

Quid pro quo: Friendly information exchange between rivals

ANDREAS BLUME

Department of Economics, University of Arizona

IN-UCK PARK

School of Economics, University of Bristol

We show that information exchange via disclosure is possible in equilibrium even when it is certain that whenever one party learns the truth, the other loses. The incentive to disclose results either from an expectation of disclosure being reciprocated—the quid pro quo motive—or from the possibility of learning from the rival’s failure to act in response to a disclosure—the screening motive. Alternating and gradual disclosures are generally indispensable for information exchange and the number of disclosure rounds grows without bound if the agents’ initial information becomes sufficiently diffuse; in that sense, the less informed agents are, the more they talk. Patient individuals can achieve efficiency by means of continuous alternating disclosures of limited amounts of information. This provides a rationale for protracted dialogues.

KEYWORDS. Alternating information disclosure, dynamic communication, two-sided incomplete information.

JEL CLASSIFICATION. C72, C73, D82, D83.

1. INTRODUCTION

Can two decision makers, each of whom has critical information concerning which of a number of possible actions is the correct one, share that information when both have a preference for acting on it alone? This question arises naturally in research and development races/joint ventures, when multiple government agencies collect information intended to avert a terrorist attack, or when separate researchers work on a common problem, as happened in the pursuit of a proof of Fermat’s last theorem. While there may be a common benefit to making the correct decision, e.g., when a new technology is developed, a terrorist is arrested in the planning stage of an attack, or a chosen proof strategy yields results, the desire to be the principal beneficiary of an invention or to receive primary credit may stand in the way of information sharing.

Andreas Blume: ablume@arizona.edu

In-Uck Park: i.park@bristol.ac.uk

We are grateful to two anonymous referees for their advice. We thank Ken Binmore, Alessandro Bonatti, Robert Evans, Francesco Giovannoni, Piero Gottardi, Philippe Jehiel, Eric Maskin, and seminar participants at EUI, Korea University, the 6th CSEF-IGIER Symposium on Economics and Institutions, and 2015 European Summer Symposium in Economic Theory in Gerzensee for helpful comments at various stages of this project.

Suppose, for example, that each of two rival intelligence agencies conducted an independent investigation of a crime and came up with a list of multiple suspects. If they knew that combining their information would reveal who is responsible, would they voluntarily share their information, even when both are motivated to be the first to identify the true culprit? If they did, in what manner would/should the information exchange take place? This paper provides some new insights into these questions by delineating the key factors that incentivize information exchange in such environments. The main findings are (i) disclosures are made in anticipation of obtaining information in return, (ii) due to the risk from disclosing too much information, information sharing is necessarily gradual, requiring multiple rounds of alternating disclosures, (iii) the necessary number of disclosure rounds to guarantee that the truth will eventually be discovered grows without bound as initial information becomes more diffuse, and (iv) irrespective of the initial information, as long as the payoff from taking the correct action is at least double the disutility from the rival taking the correct action, there always exists an equilibrium in which information exchange continues until the truth is discovered.

In this paper we investigate the case where monetary incentives are not available and instead individuals are motivated by concerns for the future. Then there are two possible reasons for providing information, the *quid pro quo* reason, which arises from the expectation that information will be disclosed by the other party if (and only if) information is first disclosed to that party, and the *screening* reason, which relies on the fact that information may be gleaned from others not acting on information provided. The quid pro quo reason is familiar from many dynamic environments in which, in equilibrium, individuals forgo short-term gain in the interest of future payoffs, and, in particular, is related to incremental exchange, incremental public goods provision, and turn taking. The screening reason, as the name suggests, is reminiscent of dynamic screening settings where, for example, a seller extracts information about a buyer's valuation for an object by tempting the buyer with a sequence of price offers. We show that the combination of quid pro quo and screening motivations generates intertemporal incentives that may counter the detrimental effects of the desire for primacy.

Our focus is on the exchange of information via disclosure. We represent (payoff-relevant) information as a subset of some (payoff) state space. Initially, two agents independently and privately learn a finite set, their "possibility set," to which the true state belongs. To avoid degeneracy, we assume that it is common knowledge that combining their information is useful in the sense that it reveals the state of the world without error, i.e., the true state of the world is the unique common element of both agents' possibility sets.

Each period, agents make one of two kinds of choices: either they take an action or they make a disclosure. For simplicity, we identify the space of actions with the state space. Each agent's objective is to take the action that corresponds to (is optimal in) the true state of the world. Having the other agent take the correct action is less desirable than no action being taken but not as damaging as taking a wrong action. If and when the correct action is taken, the game ends.

Agents disclose information by revealing states in their possibility sets. They need not disclose fully but must be truthful. Thus, a disclosure decision amounts to picking a subset of the undisclosed elements in one's possibility set. To highlight the role of disclosure, we shut down all other avenues for communication. For this reason we assume that agents have a uniform prior over the state space and lack a common language for the undisclosed elements of the state space, so that the only property of a disclosed set that matters is its size, not the identity of its elements.

With each disclosure an agent risks revealing the true state, thereby giving the other agent an opportunity to identify the true state and act on it. For any agent to disclose, therefore, there should be a prospect for him to be able to identify the true state in the future, for instance, because the other agent is expected to disclose in return. However, this quid pro quo reason is not enough to initiate information exchange because there can only be a finite number of disclosures and the last disclosure cannot be motivated by this reason. The aforementioned screening reason comes to the rescue here: If the one to disclose last disclosed all but one element in his possibility set, he retains the prospect of identifying the undisclosed element as the true state should the other agent not end the game after the last disclosure. This reasoning illuminates some key equilibrium features: (a) each disclosure must be motivated by a future prospect of obtaining enough information in return; (b) once started, the agents take turns in disclosing information until the true state of the world is identified by one of the agents; and (c) since disclosing too much information at once is too risky, communication necessarily takes the form of prolonged dialogues during which both agents become increasingly informed.

We study equilibrium behavior as the time delay between choices vanishes. We begin by outlining equilibrium behavior in an example in which agents have relatively accurate information, i.e., their possibility sets contain no more than three elements. In the baseline example, it never takes more than four periods to identify the correct action. We use the example to illustrate the screening and quid pro quo motives, the role of our assumptions on payoffs, and to explore the impact of varying our assumptions about the agents' information.

For the general case, in which the agents' possibility sets may contain any number of states, we construct Markov equilibria that exhibit a maximum quid pro quo flavor. When agents are equally informed, one of them starts by disclosing a single element of his possibility set, after which agents alternate in disclosing pairs of states until one of them identifies the true state. When agents are unequally informed, the one starting with the larger possibility set discloses proportionally more information in alternating rounds until the true state is identified. The number of disclosure rounds grows without bound as the agents become less informed at the outset. Nonetheless, the equilibria converge to efficiency as the agents become infinitely patient, because the disclosure continues without delay until the true state is identified.

Following the seminal papers by [Grossman \(1981\)](#) and [Milgrom \(1981\)](#) on disclosure and [Crawford and Sobel \(1982\)](#) on cheap talk, an extensive literature has developed

on communication with costless messages. The disclosure strand of this literature permits senders to withhold information, but does not allow them to lie.¹ The cheap-talk literature places no constraints on sender messages.²

Multi-round communication in sender–receiver settings has been studied by [Forges \(1990a\)](#), [Amitai \(1996\)](#), [Aumann and Hart \(2003\)](#), [Krishna and Morgan \(2001, 2004\)](#), [Goltsman, Hörner, Pavlov, and Squintani \(2009\)](#), [Forges and Koessler \(2008\)](#), [Esö and Fong \(2010\)](#), and [Golosov, Skreta, Tsyvinski, and Wilson \(2014\)](#). There has been work on mediated communication by [Myerson \(1982\)](#) and [Forges \(1986\)](#) and recently in the Crawford–Sobel environment by [Goltsman et al. \(2009\)](#). Another line of papers characterize the set of equilibrium outcomes obtainable when static games of three or more players are augmented by unmediated communication, as in [Forges \(1990b\)](#), [Barany \(1992\)](#), [Ben-Porath \(2003\)](#), and [Gerardi \(2004\)](#). A general message from this literature is that with three or more players, one can find communication protocols for which the set of Nash equilibrium outcomes coincides with the set of equilibrium outcomes that can be achieved by mediation.

Single-round communication between multiple, privately informed players has been studied by [Fried \(1984\)](#), [Farrell and Gibbons \(1989\)](#), [Matthews and Postlewaite \(1989\)](#), [Okuno-Fujiwara, Postlewaite, and Suzumura \(1990\)](#), [Park \(2002\)](#), [Goltsman and Pavlov \(2014\)](#), and [Halonen-Akatwijuka and Park \(2021\)](#), among others. The present paper contributes to a small but growing literature on multi-round information exchange between privately informed parties with conflicting interests. [Stein \(2008\)](#) examines an environment in which competing players engage in continued exchange of newly developed ideas driven by the fact that future ideas can only be discovered if current ideas are shared. [Rosenberg, Solan, and Vieille \(2013\)](#) study repeated games with incomplete information and show that two players facing completely unrelated decision problems can engage in mutually beneficial information exchange. [Hörner and Skrzypacz \(2016\)](#) study the acquisition and gradual sale of information when there is no outside enforcement. [Dziuda and Gradwohl \(2015\)](#) and [Augenblick and Bodoh-Creed \(2018\)](#) examine information sharing between two agents who try to discover whether they can be jointly productive, but incur a loss that is increasing in how much information they reveal; the papers differ in the representation of information and assumptions about timing.

In [Stein \(2008\)](#), payoffs are complementary; in [Rosenberg, Solan, and Vieille \(2013\)](#), they are independent. In contrast, the current paper, like [Hörner and Skrzypacz \(2016\)](#), considers environments where payoffs are negatively correlated. Unlike in [Rosenberg, Solan, and Vieille \(2013\)](#), in the present paper each player's information is useful for both players. As a result, although we share with [Rosenberg, Solan, and Vieille \(2013\)](#) the insight that each disclosure is motivated by the anticipation of receiving information in return, owing to the aforementioned screening reason, we obtain full disclosure in finite time. In [Hörner and Skrzypacz \(2016\)](#), only one party (seller) has information valuable for the other (buyer) and, thus, although gradual disclosure in multiple rounds

¹Contributions include [Milgrom and Roberts \(1986\)](#), [Shin \(1994\)](#), [Seidmann and Winter \(1997\)](#), and [Glazer and Rubinstein \(2004\)](#).

²[Green and Stokey \(1980, 2007\)](#), [Austen-Smith \(1990\)](#), [Baliga and Morris \(2002\)](#), [Matthews \(1989\)](#), [Morgan and Stocken \(2003\)](#), [Ottaviani and Sørensen \(2006\)](#), and [Blume, Board, and Kawamura \(2007\)](#).

increases the total price by enhancing buyer's trust, it is not necessary for trade. In contrast, gradualism is generally indispensable for any information exchange in our setting because the risk of losing from each disclosure needs to be kept small enough to be offset by future prospects of winning from returned information in order for the process to be viable. In both [Dziuda and Gradwohl \(2015\)](#) and [Augenblick and Bodoh-Creed \(2018\)](#) there is a positive probability that both parties benefit ex post from sharing information. Our environment also differs in that we have no explicit cost of providing information in the payoff function. [Dziuda and Gradwohl \(2015\)](#) find that both gradual and immediate exchange may be possible and optimal. In Augenblick and Bodoh-Creed, information sharing is necessarily gradual.

The need for protracted information exchange that arises in our environment resembles incremental contributions studied in the public goods literature, where they help overcome the free-riding problem stemming from a lack of commitment technology ([Admati and Perry \(1991\)](#), [Marx and Matthews \(2000\)](#)). In the public goods environment with two agents, introducing the ability to offer one-sided commitments to reciprocate a contribution of the opponent with a contribution of one's own would remove the need for incremental contributions. In our setting, in contrast, introducing the ability to offer similar one-sided commitments to reciprocate information disclosures would not be enough, unless it were paired with the ability to commit not to act on the information until after information has been provided by both sides. Moreover, the learning component, which is essential in our setting and gives rise to the screening motive, is absent in the public goods setting. [Compte and Jehiel \(2004\)](#) identify an alternative source of gradualism in public goods and bargaining environments, namely the fear of raising one's opponent's termination option value too much by large concessions, which acts as a lower bound of equilibrium payoff. This aspect is not present in our setting.

The next section introduces the key elements of our environment through an example. Section 3 describes the model and equilibrium concept. Section 4 lays down some fundamental insights common to all equilibria, including the need for protracted dialogues. Section 5 defines Markov equilibrium and demonstrates that in general environments there are Markov equilibria with alternating and gradual disclosure that are asymptotically efficient. Section 6 contains some concluding remarks, followed by technical proofs in the [Appendix](#).

2. AN ILLUSTRATION

We begin by describing (perfect Bayesian) equilibrium behavior in a simple example. The example illustrates the role of the quid pro quo and screening motives in equilibrium, as well as our assumptions about payoffs and shared information.

Two investigative journalists from competing publishers are working on identifying a factory responsible for a chemical spill along a river. At any point in time (modeled as discrete periods) each has the option to disclose some or all of his or her information, to publish an article accusing one of the factories, or to do nothing. One has access to data that allow him to narrow the set of possible factories using the time of the spill and

the other has access to data that allow her to narrow the set on the basis of the types of chemicals used by the factories. From their prior disclosure history, it is commonly known between them that each has been able to narrow the set of candidates down to two. They also know, perhaps through a third party that has access to both of their information but is legally prevented from sharing details, that by combining their information they would be able to identify the factory; it is the unique factory suspected by both journalists.³ The game continues until the correct factory has been identified and published.

Publishing the correct identity of the factory would be a scoop worth a payoff $\alpha > 0$. Getting scooped by the other journalist would be a disappointment, resulting in a payoff $\beta < 0$. While it is disappointing to get scooped, it is not devastating, i.e., $\alpha + 2\beta > 0$. Simultaneously publishing the identity of the correct factory amounts to an equal chance of achieving the scoop and getting scooped, which results in a positive payoff, as getting scooped is not devastating. Going ahead with publishing an accusation of the wrong factory would be an embarrassment (with a loss of reputation and the possibility of getting fired), yielding a payoff γ that satisfies $(2\alpha + \gamma)/2 < \beta$, so that making a false accusation is far worse than getting scooped.

The screening motive

In the scenario just described, with a sufficiently large discount factor $\delta \in (0, 1)$, there is an equilibrium in which one of the journalists, say journalist 2, discloses one of her remaining two candidates to the other journalist, journalist 1. If the disclosed factory is in journalist 1's set of remaining candidates, i.e., his *possibility set*, then journalist 1 realizes that it must be the correct factory, since it is the one suspected by both journalists. In that case journalist 1 goes ahead and publishes an article naming the factory that caused the spill. If the candidate factory revealed by journalist 2 is not in journalist 1's possibility set, then journalist 1 remains uncertain and refrains from publishing an article naming a factory because it is too risky. Hence, if following journalist 2's disclosure journalist 1 does not publish, journalist 2 infers that the factory from her possibility set that she did not disclose must be the one having caused the spill and publishes an article correctly identifying that factory.

Assuming that journalist 1 remains inactive until journalist 2 makes a disclosure, journalist 2 has no incentive to postpone her disclosure. A postponement would simply delay receiving the payoff $(\delta\beta + \delta^2\alpha)/2$, which is positive for sufficiently large δ . Journalist 1 has an incentive to wait for journalist 2 to disclose, since by doing so his payoff is $(\delta\alpha + \delta^2\beta)/2$, whereas taking a chance by accusing a factory without further information is too risky, and disclosing a factory of his own only increases the chance of getting scooped by his competitor. The key to this construction is that journalist 2, by disclosing one of her candidates, gains information in the event that journalist 1 does not publish

³Note also that if the two journalists arrive at their starting information by individually narrowing their possibility sets down to a small remainder from a common large set of candidates (e.g., all factories in a region), the probability is high that there is a unique factory suspected by both. For more on this, see footnote 5 below.

in the following period. By disclosing, journalist 2 screens journalist 1 for that information. This is what we refer to as the *screening motive* for disclosure.

The quid pro quo motive

Change the scenario slightly by assuming that journalist 1 has three instead two factories in his possibility set. Everything else remains the same and, as before, this structure is common knowledge between the journalists. Then there is an equilibrium in which journalist 1 discloses one of his three candidates and journalist 2 does nothing in that period and after any history in which journalist 1 has not made a prior disclosure. In the following period, if the factory that journalist 1 disclosed is in journalist 2's possibility set, journalist 2 publishes the identity of the factory, ending the game. Otherwise, the game continues as in the description of the screening motive above, with journalist 2 unilaterally disclosing one of her candidates in the period following journalist 1's disclosure.

Journalist 1's payoff in this equilibrium is $u_1 = (\delta\beta + \delta^2\alpha + \delta^3\beta)/3 > 0$.⁴ Given his lack of information, there is no incentive for journalist 1 to publish initially. Journalist 1 would also not gain from postponing disclosure until a later period since that would mean receiving a discounted value of the payoff u_1 . Finally, journalist 1's payoff from disclosing two instead of only one of his candidates would be $(2\delta\beta + \delta^2\alpha)/3$, and therefore smaller than u_1 . Thus journalist 1 has an incentive to disclose one of his candidates in exchange for journalist 2 disclosing one of her candidates in the following period. This is what we refer to as the *quid pro quo motive*.

There are other equilibria, but every equilibrium entails disclosure(s) eventually leading to publication of the correct factory, that is, productive information sharing necessarily arises. First, "no disclosure ever" is unsustainable because then it would be a profitable deviation for journalist 2 to disclose a candidate due to the screening motive. Second, if the first disclosure did not reveal the correct factory to either journalist, disclosing for the screening purpose at that point is (weakly) even more profitable.

We next use our example to illustrate that two assumptions we make about agents' information in the paper are sufficient but not necessary for information exchange leading to discovery of the truth.

Uncertainty about the size of possibility sets

Now suppose that journalist 1 has either two or three candidates left in his possibility set and that journalist 2 is uncertain about which is the case. Maintain all other assumptions, including that it is commonly known that the journalists jointly know the truth.

Then for sufficiently large δ there exists an equilibrium in which journalist 1 discloses one of the candidates from his possibility set. If his possibility set contains two candidates, the incentive to disclose is driven by the screening motive, and otherwise by the quid pro quo motive. The strategies mirror those of the case just described. Notice

⁴For sufficiently large δ and assuming, as we do, that $\alpha + 2\beta > 0$.

that when the factory disclosed by journalist 1 does not belong to journalist 2's possibility set, journalist 2 discloses one of her own candidates. In the event that journalist 1 starts out with two candidates in his possibility set, this disclosure is irrelevant (since journalist 1 publishes the true identity of the polluting factory in the subsequent period anyway), and otherwise it is a screening move by journalist 2.

Uncertainty about joint information

Return to the environment where it is commonly known that both journalists have two candidates remaining in their possibility sets. Now, however, consider the case that the journalists no longer jointly know the truth: instead of knowing that the intersection of their possibility sets is a singleton, they cannot rule out that their possibility sets coincide. That is, in this particular example, the journalists think it equally likely that they have only one candidate in common, which is the polluting factory, and that they have both candidates in common, one of which is the polluting factory.

Suppose the journalists behave as follows: Journalist 1 discloses one candidate in period 1; in period 2, journalist 2 publishes the identity of the disclosed factory if it belongs to her possibility set and does nothing otherwise; in period 3, both journalists publish the remaining candidate if journalist 2 published the identity of the disclosed factory, which turned out to be wrong, while journalist 1 publishes his remaining candidate if journalist 2 did nothing in period 2. It can be shown (see the [Appendix](#)) that this is equilibrium behavior for sufficiently large δ provided that γ is in a small neighborhood of $3(\beta - \alpha)/2$. Notice that in this equilibrium, journalist 2 sometimes goes ahead with publication without being entirely certain about the identity of the polluting factory. The constraint on γ captures that with γ too large (too near to β), there would be too strong an incentive to publish right at the outset without any additional information, and with γ too small (in the direction of $-\infty$), journalist 2 would be deterred from publication even after seeing a disclosure in her possibility set. In summary, if the stated constraints are satisfied, the equilibrium exists and exhibits screening. In the event that journalist 1's disclosure belongs to journalist 2's possibility set, journalist 2 has enough information to be willing to publish and risk potential failure.

One can show that if γ becomes extremely small (in the [Appendix](#) we consider $\gamma = -\infty$), there does not exist *any* equilibrium in which the identity of the polluting factory is revealed. The journalists have to be sufficiently optimistic about the value of combining their information for the screening and quid pro quo motives to lead to information sharing. In our main analysis this is ensured by the assumption that players are jointly able to identify the truth.

3. MODEL

There is a finite set Ω of (payoff) states. Two agents, 1 and 2, are interested in identifying the *true state*. At the beginning each agent $i \in \{1, 2\}$ privately learns a subset of the state space, denoted by $S_i \subset \Omega$ and referred to as his *possibility set*, that contains the true state.

For both i , $\#(S_i) = \nu_i > 1$ and $\#(S_1 \cap S_2) = 1$. Thus, agents can jointly but not individually identify the true state. Define

$$\mathcal{S}(\nu_1, \nu_2) := \{(R_1, R_2) \subset 2^\Omega \times 2^\Omega \mid \#(R_i) = \nu_i \text{ and } \#(R_1 \cap R_2) = 1\}.$$

The set $\mathcal{S}(\nu_1, \nu_2)$ is the set of “states of the world” in the usual sense, as it determines both which payoff state is the true state and what information players have. We will use “state” throughout to refer to payoff states. The game begins with nature drawing a pair (S_1, S_2) from a uniform distribution on $\mathcal{S}(\nu_1, \nu_2)$. The lone element of $S_1 \cap S_2$ becomes the true state, which will be denoted by ω^* .⁵ The manner in which possibility sets and the true state are determined is assumed to be common knowledge between the two agents, as is the remainder of the description of the game below.

Notice that according to this formulation all elements of Ω play a symmetric role in the determination of players’ initial information and in the selection of the true state. Thus a priori the names of states do not matter; this will enable us later to restrict attention to players using strategies that treat states identically as long as they have not been distinguished by the history of play.

After learning their possibility sets privately, the two players play a potentially infinite-horizon game as described below. In each period $t = 1, 2, \dots$, as long as the game has not ended by then, each player i has the option to make a *move*, which is either a *disclosure* of a nonempty subset of S_i (of elements that have not been disclosed already), or an *action*, where each player’s set of possible actions coincides with the state space, Ω . Alternatively, either player may opt to do nothing. The two players’ moves are simultaneous in each period. The game ends when either player takes an action that is ω^* .⁶

Formally, the set of possible choices in period 1 is $C_i = 2^{S_i} \cup S_i$ for player $i = 1, 2$, where $D \in 2^{S_i} \setminus \{\emptyset\}$ denotes disclosing a nonempty subset D of S_i , $\omega \in S_i$ denotes taking the action $\omega \in S_i$, and $\emptyset$ denotes *doing nothing*.⁷ To avoid confusion between disclosing $\{\omega\}$ and taking the action ω , we denote the latter as $\langle \omega \rangle$ in the sequel. Also, for ease of terminology in exposition, doing nothing is considered a choice but not a move. The outcome of period 1, denoted by c^1 , records the choices made by the two players in period 1, that is, $c^1 = (c_1^1, c_2^1) \in C_1 \times C_2$.

Recursively, conditional on the game not having ended, a public history at the beginning of period $t = 1, 2, \dots$, denoted by $h^t = (c^1, \dots, c^{t-1})$, records how the game has

⁵ This assumption could be relaxed to requiring that the probability of having a single candidate in common is sufficiently high. If we only impose that the truth is in each possibility set, then with high probability it is the unique element in the intersection if the state space is large and players are well informed: Fix ν_i and let S_i be drawn from a uniform distribution over subsets of Ω of size ν_i that contain ω^* . For any $\omega \neq \omega^*$, $\text{Prob}(\omega \in S_i) \rightarrow 0$ as $|\Omega| \rightarrow \infty$, and $\omega \in S_1$ and $\omega \in S_2$ are independent events. Hence, $\text{Prob}(\omega \in S_1 \cap S_2) \rightarrow 0$.

⁶ We have in mind various situations where the truth is worth searching for until it is discovered by someone, at which point it becomes common knowledge and the search is over for everyone. One player having taken a wrong action is no reason to abandon the gathering of information. In fact, the missteps of one player may provide exactly the information needed to discover the truth.

⁷ The restriction of player i ’s set of actions to S_i , rather than allowing the entire state space, is for convenience and without loss of generality because taking an action outside of S_i is strictly dominated due to assumption (1) below.

been played prior to period t . For completeness, define $h^1 = \emptyset$. Player i 's private history $h_i^t = (S_i, h^t)$ combines the public history with player i 's private information about his possibility set S_i . Let $\mathbb{D}_j^t$ denote the set of all elements disclosed by player $j \in \{1, 2\}$ according to h^t , and let $\mathbb{A}_j^t$ denote the set of actions taken by j . Then, given any private history $h_i^t = (S_i, h^t)$, player i 's information set is given by

$$I(h_i^t) := \{(R_1, R_2, h^t) | R_i = S_i, R_{-i} \supset (\mathbb{D}_{-i}^t \cup \mathbb{A}_{-i}^t), \#(R_{-i}) = \nu_{-i}, \#(R_1 \cap R_2) = 1\}.$$

Any information set of player i is a maximal set of histories that player i cannot distinguish by what he has learned during the course of the game. A history h_i^t is an *extension* of h_i^τ if the two coincide prior to period τ ; it is a *simple extension* if no move took place from period τ onward.

The set of possible choices for player i in period t with a private history $h_i^t = (S_i, h^t)$ is $C(h_i^t) = 2^{S(h_i^t)} \cup S_i$, where $S(h_i^t) = S_i \setminus (\mathbb{D}_i^t \cup \mathbb{A}_i^t)$ is the subset of S_i that consists of the elements that i has not yet disclosed or taken as an action. The outcome of period t is $c^t = (c_1^t, c_2^t) \in C(h_1^t) \times C(h_2^t)$.

If player i alone takes an action in period t and that action is ω^* , then his payoff is $\alpha > 0$ in that period while his opponent, denoted by $-i$, receives a payoff $\beta < 0$; if both players take action ω^* in the same period, each receives the payoff $(\alpha + \beta)/2$; if player i takes action $\omega \neq \omega^*$, then i 's payoff equals $\gamma < 0$ and $-i$'s payoff is zero.⁸ In any period t in which no action is taken, players receive a payoff of zero. Players have a common discount factor, $\delta \in (0, 1)$, and maximize the expected presented discounted sum of per-period payoffs.

Taking the correct action ω^* is socially desirable, $\alpha + \beta > 0$, even if costly to the player who is not the one taking it, $\beta < 0$. Taking an incorrect action, $\omega \neq \omega^*$, is worse than being preempted, $\gamma < \beta$, and so much so that a player would reject an equal probability chance of taking the correct or an incorrect action, $\alpha + \gamma < 0$. Throughout of the paper, we assume the stronger condition

$$\frac{2\alpha + \gamma}{2} < \beta < 0 < \alpha + 2\beta. \quad (1)$$

The first inequality ensures that a player would reject an equal probability chance of taking the correct or an incorrect action, even if it guaranteed him identifying the correct action in the immediate following period; as a consequence, no player would want to try to preempt his opponent by taking an action when his posterior is uniform over a nonsingleton set of states. The last inequality implies that each player i prefers that the true state becomes known provided that his chance of taking the correct action before player $-i$ is at least one-third.⁹

We now define a player's strategy by specifying a choice for every possible private history.¹⁰ Specifically, a *planned choice* of player i in period t with private history h_i^t is

⁸Although set at 0 for expositional ease, this payoff is unimportant for our result because no player would take an action $\omega \neq \omega^*$ in equilibrium due to the assumption (1) below.

⁹It is used to ensure that there are nontrivial equilibria for general (ν_1, ν_2) .

¹⁰Note that each player's possibility set S_i can be perceived as his private type and thus, each player's strategy may be described as type-contingent choice for every possible public history. The current approach is equivalent to this.

$\sigma_i(h_i^t) \in \Delta C(h_i^t)$, where ΔX is the set of all probability distributions over the set X ; it is a *pure* planned choice if $\sigma_i(h_i^t)$ assigns probability 1 to a single element of $C(h_i^t)$. A *strategy* of player i , denoted by σ_i , is a collection of planned choices, one for each and every possible history (of any length). Given a strategy σ_i and a history h_i^t , a “continuation strategy” of player i is σ_i restricted to h_i^t and all possible extensions of it. Note that, given S_i , only those histories are possible according to which player i neither discloses elements outside of S_i nor takes them as actions, which we take for granted.

We assume, in the spirit of Crawford and Haller (1990), that there is no common labeling of the elements of Ω . As a result, from player $-i$'s perspective, player i 's behavior treats elements of Ω identically as long as they are not distinguished by the history of play. Here, ω and ω' are undistinguished by history h_i^t if $\omega \in S_i \Leftrightarrow \omega' \in S_i$; $\omega \in D_j^\tau \Leftrightarrow \omega' \in D_j^\tau$ for all $\tau < t$ and $j = 1, 2$, where D_j^τ denotes the set of states disclosed by player j in period τ ; and $\omega, \omega' \notin \mathbb{A}_j^t$, $j = 1, 2$. Formally, adopting the perspective of player $-i$, this means that for each player i , his strategy σ_i is invariant under permutations of the elements of the state space. Denote a permutation of the state space Ω by π and the set of all such permutations by Π . For every private history h_i^t , let $\pi(h_i^t)$ stand for the private history obtained by renaming the elements of h_i^t according to the permutation π , and for every choice $c \in C(h_i^t)$, let $\pi(c)$ stand for the choice obtained by renaming the elements of c according to π . Then we have *no common labeling* (NCL) if

$$\sigma_i(h_i^t)(c) = \sigma_i(\pi(h_i^t))(\pi(c)) \quad \forall c \in C(h_i^t)$$

for all $h_i^t = (S_i, h^t)$ and $\pi \in \Pi$. On the equilibrium path, NCL implies that players assign equal probability to states that have not been distinguished by history. We maintain that property of beliefs also off the equilibrium path.¹¹ We take it for granted throughout the paper that all strategies and beliefs satisfy this property.

A perfect Bayesian equilibrium is a strategy–belief pair (σ, μ) consisting of a strategy profile $\sigma = (\sigma_1, \sigma_2)$ and a belief system μ that assigns a belief to every information set, with the property that strategies are sequentially rational given beliefs and beliefs are derived from Bayes' rule where possible. We strengthen the latter requirement by requiring that for any beliefs given some history, beliefs for continuation histories are also derived from Bayes' rule where possible, as in Fudenberg and Tirole (1991).

DEFINITION 1. A strategy–belief pair (σ, μ) forms a *perfect Bayesian equilibrium* (PBE) if at every possible private history h_i^t , (i) σ_i is a best response of player i given σ_{-i} and μ , and (ii) for all possible extensions of h_i^t , the belief assigned by μ is obtained from σ by Bayes' rule based on $\mu(I(h_i^t))$, where possible.

Under the NCL assumption, whenever player i makes a disclosure, player $-i$ updates his posterior belief about the true state by dismissing the disclosed states and concentrating beliefs on the remaining states, unless he finds one of the disclosed states, say ω , in his possibility set S_{-i} ; in the latter case, $\omega = \omega^*$ and it is clearly optimal for player

¹¹This property of beliefs would be satisfied if, in the spirit of sequential equilibrium, we had beliefs be generated by trembles that have full support on NCL strategies.

— i to take action ω in the next period (as will be formalized shortly). A straightforward consequence of the NCL assumption is, therefore, that each player continues to assign the same probability of being the true state to each of the elements in his possibility set that are undistinguished by history. When player i considers a disclosure after history h_i^t (as ω^* has not been identified by then), all that matters strategically is how many elements to disclose, not their identities, since all elements of $S(h_i^t)$ are undistinguished by history.

The role of the NCL assumption is to emphasize the hard-information nature of our model: player i cannot indirectly communicate information about the elements in S_i ; all that player $-i$ learns about S_i from a disclosure D_i by player i is that $D_i \subset S_i$. This also implies that at the disclosure stage, player i 's only relevant decision concerns how many (further) elements of S_i to disclose. In the sequel, therefore, we represent a disclosure move by the number of the elements to disclose (irrespective of their identities), i.e., we can write the set of available choices following private history h_i^t as $C(h_i^t) = \{1, \dots, \#(S(h_i^t))\} \cup S_i \cup \{\emptyset\}$.

4. GENERAL PROPERTIES OF PBE

In this section, we establish some core properties that pertain to all equilibria. The overall picture that emerges from these results is that a player will take an action if and only if he is certain of it being correct, either because the other player has revealed the true state or failed to respond to a revelation of all but one state; a player will generally not reveal all of the remaining elements of his possibility set at once; equilibrium continuation payoffs are bounded from below by zero; and, as a result, there is a limit on the size of disclosed sets. Therefore, disclosure must be gradual, involves reciprocation, and the length of time required to guarantee finding the true state grows without bound, as the initial uncertainty, represented by the size of the initial possibility sets, increases.

Recall that $S(h_i^t) = S_i \setminus (\mathbb{D}_i^t \cup \mathbb{A}_i^t)$ is the subset of S_i consisting of the elements that player i has not yet disclosed or taken as an action according to h_i^t . For brevity, we use S_i^t as a shorthand for $S(h_i^t)$ and refer to it as player i 's *remaining possibility set* when no confusion arises. The following two classes of histories are of special interest:

$$H_i^*(\omega) := \{h_i^t | \{\omega\} = S_i \cap \mathbb{D}_{-i}^t\}$$

$$H_i^\circ(\omega) := \{h_i^t | \{\omega\} = S_i^t \text{ and player } i \text{ disclosed no element in period } t-1\}$$

The class $H_i^*(\omega)$ consists of all private histories of player i in which he has identified ω to be the true state, $\omega = \omega^*$, because his opponent has disclosed it as being also in his own possibility set. The class $H_i^\circ(\omega)$ consists of all private histories of player i in which he has disclosed every state in his initial possibility set with the exception of the state ω , and even though his opponent had a chance to make a move, the game has not ended; note that from this, player i can infer that ω must be the true state. The next two lemmas state that player i always takes an action ω that has been identified as the true state in one of these two manners. Furthermore, player i never takes an action otherwise. We

define

$$H_i^* := \bigcup_{\omega \in S_i} H_i^*(\omega) \quad \text{and} \quad H_i^\circ := \bigcup_{\omega \in S_i} H_i^\circ(\omega).$$

LEMMA 1. *In any PBE, $\sigma_i(h_i^t)(\langle \omega \rangle) = 1$ for all $h_i^t \in H_i^*(\omega)$, and $\sigma_i(h_i^t)(\langle \omega \rangle) = 1$ for all $h_i^t \in H_i^\circ(\omega)$ on the equilibrium path.*¹²

PROOF. It is obvious that $\sigma_i(h_i^t)(\langle \omega \rangle) = 1$ for all $h_i^t \in H_i^*(\omega)$. This implies that if $h_i^t \in H_i^\circ(\omega)$ is along the equilibrium path, then $\omega = \omega^*$, since otherwise ω^* must have been disclosed by player i and, thus, the other player must have ended the game by taking action ω^* . Hence, $\sigma_i(h_i^t)(\langle \omega \rangle) = 1$ if $h_i^t \in H_i^\circ(\omega)$ is on the equilibrium path. $\square$

LEMMA 2. *In any PBE, on the equilibrium path,*

- (a) *if $\#(S_i^t) \geq 2$ and $h_i^t \notin H_i^*$, then $\sigma_i(h_i^t)(\langle \omega \rangle) = 0 \forall \omega \in S_i$*
- (b) *if $\#(S_i^t) = 1$ and $h_i^t \notin H_i^* \cup H_i^\circ$, then $\sigma_i(h_i^t)(\langle \omega \rangle) = 0 \forall \omega \in S_i$.*

PROOF. Being on the equilibrium path, ω^* has not been disclosed by player i before period $t - 1$, for if it had been, then the game would have ended according to Lemma 1. By the NCL property, therefore, any of the elements that have not been disclosed by player i before period $t - 1$ is equally likely to be ω^* . Let n be the number of such elements. Note that $n \geq 2$ because $\#(S_i^t) \geq 2$ for (a), and $\#(S_i^t) = 1$ and $h_i^t \notin H_i^\circ$ for (b). Thus, player i 's payoff in the continuation from taking any of these elements as an action in period t is bounded from above by what he obtains if he wins in the next period when the action taken is not ω^* , i.e., $[(n - 1)(\gamma + \delta\alpha) + \alpha]/n < [(n - 1)\gamma + n\alpha]/n \leq (\gamma + 2\alpha)/2$, which is less than the lower bound, β , of the payoff from doing nothing forever due to (1). Since the payoff from taking any other element of S_i as an action is bounded above by $\gamma + \delta\alpha < \beta$, we conclude that taking any element in period t as an action is suboptimal for player i . $\square$

The next result states that a player does not disclose all the remaining elements in his possibility set as long as his opponent has two or more elements undisclosed and, therefore, it remains uncertain whether his opponent has identified ω^* .

LEMMA 3. *In any PBE, on the equilibrium path, $\sigma_i(h_i^t)(S_i^t) = 0$ if $\#(S_{-i}^t) \geq 2$.*

PROOF. In the contingency that player $-i$ knows ω^* , doing nothing is trivially no worse than disclosing S_i^t for player i . In the alternative contingency, which has a positive probability when $\#(S_{-i}^t) \geq 2$, let $d \in \{0, 1, \dots\}$ denote the number of elements that player $-i$ discloses in period t . By doing nothing forever unless player i knows ω^* for sure,

¹²The qualification “on the equilibrium path” is needed here because if in period $t - 1$, in which player $-i$ had a chance to respond, player $-i$ made an unexpected move, PBE permits player i to believe in period t that she did disclose ω^* at some earlier point in time, without player $-i$ having responded optimally by taking the action ω^* . In that case, player i would believe in period t that $\omega \neq \omega^*$.

player i would get an expected payoff strictly above $\delta[d\alpha + (\#(S_{-i}^t) - d)\beta]/\#(S_{-i}^t)$ because player $-i$ would never take an action by Lemma 2 unless history evolved so that it is in H_{-i}° . If player i discloses S_i^t in period t , on the other hand, player $-i$ would take action ω^* in period $t + 1$ with probability 1 and, thus, player i 's expected payoff would be $\delta[d(\alpha + \beta)/2 + (\#(S_{-i}^t) - d)\beta]/\#(S_{-i}^t) \leq \delta[d\alpha + (\#(S_{-i}^t) - d)\beta]/\#(S_{-i}^t)$. Hence, disclosing S_i^t in period t is dominated by doing nothing forever unless player i knows ω^* for sure. $\square$

The next lemma addresses a situation for a player i in period t who has disclosed all but one element of his possibility set and does not know whether the sole remaining element corresponds to the true state or not. It says that unless the other player is in the same situation, player i makes no move and instead waits to see whether the other player ends the game by taking action ω^* , anticipating that if player $-i$ does not end the game, he will be in a position to end it himself by making the optimal choice (Lemma 1).

LEMMA 4. *In any PBE, on the equilibrium path, if $\#(S_i^t) = 1 < \#(S_{-i}^t)$ and $h_i^t \notin H_i^* \cup H_i^\circ$, then $\sigma_i(h_i^t)(\emptyset) = 1$.*

PROOF. Since player i does not disclose his remaining element in period t by Lemma 3, it suffices to show that he does not take any action in period t either. If $h_{-i}^t \in H_{-i}^*$, player $-i$ will take action ω^* in period t by Lemma 1. In this contingency, it is trivially optimal for player i to do anything other than taking some action ω . In the other contingency that $h_{-i}^t \notin H_{-i}^*$, which has positive probability, player $-i$ will not take any action in period t by Lemma 2(a). If player $-i$ does not end the game in period t , therefore, player i will correctly infer that his remaining element must be ω^* .

If player i discloses the lone remaining element in his possibility set in period t , in the contingency that $h_{-i}^t \notin H_{-i}^*$, both players will take action ω^* in period $t + 1$. If player i does not disclose his lone remaining element in period t , in the same contingency, player i is certain to take action ω^* , but player $-i$ will take action ω^* only with probability strictly less than 1 given that $\#(S_{-i}^t) \geq 2$. This proves that $\sigma_i(h_i^t)(\emptyset) = 1$. $\square$

It is generally not optimal for a player to disclose his entire remaining possibility set (as shown in Lemma 3). Additional constraints on the size of disclosed sets follow from the fact that equilibrium continuation payoffs are bounded from below by zero, as shown in the next lemma. As a result, player i will avoid making large disclosures, which would result in a high probability of the opponent discovering the true state, taking an action, and leaving player i with a negative payoff, as shown in Theorem 1 below.

LEMMA 5. *In any PBE, after any private history h_i^t with $\#(S_j^t) \geq 2$ for $j = 1, 2$, player i 's expected payoff conditional on $h_{-i}^t \notin H_{-i}^*$, is no less than 0.*

Note that for this observation, a player's expectation is taken conditional on information that is not available to him. The observation is of interest because the event that player $-i$ has already discovered ω^* is irrelevant for player i 's disclosure decision.

PROOF OF LEMMA 5. Denote the PBE strategy profile by σ . Following h_i^t , consider letting player i adopt the strategy $\tilde{\sigma}_i$, instead of never disclosing any elements and taking an action if and only if that action is revealed to be ω^* . Since player $-i$'s strategy σ_{-i} , by assumption, is sequentially rational, at those private histories of player $-i$ that are consistent with player i using strategy $\tilde{\sigma}_i$, player $-i$ will never take an action unless player $-i$ has disclosed all but one of the elements in his possibility set and player i had an opportunity to respond. Therefore, given player i 's strategy $\tilde{\sigma}_i$, neither player will ever take an action and, therefore, both players receive a payoff of 0, or player $-i$ makes disclosures before taking an action. Each time player $-i$ makes a disclosure there is a chance that he discloses ω^* , in which case player i receives a positive payoff. Only in the event that player $-i$ has disclosed all but one element and player i had a chance to respond will player $-i$ take an action. At the moment player $-i$ makes the final disclosure that leaves him with one undisclosed element, player i 's expected payoff conditional on player $-i$ disclosing all but one of his remaining K elements is $\delta\alpha(K-1)/K + \delta^2\beta/K > 0$. Therefore, unless player $-i$ has already identified ω^* , player i adopting strategy $\tilde{\sigma}_i$ following private history h_i^t always results in a nonnegative payoff for player i . $\square$

We now use the observation that players' equilibrium continuation payoffs are bounded from below by 0 to show that disclosure sizes are bounded and, as a consequence, that ensuring discovery of the truth requires that the number of disclosure rounds grows without bound when players' initial information is made sufficiently imprecise.

THEOREM 1. *For any integer M , if v_1 and v_2 are sufficiently large, disclosure goes on for M or more rounds with positive probability in any PBE in which the true state is revealed with certainty.*

In the next section, we will show that there always are equilibria in which the true state is revealed with certainty, provided that players are sufficiently patient. Jointly these two results establish that with patient players, information exchange is possible, but requires protracted dialogues.

PROOF OF THEOREM 1. Consider any private history $h_i^t \notin H_i^*$ of player i after which $n_j \geq 2$, $j = 1, 2$, elements remain in players' possibility sets. With positive probability $h_{-i}^t \notin H_{-i}^*$; otherwise player i 's disclosure decision is irrelevant. Consider player i 's problem of how many elements, K_i , to disclose in period t following h_i^t . There are two possibilities for player i to consider: one is that player $-i$ contemporaneously discloses ω^* ; conditional on that event, the unique optimal choice of K_i would be 0. Hence player i only possibly discloses $K_i > 0$ elements in consideration of the possibility of being in the second case, in which player $-i$ does not disclose the true element in period t . In that case, player i 's payoff from disclosing K_i elements is bounded from above by

$$\frac{K_i}{n_i}\beta + \left(1 - \frac{K_i}{n_i}\right)\alpha.$$

Since in the (positive probability) event that player i 's disclosure decision is relevant, i.e., $h_{-i}^t \notin H_{-i}^*$, player i can guarantee a payoff no less than 0 by Lemma 5, in order for player i to be willing to disclose K_i elements, it has to be the case that

$$\frac{K_i}{n_i}\beta + \left(1 - \frac{K_i}{n_i}\right)\alpha \geq 0. \quad (2)$$

Let N be the smallest (integer) value of n such that $\alpha + n\beta < 0$. Then the condition in (2) amounts to

$$\frac{\frac{K_i}{n_i}}{1 - \frac{K_i}{n_i}} < N \iff \frac{K_i}{n_i} < \frac{N}{N+1}.$$

Therefore, neither player will ever disclose a fraction $N/(N+1)$ or more of the elements of his remaining possibility set. Hence, after M disclosure rounds, player i 's remaining possibility set contains at least $(1/(N+1))^M \nu_i$ elements, provided $(1/(N+1))^M \nu_i \geq 1$, which can be ensured by choosing ν_i sufficiently large. Choose ν so that $(1/(N+1))^M \nu \geq 2$. Then, for any $\nu_j \geq \nu$, $j = 1, 2$, as long as neither player discloses the true element in any of the first M disclosure rounds, neither player will take an action by Lemma 2. Hence, for any $\nu_j \geq \nu$, $j = 1, 2$, with positive probability there are at least M disclosure rounds. $\square$

5. ASYMPTOTICALLY EFFICIENT MARKOV EQUILIBRIA

In this section, we describe equilibria that satisfy four desiderata: they are in Markov strategies; they exhibit a natural quid pro quo pattern of information exchange; they exist for all initial sizes of players' possibility sets, provided the discount factor is large enough; and, they are efficient in the limit as players become infinitely patient. We briefly discuss inefficiencies that may arise from, in addition, imposing the condition that equilibria be symmetric.

The consideration of Markov strategies and equilibria requires an appropriate state space. As shown in the previous section, the key variables that govern agents' decisions are how many elements each player has disclosed and whether or not a player's private history satisfies $h_i^t \in H_i^* \cup H_i^\circ$. This inspires our definition of a *Markov state* of a private history h_i^t as a tuple

$$(\#(S(h_i^t)), \nu_{-i} - \#(\mathbb{D}_{-i}^t \cup \mathbb{A}_{-i}^t), \mathbf{1}(h_i^t)) \in \mathbb{N} \times \mathbb{N} \times \{0, 1\},$$

where $\mathbf{1}(h_i^t)$ is an indicator function such that

$$\mathbf{1}(h_i^t) = 1 \quad \text{if } h_i^t \in H_i^* \cup H_i^\circ \quad \text{and} \quad \mathbf{1}(h_i^t) = 0 \quad \text{if } h_i^t \notin H_i^* \cup H_i^\circ.$$

A strategy σ_i is a *Markov strategy* if $\sigma_i(h_i^t)$ depends only on the Markov state of h_i^t for every h_i^t .

DEFINITION 2. A PBE $(\sigma_1, \sigma_2, \mu)$ is a *Markov equilibrium* if σ_1 and σ_2 are Markov strategies. It is *symmetric* if $\sigma_1(h_1^t) = \sigma_2(h_2^t)$ whenever h_1^t and h_2^t have identical Markov states.

We begin by describing an equilibrium strategy profile, σ^* , that satisfies our desiderata whenever the sizes of the possibility sets do not differ by too much. What is “too much” depends on the parameters of the game. When following the profile σ^* , players eventually establish a routine of alternating disclosure of two elements each, after every step leaving a gap of 1 between the remaining possibility sets, which we refer to as the *quid pro quo pattern*. If the difference in sizes of the initial possibility sets is not too large, then initially the player with the larger possibility set bridges the gap by disclosing either the difference or the difference plus 1 depending on the identity of the player and whether the size of the smaller possibility set is odd or even. If the gap is too large, doing so is excessively costly and there is no disclosure unless the smaller information set is so small that it is beneficial for its holder to disclose all but one element for the screening motive. Existence can be guaranteed for any pair of initial possibility sets provided players are sufficiently patient. Efficiency, in the limit as the discount factor converges to 1, is achieved whenever the initial possibility sets are of similar size or one of them is small enough. In order to achieve efficiency also in cases where the difference in the sizes of the initial possibility sets is large, the strategy profile σ^* can be amended by having players exchange information in a way that gradually bridges the gap.

Here we describe the key elements of the strategy profile σ^* ; the details are in the [Appendix](#). Player i acts on the element ω whenever the history of play reveals ω as being ω^* , that is,

$$\sigma_i^*(h_i^t)(\langle \omega \rangle) = 1 \quad \text{if } h_i^t \in H_i^*(\omega) \cup H_i^o(\omega).$$

Otherwise player i decides whether and how many elements to disclose as follows. Let $n = \#(S_1^t)$ and $k = \#(S_2^t)$, and call the corresponding game an $n \times k$ game. A player with only one element in his possibility set does not disclose. If $n = k > 1$ and n is odd, then player 1 discloses one element, and player 2 discloses none; if $n = k > 1$ and n is even, the roles are reversed. This way, whenever the game starts with equal-sized possibility sets, players reach a state in which the sizes differ by 1 and it becomes possible to alternately disclose two, while maintaining the size difference, as follows: If $n = k + 1 > 2$ and n is even, then player 1 discloses two elements and player 2 discloses none; if $k = n + 1 > 2$ and k is odd, the roles are reversed. This implements the quid pro quo pattern of alternating disclosure of two elements. Note that following the pattern is tied to the players' identities. If $n = k + 1 > 2$ and n is odd, then player 1 discloses only one element, and likewise if $k = n + 1 > 2$ and k is even, player 2 discloses only one element:¹³ doing so leads to a state with an equal number of elements, prompting the opponent to disclose one as prescribed above, which then leads to a state that initializes the quid pro quo pattern.

In order to describe σ^* also for states in which the sizes of the remaining possibility sets differ by more than one element, it proves useful to recursively define the payoff $\phi(m)$ that a player receives when m elements remain in his opponent's possibility set, $m - 1$ remain in his set, and starting from his opponent, the players alternate disclosing

¹³It is for ease of exposition that we keep this particular pattern and the one above that player 1 (2) discloses one element when they have the same odd (even) number of elements. Clearly, there also exist equilibria in which the two players swap their roles.

two elements until action ω^* is taken by one player either because it was disclosed by the other player or because he disclosed all but one element and yet the other player did not end the game thereafter. Similarly, we define the payoff $\psi(m)$ that a player receives when m elements remain in his set, $m - 1$ elements remain in his opponent's set, and starting with himself, players alternate disclosing two elements until action ω^* is taken by one player as explained above. Thus,

$$\begin{aligned}\phi(3) &:= \delta\left(\frac{2}{3}\alpha + \frac{\delta}{3}\beta\right), & \psi(3) &:= \delta\left(\frac{2}{3}\beta + \frac{\delta}{3}\alpha\right) \\ \phi(m) &:= \delta\left(\frac{2}{m}\alpha + \frac{m-2}{m}\psi(m-1)\right), & m &\geq 4 \\ \psi(m) &:= \delta\left(\frac{2}{m}\beta + \frac{m-2}{m}\phi(m-1)\right), & m &\geq 4.\end{aligned}$$

Consider the case $n > k > 2$, in which player 1 has more elements left in his possibility set than player 2. If, in addition, k is odd, then player 1 has the opportunity to trigger the quid pro quo pattern by disclosing the difference plus 1, that is, $n - k + 1$, provided player 2 does not make any disclosure. A necessary condition for this to be optimal is that the chance of receiving the continuation payoff $\phi(k)$ with probability $(k - 1)/n$ compensates for the chance of receiving the payoff $\beta < 0$ with probability $(n - k + 1)/n$ when player 2 learns the identity of ω^* from the disclosure. This motivates the following (partial) specification of σ^* in this case: if $n > k > 2$ and k is odd, then $\sigma_1^*(h_1^t)(n - k + 1) = 1 = \sigma_2^*(h_2^t)(\emptyset)$ provided that the profitability condition $(n - k + 1)\beta + (k - 1)\phi(k) > 0$ holds.

If $n > k > 2$ but k is even, then player 1 can trigger the quid pro quo pattern by disclosing the difference, that is, $n - k$ elements. This is in anticipation of player 2 disclosing a single element in the following period, after which both players alternate disclosing two elements. A necessary condition for it to be optimal for player 1 to disclose $n - k$ elements in this case is that the expected payoff from receiving β with probability $(n - k)/n$, $\delta\alpha$ with probability $1/n$, and $\delta\psi(k)$ with probability $(k - 1)/n$ is positive. This motivates: If $n > k > 2$ and k is even, then $\sigma_1^*(h_1^t)(n - k) = 1 = \sigma_2^*(h_2^t)(\emptyset)$ provided that the profitability condition $(n - k)\beta + \delta(\alpha + (k - 1)\psi(k)) > 0$ holds. In this as well as the previous case, if the profitability condition does not hold, then the player with the larger possibility set stays put and the opponent either stays put as well or discloses all but one provided this is profitable according to the screening motive.

The construction of σ^* is similar for states in which player 2 has more elements remaining in his possibility set than player 1. In all of these case, if the conditions for profitable disclosure of the difference (or the difference plus 1) are met, the player with the larger possibility set discloses the difference (or the difference plus 1) and thus triggers the quid pro quo pattern. This process leads to the discovery of ω^* in finite time and, therefore, an efficient outcome in the limit as players become infinitely patient.

We briefly summarize the key elements of verifying that if both players adopt this strategy we have an equilibrium (the details are in the [Appendix](#)): When the strategy σ^* prescribes that the player with the larger possibility set discloses some number of elements, then this number is equal to the size difference plus 1 (or the size difference).

To disclose fewer elements and following the strategy thereafter would be suboptimal, because following such a disclosure, the strategy σ^* prescribes that the same player discloses again, because the order of the possibility set sizes has remained the same; because of discounting, the player with the larger possibility set prefers disclosing a given number of elements all at once rather than disclosing that same number in multiple installments. To disclose more would be suboptimal because it would result in giving away too much information too quickly.

A player who according to the strategy σ^* is designated to disclose when having the smaller possibility set is meant to disclose all but one of the elements of that set, and the expected payoff from doing so is positive. To disclose less and thereafter continuing to follow σ^* would be suboptimal because the same player would be called upon to disclose again, since the size gap would have increased, lessening the incentive of the other player to make the minimal order-reversing disclosure, and, as before, because of discounting, it is preferable to disclose all but one element all at once rather than disclosing that same number in multiple installments. Hence, if the player with the larger possibility set does not disclose, it is optimal for the player with the smaller possibility set to disclose all but one as long as the payoff from disclosing some number of elements is positive.

Having verified the optimality of disclosures stipulated by σ^* , we now turn to the optimality of nondisclosure: There are the following four cases in which a player i is supposed not to disclose.

- (i) Player i has the larger possibility set and player $-i$ does not disclose either. In this case, if instead player i disclosed such a small number of elements that he remains in this case, then the direct payoff impact of the initial disclosure is negative and there are no other consequences. If he discloses more, but short of reversing or equalizing the order, then σ^* prescribes that thereafter he makes the minimal order-reversing/equalizing disclosure. If this were profitable, then by discounting he would have been better off making the minimal order-reversing/equalizing disclosure at the outset, but a defining characteristic of the present case is that making such a disclosure is not profitable. Disclosing even more is not profitable because of the negative payoff impact of parting with more information too soon.
- (ii) Player i has the larger possibility set and player $-i$ discloses all but one of his elements. In this case, the immediate payoff consequence of disclosure is sufficiently negative to make disclosure unattractive.
- (iii) Player i has the smaller possibility set and player $-i$ does not disclose. If instead player i disclosed a sufficiently small number, then this would take us back to the same case. If player i disclosed a number of elements large enough that σ^* then would prescribe to disclose all but one, and if doing so were profitable, then by discounting, it would be even more attractive to disclose all but one immediately, yet doing so is unprofitable under the conditions of the case in question.

- (iv) Player i has the smaller possibility set and player $-i$ makes the minimal order-reversing or order-equalizing disclosure. Suppose a deviation of player i disclosing as well were profitable. Then player i would do even better by postponing that disclosure until the next period, because the resulting state would be the same and he could first take advantage of player $-i$'s disclosure, assuming players are patient enough.

If the difference in the sizes of the possibility sets is large, the profitability conditions may fail, and there may not be any further disclosure under σ^* . To establish asymptotic efficiency for all initial sizes of the possibility sets, we amend the profile σ^* . Whenever there is a large gap in the sizes of the initial possibility sets, the amended profile prescribes a gradual closing of the size gap. For example, in an 11×5 game, counting his 11 elements as 5 *units* of two elements and one residual, player 1 first discloses three elements (leaving 1 less unit undisclosed than his opponent's elements); then player 2 discloses two elements, followed by player 1 disclosing 2 units (four elements), and so on, until ω^* is identified. This gives us our desired result, which is stated below and proved in the [Appendix](#).

THEOREM 2. *For every $n \times k$ game there is a $\underline{\delta} \in (0, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a Markov equilibrium in which information disclosure takes place in every round until ω^* is identified. Consequently, the game ends in no more than $n + k + 1$ rounds and efficiency is achieved in the limit as $\delta \rightarrow 1$.*

For any (n, k) , in the equilibrium constructed for Theorem 2, disclosure starts without delay and continues until ω^* is identified, which takes place within a finite number of periods. At the same time, our earlier result (Theorem 1) implies that the maximum number of information exchanges that may take place in any such equilibrium increases without bound as n and k increase. In summary, if players are patient or, equivalently, disclosures can be made at a rapid rate, information exchange can be made efficient even though it requires protracted rounds of alternating information provision.

In addition to having equilibria satisfy the desiderata of being in Markov strategies, displaying a quid pro quo pattern of information exchange, and existence for all initial states, one might want to impose symmetry. This, together with the Markov condition, implies that players with identically sized possibility sets must behave identically. To achieve the quid pro quo pattern of information exchange, it is, therefore, necessary to break the initial symmetry of the state. This implies that players must randomize. When the initial sizes of the possibility sets are identical, players engage in a war of attrition, which is a potential source of inefficiency.¹⁴

6. CONCLUDING REMARKS

Our interest in this paper has been in understanding the interaction between rivals who compete to be the first to learn the truth, but depend on each other's information to be

¹⁴We construct such equilibria in our working paper (Blume and Park (2020)) and show that for identical odd-sized initial possibility sets, they are inefficient even when players become infinitely patient.

able to do so. We found that in such situations, information exchange is possible even if it is certain that only one party gains *ex post* while the other loses. With patient players, one can guarantee that the truth will be found. Finding the truth with certainty necessitates that there is a positive probability of multiple rounds of information exchange, with the number of rounds growing without bound as each side is made less informed.

A key driver of our analysis is the screening motive that sometimes makes players willing to disclose all but one of the elements of their possibility sets in hoping to infer that it is the true state from a failure of the other player ending the game. This argument makes use of our assumption that the two players' possibility sets have exactly one element in common, but this is not critical. The analysis continues to hold as long as the probability of there being a single common element is sufficiently high.

In the equilibrium specified in Section 5, for example, the players' equilibrium strategy specifies a uniquely optimal choice for every Markov state. When the two players' initial possibility sets may have multiple elements in common with a small enough probability, the expected payoff from each feasible choice at every Markov state is a continuous function of that probability (presuming the other player behaves according to the equilibrium strategy), thus preserving optimality.¹⁵ Efficiency is impaired because the common element disclosed by the other player may not be the true state of the world, albeit with a small probability.

One may wonder what would happen if players may not disclose elements in their possibility set S_i , but only those in the complement of S_i . In this case, no information disclosure may take place that leads to identification of the true state ω^* with a positive probability. To see this, suppose to the contrary that an equilibrium existed where ω^* is taken with a positive probability. Given a finite state space Ω , there is a finite sequence of disclosures after which ω^* is taken with a positive probability. Suppose player 1 is to disclose in the last stage of this disclosure sequence: by doing so, player 1 risks losing the game in the case that all remaining elements in $S_2 \setminus \{\omega^*\}$ are disclosed, without anything to gain even if player 2 did not take ω^* , because that fact would not eliminate any element of S_1 from being the true state. Thus, player 1 should find it suboptimal to disclose anything in the last stage of disclosure, upsetting the presumed equilibrium. That is, the screening motive is absent and the last disclosure is unsustainable when players may only disclose elements outside of their possibility sets, precluding any beneficial information exchange.

We have considered disclosures backed by "hard" evidence. One might also wonder what would happen if the players communicated by cheap talk. In that case, the no common labeling assumption needs to be abandoned in order for there to be a language for communication. Then, when a player has two or three elements in his possibility set, the screening motive continues to incentivize the agent to truthfully disclose all but one element, because the expected payoff from doing so exceeds that from falsely disclosing

¹⁵Additional Markov states arise when a player discloses multiple elements that are also in the remaining possibility set of the other player. In such states, it is optimal for the latter player to do nothing in all future periods. Arising with such a small probability, these additional states do not affect optimality of the equilibrium strategy.

an element outside his possibility set. In particular, the strategy profile described in Section 5 remains an equilibrium when both players have two elements in their possibility sets. But incentives do not work the same way for disclosures that take place before the last disclosure. When both players have three elements, for instance, the strategy profile in Section 5 ceases to be an equilibrium because, conditional on the opponent behaving according to it, player 1 would benefit by falsely disclosing first an element outside of his possibility set, as it would not expose him to an immediate risk of losing yet induces the opponent to disclose two elements truthfully, increasing his expected payoff. By the same token, any longer sequence of truthful disclosures by cheap talk is unsustainable. The question of when cheap talk is enough to induce information exchange between rivals, we leave for future work.

APPENDIX

A.1 Proofs for Section 2 (Example)

Uncertainty about joint information Suppose that $\Omega = \{\omega, \omega', \omega''\}$, $\#(S_i) = 2$ for $i = 1, 2$, ω^* is drawn from a uniform distribution on Ω , and each S_i is independently drawn from a uniform distribution over all S_i that contain ω^* (thus, $S_1 = S_2 \ni \omega^*$ is possible).

A. An equilibrium in which ω^* is taken with certainty

Consider a strategy profile that prescribes the following behavior: Player 1 discloses one element, ω_1 , and player 2 does nothing in period 1; in period 2, player 1 does nothing and player 2 acts on ω_1 if $\omega_1 \in S_2$, but does nothing if $\omega_1 \notin S_2$; in period 3, both players act on the remaining element if player 2 acted on ω_1 , which turned out to be wrong, while player 1 acts on $\omega^* \neq \omega_1$ if player 2 did nothing in period 2.

The strategy specified for period 3 is clearly optimal. Below we verify optimality of the strategies specified for periods 1 and 2 for an open set of parameter values.

Each S_i consists of ω^* and one of the other two elements of Ω with equal probability. Since S_1 and S_2 are drawn independently, it is routinely verified that

$$\omega_1 \in S_2 \quad \text{with probability } 3/4 \text{ and } \omega_1 \notin S_2 \text{ with probability } 1/4.$$

Thus, the probabilities of possible events in period 1 are

$$\begin{aligned} (\text{with prob } 3/4) \quad \omega_1 \in S_2 &\Rightarrow \begin{cases} S_1 = S_2 & \text{with probability } 2/3 \\ S_1 \cap S_2 = \{\omega_1 = \omega^*\} & \text{with probability } 1/3 \end{cases} \\ (\text{with prob } 1/4) \quad \omega_1 \notin S_2 &\Rightarrow \text{either element of } S_2 \text{ is } \omega^* \text{ with equal probability.} \end{aligned}$$

Therefore, player 2's payoffs from choices in period 2 are

$$(\text{with prob } 3/4) \quad \omega_1 \in S_2 \Rightarrow \begin{cases} \frac{2\alpha + \gamma + \delta \frac{\alpha + \beta}{2}}{3} & \text{from acting on } \omega_1 \\ \frac{\beta + \delta(\alpha + \beta)}{3} & \text{from doing nothing} \end{cases}$$

$$(\text{with prob } 1/4) \quad \omega_1 \notin S_2 \Rightarrow \begin{cases} \frac{\alpha + \gamma + \delta \frac{\alpha + \beta}{2}}{2} & \text{from acting on an element of } S_2 \\ \delta\beta & \text{from doing nothing.} \end{cases}$$

For optimality of player 2's strategy in period 2, we need

$$\frac{2\alpha + \gamma + \delta \frac{\alpha + \beta}{2}}{3} \geq \frac{\beta + \delta(\alpha + \beta)}{3} \quad \text{and} \quad \frac{\alpha + \gamma + \delta \frac{\alpha + \beta}{2}}{2} \leq \delta\beta. \quad (3)$$

Given player 2's strategy in period 2, it is clearly optimal for player 1 to do nothing in period 2 because he has no additional information at that point and is engaging in screening through player 2's response.

If $\delta = 1$, the two inequalities of (3) hold as equalities when $\gamma = 3(\beta - \alpha)/2$. For $\delta < 1$, they get relaxed at $\gamma = 3(\beta - \alpha)/2$. Therefore, the period 2 strategies are optimal for γ in a small neighborhood of $3(\beta - \alpha)/2$.

Turning to period 1, we verify optimality of player 2 doing nothing for γ near $3(\beta - \alpha)/2$. The only deviation to consider is for her to disclose one element herself as well, say ω_2 . In this case, as S_1 and S_2 are drawn independently, we deduce

$$(\text{with prob } 3/8) \quad \omega_1 = \omega_2 \quad \begin{cases} = \omega^* & \text{with probability } 2/3 \\ \neq \omega^* & \text{with probability } 1/3 \end{cases}$$

$$(\text{with prob } 5/8) \quad \omega_1 \neq \omega_2 \quad \& \quad \begin{cases} \omega^* = \omega_1 \in S_2 & \text{with probability } 2/5 \\ \omega^* = \omega_2 \in S_1 & \text{with probability } 2/5 \\ S_2 \not\ni \omega_1 \neq \omega^* \neq \omega_2 \notin S_1 & \text{with probability } 1/5. \end{cases}$$

In the contingency that $\omega_1 = \omega_2$, the continuation equilibrium is as follows: in period 2, player 2 acts on $\omega_1 = \omega_2$ and player 1 does nothing; in period 3, both players act on ω^* if player 2 acted on $\omega_1 = \omega_2 \neq \omega^*$. The optimality in period 3 is clear. To see optimality in period 2, given that player 1 does nothing, observe that player 2 gets a payoff of $[2\alpha + \gamma + \delta(\alpha + \beta)/2]/3$ from acting on $\omega_1 = \omega_2$, a lower payoff of $[\alpha + 2(\gamma + \delta(\alpha + \beta)/2)]/3$ from acting on the other element, and a delayed payoff of the former from doing nothing. Also, given that player 2 acts on $\omega_1 = \omega_2$, player 1 gets a payoff of $[2\beta + \delta(\alpha + \beta)/2]/3$ from doing nothing, a lower payoff of $[\alpha + \beta + \gamma + \delta(\alpha + \beta)/2]/3$ from acting on $\omega_1 = \omega_2$ as well, and an even lower payoff of $(\alpha + 2\gamma)/3$ from acting on the other element.

In the contingency that $\omega_1 \neq \omega_2$, consider the following behavior: in period 2, player 2 acts on ω_1 if $\omega_1 \in S_2$, but does nothing if $\omega_1 \notin S_2$, and player 1 does nothing; in period 3, both players act on ω^* if player 2 acted wrongly previously, and player 1 acts on his undisclosed element if player 2 did not act previously. The optimality in period 3 is clear. To see optimality in period 2, observe that conditional on $\omega_1 \neq \omega_2$, we have $\omega^* = \omega_i \in S_{-i}$ with probability $2/5$ for each i as stated above, but $\omega^* \neq \omega_i \in S_{-i}$ with probability $1/5$ and in this latter case, we have $\omega^* = \omega_{-i} \in S_i$. Therefore, observing $\omega_1 \in S_2$, player 2 gets a payoff of $[2\alpha + \gamma + \delta(\alpha + \beta)/2]/3$ from acting on ω_1 , but lower payoffs of $\delta(2\alpha + \gamma)/3$ and $\delta[\beta + \delta(\alpha + \beta)]/3$, respectively, from doing nothing now and acting on ω_1 in the next

period, and from doing nothing now and in the next period. Observing $\omega_1 \notin S_2$, player 2 gets a payoff of $\delta\beta$ from doing nothing and strictly less than $[\alpha + \gamma + \delta(\alpha + \beta)/2]/2 < \delta\beta$ by acting on one element. Moving to optimality of player 1, observing $\omega_2 \in S_1$, player 1 gets a payoff of $(2\alpha + \gamma)/3$ from acting on ω_2 , but a larger payoff of $\delta[\beta + \delta(\alpha + (\alpha + \beta)/2)]/3$ from doing nothing now (and acting correctly next period if player 2 either acts wrongly or does not act in this period). Observing $\omega_2 \notin S_1$, player 1 gets a payoff of $\delta(\beta + \delta\alpha)/2$ from doing nothing and strictly less than $(\alpha + \gamma + \delta\alpha)/2$ by acting on one element.

Combining the two contingencies above, player 2 gets a payoff of $[2\alpha + \gamma + \delta(\alpha + \beta)/2]/3$ from acting on ω_1 if $\omega_1 \in S_2$ and gets a payoff of $\delta\beta$ from doing nothing if $\omega_1 \notin S_2$. This is the same as what she gets by doing nothing in period 1, which, therefore, is optimal for player 2.

Finally, postulating that both players would expect the original equilibrium to prevail from next period on in case nothing happened until some period (off equilibrium), it is optimal for player 1 to disclose one element in period 1 if his expected payoff is positive, which is indeed the case,

$$\frac{2\beta + \delta \frac{\alpha + \beta}{2}}{4} + \delta \frac{\alpha}{4} \rightarrow \frac{3\alpha + 5\beta}{8} > 0 \quad \text{as } \delta \rightarrow 1,$$

where the inequality follows from the last inequality of assumption (1) in the main text.

Therefore, for $\delta < 1$ large enough, the strategy profile above is an equilibrium if γ is in a small neighborhood of $3(\beta - \alpha)/2$. The key is that the screening works (albeit crudely).¹⁶

B. No action is ever taken in any equilibrium if $\gamma = -\infty$

We now assume $\gamma = -\infty$ so that a player only acts if he is certain about the identity of ω^* , and show that in this case, the probability of players identifying ω^* is nil in every equilibrium (hence, no action is ever taken, which we refer to as a “trivial” equilibrium for short). In doing so, we repeatedly make use of a technical observation about the equivalence of all equilibria with a subclass of equilibria, in which no player ever discloses all his remaining elements.

OBSERVATION 1. For every equilibrium in which a player discloses all his remaining elements with positive probability after some history on the equilibrium path, there is another (equivalent) equilibrium in which there is no history after which that player discloses all his remaining elements, and that is otherwise identical.

The intuition for this observation is simple: Disclosing all remaining elements provides the opponent with all available information. Being the first player to do so therefore entails a payoff loss, unless the opponent is already in a position to identify ω^* , in which case it does not matter.

¹⁶We note that the first inequality of assumption (1) in the main text is not satisfied, but it is not needed for this equilibrium.

PROOF OF OBSERVATION 1. Consider an equilibrium in which some player discloses all his remaining elements (i.e., either discloses both elements or the single remaining element) in some period with positive probability along the equilibrium path. Let t be the first such period and let player 1 be the player who does so in that period (allowing for the possibility that 2 does so as well).

Consider player 1's payoff from disclosing all remaining elements in period t . In the event that player 2 has identified ω^* as of the beginning of period t , player 1's payoff is β regardless of what he does (as he will not act in period t). In the alternative event that 2 has not identified ω^* , if player 1 discloses all remaining elements in period t , then 2 will either (i) identify ω^* and act on it in the next period or (ii) learn that $S_1 = S_2$ and not act ever (note that player 2 cannot trick player 1 into taking an action in equilibrium in the contingency that $S_1 = S_2$, since player 2's inaction reveals to player 1 that $S_1 = S_2$).

Player 1 understands that if (ii) occurs, then his payoff is 0 regardless of whether or not he discloses in period t . Thus, for player 1 to find it optimal to disclose all remaining elements in period t , it must be the case that in the contingency that $S_1 \neq S_2$ (and player 2 has not identified ω^* yet), player 2 would identify ω^* and act on it in the next period even if 1 did not disclose any element in period t . This is feasible only if player 2 can infer what ω^* is from the fact that player 1 did not act in period t (because 2's identification of ω^* is independent of the contents of remaining elements of player 1). This implies that player 2 will have identified ω^* at the end of period t if (and only if) $S_1 \neq S_2$ regardless of whether player 1 disclosed all remaining elements or not. This further implies that player 2 would not disclose any element in period t if there is any chance player 1 could identify ω^* from his disclosure. Therefore, it continues to be an equivalent equilibrium to re-specify player 1's strategy in period t as not disclosing any element. $\square$

First, consider the contingency that both players disclose one element simultaneously in some period t along the equilibrium path, when there has not been a prior disclosure or action. Let ω_i denote the element disclosed by player $i = 1, 2$.

- (i) If $\omega_1 = \omega_2$, both players assign positive probability to both elements of theirs being ω^* and do not act in period $t + 1$; hence, their inaction in period $t + 1$ reveals no additional information for either player to act subsequently and, by Observation 1, nothing happens afterward.
- (ii) Consider the case that $\omega_1 \neq \omega_2$. Regardless of whether $\omega_2 \in S_1$ or not, player 1 assigns positive probability to both elements of his being ω^* and, thus, does not act in period $t + 1$. By symmetry, the same holds for player 2. Hence, their inaction in period $t + 1$ reveals no additional information for either player to act subsequently and, by Observation 1, nothing happens afterward.

Next, assuming no prior disclosures or actions, suppose player 1 unilaterally discloses one element ω_1 . Then

$$\omega_1 \in S_2 \quad \text{with probability } 3/4; \quad \text{and} \quad \omega_1 \notin S_2 \quad \text{with probability } 1/4.$$

Following player 1's disclosure, nothing happens until player 2 responds by disclosing one element of his own, by Observation 1 and because neither player has the necessary information to act.

Suppose, therefore, that there is a first period, t^* , in which player 2 counter-discloses an element ω_2 with a positive probability. For that period, let p be the probability with which player 2 discloses if $\omega_1 \in S_2$ and let q be the probability with which player 2 discloses if $\omega_1 \notin S_2$. By the definition of period t^* , $p + q > 0$.

We start with the case that $\omega_2 \neq \omega_1$ even if $\omega_1 \in S_2$. Subsequently, we consider the possibility that agent 2 discloses $\omega_2 = \omega_1$ when $\omega_1 \in S_2$.

Then

$$\text{(with prob } 3/4) \quad \omega_1 \in S_2 \longrightarrow_p \omega_2 \begin{cases} \in S_1 & \text{with probability } 2/3 \\ \notin S_1 & \text{with probability } 1/3 \end{cases} \quad (4)$$

$$\text{(with prob } 1/4) \quad \omega_1 \notin S_2 \longrightarrow_q \omega_2 \begin{cases} \in S_1 \\ \notin S_1 \end{cases} \quad \text{with equal probability.} \quad (5)$$

Therefore, upon observing whether $\omega_2 \in S_1$ or not, player 1's posterior belief about ω_1 belonging to player 2's possibility set S_2 at the end of period t^* equals

$$\left(\text{with prob } \frac{4p+q}{8} \right) \quad \omega_2 \in S_1 \longrightarrow \begin{cases} \text{Prob}(\omega_1 \in S_2) = \frac{4p}{4p+q} \\ \text{Prob}(\omega_1 \notin S_2) = \frac{q}{4p+q} \end{cases} \quad (6)$$

$$\left(\text{with prob } \frac{2p+q}{8} \right) \quad \omega_2 \notin S_1 \longrightarrow \begin{cases} \text{Prob}(\omega_1 \in S_2) = \frac{2p}{2p+q} \\ \text{Prob}(\omega_1 \notin S_2) = \frac{q}{2p+q} \end{cases} \quad (7)$$

$$\left(\text{with prob } 1 - \frac{3p+q}{4} \right) \quad \text{no disclosure by 2} \longrightarrow \begin{cases} \text{Prob}(\omega_1 \in S_2) = \frac{3(1-p)}{4-3p-q} \\ \text{Prob}(\omega_1 \notin S_2) = \frac{1-q}{4-3p-q} \end{cases} \quad (8)$$

With $\gamma = -\infty$, player 1 does not take an action in period $t^* + 1$ unless he identifies ω^* with probability 1. Player 1 taking an action in period $t^* + 1$ in the event of a counter-disclosure ω_2 in period t^* requires that either $\omega_2 \in S_1$ and $\text{Prob}(\omega_1 \notin S_2) = 1$, or $\omega_2 \notin S_1$ and $\text{Prob}(\omega_1 \in S_2) = 1$ or $\text{Prob}(\omega_1 \notin S_2) = 1$. Player 1 taking an action in period $t^* + 1$ in the event that there is no counter-disclosure in period t^* requires that $p = 1$ and $q < 1$. If these conditions are not satisfied, then, by Observation 1, player 1 will neither act nor disclose in period $t^* + 1$ or in later periods before player 2 either acts or discloses.

If nondisclosure by player 2 in period t^* is a zero-probability event (i.e., $p = q = 1$), then player 1 never acts following player 2's disclosure in period t^* (as none of the conditions above holds) and, therefore, player 1's inaction in period $t^* + 1$ reveals no additional information for player 2 to act subsequently, either; moreover, by Observation 1, nothing happens afterward as well. For a nontrivial equilibrium, we consider the possibility that $p + q < 2$.

Suppose that $p = 1 > q$. Then, upon observing $\omega_1 \notin S_2$, player 2 loses for sure after not disclosing (because player 1 infers correctly that $\omega_1 \notin S_2$ and, thus, acts on his remaining element, which is ω^*); by disclosing instead (thus, inducing player 1 to believe that $\text{Prob}(\omega_1 \in S_2) > 0$), player 2 avoids losing with a positive probability (at least when $\omega_2 \in S_1$), in which case he can guarantee a payoff of at least 0 by never acting or disclosing subsequently. Hence, player 2 should disclose for sure when $\omega_1 \notin S_2$, which contradicts $1 > q$.

If $p < 1$, then nondisclosure by 2 does not lead to an immediate loss and, therefore, 2 can guarantee a payoff of 0 by never disclosing subsequently. Hence, 2's payoff from disclosing must be nonnegative. If $p = 0 < q$, then upon observing $\omega_1 \notin S_2$, player 2 loses for sure by disclosing one element, because player 1 correctly infers that $\omega_1 \notin S_2$ from 2's disclosure and acts on ω^* , contradicting $q > 0$. If $q = 0 < p < 1$, by disclosing one element ω_2 upon observing $\omega_1 \in S_2$, player 2 faces losing if $\omega_2 \notin S_1$ (because then player 1 would act on $\omega_1 = \omega^*$), whereas nothing happens if $\omega_2 \in S_1$ because then player 1 does not act and from this player 2 infers that $\omega_2 \in S_1$ and, thus, that $S_1 = S_2$. This contradicts $p > 0$ because player 2 can guarantee a payoff of 0 by not disclosing (as noted above). Finally, if $p, q > 0$, then player 1 never acts after 2's disclosure and this inaction reveals no further information to player 2, and nothing happens afterward (which would constitute a trivial equilibrium).

We now consider the possibility that agent 2 discloses $\omega_2 = \omega_1$ when $\omega_1 \in S_2$. In this case, nothing happens subsequently for the same reasoning that we used in (i) above where we considered the case that both agents disclose one element each simultaneously and, therefore, both get a payoff of 0.

Denoting by $r > 0$ the probability with which agent 2 discloses $\omega_2 = \omega_1$ when $\omega_1 \in S_2$, the probabilities in (4) remain valid subject to $p < 1$, but together with the additional contingency that agent 2 discloses $\omega_2 = \omega_1$ with probability r in which case both get a payoff of 0. The probabilities (5)–(7) remain intact, and (8) changes to

$$\left(\text{with prob } 1 - \frac{3(p+r)+q}{4} \right)$$

$$\text{no disclosure by 2} \longrightarrow \begin{cases} \text{Prob}(\omega_1 \in S_2) = \frac{3(1-(p+r))}{4-3(p+r)-q} \\ \text{Prob}(\omega_1 \notin S_2) = \frac{1-q}{4-3(p+r)-q}. \end{cases}$$

- (a) If $r = 1$ in equilibrium, then 2 loses for sure if $\omega_1 \notin S_2$ because agent 1, inferring correctly that $\omega_1 \notin S_2$ whenever 2 fails to disclose ω_1 , would act on $\omega^* \in S_1 \setminus \{\omega_1\}$ next period. But then, when $\omega_1 \in S_2$, agent 2 can induce agent 1 act on $S_1 \setminus \{\omega_1\}$ by deviating (disclose $\omega_2 \neq \omega_1$ or none). From (4), agent 2 has 1/3 chance of losing in the next period, but ties in the subsequent period otherwise, generating an expected payoff of $(1/3) \times (\alpha + 2\beta)/3 > 0$ as $\delta \rightarrow 1$, which is better than the payoff of 0 that 2 gets by disclosing ω_1 , a contradiction.
- (b) If $r \in (0, 1)$ and $p = 0$, then since player 2 can get 0 by no disclosure but would lose by disclosing when $\omega_1 \notin S_2$ (for the same reason as above), we must have $q = 0$.

That is, player 2 either discloses $\omega_2 = \omega_1$ or none; both of these cases would lead to a trivial equilibrium, as nothing further happens.

- (c) If $r \in (0, 1)$ and $p > 0$, then player 2 must get 0 from disclosing $\omega_2 \neq \omega_1$ when $\omega_1 \in S_2$, which is possible only if $q > 0$ according to (4) and (5). This means that nothing happens subsequent to agent 2 disclosing $\omega_2 \neq \omega_1$ (as well as $\omega_2 = \omega_1$). Consequently, nothing should happen after no disclosure by player 2 if that happens in equilibrium, again leading to a trivial equilibrium. This completes the proof.

A.2 Proof of Theorem 2

In the first stage, we prove that the Markov strategy profile σ^* described in Section 5 is indeed a PBE. In the second stage, we modify σ^* so that, regardless of the sizes of the possibility sets, ω^* is necessarily taken in a finite number of periods and efficiency achieved asymptotically as $\delta \rightarrow 1$.

Stage 1: σ^* is a PBE

We start by giving a complete description of the strategy–belief pair (σ^*, μ^*) using the functions ϕ and ψ :

$$\sigma_i^*(h_i^t)(\langle \omega \rangle) = 1 \quad \text{if } h_i^t \in H_i^*(\omega) \cup H_i^\circ(\omega);$$

and for all other $h_i^t \in H_i$, denoting $n = \#(S_1^t)$ and $k = \#(S_2^t)$,

- (i) $\sigma_1^*(h_1^t)(\emptyset) = 1 \quad \text{if } k = 1$
 $\sigma_2^*(h_2^t)(\emptyset) = 1 \quad \text{if } n = 1$
- (ii) $\sigma_1^*(h_1^t)(1) = 1 = \sigma_2^*(h_2^t)(\emptyset) \quad \text{if } n = k \geq 2 \text{ and } n \text{ is odd}$
 $\sigma_2^*(h_2^t)(1) = 1 = \sigma_1^*(h_1^t)(\emptyset) \quad \text{if } n = k \geq 2 \text{ and } n \text{ is even}$
- (iii) $\sigma_1^*(h_1^t)(2) = 1 = \sigma_2^*(h_2^t)(\emptyset) \quad \text{if } n = k + 1 > 2 \text{ and } n \text{ is even}$
 $\sigma_2^*(h_2^t)(2) = 1 = \sigma_1^*(h_1^t)(\emptyset) \quad \text{if } n + 1 = k > 2 \text{ and } k \text{ is odd}$
- (iv) $\sigma_1^*(h_1^t)(n - 2) = 1 = \sigma_2^*(h_2^t)(\emptyset) \quad \text{if } n > k = 2 \text{ and } (n - 2)\beta + \delta \frac{\alpha + \delta\beta}{2} > 0$
 $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(1) \quad \text{if } n > k = 2 \text{ and } (n - 2)\beta + \delta \frac{\alpha + \delta\beta}{2} \leq 0$
- (v) $\sigma_2^*(h_2^t)(k - 1) = 1 = \sigma_1^*(h_1^t)(\emptyset) \quad \text{if } n = 2 < k \text{ and } (k - 1)\beta + \delta\alpha > 0$
 $\sigma_2^*(h_2^t)(\emptyset) = 1 = \sigma_1^*(h_1^t)(1) \quad \text{if } n = 2 < k \text{ and } (k - 1)\beta + \delta\alpha \leq 0$
- (vi) if $n > k > 2$ and k is odd,
 $\sigma_1^*(h_1^t)(n - k + 1) = 1 = \sigma_2^*(h_2^t)(\emptyset) \quad \text{if } (n - k + 1)\beta + (k - 1)\phi(k) > 0$
 $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(k - 1)$

- if $(n - k + 1)\beta + (k - 1)\phi(k) \leq 0$ & $(k - 1)\beta + \delta\alpha > 0$
- $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(\emptyset)$
- if $(n - k + 1)\beta + (k - 1)\phi(k) \leq 0$ & $(k - 1)\beta + \delta\alpha \leq 0$
- (vii) if $n > k > 2$ and k is even,
- $\sigma_1^*(h_1^t)(n - k) = 1 = \sigma_2^*(h_2^t)(\emptyset)$ if $(n - k)\beta + \delta(\alpha + (k - 1)\psi(k)) > 0$
- $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(k - 1)$ if $(n - k)\beta + \delta(\alpha + (k - 1)\psi(k)) \leq 0$
- and $(k - 1)\beta + \delta\alpha > 0$
- $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(\emptyset)$ if $(n - k)\beta + \delta(\alpha + (k - 1)\psi(k)) \leq 0$
- and $(k - 1)\beta + \delta\alpha \leq 0$
- (x) if $2 < n < k$ and n is even,
- $\sigma_2^*(h_2^t)(k - n + 1) = 1 = \sigma_1^*(h_1^t)(\emptyset)$ if $(k - n + 1)\beta + (n - 1)\phi(n) > 0$
- $\sigma_2^*(h_2^t)(\emptyset) = 1 = \sigma_1^*(h_1^t)(n - 1)$
- if $(k - n + 1)\beta + (n - 1)\phi(n) \leq 0$ & $(n - 1)\beta + \delta\alpha > 0$
- $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(\emptyset)$
- if $(k - n + 1)\beta + (n - 1)\phi(n) \leq 0$ & $(n - 1)\beta + \delta\alpha \leq 0$
- (xi) if $2 < n < k$ and n is odd,
- $\sigma_2^*(h_2^t)(k - n) = 1 = \sigma_1^*(h_1^t)(\emptyset)$ if $(k - n)\beta + \delta(\alpha + (n - 1)\psi(n)) > 0$
- $\sigma_2^*(h_2^t)(\emptyset) = 1 = \sigma_1^*(h_1^t)(n - 1)$
- if $(k - n)\beta + \delta(\alpha + (n - 1)\psi(n)) \leq 0$ & $(n - 1)\beta + \delta\alpha > 0$
- $\sigma_1^*(h_1^t)(\emptyset) = 1 = \sigma_2^*(h_2^t)(\emptyset)$
- if $(k - n)\beta + \delta(\alpha + (n - 1)\psi(n)) \leq 0$ & $(n - 1)\beta + \delta\alpha \leq 0$.

The belief profile μ^* specifies for each h_i^t a posterior belief (i) that the true state is ω with certainty if $h_i^t \in H_i^*(\omega) \cup H_i^c(\omega)$, and (ii) that each element in the remaining possibility set $S(h_i^t)$ is equally likely to be the true state otherwise. It is clear from σ^* and NCL assumption that μ^* conforms to Bayes rule whenever possible. To facilitate exposition, we represent the PBE by σ^* only in the sequel, taking it for granted that it is accompanied by μ^* specified above.

By an $N \times N$ game, we refer to the game that starts with possibility sets of cardinality $\#(S_1) = \#(S_2) = N$. We verified in Section 2 that σ^* is a PBE for the 2×2 game. The proof proceeds by induction on N , which will establish that σ^* is a PBE for $N \times N$ games for all $N \geq 2$. Then, it follows that σ^* is a PBE for any game that starts with any cardinality $\#(S_1) \geq 2$ and $\#(S_2) \geq 2$, because it is a continuation game of $N \times N$ game where $N = \max\{\#(S_1), \#(S_2)\}$, completing the proof.

We now proceed with the induction on N for $N \times N$ games. As an induction hypothesis, we assume that σ^* is a PBE for $N \times N$ games for some $N \geq 2$. Then we verify below

that σ^* is a PBE for $(N + 1) \times (N + 1)$ games as well. We present the verification for the case that $N + 1$ is an odd number. The other case is verified analogously. By the induction hypothesis, it suffices to verify optimality of σ^* in continuation games in which at least one of the players starts with $N + 1$ elements, as is done below.

Step 1: Continuation games with $n = \#(S'_1) < N + 1 = k = \#(S'_2)$. First, consider the case that player 2 is supposed to disclose either $d = k - n$ or $d = k - n + 1$ elements according to σ^* and player 1 none. Clearly, it is suboptimal for player 2 to disclose less because then she should disclose the remaining number of elements in the next period according to σ^* (which is optimal by the induction hypothesis), which is worse than disclosing all in one instalment due to discounting. Consider player 2 disclosing more, say d' elements where $d < d' < k$. (i) If player 1 would disclose none in the next period according to σ^* , then either player 2 discloses all but one or none, in either case of which player 2's payoff would be negative. (ii) If $k - d'$ is odd and player 1 would disclose $n - k + d' + 1$ elements in the next period according to σ^* , the same situation would be reached if player 2 disclosed d and followed σ^* (unless the game ended before then), which is better for player 2 because her risk of losing from disclosing ω^* in the first d' elements with probability $\frac{d'}{k}$ would be spread over multiple periods interspersed with her probability of winning from player 1's disclosure of $n - k + d' + 1$ elements. (iii) If $k - d'$ is even and player 1 would disclose $n - k + d'$ elements in the next period according to σ^* , the same situation would be reached if player 2 disclosed d and followed σ^* until $\#(S'_1) = k - d'$ and $\#(S'_2) = k - d' + 1$, and then disclosed one element (unless the game ended before then), which is better for player 2 by the same reasoning as above; however, player 2 disclosing d and following σ^* all the way is even better by the induction hypothesis. This verifies player 2 disclosing according to σ^* is optimal in the considered continuation game.

To check player 1's optimality of doing nothing, consider him disclosing some, say $\ell > 0$, elements in the first period of the continuation game, so that the players will start the next period with $\#(S'_1) = n - \ell$ and $\#(S'_2) = k - d$. Note that the same situation would have been reached if player 1 disclosed none in the first period, then disclosed ℓ elements in the next period. In either case, the subsequent payoffs are identical for the two agents, denoted by $v_i \geq 0$ for player $i = 1, 2$. Thus, the expected payoff of player 1 from disclosing ℓ elements in the first period is $(d/k)(1 - \ell/n)\alpha + (1 - d/k)(\ell/n)\beta + (d/k)(\ell/n)(\alpha + \beta)/2 + \delta(1 - d/k)(1 - \ell/n)v_1$ and that from disclosing ℓ in the next period is $(d/k)\alpha + \delta(1 - d/k)(\ell/n)\beta + \delta^2(1 - d/k)(1 - \ell/n)v_1$ which is clearly higher and positive for large δ . This verifies player 1 disclosing none (as σ^* prescribes) is optimal in the considered continuation game.

Second, consider the case that player 1 is supposed to disclose $n - 1$ elements and player 2 none according to σ^* . For player 1 disclosing less is suboptimal because he should disclose the remaining number of elements in the next period. Given that player 1 discloses $n - 1$ elements, it is clearly optimal for player 2 to disclose none.

Last, consider the case that both players are supposed to do nothing according to σ^* . The optimality in this case is clearly spelled out in Section 5.

Step 2: Continuation games with $n = N + 1 > k$. Optimality of σ^* in these continuation games is verified by an argument analogous to that used in Step 1; hence, it is omitted here.

Step 3: The game with $n = k = N + 1$. Recall that σ^* prescribes that player 1 discloses one element and player 2 none as $n = k$ is odd, generating payoffs $v_1^* = \delta\beta/n + \delta^2(1 - 1/n)\phi(n) > 0$ for player 1 and $v_2^* = \delta\alpha/n + \delta^2(1 - 1/n)\psi(n) > 0$ for player 2. Note that $v_1^* + v_2^* \rightarrow (\alpha + \beta)/2$ as $\delta \rightarrow 1$.

For player 1, disclosing none is clearly suboptimal due to discounting. Consider player 1 disclosing more, say $d > 1$ elements. There are three possibilities in the next period as follows.

- (i) d is odd and player 2 is supposed to disclose $d + 1$ elements according to σ^* , so that the subsequent continuation game starts with $\#(S_1^t) = n - d$ and $\#(S_2^t) = n - d - 1$. The same continuation game could have been reached if player 1 disclosed one element in the first period and followed σ^* (unless the game ended before then), which would be better for player 1 because his risk of losing from disclosing ω^* in the first d elements with probability d/n would be spread over multiple periods interspersed with his probability of winning from player 2's disclosure of $d + 1$ elements.
- (ii) d is even and player 2 is supposed to disclose d elements according to σ^* , so that the subsequent continuation game starts with $\#(S_1^t) = n - d = \#(S_2^t)$. The same continuation game could have been reached if player 1 disclosed one element in the first period and followed σ^* until $\#(S_1^t) = n - d + 1$ and $\#(S_2^t) = n - d$, and then disclosed one element (unless the game ended before then), which would be better for player 1 by the same reasoning as above; however, player 1 disclosing one element and following σ^* all the way is even better by the induction hypothesis.
- (iii) d is large enough that player 2 is supposed to disclose no element in the next period according to σ^* . Then player 1 is supposed to disclose either all but one element or no element. In either case, player 1's expected payoff is negative and worse than that if he followed σ^* , i.e., $v_1^* > 0$.

Thus, we have shown that disclosing one element is uniquely optimal for player 1 (as σ^* stipulates) in the game starting with $n = k = N + 1$ when δ is large enough. To show that disclosing none is optimal for player 2, consider player 2 disclosing some $\ell > 0$ elements as well, so that in the next period a continuation game starts with $\#(S_1^t) = N$ and $\#(S_2^t) = N - \ell \leq N$. By the induction hypothesis, σ^* is optimal in this continuation game, generating the continuation payoffs, say $v_i \geq 0$ for player $i = 1, 2$. Thus, the expected payoff of player 2 from disclosing ℓ elements in the first period is $(1/n)(1 - \ell/n)\alpha + (1 - 1/n)(\ell/n)\beta + (1/n)(\ell/n)(\alpha + \beta)/2 + \delta(1 - 1/n)(1 - \ell/n)v_2$. If player 2 disclosed none in the first period and disclosed ℓ elements in the next period, hence subsequently starting the continuation game with $\#(S_1^t) = N$ and $\#(S_2^t) = N - \ell \leq N$, her expected payoff is $\alpha/n + \delta(1 - 1/n)(\ell/n)\beta + \delta^2(1 - 1/n)(1 - \ell/n)v_2$, which is clearly higher and positive for large δ . This verifies player 2 disclosing none according to σ^* being optimal in the game starting with $n = k = N + 1$.

Steps 1–3 above establish the induction argument that σ^* is a PBE in the $(N + 1) \times (N + 1)$ game as well. Clearly, the equilibrium payoffs of players 1 and 2 in the 2×2 game

are, respectively,

$$v_1^* = \delta \frac{\alpha + \delta\beta}{2} \quad \text{and} \quad v_2^* = \delta \frac{\beta + \delta\alpha}{2},$$

both of which converge to $(\alpha + \beta)/2$ as $\delta \rightarrow 1$. In the 3×3 game, the limit equilibrium payoffs of players as $\delta \rightarrow 1$ are the same as those when player 1 discloses one in the first period, player 2 discloses one in the next period, and then they follow σ^* in the 2×2 game, i.e.,

$$v_1^* = \frac{\beta + 2\left(\alpha + 2\frac{\alpha + \beta}{2}\right)/3}{3} = \frac{4\alpha + 5\beta}{9} \quad \text{and} \quad v_2^* = \frac{\alpha + 2\left(\beta + 2\frac{\alpha + \beta}{2}\right)/3}{3} = \frac{5\alpha + 4\beta}{9}.$$

Recursively, the limit equilibrium payoffs of players as $\delta \rightarrow 1$ in $N \times N$ game are calculated as follows: if N is even,

$$\begin{aligned} v_1^* &= \frac{\beta + (N-1)\left(\alpha + (N-1)\frac{((N-1)^2 + 1)\alpha + ((N-1)^2 - 1)\beta}{2(N-1)^2}\right)/N}{N} \\ &= \frac{\alpha + \beta}{2} = v_2^* \end{aligned} \quad (9)$$

$$\text{if } N \text{ is odd, } \left\{ \begin{aligned} v_1^* &= \frac{\beta + (N-1)\left(\alpha + (N-1)\frac{\alpha + \beta}{2}\right)/N}{\frac{(N^2 - 1)\alpha + (N^2 + 1)\beta}{N^2}} \\ v_2^* &= \frac{\alpha + (N-1)\left(\beta + (N-1)\frac{\alpha + \beta}{2}\right)/N}{\frac{(N^2 + 1)\alpha + (N^2 - 1)\beta}{N^2}}. \end{aligned} \right. \quad (10)$$

In either case, $v_1^* + v_2^* = \alpha + \beta$, affirming asymptotic efficiency.

Stage 2: Modifying σ^*

Assume without loss of generality that $\#(S_1) = n \geq k = \#(S_2)$. According to σ^* , information disclosure takes place with certainty in every period until ω^* is identified provided that one of the players discloses some elements according to σ^* in period 1, in which case efficiency is achieved asymptotically as $\delta \rightarrow 1$ because ω^* is taken for sure within a finite periods (and no other action is taken). Below we modify σ^* for the remaining cases so that asymptotic efficiency prevails in those cases as well. Note from σ^* that these remaining cases satisfy $n - k \geq 2$ and $((k-1)\beta + \alpha)/k \leq 0$. As the latter inequality implies $k \geq |\alpha/\beta| + 1 > 3$, we consider $k \geq 4$ below.

The guiding principle of the modification is for the two players to get in a minimal number of steps to a stage from which σ^* is followed with player 1's disclosures counted in units, where a *unit* is some fixed number of elements. Since the exact manner in

which the players do so differs depending on the values of n and k , so does the precise description of the modified equilibrium. We describe and verify the modified equilibrium in six cases separately below.

Let m be the largest integer such that $mk \leq n$ so that $0 \leq r = n - mk < k$.

Case 1: $m \geq 1$ and k is odd, but not $m = 1$ and $r > \hat{r} := (k^2 - 3)/(2k)$. Consider the following Markov strategy along the equilibrium-path.

(*) Player 1 discloses $r + m$ elements and player 2 does nothing in period 1; then starting with player 2, the two players alternate disclosing 2 and $2m$ elements each, respectively, in alternating periods (player 2 in even periods and player 1 in odd periods) until ω^* is identified and taken as an action. Note that since k is odd, the last period of possible disclosure is period $k - 1$ in which player 2 would disclose all but one element.

Next, we described off-equilibrium strategies in (**) below based on the following principle: the players get back to the equilibrium path as quickly as possible if doing so gives a positive expected payoff, with the added incentive feature that the player who deviated by delaying the equilibrium exchange process bears the cost of getting back on the equilibrium path. We note that strategies in (**) are described in a way also applicable to subsequent Cases 2–6, where the on-path strategy (*) is different.

(**) Let n' and k' denote sizes of the remaining possibility sets of players 1 and 2, respectively, at the beginning of an off-equilibrium period in the sense that (n', k') does not arise in any period following (*). Let X_i , $i = 1, 2$, denote the set of all possible sizes of the remaining possibility set that player i starts with in some period along the equilibrium path according to (*) above. For each $n' \geq 1$, let $\bar{n}'$ denote the smallest number in X_1 subject to being equal to or larger than n' , and let $\underline{n}'$ denote the largest number in X_1 that is strictly lower than n' if it exists and $\underline{n}' = n'$ otherwise. For each $\bar{n}' \in X_1$, let $k(\bar{n}')$ denote the number of elements with which player 2 ends the period in which player 1 starts with $\bar{n}'$ and discloses no element according to (*). Analogously, for each $k' \geq 1$, let $\bar{k}'$ denote the smallest number in X_2 subject to being equal to or larger than k' , and let $\underline{k}'$ denote the largest number in X_2 that is strictly lower than k' if it exists and $\underline{k}' = k'$ otherwise. For each $\bar{k}' \in X_2$, let $n(\bar{k}')$ denote the number of elements with which player 1 ends the period in which player 2 starts with $\bar{k}'$ and discloses no element according to (*).

- (i) If $k' \geq 2$ and $n' > n(\bar{k}')$, then player 1 discloses $n' - n(\bar{k}')$ elements if his expected payoff from doing so, followed by player 2 disclosing $k' - \underline{k}'$ elements if $k' - \underline{k}' > 0$ and they follow (*), is positive, and player 2 does nothing; otherwise, player 1 does nothing and player 2 discloses all but one element if her payoff from doing so is positive, and does nothing otherwise.
- (ii) If $k' \geq 2$ and $\min X_1 \leq n' \leq n(\bar{k}')$, then player 2 discloses $k' - k(\bar{n}')$ elements if her expected payoff from doing so, followed by player 1 disclosing $n' - \underline{n}'$ elements if $n' - \underline{n}' > 0$ and they follow (*), is positive, and player 1 does nothing; otherwise, player 2 does nothing and player 1 discloses all but one element if his payoff from doing so is positive, and does nothing otherwise.
- (iii) If $k' \geq 2$ and $1 < n' < \min X_1$, then player 2 discloses $k' - 1$ elements if her expected payoff from doing so, is positive, and player 1 does nothing; otherwise,

player 2 does nothing and player 1 discloses all but one element if his payoff from doing so is positive, and does nothing otherwise.

- (iv) If $k' = 1$ or $n' = 1$, both players disclose no element and take action ω^* if identified.

We now verify optimality of (*) and then that of (**). For this, we analyze every possible period in the game, presuming that neither player has identified ω^* at that point, which is taken for granted in the sequel. Also, we take it for granted that $\delta < 1$ is large enough for the argument to be valid whenever pertinent.

Along the equilibrium path, note that player 2 in even periods with k' elements faces a continuation equilibrium that is equivalent to σ^* when $\#(S_1) = k' - 1$ and $\#(S_2) = k'$, in terms of probabilities of winning and losing in subsequent periods, so his payoff is $\psi(k')$. Unless $k' = 1$, therefore, disclosing two elements is optimal for her because disclosing less will subject her to disclosing more in subsequent periods by (ii) above, and disclosing more would expose her to a higher risk of losing before getting back to the equilibrium path at some future point by (i). If $k' = 1$, optimality of taking action ω^* as soon as identified is trivial. In odd periods along the equilibrium path, doing nothing is clearly optimal for player 2 for the same reason that it is so in σ^* because the opponent is expected to disclose a positive number of elements (or take action ω^* if identified).

An analogous logic verifies optimality of player 1 along the equilibrium path from period 2 onward. In period 1, as $\delta \rightarrow 1$, player 1's expected payoff converges to that of disclosing $r + m$ elements in two consecutive installments: r elements first; then m . This payoff conditional on the first r elements disclosed not containing ω^* , converges to $v_1^* = ((k^2 - 1)\alpha + (k^2 + 1)\beta)/(2k^2)$ in (10). Therefore, player 1's payoff in period 1 according to (*) converges (as $\delta \rightarrow 1$) to

$$\frac{r\beta + mkv_1^*}{r + mk} = \frac{1}{r + mk} \left(\frac{m(k^2 - 1)\alpha + (2kr + mk^2 + m)\beta}{2k} \right),$$

which is positive by (1) if $m \geq 2$, or $m = 1$ and $r \leq \hat{r} = \frac{k^2 - 3}{2k}$, because given $k \geq 4$,

$$\begin{aligned} & 2m(k^2 - 1) - (2kr + mk^2 + m) \\ &= m(k^2 - 3) - 2kr \begin{cases} \geq 2(k(k - r) - 3) > 0 & \text{if } m \geq 2 \\ = k(k - 2r) - 3 > 0 & \text{if } m = 1, r \leq \hat{r}. \end{cases} \end{aligned}$$

In period 1, disclosing less than $r + m$ elements only delays the process by (i), and disclosing more is suboptimal by the same reason as above, verifying optimality of player 1 along the equilibrium path. Given this, doing nothing is clearly optimal for player 2 in period 1.

As mentioned, the off-equilibrium strategy (i)–(iv) stipulates that the players get back to the equilibrium path as quickly as possible if doing so gives a positive expected payoff, with the added incentive feature that the player who deviated by delaying the equilibrium exchange process bears the cost of getting back on the equilibrium path

by disclosing first. The optimality of the off-equilibrium strategy can be verified analogously to how it is done for σ^* (in Stage 1 above) and, hence, is omitted here.

Case 2: $m \geq 1$ and k is even, but not $m = 1$ and $r > \hat{r}$. Consider the following Markov strategy along the equilibrium path, presuming that $r \neq 0$ (the analysis is the same when $r = 0$ except that the first period below is redundant):

(*) Player 1 discloses r elements and player 2 does nothing in period 1; player 2 discloses one element and player 1 does nothing in period 2; then starting with player 1, the two players alternate disclosing $2m$ and 2 elements each period, respectively, until ω^* is identified and taken as an action. Note that player 2 discloses all but one in the last potential period of disclosure.

The same description of off-equilibrium strategy (**) applies here. The optimality is verified in the same manner as above, except for period 1, which we explain below.

Since player 1's continuation payoff as of period 2 is v_1^* in (9), his payoff in period 1 converges to

$$\frac{r\beta + mk(\alpha + \beta)/2}{r + mk} = \frac{1}{r + mk} \left(\frac{mk\alpha + (2r + mk)\beta}{2} \right),$$

which is positive if $m \geq 2$, or $m = 1$ and $r \leq \hat{r}$, because

$$2mk - 2r - mk = mk - 2r \begin{cases} \geq 2(k - r) > 0 & \text{if } m \geq 2 \\ = k - 2r > 0 & \text{if } m = 1, r \leq \hat{r} < k/2. \end{cases}$$

In period 1, therefore, disclosing r elements is optimal for player 1 because disclosing less only delays the process by (ii) of (**), and disclosing more is suboptimal by the same reason as above. Given this, doing nothing is clearly optimal for player 2 in period 1.

Below we consider cases in which $m = 1$ and $r > \hat{r} = (k^2 - 3)/(2k)$. As $\hat{r} > k/3$, we use $k/3$ instead of $\hat{r}$ when it suffices and simplifies calculation.

Case 3: $m = 1$ so that $n = k + r < 2k$, $r > \hat{r}$, n is even, and $n/2$ is odd. Consider the following Markov strategy along the equilibrium path.

(*) Player 2 discloses $k - n/2$ elements in period 1; player 1 discloses two elements in period 2; then starting with player 2, the two players alternate disclosing 2 and 4 elements each period, respectively, until ω^* is identified and taken as an action. Player 2 discloses all but one element in the last potential period of disclosure.

The same description of off-equilibrium strategy (**) applies here. The optimality is verified in the same manner as before, apart from some minor modification of details for periods 1 and 2, explained below.

In period 2, players start a continuation equilibrium that has been described and verified in Case 1 above where $m = 2$ and $r = 0$. The limit continuation payoffs are v_1^* and v_2^* in (10). Thus, player 2's payoff in period 1 converges to

$$\frac{(k - n/2)\beta + n \cdot v_2^*/2}{k} = \frac{1}{2k} \left(\frac{(n^2 + 4)\alpha + (4kn - n^2 - 4)\beta}{2n} \right),$$

which is positive because

$$2n^2 + 8 - (4kn - n^2 - 4) \geq -k^2 + 2kr + 3(r^2 + 4)|_{r=k/3} = 12 > 0.$$

In period 1, therefore, disclosing $k - n/2$ elements is optimal for player 2 because disclosing less only delays the process by (ii) of (**), and disclosing more is suboptimal by the same reason as above.

Case 4: $m = 1$ so that $n = k + r < 2k$, $r > \hat{r}$, and both n and $n/2$ are even. Consider the following Markov strategy along the equilibrium path.

(*) Player 2 discloses $k - n/2 + 1$ elements in period 1; then starting with player 1, the two players alternate disclosing 4 and 2 elements each period, respectively, until ω^* is identified. Player 2 discloses all but one in the last period of potential disclosure.

The same description of off-equilibrium strategy (**) applies here. The optimality is verified in the same manner as before, apart from a minor modification of details for period 1, explained below.

Player 2's payoff in period 1 is equivalent to that when he discloses $k - n/2$ elements and then another element before the alternation starts. Therefore, player 2's payoff in period 1 converges to

$$\frac{(k - n/2)\beta + n(\alpha + \beta)/2}{k} = \frac{1}{2k} \left(\frac{n\alpha + (4k - n)\beta}{2} \right),$$

which is positive because

$$2n - (4k - n) \geq -k + 3r|_{r=k/3} = 0.$$

In period 1, therefore, disclosing $k - n/2 + 1$ elements is optimal for player 2 because disclosing less only delays the process by (ii) of (**), and disclosing more increases the risk of losing.

Case 5: $m = 1$ so that $n = k + r < 2k$, $r > \hat{r}$, n is odd and $\kappa = (n - 1)/2$ is even. Note that $\kappa \geq 3$.

Consider the following Markov strategy along the equilibrium path.

(*) Player 2 discloses $k - \kappa - 1$ elements in period 1; player 1 discloses one element in period 2; player 2 discloses two in period 3; then starting with player 1, the two players alternate disclosing 4 and 2 elements each period, respectively, until ω^* is taken as an action. Player 2 discloses all but one in the last period of potential disclosure.

The same description of off-equilibrium strategy (**) applies here. The optimality is verified in the same manner as before, apart from some minor modification of details for periods 1, 2, and 3, explained below.

Note that player 2's payoff in period 3 converges to

$$U = \frac{\beta + \kappa(\alpha + \beta)/2}{\kappa + 1} = \frac{\kappa\alpha + (\kappa + 2)\beta}{2(\kappa + 1)},$$

which is positive because $\kappa \geq 3$. Hence, disclosing two is optimal for her in period 3 for the now usual reasoning. Then player 1's payoff in period 2 converges to

$$U' = \frac{\beta + (n - 1)(\alpha + \beta - U)}{n} = \frac{1}{n} \left(\frac{(n - 1)(\kappa + 2)\alpha + ((n + 1)\kappa + 2)\beta}{2(\kappa + 1)} \right),$$

which is positive because

$$\begin{aligned} 2(n-1)(\kappa+2) - ((n+1)\kappa+2) &\geq (k^2 + 2k(r+2) + r^2 + 4r - 9)/2|_{r=k/3} \\ &= 8k^2/9 + 8k/3 - 9/2 > 0. \end{aligned}$$

Hence, it is optimal for him to disclose one in period 2 again for the usual reasoning.

Player 2's payoff in period 1 converges to

$$\begin{aligned} &\frac{(k - \kappa - 1)\beta + (\kappa + 1)(\alpha + \beta - U')}{k} \\ &= \frac{1}{k} \left((k - \kappa - 1)\beta + (\kappa + 1) \frac{((n+1)\kappa+2)\alpha + (n-1)(\kappa+2)\beta}{2n(\kappa+1)} \right) \\ &= \frac{1}{k} \left(\frac{(\kappa+1)((n+1)\kappa+2)\alpha + (\kappa+1)((n-1)(\kappa+2) + 2n(k - \kappa - 1))\beta}{2n(\kappa+1)} \right), \end{aligned}$$

which is positive because

$$\begin{aligned} &2(\kappa+1)((n+1)\kappa+2) - (\kappa+1)((n-1)(\kappa+2) + 2n(k - \kappa - 1)) \\ &= (-k^3 + k^2(r-1) + k(5r^2 + 2r + 9) + 3(r^3 + r^2 + 3r + 3))/4 \\ &> (-k^3 + k^2(r-1) + k(5r^2 + 2r + 9) + 3(r^3 + r^2 + 3r + 3))/4|_{r=\hat{r}} \\ &= 3(3k^6 + 2k^5 + 3k^4 + 4k^3 + 21k^2 + 18k - 27)/(32k^3) > 0, \end{aligned}$$

where the first inequality follows from the derivative of the left-hand side with respect to r being

$$(k^2 + 2k + 10kr + 9r^2 + 6r + 9)/4 > 0.$$

In period 1, therefore, disclosing $k - \kappa - 1$ elements is optimal for player 2 because disclosing less only delays the process by (ii) of (**), and disclosing more increases the risk of losing.

Case 6: Finally, suppose $m = 1$ so that $n = k + r < 2k$, $r > \hat{r}$, and both n and $\kappa = (n-1)/2$ are odd. Consider the following Markov strategy along the equilibrium-path:

(*) Player 1 discloses one element in period 1; player 2 discloses $k - \kappa$ elements in period 2; player 1 discloses two elements in period 3; then starting with player 2, the two players alternate disclosing 2 and 4 elements each period, respectively, until ω^* is taken as an action. Player 2 discloses all but one in the last period of potential disclosure.

The same description of off-equilibrium strategy (**) applies here. The optimality is verified in the same manner as before, apart from some minor modification of details for periods 1 and 2, explained below.

In period 3, player 1's continuation payoff converges to $v_1^* > 0$ in (10), where $N = \kappa$; hence, disclosing two elements is optimal for the now usual reasoning. Since player 2's continuation payoff of period 3 converges to $v_2^* > 0$ in (10), where $N = \kappa$, her expected payoff in period 2 converges to

$$U'' = \frac{(k - \kappa)\beta + \kappa \cdot v_2^*}{k} = \frac{1}{k} \left(\frac{(\kappa^2 + 1)\alpha + (2k\kappa - \kappa^2 - 1)\beta}{2\kappa} \right),$$

which is positive because

$$\begin{aligned} 2(\kappa^2 + 1) - (2k\kappa - \kappa^2 - 1) &\geq \frac{-k^2 + 2k(r-1) + 3(r^2 - 2r + 5)}{4} \Big|_{r=\hat{r}} \\ &= \frac{3k^4 - 20k^3 + 30k^2 + 36k + 27}{16k^2} > 0. \end{aligned}$$

Thus, disclosing $k - \kappa$ elements in period 2 is optimal for player 2 by the now usual reasoning.

Therefore, player 1's payoff in period 1 converges to

$$\frac{\beta + (n-1)(\alpha + \beta - U'')}{n} = \frac{1}{n} \left(\frac{(n-1)(2k\kappa - \kappa^2 - 1)\alpha + ((n-1)(\kappa^2 + 1) + 2k\kappa)\beta}{2k\kappa} \right),$$

which is positive because

$$\begin{aligned} \Delta &= 2(n-1)(2k\kappa - \kappa^2 - 1) - ((n-1)(\kappa^2 + 1) + 2k\kappa) \\ &= (5k^3 + k^2(7r - 11) - k(r^2 + 2r + 9) - 3(r^3 - 3r^2 + 7r - 5))/4 \end{aligned}$$

is concave in r with a positive slope at $r = 1$ as

$$\frac{\partial \Delta}{\partial r} \Big|_{r=1} = \frac{7k^2 - 2k - 2kr - 3(3r^2 - 6r + 7)}{4} \Big|_{r=1} = \frac{7k^2}{4} - k - 3 > 0$$

and $\partial^2 \Delta / \partial r^2 = (9 - k - 9r)/2 < 0$, and Δ is positive both at $r = 1$ and $r = k - 1$ for every $k \geq 4$:

$$\Delta|_{r=1} = k(5k^2 - 4k - 12)/4 > 0 \quad \text{and} \quad \Delta|_{r=k-1} = 2(k^3 - 7k + 6) > 0.$$

In period 1, therefore, disclosing one element is optimal for player 1, because disclosing none only delays the process by (ii) of (**), and disclosing more increases the risk of losing.

In the equilibrium constructed above for each of the six cases, in every period, one player discloses at least one element for sure. Therefore, ω^* is identified and taken as an action in period $n + k + 1$ at the latest, and efficiency is achieved as $\delta \rightarrow 1$. This completes the proof. $\square$

REFERENCES

- Admati, Anat R. and Motty Perry (1991), "Joint projects without commitment." *Rev. Econ. Stud.*, 58, 259–276. [1187]
- Amitai, Mor (1996), "Cheap-talk with incomplete information on both sides." DP 90, Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem. [1186]

- Augenblick, Ned and Aaron Bodoh-Creed (2018), “To reveal or not to reveal: Privacy preferences and economic frictions.” *Games and Economic Behavior*, 110, 318–329. [1186, 1187]
- Aumann, Robert and Sergiu Hart (2003), “Long cheap talk.” *Econometrica*, 71, 1619–1660. [1186]
- Austen-Smith, David (1990), “Information transmission in debate.” *Amer. Jour. Pol. Sci.*, 34, 124–152. [1186]
- Baliga, Sandeep and Stephen Morris (2002), “Co-ordination, spillovers, and cheap talk.” *Jour. Econ. Theory*, 105, 450–468. [1186]
- Barany, Imre (1992), “Fair distribution protocols or how players replace fortune.” *Math. Oper. Res.*, 17, 327–340. [1186]
- Ben-Porath, Elchanan (2003), “Cheap talk in games with incomplete information.” *Jour. Econ. Theory*, 108, 45–71. [1186]
- Blume, Andreas, Oliver Board, and Kohei Kawamura (2007), “Noisy talk.” *Theoretical Economics*, 2, 395–440. [1186]
- Blume, Andreas and In-Uck Park (2020), “Quid pro quo: Friendly information exchange between rivals.” Economics Discussion Paper 20/733, University of Bristol, UK. [1202]
- Compte, Olivier and Philippe Jehiel (2004), “Gradualism in bargaining and contribution games.” *Rev. Econ. Stud.*, 71, 975–1000. [1187]
- Crawford, Vincent P. and Hans Haller (1990), “Learning how to cooperate: Optimal play in repeated coordination games.” *Econometrica*, 58, 571–595. [1193]
- Crawford, Vincent P. and Joel Sobel (1982), “Strategic information transmission.” *Econometrica*, 50, 1431–1451. [1185]
- Dziuda, Wioletta and Roben Gradwohl (2015), “Achieving cooperation under privacy concerns.” *American Economic Journal: Microeconomics*, 7, 142–173. [1186, 1187]
- Esö, Peter and Yuk-fai Fong (2010), “Wait and see: A theory of communication over time.” Report. [1186]
- Farrell, Joseph and Robert Gibbons (1989), “Cheap talk can matter in bargaining.” *Jour. Econ. Theory*, 48, 221–237. [1186]
- Forges, Francoise (1986), “An approach to communication equilibria.” *Econometrica*, 54, 1375–1385. [1186]
- Forges, Francoise (1990a), “Equilibria with communication in a job market example.” *Quart. Jour. Econ.*, 105, 375–398. [1186]
- Forges, Francoise (1990b), “Universal mechanisms.” *Econometrica*, 58, 1341–1354. [1186]
- Forges, Francoise and Frederic Koessler (2008), “Long persuasion games.” *Jour. Econ. Theory*, 143, 1–35. [1186]

- Fried, Dov (1984), “Incentives for information production and disclosure in a duopolistic environment.” *Quart. Jour. Econ.*, 99, 367–381. [1186]
- Fudenberg, Drew and Jean Tirole (1991), “Perfect Bayesian equilibrium and sequential equilibrium.” *Journal of Economic Theory*, 53, 236–260. [1193]
- Gerardi, Dino (2004), “Unmediated communication in games with complete and incomplete information.” *Jour. Econ. Theory*, 114, 104–131. [1186]
- Glazer, Jacob and Ariel Rubinstein (2004), “On optimal rules of persuasion.” *Econometrica*, 72, 1715–1736. [1186]
- Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson (2014), “Dynamic strategic information transmission.” *Jour. Econ. Theory*, 151, 304–341. [1186]
- Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani (2009), “Mediation, arbitration and negotiation.” *Jour. Econ. Theory*, 144, 1397–1420. [1186]
- Goltsman, Maria and Gregory Pavlov (2014), “Communication in Cournot oligopoly.” *Jour. Econ. Theory*, 153, 152–176. [1186]
- Green, Jerry R. and Nancy L. Stokey (1980), “A two-person game of information transmission.” Discussion Paper, No. 418, Northwestern University, Evanston, IL. [1186]
- Green, Jerry R. and Nancy L. Stokey (2007), “A two-person game of information transmission.” *Jour. Econ. Theory*, 135, 90–104. [1186]
- Grossman, Sanford J. (1981), “The informational role of warranties and private disclosure about product quality.” *Jour. Law Econ.*, 24, 461–483. [1185]
- Halonen-Akatwijuka, Maija and In-Uck Park (2021), “Coordinating public good provision by mediated communication.” *American Economic Journal: Microeconomics*, 13, 311–342. [1186]
- Hörner, Johannes and Andrzej Skrzypacz (2016), “Selling information.” *Jour. Polit. Econ.*, 124, 1515–1562. [1186]
- Krishna, Vijay and John Morgan (2001), “A model of expertise.” *Quart. Jour. Econ.*, 116, 747–775. [1186]
- Krishna, Vijay and John Morgan (2004), “The art of conversation: Eliciting information from informed parties through multi-stage communication.” *Jour. Econ. Theory*, 117, 147–179. [1186]
- Marx, Leslie and Steven Matthews (2000), “Dynamic voluntary contributions to a public project.” *Rev. Econ. Stud.*, 67, 327–358. [1187]
- Matthews, Steven (1989), “Veto threats: Rhetoric in a bargaining game.” *Quart. Jour. Econ.*, 104, 347–370. [1186]
- Matthews, Steven and Andrew Postlewaite (1989), “Pre-play communication in two-person sealed-bid double auctions.” *Jour. Econ. Theory*, 48, 238–263. [1186]

- Milgrom, Paul (1981), “Good news and bad news: Representation theorems and applications.” *Bell Jour. Econ.*, 12, 380–391. [1185]
- Milgrom, Paul and John Roberts (1986), “Relying on the information of interested parties.” *RAND Jour. Econ.*, 17, 18–32. [1186]
- Morgan, John and Philip C. Stocken (2003), “An analysis of stock recommendations.” *RAND Jour. Econ.*, 34, 183–203. [1186]
- Myerson, Roger (1982), “Optimal communication mechanisms in generalized principal-agent problems.” *Jour. Math. Econ.*, 10, 67–81. [1186]
- Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura (1990), “Strategic information revelation.” *Rev. Econ. Stud.*, 57, 25–47. [1186]
- Ottaviani, Marco and Peter N. Sørensen (2006), “Professional advice.” *Jour. Econ. Theory*, 126, 120–142. [1186]
- Park, In-Uck (2002), “Cheap talk coordination of entry by privately informed firms.” *RAND Jour. Econ.*, 33, 377–393. [1186]
- Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2013), “Strategic information exchange.” *Games Econ. Behav.*, 82, 444–467. [1186]
- Seidmann, Daniel J. and Eyal Winter (1997), “Strategic information transmission with verifiable messages.” *Econometrica*, 65, 163–169. [1186]
- Shin, Hyun Song (1994), “The burden of proof in a game of persuasion.” *Jour. Econ. Theory*, 64, 253–264. [1186]
- Stein, Jeremy C. (2008), “Conversations among competitors.” *Amer. Econ. Rev.*, 98, 2150–2162. [1186]

Co-editor Marina Halac handled this manuscript.

Manuscript received 1 December, 2020; final version accepted 22 August, 2021; available online 3 September, 2021.