BME Faculty of Economic and Social Sciences

Budapest University of Technology and Economics

QSMS DISCUSSION PAPERS

Kemal Kivanç Aköz & Arseniy Samsonov

Bargaining over Information Structures

QSMS-DP-23-001

Budapest, January 2023

Research Centre in Quantitative Social and Management Sciences
Department of Economics and Social Sciences (GTK)
Budapest University of Technology and Economics (BME)

Bargaining over Information Structures*

Kemal Kıvanç Aköz† Arseniy Samsonov‡

January 5, 2023

Abstract

How transparent are informational institutions if their founders have to agree on the design? We analyze a model where several agents bargain over persuasion of a single receiver. We characterize the existence of an agreement that is beneficial for all agents relative to some fixed benchmark information structure, when the preferences of agents are state-independent, and provide sufficient conditions for general preferences. We further show that a beneficial agreement exists if, for every coalition of a fixed size, there is a belief that generates enough surplus for its members. Next, we concentrate on "agent-partitional" environments, where for each agent there is a state where the informed decision of the receiver benefits her the most. In these environments, we define endorsement rules that fully reveal all such "agent-states". Endorsement rules are Pareto efficient when providing information at all agent-states generates enough surplus, and they correspond to a Nash Bargaining solution when the environment is also symmetric. We provide two political economic applications of our model. In a running example, we discuss the implication of our model to bargaining of authoritarian elites over media policy. The last section applies the model to an electoral campaign in a multiparty democracy.

JEL classification: C71 , D82
Keywords: Persuasion, Bargaining solution, Efficiency

*We would like to thank Georgy Egorov, Gregory Sasso, Steven Kivinen, Ming Li, Alex Suzdaltsev, the participants of the seminars at the Budapest University of Technology and Economics, Central European University, Concordia University, Corvinus University of Budapest, CERGE-EI, and HSE University for their useful comments.

†Faculty of Economic Sciences, HSE University, Russia (e-mail:kakoz@hse.ru).
‡QSMS research centre, Faculty of Economic and Social Sciences, Budapest University of Technology and Economics, Műegyetem rkp. 3., H-1111 Budapest, Hungary. E-mail: samsonov.arseniy@gtk.bme.hu

1 Introduction

A group of agents negotiates how to set up an information institution to influence the action of a decision-maker. If agents agree, a signal is generated at each state of the world according to some probability distribution. If they cannot agree, full, partial, or no information is revealed depending on the context. For instance, representatives of an industry and public officials design a disclosure policy, such as stress testing in the financial sector or disclosure of the health effects of tobacco products, in which case nothing becomes public unless regulations are in place. Another application is the bargaining between political elites over how much to limit (encourage) media freedom, in which case the media remain free (censored) unless the elites agree. For example, after Stalin's death, the Communist Party of the Soviet Union did not have a single leader, and its members competed for power. Some, including Khruschev (the General Secretary of the Party at the time), tried to ride the wave of exposing Stalin's crimes. In 1956, Khruschev denounced Stalin's repressions, which sparked a debate in Soviet society. Later that year, a group of Politburo members wrote a letter to Party sections that suggested limiting the criticism of the Soviet system. They agreed on a single line that preserved the regime's stability [Rowe, 1964]. After this, proponents of reform who called for a debate about repressions disappeared from newspapers and journals, and the official rhetoric became more conservative [Loewenstein, 2006].

In this paper, we analyze a generalization of the canonical Bayesian Persuasion setup by Kamenica and Gentzkow [2011] where multiple agents bargain over the information structures for the persuasion of a single receiver. When is there an agreement that benefits all agents relative to the expected disagreement outcome? A beneficial agreement exists if there is a distribution over beliefs, not necessarily Bayesian plausible, where for each agent the expected payoff is greater than the disagreement one. Furthermore, there is no such agreement if and only if there exist welfare weights such that the weighted sum of expected payoffs is always lower than the corresponding sum at the disagreement stage. How much transparency would agents agree on? We address this question for "agent-partitional" environments, where for each agent there is a state of the world where the informed decision of the receiver benefits her the most. The complete revelation of all such "agent-states" is Pareto efficient when the informed decision of the receiver at each of the agent-states generates enough total surplus.

More formally, we consider an environment with a finite set of agents, a single receiver, finite state space, and a finite set of actions available to the receiver. An information structure specifies a probability distribution over a given set of signals at each state of the world. If agents agree on an information structure, the receiver observes the signal, forms a Bayesian posterior, and takes an action. The corresponding outcome affects all agents in potentially different ways. If agents cannot agree, the game enters a disagreement stage, where a benchmark information structure prevails. The benchmark information structure might be a result of a competitive persuasion game played by the agents [Gentzkow and Kamenica, 2016, 2017], or it could be that none of the agents is able to provide any information without agreeing with others. We assume that the benchmark information structure induces beliefs on which the receiver takes a unique optimal action. We call such information structures *discrete*. We concentrate on environments where no information and full revelation are discrete.[1]

In section 3, we first assume that the agents' preferences are independent of the states and prove that the existence of a beneficial agreement is equivalent to the existence of a beneficial distribution over posterior beliefs. The proof for the sufficiency establishes the existence of a Bayesian plausible distribution over beliefs by using the assumption that the benchmark information structure is discrete. We also show that the non-existence of the beneficial agreement is equivalent to the existence of welfare weights over the agents such that the weighted sum of agents' expected payoffs does not exceed the disagreement payoff. We further prove that there is a beneficial agreement whenever for each coalition of a fixed size there is a belief such that the expected payoff is higher than the disagreement one for all coalition members, lower for others, and the overall surplus is positive. For general preferences, we prove the existence of a beneficial agreement whenever there is information to share by the agents.

In section 4, we concentrate on environments, in which for each agent there is a state of the world where the informed decision of the receiver generates the highest possible payoff for her. We call such states as *agent-states*, and the other states, where this is not true, the *agent-neutral* states. We also define *agent-actions* and *neutral actions* as the informed decisions at the agent-states and the agent-neutral states, respectively. Within such environments, bargaining among the agents could boil down to the agreement over the states where more information should be available to the receiver. We define *endorsement rules* that fully reveal all agent-states, while implementing some censorship at the agent-neutral state. Our results show that there is a Pareto efficient endorsement rule when the informed decision of the receiver at the agent-states generates a higher total surplus than her other actions while all agents prefer agent-actions to neutral ones. When the agents' payoffs are symmetric, and the prior belief is not biased towards any single agent-state, the symmetric endorsement rule is efficient whenever it is beneficial against the no-information benchmark. This, along with the proof that symmetric endorsement rule satisfies scale invariance and independence of irrelevant alternatives, enables us to interpret endorsement rules as a Nash solution to the bargaining problem among the agents.[2]

As an illustration of our model, consider the following example. A dictator represents an elite faction and rules a country. Another elite faction holds key positions in the bureaucracy and may contest the dictator's power. There are three potential leaders: the dictator, the insider, who represents the second elite faction, and an opposition leader. Citizens would like to support a competent leader but do not know which one is the most competent. The prior belief of the citizens that the opposition leader is the most competent is 0.4, while the corresponding probability for each of the other candidates is 0.3. If the citizens support the dictator, she stays in power. The dictator and the insider negotiate over the design of the transparency of political institutions that provide or censor information about the competency of potential leaders.

Assume the following payoff structure. If citizens revolt against the dictator and support another

[1] Note that in some cases, an individually optimal information structure makes the receiver indifferent between multiple actions for some signals. Then, our assumption of discrete benchmark information structure can be interpreted as the condition that no individual monopolizes the disagreement stage by implementing her own individually optimal information structure.

[2] We further show that Kalai-Smorodinsky and egalitarian bargaining solutions also coincide with the symmetric endorsement rule in symmetric environments.

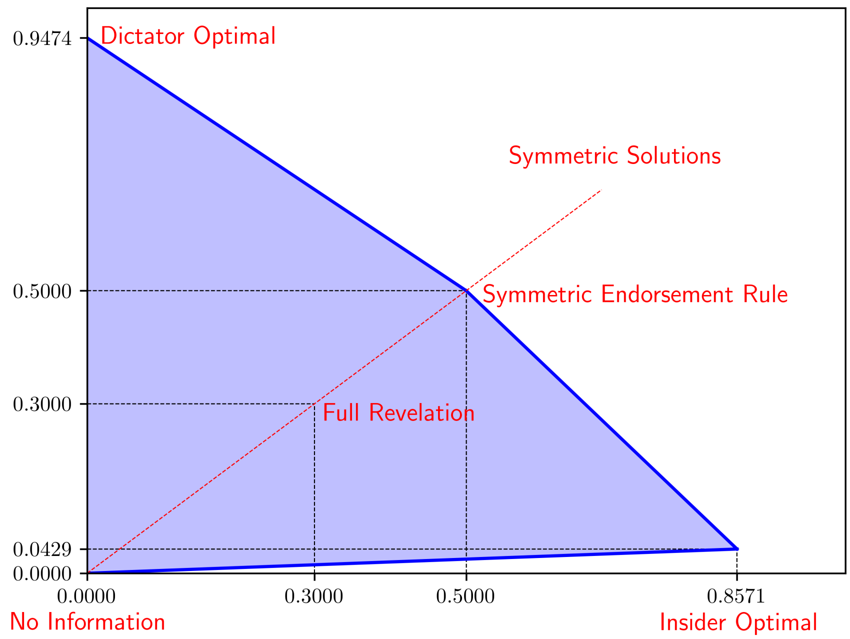

Figure 1: The horizontal axis is the Insider's expected payoff that some information structure can generate, and the vertical axis shows the same for the Dictator. The symmetric endorsement rule is efficient and corresponds to the Nash Bargaining solution.

leader, they incur a cost related to regime change. The cost of revolting is 0.05, while the benefit of supporting the most competent leader is 1. If a leader gets to rule, she gets 1, and 0 otherwise. Because of its high probability of being competent and low cost of revolting, the citizens support the opposition if they do not receive further information. In this case, the dictator and the insider get 0. Full information revelation makes each leader take power whenever she is competent. As a result, it yields each of the negotiating factions an expected payoff of 0.3. The dictator prefers an information structure that endorses her with full probability whenever the opposition leader is not the most competent and endorses the opposition with probability 0.0526. The insider prefers a similarly biased information structure. Note that the information structure must sometimes endorse the opposition to be credible. If the dictator and the insider have the same bargaining power, the Nash bargaining solution is a relevant approach. A symmetric endorsement rule generates an endorsement for the dictator with full probability when she is the most competent and with probability 0.5 when the opposition is the most competent. It works symmetrically for the insider. The symmetric endorsement rule benefits the dictator and the insider regardless of whether the disagreement stage leads to no information or full information revelation. The bargaining set of payoff profiles for this example is illustrated in Figure 1.

This example reveals some of the insights that the model can generate. The symmetric agreement between the two elite factions keeps the opposition from taking over the regime, while none of the individually optimal information structures does. This agreement benefits both elite factions compared to the disagreement scenarios that could lead to full or no information. Furthermore, the resulting information structure is fully transparent in the states where either of the two elite factions is competent and fully censors information in the remaining case. This type of partial transparency is novel compared to some of the other mechanisms noted in the literature, such as the trade-off between credibility and persuasion [Boleslavsky et al., 2021, Gehlbach and Sonin, 2014] or lack of enough resources for censorship [Guriev and Treisman, 2020].

In section 5, we discuss an application of our model to multiparty elections. We consider a scenario where multiple political parties contemplate a joint campaign against the incumbent. We show that there is a way for them to coordinate their electoral campaigns to maximize the probability of the incumbent's defeat. With the electoral campaign, each political party wins the upcoming election whenever its policy platform is the best for voters. When the parties have substantial ideological differences, such coordination might be easier than colluding on a single policy platform.

Our model is close to the recent work by Doval and Smolin [2021], which analyzes the set of payoff profiles of types of players that some information structure can generate. The main application is the interim payoffs that a single sender with multiple types might have. They also discuss the population interpretation of their framework, where each state of the world corresponds to a type of player in the population, and the prior belief corresponds to the population distribution of different types. One can view our setup as a generalization of the second approach to the model in Doval and Smolin [2021], where we explicitly distinguish the state space and the agent types. This distinction enables us to treat the beliefs and the social welfare weights as separate objects and to capture a wide range of applications.

Our paper is related to the literature on competition among multiple senders over information structures. The canonical models of competitive persuasion by Gentzkow and Kamenica [2016] and Gentzkow and Kamenica [2017] show that an increase in competition leads to more informative outcomes. However, consecutive work by Li and Norman [2018] and Li and Norman [2021] show that this result might depend on some factors, most importantly, on the timing of persuasion by the senders [See also Bhattacharya and Mukherjee, 2013, Board and Lu, 2018, Hoffmann et al., 2020, Au and Kawai, 2020, for some other related models]. The disagreement stage with partial or full revelation as the benchmark information structure may capture the outcomes from some competitive persuasion game in case of disagreement. In this sense, we consider environments where the senders can collude before entering the competition.

The rest of the text is organized as follows. In section 2 we lay out our model and establish the background for the analysis that follows. In section 3 we explore the conditions for the existence of beneficial agreements. In section 4.1 we prove the results that establish the relation between symmetry and transparency; while in section 4.2 we consider asymmetric environments. Finally, in section 5 we discuss applications of our model. Section 6 concludes.

2 Model

Our setup focuses on a decision by a group of agents on how to provide information to a receiver. To formalize this situation, it is useful to introduce a social planner who proposes an information structure. Consider a persuasion game between a social planner and a receiver where the planner chooses an information structure that a finite set of agents can agree on. Let $N = \{1, \dots, n\}$ be the finite set of agents. There is a single receiver r. The state space is finite and denoted by Ω, and let $\Delta(\Omega)$ be the set of beliefs, $\mu \in \Delta(\Omega)$ denote any belief. We assume a symmetric information environment and let $\mu_0 \in \Delta(\Omega)$ be the common prior among the planner, the agents, and the receiver. With an abuse of notation, let ω also denote the degenerate belief that the state is ω with probability 1. Any belief $\mu \in \Delta(\Omega)$ is called **interior** if $\mu(\omega) > 0$ for any $\omega \in \Omega$. We assume that the prior belief, μ_0, is interior.

We assume that the set of pure actions that the receiver can take is finite and is denoted as A. Let $u_r : A \times \Omega \to \mathbb{R}$ be the payoff function for the receiver and for any $i \in \{1, \dots, n\}$ $u_i : A \times \Omega \to \mathbb{R}$ be the payoff function for the agent i. Define the expected payoff of the receiver for playing a when her belief is μ as

$$v_r(a, \mu) \equiv \sum_{\omega \in \Omega} \mu(\omega) u_r(a, \omega).$$

For any agent $i \in N$, $v_i(a, \mu)$ is similarly defined. For any belief $\mu \in \Delta(\Omega)$ $a(\mu)$ denotes any receiver-optimal action at belief μ. That is, $a(\mu) \in \arg\max_{a' \in A} v_r(a', \mu)$. Note that for any two $a(\mu), a'(\mu) \in A$ $v_r(a(\mu), \mu) = v_r(a'(\mu), \mu)$; however, the same may not hold for the agents. Therefore, whenever the receiver is indifferent among multiple actions under some beliefs, the agents' payoffs might be non-trivially dependent on the particular choice of the receiver. Now, fix the receiver's optimal action profile $\{\bar{a}(\mu)\}_{\mu \in \Delta(\Omega)}$. Then, we can define $\bar{v}_i(\mu) \equiv v_i(\bar{a}(\mu), \mu)$ as the value of belief $\mu \in \Delta(\Omega)$ for agent $i \in N$ given the action profile $\{\bar{a}(\mu)\}_{\mu \in \Delta(\Omega)}$ of the receiver.

Information Structure

An information structure π chooses a probability distribution at every state $\omega \in \Omega$ over a signal space. After observing the realized signal, the receiver updates her belief according to the Bayesian rule and chooses an optimal action. It is without loss of generality to assume that the social planner chooses a probability distribution over actions at each state of the world.[3] As the action space is finite we can represent such an information structure as a $|A| \times |\Omega|$ probability matrix such that the sum of each column is 1.

We call an information structure *incentive compatible* if the receiver finds it optimal to follow the recommendations generated by the information structure. Fix any information structure π. For any action $a \in A$, let μ_a denote the posterior belief after observing the recommendation a, which can be calculated as for any $\omega \in \Omega$

$$\mu_a(\omega) = \frac{\pi(a, \omega)\mu_0(\omega)}{\sum_{\omega \in \Omega} \pi(a, \omega')\mu_0(\omega')},$$

whenever $\pi(a, \omega) > 0$ for some $\omega \in \Omega$. Then, an information structure is incentive-compatible if for any $a, a' \in A$ and $\omega \in \Omega$, $\pi(a, \omega) > 0$ implies that $v_r(a, \mu_a) \geq v_r(a' \mu_a)$. The following remark immediately follows from construction.

[3]Let Z denote any finite public signal space. Then, an information structure is a function $\hat{\pi} : Z \times \Omega \to [0, 1]$ such that for any state $\omega \in \Omega$ $\sum_Z \hat{\pi}(z, \omega) = 1$. We can define posterior beliefs $\{\mu_{\hat{\pi}(z)}\}_{z \in Z}$ that can be generated by the information structure $\hat{\pi}$. Fix the action profile of the receiver. Then, for any $a \in A$ define $\pi(a, \omega) \equiv \sum_{\{z \in Z | a = a(\mu_{\hat{\pi}(z)})\}} \hat{\pi}(z, \omega)$.

Remark 1 *Let* Π *be the set of incentive-compatible information structures.* Π *is a compact and convex subset of* $[0,1]^{|A|\times|\Omega|}$.

We can describe the timeline of the persuasion game as follows:

- The social planner publicly chooses and commits to an information structure π given the common prior μ_0.
- The nature draws the state of the world ω according to the common prior μ_0 and a recommendation $a \in A$ realizes according to π.
- The receiver observes both the information structure π and the recommendation a, then she updates her belief and chooses an optimal action.

An incentive-compatible information structure does not necessarily mean that the receiver always follows a recommendation. The receiver may break ties randomly or according to any rule. To make sure that the set of expected payoffs is compact, one can concentrate on sender-preferred equilibria, where the receiver breaks ties in favor of the sender's objective function [See Kamenica and Gentzkow, 2011]. As we have multiple objective functions in our framework, it is not possible to make a similar assumption. Instead, we assume that the receiver follows the recommendations of any incentive-compatible information structure.[4]

Assumption 1 *Let* π *be any incentive-compatible information structure. Whenever the receiver observes the choice of the information structure* π*, she follows any recommendation generated by* π.

For any incentive-compatible information structure π, Assumption 1 enables us to write the expected payoff of the receiver (and any agent) directly in terms of recommendation probabilities as follows:

$$\mathbb{E}_\pi \bar{v}_r(\mu_\pi) \equiv \sum_{\omega\in\Omega} \mu_0(\omega) \sum_{a\in A} \pi(a,\omega) u_r(a,\omega).$$

As proved by Kamenica and Gentzkow [2011], the expected payoffs for a given information structure can be written as expectations over a Bayesian plausible distribution over the posteriors. For any information structure π, let $\tau_\pi \in \Delta(\Delta(\Omega))$ be the induced posterior distribution.[5] Then

$$\mathbb{E}_\tau \bar{v}_r(\mu) = \sum_{\mu\in supp(\tau)} \tau(\mu)\bar{v}_r(\mu),$$

and $\mathbb{E}_\tau \bar{v}_r(\mu) = \mathbb{E}_\pi \bar{v}_r(\mu_\pi)$ if τ is induced by π.

Each agent has a favorite information structure. However, agents have to agree on the information structure they implement. If they cannot agree on an information structure, the game enters into a disagreement stage. The outcome of the stage for each agent $i \in N$ defines her disagreement payoff, d_i. If agents fail to generate any credible information in case they disagree, $d_i = \bar{v}_i(\mu_0)$. However, it might also be possible in some contexts that the agents enter into a competitive disclosure game when

[4] An alternative way is to assume that the social planner can control the tie-breaking rule that the receiver follows. See Doval and Smolin [2021] for a discussion.

[5] Note that $supp(\tau_\pi) = \{\mu_a | a \in A \text{ and } \pi(a,\omega) > 0 \text{ for some } \omega \in \Omega\}$ and $\tau_\pi(\mu_a) = \sum_{\omega\in\Omega} \mu_0(\omega)\pi(a,\omega)$.

they cannot agree. If that game results in full revelation, $d_i = \sum_{\omega \in \Omega} \mu_0(\omega)\bar{v}_i(\mu_\omega)$. In general, we can define a default information structure π_0, and the induced Bayesian plausible posterior distribution τ_0, as the outcome of the disagreement stage, and define the disagreement payoffs as $d_i \equiv \mathbb{E}_{\tau_0}\bar{v}_i(\mu)$.

The Bargaining Set

The agreement among agents follows a given bargaining solution. The bargaining set consists of payoff vectors that can be achieved by choosing an information structure and letting the receiver choose a receiver-optimal action. Fix any action profile $(\bar{a}(\mu))_{\mu \in \Delta(\Omega)}$ of the receiver. Let $B \subseteq \mathbb{R}^n$ denote the bargaining set, which is defined as

$$B \equiv \{(x_i)_{i \in N} \in \mathbb{R}^n | \exists \pi \in \Pi \;\; x_i = \mathbb{E}_\pi \bar{v}_i(\mu_\pi) \;\; \forall i \in N\}$$

With Assumption 1, it is possible to show that B is a compact and convex set (See Lemma 1 in Appendix A). Furthermore, following Doval and Smolin [2021], we can interpret the bargaining set B as the convex cone of a vector-valued payoff function and extend the concavification result of Kamenica and Gentzkow [2011] to our framework (See Proposition 4 in Appendix A).

A bargaining solution $F : \mathcal{B} \to B$, where $\mathcal{B}$ is the non-empty subsets of B, picks a vector of payoffs from each subset of B. Our main focus is agent-efficient bargaining solutions, where the solution should be Pareto efficient among the agents. Nash bargaining solution is an important example of an efficient bargaining solution.

3 Beneficial Agreement

We say that there is a *beneficial agreement* whenever there is an information structure that strictly benefits all agents relative to their disagreement payoffs. Clearly, if there is a Bayesian plausible distribution of posterior beliefs that generate a higher payoff to all agents, then there is beneficial agreement. However, it is possible to drop the requirement of Bayesian plausibility under some conditions on the default information structure. Kamenica and Gentzkow [2011] show a similar result when the receiver's preference is *discrete* at the prior; that is, there is a unique optimal action at the prior belief. We generalize this notion as follows.

Definition 1 *Let $\pi \in \Pi$ be any incentive-compatible information structure. π is called **discrete** if for any $a \in A$ $\pi(a, \omega) > 0$ for some $\omega \in \Omega$ implies that $\{a\} = \arg\max v_r(a(\mu_a), \mu_a)$.*

Discrete information structures induce posterior beliefs where the receiver's preferences are discrete. The set of discrete information structures is a dense subset of the incentive-compatible information structures[6]; however, it may not include some of the interesting information structures in general. Assumption 2 below ensures that full revelation and no information are both discrete.

Assumption 2 *At every state $\omega \in \Omega$, $|\arg\max_{a \in A} v_r(a, \mu_\omega)| = 1$ and $|\arg\max_{a \in A} v_r(a, \mu_0)| = 1$.*

[6]Note that we can arbitrarily approximate any incentive-compatible information structure with a discrete one in terms of the posterior beliefs they induce. However, note that since the receiver might be indifferent in multiple actions at some beliefs that are induced by some interesting information structures, the approximation in terms of beliefs does not necessarily translate into approximation in payoffs.

When the preferences of agents do not depend on states but only on actions, we can fully characterize the existence of a beneficial agreement in terms of finite lotteries over posterior beliefs, whether or not the lottery is Bayesian plausible or not. Assumption 3 formalizes the notion of state independence.

Assumption 3 *For any $a \in A$, $\omega, \omega' \in \Omega$, and $i \in N$ $u_i(a,\omega) = u_i(a,\omega')$.*

Our first main result in this section is as below.

Theorem 1 *Suppose that Assumptions 1 and 3 hold, the default information structure π_0 is discrete, and fix any optimal action profile $(\bar{a}(\mu))_{\mu \in \Delta(\Omega)}$ by the receiver. Then the following are equivalent.*

(i) *There is a beneficial agreement.*

(ii) *There is a finite lottery over beliefs $(\lambda_l, \mu_l)_{l=1}^h$ for some finite positive integer h such that for any $l = 1, \ldots, h$ $\lambda_l \in [0,1]$ and $\sum_l \lambda_l = 1$ and for any agent $i \in N$*

$$\sum_l \lambda_l \bar{v}_i(\mu_l) > d_i. \tag{1}$$

(iii) *There do not exist non-negative numbers $y_1, ..., y_n$ which are not all equal to 0 such that for any belief $\mu \in \Delta(\Omega)$*

$$\sum_{i \in N} y_i(\bar{v}_i(\mu) - d_i) \leq 0. \tag{2}$$

All the proofs are in Appendix A.

The first critical argument in the proof of Theorem 1 is to show that (ii) implies (i); that is, the existence of any finite lottery that satisfies the inequality (1) is sufficient for a beneficial agreement. The proof follows similar steps as the proof of Proposition 2 by Kamenica and Gentzkow [2011]. First, state-independence and the fact that π_0 is discrete allow us to find an approximate interior information structure $\hat{\pi}_0$ such that the receiver's optimal actions do not change, and $\hat{\pi}_0$ leads to exactly the same expected payoffs as π_0 for all agents. As each belief induced by $\hat{\pi}_0$ is interior, we can find beliefs that are close enough to the beliefs induced by $\hat{\pi}_0$ so that we can express all beliefs induced by $\hat{\pi}_0$ as convex combinations of a belief in the support of the lottery and some other belief. This enables us to construct a Bayesian plausible lottery over beliefs that leads to an expected payoff higher than the one at the default information structure π_0.

The second part of the proof of Theorem 1 is to show that either statement (ii) or the negation of statement (iii) holds. For any finite collection of beliefs we can construct a matrix of expected payoffs. Then, by an equivalent statement of Farkas' Lemma [Perng, 2017], either there exists a vector of beliefs such that statement (ii) holds, or there exists a set of weights over agents such that the inequality (2) holds.

Theorem 1 enables us to find simple sufficient conditions for the existence of beneficial agreement. Corollary 1 below shows that if for each agent the average payoff of full revelation is positive, then there is a beneficial agreement.

Corollary 1 *Suppose that Assumptions 1 and 3 hold, the default information structure* π_0 *is discrete, and fix any optimal action profile* $(\bar{a}(\mu))_{\mu\in\Delta(\Omega)}$ *by the receiver. If for any agent* $i \in N$ $\sum_{\omega\in\Omega} \frac{1}{|\Omega|} u_i(a(\omega),\omega) > d_i$*, then there is a beneficial agreement.*

The example below illustrates Corollary 1.

Example 1 *Suppose that* $n = 2$*,* $\Omega = \{\omega_1, \omega_2, \omega_0\}$*, and* $A = \{a_1, a_2, a_0\}$ *with* $u_r(a_i, \omega_j) = 1$ *if* $i = j$*, and* 0 *otherwise.* $\mu_0 = (0.2, 0.2, 0.6)$*. For each* $i, j \in N$ *with* $i \neq j$*,* $u_i(a_i) = 3$*,* $u_i(a_j) = -1$*, and* $u_i(a_0) = -2$*.*

For the example below, no information leads to -2 for all agents, while the full revelation leads to -0.4 as the ex-ante expected payoff. However, the average revelation payoff is 0, which by Corollary 1 implies that there is a beneficial information structure whether the default information structure is no information or full revelation. Set $\pi(\cdot, \omega_0) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ and $\pi(a_i, \omega_i) = 1$ for each $i \in N$. The expected ex-ante payoff corresponding to π is 0.8 for all agents.

The inequality (2) in Theorem 1 can be interpreted as the existence of a *linear conflict of interest* between the agents. Whenever that is the case, there is no beneficial agreement. A linear conflict of interest that prevents beneficial agreement can also arise among a subset of agents. The following corollary points out an immediate but useful implication of Theorem 1.

Corollary 2 *Suppose that Assumptions 1 and 3 hold, the default information structure* π_0 *is discrete, and fix any optimal action profile* $(\bar{a}(\mu))_{\mu\in\Delta(\Omega)}$ *by the receiver. If there are two agents* $i, j \in N$ *such that for any belief* $\mu \in \Delta(\Omega)$*,* $\bar{v}_i(\mu) - d_i = -(\bar{v}_j(\mu) - d_j)$*, then there is no beneficial agreement.*

The example below illustrates Corollary 2 and proves that the existence of excess total surplus at every belief is not sufficient for the existence of a beneficial agreement.

Example 2 *Consider an environment with* $\Omega = \{\omega_0, \omega_1, \omega_2, \omega_3\}$*,* $A = \{a_0, a_1, a_2, a_3\}$ *and* $N = \{1, 2, 3\}$*. The payoff function of the receiver is such that* $\bar{a}_{\omega_i} = a_i$*. The payoff profile of the agents satisfy for any belief* $\bar{v}_1(\mu) - d_1 = -(\bar{v}_2(\mu) - d_2)$*,*

$$\bar{v}_3(\mu) - d_3 = \max\{\bar{v}_1(\mu) - d_1, \bar{v}_2(\mu) - d_2\},$$

and, $u_1(a_1, \omega_1) > d_1$*, while* $u_2(a_2, \omega_2) > d_2$*.*

In Example 2, in every belief $\mu \in \Delta(\Omega)$ at which at least one agent expects to receive a payoff different than the disagreement payoff, the total payoff exceeds the total disagreement payoff. However, the "zero-sum" nature of the interests between the first two agents, prevents the existence of a beneficial agreement. Theorem 2 below provides a general sufficient condition for the existence of the beneficial agreement. The sufficient conditions imply that for every coalition of a fixed size among agents, there exists a belief at which the coalition "wins" relative to no agreement. An implication of this condition is that there are no groups of agents who are in conflict with each other at all possible states.

Theorem 2 *Suppose that Assumptions 1 and 3 hold, the default information structure π_0 is discrete, and fix any optimal action profile $(\bar{a}(\mu))_{\mu \in \Delta(\Omega)}$ by the receiver. If there exists $k \in \{1, \ldots, n\}$ such that we can find a set of beliefs $M \subset \Delta(\Omega)$ and for each $\phi \in \Phi(k)$, where $\Phi(k)$ denotes the set of subsets of $\{1, \ldots, n\}$ of size k, there is $\mu^\phi \in M$ that satisfies*

$$\forall i \in \phi \quad \bar{v}_i(\mu^\phi) > d_i$$
$$\forall i' \notin \phi \quad \bar{v}_{i'}(\mu^\phi) < d_{i'}$$
$$\forall \phi \in \Phi(k) \quad \sum_{i'=1}^{n} \left(\bar{v}_{i'}(\mu^\phi) - d_{i'}\right) > 0,$$

there is a beneficial agreement.

The proof of Theorem 2 is an application of the Separating Hyperplane Theorem. We first prove the general Lemma 4 in Appendix A. For any fixed integer $k \in \{1, \ldots, n\}$, for any subset of $\{1, \ldots, n\}$ of size k suppose that there is a vector in $\mathbb{R}^n$ such that the components of the vector corresponding to the subset of $\{1, \ldots, n\}$ are positive while others are negative, and the sum of components is positive. Then, if there does not exist a convex combination of these vectors that is strictly positive, by the Separating Hyperplane Theorem we can separate the set of convex combinations of these vectors from the set of non-negative vectors. But this contradicts the fact that the sum of components of all vectors is positive. Then, we apply Theorem 1 to show that there exists a beneficial agreement.

To illustrate the Theorem 2 above, consider a variant of the example that we discussed in the Introduction. Suppose that there are three agents: the dictator, the military regime insider, and the civilian regime insider, and a non-strategic opposition. The citizens would like to support the competent agent or the opposition but do not know which one is. The dictator is competent with probability 0.3, the military and civilian are each competent with probability 0.15, and the opposition is competent with probability 0.4. As before, supporting a competent leader yields a payoff of 1 to the citizens; and supporting the civilian, military, or opposition involves a cost of 0.05 to the citizens.

	The Leader in Power			
payoffs	dictator	civilian	military	opposition
dictator	1	0	0.5	0
civilian	0	1	-5	0
military	-5	0	1	-8
total	-4	1	-3.5	-8

Suppose that the benchmark information structure is full revelation. Then the expected disagreement payoffs are 0.375 for the dictator, -0.6 for the civilian insider, and -4.55 for the military insider; and the total is -4.75. Consider all pairs of agents. It is clear from the table that the pair (dictator, civilian) gains relative to the benchmark if the dictator is in power while the military loses relative to the benchmark. The pair (civilian, military) gains if civilian is in power while the dictator loses. Finally, the pair (military, dictator) gains if the military is in power while the civilian loses. In all cases, the sum of payoffs exceeds the sum of payoffs under the benchmark. Hence, by the previous Theorem, there exists a beneficial information structure relative to full revelation.

If for each agent $i \in N$ there is a unique state of the world such that i is the unique winner when the receiver learns the state but the total amount of loss is less than the win, there is a beneficial agreement.

Corollary 3 *Suppose that Assumptions 1 and 3 hold, the default information structure π_0 is discrete, and fix any optimal action profile $(\bar{a}(\mu))_{\mu \in \Delta(\Omega)}$ by the receiver. Assume that $|\Omega| \geq n$. If there exists an injection $f : N \to \Delta(\Omega)$ such that for any $i \in N$ $u_i(\bar{a}(f(i))) > d_i$ but $u_j(\bar{a}(f(i))) < d_j$ for any $j \neq i$ and*

$$\sum_{j=1}^{n} u_j(\bar{a}(f(i))) - d_j > 0,$$

there is a beneficial agreement.

3.1 State-Dependent Payoffs

When the payoffs of agents are allowed to be state-dependent, the existence of any lottery of beliefs with a higher expected payoff is no longer generally sufficient for the existence of a beneficial information agreement. In particular, when the default information structure is not interior; that is, when some of the beliefs that it generates are not interior, then Theorem 1 may fail to hold. Consider the following example.

Example 3 *Suppose that $n = 2$, $\Omega = \{\omega_1, \omega_2\}$, and $A = \{a_1, a_2\}$ with $u_r(a_i, \omega_j) = 1$ if $i = j$, and 0 otherwise. The agents' payoffs are such that $u_1(a_1, \omega_1) = 2 = u_2(a_1, \omega_1)$, $u_1(a_2, \omega_2) = 1 = u_2(a_2, \omega_2)$ and 0 otherwise. $\mu_0 = (0.5, 0.5)$, and the default information structure is full revelation.*

In the example above, the optimal information structure is full revelation, which is also the default information structure. Therefore, there is no beneficial information structure; however, there is a belief μ_{ω_1} where the payoff is 2, which is higher than the ex-ante payoff at the full revelation, which is 1.5.

Suppose that Assumption 2 holds and the default is no information. Kamenica and Gentzkow [2011] propose a condition called "information to share" for the existence of an information structure that a single sender benefits from. We generalize this condition to lotteries of beliefs and any default information structure as follows. Consider any lottery over a finite set of beliefs, which is denoted as $(\lambda_l, \mu_l)_{l=1}^{h}$ with $\sum \lambda_l = 1$ and for any $l = 1, \ldots, h$ $\lambda_l \in (0, 1)$. We say that a lottery over beliefs has **further information to share** relative to some incentive compatible information structure $\pi \in \Pi$ if for any agent $i \in N$

$$\sum_{l=1}^{h} \lambda_l \bar{v}_i(\mu_l) > \sum_{\mu \in supp(\tau)} \tau(\mu) \sum_{l=1}^{h} \lambda_l v_i(a(\mu), \mu_l), \tag{3}$$

where τ is the Bayesian plausible distribution over posterior beliefs induced by π. The idea is that if all agents have some private information, represented by the lottery over beliefs $(\lambda_l, \mu_l)_{l=1}^{h}$, none of them would veto sharing this information with the receiver; because the expected payoff under not sharing this information (right-hand side of the inequality (3)) is lower than the expected payoff of sharing it (left-hand side of the inequality (3)). The Proposition 1 result is the corresponding sufficient condition of the existence of a beneficial agreement when the default information structure is interior.

Proposition 1 *Suppose that Assumption 1 holds and fix any optimal action profile* $(a(\mu))_{\mu\in\Delta(\Omega)}$ *by the receiver. Assume that* π_0 *is discrete and interior. There is a beneficial agreement if there is a lottery over beliefs* $(\lambda_l, \mu_l)_{l=1}^{h}$ *that has further information to share relative to* π_0*.*

To illustrate the Proposition 1 above consider the following variant of the example that we discussed in the Introduction. There are two agents, the dictator and the insider, and a non-strategic opposition. The citizens would like to support the competent agent or the opposition but do not know which one is. The dictator and insider are each competent with probability 0.3 and the opposition is competent with probability 0.4. As before, supporting a competent leader yields a payoff of 1 to the citizens; and supporting the civilian, military, or opposition involves a cost of 0.05 to the citizens. Now, we assume that the insider has state dependent preferences. In particular, the insider gets 1 if dictator is in power and he is also competent; but gets -1 if the dictator is incompetent. The dictator gets 1 only if he is in power and 0 otherwise. Finally, if opposition is in power, the dictator gets -1, while the insider gets 0.

Suppose that the benchmark information structure is full censorship; hence, the opposition yields the power in case the agents cannot agree on an information structure. Consider a trivial lottery that always puts probability 1 on the dictator being competent. If citizens hold this belief, they support the dictator. Hence, for both agents, the left-hand side of inequality (3) is equal to 1. With probability 1, censorship yields a belief equal to the prior. Hence, the right-hand side of (3) is -1 for the dictator and 0 for the insider. Proposition 1 ensures that an information structure that benefits both agents compared to censorship exists. For instance, full revelation yields both agents a higher ex-ante expected payoff than censorship.

Note that it is possible to state and prove a parallel of Theorem 2 using Proposition 1 above. If for any coalition of fixed size, there is further information such that sharing this information benefits the coalition members but hurts others; then we can find a lottery of beliefs that has further information to share as in Proposition 1.

We make a final observation before concluding this section. If there is a beneficial agreement in an environment, there should be a beneficial agreement in less antagonistic environments as well. To formalize such a comparison, for any given optimal action profile $(a(\mu))_{\mu\in\Delta(\Omega)}$, we say that an environment with the preference profile $(u_i)_{i\in N}$ is less antagonistic than an environment with $(u_i')_{i\in N}$ if for any belief $\mu \in \Delta(\Omega)$ and agent $i \in N$

$$\bar{v}_i(\mu) - \mathbb{E}_{\pi_0}\bar{v}_i(\mu_\pi) \geq \bar{v}_i'(\mu) - \mathbb{E}_{\pi_0}\bar{v}_i'(\mu_\pi),$$

where $\bar{v}_i(\cdot)$ and $\bar{v}_i'(\cdot)$ are the expected value functions corresponding to the preference profiles $(u_i)_{i\in N}$ and $(u_i')_{i\in N}$ respectively.

If an environment is less antagonistic than another, there is more surplus to share given each information structure. Remark 2 below formalizes this intuition.

Remark 2 *If there is a beneficial agreement in any environment with the preference profile* $(u_i)_{i\in N}$*, then there is also a beneficial agreement in any environment that is less antagonistic.*

4 Endorsement Rules and Symmetry

In this section, we concentrate on environments that can be completely described from the agents' perspective. That is, for each state $\omega \in \Omega$, the corresponding informed action by the receiver leads to either the best outcome for one of the agents or leads to an outcome that no agent prefers the most. We call such environments as **agent-partitional**. More formally, an agent-partitional environment is such that $\Omega = \{\omega_0, \omega_1, \dots, \omega_n\}$, $A = \{a_0, a_1, \dots, a_n\}$, and for each $i \in N$ $\arg\max_{a\in A} u_r(a, \mu_{\omega_i}) = a_i$, while for any $j, l \in \{0, 1, \dots, n\}$ $u_i(a_i, \omega_i) \geq u_i(a_j, \omega_l)$. Based on this definition, we can label the states $\{\omega_i\}_{i\in N}$ as the agent-states and the state $\{\omega_0\}$ as the neutral state. Actions can be partitioned in a similar way between agent-actions and the neutral action. As the receiver always prefers more information, for any $i \in N$ the receiver and the agent i have aligned interests regarding full information at ω_i. However, when the actions of the receiver are more important for the agents than the states, there would be a conflict of interest between the receiver and agents at the neutral state.

In such environments, a natural information structure to consider is the one that favors each agent at the corresponding state by providing full information at that state. Fix any agent-partitional environment; we call an information structure $\pi^e : A \times \Omega \to [0, 1]$ an **endorsement rule** if for any $i \in N$ $\pi^e(a_i, \omega_i) = 1$ as π^e recommends, in a credible way, the favorite action of each agent at the associated state. Note that full disclosure is also an endorsement rule; moreover, full disclosure is the only endorsement rule when there is no neutral state or action but all states and actions are agent states and actions.

Throughout this section, we impose the following restriction on the Receiver's payoffs to simplify the exposition. For any $i, j \in \{0, 1, \dots, n\}$

$$u_r(a_i, \omega_j) = \begin{cases} x_i, & \text{if } i = j; \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

where $x_i > 0$.

4.1 Symmetry and Nash Bargaining

Fairness in bargaining solutions is usually captured by symmetry requirements such as the symmetry of final payoffs or symmetry of the surplus relative to the disagreement payoff. One can interpret full disclosure as an information rule that satisfies procedural symmetry because the amount of information it provides at each state does not depend on which agent favors full information. However, the resulting payoffs may not be symmetric if the prior distribution or agents' payoffs are not symmetric. If the environment is symmetric in both aspects, then full disclosure leads to symmetric outcomes.

As we discussed in the example in the introduction, symmetry does not imply efficiency. Recall that, there, a symmetric endorsement rule leads to a higher payoff for both agents than full disclosure. We show below that a symmetric endorsement rule corresponds to a symmetric Nash Bargaining solution, which is symmetric and efficient when the environment is symmetric. We define symmetric environments below.

Definition 2 *Consider an agent-partitional environment, so $|\Omega| = |A|$. This environment is called* ***symmetric*** *if for any permutation $\sigma : N \to N$ and $i, j, l \in N$ $u_i(a_j, \omega_l) = u_{\sigma(i)}(a_{\sigma(j)}, \omega_{\sigma(l)})$, $\mu_0(\omega_i) = \mu_0(\omega_{\sigma(i)})$, and finally the receiver is indifferent among all agent states at the prior; that is, $v_r(a_i, \mu_0) = v_r(a_j, \mu_0)$.*

Note that in a symmetric environment, the restriction in equation 4 implies that the Receiver's payoff for choosing the correct action at all agent states is constant. For the results below, denote $x \equiv u_r(a_i, \omega_i)$ for any $i \in N$, and $x_0 = u_r(a_0, \omega_0)$.

In a symmetric environment, any symmetric information structure would provide the same amount of information at each of the agent-states so that no agent's best scenario is favored in terms of prior probabilities. A Pareto efficient information structure would provide enough information in at least one of agent-states to make sure that the likelihood of the best scenarios of some agents is not suboptimally low. Combining these two features proves that the symmetric endorsement rule is both symmetric in terms of payoffs and Pareto efficient among the agents. Proposition 2 below provides sufficient conditions for the symmetric endorsement rule to be symmetric and efficient. It proves further that it corresponds to the Nash Bargaining solution at the bargaining set B.

Proposition 2 *Suppose that Assumption 1 holds and consider any symmetric environment with $|\Omega| = |A| = n+1$, $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, $A = \{a_0, a_1, \ldots, a_n\}$, and for any $i \in \{0, 1, \ldots, n\}$ the optimal receiver action profile satisfies $a(\omega_i) = a_i$ and $a(\mu_0) = a_0$.*

Then the symmetric Nash, Kalai-Smorodinsky, and egalitarian bargaining solutions correspond to a symmetric endorsement rule π if

- $\forall j \neq i : \sum_k u_k(a_i, \omega_i) \geq \sum_k u_k(a_i, \omega_j)$
- $\forall i \neq 0 : \sum_k u_k(a_i, w_0) \geq \sum_k u_k(a_0, w_0)$

Moreover, for any $i \in N$,

$$\pi_{i0} = \begin{cases} \frac{1}{n}, & \text{if } \frac{x(1-\mu_0)}{x_0\mu_0} \geq 1 \\ \frac{x(1-\mu_0)}{nx_0\mu_0}, & \text{otherwise} \end{cases}$$

When there is no agent-neutral state and action, the symmetric endorsement rule reduces to full disclosure.

Corollary 4 *Suppose that Assumption 1 holds and consider any symmetric environment with $|\Omega| = |A| = n$, $\Omega = \{\omega_1, \ldots, \omega_n\}$, $A = \{a_1, \ldots, a_n\}$, and for any $i \in \{0, 1, \ldots, n\}$ the optimal receiver action profile satisfies $a(\omega_i) = a_i$, and she uniformly randomizes over all actions at μ_0. Full disclosure corresponds to the symmetric Nash bargaining solution when an environment is symmetric* *and when full revelation is beneficial against a no-information benchmark, which is equivalent to the condition below that must hold for any $i, j, l \in N$ with $i \neq j \neq l$*

$$u_i(a_i, \omega_i) + (n-1)u_i(a_j, \omega_j) > u_i(a_j, \omega_i) + u_i(a_i, \omega_j) + (n-2)u_i(a_j, \omega_l). \tag{5}$$

4.2 Efficiency in asymmetric environments

When we allow for asymmetric environments, requiring symmetry from an agreement may not be straightforward because an information structure that ensures either equal outcomes or equal surplus relative to the disagreement payoffs may not provide the same amount of information at each agent state. In other words, there could be a tension between consequential symmetry and a procedural one. Nevertheless, the appeal of a Pareto efficient agreement may extend beyond the symmetric environments. In this section, we provide sufficient conditions for the Pareto efficiency of endorsement rules.

Consider any environment with $|\Omega| = |A| = n+1$, $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, $A = \{a_0, a_1, \ldots, a_n\}$, and for any $i \in \{0, 1, \ldots, n\}$ the optimal receiver action profile satisfies $a(\omega_i) = a_i$. For any $l, j \in \{0, 1, \ldots, n\}$, let $\Gamma_{lj} \equiv \sum_{i \in N} u_i(a_l, \omega_j)$.

Assumption 4 *Consider any environment with $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, $A = \{a_0, a_1, \ldots, a_n\}$, and for any $i \in \{0, 1, \ldots, n\}$ the optimal receiver action profile satisfies $a(\omega_i) = a_i$. The payoff profile of agents satisfy for any $i \in N$, $j, l \in \{0, 1, \ldots, n\}$ with $l \neq j$ $u_i(a_l, \omega_j) > u_i(a_0, \omega_j)$ if $l \neq 0$, and $\Gamma_{ii} \geq \Gamma_{lj}$.*

Note that Assumption 4 implies that any information structure that incentivizes the receiver to play any action other than a_0 with some positive probability is beneficial against the no-information benchmark. Moreover, there is always an endorsement rule that is beneficial against the full information benchmark. The existence of a Pareto efficient and beneficial endorsement rule is guaranteed when there is additional surplus to share when the agent states are fully revealed.

Theorem 3 *Consider any environment with $|\Omega| = |A| = n+1$, $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, $A = \{a_0, a_1, \ldots, a_n\}$, and suppose that Assumption 1 holds. For the receiver, assume that $x_0\mu_0(\omega_0) > x_i\mu_0(\omega_i)$.*

If Assumption 4 also holds, there is an incentive-compatible endorsement rule π^e that is Pareto efficient and beneficial relative to no or full information benchmarks.

It is possible to provide tighter conditions for the efficiency of endorsement rules when agents' payoffs are state-independent. The first part of Theorem 4 below shows that the endorsement rule is Pareto-efficient if there are social welfare weights so that the corresponding weighted sum of the payoffs across the Receiver's informed decisions at all agent-states are equal to each other. Note that this condition does not require that for every non-neutral action ($a \neq a_0$) there is an agent that finds this action optimal. Instead, this condition requires that each such action generates enough surplus for every agent so that the agents may collectively agree to recommend this action with positive probability along with other agent-actions. The second part of Theorem 4 shows that an endorsement rule is efficient when the sum of agents' payoffs at every agent-action is greater than the sum of their second-best payoffs. This rules out the existence of an agent-action that is almost optimal for all agents, while all other actions hurt all agents except one.

Theorem 4 *Consider any environment with $|\Omega| = |A| = n+1$, $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, $A = \{a_0, a_1, \ldots, a_n\}$, and suppose that Assumption 1 and 3 hold with the additional restriction that $u_i(a_j) > u_i(a_0)$ for any $i, j \in N$. For the receiver, assume that $x_0\mu_0(\omega_0) > x_i\mu_0(\omega_i)$. Then there exists an endorsement rule that is incentive-compatible, Pareto-efficient, and beneficial relative to no and full information benchmarks, if*

(i) *there exist strictly positive weights $\{y_i\}_{i\in N}$ such that for any $j, j' \in N$*

$$\sum_i y_i u_i(a_j) = \sum_i y_i u_i(a_{j'}),$$

(ii) *or if for any $j \in N$ $u_i(a_i) > u_i(a_j)$ and*

$$\sum_i u_i(a_j) > \sum_i u_i^{sb},$$

where for any $i \in N$ $u_i^{sb} \equiv \max_{j\neq i} u_i(a_j)$.

A special case for the conditions for Theorem 4 is the one that the payoffs of each agent are constant at every agent-action that is not her optimal one.

Corollary 5 *Consider any environment with $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, $A = \{a_0, a_1, \ldots, a_n\}$, and suppose that Assumption 1 and 3 hold with the additional restriction that for any $i, j, j' \in N$ $u_i(a_i) > u_i(a_j) = u_i(a_{j'}) > u_i(a_0)$, and $x_0\mu_0 > x_i\mu_i$ holds for the Receiver. Then, there exists an endorsement rule that is incentive-compatible, Pareto-efficient, and beneficial relative to no and full information benchmarks.*

When there are only two agents, the conditions in Theorem 4 reduce to that for any $i, j \in N$ with $i \neq j$ $u_i(a_i) > u_i(a_j) > u_i(a_0)$. However, there is no straightforward generalization when there are more than two agents.

Example 4 *Consider an environment with $n = 3$, $\Omega = \{\omega_0, \omega_1, \omega_2, \omega_3\}$, $A = \{a_0, a_1, a_2, a_3\}$. Suppose that the payoffs of the agents are as follows: for any $i \in N$ $u_i(a_i) = 5$, $u_1(a_2) = u_3(a_2) = 4$, $u_1(a_3) = u_2(a_1) = u_2(a_3) = u_3(a_1) = 1$, while $u_i(a_0) = 0$. For the receiver, $u_r(a_i, \omega_j) = 1$ if $i = j$ and 0 otherwise. The prior belief is $\mu_0 = (0.28, 0.24, 0.24, 0.24)$.*

In the example above, any endorsement rule provides no greater than 3.8 as a payoff to agent 1 or 3 but agent 2's individually optimal information structure provides 3.84 to both agents and clearly more to agent 2.[7] Intuitively, the agent-action a_2 generates so much payoff to all agents that it is easier to agree on giving agent 2 full bargaining power than implementing each agent action with equal probability.

[7] The best endorsement rule for agent 1 recommends action 1 with full probability at state 0 and leads to $5(0.24 + 0.28) + 0.24(4 + 1) = 3.8$, while the agent 2-optimal incentive compatible information structure recommends action 2 with probability $\frac{0.24}{0.28}$, which leads to $4 * 0.24 * 4 = 3.84$ as payoff to player 1.

5 An Application to Multiparty Elections

The role of information in elections when there are at most two strategic candidates has been extensively explored in the literature.[8] However, the literature on competition in multiparty elections is rather limited [See Shepsle, 2012]. In this section, we provide a simple model of informational competition among multiple parties. With the help of this discussion, it is possible to detect a mechanism where the opposition parties may prevail against an incumbent party with a pre-election coordination [Golder, 2006] on political campaigns.

Consider an electoral political regime, where there are $n+1$ political parties indexed as $\{0, ..., n\}$, which are in competition with each other for the upcoming presidential election. Suppose that party 0 is the incumbent, and parties $1, ..., n$ are in opposition. For simplicity assume that the electoral rule is simple plurality, and so the party that gets the highest share of votes wins the election[9]. Consider a scenario, where the opposition parties cannot win the election against the incumbent party without cooperating with each other because they have smaller voter bases compared to the incumbent party. Recent relevant examples of pre-electoral cooperation among opposition parties across the world include the one in Israel in 2019 - present [Makovsky, 2022], in Hungary in 2022 [Scheppele, 2022], and the one in Turkiye for the upcoming 2023 elections [Pitel, 2022]. For each party $i \in \{0, 1, \ldots, n\}$ let $\gamma_i \in (0, 1)$ denote the population share of partisan voters who always support party i; while, let γ_r be the population share of swing voters who respond to the information revealed during the election campaign period. Assume $\sum_{i=1}^{n} \gamma_i < \gamma_0 < \gamma_i + \gamma_r$ for any opposition party i so that if an opposition party persuades all of the swing voters, it can win the election. Assume that the payoff of an opposition party $i \in \{1, \ldots, n\}$ is 1 if it wins the election, $\beta_i < 1$ if some other opposition party wins, and 0 is the incumbent party wins.

Let $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$ be the states of the world, where the realization of ω_i means that party i's policy is the best for the swing voters. Assume that there is a common prior belief, μ_0, among the swing voters that favors the incumbent party so that if no further information is revealed during the campaign period, the swing voters will support the incumbent party. If the opposition parties cannot agree on a joint campaign, the incumbent party dominates the narratives formed during the campaign period and prevents the revelation of any significant information during the campaign period, which leads to its electoral victory. To prevent that outcome the opposition parties can run a joint election campaign that minimizes the probability that the incumbent wins the election. We concentrate on the case where the election campaign generates public information and the voting behavior of swing voters are symmetric so that we can treat them as a single receiver.

The proposition below summarizes some of the implications of our general model above for this environment.

Proposition 3 . *Suppose that Assumptions 1 and 3 hold and fix any symmetric action profile of the swing voters* $a(\mu)_{\mu \in \Delta(\Omega)}$. *For any* $i, j \in \{1, \ldots, n\}$,

[8] Some of the recent studies that use the Bayesian persuasion framework are Alonso and Câmara [2016a,b] Schnakenberg [2017], Bardhi and Guo [2018], Chan et al. [2019].

[9] Alternatively, one can think of a legislative election, where the party who gets the highest share of votes gets a disproportionate advantage in terms of the share of seats in the legislative body.

(i) *if* $\beta_i < 0$ *but* $1 + \sum_{i=1}^{n} \beta_i > 0$*, then there is a joint campaign that benefits all opposition parties against the no-information benchmark.*

(ii) *if* $\beta_i > 0$*, then there is a joint campaign, which endorses the opposition party* i *at the state* ω_i*, that is beneficial against the no-information benchmark and Pareto efficient among the opposition parties.*

(iii) *if* $\beta_i = \beta_j > 0$*,* μ_0 *assigns equal likelihood to each of the states* $\{\omega_1, \ldots, n\}$*, and the swing voters are ex-ante indifferent among the opposition parties, the endorsement rule above corresponds to the symmetric Nash Bargaining rule among the opposition parties.*

We skip the proof as the parts (i) to (iii) of Proposition 3 directly follow from Corollary 3, Theorem 4, and Proposition 2 respectively.

Proposition 3 above shows a way out for the opposition parties without forming a full coalition by supporting a single candidate against the incumbent. For instance, when $\beta_i < 0$ for any opposition party i, it might be hard for the opposition parties to form a joint electoral platform against the incumbent and rule as a coalition party even after winning the election. Instead, a pre-electoral cooperation among the opposition parties on an informative endorsement rule can substantially improve the electoral outcome for them. With an endorsement rule each opposition party can win the election with some positive probability. Furthermore, since this type of a campaign provides full information at all states in $\{\omega_1, \ldots, \omega_n\}$, it significantly improves the welfare of the swing voters.

6 Conclusion

This paper analyzes how a group of agents collectively decides on ways to persuade a receiver. We prove results that characterize the possibility of agreement between agents. The overarching conclusion is that agreement is possible when every player can significantly gain from some belief of the receiver compared to the status quo. We also study more specific situations where every agent has a state that she would like to disclose to the receiver because the correct action in this state benefits her. The receiver is trying to guess the state. In such cases, an information structure that reveals agent-preferred states is Pareto-efficient and incentive compatible. Also, if the status quo action does not correspond to any of such states, this rule benefits all agents relative to complete censorship. We call this agreement the "endorsement rule." Finally, we discuss some political economy applications. However, our model pertains to many other scenarios, including but not limited to industrial and financial regulation [Goldstein and Leitner, 2018], and contest design [Zhang and Zhou, 2016, Antsygina and Teteryatnikova, 2022].

We assume that once the agents agree on an informational design, none of them can deviate. More specifically, agents can not share extra information on top of what the group decided to provide. This assumption is natural in authoritarian settings, where informational competition is restricted. However, in other situations, it might be less plausible, and it is interesting which informational agreements are robust to ex-post individual deviations. The answer is straightforward given the assumptions of Theorem 3. In this setting, every agent has a state in which the receiver's optimal action gives her the highest possible benefit. These states differ across agents. Finally, there is a state such that all agents would like to avoid the corresponding action. Given these assumptions, the endorsement rule is robust to individual deviations. If a state arises that some agent prefers, the rule will recommend the receiver to take the corresponding action. Hence, the receiver will think that the true state is either this one or the one that all agents want to avoid. Hence, no agent can persuade the receiver to take her preferred action contrary to the recommendation. Extra information could make the receiver choose the worst action for all agents. However, no agent would share it. Conversely, if the information structure is not the endorsement rule, it misleads the receiver in a state that some agent would prefer to disclose. Hence, no other information structure in this setting is robust to individual deviations. Identification of environments where endorsement rules are the only information structures robust to ex-post deviations is an interesting open question.

References

Ricardo Alonso and Odilon Câmara. Persuading voters. *American Economic Review*, 106(11): 3590–3605, 2016a.

Ricardo Alonso and Odilon Câmara. Political disagreement and information in elections. *Games and Economic Behavior*, 100:390–412, 2016b.

Anastasia Antsygina and Mariya Teteryatnikova. Optimal information disclosure in contests with stochastic prize valuations. *Economic Theory*, pages 1–38, 2022.

Pak Hung Au and Keiichi Kawai. Competitive information disclosure by multiple senders. *Games and Economic Behavior*, 119:56–78, 2020.

Arjada Bardhi and Yingni Guo. Modes of persuasion toward unanimous consent. *Theoretical Economics*, 13(3):1111–1149, 2018.

Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences.* Computer science and applied mathematics. Elsevier Inc, Academic Press, 1979.

Sourav Bhattacharya and Arijit Mukherjee. Strategic information revelation when experts compete to influence. *The RAND Journal of Economics*, 44(3):522–544, 2013.

Simon Board and Jay Lu. Competitive information disclosure in search markets. *Journal of Political Economy*, 126(5):1965–2010, 2018.

Raphael Boleslavsky, Mehdi Shadmehr, and Konstantin Sonin. Media freedom in the shadow of a coup. *Journal of the European Economic Association*, 19(3):1782–1815, 2021.

Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

Jimmy Chan, Seher Gupta, Fei Li, and Yun Wang. Pivotal persuasion. *Journal of Economic theory*, 180:178–202, 2019.

Laura Doval and Alex Smolin. Information payoffs: An interim perspective. *arXiv preprint arXiv:2109.03061*, 2021.

Scott Gehlbach and Konstantin Sonin. Government control of the media. *Journal of public Economics*, 118:163–171, 2014.

Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2016.

Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.

Sona Nadenichek Golder. *The logic of pre-electoral coalition formation*. Ohio State University Press, 2006.

Itay Goldstein and Yaron Leitner. Stress tests and information disclosure. *Journal of Economic Theory*, 177:34–69, 2018.

Sergei Guriev and Daniel Treisman. A theory of informational autocracy. *Journal of public economics*, 186:104158, 2020.

Florian Hoffmann, Roman Inderst, and Marco Ottaviani. Persuasion through selective disclosure: implications for marketing, campaigning, and privacy regulation. *Management Science*, 66(11): 4958–4979, 2020.

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

Fei Li and Peter Norman. On bayesian persuasion with multiple senders. *Economics Letters*, 170: 66–70, 2018.

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

Karl E Loewenstein. Re-emergence of public opinion in the soviet union: Khrushchev and responses to the secret speech. *Europe-Asia Studies*, 58(8):1329–1345, 2006.

David Makovsky. Israeli elections, round five: A game of inches, Sep 2022. URL `https://www.washingtoninstitute.org/policy-analysis/israeli-elections-round-five-game-inches`.

Cherng-tiao Perng. On a class of theorems equivalent to farkas's lemma. *Applied Mathematical Sciences*, 11(44):2175–2184, 2017.

Laura Pitel. Turkish opposition's alliance against erdogan battles rifts, Sep 2022. URL `https://www.ft.com/content/5b391b66-834d-4d2b-8f0c-b8d6f64392ee`.

Donnalee Rowe. Khrushchev's secret speech and the aftermath. 1964.

Kim Lane Scheppele. How viktor orbán wins. *Journal of Democracy*, 33(3):45–61, 2022.

Keith E Schnakenberg. Informational lobbying and legislative voting. *American Journal of Political Science*, 61(1):129–145, 2017.

Kenneth Shepsle. *Models of multiparty electoral competition*. Routledge, 2012.

William Thomson. Cooperative models of bargaining. *Handbook of game theory with economic applications*, 2:1237–1284, 1994.

Jun Zhang and Junjie Zhou. Information disclosure in contests: A bayesian persuasion approach. *The Economic Journal*, 126(597):2197–2217, 2016.

A Proofs

A.1 Bargaining Set

Lemma 1 *Suppose that Assumption 1 hold. The bargaining set B is compact and convex.*

Proof B is bounded as A is finite and Π is compact by Remark 1. Now, take any converging sequence of incentive compatible information structures $\{\pi_l\}_{l\geq 1}$ with $\lim \pi_l = \bar{\pi} \in \Pi$. Then, for any agent $i \in \{1, \ldots, n\}$

$$\lim \mathbb{E}_{\pi_l} \bar{v}_i(\mu_a) = \sum_{a\in A}\sum_{\omega\in\Omega} \lim \pi_l(a,\omega)\mu_0(\omega)u_i(a,\omega) = \mathbb{E}_{\bar{\pi}}\bar{v}_i(\mu_a),$$

where the second equality follows from Assumption 1. Finally, to prove convexity take any two payoff vectors v, v' such that

$$v = (\mathbb{E}_{\pi}\bar{v}_i(\mu_a))_{i\in\{1,\ldots,n\}} \quad \text{and} \quad v' = (\mathbb{E}_{\pi'}\bar{v}_i(\mu_a))_{i\in\{1,\ldots,n\}}$$

for some two incentive compatible information structures π and π'. Fix an arbitrary $\lambda \in (0,1)$. Consider the following information structure π_λ that applies π with probability λ and π' with the remaining probability. Clearly, π_λ is an incentive compatible information structure and for any $i \in \{1, \ldots, n\}$

$$\mathbb{E}_{\pi_\lambda}\bar{v}_i(\mu_a) = \lambda\mathbb{E}_{\pi}\bar{v}_i(\mu_a) + (1-\lambda)\mathbb{E}_{\pi'}\bar{v}_i(\mu_a),$$

which proves the convexity of B. ■

Fix any optimal action profile $\{\bar{a}(\mu)\}_{\mu\in\Delta(\Omega)}$ by the receiver. Recall that $v_i(\mu) = \mathbb{E}_\mu u_i(a_\mu, \omega)$. Let

$$\hat{v}_i(\mu) = \sum_{\omega\in\Omega} \frac{\mu(\omega)}{\mu_0(\omega)} u_i(a_\mu, \omega),$$

and $\hat{v}(\mu) = (\hat{v}_1(\mu)), \ldots, \hat{v}_n(\mu))$.

Proposition 4 *Suppose that Assumption 1 holds. Then,*

$$B = \{x \in \mathbb{R}^n | (\mu_0, x) \in co(graph(\hat{v}))\}$$

Proof The proof follows the same steps as the arguments provided by Doval and Smolin [2021] but we provide the arguments below for completeness.

Take any payoff profile $v \in B$. By definition, there exists an incentive compatible information structure π such that for any agent $i \in \{1, \ldots, n\}$ $x_i = \mathbb{E}_\pi \bar{v}_i(\mu_\pi)$ and π leads to a Bayesian plausible distribution over beliefs. To see this let $Pr_\pi(a) = \sum_{\omega \in \Omega} \pi(a, \omega)\mu_0(\omega)$, and the expected posterior belief induced by π at any state $\omega \in \Omega$

$$\mathbb{E}_\pi \mu(\omega) = \sum_{a \in A} Pr_\pi(a)\mu_a(\omega) = \sum_{a \in A} Pr_\pi(a) \frac{\pi(a, \omega)\mu_0(\omega)}{Pr_\pi(a)} = \mu_0(\omega) \sum_{a \in A} \pi(a, \omega) = \mu_0(\omega).$$

Now, for any agent i

$$\begin{aligned} x_i &= \mathbb{E}_\pi \bar{v}_i(\mu_\pi) = \sum_{a \in A} \sum_{\omega \in \Omega} \pi(a, \omega)\mu_0(\omega) u_i(\mu_a, \omega) \\ &= \sum_{a \in A} Pr_\pi(a) \sum_{\omega \in \Omega} \frac{\mu_0(\omega)\pi(a, \omega)}{Pr_\pi(a)} \frac{u_i(a, \omega)}{\mu_0(\omega)} = \sum_{a \in A} Pr_\pi(a) \sum_{\omega \in \Omega} \frac{\mu_a(\omega)}{\mu_0(\omega)} u_i(a, \omega) \\ &= \sum_{a \in A} Pr_\pi(a) \hat{v}_i(\mu_a), \end{aligned}$$

which implies that $(\mu_0, v) \in co(graph(\hat{v}))$, where the weights are given by $(Pr_\pi(a))_{a \in A}$. Reversely, for any element of $co(graph(\hat{v}))$, we can find a probability distribution over posterior beliefs and the corresponding values of $\hat{v}$. Then, the calculation above enables us to interpret the weights as probabilities given by an incentive compatible information structure, which further implies that the corresponding payoff profile is in B. ■

A.2 Beneficial Agreement

We first prove the first part of Theorem 1 as stated by the following Lemma.

Lemma 2 *Suppose that Assumptions 1 and 3 hold, the default information structure π_0 is discrete, and fix any optimal action profile $(\bar{a}(\mu))_{\mu \in \Delta(\Omega)}$ by the receiver. There is beneficial agreement if and only if there is a lottery over beliefs $(\lambda_l, \mu_l)_{l=1}^h$ such that for any $l = 1, \ldots, h$ $\lambda_l \in [0, 1]$ and $\sum_l \lambda_l = 1$ and for any agent $i \in N$*

$$\sum_l \lambda_l \bar{v}_i(\mu_l) > d_i. \tag{6}$$

Proof Fix any discrete and incentive compatible information structure π_0, and let τ_0 be the induced distribution over posterior beliefs with support $\{\mu_{0k}\}_{k=1,\ldots,g}$.

($\Leftarrow$) Fix $\varepsilon > 0$ small enough so that for each $k = 1, \ldots, g$ the belief $\hat{\mu}_k = \varepsilon\mu_0 + (1-\varepsilon)\mu_{0k}$ satisfies $\arg\max\{\bar{v}_r(\hat{\mu}_k)\} = \arg\max\{\bar{v}_r(\mu_{0k})\}$, and $\hat{\mu}_k$ is interior as μ_0 is interior. Note that such ε exists because π_0 is discrete and so at each μ_{0k} there is a unique receiver-optimal action.

Let's define $\hat{\tau}_0$ as for each $k = 1, \ldots, g$ $\hat{\tau}_0(\hat{\mu}_k) = \tau_0(\mu_{0k})$. Note that τ_0 is Bayesian plausible by construction as it is induced by an information structure π_0. Then,

$$\sum_{k=1}^{g} \hat{\tau}(\hat{\mu}_k)\hat{\mu}_k = \sum_{k=1}^{g} \hat{\tau}(\hat{\mu}_k)(\varepsilon\mu_0 + (1-\varepsilon)\mu_{0k})$$
$$\varepsilon\mu_0 + (1-\varepsilon)\sum_{k=1}^{g} \tau_0(\mu_{0k})\mu_{0k} = \mu_0.$$

Note further that for any agent $i \in N$,

$$\sum_{k=1}^{g} \hat{\tau}_0(\hat{\mu}_k)\bar{v}_i(\hat{\mu}_k) = \sum_{k=1}^{g} \hat{\tau}_0(\hat{\mu}_k)u_i(a(\hat{\mu}_k)) = \sum_{k=1}^{g} \tau_0(\mu_k)u_i(a(\mu_k)),$$

where the first equality holds by Assumption 3, and the second holds as $\hat{\tau}_0(\hat{\mu}_k) = \tau_0(\mu_k)$ and $a(\hat{\mu}_k) = a(\mu_k)$ for any $k = 1, \ldots, g$.

Recall that $(\lambda_l, \mu_l)_{l=1}^{h}$ is the lottery over beliefs that yields a higher expected payoff for all agents compared to the disagreement payoff. For any belief $\hat{\mu}_k \in supp(\hat{\tau}_0)$, μ_l define $\eta_{lk} \in \Delta(\Omega)$ such that

$$\gamma_{lk}\mu_l + (1-\gamma_{lk})\eta_{lk} = \hat{\mu}_k,$$

and $a(\eta_{lk}) = a(\hat{\mu}_k)$. Note that such beliefs $\{\eta_{lk}\}$ and probabilities $\{\gamma_{lk}\}$ exist as each $\hat{\mu}_k$ is interior and the information structure defined by $\hat{\tau}_0$ is discrete.

For all l define $\gamma_l' \equiv \min_k \gamma_{lk}$. Also define $\epsilon \equiv \min_l \gamma_l'$. Consider an information structure π as follows. For any k, with probability $\hat{\tau}_0(\hat{\mu}_k)\frac{\epsilon\lambda_l}{\gamma_l'}$, it runs the lottery $((\gamma_{lk}, 1-\gamma_{lk}), (\mu_l, \eta_{lk}))$, and with probability $\hat{\tau}_0(\hat{\mu}_k)(1 - \sum_l \frac{\epsilon\lambda_l}{\gamma_l'})$ the belief $\hat{\mu}_k$ is realized. Observe the following

- $\frac{\epsilon}{\gamma_l'} < 1$ and $\sum_l \lambda_l = 1 \Rightarrow \sum_l \frac{\epsilon\lambda_l}{\gamma_l'} < 1$
- The belief distribution induced by the information structure π is Bayesian plausible as

$$\sum_k \hat{\tau}_0(\hat{\mu}_k)\left[\sum_l \frac{\epsilon\lambda_l}{\gamma_l'}(\gamma_{lk}\mu_l + (1-\gamma_{lk})\eta_{lk}) + \left(1 - \sum_l \frac{\epsilon\lambda_l}{\gamma_l'}\right)\hat{\mu}_k\right] =$$
$$\sum_k \hat{\tau}_0(\hat{\mu}_k)\left[\sum_l \frac{\epsilon\lambda_l}{\gamma_l'}\hat{\mu}_k + \left(1 - \sum_l \frac{\epsilon\lambda_l}{\gamma_l'}\right)\hat{\mu}_k\right] = \sum_k \hat{\tau}_0(\hat{\mu}_k)\hat{\mu}_k = \mu_0$$

Note that the last equality follows by construction of $\hat{\tau}_0$.

- For any agent i

$$\begin{aligned}
\mathbb{E}_\pi[u_i(a)] &= \sum_k \hat{\tau}_0(\hat{\mu}_k)\left[\sum_l \frac{\epsilon\lambda_l}{\gamma_l'}(\gamma_l' u_i(a(\mu_l)) + (1-\gamma_l')u_i(a(\hat{\mu}_k))) + \left(1 - \sum_l \frac{\epsilon\lambda_l}{\gamma_l'}\right) u_i(a(\hat{\mu}_k))\right] \\
&= \sum_k \hat{\tau}_0(\hat{\mu}_k)\left[\epsilon \sum_l \lambda_l u_i(a(\mu_l)) + u_i(a(\hat{\mu}_k))\left(\sum_l \frac{\epsilon\lambda_l(1-\gamma_l')}{\gamma_l'} + 1 - \sum_l \frac{\epsilon\lambda_l}{\gamma_l'}\right)\right] \\
&= \sum_k \hat{\tau}_0(\hat{\mu}_k)\left[\epsilon \sum_l \lambda_l u_i(a(\mu_l)) + u_i(a(\hat{\mu}_k))\left(1 - \epsilon\sum_l \lambda_l\right)\right] \\
&= \epsilon \sum_l \lambda_l u_i(a(\mu_l)) + \sum_k \hat{\tau}_0(\hat{\mu}_k) u_i(a(\hat{\mu}_k))\,(1-\epsilon) > \sum_k \hat{\tau}_0(\hat{\mu}_k) u_i(a(\hat{\mu}_k)) = d_i
\end{aligned}$$

  Hence, the information structure π is beneficial.

($\Rightarrow$) If there is a beneficial agreement, then there is a beneficial incentive compatible information structure π and a corresponding Bayesian plausible distribution τ over beliefs. ■

For the second part of Theorem 1, we prove the following Lemma.

Lemma 3 *Suppose that Assumptions 1 and 3 hold, the default information structure π_0 is discrete, and fix any optimal action profile $(a_\mu)_{\mu\in\Delta(\Omega)}$ by the receiver. Exactly one of the following is true:*

(i): *There is a beneficial agreement.*

(ii): *There exist non-negative numbers $y_1, ..., y_n$ which are not all equal to 0 such that for any belief $\mu \in \Delta(\Omega)$*

$$\sum_{i\in N} y_i(\bar{v}_i(\mu) - d_i) \leq 0. \tag{7}$$

Proof Applying Theorem 21 in Perng [2017], we get that for any matrix X exactly one of the following is true

- There is a non-negative vector λ such that $X\lambda > 0$.
- There is a non-negative and non-zero vector y such that $X^T y \leq 0$,

where X^T is the transpose of the matrix. For any finite collection of beliefs $\eta = \{\mu_i\}_{i=1}^h$, set

$$X(\eta) \equiv (x_{ij}) = \bar{v}_i(\mu_j) - d_i$$

In other words, an entry at ith row and jth column is the difference between agent i's expected payoff given belief μ_j and her disagreement payoff.

Suppose that there exists a beneficial information structure. By Lemma 2, there exists a finite collection of beliefs $\eta = \{\mu_i\}_{i=1}^h$ and a vector of non-negative numbers λ such that for any agent i

$$\sum_j \lambda_j \bar{v}_i(\mu_j) > d_i \Leftrightarrow \sum_j \lambda_j(\bar{v}_i(\mu_j) - d_i) > 0$$

This statement is equivalent to there being a non-negative vector $\lambda = (\lambda_1, ..., \lambda_n)$ such that $X(\eta)\lambda > 0$. Then, Theorem 21 in Perng [2017] implies that there is no non-negative and non-zero vector y such that $X(\eta)^T y \leq 0$. Hence, for any non-negative numbers $y_1, ..., y_n$ that are not all equal to 0 there exists a belief $\mu_j \in \eta$ such that

$$\sum_i y_i(\bar{v}_i(\mu_j) - d_i) > 0$$

and the second condition is not satisfied.

Suppose now that there are non-negative numbers $y_1, ..., y_n$ that are not all equal to 0 such that for any agent j and belief μ

$$\sum_j y_j(\bar{v}_i(\mu_j) - d_j) \leq 0$$

Then for any collection of beliefs $\eta = \{\mu_j\}_{j=1}^h$, $X(\eta)^T y \leq 0$ where $y = (y_1, ..., y_n)$. The Theorem 21 in Perng [2017] implies that there is no non-negative vector λ such that $X(\eta)\lambda > 0$. Hence, there is no beneficial agreement by Lemma 2. ■

The proof of Theorem 2 follows from the Lemma 4 below.

Lemma 4 *Fix any $k \in \{1, \ldots, n\}$. Let $\Phi(k)$ be set of all subsets of $\{1, \ldots, n\}$ of size k. For each $\phi \in \Phi(k)$, fix a single vector $s^\phi \in \mathbb{R}^n$ such that*

$$\sum_{j=1}^{n} s^\phi(j) > 0$$
$$\forall j \in \phi \quad s^\phi(j) > 0 \quad \text{but} \quad \forall j \notin \phi \quad s^\phi(j) < 0.$$

Let $S = \{s^\phi\}_{\phi \in \Phi(k)}$. Then, there exists $s \in co(S)$ such that $s > 0$.

Proof: Fix any $k \in \{1, \ldots, n\}$ and let S be constructed as in the hypothesis. Note that for any $\phi \in \Phi(k)$, there is at most one $s^\phi \in S$ that satisfies the properties stated in the hypothesis. Equivalently, for any ϕ, ϕ' with $\phi \neq \phi'$ and $s^\phi, s^{\phi'} \in S$ there is $j \in \{1, \ldots, n\}$ such that $s^\phi(j) > 0 > s^{\phi'}(j)$.

Suppose for a contradiction that $\mathbb{R}^n_+ \cap co(S) = \emptyset$. As $\mathbb{R}^n_+$ is a convex set, by the Hyperplane Separation Theorem (Boyd et al. [2004]) there exists $c \in \mathbb{R}$, $\nu \in \mathbb{R}^n \backslash \{\vec{0}\}$ such that for any $x \in \mathbb{R}^n_+$ and $y \in co(S)$ $\langle x, \nu \rangle \geq c \geq \langle y, \nu \rangle$.

Now, we prove a couple of claims.

Claim 1: $c \leq 0$.

Proof For any $\lambda > 0$, $\lambda\vec{1} \in \mathbb{R}^n_+ \Rightarrow \langle \lambda\vec{1}, \nu \rangle \geq c \Rightarrow \lim_{\lambda \to 0} \langle \lambda\vec{1}, \nu \rangle = 0 \geq c$. □

Claim 2: $\nu \geq 0$.

Proof Suppose to the contrary that there exists $i \in \{1, \ldots, n\}$ with $\nu(i) < 0$. Fix $\lambda > 0$ and $t \in \mathbb{R}^n_+$ with $t(i) = \lambda$ but for any $j \neq i$ $t(j) = \varepsilon > 0$ for some ε small enough. Then, $\lim_{\lambda \to \infty} \langle t, \nu \rangle = -\infty < c$, a contradiction. □

Given the claims above, WLOG $\nu(1) \geq \nu(2) \geq \ldots \geq \nu(n)$ and $\phi \in \{1, \ldots, k\}$. Observe that for

$\phi = \{1, \ldots, k\}$,

$$\langle s^\phi, \nu \rangle = \sum_{j=1}^{n} s^\phi(j)\nu(j) = \sum_{j=1}^{k} s^\phi(j)\nu(j) + \sum_{j=k+1}^{n} s^\phi(j)\nu(j).$$

Case 1: $\nu(k) > 0$. As $\forall j \leq k$ $s^\phi(j) > 0$ and $j > k$ $s^\phi(j) < 0$,

$$\sum_{j=1}^{k} s^\phi(j)\nu(j) \geq \sum_{j=1}^{k} s^\phi(j)\nu(k), \text{ and}$$

$$\sum_{j=k+1}^{n} s^\phi(j)\nu(j) \geq \sum_{j=k+1}^{n} s^\phi(j)\nu(k)$$

$$\Rightarrow \langle s^\phi, \nu \rangle \geq \nu(k) \sum_{j=1}^{n} s^\phi(j) > 0 \geq c,$$

a contradiction.

Case 2: $\nu(k) = 0$. Then $\forall j > k$ $\nu(j) = 0$ and since $\nu \neq \vec{0}$ $\nu(1) > 0$. $\Rightarrow \sum_{j=1}^{k} s^\phi(j)\nu(j) > 0 = \sum_{j=k+1}^{n} s^\phi(j)\nu(j) \Rightarrow \langle s^\phi, \nu \rangle > 0 \geq c$, a contradiction. ■

Proof of Theorem 2

Suppose that there is k and $M \subset \Delta(\Omega)$ as described in the hypothesis. By Lemma 4 there exists a collection of weights $\{\lambda_l\}_{l=1,\ldots,|\Phi(k)|}$ with $\sum_l \lambda_l = 1$ such that for each agent $i = 1, \ldots, n$

$$\sum_l \lambda_l \bar{v}_i(\mu_l) - d_i > 0.$$

Then, by Theorem 1, there is a beneficial agreement. ■

Proof of Proposition 1

Fix any discrete, interior, and incentive compatible information structure π_0, and let τ_0 be the induced distribution over posterior beliefs with support $\{\mu_{0k}\}_{k=1,\ldots,g}$. For each μ_l, we can find $\gamma_{kl} \in (0,1)$ and η_{kl} with $a(\eta_{kl}) = a(\mu_k)$ such that

$$\mu_k = \gamma_{kl}\eta_{kl} + (1 - \gamma_{kl})\mu_l.$$

Fix any $\epsilon \in (0, \min_{k,l}(1 - \gamma_{kl}))$. Consider the following information structure that implements μ_k with probability

$$\tau_0(\mu_k)\left(1 - \sum_{l=1}^{h} \frac{\epsilon\lambda_l}{1 - \gamma_{kl}}\right),$$

and implements η_{kl} with $\tau_0(\mu_k)\frac{\epsilon\lambda_l}{1-\gamma_{kl}}\gamma_{kl}$, while μ_l with $\tau_0(\mu_k)\frac{\epsilon\lambda_l}{1-\gamma_{kl}}(1 - \gamma_{kl})$. Note that there is such an information structure, since the induced distribution over the beliefs is Bayesian plausible.

Now, the expected payoff of any agent $i \in N$ under this information structure is

$$\sum_{k=1}^{g} \tau_0(\mu_k) \sum_{l=1}^{h} \left(\frac{\epsilon\lambda_l}{1-\gamma_{kl}} \left(\gamma_{kl}\mathbb{E}_{\eta_{kl}} u_i(a(\eta_{kl}),\omega) + (1-\gamma_{kl})\mathbb{E}_{\mu_l} u_i(\mu_l,\omega)\right)\right) + \sum_{k=1}^{g} \tau_0(\mu_k) \left(1 - \sum_{l=1}^{h} \frac{\epsilon\lambda_l}{1-\gamma_{kl}}\right) \mathbb{E}_{\mu_k} u_i(a(\mu_k),\omega). \quad (8)$$

The second terms with $\mathbb{E}_{\mu_l} u_i(\mu_l,\omega)$ simplifies and satisfies the following inequality by the assumption in the hypothesis

$$\epsilon \sum_{k=1}^{g} \tau_0(\mu_k) \sum_{l=1}^{h} \lambda_l \mathbb{E}_{\mu_l} u(a(\mu_l),\omega) \ = \ \epsilon \sum_{l=1}^{h} \lambda_l \mathbb{E}_{\mu_l} u(a(\mu_l),\omega) \ > \ \epsilon \sum_{k=1}^{g} \tau_0(\mu_k) \sum_{l=1}^{h} \lambda_l \mathbb{E}_{\mu_l} u(a(\mu_k),\omega).$$

Then, combining the right-hand side of the inequality above with the first term in the expression (8) yields the following

$$\sum_{k=1}^{g} \tau_0(\mu_k) \sum_{l=1}^{h} \left(\frac{\epsilon\lambda_l}{1-\gamma_{kl}} \left(\gamma_{kl}\mathbb{E}_{\eta_{kl}} u_i(a(\eta_{kl}),\omega) + (1-\gamma_{kl})\mathbb{E}_{\mu_l} u_i(a(\mu_k),\omega)\right)\right),$$

but as $a(\eta_{kl}) = a(\mu_k)$ and expectation is linear in beliefs when the receiver's action is fixed we get

$$\sum_{k=1}^{g} \tau_0(\mu_k) \left(\sum_{l=1}^{h} \frac{\epsilon\lambda_l}{1-\gamma_{kl}}\right) \mathbb{E}_{\mu_k} u_i(a(\mu_k),\omega).$$

Combining the expression above with the third term in the expression (8) shows that the expected payoff for agent i is strictly greater than the following

$$\sum_{k=1}^{g} \tau_0(\mu_k) \left(\sum_{l=1}^{h} \frac{\epsilon\lambda_l}{1-\gamma_{kl}}\right) \mathbb{E}_{\mu_k} u_i(a(\mu_k),\omega) + \sum_{k=1}^{g} \tau_0(\mu_k) \left(1 - \sum_{l=1}^{h} \frac{\epsilon\lambda_l}{1-\gamma_{kl}}\right) \mathbb{E}_{\mu_k} u_i(a(\mu_k),\omega) = \mathbb{E}_{\tau_0}\mathbb{E}_{\mu} u_i(a(\mu),\omega).$$

This establishes that this information structure is beneficial compared to π_0. ■

A.3 Endorsement Rules

Proof of Proposition 2

Let's first define a bargaining solution over B that corresponds to the symmetric endorsement rule. For any compact and convex subset $B' \subseteq B$, let

$$f^{ser}(B') = \begin{cases} \vec{\mathbb{E}}_{ser} & \text{if } \vec{\mathbb{E}}_{ser} \in B' \\ \vec{\mathbb{E}}_{NB} & \text{if } \vec{\mathbb{E}}_{ser} \notin B', \end{cases}$$

where $\vec{\mathbb{E}}_{ser}$ is the vector of payoffs to all agents corresponding to the symmetric endorsement payoff and $\vec{\mathbb{E}}_{NB}$ is the payoff vector that corresponds to the Nash bargaining solution from any subset B'. We show below that the solution $f^{ser}(B')$ satisfies the axioms that characterize the Nash bargaining solution. These axioms are Pareto efficiency, symmetry, scale invariance, and contraction independence (independence of irrelevant alternatives) as described in Thomson [1994]. Symmetry, scale invariance, and contraction independence are trivially satisfied as the payoffs at symmetric endorsement rule are symmetric, linear in payoff parameters, and constant over all the subsets B' as long as $\vec{\mathbb{E}}_{ser} \in B'$. We prove the Pareto efficiency below.

To simplify the notation, write $\mu_i \equiv \mu_0(\omega_i)$. Fix an arbitrary information structure $\pi \in \Pi$. Write the sum of agents' expected payoffs as

$$\sum_{j,i} \mu_i \pi_{ji} \sum_k u_k(a_j, \omega_i) = \sum_{i \neq 0} \mu_i \pi_{ii} \sum_k u_k(a_i, \omega_i) + \sum_{i \neq 0} \mu_i \sum_{j \neq i,0} \pi_{ji} \sum_k u_k(a_j, \omega_i) +$$
$$\sum_{i \neq 0} \mu_i \pi_{0i} \sum_k u_k(a_0, \omega_i) + \mu_0 \sum_{j \neq 0} \pi_{j0} \sum_k u_k(a_j, \omega_0) +$$
$$\mu_0 \pi_{00} \sum_k u_k(a_0, \omega_0)$$

Using the properties of the symmetric environment, write

$$v_1 \equiv \sum_k u_k(a_i, \omega_i)$$

where $i \neq 0$ and

$$v_2 \equiv \sum_k u_k(a_j, \omega_i)$$

where $j \neq 0, i$ and $i \neq 0$

$$v_3 \equiv \sum_k u_k(a_0, \omega_i)$$

where $i \neq 0$ and

$$v_4 \equiv \sum_k u_k(a_j, \omega_0)$$

where $j \neq 0$ and

$$v_5 \equiv \sum_k u_k(a_0, \omega_0)$$

Recall that for any $i, j \neq 0 : \mu_i = \mu_j$ and therefore $\mu_i = \frac{1-\mu_0}{n}$. The expression above becomes

$$\sum_{i \neq 0} \frac{(1-\mu_0)}{n} \pi_{ii} v_1 + \sum_{i \neq 0} \frac{(1-\mu_0)}{n} (1 - \pi_{ii} - \pi_{0i}) v_2$$
$$+ \sum_{i \neq 0} \frac{(1-\mu_0)}{n} \pi_{0i} v_3 + \mu_0 (1 - \pi_{00}) v_4 + \mu_0 \pi_{00} v_5 =$$
$$\frac{(1-\mu_0)}{n} \sum_{i \neq 0} [v_1 \pi_{ii} + v_2(1 - \pi_{ii} - \pi_{i0}) + v_3 \pi_{0i}] + \mu_0 [v_4 - (v_4 - v_5) \pi_{00}]$$

By assumption, $v_1 \geq \max\{v_2, v_3\}$, which implies

$$\frac{(1-\mu_0)}{n}\sum_{i\neq 0}\left[v_1\pi_{ii} + v_2(1-\pi_{ii}-\pi_{i0}) + v_3\pi_{0i}\right] + \mu_0\left[v_4 - (v_4 - v_5)\pi_{00}\right] \leq$$

$$(1-\mu_0)v_1 + \mu_0\left[v_4 - (v_4 - v_5)\pi_{00}\right]$$

Also notice that by assumption, $v_4 \geq v_5$, and therefore the expression $\mu_0\left[v_4(1-\pi_{00}) + \pi_{00}v_5\right]$ decreases in π_{00}. In the rest of the proof, we will use bounds on π_{00} to show that the symmetric endorsement rule generates the highest sum of expected payoffs for the agents.

If the recommended action is a_i for some $i \neq 0$, the Receiver prefers a_i over a_0 iff

$$x\frac{\pi_{ii}\mu_i}{\sum_k \pi_{ik}\mu_k} \geq x_0\frac{\pi_{i0}\mu_0}{\sum_k \pi_{ik}\mu_k} \Leftrightarrow x\pi_{ii}\mu_i \geq x_0\pi_{i0}\mu_0$$

Substituting $\mu_i = \frac{1-\mu_0}{n}$ yields

$$x\pi_{ii}\frac{1-\mu_0}{n} \geq x_0\pi_{i0}\mu_0 \Leftrightarrow x\pi_{ii}(1-\mu_0) \geq nx_0\pi_{i0}\mu_0$$

We proceed by analysing two cases. Suppose first that $\frac{x(1-\mu_0)}{x_0\mu_0} \geq 1$. Consider a symmetric endorsement rule such that $\pi_{00} = 0$ and for any $i \neq 0 : \pi_{i0} = \frac{1}{n}$. Substituting these values into the constraint above yields

$$x(1-\mu_0) \geq x_0\mu_0.$$

which holds by assumption. Clearly, given the symmetric endorsement rule, the receiver will not choose action a_j if the recommendation is a_i for $i, j \neq 0$. Hence, the receiver will follow all recommendations of the information structure. The sum of agents' expected payoffs will be $(1-\mu_0)v_1 + \mu_0 v_4$. Hence, a symmetric endorsement rule maximizes the sum of expected payoffs for the agents and is Pareto-efficient.

Suppose now that $\frac{x(1-\mu_0)}{x_0\mu_0} < 1$. Summing the constraints on π_{ii} and π_{0i} yields

$$\sum_{i\neq 0}\pi_{ii}\frac{x(1-\mu_0)}{n} \geq \sum_{i\neq 0}x_0\pi_{i0}\mu_0 \Leftrightarrow \frac{x(1-\mu_0)}{n}\sum_{i\neq 0}\pi_{ii} \geq x_0\mu_0(1-\pi_{00}) \Leftrightarrow$$

$$\pi_{00} \geq 1 - \sum_{i\neq 0}\frac{x(1-\mu_0)}{nx_0\mu_0}\pi_{ii} \geq 1 - \frac{x(1-\mu_0)}{x_0\mu_0}$$

It follows that

$$(1-\mu_0)v_1 + \mu_0\left[v_4 - \pi_{00}(v_4 - v_5)\right] \leq (1-\mu_0)v_1 + \mu_0\left[v_4 - \left(1 - \frac{x(1-\mu_0)}{x_0\mu_0}\right)(v_4 - v_5)\right]$$

Now, consider a symmetric endorsement rule such that for any

$$i \neq 0 : \pi_{i0} = \frac{x(1-\mu_0)}{nx_0\mu_0}, \text{ and } \pi_{00} = 1 - \frac{x(1-\mu_0)}{x_0\mu_0}.$$

Observe that for any $i \neq 0$:

$$x\pi_{ii}\mu_i = \frac{x(1-\mu_0)}{n} = x_0 \frac{x(1-\mu_0)}{x_0\mu_0 n}\mu_0 = x_0\pi_{i0}\mu_0$$

Hence, if the recommended action is a_i where $i \neq 0$, the Receiver will prefer action a_i over a_0. Observe that for any $i, j \neq 0 : \pi_{ji} = 0$ and $\pi_{0i} = 0$, therefore if the recommended action is a_i, the Receiver will prefer it over any action a_j. Hence, the π is an incentive-compatible information structure. The sum of agents' expected utilities will be exactly

$$(1-\mu_0)v_1 + \mu_0\left[v_4 - \left(1 - \frac{x(1-\mu_0)}{x_0\mu_0}\right)(v_4 - v_5)\right]$$

Hence, the symmetric endorsement rule maximizes the sum of expected payoffs and is therefore Pareto-optimal.

It follows from the argument above that an information structure given which the Receiver always stays with the prior yields a lower sum of expected payoffs for the agents compared to the information structure above. By symmetry, under both information structures, each agent gets the same payoff. Hence, the information structure above is beneficial for each agent.

We can also show that the symmetric endorsement rule corresponds to the Kalai-Smorodinsky and egalitarian bargaining solutions. To prove the first fact, consider a vector η where

$$\eta_i = \max_{\pi} \mathbb{E}_\pi \bar{v}_i(\mu_\pi)$$

By symmetry, for any i, j $\eta_i = \eta_j$ and $d_i = d_j$. Hence, the same is true for all vectors that lie on the line that connects the vectors η and $d = (d_1, \ldots, d_n)$. The Kalai-Smorodinsky solution selects the point on this line that intersects the boundary of the feasible payoff set (Thomson [1994]). The symmetric endorsement rule yields a payoff vector the coordinates of which are equal to each other. Hence, this vector belongs both to the line and the feasible payoff set. If it lies in the interior of the payoff set, then there is another point on the line where each coordinate is greater by ϵ, where ϵ is sufficiently small. This contradicts the fact that the endorsement rule maximizes the sum of agents' payoffs over all information structures. By the same argument, the endorsement rule yields a payoff vector that is a maximal point of equal coordinates among points in the feasible payoff set. Hence, it also corresponds to the egalitarian solution ((Thomson [1994]))

■

Proof of Corollary 4 Define v_1 to v_5 as in the proof of Proposition 2. Then, the inequality in (5) reduces to $v_1 \geq v_2$. A similar argument as in the proof of Proposition 2 shows that full disclosure is Pareto efficient. ■

Proof of Theorem 3 Take any incentive compatible information structure π that recommends all actions $\{a_1, \ldots, a_n\}$ with positive probability. Note that the set of such incentive compatible information structures is not empty as full-revelation information structure is always incentive compatible. Define an endorsement rule π^e such that for any $i \in N$ $\pi^e_{ii} = 1$ but for any $i \in \{0, 1, \ldots, n\}$ $\pi^e_{i0} = \pi_{i0}$.

Claim 1: π^e is also incentive compatible and beneficial relative to no-information or full-revelation benchmark.

Proof As π is incentive compatible, for any $i, i' \in \{0, 1, \ldots, n\}$

$$\begin{aligned} \mathbb{E}_{\mu^{\pi}_{a_i}} u_r(a_i, \omega) \geq \mathbb{E}_{\mu^{\pi}_{a_i}} u_r(a_{i'}, \omega) \Leftrightarrow \\ \mu_0(\omega_i)\pi_{ii}u_r(a_i, \omega_i) \geq \mu_0(\omega_j)\pi_{ij}u_r(a_j, \omega_j). \end{aligned} \tag{9}$$

For the incentive compatibility of π^e, for any $i \in \{1, \ldots, n\}$

$$\mu_0(\omega_i)u_r(a_i, \omega_i) \geq \mu_0(\omega_0)\pi_{i0}u_r(a_0, \omega_0),$$

which holds by equation 9. Therefore, the receiver has an incentive to follow any recommendation a_i, and as $\mu^{\pi^e}_{a_0}(\omega_0) = 1$, receiver also follows the recommendation a_0.

For any agent $i \in N$, π^e is beneficial relative to no-information benchmark if

$$\sum_{j\neq 0} \mu_0(\omega_j)u_i(a_j, \omega_j) + \mu_0(\omega_0) \sum_j \pi(a_j, \omega_0)u_i(a_j, \omega_0) > \sum_{j\neq 0} \mu_0(\omega_j)u_i(a_0, \omega_j) + \mu_0(\omega_0)u_i(a_0, \omega_0),$$

which holds by Assumption 4. π^e is also beneficial against full-revelation benchmark as it differs from full-revelation only at state ω_0 and at that state π^e recommends some agent-actions with positive probability. □

Claim 2: The total expected payoff of all agents under π^e is not smaller than under π.

Proof The total equally weighted expected payoffs of agents under π can be written as

$$\sum_{i\neq 0} \mu_0(\omega_i)\pi_{ii}\Gamma_{ii} + \sum_i \sum_{j\neq i,0} \mu_0(\omega_j)\pi_{ij}\Gamma_{ij} + \mu_0(\omega_0) \sum_i \pi_{i0}\Gamma_{i0},$$

while the corresponding sum under π^e is

$$\sum_{i\neq 0} \mu_0(\omega_i)\Gamma_{ii} + \mu_0(\omega_0) \sum_i \pi_{i0}\Gamma_{i0},$$

and π^e leads to a higher total sum if

$$\sum_{i\neq 0} \mu_0(\omega_i)(1 - \pi_{ii})\Gamma_{ii} \geq \sum_i \sum_{j\neq i,0} \mu_0(\omega_j)\pi_{ij}\Gamma_{ij},$$

which holds by Assumption 4 because

$$\begin{aligned} \sum_{i\neq 0} \mu_0(\omega_i)(1 - \pi_{ii})\Gamma_{ii} &\geq \min_{i\in N} \Gamma_{ii} \sum_{i\neq 0} \mu_0(\omega_i)(1 - \pi_{ii}) \geq \max_{\{i,j|i\neq j\neq 0\}} \Gamma_{ij} \sum_{i\neq 0} \mu_0(\omega_i)(1 - \pi_{ii}) \\ &= \max_{\{i,j|i\neq j\neq 0\}} \Gamma_{ij} \sum_{i\neq 0} \mu_0(\omega_i) \sum_{j\neq i} \pi_{ji} = \max_{\{i,j|i\neq j\neq 0\}} \Gamma_{ij} \sum_{j\neq 0} \sum_{i\neq j} \mu_0(\omega_j)\pi_{ij} \\ &= \max_{\{i,j|i\neq j\neq 0\}} \Gamma_{ij} \sum_i \sum_{j\neq i,0} \mu_0(\omega_j)\pi_{ij} \geq \sum_i \sum_{j\neq i,0} \mu_0(\omega_j)\pi_{ij}\Gamma_{ij}. \end{aligned}$$

This implies, π cannot Pareto dominate π^e. $\square$

Now, pick any incentive compatible information structure π that also maximizes the total expected payoffs of all agents. Then by Claim 2, there exists an endorsement rule that leads to at least the same total payoff; hence it is Pareto efficient. As by Claim 1, π^e is also beneficial and incentive compatible, the hypothesis follows. $\blacksquare$

Proof of Theorem 4 (i) A player has the same preferences if her utility function is multiplied by a positive constant. For any player k and action a_i, define a new utility function

$$\tilde{u}_k(a_i) \equiv w_k u_k(a_i)$$

For any $i, j \neq 0$

$$\sum_k \tilde{u}_k(a_i) = \sum_k \tilde{u}_k(a_i, \omega_i) = \sum_k \tilde{u}_k(a_j, \omega_i) = \sum_k \tilde{u}_k(a_j)$$

Therefore, Assumption 4 is satisfied. By Theorem 3, there exists an endorsement rule which is incentive compatible, Pareto-efficient, and beneficial relative to no disclosure.

(ii) A player has the same preferences if a constant is added to her utility function. For every player k and action a_i, define a new utility function

$$\tilde{u}_k(a_i) \equiv u_k(a_i) - u_k^{sb} - \epsilon$$

where $\epsilon > 0$. Choose ϵ sufficiently small so that for any i $\sum_k \tilde{u}_k(a_i) > 0$. Construct a matrix U where the element in ith row and kth column is $\tilde{u}_k(a_i)$. Observe that for any i, k such that $i \neq k$ $u_k(a_i) < 0$ and $u_k(a_k) > 0$. Because the non-diagonal elements are negative, U is a Z-matrix. Because no element is equal to 0, the matrix is irreducible. Let $\mathbf{1}$ define the vector where all entries are equal to 1. Because for any i $\sum_k \tilde{u}_k(a_i) > 0$, $U\mathbf{1} > 0$. By Theorem 2.7 in Chapter 6 of Berman and Plemmons [1979], U^-1 exists and has strictly positive entries. Consider the equation

$$Uw = \mathbf{1}$$

Because U^{-1} exists, we can rewrite the equation as

$$U^{-1}Uw = U^{-1}\mathbf{1} \Leftrightarrow w = U^{-1}\mathbf{1} > 0$$

Hence, w is a strictly positive vector such that for any $i, j \neq 0$

$$\sum_k \tilde{w}_k u_k(a_i) = 1 = \sum_k \tilde{w}_k u_k(a_j)$$

Therefore, the condition in (i) is satisfied and there exists an incentive compatible and Pareto-efficient endorsement rule that is beneficial relative to no disclosure. $\blacksquare$