Trading Information

Dmitry Orlov* Andrzej Skrzypacz† Pavel Zryumov‡

Current draft: 31st August, 2022

Abstract

We analyze a class of dynamic games of information exchange between two players. Each agent possesses information about a binary state that is of interest to the other player and cares about the other player's actions. Preferences are additively separable over own and the other player's actions. We fully characterize the set of equilibrium payoffs that can be sustained in such games and construct equilibria that achieve those payoffs. We show that gradual information exchange dominates static (one-shot) communication. Moreover, the whole set of outcomes that Pareto-dominate static communication can be supported in equilibrium.

Keywords: Bayesian persuasion, stochastic games, real options.

1 Introduction

In many settings, economic agents act as both senders and receivers of valuable information. Efficient informational cooperation is needed to make correct decisions and achieve the best outcomes. Such cooperation may, however, reduce the ability of an agent to use his information as a means of influence to advance his private interests. For example, consider two firms that are members of standard-setting organizations. Firm A possesses private information about the state of its technology and would like to induce Firm B to take certain actions (for example, adopt A's technology in B's product design). Firm B would like to learn as much as possible about A's technology to make the right decision. Yet, it would not be in A's interest to reveal all of his know-how. The

*University of Wisconsin-Madison
†Stanford University
‡University of Rochester Simon Business School

optimal information disclosure by A is known as the optimal Bayesian persuasion. The main insight is that if B also has information that is useful to A, then we can construct equilibria with gradual information disclosure that Pareto-dominate just A optimally persuading B and vice versa. The main takeaways are that information exchange benefits from dynamics and may allow for efficient information exchange. We show, however, that the asymmetry of positions of the two players makes it harder to achieve full information sharing because full information exchange requires that both players continue having sticks and carrots as the information is gradually revealed.

In our model, two parties, A and B, have access to two distinct pieces of information. Player A has access to information θ_B needed by player B to make a decision. Similarly, player B has access to information θ_A needed by player A. Each player can communicate some or all of their information to their counterparty at any time. Information communicated by player A allows player B to take a more informed action. Additional signals communicated by player A generate an *allocation gain* accrued to player B. Player A, however, may have his own preferences over player B's actions – revealing more information could tilt B's action away from the one privately preferred by A and, thus, generate a *persuasion loss* for player A. Similarly, information communicated by player B helps player A take a more informed action, generating an allocation gain for him, but may also result in a persuasion loss for player B if such actions lower his value. We analyze the setting in which allocation gains and persuasion losses are additively separable across states θ_A and θ_B and allow the players to exchange arbitrary credible signals in the spirit of Bayesian persuasion. Our first result shows that when the players can communicate only once, the static Nash equilibrium minimizes the persuasion losses of both players. When allocation gains from sharing all information outweigh persuasion losses, the static equilibrium is inefficiently opaque. Our second result shows that efficient full information sharing can be restored with sequential communication. In equilibrium, the players gradually reveal all information over time while maintaining enough residual uncertainty about their piece of information to keep their counterparty engaged. Finally, we show that any Pareto efficient allocation that dominates the static Nash equilibrium can be achieved in a dynamic communication equilibrium. This allows us to characterize the payoff set of all dynamic equilibria as well as global properties of the constrained-efficient equilibria.

In our model, there are two independent sources of uncertainty: θ_A and θ_B. There are also two players, A and B who take actions a and b. Players A and B are both receivers and senders of information. Player's A utility $u_A(\theta_A, a)$ from taking action a depends on the state θ_A that is observed only by player B. Player B is in a similar predicament in that his utility $u_B(\theta_B, b)$ of taking action b depends on state θ_B observed by player A. Player A derives a private benefit $v_A(\theta_B, b)$ from player B's actions, making it unclear as to whether sharing information about θ_B is in his best interests. Similarly, player B derives a private benefit $v_B(\theta_A, a)$ from player A's action a. Throughout the paper, we assume that players' A and B preferences are additively separable in their own and others' actions, and are given by $u_A + v_A$ and $u_B + v_B$, respectively.

Suppose the agents could only communicate once before taking their actions at $t = 1$. In equilibrium, the agents choose signals that maximize their private benefits v_A and v_B, and do not take into account the allocation gains associated with more informed actions in u_A and u_B. If, however, the allocation gains associated with a more informative signal exceed the persuasion losses relative to this privately optimal outcome, both players may prefer a more informative outcome. The static equilibrium is strictly inefficient in this case. The players cannot sustain this more informative outcome as they would unilaterally defect to the static Nash signal. Their incentives are similar to that of the Prisoner's dilemma. We show that this incentive problem can be ameliorated via sequential communication.

Suppose the agents could communicate dynamically at any time $t \in [0, 1]$. Dynamic communication supports information sharing by rewarding the players through continued reciprocation and punishing their deviations by switching to the static Nash outcome, which is inefficiently opaque. We show that if full information sharing Pareto dominates the static Nash outcome ex-ante and is efficient among other possible information structures, then full information can be achieved as a dynamic equilibrium outcome. We explicitly construct an equilibrium wherein agents communicate increasingly informative signals sequentially. Uncertainty over the remaining information acts as a threat point that keeps both players engaged in equilibrium communication, and as a result, it has to be balanced across players. If one of the players disclosed too much information at once, then the threat of withholding the remaining information becomes weak, and the other player would defect to the static outcome. In this sense, our game of strategic communication bears a resemblance to a nuclear disarmament problem: both players might prefer a fully informative outcome in the static game, but neither would unilaterally share all information. Similar to how most countries prefer a world without nuclear weapons but neither would unilaterally give up its nuclear capabilities. Progress towards efficiency in both cases can be made over time. In our model, players gradually reduce their uncertainty, similar to a gradual reduction of nuclear capabilities.

The efficiency of full information sharing guarantees that along the equilibrium path, at least one of the player's incentive constraints is slack; that is, at least one of the players strictly prefers a fully informative outcome to the static Nash equilibrium. This allows communication to continue in the most crucial moments: when the incentive constraint of player A is tight, he can reduce uncertainty about θ_B and consequently relax his own incentive constraint without violating the incentive constraint of player B. The irreversibility of information sharing allows player A to credibly reduce his incentives to deviate in the future. The same argument applies to player B. We show that this equilibrium structure is not unique to the fully informative outcome and extends to all efficient signals – at least one of the agents always strictly prefers reaching the efficient frontier relative to the static Nash outcome. As a result, every Pareto optimal signal that dominates static Nash outcome can be supported as a dynamic communication equilibrium outcome. This allows us to characterize the payoff set of all feasible equilibria of this stochastic game.

We use the characterization of the equilibrium payoff set to provide comparative statics of the most efficient dynamic equilibrium, which we refer to as the constrained efficient one. We show that as long as players are sufficiently symmetric, both in terms of preferences and priors, then efficiency can be sustained. If, however, one of the parties has a substantial information advantage at the outset, it reduces the informativeness of the equilibrium outcome, generating a loss relative to the social optimum. Similarly, an increase in the private benefits of the agents reduces their interest in socially optimal information sharing, thus reducing transparency and destroying value. Such increases in the private benefits can lead to reductions in players' joint surplus despite a uniform increase in their payoff functions.

1.1 Related Literature

Our model lies in the class of Bayesian persuasion models, pioneered by Kamenica and Gentzkow (2011), and reviewed in Kamenica (2019). Our dynamic environment contributes to three strands of this literature: dynamic persuasion, persuasion by multiple senders, and persuasion of multiple audiences. Our limited commitment setting also contributes to the literature on dynamic hold-up problems.

Hörner and Skrzypacz (2016) consider the dynamic problem of selling information without commitment when one of the parties has access to money, showing that the hold-up problem can be ameliorated by gradual communication in exchange for transfers. We extend this argument by showing that information can be an effective means of payment in itself, as long as agents can share it gradually. This is highly applicable in industries where direct transfers among firms may fall under antitrust regulatory scrutiny. The irreversibility of information shared by an agent in our setting is similar to the irreversibility of effort in Ely and Szydlowski (2020), who consider a model with commitment. Our paper differs from the latter paper as the argument relies on the irreversibility of information flows, and sufficient uncertainty must be maintained for cooperation to succeed. Such gradual communication bears similarity to dynamic investment in the dynamic hold-up problem of Che and Sákovics (2004). We show that efficient data sharing is more than the gradual sharing of the data points to support cooperation – it may require large chunks of information to be shared in order to generate sufficient fluctuations in beliefs within a period. This makes dynamic information sharing conceptually different than dynamic investment. Nevertheless, it is possible to use information as an efficient means of barter as long as players maintain sufficient uncertainty along the path of play. Methodologically, our analysis of Bayesian persuasion in continuous time builds on Orlov, Skrzypacz, and Zryumov (2020), who show that strategies can be modeled directly as stochastic processes and filtrations. We expand on this work by allowing non-Markov equilibria, which are essential for dynamic equilibria to sustain the set of efficient allocations.

The literature on multiple senders has largely focused on whether competition in information among senders is welfare improving. Gentzkow and Kamenica (2017) provide sufficient conditions on the space of available signals for which increasing the number of senders weakly increases the amount of information available to the public. Our focus is to show that dynamic communication can achieve social efficiency even if players only have access to their own pieces of information and are restricted from sending correlated messages. Li and Norman (2021) study a sequential persuasion model of multiple senders. They show that such sequential communication generates weakly lower signal informativeness than simultaneous persuasion by the senders. Our findings are very different in that by allowing the agents to communicate over multiple rounds, we allow for sequential communication that leads to the most informationally efficient outcomes, which dominate static communication.

There have been a considerable number of papers in the Bayesian persuasion literature focused on multiple receivers. Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) consider a model in which a sender communicates with a privately informed receiver, making him act as if he were facing a distribution of receivers. Orlov, Zryumov, and Skrzypacz (forthcoming) consider a stress test model in which the regulator communicates with both banks and investors. Li, Szydlowski, and Yu (2021) consider a model in which an incumbent firm communicates with both a customer and a potential entrant. Inostroza (2019) considers a stress test model with long- and short-term investors. In all of these settings, a receiver of the information is only responsible for their actions, while the sender is only choosing their information. Our model differs conceptually from this literature as each player is both the sender and receiver of information, making them consider the costs and benefits of informed cooperation in a way that is new to the Bayesian persuasion literature.

The rest of the paper is structured as follows. Section 2 provides a simple example of how dynamic communication improves upon the static Nash outcome. This example turns out to be highly representative both in terms of economics and also in the dynamic equilibrium construction of the general model we consider in Section 3. Section 4 provides sufficient conditions for the feasibility of full information sharing in a dynamic equilibrium. Section 5 extends this argument to constrained efficient signals. Section 6 concludes.

2 Binary Action Example

To illustrate the key ideas of the paper consider an example featuring two players, A and B, who simultaneously act as senders and receivers of information. Each player, A and B, would like to match their action, a and b, to their payoff relevant state, $\theta_A \in \{0,1\}$ and $\theta_B \in \{0,1\}$ respectively. However, player A does not observe the state θ_A, instead she starts with a prior $q = \mathrm{P}(\theta_A = 1)$. As a result, she needs to rely on communication by player B, who privately observes θ_A and can share credible signals about it in the spirit of Bayesian persuasion. Similarly, player B starts with a prior $p = \mathrm{P}(\theta_B = 1)$ and needs to rely on player A's communication to learn θ_B. In addition to deriving utility from matching their own action to the state, each player derives utility δ if the other player takes a preferred action 1.

Formally, the preferences of players A and B are given by

$$u_A = \mathbb{1}\{a = \theta_A\} + \delta \cdot \mathbb{1}\{b = 1\}, \qquad u_B = \mathbb{1}\{b = \theta_B\} + \delta \cdot \mathbb{1}\{a = 1\}.$$

Static communication. If the players A and B could only communicate once, they would choose signals $\tilde{p}^N$ and $\tilde{q}^N$ respectively that maximize the probability of their opponent taking the high action 1.[1] If the starting prior p about θ_B exceeds $1/2$, then player B takes the high action with certainty, and, expecting that, player A does not communicate any information and collects utility $\delta \cdot 1$ from player B's actions. If the starting prior $p < 1/2$, then player B would take the low action $a = 0$ in the absence of additional information. Player A can improve upon that outcome by disclosing $\theta_B = 0$ with probability $1 - 2p$ and nothing otherwise, leading to player B taking action $b = 0$ with probability $1 - 2p$ and action $b = 1$ with probability $2p$. Player A then derives utility $\delta \cdot 2p$, which is an improvement over the status quo.[2] Similarly, player B shares no information with player A if $q \geq 1/2$, and reveals $\theta_A = 0$ with probability $1 - 2q$ if $q < 1/2$. Figure 1 illustrates the optimal signals in a 2-dimensional plot as a function of the initial priors. Communication by player A (in blue) moves belief p along the horizontal axis, while communication by player B (in orange) moves belief q along the vertical axis.

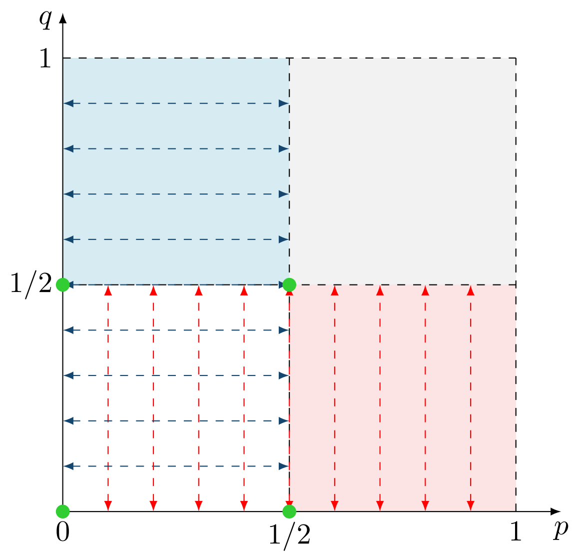

Figure 1: Optimal static persuasion from the starting prior $p_0 = q_0 = 1/4$.

The expected payoffs for players A and B in this static equilibrium are

$$u_A^N(p, q) = \max(q, 1 - q) + \delta \cdot \min(2p, 1), \qquad u_B^N(p, q) = \max(p, 1 - p) + \delta \cdot \min(2q, 1). \quad (1)$$

In the static equilibrium, each player maximizes her own private benefit of persuasion without taking into consideration the benefits of a more informative signal for the other player. When the

[1] We assume that such communication is done simultaneously, but the same argument applies if players move sequentially as the second mover faces a hold-up problem and would choose signals that would be optimal if the moves were simultaneous.

[2] Player A's optimal signal is identical to the trial example of Kamenica and Gentzkow (2011).

persuasion loss from a more informative signal of the sender is dominated by the allocation gain of the receiver, i.e., $\delta < 1$, the equilibrium is inefficiently opaque. In this case, the players are faced with a prisoner's dilemma like incentives: while both parties benefit from a more informative signal, i.e., both prefer cooperation, each player has strict incentives to defect to the opaque signal $\tilde{p}$ and $\tilde{q}$. We next show that the efficient outcome can be restored when the players can communicate dynamically. Sequential communication allows for gradual back-and-forth information revelation that rewards sharing and punishes withholding valuable information.

Dynamic communication. Suppose the agents could commit to sharing all of the information about θ_A and θ_B. This way, both players would always match their action to the state and then benefit from the high action of their opponent only when the corresponding state θ equals to 1. The resulting expected utilities for both agents are

$$u_A^F(p,q) = 1 + \delta \cdot p, \qquad u_B^F(p,q) = 1 + \delta \cdot q. \tag{2}$$

Comparing (2) with (1) we see that both players prefer full information over the static equilibrium whenever the priors (p, q) belong to set S^{FI}

$$S^{FI} \stackrel{def}{=} \left\{ (p,q) : \quad \delta \cdot \min(p, 1-p) \leq \min(q, 1-q) \leq \frac{1}{\delta} \cdot \min(p, 1-p) \right\}.$$

Set S^{FI} is illustrated by the blue shaded area in Figure 2 and is non-empty whenever $\delta \in [0, 1]$. Whenever $(p, q) \in S^{FI}$, the static equilibrium is not only inefficient, but it is Pareto dominated by the fully informative outcome. In what follows, we show that efficient full information sharing can be achieved starting from any point in S^{FI} when players can communicate sequentially.

Consider the following sequential communication protocol. For a starting prior (p, q) in set S^{FI}, agent A designs a signal at $t = 1/2$ about θ_B so that the posterior belief $P_{1/2}$ jumps either to the left or right contour of set S^{FI} along the horizontal axis. Then, at $t = 3/4$, agent B communicates enough information about θ_A as to bring the posterior belief $Q_{3/4}$ to either the top or bottom contour of set S^{FI}.[3] Repeating this process for every $t = 1 - \frac{1}{2^k}$ we see that the beliefs become degenerate at $t = 1$, which corresponds to full information sharing by both agents. To ensure that both parties follow the proposed protocol, deviations from it are punished by switching to

[3] This is always feasible given the star shape of set S^{FI} – a property that we formally define as contour convexity in Section 4.

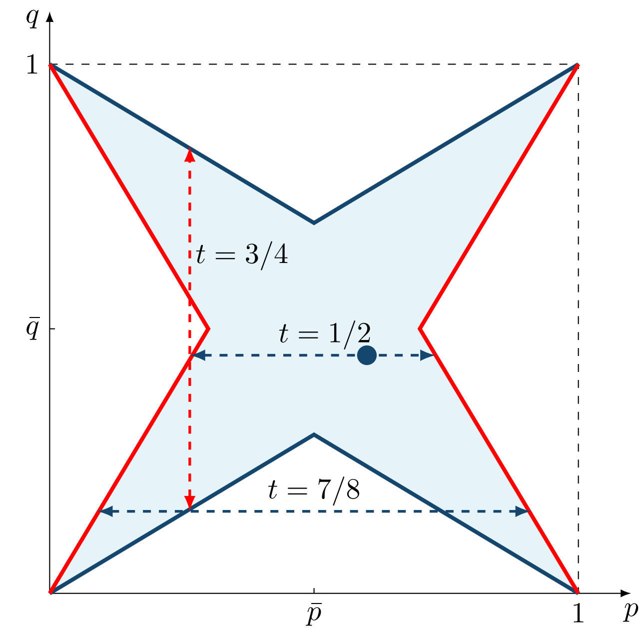

Figure 2: Dynamic Equilibrium that achieves full information outcome.

the static equilibrium communication – signals $\tilde{p}^N$ and $\tilde{q}^N$ described earlier. Given such a threat, the strategies above constitute an equilibrium since the posterior beliefs never exit the region S^{FI} where the full information outcome dominates the static equilibrium outcome for both agents.

Dynamic communication allows for gradual information revelation. In the constructed equilibrium, each player sends informative signals over time, provided that the counterparty reciprocates. To keep player B engaged, player A maintains some residual uncertainty about θ_B and decreases it only in return for a lower uncertainty about θ_A. Hence, full information is achieved gradually along the path of declining uncertainty about both θ_A and θ_B, despite the players being unable to perform simultaneous correlated tests.

Pareto efficiency of the fully informative signal ensures that if both players prefer fully informative allocation prior to any communication, then their incentives to share increasingly informative signals can be sustained everywhere along the equilibrium path. This turns out to be quite general. In what follows, we introduce a general model with arbitrary actions and a rich set of preferences and show that ex-ante efficient outcomes can be sustained as dynamic equilibria as long as they dominate the static Nash outcomes for both players. This also allows us to characterize equilibrium behavior globally and provide comparative statics of these constrained efficient outcomes.

3 Model

3.1 Setup

Players and Payoffs. Two players A and B choose action take an irreversible action $a \in \mathbb{A}$ and $b \in \mathbb{B}$ respectively, where $\mathbb{A}$ and $\mathbb{B}$ are Polish spaces. The state of the world is captured by two independent variables $\theta_A, \theta_B \in \{0, 1\}$ that affect the players' payoffs.

The payoff of player A, $u_A(\theta, a)$, (player B, $u_B(\theta_B, b)$) from taking an irreversible action a (b) depends directly only on one part of the state, namely θ_A (θ_B). Player A privately observes θ_B, and player B privately observes θ_A, i.e., neither player knows her action relevant part of the state. Instead each player A (B) holds a commonly know prior of her action relevant state $q_0 = \mathrm{P}(\theta_A = 1)$ and $p_0 = \mathrm{P}(\theta_B = 1)$.

In addition to utility from their own actions u_A and u_B, each player receives utility from the action taken by the other player. That is, total utility is given by

$$U_A = u_A(\theta_A, a) + v_A(\theta_B, b), \qquad U_B = u_B(\theta_B, b) + v_B(\theta_A, a). \tag{3}$$

We study a game that consists of two distinct stages: information sharing and action. The action stage occurs at the final moment of the game $t = 1$, and the players have to choose their actions a and b simultaneously. Communication occurs for $t \leq 1$ and results in time $t = 1$ posteriors q about state θ_A and p about state θ_B. Given the posteriors q and p the optimal actions of the agents A and B $a^*(q)$ and $b^*(p)$ are given by[4]:

$$a^*(q) \stackrel{def}{=} \arg\max_{a \in \mathbb{A}} \mathrm{E}_q \left[u_A(\theta_A, a) \right], \qquad b^*(p) \stackrel{def}{=} \arg\max_{b \in \mathbb{B}} \mathrm{E}_p \left[u_B(\theta_B, b) \right]. \tag{4}$$

If there are multiple actions that maximize the payoff of player i we, without loss, choose the action that benefits player j to maintain the continuity of the game. Thus, we can think of $a^*(q)$ and $b^*(p)$ as functions rather than correspondences.

[4]We assume that the utility functions u_A and u_B are such that arg max is well defined for any $p, q \in [0, 1]$.

3.2 Static Strategies and Equilibrium Concept

First, we define the strategies of the static benchmark. The information sharing stage of the static game occurs at $t = 0$. During the information sharing stage, the players can simultaneously share arbitrary information about the privately known state via the Bayesian persuasion technology. We summarize the strategy of each player using the posterior belief that it induces.

Definition. *A static information sharing strategy of Player A (B) is a single random variable* $\tilde{p} \in [0,1]$ *(*$\tilde{q} \in [0,1]$*), that denotes the posterior belief about* θ_A *(*θ_B*) conditional on the observed signal, that satisfies Bayes plausibility constraint, i.e.,* $\mathrm{E}[\tilde{p}] = p_0$ *(*$\mathrm{E}[\tilde{q}] = q_0$*).*

We define a Static Nash equilibrium as a pair of information sharing and action strategies that are mutual best responses.

Definition (Static Nash Equilibrium)**.** *A Static Nash equilibrium is a pair of information sharing strategies* $(\tilde{p}^*, \tilde{q}^*)$ *such that*

(A) signal $\tilde{p}^*$ *is a best response to action* $b^*(\tilde{p})$ *chosen by player B:*

$$\tilde{p}^* \in \arg\max_{\tilde{p}} \mathrm{E}\left[u_A(\theta_A, a^*(\tilde{q}^*)) + v_A(\theta_B, b^*(\tilde{p}))\right];$$

(B) signal $\tilde{q}^*$ *is a best response to the action plan* $a^*(\tilde{q})$ *chosen by player A:*

$$\tilde{q}^* \in \arg\max_{\tilde{q}} \mathrm{E}\left[u_B(\theta_B, b^*(\tilde{p}^*) + v_B(\theta_A, a^*(\tilde{q}))\right].$$

We can see that agent A's choice of signal $\tilde{p}$ does not affect the equilibrium signal $\tilde{q}^*$ and action $a^*(\tilde{q}^*)$. Consequently, the sole purpose of signal $\tilde{p}$ is to maximize the expected value of player B's action to player A, i.e., v_A. If there are two different signals $\tilde{p}^*$ and $\tilde{p}^{*'}$ that are optimal for agent A, it means they both maximize the expected value of player B's action to player B, and thus generate the same payoff for player A, if we hold $\tilde{q}^*$ fixed. However signals $\tilde{p}^*$ and $\tilde{p}^{*'}$ may carry different values for player B who prefers more information to less. Based on this argument, the set of static Nash equilibria forms a lattice.

Definition. *A static Nash equilibrium is* ***minimal*** *if the agents' payoffs under it are (weakly) Pareto dominated by all other static Nash equilibria. Denote the minimal static Nash equilibrium as* $(\tilde{p}^N, \tilde{q}^N)$*.*

In many instances, like the example we consider in Section 2, the static Nash equilibrium is unique, but Definition 3.2 is useful whenever there is multiplicity given agents' preferences.

3.3 Dynamic Strategies and Equilibrium Concept

Next, we allow for multiple rounds of communication and expand our definition of the strategies to the dynamic setting. To do so, we allow the agents, without loss, to communicate any time $t \in [0, 1]$, continuously if needed. Within every "period" of communication, the agents first simultaneously provide informative signals about θ_A and θ_B by conducting a test. The test induces a posterior over θ_A and θ_B from some distribution subject to the martingale constraint that the average posterior belief has to be equal to the prior. The agents can commit within a period to an arbitrary distribution, but they cannot commit to future signals. Heuristically, the sequence of events in a short period of time dt is shown in Figure 3. The agents' information sharing strategy is a function of the past history, i.e. information shared about θ_A and θ_B *up to* time t.

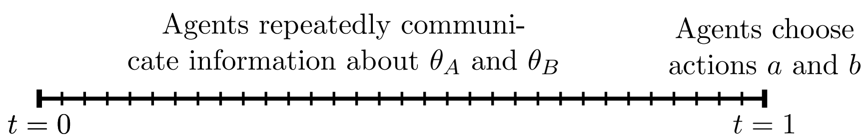

Figure 3: Timing of Events

In our model, the agents control information flow about θ_A and θ_B, which affects their posterior beliefs over time. Denote by $\mathcal{F}_t$ all information available to the players at time t.[5] Denote by P_t and Q_t to be the posterior belief about θ_A and θ_B respectively, given by

$$Q_t \overset{def}{=} \mathrm{P}(\theta_A = 1 \,|\, \mathcal{F}_t)\,, \qquad P_t \overset{def}{=} \mathrm{P}(\theta_B = 1 \,|\, \mathcal{F}_t)\,.$$

We allow agents to continuously generate informative signals whose distributions are contingent on the history $\mathcal{F}_t$ and realizations of θ_A and θ_B. We require only that information disclosed by the agent A (B) at time t is independent of future increments of Q (P) to ensure that the belief process P (Q) is "not forward-looking", i.e., it does not foresee the future evolution of Q (P).

Definition (Dynamic Strategy)**.** *Belief processes $P = (P_t)_{t\in[0,1]}$ and $Q = (Q_t)_{t\in[0,1]}$ are admissible martingales if*

[5]Formally, $\mathcal{F}_t$ is the σ-algebra generated by all the signals disclosed by the agent. For technical reasons, we require filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ to be complete and right-continuous. See **?**, Appendix E.1 for details.

(i) they take values in $[0,1]$*;*

(ii) they are right-continuous with left limits (rcll) with respect to the natural filtration of (P,Q)*;*

(iii) P_t *is independent of* Q*'s future innovation paths* $\{Q_{t+s} \,|\, Q_t\}_{s\in[0,1-t]}$ *for every* $t \in [0,1]$ *and vice versa, i.e,* Q_t *is independent of the future innovations in the path of* P*.*

Agents' strategies are admissible belief processes (i.e., instead of modeling messages the agents send, we represent the strategies directly in terms of the posterior beliefs). We require the belief process to be right-continuous to capture the idea that the agents disclose information in every period. After observing the information shared prior to and including time $t = 1$, each player takes an irreversible action and receives a payoff that corresponds to $\alpha^*(Q_1)$ and $\beta^*(P_1)$.

3.4 Dynamic Equilibrium Definition

We formulate and analyze the bilateral communication game in continuous time for tractability. A subgame equilibrium is a pair of communication strategies $P^* = (P_t^*)_{t\in[0,1]}$ and $Q^* = (Q_t^*)_{t\in[0,1]}$ that are mutual best responses. We capture the notion of subgame perfection by requiring the equilibrium belief processes to be Markov along the path, even if each agent can, in principle, deviate to non-Markov policies. As the players observe the choice of signals made by their counterparties, making this a game of perfect public monitoring, they can respond to such deviations by switching away from the on-path equilibrium behavior.

To capture the idea of how a deviation is detected in this continuous time game, suppose at time t the history of the game is $(\mathcal{F}_s)_{s\leq t}$, where $\mathcal{F}_s = \sigma\left\{(u, P_u, Q_u)_{u\leq s}\right\}$ is the σ-algebra generated by the belief processes $(P_u)_{u\leq t}$ and $(Q_u)_{u\leq t}$ chosen by the players up to that time. If the agents play their on-path equilibrium strategies P^* and Q^*, then this would correspond to an on-path filtration $\mathbb{F}^* = (\mathcal{F}_t^*)_{t\geq 0}$.

Definition (Identical Filtrations)**.** *We say that two filtrations are identical up to a random time* τ*, measurable with respect to their union, if* $\mathcal{F}_{\hat{\tau}} = \mathcal{F}^*_{\hat{\tau}}$ *for any stopping time* $\hat{\tau} \leq \tau$*.*

At the random time τ, when the agents observe a deviation from the expected path of play, we assume that the equilibrium switches to the *minimal* static Nash equilibrium upon observing such a deviation. We impose this for simplicity and without loss as the minimal static Nash equilibrium achieves the lowest payoffs across all dynamic equilibria and poses the most efficient threat point.

Definition (Markov Threat Equilibrium)**.** *A Perfect Bayesian Equilibrium is a pair of communication processes* $(P_t^*, Q_t^*)_{t\in[0,1]}$ *measurable with respect to* $\mathbb{F}$*, such that*

1. *The joint belief process* $(P_t^*, Q_t^*)_{t\geq 0}$ *is Markov.*

2. *Player A's strategy* P^* *is optimal given player B's strategy* Q^**:*

$$P^* \in \arg\max_{\hat{P}} \mathrm{E}\left[\, u_A\big(\theta_A, \alpha(Q_1)\big) + \delta_A \cdot \beta^*\big(\hat{P}_1\big) \;\Big|\; \hat{P}_0 = p_0,\ Q_0^* = q_0\right].$$

   where $Q_1 = \tilde{q}^N(Q_\tau^*) \cdot \mathbb{1}\{\tau \leq 1\} + Q_1^* \cdot \mathbb{1}\{\tau > 1\}$ *and* $\tau =$

3. *Player B's strategy* Q^* *is optimal given player A's strategy* P^**:*

$$Q^* \in \arg\max_{\hat{Q}} \mathrm{E}\left[\, u_B\big(\theta_B, \beta(P_1)\big) + \delta_A \cdot \alpha^*\big(\hat{Q}_1\big) \;\Big|\; \hat{Q}_0 = q_0,\ P_0^* = p_0\right].$$

   where $P_1 = \tilde{p}^N(P_\tau^*) \cdot \mathbb{1}\{\tau \leq 1\} + P_1^* \cdot \mathbb{1}\{\tau > 1\}$

It is worth highlighting why we define the deviation time with regard to filtrations. The agents' strategies, described in Definition 3.3, require the agents to choose distributions of belief processes rather than actual paths. The distributions chosen are observed by both agents. A deviation by the agent is thus a deviation in the distribution of the observed process. While a distribution is a highly non-parametric object, the probabilities conditional on the filtration imposed by the corresponding belief process are readily available. Thus, we can think of information in the sense of filtrations, i.e., knowledge sets, similar to Green and Stokey (1978), and can define the first deviation in those terms for arbitrary belief distributions chosen by the agents.

4 Full Information Equilibria

4.1 Socially Optimal Signals

First, consider a social planner who chooses signals $\tilde{p}$ and $\tilde{q}$ about θ_B and θ_A to maximize the joint welfare of both players. In this benchmark, players retain their decision-making roles, and choose actions $a^*(\tilde{p})$ and $b^*(\tilde{q})$ contingent on their posterior beliefs.

Definition (Socially optimal signals.)**.** *Signals* $\tilde{p}^{FB}$ *and* $\tilde{q}^{FB}$ *maximize players' welfare if*

$$(\tilde{p}^{FB}, \tilde{q}^{FB}) \in \arg\max_{\tilde{p},\tilde{q}} \Big\{ \overbrace{\mathrm{E}_{p,q}\left[u_A(\theta_A, a^*(\tilde{q}) + v_A(\theta_B, b^*(\tilde{p}))\right]}^{\textit{Player A's expected payoff}} \quad + \quad \overbrace{\mathrm{E}_{p,q}\left[u_B(\theta_B, b^*(\tilde{p}) + v_B(\theta_A, a^*(\tilde{q}))\right]}^{\textit{Player B's expected payoff}} \Big\} \tag{5}$$

By rearranging terms in (5) we see that the problem is separable across signals $\tilde{p}$ and $\tilde{q}$. By applying the now standard concavification approach of Kamenica and Gentzkow (2011), socially optimal signals $\tilde{p}^{FB}$ and $\tilde{q}^{FB}$ must satisfy

$$\begin{aligned} \overbrace{\mathrm{E}_p\left[u_B\left(\theta_B, b^*\left(\tilde{p}^{FB}\right)\right) + v_A\left(\theta_B, b^*\left(\tilde{p}^{FB}\right)\right)\right]}^{\textit{players' joint payoff specific to } \theta_B \textit{ and } \tilde{p}} &= \mathrm{cav}_p\left[\mathrm{E}_p\left[u_B(\theta_B, b^*(p)) + v_A(\theta_B, b^*(p))\right]\right](p), \\ \underbrace{\mathrm{E}_q\left[u_A\left(\theta_A, a^*\left(\tilde{q}^{FB}\right)\right) + v_B\left(\theta_A, b^*\left(\tilde{q}^{FB}\right)\right)\right]}_{\textit{players' joint payoff specific to } \theta_A \textit{ and } \tilde{q}} &= \mathrm{cav}_q\left[\mathrm{E}_q\left[u_A(\theta_A, a^*(q)) + v_A(\theta_B, b^*(q))\right]\right](q). \end{aligned} \tag{6}$$

The optimal signals $\tilde{p}^{FB}$ and $\tilde{q}^{FB}$ are solutions to independent maximization problems, each focused on its own source of uncertainty and how it affects the players' joint payoffs. This will be in sharp contrast to both static and dynamic incentives of the players that we proceed to consider.

4.2 Static Nash Equilibrium

The static Nash equilibrium is obtained by considering the static incentive of each agent to share information given the best responses a^* and b^* conditional on this information. The optimal signals can be characterized by the concavification approach of Kamenica and Gentzkow (2011).

Lemma 1 (Static Nash Equilibrium)**.** *The pair of signals* $\tilde{p}$ *and* $\tilde{q}$ *constitute a static Nash equilibrium if and only if and only if*

$$\begin{aligned} \mathrm{E}_{p_0}\left[v_A\left(\theta_A, b^*\left(\tilde{p}\right)\right)\right] &= \mathrm{cav}_p\left[p \cdot v_A(1, b^*(p)) + (1-p) \cdot v_A(0, b^*(p))\right](p_0), \\ \mathrm{E}_{q_0}\left[v_B\left(\theta_B, a^*\left(\tilde{q}\right)\right)\right] &= \mathrm{cav}_q\left[q \cdot v_B(1, a^*(q)) + (1-q) \cdot v_B(0, a^*(q))\right](q_0). \end{aligned} \tag{7}$$

Moreover, the minimal static Nash equilibrium signals $\tilde{p}^N$ *and* $\tilde{q}^N$ *are the Blackwell least informative solutions to* (7)*.*

The exact nature of the optimal signal $\tilde{p}$ in (7) stems from concavification of $\mathrm{E}_p\left[v_A(\theta_A, b^*(p))\right]$ of player A along belief p, illustrated in Figure 4. When the concavification can be implemented with

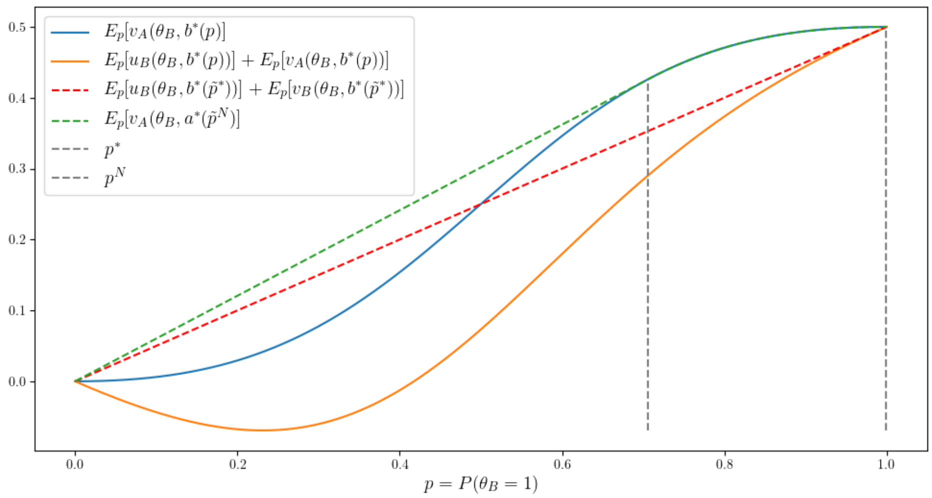

Figure 4: Concavification of player A's value function and the social planner's objective with respect to the belief about $p = \mathrm{P}(\theta_B = 1)$. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$.

multiple supporting points, then the minimal signal $\tilde{p}^N$ picks the implementation of the concave hull that relies on the greatest number of points to minimize the information shared with agent B. Comparing the optimal signals obtained in (7) and the social planner's optimal signals in (6), we see that the static Nash equilibrium signal $\tilde{p}^N$ does not take into account the value of information for player B u_B, beyond the effect of his best response action $b^*(p)$ on A's utility through v_A. This makes the static Nash equilibrium potentially inefficient relative to the social optimum in (5) as a social planner would wish to communicate weakly more information relative to $\tilde{p}^N$, as can be seen in Figure 4 by comparing the blue and orange lines. The optimal static signal $\tilde{q}^N$ chosen by player B is inefficiently opaque for the same reasons.[6]

4.3 Full Information Sharing

Information about state θ_i always benefits player i as it lets him improve his action choice. However, the incentives of the players are misaligned as players differ in the utilities they derive from specific actions. The resulting static Nash equilibrium may be inefficient. We show that by considering dynamic communication, the agents can sustain full information sharing whenever it is ex-ante

[6]This argument applies to all static Nash equilibria, and while we use the minimal static equilibrium to construct the dynamic equilibrium set, we compare the efficiency gains relative to the static Nash equilibrium that generates the highest payoffs to both agents.

efficient and dominates the static Nash payoffs for both agents. This is possible by focusing on communication strategies that maintain sufficient uncertainty about the respective states θ_A and θ_B to ensure that the continuation value of obtaining full information about θ dominates the value from a deviation to the static Nash outcome.

We go on to introduce allocation gain functions $AG_i(\cdot)$ and persuasion loss functions $PL_i(\cdot)$ for each player $i \in \{A, B\}$, which provide a useful perspective on both the efficient outcomes as well as incentive compatibility in the dynamic equilibrium.

Allocation gains. Player A benefits *receiving* full information about θ_A rather than the imperfectly informative signal $\tilde{q}^N$ as it allows him to tailor his action a^* to the state more precisely. We define the allocation gain obtained by player A from receiving full information about θ_A relative to receiving the static Nash signal $\tilde{q}^N$ as

$$AG_A(q) \overset{def}{=} \underbrace{\mathrm{E}_q\left[u_A(\theta_A, a^*(\theta_A))\right]}_{\substack{\textit{utility from own action} \\ \textit{under full information}}} - \underbrace{\mathrm{E}_q\left[u_A\left(\theta_A, a^*\left(\tilde{q}^N\right)\right)\right]}_{\substack{\textit{utility from own action} \\ \textit{under static Nash signal}}}. \tag{8}$$

The expected value to player A under the static Nash outcome is weakly convex in beliefs as more information is beneficial to make decisions absent full information. This results in $AG_A(q)$ being a weakly positive and concave function with $AG_A(0) = AG_A(1) = 0$. Similar to $AG_A(q)$ for player A, define player B's allocation gain $AG_B(p)$ as the benefit of knowing θ_B fully when taking action relative to the static Nash signal $\tilde{p}^N$.

Persuasion losses. Player A loses as a sender from *sharing* all information about θ_B relative to the strategically chosen signal $\tilde{p}^N$. Define the persuasion loss of player A as the difference of expected utilities from sharing a fully informative signal about θ_B relative to the statically optimal signal $\tilde{p}^N$ as

$$PL_A(p) \overset{def}{=} \underbrace{\mathrm{E}_p\left[v_A(\theta_B, b^*(\theta_B))\right]}_{\substack{\textit{utility from other's action} \\ \textit{under full information}}} - \underbrace{\mathrm{E}_p\left[v_A\left(\theta_B, b^*\left(\tilde{p}^N\right)\right)\right]}_{\substack{\textit{utility from other's action} \\ \textit{under static Nash signal}}}. \tag{9}$$

By fully revealing θ_B, player A loses $PL_A(p)$ relative to if he were to unilaterally deviate to strategically static persuasion of player B. By construction, $PL_A(p)$ is weakly negative and convex in p with $PL_A(0) = PL_A(1) = 0$. As the value obtained from persuasion is identical across all static Nash equilibria for the agents, it is without loss to use the minimal static signal $\tilde{p}$ in the definition of $PL_A(p)$. Similar to $PL_A(p)$ for player A, define the persuasion loss $PL_B(q)$ to player B as the loss from fully revealing θ_A relative to sending a statically optimal $\tilde{q}^N$.

Efficiency of full information. From a planner's perspective, full revelation of θ_A is beneficial when player A's allocation gain $AG_A(q)$ exceeds player B's persuasion loss $PL_B(q)$, captured by $AG_A(q) + PL_B(q) \geq 0$. Similarly, full information revelation about θ_B is beneficial if $AG_B(p) + PL_A(p) \geq 0$. However, a stronger condition is necessary for full information to be efficient across all possible signals.

Lemma 2 (Full information efficiency). *Full information disclosure of θ_A and θ_B are solutions to the social planner's problem, i.e., $\tilde{p}^{FB} = \theta_B$ and $\tilde{q}^{FB} = \theta_A$ are solutions to* (5)*, if and only if $AG_B(p) + PL_A(p) \geq 0$ and $AG_A(q) + PL_B(q) \geq 0$ for all $p, q \in [0, 1]$.*

Lemma 2 states that full information dominates any other information sharing rule if the allocation gain is higher than the persuasion loss for both states for all possible priors. If there does exists a p' such that $AG_B(p') + PL_A(p') < 0$, then there exist a signal $\tilde{p}^*$ that would lead to an overall improvement relative to sharing full information. Put differently, even if full information dominates the static Nash outcome at a given p, to be globally optimal across all signals about θ_A and θ_B, it has to dominate it for all starting priors. A simple corollary of Lemma 2 hold is that if its conditions hold, then full information also dominates all static Nash equilibria from the perspective of maximizing the joint welfare of the agents.

Incentive compatibility. Consider now the players' private interests over whether to fully share their information or deviate to the static Nash equilibrium, which constitutes the outside option for both players. Player A finds the full information outcome individually rational as long as his allocation gain $AG_A(q)$ exceeds the persuasion loss $PL_A(p)$, i.e., $AG_A(q) + PL_A(p) \geq 0$. Similarly, full information outcome is individually rational for player B if and only if $AG_B(p) + PL_B(q) \geq 0$.[7]

If $AG_A(q) + PL_A(p) \geq 0$ and $AG_B(p) + PL_B(q) \geq 0$, both players prefer a fully informative outcome to the minimal static Nash equilibrium. However, a fully informative outcome cannot be achieved in a single step. Full information can be sustained as a dynamic equilibrium only if the players continue to prefer full information revelation along the equilibrium path of beliefs $(P_t^*, Q_t^*)_{t\in[0,1]}$.

[7] These individual rationality constraints compare the value of revealing information across prior beliefs p and q about states θ_B and θ_A, unlike the planner who evaluated the value of revealing information across players A and B.

To this end, denote by S^{FI} to be the set of beliefs (p, q) for which the full information is a Pareto improvement over the Static Nash Equilibrium for both players

$$S^{FI} \overset{def}{=} \{(p, q) : \quad AG_A(q) + PL_A(p) \geq 0 \quad \text{and} \quad AG_B(p) + PL_B(q) \geq 0\}. \tag{10}$$

Set S^{FI} captures the set of incentive-compatible beliefs at which both agents prefer a fully informative outcome over the static Nash outcome, and hence they might be compelled to follow an equilibrium that results in full information revelation. The definition of set S^{FI} in (10) is identical to its definition in Section 2 for binary preferences. In that example, set S^{FI} took a star-shape as depicted in Figure 2. For more general preferences, the shape of set S^{FI} can look quite different, and our next step is to formalize the properties of set S^{FI} that are necessary to achieve a fully informative dynamic equilibrium outcome.

For a set $S \subseteq \mathbb{R}^2$ define its top and bottom contours as

$$\begin{aligned} T(S) &\overset{def}{=} \left\{(p, q) \in S \quad \text{such that} \quad \forall (p, q') \in S \quad \Rightarrow \quad q' \leq q\right\}, \\ B(S) &\overset{def}{=} \left\{(p, q) \in S \quad \text{such that} \quad \forall (p, q') \in S \quad \Rightarrow \quad q' \geq q\right\}. \end{aligned} \tag{11}$$

The top contour $T(S)$ of set S consists of points (p, q) in S with the highest q for a given p. Similarly, the bottom contour $B(S)$ of set S consists of points (p, q) in S with the lowest q for a given p. Economically, the top and bottom contours of a set S capture the beliefs at which the value of obtaining full information for agent A is equal to his value in the minimal static Nash equilibrium. Similarly, define the left and right contours of a set S as

$$\begin{aligned} L(S) &\overset{def}{=} \left\{(p, q) \in S \quad \text{such that} \quad \forall (p', q) \in S \quad \Rightarrow \quad p' \geq p\right\}, \\ R(S) &\overset{def}{=} \left\{(p, q) \in S \quad \text{such that} \quad \forall (p', q) \in S \quad \Rightarrow \quad p' \leq p\right\}. \end{aligned} \tag{12}$$

Economically, the left and right contours of a set S capture the beliefs at which the value of obtaining full information for agent B is equal to his value in the static Nash equilibrium.

Definition (Contour convexity)**.** *Set S is **contour convex** if*

- $\forall (p, q) \in T(S) \cup B(S)$ *there exist* $p^L < p < p^R$ *such that* $\left(p^L, q\right) \in L(S)$ *and* $(p^R, q) \in R(S)$*;*
- $\forall (p, q) \in L(S) \cup R(S)$ *there exist* $q^B < q < q^T$ *such that* $(p, q^B) \in B(S)$ *and* $(p, q^T) \in T(S)$*.*

Geometrically, contour convexity states that the upper and lower contours are convex combinations along the horizontal axis p of the left and right contours. Similarly, the left and right contours are convex combinations along the vertical axis q of the top and bottom contours.

Lemma 3 (Contour property)**.** *Set S^{FI} is contour convex if full information sharing is efficient, i.e., if $AG_B(p) + PL_(p) \geq 0$ and $AG_A(q) + PL_B(q) \geq 0$ for all $p, q \in [0, 1]$.*

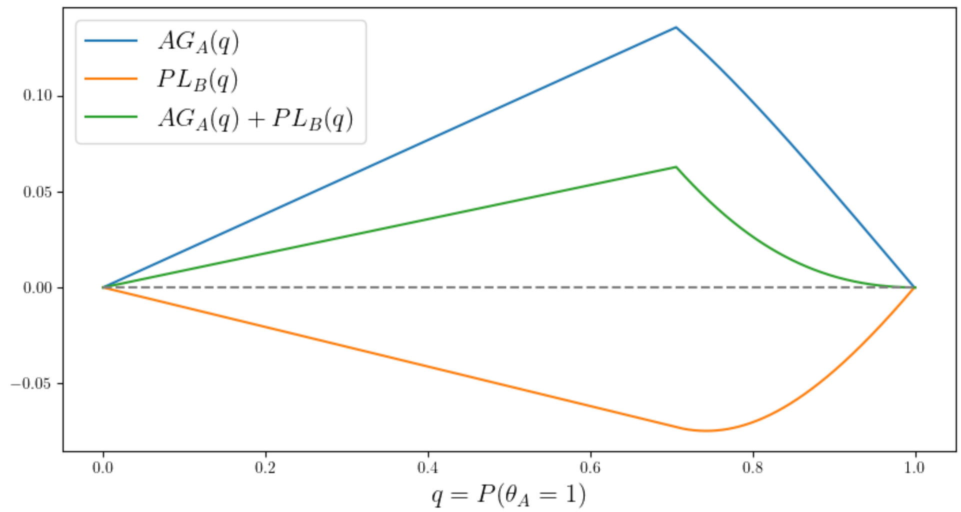

Figure 5: Construction of q^B and q^T for any $(p, q) \in L(S^{FI})$. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$.

Consider, for example, the left contour of the set S^{FI}. For any point (p, q) on the left contour the player B is indifferent between full information and Nash outcomes, i.e, $AG_B(p) + PL_B(q) = 0$. Any communication by player A would change belief p and result in outcomes with posterior $p' < p$ where the allocation gain $AG_B(p')$ for player B is lower than his persuasion loss $PL_B(q)$ precluding him from sharing full information. Hence, any further information revelation is possible in this case if and only if the player B can disclose some information without violating player A's incentive constraint. The latter is possible only if the player A strictly prefers full information sharing to the Nash outcome. When sharing full information is efficient (condition of Proposition 1) we have

$$AG_B(p) + PL_B(q) = 0 \qquad \Rightarrow \qquad AG_A(q) + PL_A(p) > 0. \tag{13}$$

As a result, information disclosure by player B that changes q, can (i) relax his own incentive compatibility constraint by lowering the persuasion loss $PL_B(q)$ and (ii) avoid violating player A's incentive compatibility constraint through simultaneously lowering $AG_A(q)$.

Now that we know that some communication by player B is feasible at the left contour, we can explicitly construct the signal that would take us to the upper and lower contours using the procedure illustrated in Figure 5. Communication by player B changes belief q without affecting belief p. At the upper and lower contours the individual rationality constraint of the player A is binding, hence we can solve for $\tilde{q}$ such that $AG_A(\tilde{q}) = PL_A(p)$. The concavity of allocation gain implies that such an equation would have two solutions $q^B < q^T$. Finally, since $AG_A(q) > PL_B(q) = AG_B(p) > PL_A(p) = AG_A(q^B) = AG_A(q^T)$ and AG_A is quasi-concave it must be that $q \in (q^B, q^T)$.

Proposition 1 (Full information implementation.)**.** *Full information sharing can be sustained as a dynamic equilibrium outcome if it is ex-ante efficient and the starting prior* $(p, q) \in S^{FI}$.

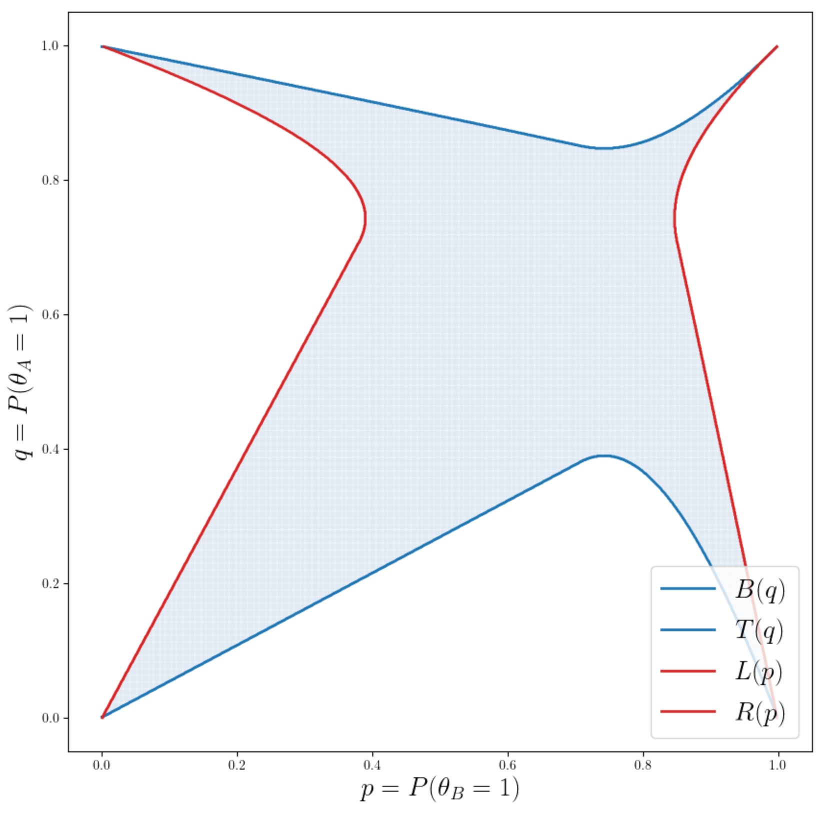

Figure 6: Dynamic Equilibrium that achieves full information outcome. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$, $q_0 = \mathrm{P}(\theta_A = 1) = 0.5$, $p_0 = \mathrm{P}(\theta_B = 1) = 0.5$.

In order to support the full information outcome as a dynamic equilibrium, it should Pareto dominate Static Nash equilibrium not only at the prior (p, q) but everywhere along the (stochastic) equilibrium path of play. Proposition 1 shows that as long as sharing full information is efficient and the starting prior belongs to set S^{FI}, then such a path exists, meaning that the ex-ante individual rationality constraint is sufficient for on-path individual rationality to be satisfied.

The game begins with the starting prior $(p, q) \in S^{FI}$ about (θ_B, θ_A). Denote $P_0 = p$ and $Q_0 = q$. Similar to the example of Section 2, construct the players' strategies sequentially using the contours of the set.

- Time $t = 1/2$. Player A shares a binary signal about θ_B with posteriors $p^L_{1/2}$ and $p^R_{1/2}$ such that the posterior beliefs fall on either the left- or right- contours of set S^{FI}, i.e., that $\left(p^L_{1/2}, Q_0\right) \in L(S^{FI})$ and $\left(p^R_{1/2}, Q_0\right) \in R(S^{FI})$. Since the initial belief belongs to set S^{FI}, then such signals are well defined, and belief processes (P_t, Q_t) are equal to (P_0, Q_0) for $t \in [0, 1/2)$ and

$$P_{1/2} = \begin{cases} p^R_{1/2} & \textit{with probability} \quad \dfrac{P_0 - p^L_{1/2}}{p^R_{1/2} - p^L_{1/2}}, \\ p^L_{1/2} & \textit{with probability} \quad \dfrac{p^R_{1/2} - P_0}{p^R_{1/2} - p^L_{1/2}}, \end{cases} \qquad Q_{1/2} = Q_0. \tag{14}$$

  Signal (14) is illustrated as an orange dashed lines in Figure 6 as the first step of communication at $t = 1/2$.

- Time $t = 3/4$. Player B shares a signal that depends on the first round of communication. He shares a binary signal about θ_A with posteriors $q^T_{3/4}$ and $q^B_{3/4}$ that fall either on the top or bottom contour of set S^{FI}, i.e., $\left(P_{3/4}, q^T_{3/4}\right) \in T(S^{FI})$ and $\left(P_{3/4}, q^B_{3/4}\right) \in B(S^{FI})$. Such signals exist since set S^{FI} satisfies 4.3. Belief process (P_t, Q_t) is then equal to $(P_{1/2}, Q_{1/2})$ for $t \in [1/2, 3/4)$ and

$$P_{3/4} = P_{1/2}, \qquad Q_{3/4} = \begin{cases} q^T_{3/4} & \textit{with probability} \quad \dfrac{Q_{1/2} - q^B_{3/4}}{q^T_{3/4} - q^B_{3/4}}, \\ q^B_{3/4} & \textit{with probability} \quad \dfrac{q^T_{3/4} - Q_{1/2}}{q^T_{3/4} - q^B_{3/4}}. \end{cases} \tag{15}$$

  Signal $Q_{3/4}$ is illustrated as a green dashed line in Figure 6.

- Time $t = 1 - \frac{1}{2^k}$. If k is odd, it is player A's turn to communicate, while if k is even, then it is player B's turn to communicate. Player A's signals bring the posterior belief to the left and right contours of set S^{FI}, while player B's signals bring the posterior belief to the top and bottom contours of set S^{FI}.

Such iterative progression gradually expands the posterior belief towards the four corners of set S^{FI}, which then correspond to degenerate priors $\lim_{t \to 1} \mathrm{E}\left[\theta_B \mid P_t\right] \overset{P-a.s.}{=} \theta_B$ and $\lim_{t \to 1} \mathrm{E}\left[\theta_A \mid Q_t\right] \overset{P-a.s.}{=} \theta_A$. This specifies candidate belief processes $(P_t^*, Q_t^*)_{t \in [0,1)}$ along the equilibrium path. The fact that beliefs stay within set S^{FI} ensures that their value of a fully informative outcome (continuation play) Pareto dominates the expected payoff of the minimal static Nash equilibrium (deviation payoff), thus ensuring that intertemporal individual rationality constraints are satisfied along the equilibrium path.

4.4 Discussion

An important feature of the construction is that, despite information arriving unpredictably, there is always sufficient uncertainty about both states θ_A and θ_B in order to make the players be willing to trade their information with one another, rather than deviate to the static Nash outcome. Moreover, the actions of the agents in our construction are complementary – communication by agent i ensures and reinforces agent j's incentives to continue the communication process. Such strategic and dynamic management of opacity allows the players to attain the welfare-maximizing optimum. We show in the next section that this argument extends to the entire Pareto frontier, meaning that it is possible to implement all centralized payoffs via dynamic equilibria and characterize the set of all equilibrium payoffs.

It is important to note that the dynamic equilibrium we construct in Sections 2 and 4 is not unique. Our approach was to structure communication in the most intuitive way to occur on dates $t = 1 - \frac{1}{2^k}$, but we could have, equivalently, appealed to signals that arrive at random times or via diffusion processes. The intuition behind such alternative constructions that implement full information remains the same, however – the belief path must remain within set S^{FI}, and, thus, rely on jumps in beliefs to preserve the approximate symmetry in residual uncertainties about θ_A and θ_B.

5 Equilibrium Payoff Set and Dynamic Implementation

Full information maximizes welfare whenever the conditions of Lemma 2 are satisfied. Proposition 1 states that as long as players prefer this socially efficient, fully informative outcome to the minimal static Nash equilibrium, the worst possible equilibrium of the game, it can be supported as a dynamic equilibrium outcome. If, however, the initial prior $(p, q) \notin S^{FI}$, the ex-ante individual rationality constraints of one of the agents are binding and full information cannot be implemented even if it is ex-ante efficient. In this section, we establish two results. First, we characterize the constrained efficient dynamic equilibrium outcome for priors $(p, q) \notin S^{FI}$, i.e., when efficient full information sharing cannot be implemented in equilibrium. Second, we show that even if full information is not ex-ante efficient, i.e., if conditions of Lemma 2 are not satisfied, the signals that are efficient and that dominate the static Nash outcome can still be supported as a dynamic equilibrium outcome. These two steps permit the characterization of all attainable dynamic equilibrium payoffs in the model. We then build upon these two results and provide comparative statics of how players' payoffs respond to changes in preference parameters and all ex-ante priors in constrained-efficient equilibria.

5.1 Information Pareto Frontier

Consider the set of feasible payoffs corresponding to signals $\tilde{p}$ and $\tilde{q}$ whilst the agents choose their contingent actions $a^*(\tilde{q})$ and $b^*(\tilde{p})$. In addition to welfare maximizing signals, as defined in Section 4.1, consider payoffs that are not dominated by payoffs under alternative signals and, thus, constitute the Pareto frontier.

Definition (Pareto frontier)**.** *A pair of player utilities (u_A, u_B) lies on the Pareto frontier if there exists a signal $(\tilde{p}^*, \tilde{q}^*)$ that generates expected utilities (u_A, u_B) for the players A and B and there do not exist signals $(\tilde{p}, \tilde{q})$ which generate strictly higher payoffs for both players.*

Similar to Section 4.1, it is without loss to consider independent static signals $\tilde{p}$ and $\tilde{q}$, rather than dynamic belief processes (P_t, Q_t), due to the separability of the agents' preferences across states. Consequently, the Pareto frontier is an upper boundary on the set of equilibrium payoffs that can be attained by the players in any dynamic equilibrium.

To obtain the signals that constitute the extreme points of the Pareto frontier, we can use the tangency approach: a pair of utilities (u_A, u_B) lies on the Pareto frontier if there exists a weight

$w \in [0, 1]$ such that $w \cdot u_A + (1 - w) \cdot u_B$ is equal to

$$\max_{\tilde{p},\tilde{q}} \Big\{ w \cdot \underbrace{\mathrm{E}\left[u_A(\theta_A, a^*(\tilde{q})) + v_A(\theta_B, b^*(\tilde{p}))\right]}_{\textit{player A's expected utility}} + (1 - w) \cdot \underbrace{\mathrm{E}\left[u_B(\theta_B, b^*(\tilde{p})) + v_B(\theta_A, a^*(\tilde{q}))\right]}_{\textit{player B's expected utility}} \Big\}. \quad (16)$$

The separability of (16) in states θ_A and θ_B and signals $\tilde{q}$ and $\tilde{p}$ respectively provides a direct way to compute signals $\tilde{p}^w$ and $\tilde{q}^w$ that implement payoffs (u_A, u_B) by adjusting (6) for the unequal weighting of the players' expected utilities. Any public randomization required to implement the non-extreme points of the Pareto frontier can be encoded into the signals themselves and, consequently, all points on the Pareto frontier can be implemented by a pair of signals $(\tilde{p}, \tilde{q})$ and not requiring additional randomization.

5.2 Implementation of the Pareto Frontier in a Dynamic Equilibrium

The Pareto frontier described in Section 5.1 provides an upper bound on the set of attainable payoffs in any dynamic equilibrium. The following proposition shows that the converse is also true as long as the agents find it ex-ante individually rational relative to the minimal static Nash equilibrium.

Proposition 2 (Pareto frontier implementation). *A payoff pair (u_A, u_B) on the Pareto frontier, and the corresponding signals $(\tilde{p}, \tilde{q})$, constitutes a dynamic equilibrium outcome if and only if it dominates the players' payoffs from the minimal static Nash equilibrium, i.e., $(u_A, u_B) \geq (u_A^N, u_B^N)$.*

It is clear that for a pair of utilities to be consistent with an equilibrium outcome, it is necessary that they dominate the worst payoff that the agents can attain in any equilibrium. That lower bound is equal to the minimal static Nash equilibrium as described in Section 4.2. In what follows, we show that it is also sufficient. To this end, fix a pair of signals $\tilde{p}^*$ and $\tilde{q}^*$ that implement a pair of utilities on the Pareto frontier.

Define the allocation gain for player A of signal $\tilde{q}^*$ relative to $\tilde{q}^N$ is, similar to (8), as

$$AG_A^*(q) \stackrel{def}{=} \mathrm{E}_q\left[u_A(\theta_A, a^*(\tilde{q}^*))\right] - \mathrm{E}_q\left[u_A(\theta_A, a^*(\tilde{q}^N))\right].$$

Similar to $AG_A^*(q)$, define $AG_B^*(p)$ to be the allocation gain to player B from signal $\tilde{p}^*$ relative to the static Nash signal $\tilde{p}^N$. Signals $\tilde{p}$ and $\tilde{q}$ are weakly more informative than the static Nash signals $\tilde{p}^N$ and $\tilde{q}^N$ since the social planner always internalizes the receiver's value at least to some degree.

For this reason, the allocation gain functions $AG_A^*(q)$ and $AG_B^*(p)$ are weakly concave whenever they are positive.

Define persuasion loss to player A from communicating signal $\tilde{p}^*$ over $\tilde{p}^N$, similar to (9), as

$$PL_A^*(p) \overset{def}{=} \mathrm{E}_p\left[v_A(\theta_B, b^*(\tilde{p}^*))\right] - \mathrm{E}_p\left[v_A\left(\theta_B, b^*\left(\tilde{p}^N\right)\right)\right].$$

The persuasion loss $PL_B^*(q)$ to player B from signal $\tilde{q}^*$ relative to q^N is defined similarly.

Lemma 4 (Efficient signals)**.** *Signals $\tilde{p}^*$ and $\tilde{q}^*$ lie on the Pareto frontier if and only if there exists a weight $w \in [0,1]$ such that $w \cdot AG_A^*(q) \geq (1-w) \cdot PL_B^*(q)$ and $(1-w) \cdot AG_B^*(p) \geq w \cdot PL_A^*(p)$ for all $q, p \in [0,1]$.*

Lemma 2 showed that allocation gains must globally dominate persuasion losses for full information to be a welfare maximizing outcome. Lemma 4 substantially expands this earlier result by showing that this ranking, appropriately weighted, is always satisfied for efficient signals. We see that efficiency, again, compares the gains and losses across agents, even though the players have an asymmetric weight now.

For signals $\tilde{p}^*$ and $\tilde{q}^*$ to be an individually rational equilibrium outcome for player A his allocation gain from signal $\tilde{q}^*$ must exceed his persuasion loss under signal $\tilde{p}^*$, i.e., $AG_A^*(q) + PL_A^*(p) \geq 0$. Similarly, for player B to find these signals individually rational it must be the case that $AG_B^*(p) + PL_B^*(q) \geq 0$. Similar to set S^{FI}, denote set S^* to be the set of prior beliefs for which signal $(\tilde{p}^*, \tilde{q}^*)$ dominates the minimal static Nash payoff:

$$S^* \overset{def}{=} \{(p,q): \quad AG_A^*(q) + PL_A^*(p) \geq 0 \quad \text{and} \quad AG_B^*(p) + PL_B^*(q) \geq 0\}.$$

Set S^* is illustrated in Figure 7 is different from set S^{FI} as signals $\tilde{p}^*$ and $\tilde{q}^*$ are not fully informative. Similar to the previous argument, however, intertemporal individual rationality for players A and B requires that for an equilibrium $(P_t^*, Q_t^*)_{t\in[0,1]}$ to implement signals $\tilde{p}^*$ and $\tilde{q}^*$ it must be the case $(P_t^*, Q_t^*) \in S^*$ for all t. The following Lemma provides the final step of the argument underlying Proposition 2.

Lemma 5. *If signals $\tilde{p}^*$ and $\tilde{q}^*$ are Pareto-efficient, then set S^* is contour convex.*

Analogously to Section 4 Pareto efficiency of signals $\tilde{p}^*$ and $\tilde{q}^*$ implies that only one's player incentive constraint can be tight at the boundary of the set S^* (with the exception of four corners).

This property ensures that the player whose incentive constraint is tight can continue the path of information revelation without violating the counterparty's incentive constraint.

Contour convexity of set S^* allows us to construct the dynamic equilibrium in the same way we did in Sections 2 or 4.3: the players take alternate turns and shift beliefs between the contours of set S^* until its extreme points are reached in the limit at $t \to 1$.

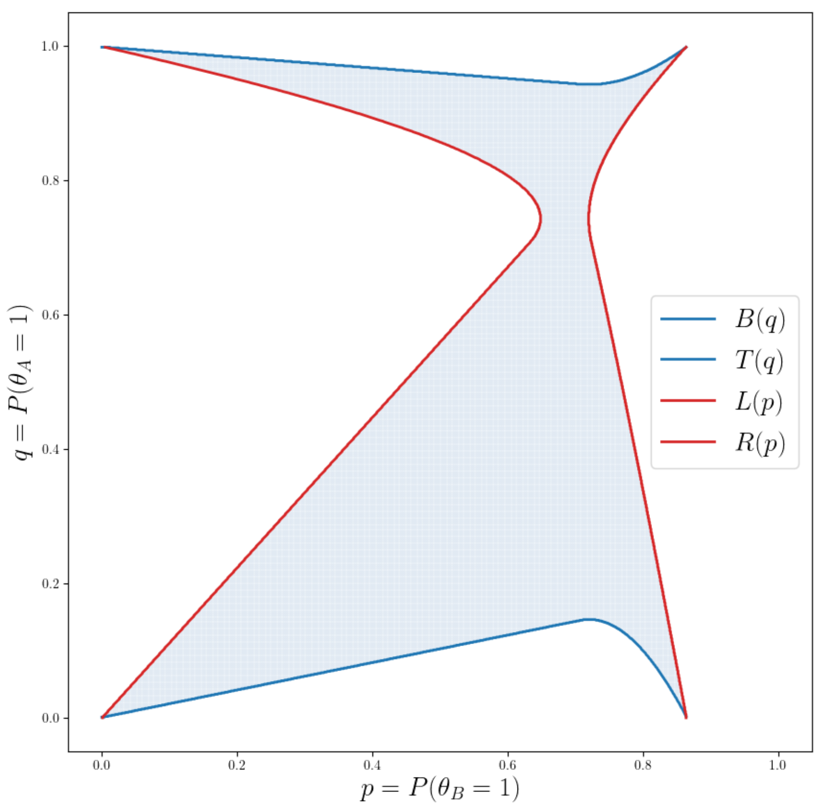

Figure 7: Set S^* and the first three communication steps implementing signal $(\tilde{p}, \theta_A)$ as a dynamic equilibrium. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$. Priors $q_0 = \mathrm{P}(\theta_A) = 0.1$, $p_0 = \mathrm{P}(\theta_B = 1) = 0.8$. Constrained efficient signals put weight $w = 0.6$ on player A and are given by binary signals $\tilde{p}^* \in \{0, 0.86\}$, $\tilde{q}^* \in \{0, 1\}$.

The shape of set S^* is illustrated in Figure 7. We can see that the contours of set S^* satisfy the contour convexity property 4.3. Using these properties, we now construct the dynamic equilibrium that supports signals $\tilde{p}$ and $\tilde{q} = \theta_A$.

5.3 Dynamic Equilibrium Payoff Set

The characterization of the Pareto frontier enables us to generate a complete characterization of the set of payoffs attainable under dynamic equilibria.

Corollary 1 (Equilibrium set)**.** *The set of payoffs that can be attained in dynamic equilibria is equal to payoffs that exceed the minimal static Nash outcome but are weakly dominated by the Pareto frontier.*

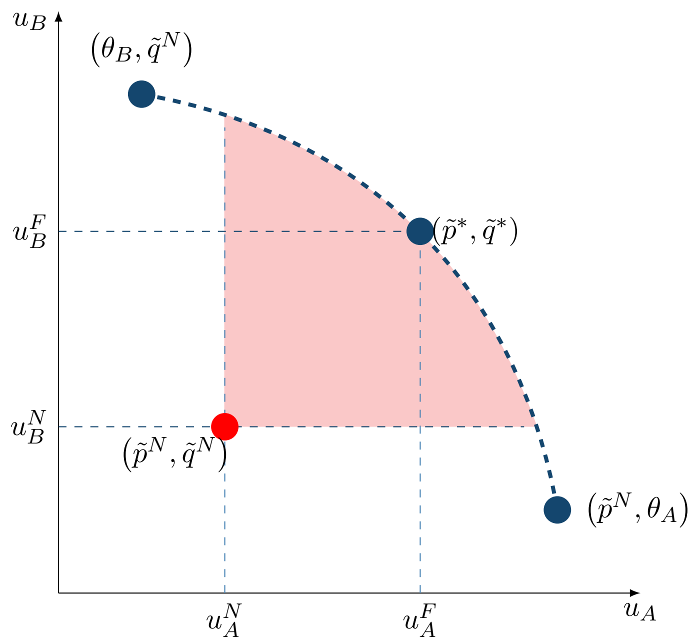

Figure 8: The set of feasible payoffs (in red) is the set of all possible payoffs that dominate the minimal static Nash equilibrium. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$. Priors $p = 0.5$, $q = 0.5$.

The payoff in any equilibrium has to weakly exceed the payoff under the static Nash outcome, as follows from Lemma 1. In addition, Corollary 1 states that it is sufficient to consider points that lie below the Pareto frontier. The intuition is that the corners of the Pareto frontier that the top left corner of the Pareto frontier is associated with signal $(\theta_B, \tilde{q}^N)$ that is dominated by the static Nash outcome for player A. Similarly, the bottom right corner of the Pareto frontier is implemented with signal $(\tilde{p}^N, \theta_A)$ and is dominated by the minimal static Nash outcome for player B. Consequently, the Pareto frontier is the relevant upper boundary on the set of payoffs attainable as dynamic equilibrium outcomes. The implementation of a specific point in the feasible set is an outcome of public randomization between the static Nash equilibrium and an equilibrium on the Pareto frontier, which we have constructed during the proof of Proposition 2.

Definition. *A dynamic equilibrium* $(P_t^*, Q_t^*)_{t\geq 0}$ *is* ***constrained efficient*** *if it maximizes the sum of expected payoffs of players A and B among all other dynamic equilibria.*

If the socially efficient outcome dominates the static Nash payoff for both agents, then the constrained efficient equilibrium implements the social planner's preferred outcome. If, however, the socially efficient outcome cannot be implemented as an equilibrium, the constrained efficient equilibrium is the best dynamic equilibrium from the social planner's perspective. By considering the constrained efficient equilibria, we can understand what happens in the model as we change model parameters and exit the regions when the first-best allocation is feasible.

5.4 Comparative Statics

In this section, we compare the players' expected utilities under the socially efficient outcome, the static Nash outcome, and the constrained efficient equilibrium outcome. We show that, to attain efficiency, there has to be sufficient uncertainty about the states and the private benefits from persuasion cannot be too large.

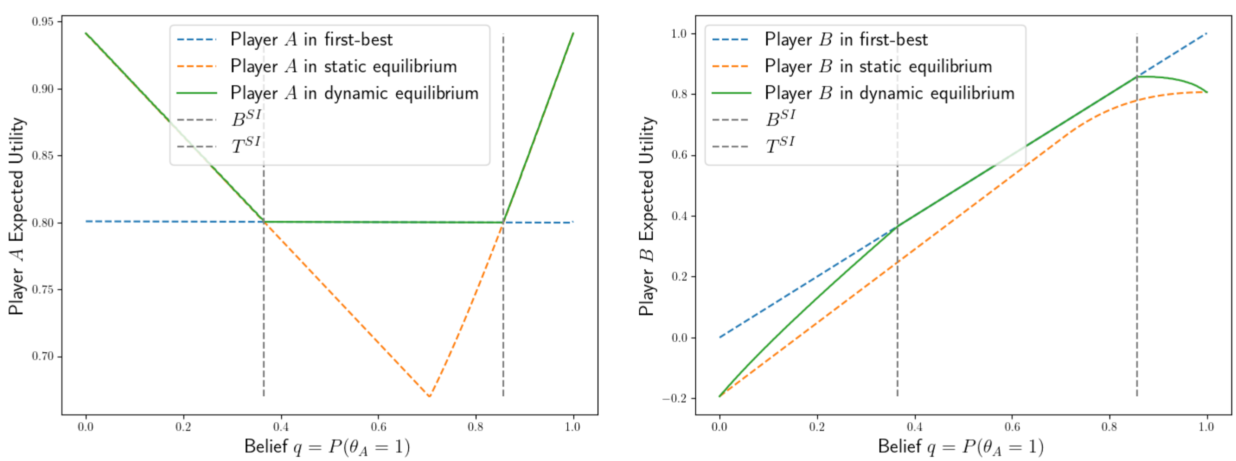

(a) Player A expected payoff.

(b) Player B expected payoff.

Figure 9: Expected welfare of the players in the most efficient dynamic equilibrium relative to static Nash and first best as a function of the starting prior $q = \mathrm{P}(\theta_B = 1)$. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$, $p_0 = \mathrm{P}(\theta_B = 1) = 0.8$.

Figure 9 considers the implications of ex-ante uncertainty about θ_A, captured by varying the prior belief $q = \mathrm{P}(\theta_A = 1)$. The socially efficient signal involves full transparency. When belief q is very low or very high, however, there is very little residual uncertainty about θ_A. For these beliefs, player

A's allocation gain from learning θ_A is lower than the persuasion loss associated with disclosing full information about θ_B. As we can see in Figure 9a, the static Nash outcome dominates the efficient one for player A. The constrained efficient equilibrium takes this into account and selects signals $\tilde{p}^*$ and $\tilde{q}^*$ that leave player A with the same expected value as his static Nash outcome. They are, however, strictly more efficient for player B, as can be seen in Figure 9b – player B's expected payoff is strictly higher than his static Nash outcome for all $q \in [0, 1]$. It is also, however, also strictly lower than his payoff in the socially optimal outcome when beliefs q are either very high or very low. When beliefs about q are intermediate, then player A is willing to trade the information he has about θ_B in exchange for learning about θ_A, resulting in socially optimal payoffs for both agents.

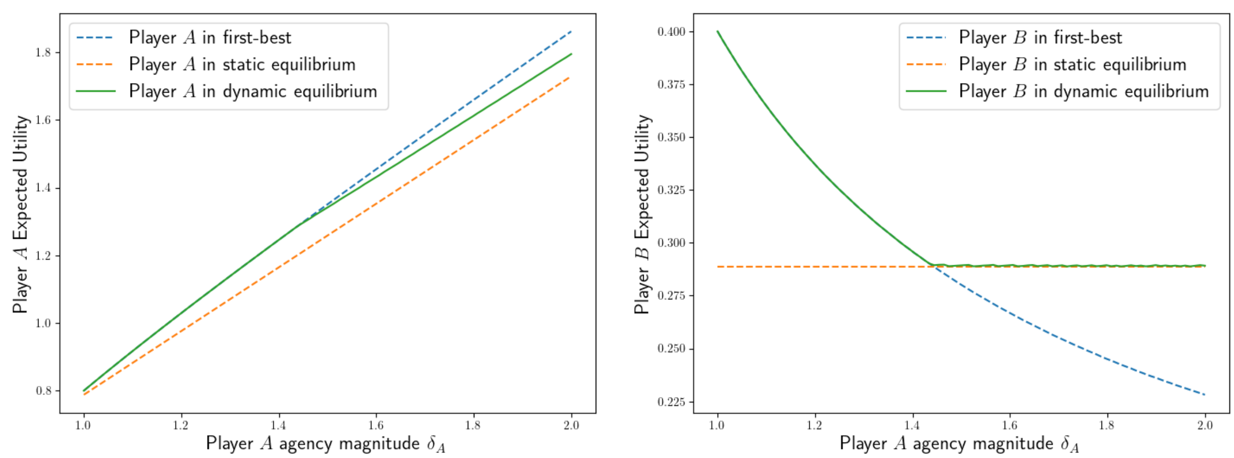

(a) Player A expected payoff.

(b) Player B expected payoff.

Figure 10: Expected welfare of the players in the most efficient dynamic equilibrium relative to static Nash and first best as a function of player A's private bias magnitude δ_A. Parameters: $u_A(\theta_A, a) = (\theta_A - a)^{3/2}$, $v_A(\theta_B, b) = \delta_A \cdot b$, $u_B(\theta_B, b) = (\theta_B - b)^{3/2}$, $v_B(\theta_A, a) = a$, $p_0 = \text{P}(\theta_B = 1) = 0.8$, $q = \text{P}(\theta_A = 1) = 0.4$.

Figure 10 considers the implications of the private benefits of player A on the efficiency of the equilibrium outcome. We do so by scaling up the private benefit of agent A by the parameter δ_A. On the one hand, an increase in δ_A increases the social weight on player A's preferences, making the socially optimal signal more opaque, favoring that player. On the other hand, such opacity reduces player B's allocation gains and results in his static Nash outside option becoming binding – distorting the optimal signals away from the social optimum. We see this play out in Figures 10a and 10b. For low values of δ_A, the constrained efficient equilibrium outcome coincides with the first best one. As the private benefits of player A increase, however, the individual rationality constrained of player B becomes binding, resulting in a welfare distortion.

5.5 Model Extension

For clarity of exposition, we specify the preferences of the agents as contingent on the underlying state and their actions. We can generalize this assumption by specifying their expected preferences as $U_A = u_A(q) + v_A(p)$ and $U_B = u_B(p) + v_B(q)$. By doing so, we can capture situations in which the players face not a decision problem but a game conditional on each state. This leads to only one difference in the analysis – the social optimum could, in principle, be more opaque than the static Nash equilibrium. Our analysis goes through, however. Any Pareto efficient allocation that dominates the static Nash outcome and is weakly more informative than it in the Blackwell sense can be supported in the dynamic equilibrium.

6 Conclusion

We consider the model of information sharing between two players who have information valuable to their counterparty and care about their counterparty's actions. In this sense, the players act as both senders and receivers of information. We show that the static Nash equilibrium focuses on maximizing the private benefits to the information sender and does not take into account the allocation gains accrued to the receiver. This is socially inefficient. Dynamic communication can restore most, if not all, benefits of a welfare-maximizing information structure by communicating information gradually. The equilibrium requires that players maintain sufficient uncertainty about the information they hold as to keep their counterparty engaged. Surprisingly, this is enough to implement the socially efficient outcome whenever it dominates the static Nash payoff for both players at the start of the game.

Our continuous time model allows players to communicate infinitely many steps – which is, indeed, used in equilibrium. Our qualitative results continue to hold if we allow for N rounds of communication, with an additional round allowed if either player has deviated. While such a process will not yield an efficient signal, the marginal gains decrease over time. This means that the constrained efficient equilibrium from this finite horizon model would converge to that of an infinite horizon model. While our benchmark model does not allow for monetary transfers between the players.

This is motivated by the fact that, in many instances of corporate information sharing, transfers are scrutinized by regulators as they may lead to collusion in the marketplace. Our model shows that informational collaborations are still feasible even in the absence of transfers.

References

Yeon-Koo Che and József Sákovics. A dynamic theory of holdup. *Econometrica*, 72(4):1063–1103, 2004.

Jeffrey C Ely and Martin Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2): 468–506, 2020.

Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2017.

Jerry Green and Nancy Stokey. *Two representations of information structures and their comparisons*. Number 271. Institute for Mathematical Studies in the Social Sciences, 1978.

Johannes Hörner and Andrzej Skrzypacz. Selling information. *Journal of Political Economy*, 124 (6):1515–1562, 2016.

Nicolas Inostroza. Persuading multiple audiences: An information design approach to banking regulation. *Available at SSRN 3450981*, 2019.

Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11: 249–272, 2019.

Emir Kamenica and Matthew Gentzkow. Bayesian Persuasion. *American Economic Review*, 101 (6):2590–2615, October 2011.

Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.

Fei Li and Peter Norman. Sequential persuasion. *Theoretical Economics*, 16(2):639–675, 2021.

Xuelin Li, Martin Szydlowski, and Fangyuan Yu. Hype cycles: Dynamic information design with two audiences. *Available at SSRN 3923908*, 2021.

Dmitry Orlov, Andrzej Skrzypacz, and Pavel Zryumov. Persuading the principal to wait. *Journal of Political Economy*, 128(7):2542–2578, 2020.

Dmitry Orlov, Pavel Zryumov, and Andrzej Skrzypacz. Design of macro-prudential stress tests. *Review of Financial Studies*, forthcoming.

A Appendix

A.1 Proof of Lemma 3

This is a special case of Lemma 5 for the case when $\tilde{p}^* = \theta_B$ and $\tilde{q}^* = \theta_A$.

A.2 Proof of Proposition 1

Consider the dynamic belief process $(P_t, Q_t)_{t\in[0,1]}$ described in Section 4.3 following Proposition 1. Players communicate on dates $t = 1 - \frac{1}{2^k}$ with player A communicating when k is odd, and player B communicating when k is even. At each step of the process, the sender sends a binary signal – player A sends a signal $P_t \in \{p_t^L, p_t^R\}$ (the left and right contours of set S^*) and player B sends a signal $Q_t \in \{q_t^B, q_t^T\}$ (the bottom and top contours of set S^*).

Lemma A.1 (Monotone informativeness)**.** *The signals are becoming uniformly weakly more informative with each step. If if k is odd then*

$$p^L_{1-\frac{1}{2^{k+2}}} \leq p^L_{1-\frac{1}{2^k}} < p^R_{1-\frac{1}{2^k}} \leq p^R_{1-\frac{1}{2^{k+2}}}.$$

If k is even then

$$q^B_{1-\frac{1}{2^{k+2}}} \leq q^B_{1-\frac{1}{2^k}} < q^T_{1-\frac{1}{2^k}} \leq q^T_{1-\frac{1}{2^{k+2}}}.$$

Proof. Consider the first point without loss. Contour convexity requires that $L(T(p)) \leq p \leq R(T(p))$ and $L(B(p)) \leq p \leq R(B(p))$. □

Following Lemma A.1, it follows that belief process $P_t \in \{p_t^L, p_t^R\}_{t\geq 0}$ with p_t^L weakly decreasing and p_t^R weakly increasing. Similarly $Q_t \in \{q_t^L, q_t^R\}$ with q_t^L weakly decreasing and q_t^R weakly increasing. Define

$$p^L = \lim_{t\to\infty} p_t^L, \qquad p^R = \lim_{t\to\infty} p_t^R, \qquad q^B = \lim_{t\to\infty} q_t^B, \qquad q^T = \lim_{t\to\infty} q_t^T.$$

Suppose that $\mathrm{P}\big((p^L, q^B) > 0\big) > 0$. Due to Lemma 4 and Lemma 5, it implies that allocation gains are strictly higher than persuasion losses, implying that S^{FI} has a strictly positive interior. This implies that if $(p^L, q^B) > (0,0)$, then one of the players will disclose a positive amount of information. This poses a contradiction with the fact that (p^L, q^B) is a limiting point of the belief process.

A.3 Proof of Lemma 5

Suppose signals $\tilde{p}^*$ and $\tilde{q}^*$ are on the Pareto frontier. Following Lemma 4 there exists a weight $w \in [0,1]$ such that signals $\tilde{p}^*$ and $\tilde{q}^*$ maximize (16). The separability of preferences in states allows this to be rewritten as

$$\begin{aligned} \tilde{q}^* &= \arg\max_{\tilde{q}} \Big\{ w \cdot \mathrm{E}\left[u_A(\theta_A, a^*(\tilde{q}))\right] + (1-w) \cdot \mathrm{E}\left[v_B(\theta_A, a^*(\tilde{q}))\right] \Big\}, \\ \tilde{p}^* &= \arg\max_{\tilde{p}} \Big\{ w \cdot \mathrm{E}\left[v_A(\theta_B, b^*(\tilde{p}))\right] + (1-w) \cdot \mathrm{E}\left[u_B(\theta_B, b^*(\tilde{p}))\right] \Big\}. \end{aligned} \tag{A.1}$$

Following Kamenica and Gentzkow (2011), the solution to (A.1) can be obtained by static concavification. This means that for a starting prior (p_0, q_0), the optimal signals take at most two values $\tilde{p}^* \in \{p_L, p_H\}$ and $\tilde{q}^* \in \{q_L, q_H\}$, with $p_L \le p_0 \le p_H$ and $q_L \le q_0 \le q_H$. These signals may not be unique as the value function may allow some indifference for the agents. However, similar to the minimal static Nash equilibrium, we can pick minimal signals $\tilde{p}^*$ and $\tilde{q}^*$ that minimize $p_H - p_L$ and $q_H - q_L$ among all solutions to (A.1).

Lemma A.2. *Signals $\tilde{p}^*$ and $\tilde{q}^*$ are weakly more informative than static Nash signals $\tilde{p}^N$ and $\tilde{q}^N$.*

Proof. When choosing signal $\tilde{q}$, the social planner puts a weakly positive weight w on the allocation of player A, which is strictly convex in beliefs. Consequently, the private objective of the planers is strictly more concave, and hence opaque, compare to the objective of the planner. □

Lemma A.3. *If $\tilde{q}^* = q^N$, then it must be the case that $\tilde{p}^* = p^N$.*

Proof. If $\tilde{q}^* = q^N$, then signal $\tilde{p}^*$ does not violate the individual rationality constraint of player A if and only if it maximizes their private benefits. Since $\tilde{p}^*$ and $\tilde{p}^N$ are both minimal signals, it requires that $\tilde{p}^* = \tilde{p}^N$. □

Lemma A.4 (Strict ranking of gains and losses)**.** *If the (minimal) Pareto efficient signal $\tilde{q}^* \neq \tilde{q}^N$ and $\tilde{p}^* \neq \tilde{p}^N$, then $w \cdot AG^*_A(q) + (1-w) \cdot PL^*_B(q) > 0$ for all $q \in (q_L, q_H)$ and $(1-w) \cdot AG^*_B(p) + w \cdot PL^*_A(p) > 0$ for all $p \in (p_L, p_H)$*

Proof. Suppose there exist a $q_0 \in (q_L, q_H)$ such that $w \cdot AG_A^*(q) + (1-w) \cdot PL_B^*(q) \leq 0$. Using the definition of allocation gains and persuasion losses this implies

$$w \cdot \Big[\mathrm{E}_q\left[u_A(\theta_A, a^*(\tilde{q}^*))\right] - \mathrm{E}_q\left[u_A(\theta_A, a^*(\tilde{q}^N))\right]\Big] + (1-w) \cdot \Big[\mathrm{E}_q\left[v_B(\theta_A, a^*(\tilde{q}^*))\right] - \mathrm{E}_q\left[v_B(\theta_A, a^*(\tilde{q}^N))\right]\Big] \leq 0.$$

Rewriting this condition, obtain that

$$w \cdot \mathrm{E}_q\left[u_A(\theta_A, a^*(\tilde{q}^*))\right] + (1-w) \cdot \mathrm{E}_q\left[v_B(\theta_A, a^*(\tilde{q}^*))\right] \overset{(i)}{\leq} w \cdot \mathrm{E}_q\left[u_A(\theta_A, a^*(\tilde{q}^N))\right] + (1-w) \cdot \mathrm{E}_q\left[v_B(\theta_A, a^*(\tilde{q}^N))\right].$$

If inequality (i) is strict, then it contradicts the optimality of signal $\tilde{q}^*$. If inequality (i) is an equality, then this either implies that $\tilde{q}^*$ is not a minimal signal, or that the static Nash signal is socially efficient, i.e., that $\tilde{q}^* = \tilde{q}^N$. □

For signals $\tilde{p}^*$ and $\tilde{q}^*$ the incentive compatibility set S^* is defined as

$$S^* \overset{def}{=} \{(p,q) : \quad AG_A^*(q) + PL_A^*(p) \geq 0 \quad \text{and} \quad AG_B^*(p) + PL_B^*(q) \geq 0\}.$$

If $(\tilde{p}^*, \tilde{q}^*) = (\tilde{p}^N, \tilde{q}^N)$, then $S^* = [p_L, p_H] \times [q_L, q_H]$, which is contour convex.

Consider now the more interesting case when $\tilde{p}^* \neq \tilde{p}^N$ and $\tilde{q}^* \neq \tilde{q}^N$. The concavity of $AG_A^*(q)$ implies that

$$AG_A^*(q) + PL_A^*(p) \geq 0 \qquad \Leftrightarrow \qquad q \in [B(p), T(p)],$$

where $T(p)$ and $B(p)$ are the top and bottom contours of set S^* evaluated at point p. Similarly,

$$AG_B^*(p) + PL_B^*(q) \geq 0 \qquad \Leftrightarrow \qquad p \in [L(q), R(q)],$$

where $L(q)$ and $R(q)$ are the left and right contours of set S^* evaluated at point q. Suppose there exists a point p when the top contour falls below the left contour. Stated formally, this requires that there exists a belief (p,q) such that $q \geq T(p)$ while $p \leq L(q)$. Inequality $q \geq T(p)$ implies that $AG_A^*(q) + PL_A^*(p) \leq 0$ and inequality $p \leq L(q)$ implies that $AG_B^*(p) + PL_B^*(q) \leq 0$. Consider then the following set of inequalities

$$\begin{aligned} 0 &\overset{(i)}{<} w \cdot AG_A^*(q) + (1-w) \cdot PL_B^*(q) \overset{(ii)}{\leq} -w \cdot PL_A^*(p) + (1-w) \cdot PL_B^*(q) \\ &\overset{(iii)}{\leq} -w \cdot PL_A^*(p) - (1-w) \cdot AG_B^*(p) \overset{(iv)}{<} 0, \end{aligned} \tag{A.2}$$

which presents a contradiction. Strict inequalities (i) and (iv) in (A.2) follows from Lemma A.4. Weak inequality (ii) follows from $AG^*_A(q) \leq -PL^*_A(p)$ for the conjectured prior and inequality (iii) follows from $PL^*_B(q) \leq -AG^*_B(p)$. Consequently, it must be the case that the top contour of set S^* is always higher (along the q axis) than the left contour. Similarly, the bottom contour is always below (along the q axis) of the left contour, proving the contour convexity property for the left contour of set S^*. The other contours are tackled similarly.