict attention to the two games described above.

We assume that DM is unable to commit to a decision rule. We also assume that she cannot use mediated communication. Some power to commit is, however, implied by the rules of the communication game. For example, in the G-leader game, as we shall see, DM would, in many cases, want to hear again from G after the two rounds of communication have taken place. The assumption therefore is that she has the power to commit to preventing any further communication.

3 Overview of Results

In this Section we provide an outline of our results. Fuller and more precise arguments are presented in the next Section.

If B is type 2 then there is no conflict of interest between B and DM: they both want decision b_2. Similarly there is no conflict if B is type 1 and G is type l: they both want decision b_1. Conflict only arises when B is type 1 and G is type h, in which case DM wants g and B wants b_1. Therefore we can interpret the parameter γ, the probability of type h, as B's degree of belief (assuming that he is type 1) that there is a conflict and so we refer to it as B's *conflict index*. Similarly, we can regard β as G's conflict index, since this is G's degree of belief (assuming he is type h) that B is type 1, hence that there is a conflict.

We show below that, in both games, when mixed strategy equilibria exist they are never optimal for DM. Therefore we restrict attention to pure strategy equilibria. Consider first the B-leader game. G always separates (if a c-equilibrium is being played) since, having identical payoffs to DM, he has no reason to lie. Either B separates or he pools. Is there a c-equilibrium in which he separates? Suppose B is type 1. If he reveals his type, then, with probability γ, G is type h and the decision will be g, and, otherwise, he is type l and the decision will be b_1. This gives him an expected payoff of $(1-\gamma)v_B(b_1, 1, .)$. If instead he mimics type 2 the decision will be b_2, giving him payoff $v_B(b_2, 1, .)$. Therefore a separating equilibrium exists if $\gamma \leq \hat{\gamma}$, where

$$\hat{\gamma} = 1 - \frac{v_B(b_2, 1, .)}{v_B(b_1, 1, .)}.$$

This equilibrium gives the maximum feasible payoff to DM, since all information is revealed. On the other hand, if $\gamma > \hat{\gamma}$, no information about B is revealed in a c-equilibrium. First B pools and then G separates. The decision is b_2 if $\beta < \tilde{\beta}$, b_1 if $\beta \in (\tilde{\beta}, \beta^*)$, and, if $\beta > \beta^*$, the decision is g if G is type h and b_1 if G is type l.

In effect, if B's conflict index γ is low enough then type 1 of B is confident enough to reveal his information when he speaks first - the expected benefit of getting his preferred project outweighs the risk of G's project being chosen. On the other hand, if his conflict index is above $\hat{\gamma}$ then his information cannot be extracted. Figure 1 shows which equilibria exist in each part of the parameter space.

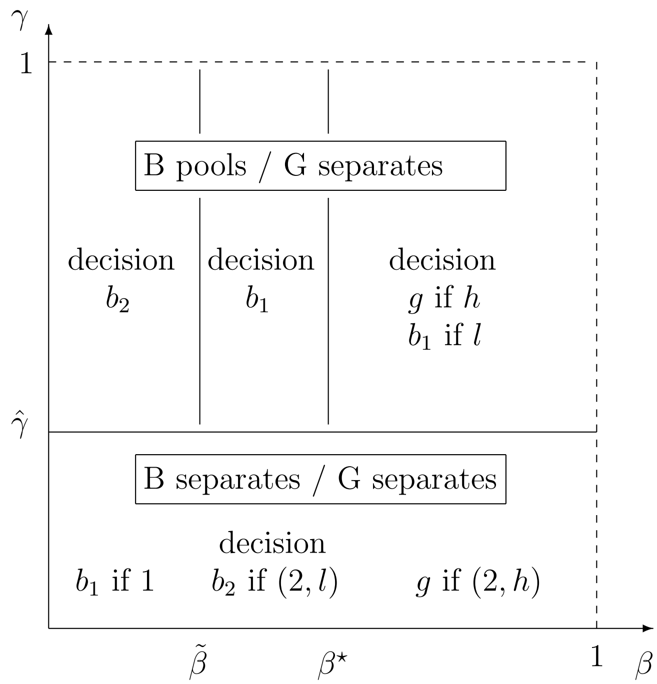

Figure 1: B-leader Game

Next, consider the G-leader game. Is it still true that G must separate in a c-equilibrium when his information is useful? Not necessarily - it depends on parameters. Consider a candidate c-equilibrium in which G separates. Once he reveals that he is type l, B will separate and the decision will be b_1 if type 1, b_2 if type 2. If, on the other hand, he is type h then there cannot be revelation by B: type 1 would mimic type 2 to get b_2 rather than g. B therefore pools. In that case the decision must be g (if it were, say, b_1, type h would certainly want to mimic type l to induce separation). Given this, will type h of G want to reveal his type? He will only do so if the probability of state 1 is high enough; if it is below a certain threshold β', the h

type would want to mimic the l type - the advantage of this is that it will extract B's information, bringing a better decision in state 2 (namely b_2) at the cost of a worse one in state 1 (b_1 rather than g). The relevant threshold is

$$\beta' = \frac{v_G(b_2, 2, .) - v_G(g, ., h)}{v_G(b_2, 2, .) - v_G(b_1, 1, .)}.$$

This equilibrium, which we call the Contingent Separating Equilibrium, exists if $\beta > \beta'$ but not if $\beta < \beta'$. Figure 2 illustrates.

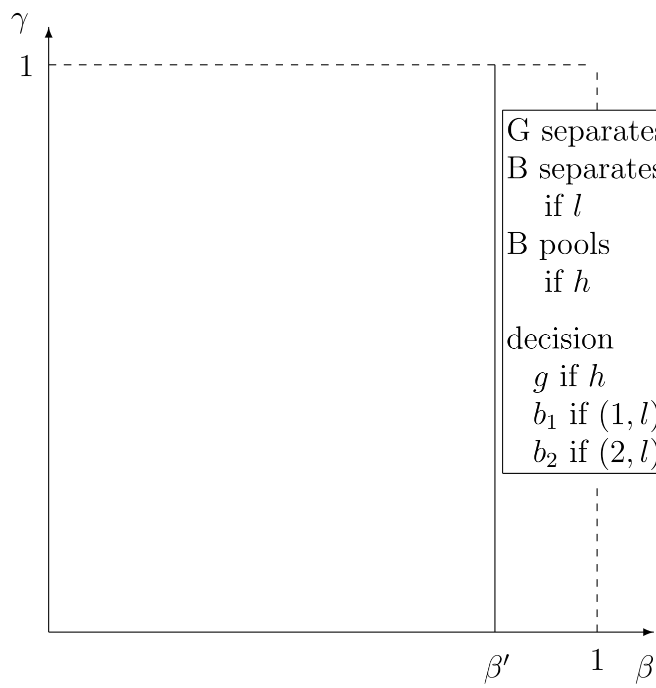

Figure 2: G-leader Game

The other possible pure strategy equilibria involve pooling by G. In this case, once G has spoken, will B want to separate? We saw above that when $\gamma > \hat{\gamma}$ B will not reveal his type if he knows that DM will discover G's information. If, on the other hand, DM learns nothing from G before taking the decision B is less concerned about revealing that he is type 1. Specifically, if γ is low enough that DM prefers b_1 to g in state 1, then B is indeed willing to reveal state 1 truthfully. The relevant threshold is

$$\tilde{\gamma} = \frac{v_G(b_1, 1, .) - v_G(g, ., l)}{v_G(g, ., h) - v_G(g, ., l)}.$$

Note that $\tilde{\gamma} > \hat{\gamma}$. Summarizing, there is an equilibrium in which G first declines to reveal his information and B then reveals his type if and only if $\gamma \leq \tilde{\gamma}$. We call this the Pooling/Separating Equilibrium.

In the third type of pure strategy equilibrium, the Pooling/Pooling Equilibrium, G pools and B then pools. This is a c-equilibrium if $\gamma > \tilde{\gamma}$. The decision is b_2 for low β and either g or b_1 for high β, depending on γ. It is also an equilibrium for $\gamma \leq \tilde{\gamma}$, but not a c-equilibrium. See Figure 3.

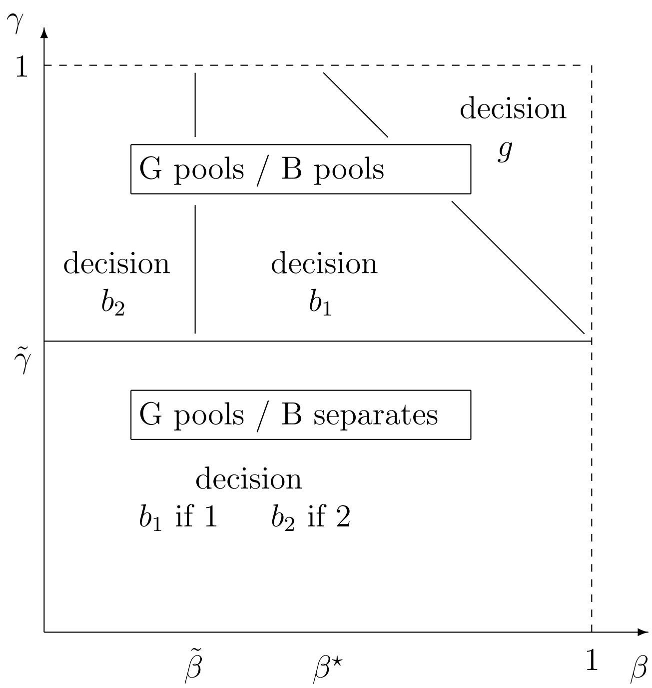

Figure 3: G-leader Game

In the G-game the Contingent Separating Equilibrium is optimal where it exists, i.e. for $\beta \geq \beta'$. It is clearly better than the Pooling/Pooling Equilibrium. As for the Pooling/Separating Equilibrium, the outcome is the same if G is type l, but if G is type h the Contingent Separating Equilibrium is better because

$$v_G(g, ., h) > \beta v_G(b_1, 1, .) + (1 - \beta) v_G(b_2, 2, .)$$

when $\beta > \beta'$.

Which of the two games should the decision maker choose? We can identify five regions of β/γ space, as shown in Figure 4.

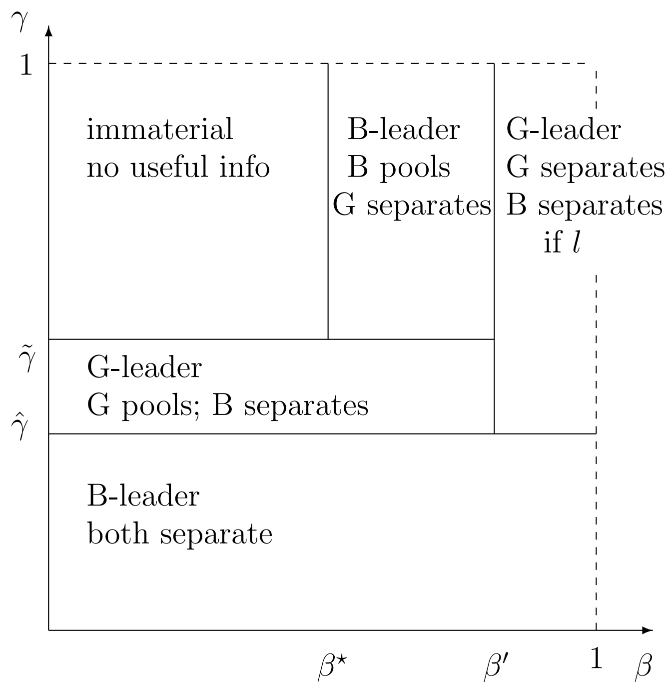

Figure 4

There are four principles which govern the choice of optimal communication structure.

Principle 1 *It is optimal to have B speak first when he is confident enough (his conflict index is low enough) to tell the truth.*

If $\gamma < \hat{\gamma}$ the B-leader game gives complete information revelation and so is optimal.

Principle 2 *One reason to have G speak first is to allow him to reveal to B that it is safe to speak truthfully.*

This is optimal when B's conflict index is not low ($\gamma > \hat{\gamma}$) and G's conflict index is high ($\beta > \beta'$), so that it is credible that G will be truthful when speaking first.

In this region DM should choose the G-leader game. The Contingent Separating equilibrium would then be played whereas, in the B-leader game, B would pool; in both cases, G's information is revealed but only in the former does B reveal anything.

Principle 3   A second reason to have G speak first is to silence him, making it safe for B to be truthful.

This applies when B's conflict index is neither low ($\gamma > \hat{\gamma}$) nor high ($\gamma < \tilde{\gamma}$), so that B could be willing to tell the truth and when, in addition, G's conflict index is not high ($\beta < \beta'$). The G-leader game is optimal in this case. The G-leader game gives a payoff of $\beta v_G(b_1, 1, .) + (1 - \beta) v_G(b_2, 2, .)$. The B-leader game in this region would give a lower payoff: if DM's decision, given type h, is g, this gives a lower payoff because $\beta < \beta'$, and, if DM would never play g DM would be better off obtaining B's information. Effectively, if it is not possible to exclude G completely, the best thing is to ask him to speak first and rely on him to reveal nothing, thereby giving B the confidence to reveal his information.

Principle 4   A second reason to have B speak first is that he cannot credibly be truthful and it is credible for G to tell the truth only if he speaks second.

If $\gamma > \tilde{\gamma}$ and $\beta \in (\beta^*, \beta')$, the B-leader game is optimal. In the G-leader game, no information would be revealed; in the B-leader game, B would pool, but G would separate, and G's information is valuable. Because G's conflict index is not high, if his information is useful, it is only credible for him to be truthful when speaking second, and B's is high so that it is never credible for B to tell the truth.

Finally, in the region in which $\gamma > \tilde{\gamma}$ and $\beta < \beta^*$, either game is optimal. Although G would reveal his information only in the B-leader game, his information is redundant because g would not be optimal even if he is the h type.

Optimality of Sequential Public Announcements

Suppose that, instead of making their announcements sequentially and publicly, the advisers make them simultaneously (or, as is equivalent, privately). In any equilibrium of this game which is optimal for DM G will tell the truth if his information is valuable. It is easy to see that B separates if $\gamma < \hat{\gamma}$ and pools if $\gamma > \hat{\gamma}$, and, therefore, that this game is outcome-equivalent to the B-leader game. If DM excludes the good adviser, i.e. makes her choice after a single message from B, then the outcome is the same as in the G-leader game in (i) the region in which $\gamma \in (\hat{\gamma}, \tilde{\gamma})$ and $\beta < \beta'$ and (ii) the region in which $\gamma > \tilde{\gamma}$ and $\beta < \beta^*$, and is otherwise worse. If she excludes the biased adviser then the outcome is the same as in the G-leader game in the region in which $\gamma > \tilde{\gamma}$ and $\beta < \beta'$, and is otherwise worse. This justifies our assertion in the previous section that sequential public announcements are always optimal.

4 Results

Here we provide more precise statements of the claims of the previous section.

Proposition 1 *In the B-leader game,*

(i) if $\gamma \leq \hat{\gamma}$ there exists a c-equilibrium in which B separates(tells the truth) and G then separates after either message; DM's expected payoff in this equilibrium is the maximum possible, namely $(1-\beta)[v_G(b_2, 2, .)] + \beta[\gamma v_G(g, ., h) + (1-\gamma)v_G(b_1, 1, .)]$.

(ii) if $\gamma > \hat{\gamma}$, all c-equilibria are payoff-equivalent to the following equilibrium. B pools (babbles) and G then separates after either message. The decision is b_2 if $\beta < \tilde{\beta}$, b_1 if $\beta \in (\tilde{\beta}, \beta^*)$, and, if $\beta > \beta^*$, b_1 if G is type l and g if G is type h.

Proof (i) Type 2 of B sends a message m_a, type 1 sends a message m_b. Type h of G sends m_h, type l sends m_l. DM chooses b_2 after m_a, for any message of G and after m_b chooses g given m_h and b_1 given m_l. DM has the appropriate point beliefs. Given these strategies, B is willing to tell the truth in state 1 if and only if $(1-\gamma)v_B(b_1, 1, .) \geq v_B(b_2, 1, .)$, i.e. $\gamma \leq \hat{\gamma}$. Optimality of the other strategies is easily checked. In each of the four states, DM's payoff is maximized. The proof of (ii) is in the Appendix. QED

Now consider the G-leader game.

Proposition 2 In the G-leader game any c-equilibrium is payoff-equivalent to one of five possible types of equilibrium.

(i) Pooling/pooling (PP) Equilibrium. In this equilibrium the two types of G pool and the two types of B pool after every message. DM chooses b_2 *if* $\beta \leq \tilde{\beta}$, b_1 *if* $\beta \in (\tilde{\beta}, \beta^*)$*, and, if* $\beta \geq \beta^*$, g *if*

$$pv_G(g, ., h) + (1-\gamma)v_G(g, ., l) \geq (1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)$$

and otherwise b_1*. This type of equilibrium exists if and only if* $\gamma > \tilde{\gamma}$*.*

(ii) Pooling/Separating (PS) Equilibrium. G pools and B separates after every message. DM chooses b_2 *if* 2 *and* b_1 *if* 1*. This exists if and only if* $\gamma \leq \tilde{\gamma}$*.*

(iii) Contingent Separating Equilibrium. G separates. B separates if G is type l *and pools if G is type* h*. DM chooses* g *if* h *and, if* l, b_2 *if* 2 *and* b_1 *if* 1*. This exists if and only if* $\beta \geq \beta'$*.*

(iv) G-semi-separating Equilibrium. G semi-separates and B plays a pure strategy. Type h *of G randomizes over a message* m_p *and a message* m_s*. Type* l *of G sends* m_s *with probability* 1*. After* m_p *B pools (babbles). After* m_s *B separates. This equilibrium exists only if* $\gamma > \tilde{\gamma}$ *and* $\beta \geq \beta'$*. DM's expected payoff is*

$$\left(\frac{p-\tilde{\gamma}}{1-\tilde{\gamma}}\right)v_G(g, ., h) + \left(\frac{1-\gamma}{1-\tilde{\gamma}}\right)[(1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)].$$

(v) G/B-Semi-separating Equilibrium. Each player semi-separates. Type h *of G randomizes over two messages* m_h *and* m_l*; type* l *sends* m_l *w.pr.1. After message* m_h *B pools (babbles); after message* m_l *B semi-separates - type* 2 *of B sends some message* m_a *with probability 1 and type* 1 *of B randomizes between* m_a *and some message* $m_b \neq m_a$*. This exists only if* $\beta \geq \beta^*$ *and* $\gamma \geq \tilde{\gamma}$*. DM's expected payoff in her most preferred equilibrium of this kind is the same as that of the equilibrium in (iv).*

Proof: In Appendix.

Next, we ask how the conflict index parameters γ and β determine DM's choice of optimal communication structure.

Proposition 3 *The B-leader game is optimal for DM if $\gamma < \hat{\gamma}$ (B and G both separate) and if $\gamma > \tilde{\gamma}$ and $\beta \in (\beta^*, \beta')$ (B pools and G then separates). The G-leader game is optimal if $\gamma > \hat{\gamma}$ and $\beta > \beta'$ (G separates; B pools if G is h and separates if G is l) and if $\gamma \in (\hat{\gamma}, \tilde{\gamma})$ and $\beta < \beta'$ (G pools and B separates). If $\gamma > \tilde{\gamma}$ and $\beta < \beta^*$ then both games are optimal.*

Proof The two types of mixed strategy equilibrium of the G-leader game, described in Proposition 2(iv) and 2(v) give the same payoff to DM. In the parameter region in which they may exist ($\gamma > \tilde{\gamma}$ and $\beta > \beta^*$) there is an equilibrium of the B-leader game in which, by Proposition 1(ii), DM gets the payoff

$$\gamma v_G(g, ., h) + (1 - \gamma)[(1 - \beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)]$$

which is greater than the mixed strategy equilibrium payoff because $v_G(g, ., h) > (1 - \beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)$ for $\beta > \beta^*$. Therefore we can ignore the mixed strategy equilibria.

If $\gamma < \hat{\gamma}$ the B-leader game equilibrium in which B separates and G then separates (Proposition 1(i)) has payoff $(1 - \beta)v_G(b_2, 2, .) + \beta[\gamma v_G(g, ., h) + (1 - \gamma)v_G(b_1, 1, .)]$. The equilibria of the G-leader game in this region are the contingent separating (if $\beta \geq \beta'$), which is worse since $v_G(b_2, 2, .) > v_G(g, ., h)$, and the strong separating, which is worse since $v_G(g, ., h) > v_G(b_1, 1, .)$.

If $\beta > \beta'$ and $\gamma > \hat{\gamma}$ then the contingent separating equilibrium of the G-leader game has payoff $\gamma v_G(g, ., h) + (1 - \gamma)[(1 - \beta)v_G(b_2, 2, .) + \beta v_G(b_1, 1, .)]$; the only c-equilibrium of the B-leader game for this region (B pools, G separates, b_1 if l, g if h), is worse since $v_G(b_2, 2, .) > v_G(b_1, 2, .)$.

If $\gamma \in (\hat{\gamma}, \tilde{\gamma})$ and $\beta < \beta'$ the strong separating equilibrium of the G-leader game gives payoff $(1-\beta)v_G(b_2, 2, .)+\beta v_G(b_1, 1, .)$. In this region the best equilibrium of the B-leader game gives either b_2 for sure (worse since $v_G(b_2, 1, .) < v_G(b_1, 1, .)$), b_1 for sure (worse since $v_G(b_1, 2, .) < v_G(b_2, 2, .)$) or (if $\beta > \beta^*$) b_1 if l and g if h, giving payoff

$\gamma v_G(g, ., h) + (1-\gamma)[(1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)]$. But $v_G(b_2, 2, .) > v_G(b_1, 2, .)$ and, since $\beta < \beta'$, $(1-\beta)v_G(b_2, 2, .) + \beta v_G(b_1, 1, .) > v_G(g, ., h)$, so this is worse.

If $\gamma > \tilde{\gamma}$ and $\beta \in (\beta^*, \beta')$ then the B-leader game has an equilibrium in which B pools and G separates, giving payoff $\gamma v_G(g, ., h) + (1-\gamma)[(1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)]$. The best equilibrium in the G-leader game (strong pooling) has payoff $max\{\gamma v_G(g, ., h) + (1-\gamma)v_G(g, ., l), (1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)\}$. The first bracketed expression is worse since $v_G(b_1, 2, .) > v_G(b_1, 1, .) > v_G(g, ., l)$; the second is worse since $v_G(g, ., h) > (1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)$ when $\beta > \beta^*$.

Finally, when $\gamma > \tilde{\gamma}$ and $\beta < \beta^*$, in the G-leader game only (i) applies, giving $max\{(1-\beta)v_G(b_2, 2, .) + \beta v_G(b_2, 1, .), (1-\beta)v_G(b_1, 2, .) + \beta v_G(b_1, 1, .)\}$, the same as in the B-leader game. QED

5 Discussion and Extensions

A Fully Informed Biased Adviser

In the previous sections each of the two advisers had some information which was not known to the other. Suppose now that the biased adviser is fully-informed. The good adviser, as before, knows the value of s_g; the biased adviser, however, knows the value of both s_b and s_g. We ask here what difference this makes to the decision-maker's choice of optimal communication protocol, and whether, or under what conditions, the decision-maker is worse off when the biased adviser is fully-informed than when, as in the previous section, he is partially-informed.

B now has four possible types: $\{1h, 1l, 2h, 2l\}$. Consider first the G-leader game. It is not difficult to see that this is essentially equivalent to the G-leader game of the previous sections. Once G has spoken, B has no incentive to reveal the G-state, if it has not been revealed, and has the same incentive as in our original G-leader game to reveal the B-state. DM's best c-equilibrium is as before. There are three regions: (i) $\beta > \beta'$, (ii) $\beta < \beta', \gamma < \tilde{\gamma}$, and (iii) $\beta < \beta', \gamma > \tilde{\gamma}$ (see Figures 2 and 3), and the profiles in each region are as before.

However, the B-leader game is now different. It can be shown that we can restrict attention to pure strategy equilibria. If $s_g = h$ then B will not reveal s_b, since, after he reveals 1, G would reveal h. The two possibilities are (a) that he pools if $s_g = h$ and separates if $s_g = l$, and (b) he pools in both cases. There is always a c-equilibrium of type (a). After B speaks, G separates. If G is type l the decision is b_2 if 2 and b_1 if 1. If type h, the decision is b_2 for $\beta < \tilde{\beta}$, b_1 for $(\tilde{\beta}, \beta^*)$, and g for $\beta > \beta^*$. The strategies are equivalent to those of the Contingent Separating Equilibrium of our original game; the difference is that it is now an equilibrium for all parameters. An equilibrium of type (b) is clearly worse than this for DM, so we ignore it.

Which communication protocol should the decision-maker select? In Region (i), it makes no difference because the Contingent Separating Equilibrium will apply in both cases. In Region (ii), the G-leader game is optimal. This gives a payoff of $(1-\beta)v_G(b_2, 2, .) + \beta v_G(b_1, 1, .)$. The B-leader game gives the same payoff in state l, but a strictly worse one in state h, since the decision is the same for 2 as for 1, and $v_G(g, ., h) < (1-\beta)v_G(b_2, 2, .) + \beta v_G(b_1, 1, .)$ because $\beta < \beta'$. In Region (iii), the B-leader game is optimal since no information is revealed in the G-leader game.

This gives us two further principles.

Principle 5 If the biased adviser if fully informed, the only motive for having the good adviser speak first is to silence him.

In the partial information case, for some parameters, it was better to have G speak first so as to tell B that it was safe to tell the truth. This no longer applies, since there is nothing to teach B. In particular, if it is feasible to exclude G there is now no positive reason to have G speak first.

Principle 6 If B's conflict index is low, DM is worse off with a fully informed than a partially informed biased adviser; if it is high, and G's is not, she is better off with a fully informed biased adviser.

For low γ, DM is worse off because it is no longer possible to extract all information. For $\gamma > \tilde{\gamma}$ and $\beta < \beta'$ it is possible to extract information from a fully-informed B, but not a partially-informed one, so DM is better off.

Two Biased Advisers

Here we show, via an example, that it may be that the decision maker is better off having two advisers whose preferences are not aligned with her own than having one biased (misaligned) and one good (aligned).

The example is the same as the one in Section 3 except that, instead of the good adviser G, there is a second biased adviser, $B2$. $B2$'s information is the same as G's was, and his preference function v_{B2} is the same as v_G except that $v_{B2}(g, ., h) = 5.5$, so that, in G-state h, he values decision g higher than DM does.

	b_1	b_2	g
1	4,4,4	3,0,0	(0,0),(5.5,3.5),(4.5,3.5)
2	5,5,5	7,7,7	(0,0),(5.5,3.5),(4.5,3.5)

We take the case in which $\gamma > \tilde{\gamma} = \frac{1}{2}$ and $\beta \in (\beta^*, \beta') = (\frac{1}{2}, \frac{5}{6})$.

In the model of Section 3, in this parameter region (see Figure 4), the optimal game is the B-leader game and, in equilibrium, B pools, G separates and DM's payoff is $\gamma v_G(g, ., h) + (1-\gamma)(\beta v_G(b_1, 1, .) + (1-\beta)v_G(b_1, 2, .)) = 4.5\gamma + (1-\gamma)(5-\beta))$.

Now suppose that $B2$ speaks first. The following is a c-equilibrium. $B2$ separates; B separates if $B2$ reveals as l and pools if he reveals as h. DM has the appropriate beliefs and chooses g if h, b_1 if $(l, 1)$ and b_2 if $(l, 2)$. It is straightforward to check that this is an equilibrium. In particular, type h of $B2$ is willing to tell the truth because this gives 5.5 whereas lying would give him $\beta(4) + (1-\beta)(7) < 5.5$ for $\beta > \frac{1}{2}$.

DM's expected payoff is

$$\gamma(4.5) + (1-\gamma)[\beta(4) + (1-\beta)(7)] > \gamma(4.5) + (1-\gamma)(5-\beta)$$

and so DM is strictly better off than before. Somewhat counter-intuitively, it is possible to get the adviser to tell the truth when he does not share the decision maker's preferences but not when he does.

Appendix

Lemma 1 Take any c-equilibrium, either in the G-leader or the B-leader game, in which B semi-separates (i.e. in which B randomizes and the equilibrium distributions over decisions are different for the two types). Then in this equilibrium type 2 *of B sends some message* m_2 *with probability 1 and type* 1 *of B randomizes between* m_2 *and some message* $m_1 \neq m_2$*. When* m_2 *is sent there is strictly positive probability of decision* b_2 *and when* m_1 *is sent there is zero probability of decision* b_2*.*

Proof Suppose that, in some c-equilibrium of either game, both types of B are indifferent between two messages m_1 and m_2. Let the equilibrium distribution over decisions after m_i $(i = 1, 2)$ is sent be $\psi^i = (\psi_1^i, \psi_2^i, \psi_3^i)$. Then, by indifference of the two types,

$$v_B(b_2, 2, .)\psi_1^1 + v_B(b_1, 2, .)\psi_2^1 = v_B(b_2, 2, .)\psi_1^2 + v_B(b_1, 2, .)\psi_2^2$$

and

$$v_B(b_2, 1, .)\psi_1^1 + v_B(b_1, 1, .)\psi_2^1 = v_B(b_2, 1, .)\psi_1^2 + v_B(b_1, 1, .)\psi_2^2,$$

which is not possible unless $\psi^1 = \psi^2$.

Let m_1 and m_2 be equilibrium messages sent with strictly positive probability by type 2. Then type 1 strictly prefers, say, m_1. This implies that m_2 reveals type 2 and so must be followed by decision b_2 with probability 1. This gives 2's highest possible payoff, so m_1 must also be followed by b_2 for sure. Hence 2 will not randomize over two messages which lead to different decision distributions and, w.l.o.g, 2 sends some message m_2 w.pr.1. Any other equilibrium message reveals 1 so, w.l.o.g, the randomization takes the given form. Since m_1 reveals B to be type 1, it must be followed by b_1 or g, since b_2 is dominated for DM if $s_b = 1$. If m_2 is followed only by b_1 or g then both the 2 and the 1 type would prefer the message giving the higher probability of b_1, so the equilibrium would be uninformative about B, hence there must be positive probability of b_2. QED

Lemma 2 *In the B-leader game, there is no c-equilibrium in which B semi-separates if $\gamma > \hat{\gamma}$.*

Proof Suppose that there is an equilibrium in which B semi-separates. By Lemma 1, it is optimal for type 1 to reveal his type by sending m_1. G must then reveal his type. If G is type h, the decision is g, and if l, the decision is b_1. This gives 1 an expected payoff of $(1-\gamma)v_B(b_1, 1, .)$ which is strictly less than $v_B(b_2, 1, .)$ since $\gamma > \hat{\gamma}$. Suppose instead that 1 sends message m_2. If DM's updated belief $\beta_G(m_2) > \tilde{\beta}$ then b_2 is strictly worse for DM than b_1, which contradicts Lemma 1. Hence $\beta_G(m_2) \leq \tilde{\beta} < \beta^*$. This implies that DM strictly prefers b_1 to g (regardless of whether G announces h or l) and so only plays b_2 or b_1 after m_2. Type 1 then gets a payoff of at least $v_B(b_2, 1, .)$, which is strictly greater than his payoff from m_1, contradicting the supposition that he randomizes between m_2 and m_1. QED

Proof of Proposition 1(ii) By Lemma 2, B does not separate if $\gamma > \hat{\gamma}$, hence either pools or separates. If he separates, G must then separate after type 1 is revealed and the equilibrium must be as in (i). This is not an equilibrium since $\gamma > \hat{\gamma}$, so B must pool. G must then separate. It is easy to see that this is an equilibrium. DM's optimal decisions are as given since DM learns nothing about B. QED

Lemma 3 *In the G-leader game, if there is a c-equilibrium in which B semi-separates then $\gamma > \tilde{\gamma}$. Type h of G randomizes over two messages m_h and m_l; type l sends m_l w.pr.1. After message m_h B pools (babbles); after message m_l B semi-separates as in Lemma 1.*

Proof In the G-leader game, if B semi-separates after some message m from G then the updated belief $\gamma_B(m)$ must be $\tilde{\gamma}$. Otherwise, by Lemma 1, when B sends m_1, DM either plays b_1 w.pr.1 or plays g w.pr.1. In the first case, 1 would strictly prefer to play m_1 and in the second would strictly prefer m_2. This contradicts Lemma 1. Take an equilibrium as described in the statement of the Lemma. W.l.o.g., assume that there is a single message m such that $\gamma_B(m) = \tilde{\gamma}$. Call this message m_1. For any message $m_2 \neq m_1$, either $\gamma_B(m_2) > \tilde{\gamma}$ or $\gamma_B(m_2) < \tilde{\gamma}$. Suppose $\gamma_B(m_2) < \tilde{\gamma}$. Then

DM strictly prefers b_1 to g whether in state 2 or 1. Hence there is an equilibrium (and this is best for DM) in which B separates, since there is no danger of g. This would give type l of G his best possible payoff, so he would strictly prefer m_2 to m_1, in which there is only semi-separation. Therefore m_1 reveals type h, which contradicts $\gamma_B(m_1) = \tilde{\gamma}$. Hence $\gamma_B(m_2) > \tilde{\gamma}$. This implies that the prior $\gamma > \tilde{\gamma}$. Also, w.l.o.g., there are two messages, m_h and m_l; l plays m_l for sure, h randomizes over m_h and m_l; $\gamma_B(m_l) = \tilde{\gamma}$ and $\gamma_B(m_h) = 1$. After m_l, B semi-separates as in Lemma 1; after m_h, B does not semi-separate. He does not separate, since type 1 would not want to reveal his type and get g. Hence he must pool. QED

Lemma 4 *Consider a two-player game in which B sends a message to DM and DM takes a decision. The prior belief of both players that $s_g = h$ is γ. Let*

$$\tilde{\gamma} = \frac{v_G(b_1, 2, .) - v_G(g, ., h)}{v_G(b_1, 2, .) - v_G(b_1, 1, .)}.$$

A separating c-equilibrium exists if and only if $\gamma \leq \tilde{\gamma}$ and a pooling c-equilibrium exists if and only if $\gamma > \tilde{\gamma}$.

Proof Take a candidate separating equilibrium. After B is revealed as 2, DM must play b_2. After B is revealed as 1, DM plays g if $\gamma > \tilde{\gamma}$. In that case, type of 1 would prefer to deviate and mimic type 2, so this is not an equilibrium. If $\gamma \leq \tilde{\gamma}$ and DM plays b_1 after 1 is revealed, it is an equilibrium. Take a candidate pooling (babbling) equilibrium. If $\gamma \leq \tilde{\gamma}$ then, as just argued there is also a separating equilibrium, so this fails to satisfy our refinement test. QED.

Proof of Proposition 2 By Lemma 3, the only equilibrium in which B mixes has the form described in (iv) and (v). In any other equilibrium, B plays a pure strategy after every message, hence must either (a) pool after every equilibrium message, or (b) separate after every equilibrium message, or (c) pool after some equilibrium messages and separate after others. Take these possibilities in turn.

(a) Suppose that B pools after every message. Then, by Lemma 4, $\gamma_B(m) > \tilde{\gamma}$ for all m; hence $\gamma > \tilde{\gamma}$. If $\beta < \beta^*$ then DM's decision is either b_2 or b_1, and independent of G's message, so the equilibrium is outcome-equivalent to one in which G pools. If $\beta > \beta^*$ then DM chooses either g or b_1 after every equilibrium message by G. Suppose different messages lead to different distributions over g and b_1. Then, since h strictly prefers g to b_1 and vice versa for l, they must separate. But this contradicts the fact that $\gamma_B(m) > \tilde{\gamma}$ for every equilibrium m. Hence the distribution is invariant to G's message and we can again assume that G pools. This gives the equilibrium described in (i). The expressions for DM's payoff follow because DM chooses b_2 on $(0, \tilde{\beta})$, b_1 on $(\tilde{\beta}, \beta^*)$ and either b_1 or g on $(\beta^*, 1)$.

(b) Suppose that B separates after every equilibrium message. By Lemma 4, $\gamma_B(m) \leq \tilde{\gamma}$ for every equilibrium message m. Hence $\gamma \leq \tilde{\gamma}$. If there is a higher probability of g after some messages than others then l and h would separate, contradicting the fact that $\gamma_B(m) \leq \tilde{\gamma}$ for every equilibrium message m. Hence DM's choice is independent of G's message (b_2 if 2 and some distribution over g and b_1 if 1), so we can assume that G pools. Lemma 5 implies that this is an equilibrium if $\gamma \leq \tilde{\gamma}$. This gives the equilibrium described in (ii). DM's expected payoff follows because if she plays g with positive probability then she is indifferent between g and b_1.

(c) Suppose that B separates after some messages and pools after others. Assume first that G plays a pure strategy, i.e. one type plays a message m_s which leads to separation and the other a message m_p which leads to pooling. A continuation equilibrium in which B pools after G has revealed as type l is dominated by a separating continuation and so fails our refinement. Therefore l sends m_s, h sends m_p; B separates after m_s and pools after m_p; after B reveals as l, DM plays b_2 if 2 and b_1 if 1. Suppose that g is played with zero probability after G reveals as h. Then type h strictly prefers to play m_s. Hence g is played with positive probability after G sends m_p and B pools. Therefore h's equilibrium payoff is $v_G(g, ., h)$ (if DM randomizes between g and b_1 at this point then she, and hence h, must be indifferent between g and b_1). This is an equilibrium as long as h prefers m_p to m_s, i.e. if

$$v_G(g, ., h) \geq (1 - \beta)v_G(b_2, 2, .) + \beta v_G(b_1, 1, .)$$

or $\beta \geq \beta'$. Note that $\beta' > \beta^*$. This gives the equilibrium described in (iii).

Now suppose that G plays a mixed strategy. At least one type of G must send both pooling (i.e. a message which leads to B pooling) and separating messages, otherwise the equilibrium is equivalent to the one just described. Take a pair of equilibrium messages for G, one pooling, m_p, and one separating, m_s, such that at least one is in the support of both h and l. If g is never subsequently played after either m_p or m_s then both types of G would strictly prefer m_s, contradicting the fact that m_p is an equilibrium message. There are two cases to consider: (1) Only b_2 or b_1, or a mixture of the two, is played after m_p; (2) g is played with positive probability after m_p.

Case (1). g is played with positive probability after m_s, so if m_s is weakly better for l than m_p then h strictly prefers m_s to m_p since their payoffs differ only if g is played and h derives a higher payoff from g than l does. If h does not send m_p then m_p reveals type l; however, m_p is followed by pooling by B, which violates our refinement condition. Therefore h sends both m_s and m_p and l only sends m_p. m_s reveals G as h. Therefore, after B has then revealed as 1, DM plays g for sure; however, 1 would get b_2 by mimicking 2, which he prefers, a contradiction.

Case (2). g is played with positive probability after m_p. Therefore $\beta \geq \beta^*$. There are two cases: 2(i), in which g is not played after m_s, and 2(ii), in which g is played with positive probability after m_s. Case 2(i): after m_s DM plays b_2 if B is 2 and b_1 if 1. l strictly prefers this to a continuation in which B pools. Therefore m_p reveals h. h gets a payoff of $v_G(g, ., h)$ from m_p (since $\beta \geq \beta^*$ only b_1 or g could be played and b_1 only if $\beta = \beta^*$ in which case h is indifferent between b_1 and g). h is indifferent between m_s and m_p, so

$$v_G(g, ., h) = (1 - \beta)v_G(b_2, 2, .) + \beta v_G(b_1, 1, .)$$

i.e. $\beta = \beta'$, and this equilibrium is payoff-equivalent to the pure strategy equilibrium described in (iii). Case 2(ii): g is played after m_s, if B is type 1. 1 will not separate if g is played with probability 1, so DM randomizes between g and b_1 is B revealed as 1. Therefore $\gamma_B(m_s) = \tilde{\gamma}$. Therefore both types send m_s, otherwise $\gamma_B(m_s)$ would be either zero or 1. g is played with positive probability after m_p, so it cannot be that

m_p reveals l, so either l sends m_s and h sends both m_s and m_p, or both send both messages.

Suppose l sends m_s and h sends both m_s and m_p. Let ϕ be the probability that DM plays b_1 after m_s is followed by B revealing as 1, and $1-\phi$ be the probability that DM plays g. m_p reveals G as h and gives h payoff $v_G(g,.,h)$ (as above). After m_s is followed by B revealing as 1, DM randomizes over g and b_1, hence her belief $\gamma_B(m_s)=\tilde{\gamma}$. This implies that the probability that h sends m_p is

$$\psi = \frac{\gamma-\tilde{\gamma}}{\gamma(1-\tilde{\gamma})}.$$

Since h is indifferent between m_p and m_s it must be that

$$v_G(g,.,h) = (1-\beta)v_G(b_2,2,.) + (1-\beta)[\phi v_G(b_1,1,.) + (1-\phi)v_G(g,.,h)],$$

hence

$$\phi = \frac{(1-\beta)(v_G(b_2,2,.)-v_G(g,.,h))}{\beta(v_G(g,.,h)-v_G(b_1,1,.))}.$$

The condition $\phi \leq 1$ gives $\beta \geq \beta'$. This gives the equilibrium in (iv). For it to be an equilibrium we also require $\phi \geq (v_B(b_2,1,.))(v_B(b_1,1,.))^{-1}$ so that 1 wants to separate after m_s. DM's payoff is (since DM is indifferent between g and b_1 when 1 reveals)

$$\gamma\psi v_G(g,.,h) + [\gamma(1-\psi)+(1-\gamma)][(1-\beta)v_G(b_2,2,.)+\beta v_G(b_1,1,.)],$$

which reduces to the expression given.

Now suppose that B plays a mixed strategy. By Lemma 3, the form of the equilibrium must be as follows. Type h of G randomizes between a message m_p (probability ψ) and a message m_{ss} (probability $1-\psi$). Type l sends m_{ss} for sure. After m_p B pools. After m_{ss} B semi-separates - 2 sends a message m_2 for sure; 1 randomizes between m_2 and m_1 (probabilities $1-\delta$ and δ respectively). After m_2 DM plays b_2 with strictly positive probability; after m_1 she plays b_2 with zero probability. By the assumption that $\beta^* > \tilde{\beta}$, DM will not play g after (m_{ss}, m_2) (if b_2 is optimal then $\beta \leq \tilde{\beta}$ and so g is strictly worse than b_1). She will not play b_2 after (m_{ss}, m_1)

because B is then revealed as 1. Therefore we can assume that after (m_{ss}, m_2) DM plays b_2 with probability ϕ and b_1 with probability $1-\phi$; and after (m_{ss}, m_1) she plays b_1 with probability ω and g with probability $1-\omega$.

Note that $\omega \in (0,1)$ because if DM plays g for sure then 1 would strictly prefer m_2 and if she played b_1 for sure 1 would strictly prefer m_1 (since $\phi > 0$). This in turn means that DM's belief after m_{ss}, $\gamma_B(m_{ss})$, equals $\tilde{\gamma}$, so that DM is willing to randomize after 1 reveals. This in turn implies that

$$\psi = \frac{\gamma - \tilde{\gamma}}{\gamma(1-\tilde{\gamma})}$$

by Bayes' Rule.

By Lemma 3 this equilibrium can exist only if $\gamma \geq \tilde{\gamma}$. To show that $\beta \geq \beta^*$, assume $\beta < \beta^*$. Then, after G reveals as h and B pools, DM plays g with probability zero. But, after m_{ss}, DM plays g with strictly positive probability since $\omega < 1$. Therefore, since h is indifferent between m_p and m_{ss}, l must strictly prefer m_p. Contradiction. Hence $\beta \geq \beta^*$.

As just shown, DM must play g with strictly positive probability after G reveals as h. This implies that h's payoff from m_p is $v_G(g,.,h)$ (if b_1 is also played then DM, and hence h must be indifferent between b_1 and g). His payoff from m_{ss} is the same. The equilibrium payoff of l is $v_G(g,.,h) - \beta\delta(1-\omega)(v_G(g,.,h) - v_G(g,.,l)$ since he gets the same payoff as h unless g is played, which happens with probability $\beta\delta(1-\omega)$. So DM's equilibrium payoff is

$$\gamma v_G(g,.,h) + (1-\gamma)v_G(g,.,h) - \beta\delta(1-\omega)(v_G(g,.,h) - v_G(g,.,l)$$

and her best equilibrium in this class is the one with the lowest δ, so that $\beta(m_2) = \tilde{\beta}$. This implies that g is optimal both after m_2 and after m_1 and so DM's payoff in this equilibrium is the same as if she played g for sure after m_p and b_1 for sure after m_{ss}. This payoff is

$$\gamma\psi v_G(g,.,h) + [\gamma(1-\psi) + (1-\gamma)][(1-\beta)v_G(b_1,2,.) + \beta v_G(b_1,1,.)]$$

which reduces to the expression given in the Lemma. QED

References

Austen-Smith, D. (1993). Interested experts and policy advice: multiple referrals under open rule. *Games and Economic Behavior*, 5, 3-43.

Boujarde, S. and B. Jullien (2011). The roles of reputation and transparency on the behaviour of biased experts. *RAND Journal of Economics*, 42 (3), 575-594.

Che, Y.-K., Dessein. W., and N. Kartik (2013). Pandering to Persuade. *American Economic Review*, 103, 47-79.

Crawford, V., and J. Sobel (1982). Strategic information transmission. *Econometrica*, 50 (6), 1431-1452.

Dewatripont M., and J. Tirole (1999). Advocates. *Journal of Political Economy*, 107 (1), 1-39.

Krishna, V., and J. Morgan (2001). A model of expertise. *Quarterly Journal of Economics*, 747-775.

Ottaviani, M., and P. Sorensen (2000). Information aggregation in Debate: who should speak first? *Journal of Public Economics*, 81, 393-421.

Ottaviani, M., and P. Sorensen (2006). Professional advice. *Journal of Economic Theory*, 126, 120-142.

Wolinsky, A. (2002). Eliciting information from multiple experts. *Games and Economic Behavior*, 41, 141-160.