Disclosure of Endogenous Information

Matthew Gentzkow and Emir Kamenica*
Stanford University and University of Chicago

March 2016

Abstract

We study the effect of disclosure requirements in environments where experts publicly acquire private information before engaging in a persuasion game with a decision maker. In contrast to settings where private information is exogenous, we show that disclosure requirements never change the set of equilibrium outcomes regardless of the players' preferences.

Keywords: persuasion; disclosure regulation; verifiable types
JEL: D82; D83; C72

*We thank Eduardo Perez-Richet for helpful comments. We thank the Neubauer Family Foundation, the William Ladany Faculty Research Fund, and the Initiative on Global Markets at the University of Chicago Booth School of Business for financial support.

1 Introduction

In many settings of economic interest, informed experts choose how much of their private information to disclose to a decision-maker (DM) who will take an action that affects the payoffs of all the players. Often, the disclosed information is verifiable, meaning that the claims made by the experts might be more or less informative but cannot be false.[1] A large literature establishes conditions under which all private information will be disclosed in equilibrium.[2] A key insight from this literature is that when experts' preferences are suitably monotonic in DM's belief or sufficiently opposed, full disclosure is an equilibrium. With a single expert who has monotonic preferences, full disclosure is the unique equilibrium outcome.[3]

Models in this literature typically take the experts' initial private information as exogenous. This is a suitable assumption in many settings. There are other settings, however, where information is symmetric at the outset and the experts choose how much private information to gather. For example, a pharmaceutical company may or may not conduct clinical trials that assess whether a drug has differential efficacy for a particular demographic group. When such information is costless and can be covertly gathered, it is a dominant strategy to become as informed as possible. When the acquisition of private information is public, however, becoming more informed may be harmful in equilibrium. If the FDA knows that a pharmaceutical company conducted a clinical trial specifically to assess a drug's side effects in children, the failure to disclose the results of this trial is likely to generate skepticism.

In this paper, we study disclosure when private information is endogenous. We consider *ex ante* symmetric information games where $n \geq 1$ experts simultaneously conduct experiments about the state of the world. More informative experiments are (weakly) more costly to the expert. DM observes which experiments are conducted, and each expert privately observes the outcome of his own experiment. The experts convey verifiable messages to DM about the outcomes.[4] DM then takes an action. We focus on pure-strategy, perfect Bayesian equilibria.[5] The outcome of the game is the joint distribution of the state of the world, DM's belief, DM's action, and all the players'

[1] In the persuasion games literature, terms *certifiable*, *verifiable*, and *provable* are typically used interchangeably. A separate literature initiated by Crawford and Sobel (1982) examines cheap talk communication.

[2] For seminal contributions, see Grossman (1981), Milgrom (1981), and Milgrom and Roberts (1986). For a recent survey, see Milgrom (2008).

[3] This result holds when it is commonly known that the expert can costlessly disclose all information. If there is uncertainty about what information the expert could have reported, there might be only partial disclosure in equilibrium (e.g., Shin 1994).

[4] In particular, we assume that each expert can send a message that proves what he observed.

[5] Related arguments can be used to analyze mixed-strategy equilibria, but focusing on pure strategies substantially simplifies the exposition.

payoffs.

We study the effect of requiring experts to fully disclose the results of their experiments. We might expect such a requirement to change the equilibrium outcomes when the usual monotonicity or opposed preferences conditions for full disclosure are not satisfied. This requirement might benefit DM if it causes more information to be revealed. It might also benefit the experts if their inability to commit to truthful disclosure reduces their equilibrium payoffs, as can happen in some cheap talk settings.

Our main result is that disclosure requirements can have no effect on the set of equilibrium outcomes and thus no effect on either DM's or the experts' payoffs. Essentially, we show that endogenous information will always be disclosed in equilibrium; if there is an equilibrium in which information is withheld, the outcome must be the same as in another equilibrium with truthful disclosure. Moreover, if strictly more informative signals are strictly more costly, there is no equilibrium where information is withheld.

The remainder of the paper is structured as follows. Section 2 covers some mathematical preliminaries. Section 3 presents the model. The statement and the proof of the main result are in section 4. We discuss the relationship to the existing literature in the final section.

2 Mathematical preliminaries

In this section, we introduce notation and mathematical concepts that are building blocks of our model. Both the notation and the particular way of formalizing signals are taken from Gentzkow and Kamenica (2016).

2.1 State space and signals

Let Ω be a finite state space. A state of the world is denoted by $\omega \in \Omega$. A belief is denoted by μ. The prior distribution on Ω is denoted by μ_0. Let X be a random variable that is independent of ω and uniformly distributed on $[0,1]$ with typical realization x. We model signals as deterministic functions of ω and x. Formally, a *signal* π is a finite partition of $\Omega \times [0,1]$ s.t. $\pi \subset S$, where S is the set of non-empty, Lebesgue measurable subsets of $\Omega \times [0,1]$. We refer to any element $s \in S$ as a *signal realization*.

With each signal π, we associate an S-valued random variable that takes value $s \in \pi$ when $(\omega, x) \in s$. Let $p(s|\omega) = \lambda\left(\{x | (\omega, x) \in s\}\right)$ and $p(s) = \sum_{\omega \in \Omega} p(s|\omega)\, \mu_0(\omega)$ where $\lambda(\cdot)$ denotes the

Lebesgue measure. For any $s \in \pi$, $p(s|\omega)$ is the conditional probability of s given ω and $p(s)$ is the unconditional probability of s.

2.2 Lattice structure

The formulation of a signal as a partition induces an algebraic structure on the set of signals. This structure allows us to "add" signals together and thus easily examine their joint information content. Let Π be the set of all signals. Let $\trianglerighteq$ denote the refinement order on Π, that is, $\pi_1 \trianglerighteq \pi_2$ if every element of π_1 is a subset of an element of π_2. The pair $(\Pi, \trianglerighteq)$ is a lattice. The join $\pi_1 \vee \pi_2$ of two elements of Π is defined as the supremum of $\{\pi_1, \pi_2\}$. For any finite set (or vector) P, we denote the join of all its elements by $\vee P$. We write $\pi \vee P$ for $\pi \vee (\vee P)$. Note that $\pi_1 \vee \pi_2$ is a signal that consists of signal realizations s such that $s = s_1 \cap s_2$ for some $s_1 \in \pi_1$ and $s_2 \in \pi_2$. Hence, $\pi_1 \vee \pi_2$ is the signal that yields the same information as observing both signal π_1 and signal π_2. In this sense, the binary operation $\vee$ "adds" signals together.

2.3 Distributions of posteriors

A *distribution of posteriors*, denoted by τ, is an element of $\Delta(\Delta(\Omega))$ that has finite support.[6] Observing a signal realization s s.t. $p(s) > 0$ generates a unique posterior belief[7]

$$\mu_s(\omega) = \frac{p(s|\omega)\,\mu_0(\omega)}{p(s)} \text{ for all } \omega.$$

We write $\langle\pi\rangle$ for the distribution of posteriors induced by π. It is easy to see that $\tau = \langle\pi\rangle$ assigns probability $\tau(\mu) = \sum_{\{s \in \pi : \mu_s = \mu\}} p(s)$ to each μ.

3 The model

3.1 The baseline game

DM has a continuous utility function $u(a, \omega)$ that depends on her action $a \in A$ and the state of the world $\omega \in \Omega$. There are $n \geq 1$ experts indexed by i. Each expert i has a continuous utility function $v_i(a, \omega)$ that depends on DM's action and the state of the world. All experts and DM share the prior μ_0. The action space A is compact.

[6] The fact that a distribution of posteriors has finite support follows from the assumption that each signal has finitely many realizations. The focus on such signals is without loss of generality under the maintained assumption that Ω is finite.

[7] For those s with $p(s) = 0$, set μ_s to be an arbitrary belief.

The timing in this *baseline game* is as follows:

1. Each expert simultaneously chooses a signal π_i, the choice of which is not observed by the other experts.

2. Each expert privately observes the realization s_i of his own signal.

3. Each expert simultaneously sends a message $m_i \in M(s_i)$ to DM.

4. DM observes the signals chosen by the experts and the messages they sent.

5. DM chooses an action.

Function $M(\cdot)$ specifies the set of messages that are feasible upon observing a given signal realization. Let $\mathscr{M} = \cup_{s \in S} M(s)$ denote the set of all possible messages. For each $m \in \mathscr{M}$, let $T(m) = \{s \in S | m \in M(s)\}$. We say that message m *verifies* s if $T(m) = \{s\}$. We assume that for each $s \in S$ there exists a unique message that verifies it.[8]

For each expert i, let $c_i : \Pi \to \bar{\mathbb{R}}_+$ denote the cost of each signal.[9] Expert i's payoff in state ω is thus $v_i(a, \omega) - c_i(\pi)$ if he chooses signal π and DM takes action a. We assume that more informative signals are more expensive: $\pi \trianglerighteq \pi' \Rightarrow c_i(\pi) \geq c_i(\pi')$ for any $\pi, \pi' \in \Pi$ and any i. This is an important assumption that is absolutely necessary for our result. If an expert can save money by becoming more informed and withholding the additional information, full disclosure will not always happen in equilibrium and thus disclosure requirements will change the set of equilibrium outcomes.

Let $\sigma_i = \left(\pi_i, (\gamma_i^\pi)_{\pi \in \Pi}\right)$ denote expert i's strategy and $\boldsymbol{\sigma}$ denote a strategy profile. A strategy for expert i consists of a signal $\pi_i \in \Pi$ and a feasible messaging policy[10] $\gamma_i^\pi : S \to \Delta(M)$ following each signal π.[11] Let $\tilde{\mu}(\boldsymbol{\pi}, \boldsymbol{m}) \in \Delta(S^n)$ denote DM's belief about the signal realizations observed by the experts given the observed signals $\boldsymbol{\pi}$ and messages $\boldsymbol{m}$. The structure of the information sets requires DM's belief about expert i's signal realization to have support in $T(m_i)$ (the set of signal realizations for which m_i was an available message). Since DM knows the experts' choices of signals, each belief about the signal realizations implies a unique belief about ω. Throughout the paper, we assume that DM has a unique optimal action at any given belief about ω, i.e.,

[8] Assuming that this message is unique is not needed for our result. However, it simplifies our notation by making the truthful messaging policy unique.

[9] $\bar{\mathbb{R}}_+$ denotes the affinely extended non-negative real numbers: $\bar{\mathbb{R}}_+ = \bar{\mathbb{R}} \cup \{\infty\}$. Allowing the cost to be infinite for some signals incorporates the cases where some expert might not have access to particular signals.

[10] A messaging policy γ_i^π is feasible if $Supp(\gamma_i^\pi(s)) \subset M(s)$ for all s.

[11] As we focus on pure-strategy equilibria, we do not introduce notation for mixed strategies in the choice of π_i.

$a^*(\mu) \equiv argmax_{a\in A} E_\mu [u(a,\omega)]$ is single-valued for all μ. By the theorem of the maximum, the fact that $a^*(\cdot)$ is single-valued implies that it is continuous.

Let $\mathcal{B}^i(\boldsymbol{\sigma}_{-i}, \tilde{\mu})$ denote the best-response correspondence for expert i, i.e., a strategy σ_i is in this set if it is a best response, at every information set, to other players' strategies $\boldsymbol{\sigma}_{-i}$ and to the belief function $\tilde{\mu}$. Expert i's information sets are the "initial" node where he selects a signal and each possible (π, s) s.t. $\pi \in \Pi$ and $s \in \pi$. Note that this best-response correspondence does not depend on DM's strategy; because DM has a unique optimal action at every belief, expert i can take her behavior (given $\tilde{\mu}$) as fixed. A pair $(\boldsymbol{\sigma}, \tilde{\mu})$ is a (perfect Bayesian) equilibrium if $\tilde{\mu}$ obeys Bayes' rule on the equilibrium path and $\sigma_i \in \mathcal{B}^i(\boldsymbol{\sigma}_{-i}, \tilde{\mu})$ for all i. We say σ_i is a pure strategy if it employs a messaging policy that is deterministic on the equilibrium path (i.e., $\gamma_i^{\pi_i}$ is deterministic). An equilibrium is a pure-strategy equilibrium if each σ_i is a pure strategy. We define the *outcome* of the game to be the joint distribution of the state of the world, DM's belief, DM's action, and all the players' payoffs.

A pure strategy σ_i defines a partition of $\Omega \times [0,1]$ with each message sent in equilibrium corresponding to one element of the partition. We denote this signal equivalent of σ_i by $r(\sigma_i)$. Note that if $\sigma_i = \left(\pi_i, (\gamma_i^\pi)_{\pi\in\Pi}\right)$, then $\pi_i \trianglerighteq r(\sigma_i)$, which implies that $c_i(\pi_i) \geq c_i(r(\sigma_i))$. Given a strategy profile $\boldsymbol{\sigma} = (\sigma_1, ..., \sigma_n)$, let $\boldsymbol{r}(\boldsymbol{\sigma})$ denote $(r(\sigma_1), ..., r(\sigma_n))$.

3.2 Disclosure requirements

We now define an alternative game with *disclosure requirements* in which experts are required to disclose their private information truthfully.

Let γ^* denote the truthful messaging policy, i.e., $\gamma^*(s)$ places probability one on the message that verifies s. Under disclosure requirements, each expert i must set $\gamma_i^\pi = \gamma^*$ for every $\pi \in \Pi$. Accordingly, we can represent each expert's strategy simply as π_i and let $\boldsymbol{\pi}$ denote a strategy profile.

Let $\mathcal{B}_{DR}^i(\boldsymbol{\pi}_{-i})$ denote the best-response correspondence for expert i under disclosure requirements. A strategy profile $\boldsymbol{\pi}$ is a pure-strategy equilibrium under disclosure requirements if and only if $\pi_i \in \mathcal{B}_{DR}^i(\boldsymbol{\pi}_{-i})$.

4 Main result

Our main result is the following:

Theorem 1. *Disclosure requirements do not change the set of pure-strategy equilibrium outcomes.*

The statement of this result does not require any assumptions about the number of experts nor about the experts' or DM's utility functions. Moreover, the theorem does not only state that there exists some equilibrium of the baseline game where all information is disclosed. Rather, it makes a stronger claim that disclosure requirements have no impact whatsoever on the entire set of equilibrium outcomes.

The remainder of this section provides a proof of theorem 1. Let Σ^{π} denote the set of strategies in the baseline game that select signal π. Let Σ^{*} denote the set of strategies in the baseline game that utilize truthful messaging on the equilibrium path, i.e., strategies of the form $\left(\pi', (\gamma^{\pi})_{\pi \in \Pi}\right)$ s.t. $\gamma^{\pi'} = \gamma^{*}$. To show that any outcome under disclosure requirements is also an outcome of the baseline game, it will suffice to establish the following lemma.

Lemma 1. *For every $\boldsymbol{\pi}$ such that $\pi_i \in \mathcal{B}^i_{DR}(\boldsymbol{\pi}_{-i})$ for all i, there exist $\boldsymbol{\sigma}$ and $\tilde{\mu}$ such that*

(i) $\tilde{\mu}$ obeys Bayes' rule given $\boldsymbol{\sigma}$

(ii) $\sigma_i \in \mathcal{B}^i(\boldsymbol{\sigma}_{-i}, \tilde{\mu})$ for all i

(iii) $\sigma_i \in \Sigma^{*} \cap \Sigma^{\pi_i}$ for all i

Proof. We begin the proof by introducing a class of auxiliary games $G^i(\pi_S, \pi_R)$. There is a single expert with utility v_i and a DM with utility u. The timing is: (i) DM privately observes a signal realization s_R from signal π_R; (ii) the expert privately observes a signal realization s_S from signal π_S; (iii) the expert sends a message $m \in M(s_S)$ to DM; (iv) DM takes an action. Let γ denote the expert's messaging strategy and $\eta(m)$ denote DM's beliefs, given m, about s_S. An equilibrium of this game is a pair (γ, η) s.t. η obeys Bayes' rule on the equilibrium path and, at each information set s_S, γ is the best response to η. (Since DM has a unique optimal action for every belief about ω, we do not need to specify her behavior given η.) Standard arguments ensure that, given any π_S and π_R, there exists a perfect Bayesian equilibrium of $G^i(\pi_S, \pi_R)$.

To construct the requisite $\boldsymbol{\sigma}$ and $\tilde{\mu}$ in the statement of lemma 1, we begin by imposing condition (iii): for each expert i, we specify that σ_i selects π_i at the initial information set and that $\gamma^{\pi_i} = \gamma^{*}$. We also begin the construction of $\tilde{\mu}$ by imposing condition (i): $\tilde{\mu}$ follows Bayes' rule given $\boldsymbol{\pi}$ followed by truthful messaging by all of the experts.

We next construct off-equilibrium messaging policies $(\gamma_i^{\pi})_{\pi \neq \pi_i}$ for each expert. Consider expert i. Given any $\pi \in \Pi$, consider the auxiliary game $G^i(\pi, \vee_{j \neq i} \pi_j)$, which has some equilibrium (γ', η). We set γ^{π} to γ' and we set $\tilde{\mu}((\pi, \boldsymbol{\pi}_{-i}), (m_i, \boldsymbol{m}_{-i}))$ for equilibrium $\boldsymbol{m}_{-i}$ as follows: DM's belief about signal realizations observed by other experts is already pinned down (since other experts are playing truthful messaging policies), and we set DM's belief about the signal realization observed by expert i to $\eta(m_i)$. We repeat this procedure for each $\pi \neq \pi_i$. This completes the construction of σ_i and specifies $\tilde{\mu}$ on all of DM's information sets off the equilibrium path where only expert i deviates from σ_i. We can then repeat this procedure expert by expert and thus construct the entire strategy profile $\boldsymbol{\sigma}$, as well as $\tilde{\mu}$ on all of DM's information sets where only one expert deviates from $\boldsymbol{\sigma}$. We can choose an arbitrary specification of $\tilde{\mu}$ on DM's information sets where multiple experts deviate from $\boldsymbol{\sigma}$.

These $\boldsymbol{\sigma}$ and $\tilde{\mu}$ satisfy conditions (i) and (iii) of lemma 1 by construction. For any expert i, any $\pi \neq \pi_i$, and any $s \in \pi$, condition (ii) is satisfied on the information set (π, s) because γ_i^π is an equilibrium messaging policy of $G^i(\pi, \vee_{j\neq i}\pi_j)$. To show that condition (ii) is satisfied on the equilibrium path (on information sets (π_i, s)), we need to establish that γ^* is an equilibrium messaging policy of $G^i(\pi_i, \vee_{j\neq i}\pi_j)$. We know $G^i(\pi_i, \vee_{j\neq i}\pi_j)$ has some equilibrium, say (γ, η). It will suffice to show that given η, expert i's payoff from γ^* is the same as his payoff from γ following any s. Denote these payoffs by $y^*(s)$ and $y(s)$, respectively. Since (γ, η) is an equilibrium, we know

$$y(s) \geq y^*(s) \;\; \forall s \in \pi_i. \tag{1}$$

Moreover, since $\pi_i \in \mathcal{B}^i_{DR}(\boldsymbol{\pi}_{-i})$, we know

$$\sum_{\omega\in\Omega}\sum_{s\in\pi_i} y^*(s)\, p(s|\omega)\, \mu_0(\omega) \geq \sum_{\omega\in\Omega}\sum_{s\in\pi_i} y(s)\, p(s|\omega)\, \mu_0(\omega). \tag{2}$$

Otherwise, under disclosure requirements expert i could profitably deviate from π_i to the signal which "garbles π_i by γ", i.e., the signal $r\left(\pi_i, \left((\gamma^\pi)_{\pi\neq\pi_i}, \gamma^{\pi_i} = \gamma\right)\right)$. (Note that $r(\sigma_i)$ only depends on the messaging policy that σ_i specifies on the equilibrium path.) Combining inequalities (1) and (2) yields $y(s) = y^*(s)\ \forall s \in \pi_i$.

It remains to show that condition (ii) is satisfied for each expert at the initial information set where he chooses his signal. Let $\hat{v}_i(\mu) \equiv \mathrm{E}_\mu\, v_i(a^*(\mu), \omega)$. Let v_i^* be expert i's payoff under $\boldsymbol{\sigma}$ and $\tilde{\mu}$. Since $\pi_i \in \mathcal{B}^i_{DR}(\boldsymbol{\pi}_{-i})$, we know

$$v_i^* \geq \mathrm{E}_{\langle \pi\vee\boldsymbol{\pi}_{-i}\rangle}[\hat{v}_i(\mu)] - c_i(\pi) \; \forall \pi \in \Pi. \tag{3}$$

Suppose expert i deviates from $\sigma_i = \left(\pi_i, (\gamma_i^\pi)_\pi\right)$ to $\sigma_i' = \left(\pi_i', (\gamma_i^\pi)_\pi\right)$. By the construction of $\tilde{\mu}$ through η, we know the distribution of DM's posterior must be $\langle r(\sigma_i') \vee \boldsymbol{\pi}_{-i}\rangle$. Hence, this deviation yields the payoff

$$\mathrm{E}_{\langle r(\sigma_i')\vee\boldsymbol{\pi}_{-i}\rangle}\left[\hat{v}_i\left(\mu\right)\right]-c_i\left(\pi_i'\right)\quad\leq\quad\mathrm{E}_{\langle r(\sigma_i')\vee\boldsymbol{\pi}_{-i}\rangle}\left[\hat{v}_i\left(\mu\right)\right]-c_i\left(r\left(\sigma_i'\right)\right)\leq v_i^*$$

where the first inequality follows from the fact that $\pi_i' \trianglerighteq r(\sigma_i')$ implies $c_i(\pi_i') \geq c_i(r(\sigma_i'))$ and the second inequality follows from equation (3). Since the deviation yields a weakly lower payoff than v_i^*, condition (ii) is also satisfied at the initial information set. □

To show that any outcome of the baseline game is also an outcome under disclosure requirements, we begin with the following lemma.

Lemma 2. *Suppose* $(\boldsymbol{\sigma}, \tilde{\mu})$ *is a pure-strategy equilibrium of the baseline game. Then,* $r(\sigma_i) \in \mathcal{B}_{DR}^i(\boldsymbol{r}(\boldsymbol{\sigma}_{-i}))$ *for all* i*.*

Proof. Consider any expert i. His equilibrium strategy is some $\sigma_i = (\pi_i, (\gamma^\pi)_\pi)$. Let v^* denote his equilibrium payoff. Given $(\boldsymbol{\sigma}_{-i}, \tilde{\mu})$, for every signal $\pi' \in \Pi$, let $v_{\pi'}$ denote his payoff if he deviates to strategy $(\pi', (\gamma^\pi)_\pi)$, and let $v_{\pi'}^*$ denote his payoff if he deviates to strategy $(\pi', (\gamma^*)_\pi)$. Since $(\boldsymbol{\sigma}, \tilde{\mu})$ is an equilibrium, we know: (i) $v^* \geq v_{\pi'}$ for all π' (by the fact that σ_i was the best response at the initial information set); and (ii) $v_{\pi'} \geq v_{\pi'}^*$ for all π' (by the fact that σ_i was the best response at each information set (π', s)). Finally, if expert i deviates to strategy $(r(\sigma_i), (\gamma^*)_\pi)$, his payoff under this deviation must also be v^*, so $v_{r(\sigma_i)}^* = v^*$. Combining this with inequalities (i) and (ii), we obtain $v_{r(\sigma_i)}^* \geq v_{\pi'}^*$ for all π'. Since for all π', $v_{\pi'}^* = \mathrm{E}_{\langle \pi' \vee \boldsymbol{r}(\boldsymbol{\sigma}_{-i})\rangle}[\hat{v}_i(\mu)]$, this implies $r(\sigma_i) \in \mathcal{B}_{DR}^i(\boldsymbol{r}(\boldsymbol{\sigma}_{-i}))$. □

Lemma 2 shows that, given any equilibrium of the baseline game, there is an equilibrium under disclosure requirements that induces the same joint distribution of the state of the world, DM's beliefs, DM's actions, and DM's payoffs. It only remains to add experts' costs of signals to this list. These costs could only be different if, in the equilibrium of the baseline game, some expert i was utilizing a strategy $\sigma_i = \left(\pi_i, (\gamma_i^\pi)_\pi\right)$ s.t. $c_i(\pi_i) > c_i(r(\sigma_i))$. But this could not happen since it would then be profitable for expert i to deviate to a strategy in $\Sigma^{r(\sigma_i)} \cap \Sigma^*$.

This completes the proof of theorem 1.

5 Related Literature

5.1 Endogenous information and cheap talk

The most relevant existing work is analysis by Pei (2015). He establishes that if a sender endogenously and overtly acquires costly information and subsequently engages in a cheap talk game with a receiver, then in every pure-strategy equilibrium, the sender fully reveals all of his information, regardless of preferences. We establish a closely related result in a setting where the sender's messages are verifiable.

5.2 Persuasion games with verifiable types

As mentioned in the introduction, a large literature examines disclosure of exogenous private information in persuasion games, i.e., settings where informed expert(s) can send verifiable messages. Milgrom (1981) shows that full disclosure is a unique equilibrium outcome when there is a single expert who can send any verifiable message and whose preferences are monotonic (i.e., whether the expert, who knows the true state is ω^*, prefers DM to believe the state is ω or ω' does not depend on ω^*). Since this early contribution, this literature has evolved along three distinct dimensions.

Weakening monotonicity. Seidmann and Winter (1997), Giovannoni and Seidmann (2007), and Mathis (2008) show that existence of a fully revealing equilibrium can be guaranteed if we replace the monotonicity assumption with a somewhat weaker single-crossing property. Hagenbach *et al.* (2014) introduce a general model that encompasses much of this literature and establish a simple condition that is both necessary and sufficient for existence of a fully revealing equilibrium.

Weakening verifiability. Milgrom (1981) assumes that the set of messages is the power set of the experts' type. Okuno-Fujiwara *et al.* (1990) and others point out that other message spaces can be assumed and that full revelation can remain the unique outcome even if the expert cannot always verify his type. For example, it would not matter if the expert could not prove that he is a "low" type. In the spirit of these results, we put limited structure on sets $M(s)$ and only impose the key assumption that for each s there is a message that verifies it.

Introducing multiple experts. In all of the aforementioned papers, full revelation is driven by some version of the "unraveling argument" – if some types pool, at least one of them is "better" than the "average" and will prefer to reveal himself. When there are multiple experts, however, there are other forces that can lead to full revelation.[12] If for each state there is some expert who wishes to

[12] Okuno-Fujiwara *et al.* (1990) and Hagenbach and Koessler (2016) apply arguments related to unraveling in

disclose the state to DM so as to avoid her default action, full revelation is an equilibrium (Milgrom and Roberts 1986). Lipman and Seppi (1995) establish that, as long as DM knows experts have conflicting preferences, there is a full revelation equilibrium, even under limited verifiability.

In contrast to this literature, we show that if private information is endogenous and gathered overtly, full revelation is always an equilibrium for any number of experts and for any configuration of preferences (regardless of monotonicity, single-crossing, or conflict). Moreover, given any equilibrium, there is a full revelation equilibrium that induces the same outcomes.

In most models of verifiable communication, there is a set of experts who wish to influence a third party (DM). That said, some papers examine environments where experts disclose private information and then play games with each other (Okuno-Fujiwara *et al.* 1990, Hagenbach and Koessler 2016, Hagenbach *et al.* 2014). When private information is exogenous, this distinction is not particularly important – it does not matter whether the publicly disclosed information impacts experts' payoff through an action of a third-party or through the equilibrium outcome of the post-disclosure game. Our results, however, only apply to the environments where experts seek to influence a third party. Once experts' information is endogenous, the publicly disclosed information is no longer sufficient to determine the payoffs of a post-disclosure game.

Finally, existing literature considers both settings where experts are informed about a common state of the world (e.g., Milgrom and Roberts 1986, Lipman and Seppi 1995) and settings where each expert has private information only about his own type (Okuno-Fujiwara *et al.* 1990, Hagenbach and Koessler 2016, Hagenbach *et al.* 2014). Our model covers both of these case. If we let $c_i(\cdot)$ be the same for all experts (e.g., $c_i(\pi) = 0$ for all π), then all experts can become privately informed about the "common" state ω. Alternatively, suppose that $\Omega = T_1 \times ... \times T_n$ where T_i is the set of possible types of expert i. Then, we can set $c_i(\pi) = \infty$ for any π that is informative about types other other than i and thus capture settings where each expert can only get information about his own type – he cannot learn about nor disclose to DM any information about any other expert.

5.3 Bayesian persuasion with rich signal spaces

In Kamenica and Gentzkow (2011), we analyze a game where a single expert wishes to influence DM's action by choosing an observable, costless signal about the state of the world. We focus on identifying the conditions under which the expert benefits from the ability to generate such a signal and on characterizing the distribution of DM's beliefs under the optimal signal. In the context of

settings with multiple experts.

that model, we point out that the outcome would be the same if DM did not observe the signal realizations directly, but the expert could send verifiable messages. That observation is the starting point for the analysis in this paper.

In Gentzkow and Kamenica (2016), we consider a special case of the baseline game in this paper. We assume that each expert can select a signal whose realization (conditional on ω) is arbitrarily correlated with the realizations of other experts' signals.[13] In that context, we draw on the main result from this paper to simplify the analysis by reformulating the environment as one where DM directly observes the outcomes of the experts' signals.

[13]This is the case here if we set $c_i(\pi) = 0$ for all $\pi \in \Pi$. The general model in this paper also encompasses environments, however, where correlation in signal realizations across senders is not possible.

References

Crawford, Vincent P., and Joel Sobel. 1982. Strategic information transmission. *Econometrica.* 50:1431-1451.

Gentzkow, Matthew, and Emir Kamenica. 2016. Bayesian persuasion with multiple senders and rich signal spaces. Working paper.

Giovannoni, Francesco, and Daniel J. Seidmann. 2007. Secrecy, two-sided bias and the value of evidence. *Games and Economic Behavior.* 59:296-315.

Grossman, Sanford J. 1981. The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics.* 24:461-83.

Hagenbach, Jeanne, and Frederic Koessler. 2016. Full disclosure in decentralized organizations. *Economics Letters.* 139:5-7.

Hagenbach, Jeanne, Frederic Koessler, and Eduardo Perez-Richet. 2014. Certifiable pre-play communication: Full disclosure. *Econometrica.* 82:1093-1131.

Kamenica, Emir, and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review.* 101:2590-2615.

Lipman, Barton L., and Duane J. Seppi. 1995. Robust inference in communication games with partial provability. *Journal of Economic Theory.* 66:370-405.

Mathis, Jerome. 2008. Full revelation of information in Sender-Receiver games of persuasion. *Journal of Economic Theory.* 143:571-584.

Milgrom, Paul. 1981. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics.* 12:380-391.

Milgrom, Paul. 2008. What the seller won't tell you: Persuasion and disclosure in markets. *Journal of Economic Perspectives.* 22:115-131.

Milgrom, Paul, and John Roberts. 1986. Relying on the information of interested parties. *RAND Journal of Economics.* 17:18-32.

Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura. 1990. Strategic information revelation. *Review of Economic Studies.* 57:25-47.

Pei, Harry Di. 2015. Communication with endogenous information acquisition. *Journal of Economic Theory.* 160:132-149.

Seidmann, Daniel J., and Eyal Winter. 1997. Strategic information transmission with verifiable messages. *Econometrica.* 65:163-169.

Shin, H.S. 1994. News management and the value of firms. *RAND Journal of Economics*. 25:58-71.