

Wait and See: A Theory of Communication Over Time*

Péter Eső[†] and Yuk-fai Fong[‡]

January 2010

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. Delays are wasteful, and no player can commit to any action or inaction; the receiver can choose momentary inaction only if her beliefs about the continuation play rationalize that. In contrast to the results in static versions of the model, we show that when the senders commonly know the state of nature, a perfect Bayesian equilibrium exists with instantaneous, full revelation irrespective of the size and direction of the senders' biases. We show that the equilibrium outcome is robust to the introduction of a type of small noise in the senders' signals about the state.

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

JEL Numbers: D82, D83, D72, D74

*First version: November 2007. We thank Ricardo Alonso, Attila Ambrus, David Austen-Smith, Olivier Compte, Tim Feddersen, Drew Fudenberg, Li Hao, Joseph Harrington, Johannes Hörner, Andrea Prat, Philip Reny, Larry Samuelson, James Schummer, Joel Sobel, Balázs Szentes, Satoru Takahashi, seminar participants at the universities of Arizona, Arizona State, Boston, Brown, Chicago, Columbia, Harvard–MIT, Iowa State, Johns Hopkins, LSE, Northwestern, Queen's, Southern Methodist, UBC, Yale, and participants at various conferences for comments, and Renato Gomes for excellent research assistance.

[†]University of Oxford, peter.eso@economics.ox.ac.uk

[‡]Northwestern University, 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 in a truly *dynamic setting*. We investigate how the decision maker’s self-fulfilling belief in *learning more by waiting* (or taking other temporary, suboptimal actions) can induce speedier and more informative communication from biased experts.

In the model, a decision maker (she) consults multiple experts who have private information about the payoff-relevant state of nature. There is a single decision to be made, and neither the state nor the experts’ signals and biases change over time.¹ At every point in time, first the experts can send public messages, then the receiver either picks a game-ending policy, or chooses momentary inaction. When the game ends, the players receive payoffs representing state-dependent, single-peaked preferences. Conflict arises because conditional on the state, the senders’ ideal policies differ from the receiver’s. All players have strict time preferences, and play a perfect Bayesian equilibrium, requiring consistency and sequential rationality. No player has commitment power—in particular, the decision maker cannot commit to future actions (e.g., choosing a policy with a certain delay).

This is a dynamic (though not repeated), *multi-sender cheap talk game* where the receiver continuously faces the tradeoff between more communication and earlier decision, and the time spent on communication (or delay) adversely affects all players’ payoffs. These features distinguish it from traditional *static* models and also *long cheap talk*.² It will also become clear that “time” in our model is not simply a new dimension in which the receiver can threaten the senders with punishment (albeit without being able to commit); the fact that *beliefs can evolve over time* in response to communication is crucial to our analysis.

¹The state, the senders’ information and the preferences do not evolve over time so that we can focus on the *dynamics of communication* in this model.

²The “static” version where the senders have *common knowledge about the state* is studied by Gilligan and Krehbiel (1989), Krishna and Morgan (2001a,b), Battaglini (2002), Ambrus and Takahashi (2007), among others. Battaglini (2004) allows a certain type of noisy signal structure in the same setup. Multiple rounds of communication are studied by Aumann and Hart (2003), Krishna and Morgan (2004), and Ambrus and Takahashi (2007), among others.

For a concrete situation that we intend to model, think of the policy response to the financial crisis in the United States. The decision maker is the median voter in Congress; the senders are various committees, representatives and lobbyists who have relevant private information.³ The honest opinions of these experts may differ due to their information, but based on what they know it would be possible to determine Congress' ideal policy choice along multiple dimensions such as government spending, tax relief, financial regulation. However, the senders are also *biased* from the median voter's perspective due to their constituents' interests, ideology, 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—assuming that delays (or certain temporary policies) are wasteful, and that the median voter cannot explicitly commit to inaction or any future actions. Other legislative issues in the U.S. that share these characteristics range from entitlement and health care reform to the treatment of illegal immigrants.

We first study the case 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 and the players' preferences are, if public communication can take place sufficiently frequently, then there exists a perfect Bayesian equilibrium where the state of nature is immediately disclosed, and the decision maker carries out her ideal policy without delay. The senders' equilibrium strategy is to report the state at time 0, and repeat it frequently (every period in discrete time, or at given intervals in continuous time). The receiver takes an action matching the senders' report as soon as they agree at a time when they are expected to be truthful. Her off-equilibrium inaction is rationalized by the expectation (shared among all parties) that both senders will report the true state frequently and almost immediately following a disagreement. Note that the receiver does not have the power to commit to wait; she chooses inaction only if she expects to learn more by waiting. Off the equilibrium path, she believes that learning the state is imminent, hence her momentary inaction is rational.

³Legislative procedures are modeled as cheap talk games in Gilligan and Krehbiel (1989).

The result that the receiver can *immediately* extract *all* information from biased experts is remarkable because in the corresponding *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.⁴ In contrast, our result holds *without any restriction* on the shape or size of the state space, or the direction and 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.

Notice that the decision maker can fully extract the senders’ private information even though she is unable to commit to delay her decision. She can achieve the first-best by rationally maintaining a positive attitude and believing that the senders will immediately agree following any disagreement. Such beliefs are plausible in the model where the senders have common knowledge of the state of nature—they *must* be able to agree, and the receiver knows that.⁵

Experts do not always have common knowledge about the state of nature. If so, then “cross-checking” their reports and only acting when they agree may not be a sensible course of action for the decision maker. This criticism applies not only to our model under perfect observation, but also to a large chunk of the existing literature on multi-sender cheap talk. In those models, too, the senders perfectly observe the state, and the construction of fully-revealing equilibria relies on this fact.⁶ Another drawback of perfect-observation models is that they are not particularly insightful for analyzing situations with more than two senders. In that case, even without dynamics, a fully-revealing equilibrium trivially exists: If more than two senders are expected to report the state, then a single deviator

⁴The condition identified in the cited papers is discussed in Section 3.1.

⁵Similar equilibrium constructions for dynamic models have been used outside the cheap talk literature. 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 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.

⁶Models where the experts do *not* have common knowledge about the state of nature include Austen-Smith (1993), Wolinsky (2002), Battaglini (2004).

cannot prevent the receiver from learning the truth.

We extend our model to allow for noisy signals on the senders' part. The main difficulty in constructing informative equilibria in this case is that disagreement between the senders can happen on the equilibrium path, and so it is difficult to detect untruthful reporting. If the receiver threatens with costly delay as a function of the senders' disagreement and this threat induces truthful reports, then she has no incentive to carry out the punishment once she receives the senders' truthful reports and they actually disagree.

Despite these difficulties we prove that in our model with noisy signals, under certain conditions, there exists an equilibrium with *imperfect revelation* and *positive expected delay*. This equilibrium is sustained by the senders not revealing everything they know all at once, thereby motivating the receiver to wait along the equilibrium path. The delay anticipated by the senders depends on the degree of their disagreement, which in turn provides incentives for them not to misrepresent their signals. As the senders' signals become arbitrarily precise, the equilibrium outcome converges to the one found under perfect observation, that is, full revelation and no delay.⁷ Our result shows that the immediate, fully-revealing equilibrium under perfect observation is not an aberration as it is the limit of sensible equilibria as the noise in the senders' signals vanishes.

These results shed light on whether and how delay—to which the decision maker cannot commit—can induce the experts to reveal more information. The insights are applicable in our first motivating example, legislation in Congress: Repeated hearings and procedural delays in the absence of agreement may enhance communication and hence the quality of the median voter's decision.

Another similar, practical situation where our model applies is one in which a court needs to decide in a dispute between parties with opposing interests. The parties have information about the state, 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 costly delay; that can only be the result of the court's equilibrium beliefs that

⁷We prove the existence of the sequence of such equilibria under conditions involving the receiver's eagerness to learn the state of nature in comparison to the senders' biases and impatience, as well as assumptions imposed on the signal structure.

prolonged questioning will lead to information revelation and a socially beneficial outcome. Indeed it will, as our results show.⁸

A similar situation is that of law enforcement questioning multiple suspects of a crime. The suspects have correlated information about the facts, and they can be interrogated multiple times. Repeated questioning decreases the suspects' utilities as they have to forgo other activities while in custody. However, most interrogators cannot explicitly commit to delay as detaining suspects without charge is costly and may even be illegal, depending on the circumstances. We show that persistent interrogation can be a credible (and optimal) strategy despite the interrogator's inability to commit.

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, the availability of instruments that make waiting costs different would only strengthen our results.

As we already pointed out, our paper's main contribution is to the literature on multi-sender cheap talk games. This literature is partly motivated by applications in political theory, and focuses on equilibria and institutions that facilitate information transmission among experts and a decision maker.⁹ In the static version of the model that we study, in the special case of perfect observation, Battaglini (2002) and Ambrus and Takahashi (2007) derive necessary-sufficient conditions for the existence of fully-revealing equilibria. Aumann and Hart (2003) and Krishna and Morgan (2004) argue that multiple rounds of communication can in general expand the set of equilibrium outcomes. However, in the multi-sender cheap talk model that we study, under perfect observation, Ambrus and Takahashi (2007) show that large biases still make it impossible to sustain a

⁸A real-world example is the recent patent dispute between Research in Motion (RIM, maker of Blackberry phones) and NTP (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 settled in March 2006. Delay costs affected both parties (RIM risked losing business, NTP risked losing its patents to invalidation). Commentators found the settlement appropriate (see http://money.cnn.com/2006/03/03/technology/rimm_ntp).

⁹See also Austen-Smith (1990, 1993), Wolinsky (2002) and their references.

fully-revealing equilibrium even with long cheap talk. In contrast, we obtain an *unconditional possibility* result in our baseline dynamic model, in the comparable case of perfect observation.

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 helps players improve efficiency in a dynamic voting game. Compared to these papers our main contribution is to show that the receiver can credibly carry out the delay even without commitment and that costly delay can be avoided in equilibrium as the senders' signals become more precise.

Our dynamic model is marginally related to cheap talk models with money-burning (see Austen-Smith and Banks (2000), Kartik (2007)). 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, if the sender's budget is unlimited, then there exists an equilibrium where the sender signals a higher state of nature by burning more money, and all types separate. 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 where 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 additional results with imperfect observation.

2 The model

We study a dynamic, incomplete information game where the players are two experts (the senders, $i = 1, 2$) and one decision maker (the receiver, $i = 0$). Time is discrete with $\delta > 0$ increments and an infinite horizon: $t = 0, \delta, 2\delta, \dots$. Of particular interest is the case where δ is small.¹⁰

States and signals. Before the game starts, a state of nature, ω , is drawn from a closed set $S \subseteq \mathbb{R}^n$. The senders observe private signals about ω represented by random variables X_i , $i = 1, 2$. For simplicity, assume that the realizations of the signals belong to the state space as well. The joint distribution of (ω, X_1, X_2) is commonly known; neither the state nor the senders' signals change over time.

Timing, actions, strategies. At the beginning of every period, each sender (he) simultaneously sends a public message, $m_i^t \in M_i$ for $i = 1, 2$, after which the receiver (she) chooses a policy, $y^t \in Y$. The message sets contain at least all elements of S and a “null” message representing no communication, that is, $S \cup \{\emptyset\} \subseteq M_i$ for $i = 1, 2$.¹¹

The policy space Y includes at least the entire convex hull of S and action $y = \emptyset$. The assumptions on the receiver's payoffs made below ensure that for any beliefs about the state, the receiver's ideal policy falls in the convex hull of S . Policy $y = \emptyset$, referred to as (momentary, one-period) “inaction” should be interpreted as *any temporary, suboptimal action*, not supported by any beliefs of the receiver. The set Y may also include other policies in \mathbb{R}^n , e.g., the senders' ideal policies that may not be optimal for the receiver given any beliefs either.

Sender i 's strategy σ_i assigns a message m_i^t to the realization of X_i and the public history $h^t = (m_1^\tau, m_2^\tau, y^\tau)_{\tau < t}$ for all $t = 0, \delta, 2\delta, \dots$, while the receiver's strategy σ_0 maps h^t to y^t for all t .

Payoffs. We assume that as soon as the receiver picks an action in the convex hull of S (i.e., chooses a policy that is rational for her given some beliefs about the state), the game ends.¹² Each player's payoff, U_i for $i = 0, 1, 2$, depends

¹⁰We consider continuous time in Section 3.4.

¹¹We could allow, but do not require nor model, multiple rounds of communication in each period. Related to this, we discuss sequential reporting in Section 3.3.

¹²This is an innocuous simplification; an alternative specification where the game ends with the same probability no matter what action is chosen by the receiver is discussed below.

on the state, the game-ending policy, and the physical time. Assume that for $i = 0, 1, 2$, $U_i : \mathbb{R}^n \times \mathbb{R}^n \times [0, \infty) \rightarrow \mathbb{R}$ is continuous and bounded for any t .¹³ However, $\lim_{t \rightarrow \infty} U_i(\omega, y, t)$ need not be bounded. Throughout the game players maximize their expected payoffs given their available information.

The players' preferences over policies only depend on the state of nature: For all ω, y, y', t and t' in the relevant domains, $U_i(\omega, y, t) - U_i(\omega, y', t) = U_i(\omega, y, t') - U_i(\omega, y', t')$. Moreover, for all ω and t , $U_i(\omega, y, t)$ is concave and single-peaked in $y \in \mathbb{R}^n$. Hence, for all $i = 0, 1, 2$ and $\omega \in S$, there exists $y_i(\omega) \in \mathbb{R}^n$ such that $U_i(\omega, y_i(\omega), t) > U_i(\omega, y, t)$ for all $y \neq y_i(\omega)$ and all t . We normalize the receiver's ideal point as the state, $y_0(\omega) \equiv \omega$. Hence no matter what the receiver believes about the state, her best policy is indeed in the convex hull of S .

Time enters the players' utility functions in the form of disliking certain types of delays. In particular, we assume that in any state ω , all players prefer the decision maker to carry out action $y = \omega$ at time t rather than at any $t' > t$. (The senders may still dislike $y = \omega$ at t compared to other policies at t or $t' > t$.) Formally, for all i, ω , and $t' > t$, we assume $U_i(\omega, \omega, t) > U_i(\omega, \omega, t')$. This is exactly what it means that action $y = \emptyset$ ("inaction") is a suboptimal policy: All players (the receiver and both senders) prefer $y = \omega$ to be carried out earlier. This type of preferences could result from discounting of a positive utility function or a positive cost of waiting, but we do not need to make such specific assumptions.

Finally, we define the maximum delay that the receiver is willing to bear in order to learn the state of nature and carry out her corresponding ideal point. Let Δ_0 solve

$$E[U_0(\omega, \omega, t + \Delta_0)] = \max_{y \in Y} E[U_0(\omega, y, t)], \quad (1)$$

if such solution exists, or $\Delta_0 = \infty$ if it does not. By assumption, the left-hand side of equation (1) exceeds its right-hand side at $\Delta_0 = 0$ and it is decreasing in Δ_0 . If finite, Δ_0 is increasing in the receiver's patience and the variance (or other measures of diffusion) of the prior distribution of ω . 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 on ω .

¹³We define U_i on the domain $\mathbb{R}^n \times \mathbb{R}^n \times [0, \infty)$ instead of $S \times \text{co}(S) \times \{0, \delta, 2\delta, \dots\}$ for convenience, e.g., in order to accommodate the continuous-time extension in Section 3.4.

Equilibrium concept. In the game defined above we study perfect Bayesian equilibria (see Fudenberg and Tirole (1991)). Practically, this concept requires that at any point in time, given the player’s private information and the history of the play, no player has an incentive to deviate from his or her prescribed strategy. No player can commit to any action in the future; in particular, *the receiver cannot commit to delay or any policy*.¹⁴

Notes on an alternative specification of the model. We have assumed that any policy choice that is optimal for the receiver given some beliefs over the state ($y \in \text{co}(S)$) immediately ends the game, but $y = \emptyset$ (a given, suboptimal policy) does not. This is just a convenient and innocuous simplification of the model setup, and it allows us to interpret $y = \emptyset$ as “costly, one-period delay”. However, it is important to note that there are other ways to think about our model and that $y = \emptyset$ need not be interpreted as receiver inaction.

We could alternatively assume that the receiver chooses a policy $y \in Y$ every period and the game ends with the same, exogenous probability no matter whether $y \in \text{co}(S)$ or $y \in Y \setminus \text{co}(S)$. Each player’s per-period payoff, $u_i(\omega, y^t)$ for $i = 0, 1, 2$, is bounded, accumulates (with discounting), and is realized when the game ends.¹⁵ Assume that for each i , the payoff function $u_i(\omega, y^t)$ admits a unique optimum in $y^t \in Y$ given ω . In this version, the key assumption is that there exists an element of $Y \setminus \text{co}(S)$, denoted $y = \emptyset$, such that for all $i = 0, 1, 2$, $\omega \in S$ and $y \in \text{co}(S)$, $u_i(\omega, \emptyset) < u_i(\omega, y)$. The action $y = \emptyset$ is therefore just any temporary policy that is worse than the receiver’s ideal policy, for all players in all states.

We mention this alternative specification of the model because in certain applications it may be evident that such suboptimal policies exist, while it may be less clear that delays hurt all players. (For example, in some applications the experts may benefit from a prolonged decision-making process.) All of our results carry over to this alternative version of the model.

¹⁴In every period, the receiver can choose momentary inaction only if her beliefs regarding the continuation play rationalize it.

¹⁵The game could be literally infinite-horizon, in which case payoffs are obtained at $t = \infty$.

3 The perfect observation case

In this section we assume that each sender observes the state of nature, i.e., $X_i = \omega$ for $i = 1, 2$. We look for the most informative perfect Bayesian equilibrium; in particular, we determine whether or not there exists a *fully-revealing equilibrium*, where the senders report ω and the receiver carries out $y = \omega$.

For comparison, we first review the benchmark results in the *static* version of the game. There, Battaglini (2002) and Ambrus and Takahashi (2007) identified a necessary-sufficient condition for the existence of a fully-revealing equilibrium. After discussing this result we turn to our possibility results in the dynamic version of the same game. We extend the analysis to continuous time, and discuss the existence of other equilibria at the end.

3.1 The static benchmark

Assume, in Section 3.1 only, that the entire game is played at a single point in time (at $t = 0$), in two stages. First, the senders observe ω , and simultaneously send messages $m_i \in M_i$. Second, the receiver picks an action $y \in Y$. Note that in the static game, in any perfect Bayesian equilibrium, the policy chosen by the receiver must be optimal for some beliefs about the state, hence belong to $co(S)$. That is, sequential rationality rules out that $y = \emptyset$, or any other policies in $Y \setminus co(S)$, are chosen in equilibrium.

The question posed in the literature on static, multi-sender cheap talk games is whether there exists a perfect Bayesian equilibrium in which, by comparing the senders' messages, the receiver can infer the true state.

Denote the set of receiver-actions that sender i weakly prefers to $y = \omega$ in state ω by $B_i(\omega)$. With the notation introduced in Section 2, $B_i(\omega) = \{y \in co(S) : U_i(\omega, y, 0) \geq U_i(\omega, \omega, 0)\}$, $i = 1, 2$. The exact condition for the existence of a fully-revealing perfect Bayesian equilibrium is the following.

Proposition 1 (Battaglini (2002), Ambrus and Takahashi (2007)) *A fully-revealing perfect Bayesian equilibrium exists in the static game if, and only if, for all $\omega, \omega' \in S$, $co(S) \not\subseteq B_1(\omega') \cup B_2(\omega)$.*

This result is due to Battaglini (2002) for one-dimensional state spaces; Ambrus and Takahashi (2007) observed that it holds irrespective of the dimensionality of the state space. The intuition is simple. Consider a fully-revealing equilibrium and suppose that sender 1 sends a message consistent with the state being ω while sender 2 reports as if the true state were ω' . The receiver's response to such inconsistent reports, $y(\omega, \omega')$, must lie in $co(S)$. However, if $B_1(\omega') \cup B_2(\omega)$ contains $co(S)$ then at least one of the senders has a profitable deviation in some state: If $y(\omega, \omega') \in B_1(\omega')$ then sender 1 prefers to pretend it is state ω in state ω' , while if $y(\omega, \omega') \in B_2(\omega)$ then sender 2 prefers to report ω' in state ω . Therefore there is no fully-revealing equilibrium. In contrast, if the condition in Proposition 1 holds, then it is easy to construct an fully-revealing equilibrium: The senders announce the state and, if they agree, the receiver carries out $y = \omega$; if they disagree (sender 1 reports ω while sender 2 reports ω') then the receiver picks any $y \in co(S) \setminus (B_1(\omega') \cup B_2(\omega))$.




Figure 1: Non-existence of FRE in the static model

While the receiver can use the fact that the senders must be able to agree on ω , there is a limitation on her ability to use “cross-checking” to support a fully-revealing equilibrium of the static game. If the senders’ biases are sufficiently large relative to the state space (like in Figure 1), then it is impossible for the receiver to punish both of them for some disagreements, and so there does not

exist a fully-revealing equilibrium. It is important to note that the dimensionality of the state space makes no difference to the analysis or the necessary-sufficient conditions.

Multiple rounds of communication (at a single physical point in time, i.e., without costly delay) can expand the set of equilibria, as pointed out by Aumann and Hart (2003), Krishna and Morgan (2004), and, in the multi-sender cheap talk model, by Ambrus and Takahashi (2007). However, the latter authors also show that when the state space is bounded and the biases are sufficiently large, fully-revealing equilibria do not exist even if long cheap talk is allowed.

Our dynamic setup is different because (i) delays are costly, and (ii) the receiver can time her decision to leave more or less time for communication. The timing of the receiver’s action could be thought of as another dimension of her decision problem. However, the receiver’s preference for earlier action is commonly known—the “time-coordinate” of the state of nature is always zero. Therefore, in a game where a new policy dimension corresponding to “time” is added to the static model, the receiver cannot rationally choose (on or off the equilibrium path) a positive delay.¹⁶ This means that introducing a new dimension to the standard, *static* problem with the properties of “delay” would not guarantee the existence of a fully-revealing equilibrium.

3.2 Unconditional possibility in the dynamic game

The main result of this section is that in our dynamic multi-sender cheap talk game, if the period-length of discrete time is sufficiently short, then there exists an equilibrium where *all information* commonly known by the senders is *immediately* revealed. In this equilibrium, the senders report the state and the receiver picks her ideal policy right away. In case of disagreement, the receiver believes that the senders will agree in the next period, which leads her to choose momentary inaction, and in turn induces truthtelling on the senders’ part.

Proposition 2 *Assume that the senders perfectly observe the state: $X_1 = X_2 = \omega$. If $\delta \leq \Delta_0$, then there exists a perfect Bayesian equilibrium where the state*

¹⁶That is, even if the policy space is extended to $Y \times [0, \infty)$, only the policies in the subset $co(S) \times \{0\}$ are sequentially rational for the receiver.

of the world is immediately, publicly reported by both senders, and the receiver carries out $y = \omega$ without delay.

Proof. The equilibrium is constructed as follows. Sender $i \in \{1, 2\}$ reports $m_i^t = \omega$ for all t , irrespective of history. For all t such that the senders report the same state, the receiver picks the corresponding policy. If the senders disagree, the receiver chooses inaction. The outcome of the proposed strategies is immediate full revelation and $y^t = \omega$ at $t = 0$.

Neither sender can do better than to report ω after any history because the other sender reports ω and the receiver is expected to wait until they agree. If the senders disagree at t , the receiver maintains her prior beliefs about ω . Her best response is inaction because she expects to learn the state at time $t + \delta$, and $E[U_0(\omega, \omega, t + \delta)] \geq \max_{y \in Y} E[U_0(\omega, y, t)]$ by $\Delta_0 \geq \delta$ and equation (1). ■

Proposition 2 establishes that when the receiver is willing to wait one “tick of the clock” in order to make a fully-informed decision, there exists an immediate, fully-revealing equilibrium. The condition $\delta \leq \Delta_0$ is guaranteed to hold if the discrete time increments are sufficiently small. Indeed, we consider δ being arbitrarily close to zero as the most compelling assumption, because it implies that the communication between the senders and the receiver can be arbitrarily frequent, and that the receiver cannot commit to any amount of physical time delay. In the next subsection we also consider the case of truly continuous time, which is conceptually different from the case of discrete time with very short period lengths.

Since the equilibrium in Proposition 2 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. As we discussed it in Section 3.1, a fully-revealing equilibrium only exists under certain conditions in the static environment, even with multiple rounds of communication. 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.

It may be obvious to point out that our possibility result relies on the unboundedness of the time horizon. If there exists a last point in time by which the

receiver has to make a decision, then the sender(s) may be better off by silently waiting until then rather than divulging any information about the state. Of course, our results remain valid if either the stopping time is random (and the probability of continuation is sufficiently high), or if the senders grow infinitely impatient as the end of time approaches.

In the equilibrium constructed in the proof of Proposition 2, the receiver does not update her beliefs about the state off the equilibrium path. This is consistent with perfect Bayesian equilibrium, but may be unattractive on intuitive grounds. For example, the receiver may believe that in case of disagreement, only one of the senders was untruthful. This motivates the question whether the immediate, fully-revealing equilibrium outcome can be implemented such that the receiver's beliefs are *continuous* off the equilibrium path at time 0. We now show how to modify the equilibrium construction to achieve this goal.

In order to simplify the exposition of this construction, suppose that ω is uniform on $[0, 1]$, and allow *two rounds of communication* per period. Define $(\omega_A, \omega_B) \in \{0, 1\} \times [0, 0.5]$ such that $\omega \equiv 0.5\omega_A + \omega_B$. Consider the following strategies: In the first round of talk at $t = 0$, and in the first round of every period until the reports match, each sender reports ω_B . From the time their reports first agree, the senders start reporting ω_A . The receiver chooses inaction until the senders report the same $\omega_B \in [0, 0.5]$ and subsequently the same $\omega_A \in \{0, 1\}$; when they agree on both, she carries out $y = 0.5\omega_A + \omega_B$.

If for all ω_B , the receiver is willing to wait at least one period to find out whether $\omega = \omega_B$ or $0.5 + \omega_B$ conditional on knowing ω_B , then these strategies form an immediate, fully-revealing equilibrium such that the receiver's beliefs are continuous in the senders' reports at $t = 0$. To see this, suppose that at $t = 0$ the senders' reports are almost (but not exactly) equal, and assume that the receiver believes one of the senders is truthful, ensuring the continuity of beliefs. Although the receiver can essentially infer ω_B , she does not know whether $\omega = \omega_B$ or $0.5 + \omega_B$. She anticipates that in the following period the senders will both report ω_B , and then immediately ω_A , so she will learn ω by the end of the next period. By assumption, it is worth for her to wait one period.

3.3 Extension to continuous time, unbounded horizon

When time is continuous, players can move asynchronously and there is no “next period”. Indeed, the latter property causes a fundamental problem in continuous-time repeated or dynamic games: unrestricted strategies may not determine a unique outcome (see Bergin and Macleod (1993)). This issue is also present in our setup with continuous time.¹⁷

One, commonly-used solution is to require that strategies exhibit “inertia”: A player must keep playing an action for a short time before switching. The length of the “short period” can depend on the public history, the player’s signal and action. In line with this, we now give a proper definition of our multi-sender cheap talk game in *continuous time*.¹⁸

First, we define histories and strategies (mapping histories to actions) so that within any finite interval of time, each player can change his action on only countably many occasions. Let $h_i : [0, t) \rightarrow M_i$ for $i = 1, 2$ the *public history* of sender i ’s play up to time $t > 0$. We call h_i admissible if $h_i^{-1}(M_i)$ is a countable collection of intervals of the form $[\tau, \tau') \subseteq [0, t)$. Denote the set of all finite, admissible public histories by H and the restriction of $h = (h_1, h_2) \in H$ to $[0, t)$ by $h|t$. A (pure) *strategy* for sender i maps the realization of his signal and the public history to a message, $f_i : S \times H \rightarrow M_i$. We assume that f_i has *inertia*: For all $x_i \in S$, $t \geq 0$ and $h \in H$, if $f_i(x_i, h|t) = m_i$, then there exists $\delta > 0$ such that $f_i(x_i, h|\tau) = m_i$ for all $\tau \in [t, t + \delta)$.¹⁹ The receiver’s strategy is $f_0 : M_1 \times M_2 \times H \rightarrow Y$. The game ends at the earliest time t such that $y_0(m_1, m_2, h|t) \in co(S)$. Therefore, the inertia condition for the receiver is that for all $t \geq 0$, $h \in H$ and $(m_1, m_2) \in M_1 \times M_2$, if $f_0(m_1, m_2, h|t) \notin co(S)$ then

¹⁷To see this, suppose that each sender observes the state and uses the following strategy : “Play $m_i = \emptyset$ at $t = 0$ and as long as both senders were silent for all $t' < t$; send $m_i = \omega$ if either sender has sent ω at any $t' < t$.” Assume the receiver’s strategy is to choose inaction at $(t, 1)$ if the senders were silent at $(t, 0)$, and implement $y = m_1$ if they agreed on a $\omega \in S$ at $(t, 0)$. Then, for any $t^* > 0$, both senders playing $m_i = \emptyset$ for all $t \in [0, t^*]$ and reporting ω for all $t > t^*$ is consistent with their strategies making the outcome indeterminate.

¹⁸Our setup parallels that of Bergin and MacLeod (1993) for continuous-time repeated games. Somewhat stronger restrictions are imposed in the continuous-time bargaining games of Perry and Reny (1993, 1994).

¹⁹In the discrete-time version of model, δ is the length of a time period; there it cannot depend on the history of the play and the identity of the player.

there exists $\delta > 0$ such that $f_0(m'_1, m'_2, h|\tau)$ is constant for all $\tau \in [t, t + \delta)$ and $(m'_1, m'_2) \in M_1 \times M_2$. In words, the receiver cannot terminate the game at the “first moment” after a certain history. Denote F_i the set of inertia strategies for player $i \in \{0, 1, 2\}$, and let $F = \prod_{i \in \{0, 1, 2\}} F_i$.

The assumption that the strategies have inertia does not mean the players can commit to future actions (e.g., play their current action for a given period of time). This is so because F_i contains *every* strategy with arbitrarily small inertia. Whether or not commitment is feasible is determined by the equilibrium concept, which is discussed below.

The inertia assumptions are made purely for technical reasons, in order to guarantee that any combination of strategies $f = (f_0, f_1, f_2) \in F$ determines a unique outcome.²⁰ We denote the play induced by the strategy-tuple f by $\bar{h}(f)$, the physical time when it terminates by $\bar{t}(f) \in [0, \infty]$, and the receiver’s action at $\bar{t}(f) < \infty$ by $\bar{y}(f)$. We continue to use perfect Bayesian equilibrium. Formally, $f^* \in F$ is a perfect Bayesian equilibrium if for all $t \in [0, \infty)$, $i \in \{0, 1, 2\}$ and $f_i \in F_i$, $E[U_i(\omega, \bar{y}(f^*), \bar{t}(f^*)) | X_i, t] \geq E[U_i(\omega, \bar{y}(f_i, f_{-i}^*), \bar{t}(f_i, f_{-i}^*)) | X_i, t]$ with the convention $X_0 = \emptyset$.

The inertia assumption implies that if the senders disagree at $t = 0$, they cannot suddenly agree for all $t > 0$, therefore the receiver cannot rationally believe that she will learn the state *immediately* following an initial disagreement. This makes our construction of a no-delay, fully revealing equilibrium under discrete time infeasible under continuous time. The issue is resolved by letting the receiver believe that she will learn the true state at a particular point in time after the disagreement, until which she is willing to wait given her beliefs about the state. The senders know that the receiver only pays attention to them at that particular point in time, and only believes them if they agree. Since they expect each other to report truthfully at that time (and babble otherwise), the best they can do once they are off the equilibrium path is to report truthfully when the receiver expects them to do so.

The following result is the counterpart of Proposition 2 in continuous time.

²⁰At time 0, the play is uniquely determined by inertia for some period $[0, \delta)$. If t^* is the supremum of t such that play is uniquely determined on $[0, t)$, then by inertia there is a unique continuation until $t^* + \delta' > t^*$, which implies $t^* = \infty$. For more details see Bergin and Macleod (1993) and Perry and Reny (1993).

Proposition 3 *In continuous time, if the senders commonly know the state, then there exists an immediate, fully-revealing equilibrium.*

Proof. We construct the equilibrium as follows. On the equilibrium path, sender $i \in \{1, 2\}$ reports $m_i^t = \omega$ at $t = 0$, and repeats it for all $t > 0$ provided that both senders have done so at $t = 0$. If the senders report the same $m^* \in S$ at $t = 0$, then the receiver carries out $y = m^*$ (ending the game) at all $t \geq 0$ such that the game has not ended before t .

Suppose that the senders report $m_1^0 = \omega' \neq \omega'' = m_2^0$, $\omega', \omega'' \in S$, at time 0. (A report $m_i^0 \notin S$ is understood by all players as a report of a state according to a fixed, arbitrary mapping.) Denote $\delta(\omega', \omega'')$ the receiver's willingness to wait to find out the state given that she believes it is either ω' or ω'' :

$$E[U_0(\omega, \omega, \delta(\omega', \omega'')) \mid \omega \in \{\omega', \omega''\}] = \max_{y \in Y} E[U_0(\omega, y, 0) \mid \omega \in \{\omega', \omega''\}].$$

Define $T \equiv \{\tau \mid \tau = \lambda\delta(\omega', \omega'')/2, \lambda \in \mathbb{Z}^+\}$. If the senders report $\omega' \neq \omega''$ at 0, then for any history $h \in H$ and $t \in T$ they both report ω at t , while for all $h \in H$ and $t \notin T$ sender i reports m_i^t such that τ is the maximal element of $T \cap [0, t)$. At $t \in T$ and history such that the senders reported $\omega' \neq \omega''$ at 0, if the senders send the same $m^* \in S$ at t then the receiver believes that $\omega = m^*$, otherwise she believes $\omega \in \{\omega', \omega''\}$. If the senders disagree at time 0, then the receiver chooses inaction at all $t \geq 0$ unless $t \in T$ and both senders report the same $m^* \in S$ at t , in which case she picks $y = m^*$.

If all play the proposed strategies, then the game ends at $t = 0$ with payoffs $U_i(\omega, \omega, 0)$, $i = 0, 1, 2$. If the senders report $\omega' \neq \omega''$ at $t = 0$, then the receiver believes $\omega \in \{\omega', \omega''\}$, and that she will learn the truth at exactly time $\delta(\omega', \omega'')/2 < \delta(\omega', \omega'')$, hence it is optimal for her to choose inaction for all $t' \in [0, \delta(\omega', \omega'')/2)$. If the senders agree on some $m^* \in S$ at $t = \delta(\omega', \omega'')/2$ or any other $t \in T$, then, according to the receiver's beliefs, it is optimal for her to carry out $y = m^*$. If the senders disagree at $t = \delta(\omega', \omega'')/2$ or any other $t \in T$, then the receiver continues to believe $\omega \in \{\omega', \omega''\}$ and that she will learn the truth $\delta(\omega', \omega'')/2$ time later. Hence it is optimal for her to wait for the next time when she expects the senders to agree on the truth.

If the senders disagree at $t = 0$, then in the continuation they expect each

other to report ω at all $t \in T$ and the receiver to ignore all messages sent at all $t \notin T$. Therefore neither sender can induce any outcome other than $y = \omega$ at any $t \in T$, and their best response is to report ω at all $t \in T$. ■

An attractive feature of the equilibrium construction in continuous time is that the receiver's beliefs regarding the true state of nature are *naturally continuous* in the senders' reports at time 0. This is so because she believes that if the senders disagree at time 0, then exactly one of them is truthful, and that they will both agree on the truth in a period of time that is proportional to her willingness to wait to resolve the remaining uncertainty about the state.

3.4 On the existence of other equilibria

Propositions 2 and 3 establish the existence of an immediate, fully-revealing equilibrium in discrete and continuous time, provided the senders commonly observe the state. However, there are also other equilibria of these dynamic games. For example, any equilibrium of the corresponding static game can be sustained in the dynamic setup. This raises the questions, (i) whether all equilibria of the dynamic game can be characterized, and (ii) in what sense the immediate, fully-revealing equilibrium is focal.

The characterization of all equilibria remains an open question. The first and greatest impediment is that even in the *static* model, the set of equilibria is not characterized in the literature. The only class of cheap talk games where the set of equilibrium outcomes is known is the single-sender, static cheap talk model with a one-dimensional state space (Crawford and Sobel (1982)). The rest of the static literature focuses on the existence of particular equilibria, e.g., fully-revealing equilibria with multiple senders (from Krishna and Morgan (2001a) to Ambrus and Takahashi (2007)), or comparative equilibria with a single sender but a multi-dimensional state space (Chakraborty and Harbaugh (2007)). It is clear that more research needs to be done on the characterization of equilibria in all cheap talk games, static or dynamic.

It is easy to see that in our dynamic setup, with sufficiently frequent interaction (with δ close to zero or in continuous time), there exist immediate, *imperfectly-revealing* equilibria. In these equilibria, the senders simultaneously

report the element of a partition of S that the true state belongs to, and the receiver carries out her ideal policy given the information revealed by the senders. Such an equilibrium exists as long as the receiver is willing to wait one period to learn the state subject to the partitioning of S ; that is, if the partition is “not too coarse”. However, if the senders’ have strictly concave utilities in y (given ω), then *ex ante* they prefer the equilibrium with the finest partition, because the receiver’s response is unbiased in any equilibrium (on average, her policy matches the state), and the most-informative communication equilibrium minimizes the variability of the policy choice.

It is also clear that if the state space is one-dimensional and the senders’ biases have opposite signs (as in Krishna and Morgan (2001a)), any policy that one sender prefers over the receiver’s ideal point (the state) is disliked by the other. Therefore, with a one-dimensional state space and opposite-sign biases, the immediate, fully-revealing equilibrium is preferred by all players.

We summarize our observations in the following Proposition.

Proposition 4 (a) *Let $\mathcal{P} = \{P_k\}_{k \in K}$ be a partition of S . In discrete time with $\delta > 0$ assume $E[\max_{y \in Y} E[U_0(\omega, y, \delta) | \{\omega \in P_k\}]] \geq \max_{y \in Y} E[U_0(\omega, y, 0)]$; in continuous time \mathcal{P} can be arbitrary. Then, there exists a perfect Bayesian equilibrium where $k \in K$ such that $\omega \in P_k$ is immediately disclosed. *Ex ante* (before ω is realized), all parties prefer the outcome of the immediate, fully-revealing equilibrium to the outcome of any such equilibrium.*

(b) *If $S \subset R$ and $\text{sgn}(y_1(\omega) - \omega) = \text{sgn}(\omega - y_2(\omega))$ for all ω , then there is no perfect Bayesian equilibrium whose outcome both senders prefer conditional on the state, in any state, to that of the immediate, fully-revealing equilibrium.*

The latter result suggests that if the receiver is concerned about the possibility of sender collusion, or the focality of the fully-revealing equilibrium, then it may serve her interest to present policy decisions as simple, uni-dimensional problems, and seek the advice of (well-informed) experts that are known to have opposite biases. Propositions 2-3 imply that there exists an immediate, fully-revealing equilibrium even if the biases are large, while part (b) of Proposition 4 establishes that no other equilibrium dominates it for both senders.

4 Delay and disclosure with imperfect signals

In this section, we drop the assumption that the payoff-relevant state of nature is commonly known by the senders, and investigate the robustness of the immediate, fully-revealing equilibrium outcome to *noise* in the senders' signals. Equilibrium constructions that rely on "cross-checking" and the receiver knowing that the senders must be able to agree break down.

Nevertheless, in certain environments with noisy signals, we are able to construct equilibria where the senders reveal their signals with a positive expected delay along the equilibrium path. We show that as the senders' signals become arbitrarily precise, the outcome of this equilibrium *converges to full revelation with no delay*—the result found with perfect observation. First, we exhibit an instructive example to show how the construction works. Then, we present a general framework with a finite state space and imperfect observation, and prove that under certain conditions, there exists an equilibrium whose outcome converges to immediate, full revelation as the noise in the senders' signals vanishes. An extension to a continuous state space and normal signals is presented in the Appendix.

The main difficulty of constructing informative equilibria in a model with noisy signals is that the senders' reports can differ even if they are truthful. The receiver could threaten with delay as a function of the extent of the senders' disagreement, which may provide the right incentives for them to report honestly. However, once the threat of delay induces the senders to be truthful, the receiver has no incentive to carry out any promised delay regardless of the reports, and so this type of construction seems unworkable.

The solution to this problem involves two ideas. First, in the equilibria that we construct, the senders do not disclose all their information at once. The promise of learning more in the future makes the receiver willing to wait, if needed, along the equilibrium path. Second, the senders slow down information disclosure in case they notice a substantial disagreement, which provides incentives for them to report truthfully, in order to avoid delay. In equilibrium, neither sender can unilaterally accelerate disclosure because the receiver expects them to proceed at the same pace.

4.1 An instructive example with noisy signals

In this subsection only, let the state and the senders' signals be two-digit, positive integers: $\omega, X_1, X_2 \in S = \{10, \dots, 99\}$. Noise in the senders' signals is represented by $\varepsilon > 0$ such that $\Pr(X_1 = X_2 = k | \omega = k) = 1 - \varepsilon$. Assume that $E[\omega | X_1 = k, X_2 = \ell] = (k + \ell)/2$, and that (X_1, X_2) has full support on $\{(x_1, x_2) \in S^2 : |x_1 - x_2| < 5\}$. The former assumption simplifies the receiver's best response to truthful reports. The latter one rules out certain joint signal realizations; however, for any realization of (X_1, X_2) , the smallest set that is commonly known by the senders to contain ω (or the signal values) is S .

Denote the receiver's minimal willingness to wait to learn $E[\omega | X_1, X_2]$ when she knows only *the second digits* of X_1 and X_2 by D_0 . Recall that sender i 's maximal willingness to wait to get policy $y_i(\omega)$ instead of $y = \omega$ immediately is denoted by Δ_i for $i = 1, 2$. Assume that at least two rounds of communication are feasible at time 0.

We claim that if D_0 is greater than Δ_1 and Δ_2 , $\varepsilon > 0$ is small, and the time period lengths are sufficiently short, then there exists an equilibrium where the receiver learns both senders' signals and carries out her ideal policy with a positive expected delay. However, as the noise in the signals disappears, the expected delay converges to zero. The equilibrium is constructed as follows. At time $t = 0$, the senders first publicly announce the second digits of their respective signals. If these reports match, then they immediately announce $E[\omega | X_1, X_2]$. If the senders' reports about their signals' second digits differ at $t = 0$, then they do not change their reports until D_0 , and then announce $E[\omega | X_1, X_2]$. In either case, the receiver waits for agreement on $E[\omega | X_1, X_2]$ and carries out the corresponding policy.

In order to see that these strategies form an equilibrium under the given conditions, the key observation is that the senders' reports about the second digits of their respective signals establish common knowledge between them about the (claimed) realizations of X_1 and X_2 , without providing the same information to the receiver.²¹ However, the receiver still does not know the first digits of X_1, X_2 and $E[\omega | X_1, X_2]$, hence she is willing to wait until $t = D_0$ for a full disclosure.

²¹For example, if the reports are (4, 5) and $X_1 = 34$, then sender 1 knows that sender 2 claims $X_2 = 35$ (because $|X_1 - X_2| < 5$), and this is common knowledge between them.

At $t = 0$ or $t = D_0$ depending on whether or not the second digits of their signals match, the senders immediately and fully reveal $E[\omega|X_1, X_2]$, which is commonly known to them. They do so the same way that they report the commonly-known state in Propositions 2-3. The argument is completed by noting that the senders prefer to report the second digits of their signals truthfully at $t = 0$ because otherwise they risk a mismatch and a costly delay. The expected delay in equilibrium converges to zero as $\varepsilon \rightarrow 0$ because disagreements occur on the equilibrium path with vanishing probability.

This example demonstrates the key ingredients of our construction in the more general case. The senders release information at a deliberate pace in order to make it worth for the receiver to wait (expecting more information) in case they have a disagreement. They implement the delay by slowing down information disclosure in case of an initial mismatch. However, such delays rarely occur on the equilibrium path when the signals are sufficiently precise.

4.2 A general, finite state space model with noisy signals

Let $\Delta_i(\omega)$ denote the (state-contingent) willingness to wait of sender i in order to obtain his ideal point instead of the receiver's ideal point without delay:

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

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 proportional to b_i^2 . Also, define $\Delta_i = \max_{\omega \in S} \Delta_i(\omega)$ for $i = 1, 2$.

Let the state space be a finite subset of the unit interval: For a fixed $N \in \mathbb{N}$,

$$S = \{\omega_k = (2k - 1)/2^{N+1} \mid k = 1, \dots, 2^N\}.$$

The distribution of ω can be arbitrary over this set. The finiteness of the support of ω is not a restrictive assumption because we can approximate *any distribution* on $[0, 1]$ arbitrarily closely provided N is sufficiently large. Our model and arguments also immediately generalize to finite state spaces in \mathbb{R}^n .

We now define a family of information structures where the senders observe

correlated private signals, but never have common knowledge about the state of nature. The signal structures in this family are naturally (partially) ordered according to the precision of the senders' signals.

The senders' signals, (X_1, X_2) , are random variables with *full support* on S^2 . The full-support assumption guarantees that no realization of X_i rules out any realization of X_j , which makes the construction of an informative equilibrium more difficult. (The condition fails in the environment of Section 4.1.) Let ε be the lowest positive number such that for all k , the probability that the state and the other sender's signal both equal ω_k , given that one sender's signal realization is ω_k , is at least $1 - \varepsilon$:

$$\varepsilon = \inf \{ \tilde{\varepsilon} : \Pr(\omega = X_j = \omega_k | X_i = \omega_k) \geq 1 - \tilde{\varepsilon}, \forall k, \forall i \neq j \}. \quad (3)$$

For the sake of brevity, we omit the natural range of indices such as $i, j \in \{1, 2\}$, $k \in \{1, \dots, 2^N\}$. Note that (3) is not an assumption, just a *definition* of ε . All joint distributions of (X_1, X_2) conditional on ω are partially ordered by the index ε . The smaller the index ε , the more precise the senders' signals are.

We make two simplifying assumptions on the signal structure. We do not specify the joint distribution of (ω, X_1, X_2) in more detail because it is not necessary for obtaining our results.

Assumption 1 *Nothing can be inferred about the state of nature from knowing that the senders' signal realizations are different: For all $k = 1, \dots, 2^N$,*

$$\Pr(\omega = \omega_k \mid \{X_1 \neq X_2\}) = \Pr(\omega = \omega_k). \quad (4)$$

Assumption 2 *If the senders' signal realizations differ, then it is most likely that they fall near each other and the true state. For $i = 1, 2$,*

$$\Pr(|X_i - \omega| \leq 2^{-N} \mid \{X_1 \neq X_2\}) > 1 - \varepsilon. \quad (5)$$

These assumptions imply that for $\varepsilon > 0$ small and N large, the receiver's willingness to wait to find out the realizations of the senders signals, given that she knows they differ, is approximately Δ_0 . To see this, note that by (4), if the receiver learns that the senders' signals are different, then her beliefs about

the state of nature remain unchanged. Furthermore, condition (5) implies that the receiver knows, when the senders' signal realizations are different, they are likely to be near each other and the true state ω . Therefore, if the state space is sufficiently dense (N is large), then learning where the non-equal signal realizations fall is almost equivalent to learning ω . But Δ_0 is exactly the delay that the receiver is willing to incur to find out ω , given her prior beliefs about the state.

It is easy to construct examples where Assumptions 1-2 are satisfied. Fix the state distribution and $\varepsilon > 0$, and suppose that with 50-50% chance one of the senders is chosen to observe the realization of ω while the other sender observes ω with probability $1 - \varepsilon$ and an adjacent realization on the grid with probability ε . The senders do not know whether or not they observed ω . By construction, given the realization of X_i , the probability that it matches both ω and X_j is $1 - \varepsilon$, satisfying (3). Since the probability of $X_i \neq X_j$ is $1 - \varepsilon$ for all realizations of ω , nothing can be inferred about ω in case the signals differ, and so Assumption 1 holds. Finally, by construction, if X_i and X_j differ, then they are at or next to ω on the grid, hence Assumption 2 is satisfied.

Finally, we make an assumption on the availability of certain jointly observed random variables. This randomization device could be a dimension of the state about which the senders do have common knowledge, even though they only observe other aspects of the state with noise. It can also be interpreted as a pure "sunspot".

Assumption 3 *The senders (but not the receiver) jointly observe the realization of $2N$ i.i.d. uniform random binary variables.*

The role of these random variables is similar to that of the second digits of the senders' signals in Section 4.1: They allow the senders to exchange information without releasing it to the receiver using only public reports. The same could be achieved by a single, one-way private message from one sender to the other. In the analysis that follows, Assumption 3 can be dispensed with if private communication is allowed.²²

²²Private communication is usually not allowed in multi-sender cheap talk games. A notable exception is Wolinsky (2002) who argues in a very different setup that private communication between the senders can help the decision maker to elicit information.

4.3 Results in the model with noisy signals

The main result of this section is that in our environment, if the signal structure is sufficiently precise and the receiver's willingness to wait to learn ω exceeds the senders' willingness to wait to induce their ideal points, then there exists an equilibrium where the senders report their signals with a positive expected delay. However, as the noise in the signal structure disappears, the expected equilibrium delay converges to zero.

The first step towards proving this result is to show that using encoded public messages, the senders can establish common knowledge between themselves about their signal values without disclosing any information to the receiver.

Lemma 1 *Using a simultaneous, public message, the senders can establish common knowledge about their signal values between themselves while keeping the receiver uninformed.*

Proof. Denote the first N i.i.d. uniform binary random variables referred to in Assumption 3 by $\theta_1^1, \dots, \theta_N^1$ while the last N random bits by $\theta_1^2, \dots, \theta_N^2$.

Suppose that the realization of sender 1's signal is ω_k while sender 2's is ω_ℓ . Sender 1 represents $k \in \{1, \dots, 2^N\}$ as a base-2 integer with N bits ξ_1^1, \dots, ξ_N^1 ; sender 2 represents ℓ in binary form with the bits ξ_1^2, \dots, ξ_N^2 . Then, each sender i simultaneously and publicly reports the integer represented by the N bits r_1^i, \dots, r_N^i , where $r_n^i = \xi_n^i \oplus \theta_n^i$; here \oplus denotes the exclusive disjunction (xor) operation (the binary addition of two, single-bit numbers without "carrying").

Since sender $j \neq i$ knows each θ_n^i , he can infer ξ_n^i from r_n^i for $n = 1, \dots, N$. However, the receiver learns nothing about ξ_n^i by observing r_n^i but without knowing the realization of θ_n^i . This is so because $r_n^i = 1$ if and only if either $\xi_n^i = 1$ and $\theta_n^i = 0$ or $\xi_n^i = 0$ and $\theta_n^i = 1$, $\Pr(\theta_n^i = 0) = \Pr(\theta_n^i = 1) = 1/2$, and by Bayes' rule,

$$\Pr(\xi_n^i = 1 | r_n^i = 1) = \frac{(1/2) \Pr(\xi_n^i = 1)}{(1/2) \Pr(\xi_n^i = 1) + (1/2) \Pr(\xi_n^i = 0)} = \Pr(\xi_n^i = 1).$$

Similarly, $\Pr(\xi_n^i = 1 | r_n^i = 0) = \Pr(\xi_n^i = 1)$. ■

In the "communication protocol" used in the proof of this lemma the senders encode their signal values with the help of the commonly observed random binary

variables. When they publicly report the encoded signals, they can decode each other's messages while the receiver (without observing the random key) is not able to do so. Provided that they report truthfully according to the protocol, they end up with common knowledge about their signals. This protocol is used in the equilibrium construction of Proposition 5, and it will be shown that in equilibrium, the senders indeed report truthfully.

Recall that Δ_0 denotes the amount of delay that the receiver is willing to incur to find out the state of nature, while Δ_i is the maximum delay sender i is willing to endure in order to induce his ideal point instead of an immediate action matching the state. The main result of the section is the following.

Proposition 5 *Assume that time is either continuous or discrete with arbitrarily short periods, and there are at least two rounds of communication at time $t = 0$. In our environment with noisy observation, if N is sufficiently large and $\Delta_0 > \max\{\Delta_1, \Delta_2\}$, then for $\varepsilon > 0$ sufficiently small, there exists an equilibrium where the receiver learns the state of nature with an error of $\pm 1/2^N$ with a probability greater than $1 - \varepsilon$, and chooses her ideal policy accordingly. The equilibrium outcome converges to immediate, full revelation as $\varepsilon \rightarrow 0$.*

Proof. The equilibrium play consists of three stages. First, at time 0, the senders use the communication protocol of Lemma 1 to establish common knowledge about their signal realizations between themselves while keeping the receiver uninformed. Second, but still at time 0, the senders implement a delay of length d , such that $\max\{\Delta_1, \Delta_2\} < d < \Delta_0$, in case their reported signal values are different.²³ They do so by keeping their reports the same for all $t \in [0, d)$; meanwhile the receiver chooses inaction. Third, if and when the delay is over, the senders simultaneously report the receiver's ideal point conditional on the information they shared between each other at time 0, and the receiver carries out their recommendation.

The reason why this is indeed an equilibrium play is the following.

In the third stage of the game (when delay d , if incurred, is over), the senders have common knowledge about the receiver's ideal point conditional on the signal

²³For simplicity, in the equilibrium that we construct, any report by sender i at time 0 (even silence) is known to be interpreted by all players as an announcement of an integer between 1 and 2^N , which is then "deciphered" by the other sender to obtain sender i 's reported signal.

values they exchanged at time 0. Therefore, immediate full revelation of this information is an equilibrium outcome in the continuation, just like in the perfect observation case studied in Propositions 2-3. If the senders disagree in the third stage, then the receiver simply waits until they agree, believing that they will do so right away.

In the second stage of the game, if the senders' signals (as reported to each other via the coded messages in stage 1) differ, then neither sender can do better than report the same for all $t \in [0, d)$. This is so because all players believe that neither sender reports anything informative for all $t \in [0, d)$. During this period of delay, the receiver's beliefs about the state are unchanged (coincide with her prior beliefs) because, by Assumption 1, $X_i \neq X_j$ reveals no information about the state. In contrast, by Assumption 2, the receiver expects to learn ω at $t = d$ with an error of at most $\pm 1/2^N$ with probability greater than $1 - \varepsilon$. If N is large and $\varepsilon > 0$ sufficiently small, then the receiver's willingness to wait in this stage is approximately $\Delta_0 > d$. Therefore it is a best response for the receiver to remain inactive during the second stage of the proposed equilibrium play.

In the first stage, neither sender has an incentive to misrepresent his signal to the other sender. This is so because by sending a message other than the realization of X_i , sender i induces a delay d with a probability greater than $1 - \varepsilon$. The best outcome that this sender can hope for after delay d is his ideal point, $y_i(\omega)$. However, $E[U_i(\omega, y_i(\omega), \Delta_0)] < E[U_i(\omega, \omega, 0)]$ by $\Delta_0 > \Delta_i$, hence a delay d close to Δ_0 makes such a deviation unprofitable for sender i .

Finally, the expected equilibrium delay tends to zero as the noise in the senders' signals vanishes, that is, $\varepsilon \rightarrow 0$. This is so because the probability that the senders' signals differ approaches zero as ε goes to zero. ■

The significance of Proposition 5 is that it demonstrates the robustness of the immediate, fully-revealing equilibrium outcome to a small noise in the senders' signals. Notice that there is *no common knowledge* about the payoff-relevant state of nature between the two senders, no matter how small the noise in their signals is. As a result, the receiver cannot be sure to detect lying, or any deviation from the equilibrium path. Nevertheless, the equilibrium outcome with noisy signals converges to the perfect-observation limit as the noise vanishes.

The construction works as long as 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. If learning the senders' signals does not improve by much the receiver's decision,²⁴ while at least one of the senders is patient and strongly dislikes the receiver's ideal point, then the equilibrium with almost-immediate full disclosure is not "credible" in the sense that the patient sender could simply "outwait" the receiver. This argument does not apply in the perfect-observation model; immediate truthful reporting *is* and remains a perfect Bayesian equilibrium in that setting. However, in a reasonable perturbation of the model like the one of this section—with imperfect signals and delay on the equilibrium path—this type of consideration becomes relevant.

This 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. While we did assume that the senders have exclusive common knowledge of certain random variables other than the state (Assumption 3), this assumption can be eliminated by either assuming that the signals do not have full support over the state space (see the example of Section 4.1), or by allowing private communication between the senders.

Our model with noisy signals features a state distribution with finite support. This is a useful technical assumption for two reasons. First, if the signals are sufficiently precise, then a sender's signal "swamps" the prior in a finite-states model.²⁵ Second, there exists a simple punishment function (a fixed delay following any disagreement regarding which realization of ω is the most likely) that elicits truthful signal announcements. The results and the proof of Proposition 5 can be extended to a model with a continuum of states as long as these two properties generalize. In the Appendix we provide an example along these lines.

²⁴For 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.

²⁵If the state space is finite and the senders' signals are unbiased and sufficiently precise, then each sender's "best guess" for the realization of the state is the one nearest to his observed signal value. With a continuum of states and a proper prior, the sender would always "shade" his honest best guess towards the prior mean.

5 Conclusion

We have shown that any information that is commonly known to the senders can be immediately and fully revealed in a perfect Bayesian equilibrium of a dynamic, multi-sender cheap talk game. In the static version of the model such equilibria only exist under certain conditions on the state space and/or the senders' biases. The “perfectness” of the equilibrium in our dynamic setup implies that the construction does not rely on exogenous commitment to delay the receiver's decision. Instead, the equilibrium is sustained by the receiver's beliefs that even if the senders have disagreed in the past, they will agree soon enough in the future—which they must be able to do as they have common knowledge about the state.

We studied the robustness of this result to 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 immediate, full revelation as the noise in the senders' signals vanishes. The key idea there is that by “pacing” information disclosure, the senders can create incentives for themselves to make truthful (yet partial) reports, and, at the same time provide an incentive for the receiver to wait to see more information disclosed along the equilibrium path.

The main “practical” conclusion of our inquiry is that as long as the receiver's action space includes temporary, suboptimal policies (like “inaction” in case the players have time preference), she can induce information transmission from well-informed experts *over time*, even if their preferences are drastically different from her own. Frequent communication and self-fulfilling beliefs in the possibility of a building a “consensus” on the receiver's part are the key elements of the construction of dynamic, informative equilibria. We believe that these insights arising from our formal model are applicable in political decision-making processes and in other applications.

The result that immediate, fully-revealing equilibria exist in a dynamic model of communication when they do not exist in comparable 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*.²⁶ Even if such signals are available, a fully-revealing equilibrium does not exist in the static game if there is no “worst state” for the sender.²⁷ This is the case, for instance, if the state of nature is distributed on the circumference of a circle, and the receiver wants to match her action to the state, which is diametrically opposite to the sender’s ideal point. In the standard, static version of this game, there does not exist a fully-revealing equilibrium because by truthfully and verifiably revealing the state the sender causes the receiver to implement the worst possible outcome for him. However, there exists a fully-revealing perfect Bayesian equilibrium in the dynamic game where the receiver believes that the sender will send a verifiable signal about the state right away even if 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. The model could be extended to more than two senders and, among other things, the senders’ incentives to acquire information could be investigated. In realistic situations the receiver could also learn (privately or publicly) about the state, which would provide an additional incentive for her to wait along the equilibrium path. However, in all these variants of the model, our environment with either perfect observation or a small noise in the senders’ signals should really be thought of as a benchmark case. It would be interesting to study models where the state of nature evolves over time, and so do the senders’ imperfect signals about the state. In such models, unlike in ours, we would expect that optimal communication equilibria would involve non-negligible, positive delay.

²⁶A hard (or verifiable) signal, as opposed to a soft (unverifiable) one, can only be sent by the sender in a particular state of nature. See Milgrom (1981).

²⁷See Milgrom (1981), Seidmann and Winter (1997).

6 Appendix: Dynamic, Multi-Sender Cheap Talk with Normal Inference

In this appendix we discuss an application with a *continuum of states* and *noisy observation*. We show that as the noise in the senders' signals vanishes, equilibrium play converges to immediate, full revelation.

The information structure parallels that of Battaglini (2004). The state of nature is drawn from a uniform distribution on \mathbb{R} , that is, the common prior is *diffuse*.²⁸ 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:

$$U_i(\omega, y, \delta) = -(y - y_i(\omega))^2 - \delta \text{ for } i = 0, 1, 2. \quad (6)$$

By the rules of normal inference, 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 2 *Suppose that 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 $i = 1, 2$.*

²⁸This is a mathematically imprecise assumption, reflecting the limiting case as the variance of the receiver's prior grows without bounds (the prior is completely uninformative).

Proof. 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 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. The first-order condition of his problem becomes

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

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. ■

The outcome of this mechanism can be replicated in an *equilibrium* with no commitment to delay on the receiver's part. Formally:

Proposition 6 *Consider our model with diffuse ω , normal signals, and equal, opposite-sign biases. Assume that the senders can exchange a private message at $t = 0$. 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 precision of the senders' signals increases the expected delay tends to zero, and the outcome converges to that of an immediate, fully-revealing equilibrium.*

The equilibrium is constructed as follows. On the equilibrium path, at time $t = 0$, the senders 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 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 make truthful private reports to each other at time zero and follow their equilibrium strategy thereon, then the receiver learns and implements $E[\omega | X_1 = x_1, X_2 = x_2]$ with delay $d(x_1, x_2)$ as claimed. By Lemma 2, neither sender has an incentive to misrepresent his signal in the private reporting stage. The receiver has no profitable deviation either, because (due to the diffuse prior assumption) her expected payoff from taking any action without learning something about the state is $-\infty$. She needs to wait for the senders' agreement in order to take any action. Therefore the proposed strategies indeed form an equilibrium.

It is also clear that the expected value of $d(x_1 - x_2)$ goes to zero as the senders' signals become arbitrarily precise. That is, the equilibrium outcome converges to *full revelation* and *no delay*.

The main reason why we find our model in Section 4 more attractive than the setup discussed in this Appendix is that here, the state space is unbounded, the prior is diffuse, and consequently the receiver's willingness to wait to glean any information from the senders is also unbounded. This is a strong assumption both from a conceptual and a practical viewpoint.

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] Aumann, Robert, and Sergiu Hart (2003), “Long Cheap Talk,” *Econometrica*, 71, 1619-1660.
- [4] Austen-Smith, David (1990): “Information Transmission in Debate”, *American Journal of Political Science*, 34, 124-152.
- [5] Austen-Smith, David (1993): “Interested Experts and Policy Advice: Multiple Referrals under Open Rule,” *Games and Economic Behavior*, 5, 3-44.
- [6] Austen-Smith, David and Jeffrey Banks (2000): “Cheap Talk and Burned Money,” *Journal of Economic Theory*, 91, 1-16.
- [7] Ausubel, Larry and Raymond Deneckere (1987): “One is Almost Enough for Monopoly,” *Rand Journal of Economics*, 18, 255-274.
- [8] Battaglini, Marco (2002): “Multiple Referrals and Multidimensional Cheap Talk,” *Econometrica*, 70(4), 1379-1401.
- [9] Battaglini, Marco (2004): “Policy Advice with Imperfectly Informed Experts,” *Advances in Theoretical Economics* 4(1).
- [10] Bergin, James, and W. Bentley Macleod (1993): “Continuous Time Repeated Games,” *International Economic Review*, 34 (1), 21-37.
- [11] Chakraborty, Archisman and Rick Harbaugh (2007), “Comparative Cheap Talk,” *Journal of Economic Theory*, 132, 70-94.
- [12] Che, Yeon-Koo, and József Sákovics (2004): “A Dynamic Theory of Holdup,” *Econometrica*, 72(4), 1063-1103.
- [13] Crawford, Vincent and Joel Sobel (1982): “Strategic Information Transmission,” *Econometrica*, 50(6), 1431–1451.

- [14] Damiano, Ettore, Li Hao, and Wing Suen (2007): "Delay in Strategic Information Transmission", *mimeo*, 2007.
- [15] Fudenberg, Drew and Jean Tirole (1991), *Game Theory*, MIT Press.
- [16] Gilligan, Thomas W. and Keith Krehbiel (1989): "Asymmetric Information and Legislative Rules with a Heterogenous Committee," *American Journal of Political Science*, 33, 450-490.
- [17] Gul, Faruk (1987): "Noncooperative Collusion in Durable Goods Oligopoly," *Rand Journal of Economics*, 18, 249-254.
- [18] Gul, Faruk (2001): "Unobservable Investment and the Holdup Problem," *Econometrica*, 69, 343-376.
- [19] Kartik, Navin (2007): "A Note on Cheap Talk and Burned Money," *Journal of Economic Theory*, 136(1), 749-751.
- [20] Krishna, Vijay and John Morgan (2001a), "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 95(2), 435-452.
- [21] Krishna, Vijay and John Morgan (2001b), "A Model of Expertise," *Quarterly Journal of Economics*, 116(2), 747-775.
- [22] 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.
- [23] Lockwood, B. and J. P. Thomas (2002): "Gradualism and Irreversibility", *Review of Economic Studies*, 69, 339-356.
- [24] Marx, Leslie and Steven Matthews (2000): "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies*, 67, 327-358.
- [25] Milgrom, Paul (1981): "Good News and Bad News: Representation Theorems and Applications", *Bell Journal of Economics*, 21, 380-391.

- [26] Milgrom, Paul and John Roberts (1986), “Relying on the information of interested parties,” *Rand Journal of Economics*, 17, 18-32.
- [27] Perry, Motty and Philip Reny (1993), “A Non-cooperative Bargaining Model with Strategically Timed Offers,” *Journal of Economic Theory*, 59, 50-77.
- [28] Perry, Motty and Philip Reny (1994), “A Noncooperative View of Coalition Formation and the Core,” *Econometrica*, 62, 795-817.
- [29] Rubinstein, Ariel (1982): “Perfect Equilibrium in a Bargaining Model”, *Econometrica* 50 (1982), 97-110.
- [30] Rubinstein, Ariel and Asher Wolinsky (1992): “Renegotiation-Proof Implementation and Time Preferences,” *American Economic Review*, 82, 600–614.
- [31] Seidmann, Daniel and Eyal Winter (1997): “Strategic Information Transmission with Verifiable Messages”, *Econometrica*, 65(1), 163-169.
- [32] Serrano, Roberto (2007): “Bargaining”, *The New Palgrave Dictionary of Economics*, forthcoming
- [33] Wolinsky, Asher (2002): “Eliciting Information from Multiple Experts”, *Games and Economic Behavior* 41, 141-160.