

Wait and See*

Péter Eső[†] and Yuk-fai Fong[‡]
Kellogg School of Management
Northwestern University

November 2007

Abstract

We study a dynamic cheap talk model with multiple senders where the receiver can choose when to make her decision and communication can take place over time. No player has the ability to commit to any action in the future, in particular, the receiver cannot commit to delay the decision. In contrast to the results in static versions of the model, we show that when the senders have common knowledge about the state of the world, there exists an equilibrium with instantaneous, full revelation irrespective of the size and direction of the senders' biases. We show that the equilibrium is robust to the introduction of noise in the senders' signals about the state. The conditions under which the equilibrium outcome with noisy observation converges to immediate full disclosure as the noise disappears involve the size of the senders' biases and their patience.

Keywords: multi-sender cheap talk, full revelation, noisy signals, delay

JEL Numbers: D82, D83, D72, D74

*We thank David Austen-Smith, Olivier Compte, Tim Feddersen, Ehud Kalai, Li Hao, Andrea Prat, Larry Samuelson, James Schummer, Min Shi and Joel Sobel and seminar participants at Northwestern and LSE for comments and Renato Gomes for excellent research assistance.

[†]eso@kellogg.northwestern.edu.

[‡]y-fong@kellogg.northwestern.edu.

1 Introduction

Consulting with experts before making a decision takes time, and happens over time. Our goal in the present paper is to point out some fundamental consequences of considering information transmission between multiple experts and a decision maker in a dynamic setting. We investigate how the threat of delay (which no-one can commit to) can induce communication between parties with dissimilar interests.

We set up a model where a decision maker (she) can consult multiple experts who have private information about the state of nature. There is a single decision to be made, and neither the payoff-relevant state nor the experts' private signals and biases change over time. The parties can communicate—send messages back and forth, simultaneously or sequentially, publicly or privately—at every point in time. We assume that the message space is rich enough to transmit what the state of nature is. At every point in time, the receiver decides to either end the game by taking an action corresponding to a point in the state space, or choose momentary inaction. When the game ends, the players receive payoffs representing single-peaked preferences over the action space with convex loss functions. Conflict arises because the senders' ideal points differ from the receiver's conditional on the state of nature. All players have strict time preferences, that is, they prefer a given action to be carried out earlier, and their outside options are such that all are guaranteed to participate.

In this model, we study perfect Bayesian equilibria, that is, equilibria where all actions are rational at every point in time given the history of the play. In other words, no player has commitment power and all behave in a time-consistent way. In particular, the decision maker cannot commit to wait for any period of time; any delay on the equilibrium path has to be rationalized.

For a concrete example think of immigration reform in the United States. The decision maker is the median voter in Congress; the senders are various committees, individual representatives or even lobbyists who have pertinent private information on the matter.¹ The honest opinions of these experts may disagree

¹A large literature following Gilligan and Krehbiel (1987), (1989) models parliamentary institutions and legislative procedures as cheap talk games.

due to the difference in what they know, but based on all their information it would be possible to determine Congress' most-preferred policy choice. However, the senders are also *biased* from the median voter's perspective; due to their business interests, ethical considerations, etc. Therefore, a disagreement among them may also reflect their desire to shift Congress' decision in their favor. The question is how much information can be transmitted if the legislative process takes place over time—hearings followed by proposals, debates and votes—without the median voter being able to explicitly commit to future actions. Delay is wasteful for all interested parties because the social, political and economic problems stemming from illegal immigration are not getting better with time.² However, procedural delays (or threats thereof) could induce more honesty from the senders. Other legislative issues in the U.S. that share the same characteristics include entitlement reform and health care. We provide more applications below.

In our formal analysis, we first study a model where the two experts have *common knowledge* about the realization of the state of nature, while the decision maker only knows the state's prior distribution. We prove that, no matter what the state space or the senders' biases or any players' time preferences are, there exists a Perfect Bayesian Equilibrium where the state of nature is immediately and truthfully reported by both experts, and the decision maker carries out her ideal action without delay. The senders' equilibrium strategy is to report the truth continuously through time. The receiver's strategy is to carry out the action matching the senders' report as soon as their messages agree, and do nothing while their reports mismatch. This off-equilibrium response is rationalized by the expectation (shared among all parties) that both senders will report the state truthfully at all points in time following a disagreement.

Notice that the receiver does not have the power to commit to wait for any amount of time. She may choose inaction as long as she has an incentive to do so. However, since each sender expects the other sender to be truthful and the receiver to delay action until the messages agree with each other, they have no incentive to lie and postpone the decision even by an instant.

²There are other legislative issues where delay may be beneficial but we believe that is not the case with immigration reform.

The result that a receiver can *immediately* extract *all* information from biased experts with no commitment to delay is remarkable because in a static multi-sender cheap-talk model the result only obtains under certain conditions. As Battaglini (2002) and Ambrus and Takahashi (2007) have shown, the necessary and sufficient condition for the existence of a fully revealing equilibrium is that the state space be relatively large compared to the senders' biases.³ In contrast, our result holds *without any restriction* on the shape or size of the state space, nor does it depend on the direction or size of the senders' biases. The fully revealing equilibrium in our model involves no delay, therefore an outside observer would not even realize that the underlying situation is dynamic.

It is intuitive that a decision maker, by insisting that the experts agree on their advice and repeatedly asking them for clarification until they agree can do better than she could in the standard static framework. What is truly interesting, in our view, is that the decision maker can fully extract the senders' private information even if she is unable to commit to delay her decision. She can do so by rationally maintaining a positive attitude and believe that the senders will eventually agree. (This is also our practical advice to decision makers with access to well-informed experts.) Such beliefs are plausible in the model where the senders have bilateral common knowledge of the state of nature—they *must* be able to agree, and the receiver knows that.

Similar equilibrium constructions for (non-repeated) dynamic models have been used outside the cheap talk literature. Beliefs regarding offers and counteroffers in the continuation game uphold the perfect equilibrium with instantaneous agreement in Rubinstein's (1982) infinite-horizon dynamic bargaining game. In the durable-good oligopoly model of Ausubel and Deneckere (1987) and Gul (1987) the buyers' off-equilibrium beliefs discipline the sellers' dynamic pricing behavior. Marx and Matthews (2000) and Lockwood and Thomas (2002) use a related construction to overcome the free-riding problem in public good

³The conditions identified in the cited papers apply irrespective of the dimensionality of the state space, but their implications depend on it. For example, suppose that the senders' biases are state-independent and their preferences are represented by quadratic loss functions. Then, in a one-dimensional state space the necessary-sufficient conditions are that the biases are relatively small compared to the difference between the highest and lowest states, while in a multi-dimensional setting the conditions hold as long as the state space is rectangular in the biases.

provision when contributions are made over time. The investment hold-up problem is resolved by Gul (2001) using repeated contract offers and by Che and Sákovics (2004) using dynamic investment.

In the example of legislation in Congress we mentioned that the experts may honestly disagree due to the differences in what they know. In reality experts do not always have common knowledge about the state of nature—they may only have imperfect, correlated signals about it. If so then “cross-checking” the senders’ reports and only acting when they fully agree may not be a feasible course of action for the receiver. This criticism applies not only to our model, but also to a large chunk of the existing literature on multi-sender cheap talk. In those models, too, the state of nature is commonly known between the senders, and the construction of a fully revealing equilibrium hinges on this fact.⁴

We extend our multi-sender dynamic model to allow for noisy signals on the senders’ part. We prove that there exists an equilibrium with imperfect revelation and positive expected delay. However, if the senders’ signals become arbitrarily precise in a certain way, the equilibrium outcome converges to the one found under common knowledge about the state, that is, complete disclosure and no delay. We also discuss whether and how the receiver’s eagerness to learn the state of nature in comparison to the senders’ biases and impatience matters for the existence of such equilibria. Our result shows that the no-delay, fully efficient equilibrium under perfect observation is not an aberration as it may be the limit of sensible equilibria as the noise in the senders’ signals vanishes. This is what we consider the main result of the paper.

These results shed light on whether and how delay—that the decision maker cannot commit to—can induce the experts to reveal more information. The insights are applicable in our first motivating example, legislation in Congress.

Another application where our model and results are relevant is where a

⁴Papers where the experts do not have common knowledge about the state of nature include Austen-Smith (1993), Wolinsky (2002), Battaglini (2004). The latter paper extends Battaglini’s (2002) model to a setting where each sender observes the state of nature with an independent normal noise, while the state of nature has a diffuse prior over a Euclidean space. Battaglini shows that the equilibrium based on the orthogonal decomposition of the state of nature (in the base formed by the bias vectors) is robust to introducing this type of imperfect observation on the senders’ part. As the noise in the senders’ signals diminishes, the equilibrium converges to the one with full revelation in Battaglini’s (2002) model.

court needs to decide in a dispute between two parties with opposing interests. The parties have information about the state of the world, which the court wishes to match with its decision. The court can repeatedly ask the parties to provide (soft) information, but delays hurt the litigants as well as the court. The court cannot commit to impose a period of waiting on the parties; any delay can only be the result of the court's equilibrium beliefs that repeated questioning will lead to information revelation and a socially beneficial outcome. Indeed it will, as our results show. A case in point is the recent patent dispute between Research in Motion (RIM, the makers of Blackberry phones) and NTP Inc. (a patent holding company). NTP sued RIM for patent infringement in January 2000. After a series of claims, counterclaims, court decisions and reversals, the parties agreed on a settlement in March 2006.⁵ Both parties faced substantial delay costs: RIM risked losing business due to the ongoing litigation, while NTP's patents were being investigated (and partly invalidated) by the U.S. Patent Office. Commentators found the settlement appropriate given the information revealed during the suit.⁶

A similar situation is that of a police interrogator questioning multiple suspects of a crime. The suspects have correlated information about the facts of the case, and the interrogator can ask them again and again if their stories disagree too much. Repeated questioning decreases the suspects' utilities as they have to forego other activities while held in police custody. However, most interrogators cannot explicitly commit to delay. We claim that repeated questioning hoping that the suspects will release more information is indeed a good strategy for the interrogator.

In some of the applications it may be reasonable to assume that even in the absence of transfers, the decision maker could impose different waiting costs on the experts. For example, the interrogator could make one suspect's wait a lot less comfortable than another's—her tools may include anything from harassment to outright torture. While we do not explicitly allow this in our base model, we shall comment on its implications as we discuss the results.

Our paper contributes to the literature on multi-sender cheap talk games.

⁵Sources: Various news reports from The Register, www.theregister.co.uk.

⁶See http://money.cnn.com/2006/03/03/technology/rimm_ntp.

This literature is partly motivated by applications to political theory and focuses on equilibria and institutions that facilitate information transmission among experts and a decision maker; see Gilligan and Krehbiel (1987, 1989), Austen-Smith (1990, 1993), Krishna and Morgan (2001a), Battaglini (2002, 2004), Wolinsky (2002) and their references. Battaglini (2002) and Ambrus and Takahashi (2007) prove that in the static model where the senders have common knowledge about the state, large biases make it impossible to sustain a fully revealing equilibrium. Krishna and Morgan (2001b) show under similar conditions that adding multiple rounds does not induce full revelation either. In contrast, we obtain an *unconditional possibility* result in our baseline dynamic model whose signal structure (perfect observation) is comparable to many of the cited papers.

The explicit modeling of time and the possibility of costly delay in our model is reminiscent from dynamic bargaining models (see Serrano (2007) for an overview). However, our underlying game—multi-sender cheap talk—is rather different. Rubinstein and Wolinsky (1992) add a time dimension to a bilateral trading model and study renegotiation-proof contracts; Artemov (2006) considers Nash implementation with costly delay as a punishment device. Our problem is fundamentally different from these because the mechanism designer can commit to delay and other distortions, while the receiver in our cheap talk game cannot. A contemporaneous paper by Damiano, Li and Suen (2007) studies how the use of commitment to costly delay in the voting rule helps players improve efficiency in a dynamic voting game. In their model delay upon disagreement is enforced by the voting rule, while in ours it is a time-consistent rational choice on the part of the receiver. Again, the difference in our assumptions is important. The main contribution of our paper is to establish how the receiver can credibly carry out the delay even without commitment and how costly delay can be avoided in equilibrium as the senders’ signals become more precise, while Damiano, Li and Suen focus on other issues.

Our dynamic model is somewhat related, but not equivalent to, cheap talk models with money-burning (see Austen-Smith and Banks (2000)). Delay “burns” the payoffs of all parties, not just the senders’. More importantly, in our setup it is the receiver who can decide on delay, not the senders, and the receiver cannot commit to any period of delay. In contrast, in Austen-Smith and Banks (2000)

it is the informed party (the sender) who can *signal* his type by committing to verifiably reduce his payoff. In that model there exists an equilibrium where the sender signals a higher state of nature by burning more money, and all types separate.⁷ Our results are different, too: In our setup there is no private or social loss on the equilibrium path when the senders have common knowledge about the state.

The paper is structured as follows. We set up the model in Section 2. The case of perfect observation (when the experts commonly know the state) is analyzed in Section 3. We discuss noisy signals in Section 4. Section 5 concludes, and an Appendix contains omitted proofs.

2 The model

The state of nature, ω , is a random draw from a set $S \subseteq \mathbb{R}^n$. There are two experts (the senders) and one decision maker (the receiver). The senders have private information about the state of nature. Denote the signal that sender i observes by X_i for $i = 1, 2$; for simplicity let the signals also belong to the state space. Assume that the joint distribution of (ω, X_1, X_2) is commonly known.

Time runs continuously from $t = 0$ to infinity. At any point in time t , each sender (he) can send private or public messages to any other player, and the receiver (she) can take an action y , which is either a point in S , or the “null” action, i.e., doing nothing *at that instant*. (Moves by the senders and the receiver may be sequential while taking place at the same point in time.) As soon as the receiver picks a decision in S the game terminates.

Denote the utility of the receiver ($i = 0$) and the senders’ ($i = 1, 2$) at time t , in state ω , from a decision y that is to take place at time $t + \delta$ by $U_i(\omega, y, \delta)$. We assume that for all ω and δ , U_i is concave and single-peaked in y and the location of the maximum is invariant in δ . For notational simplicity we identify the receiver’s ideal point with the state, $y_0(\omega) = \omega$, and denote sender i ’s ideal point by $y_i(\omega) = \omega + b_i(\omega)$, $i = 1, 2$. We assume that for any $\delta < \infty$ each player’s utility $U_i(y, \omega, \delta)$ is bounded; the bound may vary (and diverge) as $\delta \rightarrow \infty$.

⁷See Austen-Smith and Banks (2001), Kartik (2007). If the sender’s budget is limited, there may not exist a fully revealing equilibrium.

All parties have strict time preferences: $\partial U_i(\omega, y, \delta)/\partial \delta < 0$, that is, players prefer an action to be carried out earlier. The time preference can be the result of exponential or hyperbolic discounting of a positive utility function, for example $U_i = \int_0^\delta u_i(y, \omega) e^{-r\tau} d\tau$ where $u_i > 0$ is single-peaked at $y = y_i(\omega)$, or due to a cost of waiting, for example $U_i = -(y - y_i(\omega))^2 - r\delta$, $r > 0$. The outside options of all players are set to $-\infty$, so they are guaranteed to participate.

In this game we study Perfect Bayesian Equilibria, that is, equilibria where all moves are time-consistent. In other words, no player can commit to any action in the future, in particular, the receiver cannot commit to delay.

We assume that for any two actions y and y' such that player $i \in \{0, 1, 2\}$ prefers y with no delay over y' with no delay, we can find a delay δ such that y with delay δ is equivalent to y' with no delay: $U_i(y, \omega, \delta) = U_i(y', \omega, 0)$. That is, a sufficiently long delay can compensate for any utility difference over immediate actions.

We define a couple of quantities related to the players' preferences (or trade-offs) over outcomes and delay. Δ_0 is the maximum delay that the receiver is willing to bear in order to learn the senders' signals and carry out her ideal point contingent on the state:

$$E[U_0(\omega, \omega, \Delta_0)] = E[U_0(\omega, E[\omega], 0)]. \quad (1)$$

Δ_0 is larger the more diffuse the prior distribution of ω is; Δ_0 is also increasing the receiver's patience. For example, if the receiver's utility is represented by a quadratic loss function and a linear cost of delay then Δ_0 is proportional to the variance of the prior of ω . Let $\Delta_i(\omega)$ denote the (state-contingent) willingness to wait of sender i in order to achieve his ideal point instead of the receiver's:

$$U_i(\omega, y_i(\omega), \Delta_i(\omega)) = U_i(\omega, \omega, 0). \quad (2)$$

For example, if the sender's ideal point differs from the one-dimensional state of nature by a constant b_i , has a quadratic loss function and a linear waiting cost, then $\Delta_i(\omega)$ is simply proportional to b_i^2 .

As it is customary in dynamic games we assume that the players can synchronize their moves, for example, with the help of a public clock. For certain

equilibrium constructions where we want to avoid communication (and hence the possibility of collusion) among the senders it will be useful to assume that the receiver has a randomization device with the following property: It generates a random variable that both senders simultaneously observe and have mutual common knowledge about, without the receiver observing its realization.

In the next section we first consider a model in which each sender observes the state of nature ($X_i \equiv \omega$) and the message space is rich enough for the transmission of the realization of ω . We show the existence of an equilibrium in which the state of nature is immediately revealed by the senders and so the receiver carries out $y = \omega$ without delay. Subsequently, we investigate the robustness of this fully revealing equilibrium to perturbations of the model. In particular, we investigate the case where the senders' private information stems from noisy signals about the state of the world.

3 The perfect observation case

In this section we assume that each sender observes the state of nature and the message space is rich enough for the transmission of the realization of ω . We want to find the most informative Perfect Bayesian Equilibrium.

3.1 The possibility result

The main result of this section is that in our dynamic multi-sender cheap talk game there exists an equilibrium where all information commonly known by the senders is immediately revealed.

Proposition 1 *Assume that the senders perfectly observe the state: $X_1 = X_2 = \omega$. Then there exists a Perfect Bayesian Equilibrium where the state of the world is immediately and truthfully reported by both senders, and the receiver carries out $y = \omega$ without delay.*

Proof. The senders' strategy is to report ω at all $t \geq 0$, and the receiver's strategy is to set y equal to the senders' report as soon as their messages agree, and do nothing while their reports mismatch. This off-equilibrium response is

rationalized by the expectation (shared among all parties) that both senders will report ω truthfully at all points in time following a disagreement. Notice that it is rational for the receiver to choose inaction off the equilibrium path because she believes she will learn the state immediately. ■

Notice that since the equilibrium in Proposition 1 exhibits no delay on the equilibrium path, the outcome of our dynamic game is *observationally equivalent* to a fully revealing equilibrium in a corresponding static environment.

However, as we mentioned in the Introduction, fully revealing equilibria do not always exist in a static environment. For example, if the state space is one-dimensional and the senders' biases have opposite signs, then a fully revealing equilibrium exists only if the sum of the absolute values of the senders' biases is not too large relative to the size of the set of states (see Krishna and Morgan (2001a), Battaglini (2002)). If the state space is multi-dimensional, then the receiver can rely on the clever technique of only considering a sender's report of the state in the dimensions that are orthogonal to that sender's bias at the state of nature, and induce full revelation (Battaglini (2002)). For this equilibrium construction to work, either the state space must be "rectangular" in order to avoid inconsistent reports, or the size of the biases must be limited, just like in the one-dimensional case (Ambrus and Takahashi (2007)). In contrast, our result regarding immediate full disclosure holds *without any restriction* on the shape or size of the state space, nor does it depend on the direction or size of the senders' biases.

The equilibrium constructed in Proposition 1 is striking because it works under very general conditions. As we have mentioned there are no restrictions on the state space, the senders' biases, or anyone's time preferences. In the equilibrium construction we only used public messages. It is also easy to see that the same equilibrium works even if the receiver is uninformed about the senders' biases. In the rest of the section (in fact, the rest of the paper) we discuss whether the outcome, instantaneous full revelation, is robust to certain other perturbations of the model. Perturbations that do not concern the information structure are discussed in the next subsection, while the issue of noisy signals is the topic of Section 4.

3.2 Discrete time, continuity of beliefs and credibility

The result of Proposition 1 can be derived under the assumption of discrete (rather than continuous) time as well. The condition that we need in this case is that the Receiver is better off making a decision with a one-period delay and knowing the exact state of nature rather than picking her optimal point right away given the prior on the state of nature. If this is the case then the strategies that form a Perfect Bayesian Equilibrium in the proof of Proposition 1 also work in a model with discrete time. The condition is satisfied if the discrete time period lengths are sufficiently short (i.e., the discount rate is sufficiently high), or if the prior distribution of ω is “sufficiently diffuse” (i.e., the prior has a sufficiently large variance).

We may also wonder what a rational receiver should believe if the two reports differ only slightly. Perhaps she would think that the true state of nature is very close to either report, and carry it out without waiting for a complete agreement? A similar issue, continuity of the receiver’s beliefs, is studied in the static problem by Battaglini (2002) and Ambrus and Takahashi (2007).

It is easy to construct an equilibrium with instantaneous full revelation that is immune to this type of problem. Suppose the senders’ strategies prescribe that they report the state of nature as an element of a random, zero-measure set that is nevertheless dense in the state space. If they play the equilibrium then the receiver can identify the state of nature as the common element of their reported sets, and there is zero probability that there are two or more points in the intersection. Moreover, if one sender reports a set that does not contain the true state of nature then there will be elements of the other sender’s set that are arbitrarily close to anything that he reported. Therefore, the issue whether the receiver should believe that the state of nature is close to the reports when the reports are close to each other is moot.

Finally, even though the equilibrium is a PBE, the receiver’s promise that she will wait out the senders may not be intuitively credible in certain situations. For example, the prior may be so precise that the receiver does not expect to gain much from learning the true state (Δ_0 is small). If, say, both senders know that telling the receiver the state of nature hurts them a lot more than it helps her, then mutual truth-telling may not be the focal equilibrium. While this

criticism is rather informal and hence difficult to validate or refute formally, we will come back to a similar idea towards the end of the next section.

Perhaps the most interesting issue is to what extent our construction is robust to the introduction of noise in the senders' signals. This is what we study in depth in the next section.

4 Delay and disclosure with imperfect signals

In this section we investigate the robustness of the fully revealing equilibrium of the perfect-observation model to relaxing the assumption that the senders have common knowledge about the state of nature. In particular, the environments that we consider have the common feature that both the state of nature and the senders' signals can take on a continuum of values and the probability that the realizations of senders' signals exactly match is zero (even conditional on the realization of the state). As a result, the equilibrium construction of the model of the previous section, which relies on the fact that it is commonly known the senders both know the state of nature, breaks down. However, a modified construction still works, where the senders either precisely or approximately reveal their signals, and there is a positive expected delay along the equilibrium path. In our environments, as the senders' signals become arbitrarily precise, the equilibrium outcome converges to full revelation and no delay—the outcome of the equilibrium found under perfect observation.

4.1 A special case: opposite biases, diffuse prior

In order to better explain the workings of our proposed equilibrium under imperfect signals we first develop the core ideas in a special case. Suppose that the state of nature is drawn from a uniform distribution on $\mathbb{R}$, that is, the common prior is *diffuse*. Sender i 's signal, X_i , is normal with mean ω and variance σ^2 for $i = 1, 2$. The receiver's ideal point is ω , sender 1's is $\omega + b$, while sender 2's is $\omega - b$, where $b > 0$. All players have the same quadratic loss function. Time preference arises as a result of a unit cost of waiting per unit time:

$$U_i(\omega, y, \delta) = -(y - y_i(\omega))^2 - \delta \text{ for } i = 0, 1, 2. \quad (3)$$

The restrictive assumptions are (i) diffuse prior and (ii) symmetry of the senders, that is, equal, opposite biases and symmetric signal structure. These are to be relaxed in the next subsection.

Before describing the equilibrium we make some observations. By our assumptions and the rules of updating normal variables, the receiver's ideal point as a function of the realization of the senders' signals is $E[\omega|X_1 = x_1, X_2 = x_2] = (x_1 + x_2)/2$. Each sender believes that the state of nature is distributed normally around the realization of his own signal with variance σ^2 ; they believe the other sender's signal is distributed around the state of nature with an additional normal noise that has variance σ^2 .

Purely as a thought-experiment, consider a *mechanism* where each sender reports the realization of his signal, $\hat{x}_1$ and $\hat{x}_2$ respectively, the receiver waits for a period of time $d(\hat{x}_1 - \hat{x}_2)$, and then carries out $y = (\hat{x}_1 + \hat{x}_2)/2$. The following lemma states that there exists a wait-function d such that this mechanism is incentive compatible for the senders.

Lemma 1 (Commitment to delay) *Assume that the senders observe conditionally independent normal signals about ω whose prior distribution is diffuse on $\mathbb{R}$, and that their biases are of equal size and opposite sign. Consider a relaxed problem where the receiver can commit to a delay d as a function of the difference of the senders' signal reports before carrying out her ideal action conditional on their reports. Truthful signal reports are elicited by setting $d(\hat{x}_1 - \hat{x}_2) = 2b \max\{\hat{x}_1 - \hat{x}_2, 0\}$, where $\hat{x}_i$ is sender i 's report for $i = 1, 2$.*

Proof. See the Appendix. ■

The significance of this mechanism is that its outcome can be replicated in an *equilibrium* with no commitment to delay on the receiver's part. Formally, we have the following result.

Proposition 2 *Assume that the senders observe conditionally independent normal signals about ω whose prior distribution is diffuse on $\mathbb{R}$, and that their biases are of equal size and opposite sign. Then, there exists an equilibrium where the receiver learns the realizations of the senders' signals with a positive expected delay and carries out her full-information ideal action. As the noise in the*

senders' signals diminishes the expected delay tends to zero, and the equilibrium outcome converges to immediate full disclosure and $y = \omega$.

Proof. We construct the equilibrium as follows. On the equilibrium path, at time $t = 0$, the senders simultaneously and privately report to each other their signals; denote the report of sender i to the other sender by $\hat{x}_i$, for $i = 1, 2$. Then, at every point in time (including $t = 0$) both senders send public messages. They independently babble for all $t \in [0, d(\hat{x}_1 - \hat{x}_2))$ and announce $m^* = (\hat{x}_1 + \hat{x}_2)/2$ from $t = d(\hat{x}_1 - \hat{x}_2)$ on. The receiver takes the “null” action as long as the senders’ reports disagree and carries out their report as soon as they agree. Clearly, if both senders report to each other truthfully at time zero and follow their equilibrium strategy thereon then the receiver will learn and implement $E[\omega|X_1 = x_1, X_2 = x_2]$ with delay $d(x_1, x_2)$ as claimed.

The equilibrium is sustained by the following off-equilibrium beliefs and behavior. At any $t \geq d(\hat{x}_1 - \hat{x}_2)$ neither sender has an incentive to report anything other than $m^* = (\hat{x}_1 + \hat{x}_2)/2$, given that the other reports m^* , because a deviation only leads to a longer delay. Neither sender can force an earlier resolution. This is so because neither sender can predict what the other sender reports during the period of babbling, hence he cannot match the other’s announcement. Moreover, he cannot credibly report m^* to the receiver either because he is believed to be babbling. By Lemma 1, sender i has no incentive to misreport his signal to the other at $t = 0$. Assume that if a sender fails to make his private report at $t = 0$ then the other sender babbles forever; this eliminates any incentive for a sender to deviate at $t = 0$. Finally, the receiver has no profitable deviation because the senders babble until $t = d(\hat{x}_1 - \hat{x}_2)$. By the diffuse prior assumption she infers nothing about the expected value of the state of nature (even as time passes and she learns lower bounds on $\hat{x}_1 - \hat{x}_2$) until $t = d(\hat{x}_1 - \hat{x}_2)$, and she has an incentive to wait because she learns the posterior mean of ω after finite expected delay.

Clearly, the expected value of $d(x_1 - x_2)$ goes to zero as the senders’ signals become arbitrarily precise. ■

The most important fact to note is that as the noise in the senders’ observation of the state of nature diminishes, the expected equilibrium delay goes to zero. That is, the equilibrium outcome converges to *full revelation* and *no*

delay—exactly what we found in the perfect observation case of the previous section. In other words, the outcome of the fully efficient equilibrium under perfect observation is robust to the introduction of noise in the senders’ signals (at least in the environment considered in Proposition 2).

Our analysis clearly demonstrates that the lack of common knowledge about the state is not an insurmountable impediment to eliciting timely and truthful reports from the experts. The key in the construction is that the senders’ private messages at time 0 *create mutual knowledge* of the (reported) signals and hence the posterior on ω ; the delay induces truthful bilateral reports at $t = 0$ and provides incentives for making matching announcements when the delay is over.⁸

There are other, perhaps more attractive ways to generate the same equilibrium outcome in the example; in what follows we describe an alternative. Note that the construction provided in Proposition 2 relies on a perfect synchronization device to make sure that the bilateral reports occur exactly at the same time at $t = 0$. Another criticism could be that if the senders can use private messages to create common knowledge of (x_1, x_2) among themselves then they could use such messages for collusion as well. Both issues are resolved if we use a certain randomization device instead of synchronized private communication.

Suppose that the receiver can generate random variables whose realizations only the senders can observe, while she cannot, and this is commonly known. Then, an alternative equilibrium construction is as follows. At time $t = 0$ the receiver generates a random number θ (“blindly” so that only the senders can observe) with a diffuse distribution over $\mathbb{R}$. Still at time 0, both senders publicly report their signals plus this random number, $m_i = \hat{x}_i + \theta$ for $i = 1, 2$. By taking the difference of the two reports, all players can calculate $m_1 - m_2 = \hat{x}_1 - \hat{x}_2$, while the receiver remains ignorant about the levels of $\hat{x}_1$ and $\hat{x}_2$. Both senders babble and the receiver waits until $t = d(\hat{x}_1 - \hat{x}_2)$. When the waiting period is over, both senders continuously and truthfully report θ , and the receiver carries out $(m_1 + m_2)/2 - \theta$ as soon as they agree. If the senders disagree at any $t \geq d(\hat{x}_1 - \hat{x}_2)$ then the receiver believes they will agree right away; as a result

⁸Wolinsky (2002), using a very different model, argues that allowing private communication among the senders can help the decision maker in eliciting information from them.

it is indeed an equilibrium for both senders to report θ at all $t \geq d(\hat{x}_1 - \hat{x}_2)$. All players believe that the senders babble during the waiting period, hence there is no profitable deviation for anybody during the wait. Finally, by Lemma 1 it is optimal for the senders to report the realization of their signals (plus θ) truthfully at $t = 0$. Since the receiver cannot infer anything about the state of nature from knowing the difference of the signal realizations, she cannot do anything but wait until the senders reveal θ .

Both this construction and the one presented in the proof of Proposition 2 rely on the assumptions that (i) the prior is diffuse and (ii) the senders are symmetric (their biases are of equal magnitude and their signals equally precise). The symmetry assumption makes it possible that we provide exact incentives for both senders to report truthfully with the same delay function (as it is impossible to require different wait times from the two senders). In a modified model where the decision maker could impose different costs of waiting on the experts the symmetry assumption would not be needed.⁹ The symmetry assumption can also be relaxed if the state space is discrete. In that case, a delay that is a sufficiently steep function of the absolute value of the difference of the reports would induce truthful announcements from asymmetric senders as well. This delay could be implemented in equilibrium the same way as it is done in Proposition 1. The symmetry assumption should therefore be thought of more as a technical condition rather than a substantial one.

The assumption that the prior distribution of the state is diffuse (improper) is not simply a technical one, however.¹⁰ The role this assumption plays in the equilibrium construction is the following. Due to the lack of prior each sender believes that the state of nature and the other sender's signal are distributed around his own signal realization. In the absence of this property it would be difficult to devise an incentive-compatible delay mechanism where the delay only depended on the difference of the two signal reports. Furthermore, the diffuse prior assumption is unattractive because of its mathematical imprecision (no such distribution exists) as well as because it is incompatible with the notion

⁹For example, an interrogator can treat two suspects differently during questioning and thereby inducing different costs of waiting with the same wait time.

¹⁰The diffuse prior assumption is also made in Battaglini (2004) where it is not innocuous either, although for different reasons.

of bounded state space. (Recall that in static models, if senders observe the state perfectly, bounded state spaces yield qualitatively different predictions compared to unbounded ones.) However, the most important consequence of the improper prior assumption is that the receiver is willing to wait arbitrarily long in order to learn the senders' signals. In the absence of a prior her loss from picking a random outcome is infinity, and she continues to have a posterior over the state with an infinite variance even after learning the difference in the experts signal realizations. A completely ignorant receiver who is not capable to decide without the experts' input is not a good description of many situations we would like to model.

Next, we exhibit an environment with a bounded state space and a proper prior where there exist communication equilibria with *imperfect revelation*. We show that under certain conditions the outcome converges to immediate full revelation as the senders' signals become arbitrarily precise.

4.2 Noisy signals and delay in the general case

We now consider a family of information structures where the senders observe correlated private signals about the state of nature that belongs to a compact set and all players have a proper common prior about. The information structures in this family are naturally ordered according to how precise the senders' signals about the state are. We prove that if the senders' biases and time preference are not too strong then, for a sufficiently precise signal structure there exists an equilibrium with imperfect information revelation and positive expected delay. However, as the noise in the senders' signals vanishes the equilibrium converges to one with full revelation and no delay on the equilibrium path. That is, in the limit, we obtain the outcome of the perfect-observation model. It may be interesting to point out that this result holds under certain conditions on the biases and patience of the senders. It turns out that these parameters do matter after all, albeit only in the imperfect-observation case with a proper prior.

Let the state space S be the unit interval and denote the state by ω . Our main assumption regarding the information structure is that the senders observe correlated signals regarding a *discrete approximation* (grid-coarsening) of the random variable ω . For $N = 1, 2, \dots$, we define the discrete approximation, ω_N ,

as follows.¹¹ Let ω_N be the random variable whose value equals the midpoint of the interval $[a_{k-1}, a_k) = [(k-1)/2^N, k/2^N)$ that the realization of ω belongs to for $k = 1, \dots, 2^N$. That is, when $|S| = \infty$, we let

$$\omega_N = \frac{2k-1}{2^{N+1}} \Leftrightarrow \exists k \in \{1, \dots, 2^N\} : \omega \in [a_{k-1}, a_k). \quad (4)$$

The senders' signals, X_1 and X_2 , are also distributed on $[0, 1]$ with full support for every realization of ω (or ω_N). However, we assume that the signals are correlated with each other and the state of nature, meaning that conditional on the realization of ω_N , both X_1 and X_2 are likely to be in the neighborhood of ω_N . Formally, we assume that each sender i believes the probability of X_{-i} (and also ω_N) belonging to $[a_{k-1}, a_k)$ is at least $1 - \varepsilon$ conditional on i 's signal realization falling in the same interval:

$$\Pr(\{\omega_N, X_{-i}\} \subset [a_{k-1}, a_k) \mid X_i \in [a_{k-1}, a_k)) \geq 1 - \varepsilon. \quad (5)$$

We also assume that $\Pr(X_1 \in [a_{k-1}, a_k), X_2 \in [a_{\ell-1}, a_{\ell}), k \neq \ell \mid \omega_N)$ is constant in ω_N (i.e., independent of the realization of ω_N). This implies that the fact whether or not the two senders' signals fall into the same interval $[a_{k-1}, a_k)$ is not informative about the state of nature. However, we do not specify the joint distribution of (ω, X_1, X_2) in more detail because that is not necessary for obtaining our results.

Information and signal structures that belong to a family described above are indexed by N and ε , the fineness of the grid (in case ω is not discrete with finitely many realizations) and the precision of the senders' signals concerning the grid point that the realization of ω is closest to. When we say that the signal structure converges to complete information regarding ω between the senders, we mean that $1/N$ and ε go to zero.

The signal structure proposed here exhibits the characteristics mentioned at the beginning of the section: The state of nature and the senders' signals can take a continuum of different values and there is zero probability that the

¹¹The discrete approximation would not be necessary, and the rest of our assumptions could be stated more simply, if ω had finitely many possible realizations (finite state space). However, we follow the literature by allowing a continuum of states.

senders' signals coincide. It is true that as ε approaches 0 the senders are almost sure that they observed signal realizations in the same interval $[a_{k-1}, a_k)$, but it never becomes commonly known between the two senders that their signals belong to the same partition element. The particular assumptions we made regarding the signal structure—namely, that the senders observe a noisy signal of a discrete approximation of ω (in case the state space is not finite)—helps in that it provides a natural partition with respect to which the senders' signals can be reported.

Recall from Section 2 that Δ_0 denotes the amount of delay that finding out the state of nature is worth for the receiver, and that $\Delta_i(\omega)$ is sender i 's willingness to wait in order to induce his ideal point instead of the state of nature, ω . Define the scalar Δ_i as the sender's maximum possible willingness to wait:

$$\Delta_i = \max_{\omega \in S} \{ \Delta_i(\omega) \mid U_i(\omega, y_i(\omega), \Delta_i(\omega)) = U_i(\omega, \omega, 0) \}. \quad (6)$$

The following proposition claims that if the receiver's maximum wait time exceeds those of the senders and the signal structure is sufficiently precise then there exists an equilibrium where the senders truthfully report the intervals $[a_{k-1}, a_k)$ that their signal realizations fall into. This equilibrium involves a positive expected delay. However, as the signal structure approaches complete information regarding ω between the senders, the information obtained by the receiver becomes arbitrarily precise regarding the true state of nature, and the expected delay converges to zero.

Proposition 3 *If $\Delta_0 > \max \{\Delta_1, \Delta_2\}$, N is sufficiently large and ε is sufficiently close to zero, then there exists an equilibrium where the receiver learns which interval of the form $[a_{k-1}, a_k)$, $k = 1, \dots, 2^N$, each sender's signal realization falls into. The equilibrium expected delay is positive and tends to zero as $(1/N, \varepsilon) \rightarrow 0$.*

Proof. The equilibrium strategies are as follows. At time $t = 0$, each sender simultaneously and privately reports to the other sender an integer $k \in \{1, \dots, 2^N\}$ corresponding to which interval $[a_{k-1}, a_k)$ his signal realization belongs to. If their reports disagree then the senders delay play by a fixed period of time d , to be characterized below, by publicly babbling until $t = d$. Then, either at

$t = 0$ (if the private reports agreed) or $t = d$ (if they did not), and from then on forever, both senders simultaneously and publicly announce the expected value of ω conditional on the signal values privately exchanged at time 0. Meanwhile the receiver chooses inaction as long as the senders' reports disagree, and carries out their report as soon as they agree.

If play ever reaches $t \geq d$ and the senders' reports disagree then the receiver believes that for all $t > d$ the senders will agree on the (expected) state of nature and therefore rationally chooses momentary inaction. For all $t \geq d$ neither sender has a profitable deviation because the other sender reports the expected ω truthfully and the receiver waits until both reports agree. During the delay stage (if it takes place) neither sender has a profitable deviation because all players believe that both senders babble for all $t \in [0, d)$.

The receiver is better off choosing inaction during this period rather than any action $y \in S$ as long as d is smaller than the delay that makes her indifferent between taking an action immediately, only knowing that the senders' signals fall in different $[a_{k-1}, a_k)$ intervals, and taking an action with delay d , but knowing exactly which intervals the senders signals fall into. By assumption, the receiver's best action in the former case is $E[\omega]$ because the senders' disagreement does not provide information about the state. Her best action in the latter case is arbitrarily close to ω , provided that the signal structure is sufficiently close to the senders having complete information about the state. Therefore, for N large and ε sufficiently close to zero, d can be chosen to be arbitrarily close to Δ_0 .

The senders have no incentive to deviate at time $t = 0$ given their initial private reports to each other. If their private reports disagreed then we already established that neither sender has a profitable deviation for $t \in [0, d)$. If their private reports agreed then they anticipate that in the continuation the receiver will carry out the state of nature (at least approximately, for N large and ε small). If sender i behaves as if they disagreed then he induces a delay of length at least d . The best outcome that this sender can hope for after delay d is his ideal point, $y_i(\omega)$. However, by $\Delta_0 > \Delta_i$, we have $E[U_i(\omega, y_i(\omega), \Delta_0)] < E[U_i(\omega, \omega, 0)]$, hence a delay approximating Δ_0 makes such a deviation unprofitable for sender i . The same argument establishes that neither sender can be better off by lying to the other sender when they send

each other private messages at time 0: For ε close to 0, each sender anticipates that a truthful private report almost surely induces $y = \omega$ with no delay in the continuation, while a deviation very likely induces delay, which cannot benefit him even if the receiver's eventual action is $y_i(\omega)$.

Finally, the expected delay tends to zero as the signal structure converges to the senders having complete information about the state of nature. This is so because for all N , the probability that the senders disagree on which realization of ω_N their signals fall closest to goes to zero as ε goes to 0. ■

The significance of this result is that it demonstrates the robustness of the full revelation, no-delay equilibrium outcome to small noise in the senders' signals. Notice that there is *no common knowledge* about the state of the world (or even its discrete approximation) between the two senders no matter how precise their noisy signals are. In other words, the receiver is unable to detect lying (deviation from the equilibrium path) with probability 1. Nevertheless, the equilibrium outcome with noisy signals converges to the perfect-observation limit as the noise vanishes.

The idea behind our equilibrium construction is the following. At time 0 the senders privately tell each other which possible realization of ω_N their signals fall closest to. They discipline each other by implementing a delay in case their reports disagree, which induces truthful private disclosures in the first place. The receiver is kept in the dark about all this except that she can infer from a delay that the senders' signals were not sufficiently close to each other. Nevertheless, she waits until the delay is over in order to learn the senders' initial private reports. At that time, the senders tell the truth because they have mutual knowledge about their initial reports. "Optimistic" beliefs on the receiver's part help induce immediate full revelation just like in the perfect-observation case.

The construction works because the receiver's willingness to wait to find out the senders' signals is greater than the delay that the senders are willing to suffer in order to implement their own ideal points instead of the receiver's. If this were not the case then a sender could find it profitable to deviate even if that results in a delay equal to the receiver's willingness to wait. This is related to the criticism of the perfect-observation result mentioned at the end

of Section 3: If learning the senders' signals does not improve much the receiver's decision¹² while at least one of the senders is patient and strongly dislikes the receiver's ideal point, then the equilibrium with immediate full disclosure is not "credible" in the sense that the patient sender could simply "outwait" the receiver. This argument does not invalidate the result of the perfect-information model because immediate truthful reporting *is* and remains a Perfect Bayesian Equilibrium there. However, in a reasonable perturbation of the model like the one of this subsection—with imperfect signals and delay on the equilibrium path—this type of consideration is relevant.

The equilibrium construction provided in the proof of Proposition 3 relies on private communication between the senders. This is not an attractive feature because private communication could also lead to collusion between the senders and the emergence of inefficient equilibria. In the following result we show that the use of private communication is not necessary to obtain the result of Proposition 3. In particular, if the players have access to a randomization device that is commonly observed by the two senders then the equilibrium can rely on that device instead of synchronized private messages.

Next we describe a process by which the senders can let the receiver learn whether or not they agree which interval $[a_{k-1}, a_k)$ their signals belong to without actually reporting k or the extent of their disagreement. We call this procedure (or protocol) the " N -question reporting protocol with d delay".

Proposition 4 *Under the conditions of Proposition 3, the same equilibrium outcome can be achieved without private communication between the senders, but with the use of a semi-public randomization device.*

Proof. We only describe the players' moves on the equilibrium path. The proof of Proposition 3 establishes that these actions can be supported in equilibrium.

In the description of the protocol it is convenient to use an alternative notation for the intervals $[a_{k-1}, a_k)$ for $k = 0, \dots, 2^N$. Start with $S = [0, 1]$, and define the rest recursively as follows. Suppose that for all $(j - 1)$ -element sequences of 0 and 1, that is for all $A \in \{0, 1\}^{j-1}$, we have already constructed the

¹²For example, because her prior is already precise, and/or the senders' signals are not, and/or the receiver is impatient and/or does not really care about getting her ideal point.

intervals S_A , where $j \geq 1$ is an integer. Then, define S_{A0} and S_{A1} as the lower and upper halves of S_A , the midpoint going to the upper half. For example, $S_0 = [0, 1/2)$, $S_1 = [1/2, 1]$; $S_{00} = [0, 1/4)$, $S_{01} = [1/4, 1/2)$, and so on. The intervals S_A for $A \in \{0, 1\}^N$ coincide with $[a_{k-1}, a_k)$ for $k = 1, \dots, 2^N$.

The N -question reporting protocol consists of the following four steps.

Step 1. At time 0 the receiver generates $2N$ independent binary random numbers where each realization (0 or 1) has a 50% chance, so that the realizations become commonly known by the senders while the receiver observes nothing. Denote the first N bits by $\theta_1^1, \dots, \theta_N^1$ while the last N bits by $\theta_1^2, \dots, \theta_N^2$.

Step 2. At time 0, the senders send N rounds of binary reports; their reports are simultaneous in every round and the reports of previous rounds are publicly observed. In the first round, sender i reports $m_1^i = I((x_i \geq 1/2) \oplus (\theta_1^i = 1))$, where $I(\cdot)$ is the indicator function for logical statements and $\oplus$ is the exclusive disjunction (xor) operator. That is, sender i 's message contains a coded report whether or not his signal belongs to S_0 or S_1 using θ_1^i as the secret key.

Since sender $i' \neq i$ knows θ_1^i , he can infer whether or not $x_i \geq 1/2$ from m_1^i . On the other hand, the Receiver learns nothing about whether or not $x_i \geq 1/2$ by observing m_1^i but without knowing the realization of θ_1^i . This is so because by Bayes' rule,

$$\begin{aligned} \Pr(x_i \geq 1/2 | m_1^i = 0) &= \frac{(1/2) \Pr(x_i \geq 1/2)}{(1/2) \Pr(x_i \geq 1/2) + (1/2) \Pr(x_i < 1/2)} \\ &= \Pr(x_i \geq 1/2), \end{aligned}$$

and similarly,

$$\Pr(x_i \geq 1/2 | m_1^i = 1) = \Pr(x_i \geq 1/2).$$

In round $j > 1$, if the senders' past announcements indicate that both x_1 and x_2 belong to the same interval S_A where $A \in \{0, 1\}^{j-1}$, then sender i reports $m_j^i = I((x_i \in S_{A1}) \oplus (\theta_j^i = 1))$. If the senders' past reports indicate that x_1 and x_2 do not belong to the same S_A then both senders submit the random bits θ_j^1 and θ_j^2 from round j on.

Step 3. Still at time 0, the senders simultaneously publicly announce one more bit: $m_{N+1}^i = 1$ if and only if their reports in the previous N rounds indicate that x_1 and x_2 belong to the same interval S_A for some $A \in \{0, 1\}^N$.

If they report 0, a delay of length d is implemented: both senders babble until time d , and the receiver remains idle during this period.

Step 4: Either at time $t = d$ (if delay took place) or $t = 0$ (in case of no delay) both senders simultaneously announce $\theta_1^1, \dots, \theta_N^1, \theta_1^2, \dots, \theta_N^2$. The receiver “decodes” their coded binary messages, and carries out her ideal point given the information. ■

5 Conclusion

We have shown that there exists a Perfect Bayesian Equilibrium with instantaneous, complete revelation in a multi-sender dynamic cheap talk game where the state of nature is commonly known between the senders. In the static version of the model such equilibria only exist under certain assumptions on the state space and/or the senders’ biases. The “perfectness” of the equilibrium implies that the construction does not rely on exogenous commitment to delay on the Receiver’s part. Instead, the equilibrium is sustained by the Receiver’s “optimistic beliefs” that even if the senders have disagreed in the past they will agree in the future—which they must be able to as they have common knowledge.

We studied the robustness of this result to various perturbations of the model. In particular, we have shown that in certain environments where the senders have noisy signals about the state, a sequence of Perfect Bayesian Equilibria (with endogenous delay on the equilibrium path) converges to instantaneous full revelation as the noise in the senders’ signals vanishes. The key idea there is that even if the senders do not have common knowledge about the state it is possible for them to create mutual knowledge of their signals (at the expense of some delay), and then all their information can be truthfully reported to the Receiver in a perfect equilibrium.

Our result that adding dynamics gives rise to fully revealing equilibria when they do not exist in static models could be usefully applied in other communication games as well. Consider, for example, communication between a single sender and a receiver, where the sender can send *hard information*.¹³ Even

¹³A hard signal, as opposed to a soft one, can only be sent by the sender of a particular type. See Milgrom (1981).

though hard signals are available, a fully revealing equilibrium does not exist in the static game if the sender's preferences are not monotonic in the receiver's action or, more generally, if there does not exist a "worst type" for the sender.¹⁴ This is the case if the sender's type (the state of nature) is distributed on the circumference of a circle, and the receiver wants to match her action to the true state that is diametrically opposite to the sender's ideal point. Both parties' utilities are positive and decreasing in the distance of the action from their respective ideal points. In this setup, there does not exist a fully revealing equilibrium because by truthfully revealing the state the sender causes the receiver to implement the worse possible outcome for him. However, there exists a fully revealing PBE in the dynamic game where the receiver's belief is that the sender will send a state-revealing hard signal right away as long as he has not done so in the past. This indeed induces immediate full revelation on the sender's part.

Communication models where the senders observe noisy signals of the state—like the ones studied in Section 4—are realistic and offer a multitude of future research questions. For example, it would be interesting to study realistic dynamic discovery methods under imperfect signals that can involve delay even as the noise becomes arbitrarily small.

6 Appendix

Proof of Lemma 1. Suppose that sender 2 reports truthfully, $\hat{x}_2 = x_2$, and compute sender 1's expected utility from reporting $\hat{x}_1$ when his true signal is x_1 :

$$V_1(x_1, \hat{x}_1) = \int \int \left[- \left(\omega + b - \frac{\hat{x}_1 + \omega + z}{2} \right)^2 - 2b(\hat{x}_1 - \omega - z)^+ \right] dF(\omega|x_1) dF(z|0),$$

where $F(\cdot|\mu)$ is the cdf of a normal random variable with mean μ and variance σ^2 , and the notational convention $a^+ = a\mathbf{1}_{\{a \geq 0\}}$ is used. Maximizing this in $\hat{x}_1$

¹⁴See Milgrom (1981), Seidmann and Winter (1997).

yields the first-order condition

$$\int \int \left(b + \frac{\omega - \hat{x}_1 - z}{2} - 2b \mathbf{1}_{\{\omega \leq \hat{x}_1 - z\}} \right) dF(\omega|x_1) dF(z|0) = 0,$$

which needs to hold at $\hat{x}_1 = x_1$ for incentive compatibility. (It is easy to check that the second-order condition holds.) Using the facts that $\int \omega dF(\omega|x_1) = x_1$, $\int z dF(z|0) = 0$, and $\int \mathbf{1}_{\{\omega \leq x_1 - z\}} dF(\omega|x_1) = F(x_1 - z|x_1) = F(-z|0)$, we can rewrite this condition as

$$b - 2b \int F(-z|0) dF(z|0) = 0. \quad (7)$$

The incentive constraint for sender 2 (provided sender 1 reports truthfully) can be derived similarly. His deviation payoff is

$$V_2(x_2, \hat{x}_2) = \int \int \left[- \left(\omega - b - \frac{\hat{x}_2 + \omega + z}{2} \right)^2 - 4b(\omega + z - \hat{x}_2)^+ \right] dF(\omega|x_2) dF(z|0).$$

Maximization in $\hat{x}_2$ at $\hat{x}_2 = x_2$ yields,

$$\int \int \left(-b + \frac{\omega - x_2 - z}{2} + 2b \mathbf{1}_{\{\omega \geq x_2 - z\}} \right) dF(\omega|x_2) dF(z|0) = 0,$$

which simplifies to

$$-b + 2b \int [1 - F(-z|0)] dF(z|0) = 0. \quad (8)$$

Since the integrands in equations (7) and (8) add up to 1 (for all z), by the symmetry of the normal distribution and $1 - F(-z|0) = F(z|0)$ we have

$$\int F(-z|0) dF(z|0) = \int [1 - F(-z|0)] dF(z|0) = \frac{1}{2}.$$

Hence both first-order conditions hold as claimed. ■

References

- [1] Ambrus Attila and Satoru Takahashi (2007): “Multi-sender Cheap Talk with Restricted State Spaces,” *Theoretical Economics*, forthcoming.
- [2] Artemov, Georgy (2006): “Imminent Nash Implementation”, *mimeo*
- [3] Austen-Smith, David (1990): “Information Transmission in Debate”, *American Journal of Political Science*, 34, 124-152.
- [4] Austen-Smith, David (1993): “Interested Experts and Policy Advice: Multiple Referrals under Open Rule,” *Games and Economic Behavior*, 5, 3-44.
- [5] Austen-Smith, David and Jeffrey Banks (2000): “Cheap Talk and Burned Money,” *Journal of Economic Theory*, 91, 1-16.
- [6] Ausubel, Larry and Raymond Deneckere (1987): “One is Almost Enough for Monopoly,” *Rand Journal of Economics*, 18, 255-274.
- [7] Battaglini, Marco (2002): “Multiple Referrals and Multidimensional Cheap Talk,” *Econometrica*, 70(4), 1379-1401.
- [8] Battaglini, Marco (2004): “Policy Advice with Imperfectly Informed Experts,” *Advances in Theoretical Economics* 4(1).
- [9] Che, Yeon-Koo, and József Sákovics (2004): “A Dynamic Theory of Holdup,” *Econometrica*, 72(4), 1063-1103.
- [10] Crawford, Vincent and Joel Sobel (1982): “Strategic Information Transmission,” *Econometrica*, 50(6), 1431–1451.
- [11] Damiano, Ettore, Li Hao, and Wing Suen (2007): “Delay in Strategic Information Transmission”, *mimeo*, 2007.
- [12] Gilligan, Thomas W. and Keith Krehbiel (1987): “Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures”, *Journal of Law, Economics, & Organization*, 3(2), 287-335.

- [13] Gilligan, Thomas W. and Keith Krehbiel (1989): “Asymmetric Information and Legislative Rules with a Heterogenous Committee”, *American Journal of Political Science*, 33, 450-490.
- [14] Gul, Faruk (1987): “Noncooperative Collusion in Durable Goods Oligopoly,” *Rand Journal of Economics*, 18, 249-254.
- [15] Gul, Faruk (2001): “Unobservable Investment and the Holdup Problem,” *Econometrica*, 69, 343-376.
- [16] Kartik, Navin (2007): “A Note on Cheap Talk and Burned Money,” *Journal of Economic Theory*, 136(1), 749-751.
- [17] Krishna, Vijay and John Morgan (2001a), “Asymmetric Information and Legislative Rules: Some Amendments,” *American Political Science Review*, 95(2), 435-452.
- [18] Krishna, Vijay and John Morgan (2001b), “A Model of Expertise,” *Quarterly Journal of Economics*, 116(2), 747-775.
- [19] Krishna, Vijay and John Morgan (2004), “The Art of Conversation, Eliciting Information from Experts through Multi-Stage Communication,” *Journal of Economic Theory*, 117(2), 147-179.
- [20] Lockwood, B. and J. P. Thomas (2002): “Gradualism and Irreversibility”, *Review of Economic Studies*, 69, 339-356.
- [21] Marx, Leslie and Steven Matthews (2000): “Dynamic Voluntary Contribution to a Public Project,” *Review of Economic Studies*, 67, 327–358.
- [22] Milgrom, Paul (1981): “Good News and Bad News: Representation Theorems and Applications”, *Bell Journal of Economics*, 21, 380-391.
- [23] Rubinstein, Ariel (1982): “Perfect Equilibrium in a Bargaining Model”, *Econometrica* 50 (1982), 97-110.
- [24] Rubinstein, Ariel and Asher Wolinsky (1992): “Renegotiation-Proof Implementation and Time Preferences,” *American Economic Review*, 82, 600–614.

- [25] Seidmann, Daniel and Eyal Winter (1997): “Strategic Information Transmission with Verifiable Messages”, *Econometrica*, 65(1), 163-169.
- [26] Serrano, Roberto (2007): “Bargaining”, *The New Palgrave Dictionary of Economics*, forthcoming
- [27] Wolinsky, Asher (2002): “Eliciting Information from Multiple Experts”, *Games and Economic Behavior* 41, 141-160.