J. Econ. Theory 229 (2025) 106068

Contents lists available at ScienceDirect

Journal of Economic Theory

journal homepage: www.elsevier.com/locate/jet

Information agreements ☆

Kemal Kıvanç Aköz [a,*], Arseniy Samsonov [b]

[a] Faculty of Economic Sciences, HSE University, Russia
[b] Department of Economics, Özyeğin University, Türkiye

ARTICLE INFO

JEL classification:
C71
D82

Keywords:
Persuasion
Bargaining solution
Efficiency

ABSTRACT

We define a (cooperative) informational bargaining problem, where several agents have to agree on the persuasion of a receiver. The bargaining set includes payoff vectors that can be generated by information structures and disagreement leads to an exogenous benchmark that may involve full or no information. We characterize the existence of an agreement that benefits all agents when preferences are state-independent. Our characterization yields conditions that depend only on the payoff structure but are independent of the prior beliefs in some cases. We analyze Pareto efficient information structures in two applications: selection environments, where the receiver picks the best agent, and the bargaining between a retailer platform and a regulator on consumer privacy regulation.

1. Introduction

Designing rules for sharing information—such as a company's disclosure policy, privacy regulations on online platforms, media censorship in authoritarian states, or political parties' election campaigns—often requires agreement among multiple agents with differing interests. In some cases, any arrangement must have the unanimous approval of all agents. As a result, the information-sharing rule that emerges reflects the outcome of negotiations rather than the optimal choice of a single representative. For instance, workers within a firm's department might collectively agree on an evaluation scheme that credibly recommends one of them for promotion. In an electoral setting, opposition parties with smaller voter bases than the incumbent party might still prevail if they coordinate their campaigns effectively. Finally, a retail platform might lobby a regulatory agency for privacy regulations that limit what the platform must share with other companies. Under what conditions can agents with different interests reach an information agreement?

In this paper, we generalize the canonical Bayesian Persuasion setup introduced by Kamenica and Gentzkow (2011) on two fronts. First, rather than having a single sender, we allow multiple agents to collectively agree on an information structure for persuading a single receiver. If all agents reach consensus, they commit to an information structure that specifies a probability distribution for the signal generated in each state of the world. We remain agnostic about the specific negotiation process that leads to this consensus and

☆ We are very grateful to the Editor, Juuso Välimäki, and two anonymous referees for their insightful and constructive feedback. We also thank Piotr Dworczak, Georgy Egorov, Gregory Sasso, Steven Kivinen, Ming Li, Aleksei Smirnov, Sergey Stepanov, Alex Suzdaltsev, Marek Weretka, and the participants of various seminars and conferences for their valuable comments. All remaining errors are our own.

* Corresponding author.
E-mail addresses: kakoz@hse.ru (K.K. Aköz), arseniy.samsonov@ozyegin.edu.tr (A. Samsonov).

https://doi.org/10.1016/j.jet.2025.106068
Received 9 December 2023; Received in revised form 21 July 2025; Accepted 15 August 2025

Available online 25 August 2025

therefore model any agreement as the outcome of a cooperative bargaining problem.[1] Second, if there is no agreement,[2] we allow the resulting default outcome to grant the receiver no, partial, or full information, depending on the context. For instance, in an electoral environment, if opposition parties fail to coordinate against the incumbent, they may run individual campaigns in which they attack not only the incumbent but also each other—thereby lowering their overall chances of winning due to the information revealed to voters. If, instead, the opposition parties successfully reach an agreement, they might launch a unified campaign that focuses solely on challenging the incumbent.

We first analyze the existence of a beneficial agreement in environments where agents care only about the receiver's action. We focus on settings in which the disagreement outcome is represented by a "discrete" information structure, meaning there is a unique optimal action at each belief attainable under disagreement. Our first main result (Theorem 1) shows that a beneficial agreement exists if and only if there is a finite lottery over posterior beliefs that yields strictly higher expected utility for each agent compared to their disagreement payoff. Notably, this lottery over posteriors need not be Bayesian-plausible—its mean may differ from the prior. To prove this result, we construct a Bayesian-plausible lottery over beliefs that implements a perturbation of the disagreement information structure with some probability and induces the beneficial lottery for the rest.

Second, we prove that if there exists a nonnegative weight for each agent such that, for every possible belief, the weighted sum of each agent's payoff gains (relative to the disagreement outcome) never exceeds zero, then a beneficial agreement is impossible. We establish this result using a theorem of the alternatives for linear inequalities: either there exists a finite lottery over beliefs that generates a strictly positive surplus for all agents—which implies the existence of a beneficial agreement by the previous characterization—or there is a vector of weights ensuring that the weighted sum of payoff gains never becomes positive.

One advantage of relaxing the requirement of Bayesian plausibility is that our conditions for the existence of a beneficial agreement depend on the prior belief only through the disagreement payoffs. As a result, in cases where the disagreement outcome provides full information, the criteria for existence depend solely on the payoff structure and not on the prior belief. We refer to this scenario as "beneficial censorship," since any agreement among the agents necessarily withholds some information from the receiver. We establish two main results in this setting. First, a beneficial agreement does not exist if and only if there is a set of nonnegative weights for the agents such that the weighted sum of their payoffs remains constant. Second, if each agent favors a different action by the receiver—gaining a strictly positive payoff only when that action is played (while others do not gain)—no beneficial agreement is possible. Intuitively, any information structure that deviates from full information reduces the probability of at least one agent's favored action, thereby harming that agent relative to the default outcome.[3]

We later provide a coalition-based sufficient condition. Theorem 2 establishes that a beneficial agreement exists if, for any coalition of a fixed size, there is a belief that makes the coalition members better off, makes non-members worse off, and increases the total payoff of all players relative to disagreement. Intuitively, an agreement among players with conflicting interests is feasible when they can generate a joint surplus compared to the disagreement, and each participant has a chance of securing a substantial gain. This result is particularly relevant in settings where each action by the receiver benefits a distinct group of agents.

Extending our characterization of beneficial agreements to situations where agents' preferences are state-dependent is generally not feasible. However, in specific contexts, strong results are still attainable. One such context involves negotiations between a retailer platform and a regulatory agency concerning privacy regulation. The platform aims to disclose as much information as possible about buyer types to sellers in order to maximize its commission, whereas the regulator wishes to prevent the erosion of consumer surplus through third-degree price discrimination. We show that Pareto-efficient information structures in this setting are characterized by ensuring that trade between a buyer and a seller occurs with full probability. Building on this finding, we identify a condition on the demand function under which no beneficial agreement can emerge if buyer types are partitioned into coarse categories.

A natural issue that arises in a cooperative bargaining problem is determining which bargaining solutions can be applied. Application of the most commonly used solutions—such as the Nash and Kalai–Smorodinsky solutions—typically require that the bargaining set is compact and convex. In Section 4.1, we show, under the tie-breaking assumption, that the receiver follows any recommended action whenever it is optimal for her, the bargaining set—defined as the set of all vectors of expected payoffs generated by possible information structures—is indeed compact and convex.

We explore bargaining solutions in "selection environments" where the receiver must either select one agent or make no selection and each agent prefers to be selected rather than not. This model applies to a variety of contexts—such as electoral competition, political censorship, and promotion schemes. We introduce a family of information structures that we call "endorsement rules" that recommend the selection of an agent if this selection is optimal for the receiver. When payoffs and probabilities are symmetric, the Nash and Kalai-Smorodinsky solutions coincide with an endorsement rule. In addition, such rules are Pareto efficient, better for all agents than no and full information, and give agents no incentive to reveal further information.

Example *(Agreement over Electoral Campaigns)*. Consider an electoral regime with three political parties: Left, Right, and Centrist. Each party stands for a distinct public policy, any of which could be best for a representative voter (the receiver). The voter earns a payoff of 1 if the correct party wins and 0 otherwise. Initially, it is most likely that the Centrist party is best, with prior beliefs

[1] Cooperative bargaining solutions are particularly useful when the bargaining protocol is unspecified or complex (Crawford, 1997), which is why they are frequently employed in experimental settings involving unstructured bargaining (Anbarci and Feltovich, 2013; Karagözoğlu, 2019).

[2] Our framework, where each agent has a veto power, diverges from a model featuring a social planner who maximizes a social welfare function by choosing an information structure. An unconstrained social welfare optimum may not coincide with an outcome acceptable to all agents.

[3] In the Online Appendix, we present an application of beneficial censorship where agents are connected via a network. We investigate the relation between network properties and the existence of beneficial censorship.

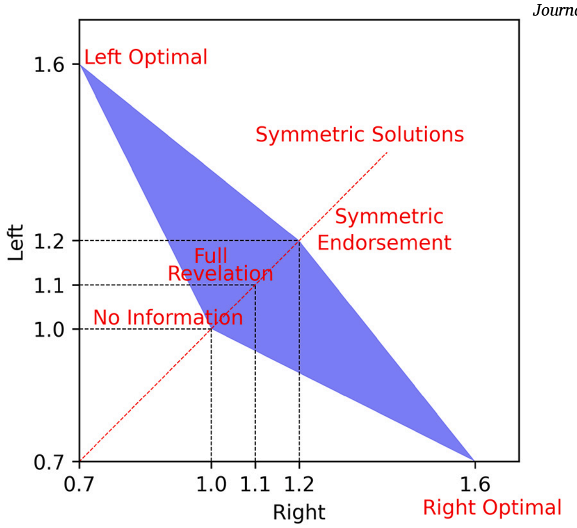

Fig. 1. The horizontal axis is the Right-wing party's expected payoff that some information structure can generate, and the vertical axis shows the same for the Left-wing party. The symmetric endorsement rule is efficient and corresponds to the symmetric Nash Bargaining solution.

$(0.1, 0.8, 0.1)$. Consequently, in the absence of further information, she would support the Centrist party. The parties themselves also do not know which policy is best but expect to learn this during the election campaign. The Left party's payoff is 3 if it wins, 1 if the Centrist party wins, and 0 if the Right party wins; the Right party's payoffs are symmetric. This setup reflects a conflict between the Left and Right parties, both of which prefer the Centrist to win over each other.

Leaders of the Left and Right parties may agree to coordinate their rallies and hold pre-electoral debates in which certain issues are deliberately excluded, aiming to prevent the likely victory of the Centrist party. If the Left and Right parties collude, they can dominate the election campaign period and prevent the Centrist party from running a campaign that effectively reaches most voters. However, if they fail to collude, a competitive disclosure game arises among all parties, resulting in full revelation as the expected equilibrium outcome. Under this disagreement scenario, each of the Left and Right parties has a prior expected payoff of $3 \times 0.1 + 1 \times 0.8 + 0 \times 0.1 = 1.1$. Is there a way for the parties to reach a mutually beneficial agreement?

The answer is yes. To see why, note first that under full revelation, the centrist party is elected with an ex-ante probability of 0.8. If it were possible to implement a different distribution over revelation outcomes—say, a uniform distribution of $(1/3, 1/3, 1/3)$—each of the Left and Right parties would have an ex-ante payoff of $3 \times 1/3 + 1 \times 1/3 + 0 \times 1/3 \approx 1.33$, which exceeds the full-revelation payoff of 1.1. Although such a uniform distribution of beliefs is not feasible given the voter's strong prior, there is still a way to construct a beneficial information structure.

Two features of this example allow us to tilt the probabilities in favor of the Left and Right parties: (1) the voter is never indifferent among multiple alternatives once she knows which party is best, and (2) both parties have state-independent preferences. Exploiting these properties, we design an information structure whose induced beliefs are close to the degenerate beliefs under full information yet shift the ex-ante probability away from the Centrist victory. Theorem 1 generalizes this insight.

In this example, the Left and Right parties can do far better than merely introducing slight probability tilts relative to full revelation. Consider the following informative campaign strategy: the two parties jointly promote a well-known political expert as an impartial referee, pledging to publicize her endorsements. The expert commits to the following endorsement rule: if the Left (Right) party is truly the best, she endorses that party; if the Centrist party is the best, she endorses the Left (Right) party with probability $\frac{1}{8}$ and the Centrist party with probability $\frac{6}{8}$. An endorsement of the Centrist party signals to voters that the Centrist is indeed the best, prompting them to support it. Meanwhile, because $\frac{0.1}{0.1+0.8(1/8)} = 0.5$ the voter will follow the endorsement to vote for the Left (or Right) party when that party is named. The expected payoff for the Left (Right) party equals $3\left(0.1 + 0.8\frac{1}{8}\right) + 1\left(0.8\frac{6}{8}\right) + 0\left(0.1 + 0.8\frac{1}{8}\right) = 1.2$, which is greater than the disagreement payoff of 1.1. Notice that this information structure is both Pareto efficient and robust to any ex-post deviations by the parties that involve revealing additional information after an endorsement. Theorem 3 establishes that such endorsement rules are the only information structures that satisfy Pareto efficiency while remaining robust to ex-post deviations across a broad class of environments in which the receiver selects one of the agents to allocate an indivisible good.

The example also illustrates our results relating to bargaining solutions. Fig. 1 displays the bargaining set for this example, i.e. a set of agreements that can be reached between the Left and Right parties, which is compact and convex.[4] The symmetric endorsement rule above corresponds to a Pareto-efficient and symmetric payoff profile, indicating that it can be supported by the Nash and Kalai–Smorodinsky bargaining solutions (Proposition 6). ⋄

[4] In the Online Appendix, we provide an algorithm for computing the set of payoff vectors that can be generated by different information structures.

1.1. Related literature

A recent influential work by Doval and Smolin (2024) examines a closely related question concerning the relationship between welfare and information. In their model, each state of the world corresponds to a type within a population, and the prior belief represents the population distribution of types. Each signal generated by an information structure induces a "segmentation" of the population, with heterogeneous preferences over possible segmentations[5]; while we consider multiple agents who share a common uncertainty over an arbitrary state space in our framework. Doval and Smolin (2024) provide a geometric characterization of the "Bayes Welfare Set," i.e., the set of payoff profiles that can be generated by an information structure, along with its Pareto frontier. Their results offer a useful framework for analyzing bargaining among heterogeneous interest groups against a no-information benchmark. In this paper, we take a related approach that enables us to focus more explicitly on the bargaining dimension, and instead of restricting attention to a no-information benchmark, we examine the strategic feasibility of agreements over information structures against an arbitrary disagreement benchmark. In Sections 3.3 and 4, we explore the connection between these two approaches in greater detail through an application to price discrimination and privacy regulation.

In our framework, multiple agents collectively agree on an information structure. We assume that agents can reach a binding agreement and commit to the chosen structure, which sets our model apart from the canonical competitive persuasion models by Gentzkow and Kamenica (2017a) and Gentzkow and Kamenica (2017b). In these models, agents cannot form binding agreements or credibly coordinate on a single information structure. Instead, each sender independently commits to an information structure, and the receiver observes multiple signals generated by these competing sources. In contrast, our model assumes that a single information structure emerges from the agents' collusion. Gentzkow and Kamenica (2017a) show that an increase in competition leads to more informative outcomes. However, subsequent work by Li and Norman (2018) and Li and Norman (2021) suggests that this result depends on several factors, particularly the timing of persuasion (See also Bhattacharya and Mukherjee, 2013; Board and Lu, 2018; Hoffmann et al., 2020; Au and Kawai, 2020, for some other related models). In our model, the agents are able to limit the competitive forces that push for more information through a binding agreement. As discussed in Section 3, they may establish a beneficial censorship scheme to strategically control the information available to the receiver. If no agreement is reached, the game enters a disagreement stage, which could resemble a competitive persuasion setting among the agents. In this sense, our model can be viewed as the first stage of a collusion game, where failure to agree results in 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, we define a non-cooperative problem and discuss the endorsement rules as a bargaining solution. Section 5 concludes.

2. Model

Let $N = \{1, \dots, n\}$ be a finite set of agents, and let there be a single receiver, denoted by r. The state space is finite and denoted by Ω, with $\Delta(\Omega)$ representing the set of possible beliefs over Ω. We denote a generic belief by $\mu \in \Delta(\Omega)$ and let $\mu_0 \in \Delta(\Omega)$ be the common prior among the agents and the receiver. With a slight abuse of notation, we also use ω to denote the degenerate belief that assigns probability 1 to the state ω. A belief $\mu \in \Delta(\Omega)$ is called **interior** if $\mu(\omega) > 0$ for all $\omega \in \Omega$. We assume that the prior belief μ_0 is interior.

We assume that the set of pure actions available to the receiver is finite and denote it by A. Let $u_r : A \times \Omega \to \mathbb{R}$ be the receiver's payoff function and for each agent $i \in N$, let $u_i : A \times \Omega \to \mathbb{R}$ be the payoff function of agent i. For any belief $\mu \in \Delta(\Omega)$ we denote by $a(\mu)$ any action that is chosen by the receiver at belief μ, while $\bar{a}(\mu)$ denotes an optimal action; that is, $\bar{a}(\mu) \in \arg\max_{a' \in A} \sum_{\omega \in \Omega} \mu(\omega) u_r(a, \omega)$. A collection of actions $\{a(\mu)\}_{\mu \in \Delta(\Omega)}$ represents any pure strategy of the receiver, while $\{\bar{a}(\mu)\}_{\mu \in \Delta(\Omega)}$ denotes an optimal pure strategy of the receiver.[6]

An information structure specifies a probability distribution over a signal space for each state $\omega \in \Omega$. After observing the realized signal, the receiver updates her belief according to the Bayesian rule and chooses an optimal action. By Revelation Principle (Myerson, 1979, 2008) it is without loss of generality to assume that the signals are recommendations of actions. As the action space is finite, any such information structure can be represented as a $|\Omega| \times |A|$ probability matrix, where each row sums to 1. We denote an information structure by π, which induces the conditional distributions over recommendations $\{\pi(\cdot|\omega)\}_{\omega \in \Omega} \subset \Delta(A)$.

An information structure is called incentive compatible (IC) if the receiver has no incentive to deviate from the recommendation it generates.[7] Fix any information structure π. For any action $a \in A$, let $\mu(\cdot|a, \pi)$ denote the posterior belief after observing the recommendation a. This posterior is given by Bayes' rule:

$$\mu(\omega|a, \pi) = \frac{\pi(a|\omega)\mu_0(\omega)}{\sum_{\omega' \in \Omega} \pi(a|\omega')\mu_0(\omega')} \quad \forall \omega \in \Omega$$

whenever $\pi(a|\omega) > 0$ for some $\omega \in \Omega$. An information structure is incentive-compatible if for any $a, a' \in A$, whenever $\pi(a|\omega) > 0$ for some $\omega \in \Omega$, the receiver weakly prefers following the recommendation:

[5] See Bergemann et al. (2015) for a related idea in the context of market segmentation of consumer types.

[6] For simplicity, we focus on pure strategies for the receiver in this paper.

[7] See Dughmi and Xu (2016) for the same formulation of a persuasion problem.

$$\sum_{\omega} \mu(\omega|a,\pi)u_r(a,\omega) \geq \sum_{\omega} \mu(\omega|a,\pi)u_r(a',\omega).$$

The ex-ante payoff for each agent under an IC information structure can be directly expressed in terms of the prior and the chosen information structure.[8] For any agent $i \in N$ the ex-ante payoff under an IC information structure π is given by:

$$\mathbb{E}(u_i|\pi) \equiv \sum_{\omega\in\Omega} \mu_0(\omega) \sum_{a\in A} \pi(a|\omega)u_i(a,\omega). \tag{1}$$

Let Π denote the set of IC information structures.

Throughout the paper, we focus on the existence of an information structure that improves the payoffs of all agents relative to the default exogenous information structure π_0. If all agents agree on such an information structure, it is implemented, and otherwise, π_0 is implemented. One interpretation is that there is a social planner who proposes an information structure $\pi \in \Pi$ to all agents. If no agent vetoes π, the social planner implements it. Otherwise, the game enters the disagreement stage, and the social planner implements π_0. Alternatively, one can think of the social planner as being able to enforce any information structure, but aiming to improve the welfare of all agents. A final interpretation is that the agents may collectively establish an institution responsible for implementing the agreed-upon information structure. In both the agreement and disagreement stages, the receiver observes the resulting information structure and the recommendation it generates, then updates her belief and selects an action accordingly.

If agents fail to agree on an information structure, the game enters a disagreement stage, and an exogenous default information structure, denoted by π_0, is implemented. The outcome of this stage determines each agent's disagreement payoff, denoted by d_i for each $i \in N$. For instance, if disagreement prevents agents from credibly generating any information, the disagreement stage corresponds to a no-information scenario. In this case, the receiver acts based on the prior belief μ_0, and the disagreement payoff for agent i is given by $d_i = \sum_{\omega\in\Omega} u_i(\bar{a}(\mu_0),\omega)$ for some $\bar{a}(\mu_0)$. Alternatively, disagreement may lead to a competitive disclosure game, in which agents reveal information strategically, resulting in full revelation. In this case, the disagreement payoff corresponds to the ex-ante payoff under full revelation: $d_i = \sum_{\omega\in\Omega} \mu_0(\omega)u_i(a(\omega),\omega)$.

3. Beneficial agreement

We begin our analysis by examining the existence of a beneficial agreement. We define a beneficial agreement as an information structure that strictly improves the payoffs of all agents relative to their disagreement payoffs. Our focus on strict dominance is motivated by strategic considerations. If an information structure benefits only a subset of agents while leaving others indifferent between agreement and disagreement, the indifferent agents might veto the agreement with positive probability in certain contexts. Strict dominance ensures robustness to such concerns, guaranteeing that all agents have a strict incentive to accept the proposed information structure.

Our main result on the existence of a beneficial agreement holds when agents' preferences depend only on actions and not on states. We formalize this condition as follows:

Assumption 1 *(State-Independence)*. For any $a \in A$, states $\omega, \omega' \in \Omega$, and agent $i \in N$ $u_i(a,\omega) = u_i(a,\omega')$.

When agents have state-independent preferences, we can simplify the notation by writing $u_i(a)$ to represent the payoff of agent $i \in N$ at any action $a \in A$. We adopt this assumption and notation throughout the analysis until Subsection 3.3, where we consider the case of state-dependent preferences.

3.1. Two characterizations under state-independence

The classical result by Kamenica and Gentzkow (2011) establishes that any information structure can be represented by a Bayesian-plausible lottery over posterior beliefs. A distribution over posteriors is Bayesian-plausible if its expectation equals the prior, ensuring consistency with Bayes' rule. Consequently, if there exists a Bayesian-plausible distribution of posterior beliefs that yields a strictly higher payoff for all agents compared to the disagreement information structure, a beneficial agreement must exist. However, under certain conditions on the disagreement information structure, the requirement of Bayesian plausibility can be relaxed. Kamenica and Gentzkow (2011) show that such a relaxation is possible when the receiver's preference is *discrete* at the prior, meaning there is a unique optimal action under the prior belief. We generalize this notion of discreteness as follows.

Definition 1. Let $\pi \in \Pi$ be any incentive-compatible information structure. We say that π is **discrete** if for any action $a \in A$ such that $\pi(a|\omega) > 0$ for some $\omega \in \Omega$, the receiver has a unique optimal action at the posterior belief $\mu(\cdot|a,\pi)$.

Discrete information structures induce posterior beliefs where the receiver's preferences are discrete, meaning she has a unique optimal action at each belief. The set of discrete information structures forms a dense subset of the set of all incentive-compatible

[8] Note that whether the receiver follows the recommendations of an IC information structure is about how the receiver breaks ties when she is indifferent. We explore this issue further in Section 4 below.

information structures.[9] However, some important information structures may not be discrete. For instance, suppose π_0 is an optimal information structure for an agent that ensures the receiver selects the best-attainable action for that agent. Then, by Proposition 5 of Kamenica and Gentzkow (2011), the receiver may be indifferent between multiple actions at the belief that induces this best action. Thus, if we assume a discrete disagreement information structure, we rule out the possibility that a single agent dominates the disagreement stage by imposing her preferred information structure.

Discrete information structures can naturally arise as disagreement outcomes in various contexts. For instance, in almost all environments, both full information and no information structures are discrete. This implies that if the disagreement stage corresponds to a setting where agents cannot prevent the flow of information or, conversely, where agents must collude to generate any credible information, the disagreement information structure is likely to be discrete. Additionally, partially informative information structures that partition the state space and send a single message for all elements within each partition are often discrete. Such information structures frequently emerge in practice. For example, credit rating agencies summarize complex economic information using a small number of letter-grade combinations. Another example is state-controlled media in autocratic regimes, which may truthfully report certain news while omitting others. In the Bayesian persuasion framework, this behavior can be captured by pooling "censored" states into a single message while sending a truthful message in all other cases.

Our main existence result below provides two characterizations for the existence of a beneficial agreement under state-independent agent preferences and a discrete disagreement information structure.

Theorem 1. *Suppose that Assumption 1 of state-independence holds and the disagreement information structure π_0 is discrete. Fix any optimal strategy $(\bar{a}(\mu))_{\mu\in\Delta(\Omega)}$ of the receiver. The following are equivalent.*

(i) *There is a beneficial agreement.*

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

$$\sum_{l} \lambda_l u_i(\bar{a}(\mu_l)) > d_i. \tag{2}$$

(iii) *For any non-negative numbers $y_1, ..., y_n$, not all 0, there exists a belief $\mu \in \Delta(\Omega)$ such that*

$$\sum_{i\in N} y_i \left(u_i(\bar{a}(\mu)) - d_i\right) > 0. \tag{3}$$

Proofs of all results are in Appendix A.

Condition (ii) in Theorem 1 establishes that a beneficial agreement exists if and only if there is a finite distribution over posterior beliefs that yields a strictly higher expected payoff for all agents. Once such a distribution is identified, we do not need to verify the Bayesian plausibility condition. For example, suppose there exists a single action by the receiver that generates a positive surplus for all agents and is rationalizable at some posterior belief. Then, it is always possible to construct an information structure that increases the ex-ante probability of selecting this action. However, condition (ii) is more general and applies even when divergent interests prevent consensus on any single action. In many applications—including the electoral campaign example discussed in the Introduction—any action that benefits some agents may simultaneously harm others relative to the disagreement outcome. Nevertheless, if there exists a probabilistic mix over actions where the gains from some actions offset the losses from others, then a beneficial agreement can still be achieved.

The proof of the equivalence between the existence of a beneficial agreement and condition (ii) follows similar steps as in Proposition 2 of Kamenica and Gentzkow (2011), constructing a Bayesian-plausible lottery over beliefs that yields higher expected payoffs for all agents. Fix any lottery $(\lambda_l, \mu_l)_{l=1}^{L}$ that promises a higher expected payoff to all agents as described in the statement. The construction proceeds as follows when the disagreement information structure is interior; that is, all beliefs it induces are of full support. For each belief μ that the information structure generates, there is a nearby belief μ' such that the receiver's optimal action remains unchanged, i.e., $\bar{a}(\mu) = \bar{a}(\mu')$. Since the agents' payoffs depend solely on the actions, their expected payoffs at μ' and μ are identical. Moreover, because μ is interior, we can express μ as a convex combination of any belief in the lottery μ_l and some μ'_l that is sufficiently close to μ. We then construct a Bayesian plausible lottery over beliefs that generates either the binary lotteries (μ'_l, μ_l) or the distribution of beliefs induced by the disagreement information structure. As the lottery $(\lambda_l, \mu_l)_{l=1}^{L}$ promises a higher expected payoff, this new Bayesian plausible lottery also promises a higher expected payoff. Finally, if the disagreement information structure is not interior, we introduce a small perturbation to ensure full support, without affecting the expected payoffs for senders, as discreteness ensures that such a perturbation does not alter optimal actions.

Condition (iii) in Theorem 1 implies that a beneficial agreement does not exist if and only if there exists a vector of weights over the agents such that the weighted sum of surpluses (relative to the disagreement payoff) is always non-positive. For example, in a zero-sum persuasion game between two agents, an agreement is impossible because one agent's gain necessarily implies the other agent's loss. In this case, assigning equal and positive weights to the two agents (and weight 0 to all others) would satisfy the inequality in condition (iii), confirming that no beneficial agreement can be reached. More generally, whenever distributional conflict among

[9] That is, for any incentive-compatible information structure, there exists a sequence of discrete information structures that approximates it arbitrarily well.

agents is severe enough that gains for some agents translate into proportional losses for others, a beneficial agreement cannot exist. In the proof, we show that either statement (ii) holds (implying a beneficial agreement exists) or the negation of statement (iii) holds (ensuring that no such vector of weights exists). For any finite collection of beliefs, we construct a matrix of expected payoffs. Then, by Ville's Alternative (Border, 2013)—a result equivalent to Farkas' Lemma—we establish that either there exists a distribution over beliefs satisfying condition (ii), ensuring a beneficial agreement, or there exists a set of weights over agents such that inequality (3) holds, confirming that no beneficial agreement can be reached.

The conditions in Theorem 1 reveal a form of partial independence from the prior belief, as μ_0 affects the outcome only through the disagreement payoffs. In some environments, this property allows us to derive existence conditions for a beneficial agreement that depend solely on the payoff structure rather than the prior belief. For example, consider the following corollary, which follows directly from condition (3).

Corollary 1. *If there exist non-negative numbers $y_1, ..., y_n$, not all 0, such that for any two actions $a, a' \in A$ $\sum_{i\in N} y_i u_i(a) = \sum_{i\in N} y_i u_i(a')$, there is no beneficial agreement.*

Corollary 1 states that if the weighted sum of payoffs, for some weights, is constant for all strategies of the receiver, then there is no beneficial agreement. This result indicates that the creation of some surplus with information is crucial for a beneficial agreement. It follows from the simple observation that in such a case all strategies are Pareto efficient among the agents; so there is no way to benefit one agent by reallocation of probabilities of actions without hurting another.

To further illustrate the implications of prior-belief independence in Theorem 1, we focus on the case where the disagreement information structure corresponds to full information.

Beneficial Censorship

Under which conditions can agents collude on an information structure when the default one provides full information? In such settings, any beneficial agreement can be interpreted as concealing some information from the receiver. One relevant scenario is media control in partially authoritarian regimes where successful censorship requires the cooperation of divided elites. When the disagreement information structure is full information and each strategy corresponds to an optimal response of the receiver to full information at some state, the sufficient condition in Corollary 1 for the non-existence of beneficial agreement also becomes necessary.

Proposition 1. *Suppose that $A = \{\bar{a}(\omega)\}_{\omega\in\Omega}$ and the disagreement information structure is full information. There exists a beneficial agreement if and only if there do not exist non-negative numbers $\{y_i\}_{i=1\in N}$, not all 0, such that the weighted sum of payoffs remains constant across all actions: $\sum_{i\in N} y_i u_i(a) = \sum_{i\in N} y_i u_i(a')$ for any $a, a' \in A$.*

Proposition 1 can be equivalently formulated as a linear algebra problem. Let U be the payoff matrix, where rows correspond to the strategies of the receiver, and columns correspond to the agents. Formally, $U = [u_{ij}]_{j,i\in N}$, where $u_{ij} = u_i(a_j)$. Then, Proposition 1 states that the existence of a beneficial agreement is equivalent to the non-existence of a non-negative and non-zero solution to $Uy = 1$ when payoffs are non-negative. The following Corollary follows immediately.

Corollary 2. *Suppose that $A = \{\bar{a}(\omega)\}_{\omega\in\Omega}$ with $|A| = k$ and the disagreement information structure is full information. Let n be the number of agents such that $k - 1 \geq n$. Then, for a generic payoff structure, a beneficial agreement exists.*

Proposition 1 implies that if no beneficial agreement exists, then the system of $k - 1$ linear equations with n unknowns must have a solution. If $k - 1 > n$, such a solution generically does not exist. If $k = n - 1$, then the conditions of Proposition 1 are $k - 1 = n$ equations of the form $\sum_{i\in N} y_i u_i(a_j) = \sum_{i\in N} y_i u_i(a_{j+1}) \Leftrightarrow \sum_{i\in N} y_i(u_i(a_j) - u_i(a_{j+1}))$ for $j \in [1, k - 1]$. The solution $y_1 = \ldots = y_n = 0$ exists and is generically unique because the number of equations and variables equals to $n = k - 1$, so the conditions of Proposition 1 are not satisfied.

Now, consider an environment where, for each possible strategy chosen by the receiver, one agent is designated as the 'winner,' receiving a strictly higher payoff than the others.

Assumption 2 *(Winners and Losers).* There is a one-to-one correspondence between agents, receiver strategies, and states of the world, such that: the state space is given by $\Omega = \{\omega_1, \ldots, \omega_n\}$, the receiver's action set is $A = \{a_1, \ldots, a_n\}$ where each action corresponds to the optimal response to a specific state: $a_i = \bar{a}(\omega_i)$ for any $i \in N$; and the winner's payoff satisfies: $\bar{u}_i \equiv u_i(a_i) > u_j(a_i)$ for any $i, j \in N$ with $i \neq j$.

To begin with, consider a *winner-takes-all* environment, where all the benefits generated by the receiver's action accrue to a single agent. In this case, we can normalize payoffs by setting $u_j(a_i) = 0$ for $i \neq j$. This simplification results in a diagonal payoff matrix U with strictly positive diagonal entries. Since U is invertible, we can define the weight vector $y = \left[\frac{1}{\bar{u}_i}\right]$. This satisfies the matrix equation $Uy = \mathbf{1}$, where $\mathbf{1}$ is the vector of ones. As a result, the weighted sum of payoffs is constant across all actions. By Proposition 1, this implies that no beneficial agreement exists relative to the full information benchmark. Intuitively, in a winner-takes-all setting with full information as the disagreement outcome, any transfer of probability from one strategy to another necessarily harms at least one agent. Since each action exclusively benefits a single agent, increasing the probability of an action a_i, which benefits agent i, must come at the expense of reducing the probability of another action, directly harming a different agent.

This observation can be generalized beyond the winner-takes-all case. We say that the "winner takes the most" if, for each agent $i \in N$

$$0 \leq \sum_{j \neq i} (u_i^{sb} - u_i(a_j)) < \bar{u}_i - u_i^{sb} \tag{4}$$

where $u_i^{sb} = \max_{j \neq i}\{u_i(a_j)\}$ is the second-highest payoff to agent i from any action. This condition states that the difference in payoffs between winning and the second-best outcome is greater than the sum of all other payoff differences. In other words, not being the winner is the most significant factor in determining payoff differences.

Proposition 2. *Suppose that Assumption 2 of winners and losers holds; and the disagreement information structure π_0 is full information and discrete. If the winner takes the most (i.e., condition (4) holds), no beneficial agreement exists.*

When condition (4) holds, each agent experiences a significant loss when their preferred action is not implemented. As a result, any information structure that shifts probability away from this action—by deviating from full information—imposes a cost that outweighs any potential gains from collusion. This result implies that in environments where agents have strongly conflicting interests, collusion to restrict information becomes increasingly unstable, as the incentives to preserve the probability of one's preferred action dominate any potential cooperative gains.

To summarize the discussion following Theorem 1, relaxing the Bayesian plausibility condition makes the existence of a beneficial agreement independent of the prior belief, except through its effect on the disagreement payoff. This independence property leads to an even stronger result when the disagreement information structure provides full information: in such cases, the existence of a beneficial agreement depends only on the payoff structure and not on the prior beliefs.

3.2. Coalitions

By Theorem 1, we know that whenever there exists a receiver action that is optimal at some belief and generates positive surplus for all agents, a beneficial agreement is possible. However, in many practical applications, agents have conflicting interests, meaning that any action that benefits some agents may simultaneously harm others. If side payments were possible, reaching an agreement would simply require finding a lottery over actions that generates a strictly positive total surplus. However, in our setting, where payoff transfers are limited to probability shifts over actions, this redistribution may not be sufficient to achieve a beneficial agreement. To illustrate this, consider the following example. Suppose there are three agents, and the disagreement payoff is 0 for all of them. Assume that the first and third agents have identical positive payoffs at every possible receiver action, while the second agent's payoff is exactly the negative of the first agent's payoff. In this case, every strategy generates positive total surplus, but a beneficial agreement is impossible because no probability adjustment can simultaneously improve the payoffs of both the first and second agents.

Theorem 2 below provides a general sufficient condition for environments where the strategies generate positive surplus for some of the agents, and negative for others.

Theorem 2. *If there exists $k \in \{1, \ldots, n\}$ such that for each $\phi \in \Phi(k)$, where $\Phi(k)$ denotes the set of subsets of $\{1, \ldots, n\}$ of size k, there is $\mu^\phi \in \Delta(\Omega)$ that satisfies*

$$\forall i \in \phi \quad u_i(\bar{a}(\mu^\phi)) > d_i, \quad \forall i' \notin \phi \quad u_{i'}(\bar{a}(\mu^\phi)) < d_{i'},$$

$$\forall \phi \in \Phi(k) \quad \sum_{i'=1}^{n} \left(u_{i'}(\bar{a}(\mu^\phi)) - d_{i'} \right) > 0,$$

there is a beneficial agreement.

Theorem 2 states that if there exists a fixed coalition size such that, for every coalition of that size, there is a receiver action that satisfies the following conditions: *(i)* members of the coalition receive positive surplus, *(ii)* non-members receive negative surplus, and *(iii)* the total surplus is positive, then a beneficial agreement exists. To illustrate this result, consider a setting where each agent has a distinct receiver action that strictly benefits her while harming others. If, at each such action, the sum of surpluses is positive, then an information structure can induce a probability mix over these actions that benefits all agents. Intuitively, by shifting the probabilities of recommendations across these actions, we can effectively redistribute the surplus among the agents, ensuring that all agents receive a net benefit. However, this redistribution is only possible if every agent has at least one action that generates a positive surplus. If there exists an agent who receives a negative surplus at every possible action, then any probability distribution over actions still results in a net negative surplus for that agent, making a beneficial agreement impossible. This argument generalizes to coalitions of a fixed size, where surplus must be transferable across all coalition members to ensure agreement feasibility.

The conditions in Theorem 2, while general, are not necessary for the existence of a beneficial agreement. There may be cases where no coalition of a fixed size satisfies the conditions outlined in the theorem, yet a beneficial agreement is still possible. For example, consider a setting where one strategy generates positive surplus for only a single agent, while another strategy generates positive surplus for all remaining agents. In this case, Theorem 2 would not apply, since no single coalition of fixed size benefits under a common belief. However, a beneficial agreement could still exist by appropriately structuring the probability distribution over actions to ensure that all agents receive a net positive payoff.

Using Theorem 2, we can analyze the existence of a beneficial agreement in settings where the receiver's actions benefit multiple agents simultaneously, rather than just a single agent at the expense of others. Such environments arise when competition among agents has interdependent effects—that is, an action favoring one agent may also generate positive spillovers for others. Examples of such settings include banks connected via interbank lending, where financial stability interventions may support multiple institutions, or firms collaborating on R&D projects, where innovation subsidies can benefit interconnected firms. The example below illustrates this idea in a voting environment.

Example *(Collusion among three parties)*. The following example, similar to the one in the Introduction, illustrates how Theorem 2 can be applied to analyze joint campaigns when multiple political parties are involved, some of which have conflicting interests. Suppose that three opposition parties—the Right party, the Religious party, and the Left party—enter into negotiations to form a pre-electoral coalition. They aim to defeat the incumbent, the Centrist party, which is not one of the persuading agents. Voters prefer to support the most competent party but lack information about which party is actually competent. The prior probabilities of competence are as follows: The Right party is competent with probability 0.3, the Religious and Left are each competent with probability 0.15, and the Centrist is competent with probability 0.4.

	The party in power				
Payoffs	Right	Left	Religious	Centrist	Benchmark
Right	1	0	0.5	0	0.375
Left	0	1	−5	0	−0.6
Religious	−5	0	1	−8	−4.55
Total	−4	1	−3.5	−8	−4.775

Suppose that the benchmark information structure is full revelation. Then the expected disagreement payoffs are $0.3 \times 1 + 0.15 \times 0 + 0.15 \times 0.5 + 0.4 \times 0 = 0.375$ for the Right, $0.3 \times 0 + 0.15 \times 1 + 0.15 \times (-5) + 0.4 \times 0 = -0.6$ for the Left, and $0.3 \times (-5) + 0.15 \times 0 + 0.15 \times 1 + 0.4 \times (-8) = -4.55$ for the Religious party; and the total is -4.775. Consider all pairs of agents. It is clear from the table that the pair (Right, Left) gains relative to the benchmark if the Right party is in power, while the Religious party loses relative to the benchmark. The pair (Left, Religious) gains if the Left party is in power while the Right party loses. Finally, the pair (Religious, Right) gains if the Religious party is in power while the Left party loses. In all cases, the sum of payoffs exceeds the sum of payoffs under the benchmark. For convenience, the above payoffs are in blue if they are greater than the benchmark and in red otherwise.[10] By Theorem 2, there exists a beneficial information structure relative to full revelation.

This concludes our analysis of the conditions for the existence of a beneficial agreement under state-independent preferences. We consider the case with state-dependent payoffs below with a focus on an application in price discrimination and privacy regulation.

3.3. State-dependence and privacy regulation

When agents' payoffs are state-dependent, the existence of a lottery over beliefs that yields a higher expected payoff is no longer sufficient to guarantee a beneficial information agreement in general. In particular, if the disagreement information structure is not interior (some of the beliefs it generates assign zero probability to certain states) then Theorem 1 may fail to hold. To illustrate this limitation, consider the following example.

Example *(Guessing Agents)*. Suppose there are two agents ($n = 2$), and two possible states of the world, $\Omega = \{\omega_1, \omega_2\}$). The receiver can take one of two actions, $A = \{a_1, a_2\}$. 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. Both agents and the receiver want to guess the correct state but the correct guess pays off more to the agents when the state is ω_1. Let the common prior be $\mu_0 = (0.5, 0.5)$, and suppose that the disagreement information structure is full revelation. Since the optimal information structure in this environment is also full revelation, no beneficial information structure exists. However, at the degenerate belief ω_1, the payoff that the receiver's optimal action generates for the agents is 2, while the ex-ante payoff at the full revelation is 1.5.

When the disagreement information structure is interior, it is possible to use an argument similar to the proof of Theorem 1 and find a sufficient condition.[11] An alternative way to check if there is a beneficial agreement is to explicitly verify whether the disagreement information structure is strictly Pareto dominated by some other incentive compatible information structure. This approach turns out to be useful for the monopolistic market setting we discuss below.

[10] For interpretation of the colors in the table, the reader is referred to the web version of this article.

[11] In the Online Appendix, we provide a sufficient condition for the existence of a beneficial agreement, which is a generalization of the condition in Theorem 1 to state dependent payoffs.

Example *(Price Discrimination and Privacy Regulation)*. Consider a market with a single seller offering a single unit of a good to a buyer.[12] The buyer's payoff function is given by $\gamma_v + v - p$, where $v > 0$ is the buyer's willingness to pay for the good, p is the price of the good, and γ_v is the additional surplus that the buyer gets from purchasing the good that is not reflected in her willingness to pay. Let $\gamma_v > 0$ for any valuation v so the consumer strictly prefers purchasing a product even if the price equals the willingness to pay.[13] We assume that there is a finite set of buyer types; $v \in \Omega \equiv \{v_1, \ldots, v_n\}$ with the ordering $0 < v_1 < \ldots < v_n$.

Normalize the cost of supply to 0 so the seller's profit is p when there is trade at a price p, and 0 otherwise. The seller does not know the type of the buyer. We consider a static trade scenario, meaning that the seller chooses a price once without the possibility of future adjustments. In this setting, it is never optimal for the seller to post a price that does not match one of the buyer's possible valuations. Therefore, we can restrict attention to a finite set of pricing actions, where the seller's action space coincides with the set of buyer valuations.

The seller relies on the information provided by a platform that has access to the consumer data. How much information the platform is willing to share with the seller depends on its incentives and the regulatory framework.[14] For illustration, we assume that the platform's incentives align with those of the seller and denote its payoff with u_{pl}. This assumption holds if the platform earns a proportional commission on sales, meaning it benefits directly from higher revenues. In some cases, platforms go beyond merely providing recommendations and instead directly set prices by utilizing buyer data to optimize revenue.

There is a regulator, who also does not know the buyer's type and can limit the information shared with the seller. Suppose that the regulator prioritizes consumer welfare. More precisely, the regulator receives the payoff $u_{re}(v, p) = \gamma_v + v - p$ at the state v and the price p, when there is trade and 0 otherwise. Under this setup, the set of feasible payoff pairs that can be generated through information structures corresponds to the set of possible consumer and producer surplus allocations, as characterized by Bergemann et al. (2015).

Now consider the following bargaining scenario. The platform and the regulator enter into a non-structured bargaining over the privacy regulation. During the negotiation phase, neither the platform nor the regulator knows the type of the buyer but they share a common prior. If no agreement is reached, the platform provides only partial or no information to the seller, in compliance with the status quo policy. If an agreement is reached, the platform commits to sharing information in accordance with the negotiated regulation. The regulator may be willing to accept a less stringent privacy regulation if the outcome is welfare-improving for consumers.[15]

At any IC information structure π the expected payoffs are given by:

$$\mathbb{E}(u_{pl}|\pi) = \sum_{v\in\Omega} \mu_o(v) \sum_{p\le v} \pi(p|v)p$$
$$\mathbb{E}(u_{re}|\pi) = \sum_{v\in\Omega} \mu_o(v) \sum_{p\le v} \pi(p|v)(\gamma_v + v - p). \tag{5}$$

Here, $\mathbb{E}(u_{pl}|\pi)$ represents the expected revenue of the platform (or seller) under information structure π, while $\mathbb{E}(u_{re}|\pi)$ represents the expected consumer welfare as prioritized by the regulator. Our first result characterizes Pareto efficient information structures.

Proposition 3. *An information structure is Pareto efficient if and only if trade occurs with full probability at all of its recommendations.*

It is straightforward to observe that any information structure that induces trade with full probability is Pareto efficient, because the only way to improve the payoffs of both the platform and the regulator is to increase the probability of trade. Among such efficient information structures, the only difference lies in how the surplus is distributed between the two agents via price recommendations. However, the necessity of full-probability trade for efficiency is less immediate. Specifically, consider an incentive-compatible (IC) information structure that, for some buyer valuation, recommends a price higher than the buyer's valuation, preventing trade from occurring. We show that there always exists an alternative information structure that strictly Pareto dominates the original one by shifting the probability mass from high-price recommendations (which prevent trade) to recommendations where price equals valuation. This alternative structure remains IC while increasing the probability of trade, thereby improving the welfare of both the platform and the regulator.

A clear example of a Pareto efficient information structure is full revelation.[16] A more (Blackwell) informative information structure among two efficient ones yields a higher payoff to the platform and a lower one for the regulator. Full revelation promises the highest payoff to the platform and it is the worst information structure for the regulator among the efficient ones.

---

[12] An equivalent interpretation is that there is a continuum of consumers of unit measure, where the seller's further information indicates a market segment, where the distribution of types differs from the full market, as in Bergemann et al. (2015).

[13] This gap between the buyer's surplus and willingness to pay could stem from financial constraints. A buyer of type v can afford and is willing to pay up to v; however, if the price exceeds v, the buyer may lack access to financial tools that would enable them to purchase the good.

[14] Airbnb provides price recommendations to hosts based on the platform data. https://www.airbnb.com/help/article/1168.

[15] Chang et al. (2022) document that Uber uses third-degree price discrimination based on information about buyers. Furthermore, they find that this practice is welfare improving for the consumers by increasing the likelihood of transactions.

[16] Note that the information design problem we consider in this case is different from a mechanism design problem as considered by Myerson and Satterthwaite (1983), where they prove that there is no efficient mechanism when there is two-sided uncertainty. In a mechanism design problem, the designer relies on screening the buyer and seller types, while in an information design problem as we consider in this paper, the platform can generate any information about the buyer valuations without relying on direct screening.

Now, we discuss the existence of beneficial agreements when the status quo regulation allows the platform to provide "coarse" information about the buyer. More specifically, we consider a setting where buyer valuations are partitioned into groups, with each partition element pooling together similar buyer valuations. These partition elements can be interpreted as "categories" of buyers, where the platform discloses only categorical information rather than precise valuations.[17]

Formally, let M be a non-trivial partition of Ω and suppose that for any i, l, j such that $i < l < j$ and $m \in M$ $v_i, v_j \in m \Rightarrow v_l \in m$. This condition ensures that the partition elements are intervals, meaning that no valuation is skipped within a category. Define a *buyer category* information structure as one in which when the buyer is of type v_i, the seller learns $m \in M$ such that $v_i \in M$. We can establish a simple condition on when there exists a buyer category information structure that is incentive-compatible and relative to which there is no beneficial agreement.

Proposition 4. *(i) If the distribution of buyer types satisfies the condition that $v \sum_{v' \geq v} \mu_0(v')$ is weakly decreasing in v, any category information structure is IC and has no beneficial agreement relative to it. (ii) If there is $i, j \in \{1, \dots, n\}$ with $i < j$ such that $v_i \sum_{v' \geq v_i} \mu_0(v') < v_j \sum_{v' \geq v_j} \mu_0(v')$, there is an IC categorical information structure, relative to which there exists a beneficial agreement.*

Categorical information structures are Pareto efficient when, within each category, the seller sets the lowest price among the buyer types in that category. As a result, all buyer types within each category purchase the product, inducing trade with full probability. This is the case when $v \sum_{v' \geq v} \mu_0(v')$ weakly decreases in v. Notably, this condition also implies that no information (i.e., complete pooling of all buyer types) is Pareto efficient, as it maximizes the consumer surplus. Proposition 4 part *(ii)* states that the condition that $v \sum_{v' \geq v} \mu_0(v')$ weakly decreases in v is necessary for all categorical information structures to be IC and Pareto efficient.

We return to this example below, after we introduce the concept of bargaining set. This will also enable us to discuss the relation of our framework to some of the closest models in the literature; namely the ones by Bergemann et al. (2015), Ichihashi and Smolin (2022) and Doval and Smolin (2024).

4. Bargaining solutions

Section 3 examined the existence of agreements that strictly benefit all agents. However, merely establishing the possibility of a beneficial agreement does not determine which outcomes will actually be realized in a given environment. In this section, we address this question by formulating a cooperative bargaining problem and analyzing the information content of various bargaining solutions.

The first step is to specify a bargaining set, $\mathcal{B} \subseteq \mathbb{R}^n$, which represents the set of payoff allocations that can be achieved through different incentive-compatible (IC) information structures. Formally,

$$\mathcal{B} \equiv \left\{ (x_i)_{i \in N} \in \mathbb{R}^n | \exists \pi \in \Pi \quad x_i = \mathbb{E}(u_i | \pi) \quad \forall i \in N \right\}.$$

The bargaining set we define is conceptually related to the "Bayesian welfare set" introduced by Doval and Smolin (2024), as both sets capture the welfare consequences of information on a heterogeneous group of agents. However, the key distinction lies in the relationship between the set of agents and the unknown states of the world. In Doval and Smolin (2024), the states of the world correspond to agent types, and the prior distribution represents the population distribution of these types. Consequently, any information structure induces beliefs about the population distribution. In contrast, our model does not impose a direct connection between agents and states of the world.

To illustrate this point, consider the "Price Discrimination and Privacy Regulation" example that we introduced above in Section 3.3. The welfare set, defined as by Doval and Smolin (2024) in this example captures the distribution of ex-ante consumer surplus across the buyer types. In contrast, the bargaining set in our framework—where the agents are a platform and a consumer-surplus-maximizing regulator—is not concerned with how consumer surplus is distributed across buyer types, but rather with how total surplus is split between profit (for the platform/seller) and aggregate consumer surplus. With this formulation, the bargaining set in our model is conceptually closer to the surplus sets considered in Bergemann et al. (2015) and Ichihashi and Smolin (2022).[18]

In cooperative bargaining theory, it is common to assume at the outset that the bargaining set is compact. However, in our framework, this assumption does not immediately hold; because the way the receiver breaks ties when indifferent between multiple actions can introduce discontinuities in agents' payoffs. The solution proposed by Kamenica and Gentzkow (2011) for a single-agent model is to select the sender-preferred equilibrium. However, for multiple agents this solution does not apply. Instead, our formulation of information structures as IC recommendation schedules presents us a solution on this front. As the receiver observes an information structure whether or not agents reach an agreement, we can assume that the receiver follows incentive compatible recommendations in both cases, as stated in the Assumption 3 below.

[17] Data broker companies, such as Experian, often provide information that enables sellers to segment markets based on characteristics of buyers. See for instance https://www.experian.com/blogs/insights/enhanced-customer-segmentation/.

[18] It is possible to construct an alternative environment where our bargaining set also captures the heterogeneity across consumers, making it closely aligned with the Bayesian welfare set in Doval and Smolin (2024). Suppose that each agent corresponds to a buyer type; and the agent that corresponds to a buyer of type $v \in \Omega$ receives the same payoff as the buyer of type v with probability $\mu(v)$ and 0 otherwise, while the seller, who is the receiver in this game, gets $p \sum_{v \geq p} \mu(v)$ at any belief $\mu \in \Delta(\Omega)$. Under this formulation, the Bayesian welfare set in Doval and Smolin (2024) can be interpreted as a bargaining set among these consumer-type agents. Indeed, Corollary 3 in Doval and Smolin (2024) can be understood as a result about the existence of beneficial agreements when the default information structure corresponds to no information.

Assumption 3 *(Tie-breaking)*. 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.

Proposition 5. *Suppose that the tie-breaking Assumption 3 holds. The bargaining set $\mathcal{B}$ is compact and convex.*

When the receiver follows the recommendations generated by the information structure, we can write the expected payoff for each IC information structure as in (1) without explicitly referring to the posterior beliefs. As the expected payoff is linear in recommendation probabilities, the set of payoffs is compact and convex, as long as the set of IC information structures is compact and convex. The set of IC information structures is compact and convex because any IC information structure can be written as a matrix that satisfies finitely many weak linear inequalities.

An informational bargaining problem can be defined as $(\mathcal{B}, d)$, where d is the vector of disagreement payoffs. The disagreement payoff vector is included in the bargaining set, $d \in \mathcal{B}$, as the disagreement information structure is assumed to be incentive compatible; $\pi_0 \in \Pi$. We maintain the tie-breaking Assumption 3 for all of the statements below to guarantee that the most of the standard solution concepts, such as the weighted Nash bargaining and Kalai-Smorodinsky solutions, are well-defined in our framework. A weighted Nash bargaining solution with a weight vector $\alpha \in \Delta(N)$ maximizes the weighted *Nash product*, which can be written for any information structure π as $(\mathbb{E}(u_1|\pi) - d_1)^{\alpha_1} \ldots (\mathbb{E}(u_n|\pi) - d_n)^{\alpha_n}$. The Kalai-Smorodinsky solution picks the maximum payoff vector $x \in \mathcal{B}$ such that $x - d$ is proportional to $(\max_\Pi \mathbb{E}(u_i|\pi) - d_i)_{i\in N}$, where $\max_\Pi \mathbb{E}(u_i|\pi)$ is the highest possible payoff that agent i gets with an IC information structure.[19]

As the bargaining set is not necessarily *d-comprehensive*,[20] the Nash bargaining solution for any fixed weight vector α is not guaranteed to be Pareto efficient. In contrast, the Kalai-Smorodinsky solution is Pareto efficient for every informational bargaining problem.

The informational content of Pareto efficient bargaining solutions varies depending on the context. For instance, in the price discrimination application we discussed in Section 3.3, full information is the best possible information structure for one of the players, so it is always efficient. However, in other contexts, as in the voting example in the Introduction, full information is not necessarily Pareto efficient. In section 4.1, we consider a class of problems that includes the election example from the Introduction. In this environment, we characterize the *endorsement rules* that satisfy Pareto efficiency and some further desirable properties.

4.1. Endorsement rules for selection environment

Information structures generate a total surplus and determine how it is distributed among the agents. Each agent typically has a preferred combination of a receiver strategy and a state of the world that maximizes her individual surplus. An information structure that recommends this strategy in the corresponding state effectively favors that agent. Now, consider an environment where each agent has a distinct preferred strategy-state pair. In such settings, some of the receiver's strategies can be interpreted as a selection process among the agents, while the recommendations produced by the information structure act as endorsements of agents. Thus, in these environments, information structures can be viewed as mechanisms for allocating endorsements. To formalize this idea, we introduce the concept of *selection environments* to capture these ideas.

Definition 2 *(Selection environment)*. A selection environment consists of a one-to-one correspondence between the states of the world $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$, and the receiver actions $A = \{a_0, a_1, \ldots, a_n\}$ with $a_i = \bar{a}(\omega_i)$ for any $i \in N \cup \{0\}$. Additionally, the following conditions hold for any $i, i' \in N$ and $j, j' \in N \cup \{0\}$

(i) the receiver guesses the correct state:

$$u_r(a_j, \omega_{j'}) = \begin{cases} x_j, & \text{if } j = j'; \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

where $x_j > 0$ and $\mu_0(x_0)x_0 > \mu_0(x_j)x_j$;

(ii) each agent gets the highest payoff when she is selected in her corresponding state: $u_i(a_i, \omega_i) \geq u_i(a_j, \omega_{j'})$;

(iii) any selection is better than no selection: $u_i(a_{i'}, \omega_j) > u_i(a_0, \omega_j)$;

(iv) correct selection generates a higher total surplus: $\sum_{l\in N} u_l(a_i, \omega_i) \geq \sum_{l\in N} u_l(a_{j'}, \omega_j)$.

In a selection environment, the receiver's actions consist of either selecting an agent or choosing an alternative action that is inferior for all agents. The receiver's primary objective is to make the correct selection when an agent is the optimal choice and to avoid selecting anyone in state ω_0, where none of the agents is the right choice. The inequality $\mu_0(x_0)x_0 > \mu_0(x_i)x_i$ for any $i \in N$ guarantees that in the absence of information, the receiver's optimal action is the inferior default action. If the receiver picks an agent,

[19] For a discussion on axiomatic definitions of bargaining solutions and a literature survey see Thomson (2022). One of the main distinctions between these two solutions is that the Nash bargaining solution satisfies the *independence of irrelevant alternatives* while the Kalai-Smorodinsky solution does not. A solution satisfies this property if contracting the bargaining set while keeping the disagreement payoff and the initial solution vector does not change the solution. In our framework, contracting $\mathcal{B}$ would mean restricting the set of IC information structures. Such a property could be desirable if for instance the provision of information was costly.

[20] In our setting $\mathcal{B}$ is *d-comprehensive* if for any beneficial information structure $\pi \in \Pi$ and any payoff vector x such that $d \leq x \leq (\mathbb{E}(u_i|\pi))_{i\in N}$, we have $x \in \mathcal{B}$.

the chosen agent receives the highest payoff in case the selection is done at the "correct" state. Agents prefer the receiver to make a selection among themselves over not selecting anyone, and the total welfare is highest when the receiver selects the best agent for the given state.

In selection environments, a natural family of information structures to consider are those that reveal the best agent whenever one exists. While full information is one such structure, there can be alternative information structures that only censor information in cases where the correct choice is no selection. We refer to such structured information policies as "endorsement rules", which ensure that the receiver is always able to identify the best agent when selection is optimal, while withholding information only in cases where selecting no one is the best option.

Definition 3. In a selection environment we call an incentive-compatible information structure $\pi^e \in \Pi$ an **endorsement rule** if for any $i \in N$ $\pi^e(a_i|\omega_i) = 1$.

We show below that a subset of the endorsement rules satisfy some properties that make them natural in selection environments from both positive and normative perspectives. The first such property is the Pareto efficiency. The second is being beneficial relative to the no and full information benchmarks. Finally, endorsement rules are stable against ex-post informational violations of the agreement. In many environments, agents may have opportunities to generate additional information at a low cost after a recommendation is already made by the agreed-upon information structure. A robust information structure ensures that no agent has an incentive to disclose additional information after a realized recommendation because either the recommendation already reveals a state or the expected payoff from any further information is not higher than the current payoff. Full revelation is trivially robust to ex-post deviations as there is nothing that agents can do ex-post.

Formally, we say that an information structure π is **robust to ex-post deviations** if for any action $a \in A$ with $\pi(a|\omega) > 0$ for some $\omega \in \Omega$ there is no other IC information structure π' and agent $i \in N$ such that

$$\sum_{\omega\in\Omega} \mu(\omega|a,\pi) \sum_{a'\in A} \pi'(a'|\omega) u_i(a',\omega) > \sum_{\omega\in\Omega} \mu(\omega|a,\pi) u_i(a,\omega), \tag{7}$$

where $\mu(\cdot|a,\pi)$ is the posterior belief generated by the recommendation a with π. After a recommendation is generated, there is no way to improve the expected payoff of any agent $i \in N$ by implementing any other information structure $\hat{\pi}$. Our main result for the selection environment is as follows.

Theorem 3. *Fix any selection environment. An information structure is robust to ex-post deviations if and only if it is an endorsement rule; and there exists an endorsement rule that is Pareto efficient and beneficial against no and full information.*

Theorem 3 establishes that the three properties of robustness to ex-post deviations, Pareto efficiency, and beneficial agreement against no and full information characterize a subset of endorsement rules. The proof follows from three key observations. First, there is an endorsement rule that maximizes the total (unweighted) sum of payoffs as it minimizes the probability of recommending the inferior strategy of no selection while being incentive compatible. Second, any endorsement rule that recommends all selection strategies with positive probability at state 0, where the correct actions is no selection, is beneficial against no and full information. Finally, any information structure that does not recommend the correct selection with full probability in one of the states where one of agents is the right selection is not robust to ex-post deviation. This is because the agent who is the correct selection can choose disclosure after the realized recommendation.

A symmetric endorsement rule that recommends the selection of all agents with equal probability is procedurally symmetric in the sense that it is independent of the identities of the agents. However, the resulting expected payoffs may still differ across agents unless the payoff structure itself is symmetric. We define a symmetric selection environment as below.

Definition 4. A selection environment is **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 the receiver is indifferent among all agent actions at the prior; that is, $\sum_\omega \mu_0(\omega) u_r(a_i, \omega) = \sum_\omega \mu_0(\omega) u_r(a_j, \omega)$. The payoff from guessing the correct state is x; $u_r(a_i, \omega_i) = x$.

This definition ensures that the structure of payoffs, prior beliefs, and the receiver's preferences exhibit full symmetry across agents. The first condition guarantees that the payoff function remains invariant under any relabeling of agents, states, and actions. The second condition ensures that the prior beliefs assign equal probability to all states corresponding to different agents. The third condition states that, at the prior belief, the receiver does not have a strict preference for any particular agent's selection.

In a symmetric environment, symmetric endorsement rules, where recommendation probabilities are equally distributed among the agents, correspond to symmetric bargaining solutions. Full revelation is such a solution, even though it is not Pareto efficient. However, there is a symmetric and Pareto efficient endorsement rule, and it corresponds to the well-known bargaining solutions.

Proposition 6. *Fix any symmetric selection environment. Any symmetric Nash, Kalai-Smorodinsky, or egalitarian bargaining solution that satisfies robustness to ex-post deviations correspond to the symmetric endorsement rule π, and the symmetric endorsement rule corresponds to symmetric Nash, Kalai-Smorodinsky, or egalitarian bargaining solutions. Moreover, for any $i \in N$ the information structure is $\pi(a_i|\omega_0) = \frac{1}{n}$ if $\frac{x(1-\mu_0)}{x_0\mu_0} \geq 1$ and $\pi(a_i|\omega_0) = \frac{x(1-\mu_0)}{nx_0\mu_0}$ otherwise.*

Symmetry and efficiency are two key properties of symmetric Nash Bargaining and Kalai-Smorodinsky solutions to be proved in our framework. Once we show that there is a symmetric endorsement rule that satisfies these two properties, it is straightforward to show that it also satisfies the remaining axioms.[21]

5. Conclusion

This paper examines how a group of agents collectively determines an information structure to persuade a receiver. We establish general conditions under which agents can reach an agreement and characterize the circumstances that make such agreements possible. The main conclusion is that an agreement is feasible when there is a finite distribution over posterior beliefs of the receiver where each agent expects to receive a payoff higher than some default option. With state-independent preferences, we are able to drop the Bayes' plausibility constraint so the condition is dependent on the prior belief only through the disagreement payoffs. This way we can investigate the possibility to reach an agreement by focusing only on the payoff structure when the default option is full information.

We also analyze a more structured setting where each agent benefits from revealing specific states because the optimal action in those states aligns with her interests. In these environments, the receiver's objective is to correctly infer the state, making information structures that reveal agent-preferred states both Pareto-efficient and incentive compatible. Furthermore, if the status quo action does not correspond to any of such states, an information structure that selectively discloses these states benefits all agents compared to complete censorship and full information. We call this agreement the "endorsement rule."

Our results have direct applications to models of party competition, where multiple parties coordinate their messaging in a joint campaign. The recent success of the opposition coalition in Poland (Applebaum, 2023) demonstrates the potential effectiveness of such strategies. Another relevant application in political economy concerns the decisions made by members of an authoritarian elite regarding how to control media narratives. Beyond political contexts, the framework applies to various economic and regulatory settings, including industrial and financial regulation (Goldstein and Leitner, 2018) and contest design (Zhang and Zhou, 2016; Antsygina and Teteryatnikova, 2022). These examples illustrate how strategic information design can influence decision-making in environments where agents have diverging interests but may still find ways to coordinate on a mutually beneficial information structure.

While the applications discussed in this paper focus on simple settings to clearly illustrate the results, applying them to richer models remains an interesting avenue for future research. One potential extension would be to explore the properties of public versus private information structures as bargaining solutions in settings with multiple receivers.

CRediT authorship contribution statement

Kemal Kıvanç Aköz: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Arseniy Samsonov:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

I hereby declare that I have no financial or material interests relating to the paper "Information Agreements".

Appendix A. Proofs

A.1. Beneficial agreement

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

Lemma 1. *Suppose that Assumption 1 of state-independence holds and the disagreement information structure π_0 is discrete. Fix any optimal strategy $(\bar{a}(\mu))_{\mu\in\Delta(\Omega)}$ by the receiver. There is beneficial agreement if and only if there is a lottery over beliefs $(\lambda_l,\mu_l)_{l=1}^{L}$ such that for any $l=1,\dots,L$ $\lambda_l\in[0,1]$ and $\sum_l\lambda_l=1$ and $\sum_l\lambda_l u_i(\bar{a}(\mu_l))>d_i$ for any agent $i\in N$.*

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

($\Leftarrow$) Fix $\varepsilon>0$ small enough so that for each $k=1,\dots,g$ the belief $\hat{\mu}_k=\varepsilon\mu_0+(1-\varepsilon)\mu_{0k}$ satisfies $\arg\max_{a'\in A}\sum_{\omega\in\Omega}\mu_k(\omega)u_r(a',\omega)=\arg\max_{a'\in A}\sum_{\omega\in\Omega}\mu_{0k}(\omega)u_r(a',\omega)$, 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.

Define the distribution $\hat{\tau}_0$ as for each $k=1,\dots,K$ $\hat{\tau}_0(\hat{\mu}_k)=\tau_0(\mu_{0k})$. Note that τ_0 is Bayesian plausible by construction as it is induced by an information structure π_0. Then,

[21] It is possible to provide tighter conditions for the efficiency of endorsement rules when agents' payoffs are state-independent. See the Online Appendix.

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

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

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

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

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

- $\frac{\epsilon}{\gamma_l'}\le 1$ and $\sum_l\lambda_l=1\Rightarrow\sum_l\frac{\epsilon\lambda_l}{\gamma_l'}\le 1$
- 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 the expected payoff $\mathbb{E}(u_i|\pi)=$

$$\sum_k\hat{\tau}_0(\hat{\mu}_k)\left[\sum_l\frac{\epsilon\lambda_l}{\gamma_l'}(\gamma_l'u_i(\bar{a}(\mu_l))+(1-\gamma_l')u_i(\bar{a}(\hat{\mu}_k)))+\left(1-\sum_l\frac{\epsilon\lambda_l}{\gamma_l'}\right)u_i(\bar{a}(\hat{\mu}_k))\right]$$
$$=\sum_k\hat{\tau}_0(\hat{\mu}_k)\left[\epsilon\sum_l\lambda_l u_i(\bar{a}(\mu_l))+u_i(\bar{a}(\hat{\mu}_k))\left(\sum_l\frac{\epsilon\lambda_l(1-\gamma_l')}{\gamma_l'}+1-\sum_l\frac{\epsilon\lambda_l}{\gamma_l'}\right)\right]$$
$$=\sum_k\hat{\tau}_0(\hat{\mu}_k)\left[\epsilon\sum_l\lambda_l u_i(\bar{a}(\mu_l))+u_i(\bar{a}(\hat{\mu}_k))\left(1-\epsilon\sum_l\lambda_l\right)\right]$$
$$=\epsilon\sum_l\lambda_l u_i(\bar{a}(\mu_l))+\sum_k\hat{\tau}_0(\hat{\mu}_k)u_i(\bar{a}(\hat{\mu}_k))(1-\epsilon)>\sum_k\hat{\tau}_0(\hat{\mu}_k)u_i(\bar{a}(\hat{\mu}_k))=d_i$$

  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 2. *Suppose that Assumption 1 of state-independence holds and the disagreement information structure π_0 is discrete. Fix any optimal strategy $(\bar{a}(\mu))_{\mu\in\Delta(\Omega)}$ by the receiver. Exactly one of the following is true:*

(i): *There is a beneficial agreement.*
(ii): *There exist non-negative numbers $y_1,\dots,y_n$ which are not all equal to 0 such that for any belief $\mu\in\Delta(\Omega)$ we have $\sum_{i\in N}y_i\left(u_i(\bar{a}(\mu))-d_i\right)\le 0$.*

Proof. Applying Ville's Alternative[22] (Corollary 16 in Border (2013)), we get that for any matrix X exactly one of the following is true

- There is a non-negative vector λ such that $X\lambda>0$.

[22] Ville's Alternative is a statement equivalent to the Farkas' Lemma. We use Ville's Alternative because one can directly apply it to the needs of the proof. The original Farkas' Lemma states that given a matrix A and vector b, either there exists a non-negative solution to the equation $Ax=b$ or there exists a vector y such that $A^Ty\ge 0$ and $b^Ty<0$.

- 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 $\eta = \{\mu_l\}_{l=1}^L$, set $X(\eta) \equiv (x_{il}) = u_i(\bar{a}(\mu_l)) - d_i$. In other words, an entry at ith row and lth column is the difference between agent i's expected payoff given belief μ_l and her disagreement payoff.

Suppose that there exists a beneficial information structure. By Lemma 1, there exists a finite collection of beliefs $\eta = \{\mu_l\}_{l=1}^L$ and a vector of non-negative numbers λ such that $\sum_l \lambda_l u_i(\bar{a}(\mu_l)) > d_i \Leftrightarrow \sum_l \lambda_l (u_i(\bar{a}(\mu_l)) - d_i) > 0$ for any agent i.

This statement is equivalent to there being a non-negative vector $\lambda = (\lambda_1, ..., \lambda_L)$ such that $X(\eta)\lambda > 0$. Then, Ville's Alternative 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 $\mu_l \in \eta$ such that $\sum_i y_i(u_i(\bar{a}(\mu_l)) - d_i) > 0$ 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 $\sum_j y_j(u_i(\bar{a}(\mu_j)) - d_j) \leq 0$ for any agent j and belief μ.

Then for any collection of beliefs $\eta = \{\mu_j\}_{j=1}^h$, $X(\eta)^T y \leq 0$ where $y = (y_1, ..., y_n)$. Ville's Alternative implies that there is no non-negative vector λ such that $X(\eta)\lambda > 0$. Hence, there is no beneficial agreement by Lemma 1. □

Proof of Proposition 1. Condition *(iii)* of Theorem 1 can be written as follows: there exists a beneficial agreement if and only if there do not exist non-negative numbers $\{y_i\}_{i=1,\dots,n}$ which are not all equal to 0 such that for any action $a \in A$

$$\Gamma_a \equiv \sum_{i \in N} y_i u_i(a) \leq \sum_{\omega \in \Omega} \mu_0(\omega) \sum_{i \in N} y_i u_i(\bar{a}(\omega))$$
$$= \sum_{i \in N} \Big[\sum_{a' \in A} y_i u_i(a') \big[\sum_{\{\omega \in \Omega | \bar{a}(\omega) = a'\}} \mu_0(\omega) \big] \Big] = \sum_{a' \in A} p_{a'} \sum_{i \in N} y_i u_i(a') = \sum_{a' \in A} p_{a'} \Gamma_{a'},$$

where $p_a = \sum_{\{\omega \in \Omega | \bar{a}(\omega) = a\}} \mu_0(\omega)$. Then, the condition for the non-existence of a beneficial agreement is that for any $a \in A$ the weighted sum Γ_a satisfies $\Gamma_a \leq \sum_{a' \in A} p_{a'} \Gamma_{a'}$, which is equivalent to $\Gamma_a = \Gamma_{a'} \equiv \Gamma$ for any two actions $a, a' \in A$ as $A = \{\bar{a}(\omega)\}_{\omega \in \Omega}$. □

Proof of Proposition 2. Whether the inequality (4) holds or not does not change if we add the same number to all payoffs. Hence, without loss of generality, we can assume that U is strictly positive. In the terms of Christensen (2019), the transpose of the payoff matrix is a "B-matrix" as all of its rows are mean-dominant in their diagonals. By Corollary 4.5. in Carnicer et al. (1999), $detU = detU^T > 0$ and therefore U is invertible. By Lemma 4 in Christensen (2019) the rows of its inverse have positive sums. This implies that the matrix equation $Uy = \mathbf{1}$ has a nonnegative and nonzero solution. Therefore, the weighted sums are constant and so by Proposition 1 there is no beneficial agreement. □

The proof of Theorem 2 follows from the Lemma 3 below.

Lemma 3. *Fix any $k \in \{1, \dots, n\}$. Let $\Phi(k)$ be set of all subsets of $\{1, \dots, n\}$ of size k. For each $\phi \in \Phi(k)$, fix a single vector $s^\phi \in \mathbb{R}^n$ such that $\sum_{j=1}^n s^\phi(j) > 0$ and*

$$\forall j \in \phi \quad s^\phi(j) > 0 \quad \textit{but} \quad \forall j \notin \phi \quad s^\phi(j) < 0.$$

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

Proof. Fix any $k \in \{1, \dots, n\}$ 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 $j \in \{1, \dots, n\}$ 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 $c \in \mathbb{R}$, $v \in \mathbb{R}^n \backslash \{\vec{0}\}$ such that for any $x \in \mathbb{R}^n_+$ and $y \in co(S)$ $\langle x, v \rangle \geq c \geq \langle y, v \rangle$.

Now, we prove a couple of claims.

Claim 1. *ccl1 $c \leq 0$.*

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

Claim 2. *$v \geq 0$.*

Proof. Suppose to the contrary that there exists $i \in \{1, \dots, n\}$ with $v(i) < 0$. 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, $\lim_{\lambda \to \infty} \langle t, v \rangle = -\infty < c$, a contradiction. □

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

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

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

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

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

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

Proof of Theorem 2. Suppose that there is k and $\{\mu^\phi\}_{\phi\in\Phi(k)} \subset \Delta(\Omega)$ as described in the hypothesis. By Lemma 3 there exists a collection of weights $\{\lambda_l\}_{l=1,\ldots,|\Phi(k)|}$ with $\sum_l \lambda_l = 1$ such that for each agent $i = 1, \ldots, n$ we have $\sum_l \lambda_l u_i(\bar{a}(\mu_l)) - d_i > 0$. Then, by Theorem 1, there is a beneficial agreement. □

Proof of Proposition 3. $\Rightarrow$ Consider any information structure π such that for some $v, p' \in \Omega$ with $v < p'$ we have $\pi(p'|v) > 0$ so that trade does not occur with full probability at v. Consider another information structure π' such that $\pi'(v|v) = \pi(v|v) + \pi(p'|v)$, $\pi'(p'|v) = 0$, and for the rest of the entries π and π' coincide. Note first that π' is also an information structure as we simply transferred probability mass between recommendations at the state v. As the term $\pi(p', v)$ does not appear in either of $\mathbb{E}(u_{pl}|\pi)$ or $\mathbb{E}(u_{re}|\pi)$ as defined in equations (5) while $\pi(v|v)$ does, the expected payoffs are strictly higher at π' for both the agents.

Now, we need to show that π' is also incentive compatible. After receiving the recommendation v, the seller should find it optimal to pick v. As π is already incentive compatible, we have for any $v'' \in \Omega$

$$\frac{v\sum_{\hat{v}\geq v}\pi(v|\hat{v})\mu_o(\hat{v})}{\sum_{\hat{v}}\pi(v|\hat{v})\mu_o(\hat{v})} \geq \frac{v''\sum_{\hat{v}\geq v''}\pi(v|\hat{v})\mu_o(\hat{v})}{\sum_{\hat{v}}\pi(v|\hat{v})\mu_o(\hat{v})} \Leftrightarrow$$
$$v\sum_{\hat{v}\geq v}\pi(v|\hat{v})\mu_o(\hat{v}) \geq v''\sum_{\hat{v}\geq v''}\pi(v|\hat{v})\mu_o(\hat{v})$$

The right-hand side of the inequality above stays same when we replace π with π' if $v'' > v$, and the left-hand side becomes higher with π'. For $v'' < v$, the inequality reduces to

$$(v - v'')\sum_{\hat{v}\geq v}\pi(v|\hat{v})\mu_o(\hat{v}) \geq v''\sum_{v''\leq\hat{v}<v}\pi(v|\hat{v})\mu_o(\hat{v}),$$

and again the right-hand side is same for π and π' and the left-hand side is larger with π'.

After receiving any recommendation $p \neq v$, the seller should find it optimal to set the price to p. As π is IC, we have $p\sum_{\hat{v}\geq p}\pi(p|\hat{v})\mu_o(\hat{v}) \geq v''\sum_{\hat{v}\geq v''}\pi(p|\hat{v})\mu_o(\hat{v})$.

Now, we need to prove that this inequality holds for any v'' at also π'. We first show that $\pi(p|\hat{v}) = \pi'(p|\hat{v})$ for any $\hat{v} \geq p$. If $p > v$, then $\hat{v} \geq p > v \Rightarrow \hat{v} \neq v$, so $\pi(p|\hat{v}) = \pi'(p|\hat{v})$. If $v < p$, then the equality is true by construction when $\hat{v} \neq v$. If $\hat{v} = v$, then $\pi(p|\hat{v})$ and $\pi'(p|\hat{v})$ are probabilities of recommending a price lower than v at v, which also do not change from π to π' by construction. Hence, the left-hand side of the inequality is same if we replace π with π'. For any $p \neq v$ and for any $\hat{v}\ \pi'(p|\hat{v}) \leq \pi(p|\hat{v})$ and therefore the right-hand side of the inequality does not increase when we replace π with π'.

$\Leftarrow$ As the price is simply a transfer between the agents, from equations (5) it is clear that the total surplus at any information structure satisfies

$$\mathbb{E}(u_{pl}|\pi) + \mathbb{E}(u_{re}|\pi) = \sum_{v\in\Omega}\mu_0(v)(\gamma_v + v)\sum_{p\leq v}\pi(p|v) \leq \sum_{v\in\Omega}\mu_0(v)(\gamma_v + v).$$

Therefore, any information structure π with full trade probability, i.e. $\sum_{p\leq v}\pi(p|v) = 1$ for any $v \in \Omega$, maximizes the unweighted sum of payoffs, which implies that it is Pareto efficient. □

Proof of Proposition 4. *(i)* By Proposition 3, in a market environment an information structure has no beneficial agreement relative to it if and only if trade always happens. Define $V = \{v_1, \ldots, v_k\}$ as a generic partition element and let $v' \in V$. Then, after observing that the buyer type is in V, the seller chooses a price equal to v_1 if and only if

$$v_1 \geq v'\frac{\mathbb{P}(\omega \in V \cap \omega \geq v')}{\mathbb{P}(\omega \in V)} \Leftrightarrow \frac{v_1}{v'} \geq \frac{\mathbb{P}(\omega \in V \cap \omega \geq v')}{\mathbb{P}(\omega \in V)} = \frac{\mathbb{P}(v' \leq \omega \leq v_k)}{\mathbb{P}(v_1 \leq \omega \leq v_k)}$$

By assumption,

$$v_1\mathbb{P}(\omega \geq v_1) \geq v'\mathbb{P}(\omega \geq v') \Leftrightarrow \frac{v_1}{v'} \geq \frac{\mathbb{P}(\omega \geq v')}{\mathbb{P}(\omega \geq v_1)} =$$

$$\frac{\mathbb{P}(v' \le \omega \le v_k) + \mathbb{P}(\omega > v_k)}{\mathbb{P}(v_1 \le \omega \le v_k) + \mathbb{P}(\omega > v_k)} \ge \frac{\mathbb{P}(v' \le \omega \le v_k)}{\mathbb{P}(v_1 \le \omega \le v_k)} \Rightarrow \frac{v_1}{v'} \ge \frac{\mathbb{P}(v' \le \omega \le v_k)}{\mathbb{P}(v_1 \le \omega \le v_k)} \Leftrightarrow$$
$$v_1 \ge v' \frac{\mathbb{P}(\omega \in V \cap \omega \ge v')}{\mathbb{P}(\omega \in V)}$$

(ii) Suppose that there is $i, j \in \{1, \dots, n\}$ with $i < j$ such that $v_i \sum_{v' \ge v_i} \mu_0(v') < v_j \sum_{v' \ge v_j} \mu_0(v')$. Then, there is i' such that $i' \ge i$ and $i' + 1 \le j$ such that $v_{i'} \sum_{v' \ge v_i} \mu_0(v') < v_{i'+1} \sum_{v' \ge v_{i'+1}} \mu_0(v')$. Then consider any categorical information structure π such that one of the partition elements is $\{v_{i'}, v_{i'+1}\}$. Then, by the supposition, and from a similar argument as above the seller sets the price $v_{i'+1}$ whenever it receives the information $\{v_{i'}, v_{i'+1}\}$, and so trade does not happen if the buyer type is $v_{i'}$. By Proposition 3, there is a beneficial agreement. □

A.2. Bargaining set

Lemma 4. *Let Π be the set of incentive-compatible information structures. Π is a compact and convex subset of $[0,1]^{|\Omega| \times |A|}$.*

Proof. *Convexity*: We know that each information structure induces a Bayesian plausible distribution over beliefs. As we can represent information structures as recommendation rules, for any two incentive-compatible information structures π and $\hat{\pi}$, we can write the supports of the induced belief distributions as $\{\mu(\cdot|a, \pi)\}_{a \in A}$ and $\{\mu(\cdot|a, \hat{\pi})\}_{a \in A}$. Fix any $\gamma \in (0, 1)$ and let $\pi_\gamma = \gamma\pi + (1-\gamma)\hat{\pi}$.

We have three cases to consider. First take any $a \in A$ such that $\pi(a|\omega) > 0$ and $\hat{\pi}(a|\omega) > 0$. Now consider the posterior belief $\mu(\cdot|a, \pi_\gamma)$, which can be written as

$$\mu(\omega|a, \pi_\gamma) = \frac{(\gamma\pi(a|\omega) + (1-\gamma)\hat{\pi}(a|\omega))\mu_0(\omega)}{\sum_{\omega' \in \Omega}(\gamma\pi(a|\omega') + (1-\gamma)\hat{\pi}(a|\omega'))\mu_0(\omega')}$$
$$= \frac{\gamma P(a|\pi)\mu(\omega|a, \pi) + (1-\gamma)P(a|\hat{\pi})\mu(\omega|a, \hat{\pi})}{\gamma P(a|\pi) + (1-\gamma)P(a|\hat{\pi})},$$

where for any information structure π', the probability of recommending any action a is $P(a|\pi') = \sum_{\omega \in \Omega} \pi'(a|\omega)\mu_0(\omega)$. Then, as $\mu(\omega|a, \pi_\gamma)$ is a convex combination of $\mu(\omega|a, \pi)$ and $\mu(\omega|a, \hat{\pi})$, and a is optimal for the receiver under both beliefs, a is also optimal under $\mu(\omega|a, \pi_\gamma)$.

As the second case, consider any $a \in A$ such that only π recommends a with positive probability but $\hat{\pi}$ does not recommend it. Then, clearly $\mu(\omega|a, \pi_\gamma) = \mu(\omega|a, \pi)$, and so since π is incentive compatible, a is optimal for the receiver under $\mu(\omega|a, \pi_\gamma)$. Finally, if neither of the two information structures recommend some $a \in A$, π_γ also does not recommend it. These arguments show that π_γ is incentive compatible.

Compactness: Take any convergent sequence of incentive compatible information structures $\{\pi_l\}_{l \ge 1}$ and denote their limit as $\lim \pi$. It is clear that $\lim \pi$ is also an information structure. For any a such that $\lim \pi(a|\omega) > 0$ for some $\omega \in \Omega$ there is $\bar{l}$ such that for any $l > \bar{l}$, $\pi_l(a|\omega) > 0$, which implies a is optimal under the belief $\mu_a(\pi_l)$. As the Bayesian update is continuous in information structures, a is also optimal at $\mu_a(\lim \pi)$. Therefore, $\lim \pi$ is also incentive compatible. □

Proof of Proposition 5. B is bounded as A is finite and Π is compact by Lemma 4. Now, take any converging sequence of incentive compatible information structures $\{\pi_l\}_{l \ge 1}$ with $\lim_{l \to \infty} \pi_l = \bar{\pi} \in \Pi$. Then, for any agent $i \in \{1, \dots, n\}$

$$\lim_{l \to \infty} \mathbb{E}(u_i|\pi_l) = \sum_{a \in A} \sum_{\omega \in \Omega} \lim_{l \to \infty} \pi_l(a|\omega)\mu_0(\omega)u_i(a, \omega)$$
$$= \sum_{a \in A} \sum_{\omega \in \Omega} \bar{\pi}(a|\omega)\mu_0(\omega)u_i(a, \omega) = \mathbb{E}(u_i|\bar{\pi}),$$

where the second equality follows from Assumption 3.

Finally, to prove convexity take any two payoff vectors v, v' such that $v = \left(\mathbb{E}(u_i|\pi)\right)_{i \in \{1,\dots,n\}}$ and $v' = \left(\mathbb{E}(u_i|\pi')\right)_{i \in \{1,\dots,n\}}$ for some two incentive compatible information structures π and π'. Fix an arbitrary $\gamma \in (0, 1)$. Consider the following information structure π_γ that applies π with probability γ and π' with the remaining probability. By Lemma 4, π_γ is an incentive compatible information structure and $\mathbb{E}(u_i|\pi_\gamma) = \gamma\mathbb{E}(u_i|\pi) + (1-\gamma)\mathbb{E}(u_i|\pi')$ for any $i \in \{1, \dots, n\}$, which proves the convexity of B. □

A.3. Endorsement rules

Proof of Theorem 3. We first prove that there is an endorsement rule that is Pareto efficient, beneficial against no and full information, and robust to ex-post deviations. To prove Pareto efficiency, take any incentive compatible information structure π that recommends all actions $\{a_1, \dots, a_n\}$ 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 $i \in \{0, 1, \dots, n\}$ $\pi^e_{i0} = \pi_{i0}$.

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 $i, i' \in \{0, 1, \dots, n\}$

$$\mathbb{E}(u_r(a_i,\omega)|\mu_{a_i},\pi) \geq \mathbb{E}(u_r(a_{i'},\omega)|\mu_{a_i},\pi) \Leftrightarrow$$
$$\mu_0(\omega_i)\pi_{ii}u_r(a_i,\omega_i) \geq \mu_0(\omega_j)\pi_{ij}u_r(a_j,\omega_j). \tag{8}$$

For the incentive compatibility of π^e, for any $i \in \{1,\dots,n\}$ we have $\mu_0(\omega_i)u_r(a_i,\omega_i) \geq \mu_0(\omega_0)\pi_{i0}u_r(a_0,\omega_0)$, which holds by equation (8). Therefore, the receiver has an incentive to follow any recommendation a_i, and as the posterior belief $\mu(\omega_0|a_0,\pi^e) = 1$, the receiver also follows the recommendation a_0.

For any agent $i \in N$, π^e is beneficial relative to the 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 Definition 2. π^e is also beneficial against the 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. Let $\Gamma_{ij} = \sum_{l\in N} u_l(a_i,\omega_j)$ for any $i, j = 0, 1, \dots, n$. 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 the conditions in Definition 2 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. □

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. By Claim 1, π^e is also beneficial and incentive compatible. It is straightforward to see that π^e is robust to ex-post deviations.

Claim 3. *An IC information structure π is robust to ex-post deviations $\Rightarrow$ π is an endorsement rule.*

Proof. Suppose that π is not an endorsement rule. Then, there is $i \in N$ and $j \in N\cup\{0\}\backslash\{i\}$ such that $\pi(a_j|\omega_i) > 0$. Let's first consider the case where $j \in N\backslash\{i\}$. Then, we can claim that agent j wants to share the following information structure π', which is defined as $\pi'(a_j|\omega_j) = 1$ and $\pi'(a_i|\omega) = 1$ for any $\omega \neq \omega_j$, after the recommendation a_i generated by π. Note first that π' is also IC. The recommendation a_j is clearly optimal for the receiver. And the recommendation a_i is also optimal as for any j'

$$x_i\frac{\mu(\omega_i|a_i,\pi)}{1-\mu(\omega_j|a_i,\pi)} \geq x_{j'}\frac{\mu(\omega_{j'}|a_i,\pi)}{1-\mu(\omega_j|a_i,\pi)},$$

which holds as π itself is also IC.

Now, the left-hand side of the inequality (7) is

$$\mu(\omega_j|a_i,\pi)u_j(a_j,\omega_j) + \sum_{\omega\neq\omega_j}\mu(\omega|a_i,\pi)u_j(a_i,\omega),$$

which is greater than the right-hand side by Definition 2. □

With this claim, the proof of Theorem 3 is complete. □

Proof of Proposition 6. The first part of the statement immediately follows from Theorem 3 as any bargaining solution that satisfies robustness to ex-post deviations has to correspond to an endorsement rule. Now, for the reverse part we define a bargaining solution that corresponds to the symmetric endorsement rule.

Let's first define a bargaining solution over B that corresponds to the symmetric endorsement rule. Let $x^{ser} = (x_i^{ser})_{i\in N}$ be the payoff vector generated by the symmetric endorsement rule. We introduce a solution that applies to any compact and convex subset $B' \subseteq B$ to show that, like the Nash bargaining solution, it satisfies the contraction independence axiom. For any compact and convex subset $B' \subseteq B$, let $f^{nb}(B') = (f_i^{nb})_{i\in N}(B')$ denote the payoff vector that corresponds to the Nash bargaining solution for the set of feasible payoffs B'. Define $f^{ser}(B') = x^{ser}$ if $x^{ser} \in B'$ and $f^{ser}(B') = f^{nb}(B')$ otherwise.

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 the symmetric endorsement rule are symmetric, linear in payoff parameters, and constant over all the subsets B' as long as $x^{ser} \in B'$.

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 $i \neq 0$; $v_2 \equiv \sum_k u_k(a_j,\omega_i)$ where $j \neq 0, i$ and $i \neq 0$; $v_3 \equiv \sum_k u_k(a_0,\omega_i)$ where $i \neq 0$; $v_4 \equiv \sum_k u_k(a_j,\omega_0)$ where $j \neq 0$; $v_5 \equiv \sum_k u_k(a_0,\omega_0)$.

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

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

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 $\mu_0[v_4(1-\pi_{00}) + \pi_{00} v_5]$ decreases in π_{00}. In the rest of the proof, we will use the 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 $x\pi_{ii}\frac{1-\mu_0}{n} \geq x_0 \pi_{i0}\mu_0 \Leftrightarrow x\pi_{ii}(1-\mu_0) \geq n x_0 \pi_{i0}\mu_0$.

We proceed by analyzing 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

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

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 $i \neq 0$: $\pi_{i0} = \frac{x(1-\mu_0)}{nx_0\mu_0}$, 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, 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 $\eta_i = \max_\pi \mathbb{E}(u_i|\pi)$.

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)). □

Appendix B. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jet.2025.106068.

Data availability

No data was used for the research described in the article.

References

Anbarci, Nejat, Feltovich, Nick, 2013. How sensitive are bargaining outcomes to changes in disagreement payoffs? Exp. Econ. 16, 560–596.
Antsygina, Anastasia, Teteryatnikova, Mariya, 2022. Optimal information disclosure in contests with stochastic prize valuations. Econ. Theory, 1–38.
Applebaum, Ann, 2023. Poland shows that autocracy is not inevitable. https://www.theatlantic.com/ideas/archive/2023/10/poland-parliamentary-election-autocracy-tusk/675656/.
Au, Pak Hung, Kawai, Keiichi, 2020. Competitive information disclosure by multiple senders. Games Econ. Behav. 119, 56–78.
Bergemann, Dirk, Brooks, Benjamin, Morris, Stephen, 2015. The limits of price discrimination. Am. Econ. Rev. 105 (3), 921–957.
Bhattacharya, Sourav, Mukherjee, Arijit, 2013. Strategic information revelation when experts compete to influence. Rand J. Econ. 44 (3), 522–544.
Board, Simon, Lu, Jay, 2018. Competitive information disclosure in search markets. J. Polit. Econ. 126 (5), 1965–2010.
Border, K.C., 2013. Alternative linear inequalities. Lecture notes. https://kcborder.caltech.edu/Notes/Alternative.pdf. (Accessed 26 February 2025).
Boyd, Stephen, Boyd, Stephen P., Vandenberghe, Lieven, 2004. Convex Optimization. Cambridge University Press.
Carnicer, Jésus M., Goodman, Tim N.T., Peña, Juan Manuel, 1999. Linear conditions for positive determinants. Linear Algebra Appl. 292 (1–3), 39–59.
Chang, Yenjae, Winston, Clifford, Yan, Jia, 2022. Does uber benefit travelers by price discrimination? J. Law Econ. 65 (S2), S433–S459.
Christensen, Finn, 2019. Comparative statics and heterogeneity. Econ. Theory 67 (3), 665–702.
Crawford, Vincent Z., 1997. Theory and experiment in the analysis of strategic interaction. In: Kreps, David, Wallis, Kenneth (Eds.), Advances in Economics and Econometrics, Seventh World Congress: Theory and Applications, Vol. I. In: Econometric Society Monographs, vol. 27. Cambridge University Press, Cambridge, U.K., pp. 206–242. Chapter 7.
Doval, Laura, Smolin, Alex, 2024. Persuasion and welfare. J. Polit. Econ. 132 (7), 2451–2487.
Dughmi, Shaddin, Xu, Haifeng, 2016. Algorithmic Bayesian persuasion. In: Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, pp. 412–425.
Gentzkow, Matthew, Kamenica, Emir, 2017a. Bayesian persuasion with multiple senders and rich signal spaces. Games Econ. Behav. 104, 411–429.
Gentzkow, Matthew, Kamenica, Emir, 2017b. Competition in persuasion. Rev. Econ. Stud. 84 (1), 300–322.
Goldstein, Itay, Leitner, Yaron, 2018. Stress tests and information disclosure. J. Econ. Theory 177, 34–69.
Hoffmann, Florian, Inderst, Roman, Ottaviani, Marco, 2020. Persuasion through selective disclosure: implications for marketing, campaigning, and privacy regulation. Manag. Sci. 66 (11), 4958–4979.
Ichihashi, Shota, Smolin, Alex, 2022. Data collection by an informed seller.
Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. Am. Econ. Rev. 101 (6), 2590–2615.
Karagözoğlu, Emin, 2019. On "going unstructured" in bargaining experiments. In: The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers, pp. 295–304.
Li, Fei, Norman, Peter, 2018. On Bayesian persuasion with multiple senders. Econ. Lett. 170, 66–70.
Li, Fei, Norman, Peter, 2021. Sequential persuasion. Theor. Econ. 16 (2), 639–675.
Myerson, Roger B., 1979. Incentive compatibility and the bargaining problem. Econometrica, 61–73.
Myerson, Roger B., 2008. Perspectives on mechanism design in economic theory. Am. Econ. Rev. 98 (3), 586–603.
Myerson, Roger B., Satterthwaite, Mark A., 1983. Efficient mechanisms for bilateral trading. J. Econ. Theory 29 (2), 265–281.

Thomson, William, 1994. Cooperative models of bargaining. In: Handbook of Game Theory with Economic Applications, vol. 2, pp. 1237–1284.
Thomson, William, 2022. On the axiomatic theory of bargaining: a survey of recent results. Rev. Econ. Des. 26 (4), 491–542.
Zhang, Jun, Zhou, Junjie, 2016. Information disclosure in contests: a Bayesian persuasion approach. Econ. J. 126 (597), 2197–2217.