

Private Private Information*

Kevin He[†] Fedor Sandomirskiy[‡] Omer Tamuz[§]

Abstract

A *private* private information structure delivers information about an unknown state while preserving privacy: An agent's signal contains information about the state but remains independent of others' sensitive or private information. We study how informative such structures can be, and characterize those that are optimal in the sense that they cannot be made more informative without violating privacy. We connect our results to fairness in recommendation systems and explore a number of further applications.

*The paper benefited from numerous suggestions and comments of our colleagues. We are grateful to Marina Agranov, Nageeb Ali, Itai Arieli, Eric Auerbach, Yakov Babichenko, Dirk Bergemann, Alexander Bloedel, Aislinn Bohren, Simina Brânzei, Benjamin Brooks, David Dillenberger, Piotr Dworzak, Federico Echenique, Jeffrey Ely, Drew Fudenberg, Hanming Fang, Wayne Gao, Kai Hao Yang, David Kempe, Toygar Kerman, Victoria Kostina, Annie Liang, Jonathan Libgober, Elliot Lipnowski, George Mailath, Moritz Meyer-ter-Vehn, Xiaosheng Mu, Luciano Pomatto, Andy Postlewaite, Doron Ravid, Ran Shorrer, Ran Spiegler, Bruno Strulovici, Ina Taneva, Caroline Thomas, Tristan Tomala, Nicolas Vieille, Leeat Yariv, and seminar participants at Caltech, University College London, UC Riverside, USC, Hong Kong Joint Theory Seminar, HSE Moscow and St. Petersburg, University of Maryland, Northwestern University, University of Pennsylvania, Purdue University, Stanford, 2021 LA Theory Conference, Pennsylvania Economic Theory Conference, Retreat on Information, Networks, and Social Economics, and Stony Brook Workshop on Strategic Communication and Learning.

[†]University of Pennsylvania. Email: hesichao@gmail.com.

[‡]Caltech. Email: fsandomi@caltech.edu. Fedor Sandomirskiy was supported by the Linde Institute at Caltech and the National Science Foundation (grant CNS 1518941).

[§]Caltech. Email: tamuz@caltech.edu. Omer Tamuz was supported by a grant from the Simons Foundation (#419427), a Sloan fellowship, a BSF award (#2018397) and a National Science Foundation CAREER award (DMS-1944153).

1 Introduction

Privacy has emerged as an important design concern in contemporary markets, characterized by the pervasive collection and exchange of information. In consequence, privacy has been the focus of a large literature in computer science (Dwork and Roth, 2014) and a more recent literature in economics (Eilat et al., 2021; Acemoglu et al., 2022; Bergemann et al., 2022; Liang and Madsen, 2023). In this paper, we study information design under a privacy constraint: The signal sent to an agent must contain no information about other agents’ private or sensitive information.

As a motivating example, consider an employer who is interested in an unknown state: whether or not an applicant is qualified for the job. A recommender (e.g., an automated recommendation system or a human letter writer) knows the applicant’s qualification but also knows a sensitive attribute, such as the applicant’s age. The recommender wishes to give as much information as possible to the employer about the applicant’s qualification, but does not want to reveal any information about the sensitive attribute.

If the state and the attribute are independent, then the recommender can simply report the state. However, if the state and the attribute are correlated, reporting the state also inadvertently reveals some information about the attribute. The recommender faces a privacy-constrained information-design problem: How to optimally give information about the state, while keeping the recommendation independent of the sensitive attribute? Intuitively, as the correlation between the sensitive attribute and the state increases, the privacy constraint becomes more restrictive and further limits how much information the recommender can provide.

Generalizing this example, we view the attribute and the recommendation as two signals about the state. Requiring independence between these two signals imposes a joint restriction on their informativeness. This paper studies the maximal informativeness of signals independent of each other and applies the results to settings like the recommender’s problem above.

In a *private private information structure*, agents observe private signals about the state, and these signals are statistically independent of each other. Thus the information available to each agent reveals nothing at all about the information available to her peers.¹ In other words, private signals are indeed completely private. Privacy

¹The existence of (non-trivial) private private information structures is not immediate. To con-

imposes a restriction on the informativeness of such information structures. For example, it is clearly impossible for two agents to both have signals that perfectly reveal the state while maintaining privacy. What is the maximum amount of information that can be conveyed through private signals? We formalize this question using the notion of Pareto optimality with respect to the Blackwell order: a private information structure is *Blackwell-Pareto optimal* if it is impossible to give more information to any agent—in the Blackwell sense—without violating privacy.² The main goal of this paper is to understand the Blackwell-Pareto frontier, allowing a designer to produce optimal private information structures. This includes, for example, producing the most informative recommendation that does not reveal information about a sensitive attribute.

Results. In the case of two agents and a binary state, we obtain a simple description of the Blackwell-Pareto frontier: Given a private information structure, denote by $F_1, F_2: [0, 1] \rightarrow [0, 1]$ the cumulative distribution functions of the two agents’ beliefs about the state. Then the structure is Blackwell-Pareto optimal if and only if $F_1(x) = 1 - F_2^{-1}(1 - x)$. This characterization allows for a straightforward test of optimality and a constructive procedure for finding optimal structures.

The Blackwell-Pareto frontier is agnostic about a particular designer’s objective. We refine our characterization of Blackwell-Pareto optimal structures by considering a designer who maximizes a welfare-like objective that is monotone and separable in the informativeness of individual signals. We show that the optimum is always attained at a structure where one agent has at most three possible signal realizations,

struct an example, suppose three agents each receive independent binary signals, and the unknown state is the majority of these three signals; similar information structures are used in the social learning literature (Gale, 1996; Çelen and Kariv, 2004a,b). Because of independence, each agent holds the same belief about others’ signal realizations when her own signal is high and when it is low. And yet, every signal is informative about the state. It is important to note that conditionally independent private signals—a common modeling tool—are not private. For instance, consider a binary state and conditionally independent binary signals that match the state with a certain probability. Here, each agent’s signal contains information about the others’ signals: when an agent observes a high signal, she becomes more confident that her peers also got high signals.

²The term Pareto optimality is usually used in the context of allocating goods to agents who have preferences over them. Here we think of an information structure as an allocation of information (under the privacy constraint) and assume that agents prefer signals that are more Blackwell informative over those that are less.

and the other has at most two possible signal realizations. Thus welfare maximization reduces to a simple finite-dimensional optimization problem.

We also study Blackwell-Pareto optimality in the general case with any number of agents, where our characterization provides a surprising connection to the field of mathematical tomography and the study of sets of uniqueness. A subset of $[0, 1]^n$ is called a *set of uniqueness* if it is uniquely determined by the densities of its projections to the coordinate axes. Understanding such sets has been an active area of research since the 1940s (Lorentz, 1949). This problem gained more prominence with the advent of tomography, a technology to reconstruct three-dimensional objects from their projections (Gardner, 1995).

We show that private private information structures for n agents can be identified with subsets of $[0, 1]^n$, and that the Blackwell-Pareto optimal ones correspond exactly to sets of uniqueness. In the two-dimensional case—which corresponds to the case of two agents and a binary state—the complete characterization of the sets of uniqueness is known (Lorentz, 1949) and leads to our characterization of the Blackwell-Pareto frontier. With more agents, we rely on later results on sets of uniqueness for $n \geq 3$ (Fishburn et al., 1990) to provide some sufficient conditions and some necessary conditions for Blackwell-Pareto optimality. With three or more states, an analogous equivalence holds between Blackwell-Pareto optimality and a generalization of sets of uniqueness that we term *partitions of uniqueness*.

Finally, we discuss what belief distributions are feasible, i.e., can be induced by private private information structures. Using an information-theoretic approach, we provide simple constraints on the informativeness of private private signals, thus obtaining necessary conditions of feasibility. These constraints indicate that one can think of private private signals as a way to divide an “information pie” among agents. However, the division creates negative externalities so that the sum of the pieces is less than the whole pie. The inequality formalizing this observation might be of independent interest in information theory.

Applications to privacy in recommendation systems and beyond. Our results can be applied to the problem of *optimal privacy-preserving recommendation*, a particular case of which is the motivating example presented above.³

³Our focus on informativeness of private private signals is also relevant to a number of other economic problems—the feasibility of reduced-form mechanisms, information design in zero-sum

In the problem of optimal privacy-preserving recommendation, the goal is to design a maximally informative signal (a recommendation) about the state under the constraint that this recommendation is independent of a given random variable correlated with the state. For example, the state can be a job applicant’s productivity type or a loan applicant’s creditworthiness, and a recommendation must be as informative about the state as possible but carry no information about sensitive attributes such as the applicant’s age or health condition.

When the state is binary (i.e., the applicant’s productivity type takes one of two values), our results on Blackwell-Pareto optimal private information imply a complete solution to the problem of optimal privacy-preserving recommendation. As we show, in this case, there is a unique optimal recommendation that Blackwell dominates any other recommendation that preserves the applicant’s privacy. Our proof is constructive and gives a simple recipe for generating this optimal recommendation. We also quantify information loss due to the privacy constraint and discuss its comparative statics.

Curiously, the uniqueness of the optimal recommendation for binary states means that knowledge of the decision-maker’s objective is unnecessary. Hence, decisions and recommendations can be decoupled, e.g., by creating an agency assessing job applicants for all prospective employers. By contrast, for three or more states, we give an example where the optimal recommendation is not unique, so delegation to an agency could lead to efficiency losses.

Our privacy requirement admits an alternative fairness interpretation within the field of algorithmic fairness, an active area of research at the interface of economics and computer science; see, e.g., [Liang et al. \(2022\)](#); [Kleinberg et al. \(2018\)](#); [Rambachan et al. \(2020\)](#); [Mehrabi et al. \(2021\)](#) and references therein. As before, consider an AI-powered recommendation system assessing, for example, a loan applicant’s creditworthiness, based on abundant available data. A major concern in designing such systems revolves around properly handling legally protected attributes (e.g., applicants’ gender or race) to avoid discrimination. The common wisdom is that banning the use of protected attributes may lead to even worse discrimination, as correlated attributes can be used as proxies ([Agan and Starr, 2018](#)). Instead, non-discriminatory behavior is to be imposed as a separate fairness constraint when designing the system

games, and causal inference—discussed in [Appendix A](#).

(Barocas et al., 2019).

Our privacy requirement turns out to be equivalent to the fairness constraint of *statistical parity*, one of the most popular fairness principles (Kleinberg et al., 2016; Aswani and Olfat, 2022; Dwork et al., 2012). Statistical parity requires that the distribution of the recommendations (e.g., the fraction of agents with a given level of credit score) has to be the same across sub-populations with different values of protected attributes.⁴ In other words, the recommendation and the collection of protected attributes have to be independent of each other, forming a private information structure. Thus, our results on optimal privacy-preserving recommendation provide a characterization of optimal recommendation systems satisfying statistical parity. Moreover, our framework allows one to quantify the efficiency loss due to the fairness constraint.

Our analysis makes two simplifying assumptions. First, we assume that the recommender can infer the state exactly from the data.⁵ Second, we impose an exact fairness requirement, as common in the algorithmic fairness literature, taking the normative approach and treating it as a hard design constraint. In this aspect, our paper differs from the positive perspective on privacy loss in economic mechanisms developed by Eilat et al. (2021), as well as from the differential privacy literature that focuses on aggregate rather than individual assessment and achieves partial privacy by adding noise to individual characteristics (Dwork and Roth, 2014).

In what follows, we stick to the privacy rather than the fairness interpretation of optimal recommendation for the sake of consistency.

More Related Literature. The question of which belief distributions can arise in private information structures was addressed in Gutmann et al. (1991) and Arieli et al. (2021). They provide a characterization for two agents under additional symmetry assumptions; we discuss the relation to our work below. Hong and Page (2009) consider the feasibility of private signals in a stylized setting with binary signals and binary states of nature. Concurrently and independently,

⁴The requirement of statistical parity is close in spirit to the concept of disparate impact from the US law, which forbids substantial statistical differences in outcomes for groups with different values of protected attributes.

⁵Extending our analysis to the case of a recommender observing a noisy signal about the underlying state is a promising direction for future research.

Cichomski (2020) studies private private structures that maximize the expected divergence between the agents’ posterior beliefs. His approach relies on a connection to the Gale-Ryser Theorem—a classical result on the existence of bipartite graphs with given degree distributions. The analysis suggests that this theorem can be used to derive the same two-agent, binary state feasibility result as our Corollary 2 in Appendix 7. Cichomski and Petrov (2022) rely on the connection to bipartite graphs to prove a tight upper bound on the probability that two beliefs induced by private private signals differ by more than some given δ . For the question of feasibility without the privacy constraint, see, e.g., Dawid et al. (1995); Burdzy and Pal (2019); Burdzy and Pitman (2020); Arieli et al. (2021); Cichomski and Osekowski (2021).

Private private signals arise as the worst-case information structure for the auctioneer in some problems of robust mechanism design: see Bergemann et al. (2017) and Brooks and Du (2021). Private private signals also appear as counterexamples of information aggregation in financial markets: see the discussion in Ostrovsky (2012) and similar observations in the computer science literature (Feigenbaum et al., 2003).

Our application to influencing competitors in zero-sum games relates to the literature of information design in games (Bergemann and Morris, 2016; Taneva, 2019; Mathevet et al., 2020). Our applications to the optimal privacy-preserving recommendation and welfare maximization have a high-level similarity to the type of questions studied by Eliaz et al. (2020), who also consider an optimization problem on random variables under an independence constraint.

As mentioned above, our work is related to the mathematics of sets of uniqueness and mathematical tomography (Lorentz, 1949; Fishburn et al., 1990; Kellerer, 1993). These techniques have been applied in economics, for example, by Gershkov et al. (2013) to show the equivalence of Bayesian and dominant strategy implementation in an environment with linear utilities and one-dimensional types.

We use a derandomization argument to show that, without loss of generality, one can assume that the signals determine the state. Chen et al. (2019) obtain a similar derandomization result in a mechanism design context.

2 Model

We consider a group of agents $N = \{1, \dots, n\}$ where each agent i has a signal s_i containing information about a state of nature ω taking value in $\Omega = \{0, 1, \dots, m-1\}$,

and all agents start with a common, full-support prior belief about the state. We call the tuple $\mathcal{I} = (\omega, s_1, \dots, s_n)$ an *information structure*. Formally, fix a standard nonatomic Borel probability space (X, Σ, \mathbb{P}) , and let ω, s_1, \dots, s_n be random variables defined on this space that take values in $\Omega \times S_1 \times \dots \times S_n$, where each S_i is a measurable space.⁶ The marginal distribution of ω is the prior over the state.

Denote by $p(s_i)$ the posterior associated with s_i . Formally, $p(s_i)$ is the random variable taking value in $\Delta(\Omega)$ given by $p(s_i)(k) = \mathbb{P}[\omega = k \mid s_i]$. In the case of a binary state (i.e., when $\Omega = \{0, 1\}$), we let $p(s_i)$ take value in $[0, 1]$ by setting $p(s_i) = \mathbb{P}[\omega = 1 \mid s_i]$.

Definition 1. *We say that $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is a private private information structure if (s_1, \dots, s_n) are independent random variables.*

Private private signals should not be confused with *conditionally* independent signals, where (s_1, \dots, s_n) are independent given ω . This is a different notion, and indeed conditionally independent signals cannot be private private, unless they are fully uninformative about the state.

As a simple example of a private private information structure for two agents and a binary state, let s_1, s_2 be independently and uniformly distributed on $[0, 1]$, and let ω be the indicator of the event that $s_1 + s_2 > 1$, as illustrated in Figure 1. The distribution of (s_1, s_2) conditioned on $\omega = 1$ is the uniform distribution on the upper right triangle of the unit square. Conditioned on $\omega = 0$, (s_1, s_2) have the uniform distribution on the bottom left triangle. Note that the posterior beliefs are $p(s_i) = s_i$ in this information structure, so both agents have strictly informative signals. While the two signals are independent, they are not conditionally independent given the state ω .

This paper focuses on characterizing the private private signals that are maximally informative for the group of agents, formalized through the concept of *Blackwell-Pareto optimality* of private private information structures. For the single-agent case ($n = 1$), recall that an information structure (ω, s) *Blackwell dominates* (ω, \hat{s}) if for every continuous convex $\varphi: \Delta(\Omega) \rightarrow \mathbb{R}$ it holds that $\mathbb{E}[\varphi(p(s))] \geq \mathbb{E}[\varphi(p(\hat{s}))]$.

This notion captures a strong sense in which s contains more information about ω than \hat{s} does: in any decision problem, an agent maximizing expected utility performs

⁶An alternative approach would be to define an information structure as a joint distribution over $\Omega \times S_1 \times \dots \times S_n$.




Figure 1: The pair of signals (s_1, s_2) is uniformly distributed on the unit square, with $\omega = 1$ in the black area and $\omega = 0$ in the white area. The induced posteriors $p(s_1), p(s_2)$ coincide with the signals.

better when observing s than when observing \hat{s} .

For more than one agent, our next definition introduces a partial order on private private information structures that captures the Blackwell dominance for each agent.

Definition 2. Let $\mathcal{I} = (\omega, s_1, \dots, s_n)$ and $\hat{\mathcal{I}} = (\omega, \hat{s}_1, \dots, \hat{s}_n)$ be private private information structures. We say that \mathcal{I} dominates $\hat{\mathcal{I}}$, and write $\mathcal{I} \geq \hat{\mathcal{I}}$, if for every i it holds that (ω, s_i) Blackwell dominates (ω, \hat{s}_i) . We say that \mathcal{I} and $\hat{\mathcal{I}}$ are equivalent if $\mathcal{I} \geq \hat{\mathcal{I}}$ and $\hat{\mathcal{I}} \geq \mathcal{I}$.

It follows from this definition that \mathcal{I} is equivalent to $\hat{\mathcal{I}}$ if and only if, for each i , the distributions of $p(s_i)$ and $p(\hat{s}_i)$ coincide. Thus we can partition the set of private private information structures into equivalence classes, with each class represented by n distributions (μ_1, \dots, μ_n) on $\Delta(\Omega)$. The first question that arises is that of feasibility: which n -tuples (μ_1, \dots, μ_n) represent some private private information structure? We address this question in §7.

Figure 2 illustrates another example of a private private information structure, where the signals are again uniform on $[0, 1]$, but each agent’s posterior belief is equally likely to be either $1/4$ or $3/4$. Thus this structure is equivalent to a structure where agents receive binary signals. More generally, a structure $(\omega, s_1, \dots, s_n)$ is always equivalent to the “direct revelation” structure $(\omega, p(s_1), \dots, p(s_n))$ in which agent i observes the posterior belief induced by s_i .

We use the concept of dominance to define Blackwell-Pareto optimality: which private private information structures provide a maximal amount of information to the agents, so that more information cannot be supplied without violating privacy?




Figure 2: The pair of signals (s_1, s_2) is uniformly distributed on the unit square, with $\omega = 1$ in the black area and $\omega = 0$ in the white area. The induced posteriors $p(s_1), p(s_2)$ are binary, and equally likely to be either $1/4$ or $3/4$.

Definition 3. We say that a private private information structure \mathcal{I} is Blackwell-Pareto optimal if, for every private private information structure $\hat{\mathcal{I}}$ such that $\hat{\mathcal{I}} \geq \mathcal{I}$, the structure $\hat{\mathcal{I}}$ is equivalent to \mathcal{I} .

In other words, \mathcal{I} is Blackwell-Pareto optimal if there is no private private information structure $\hat{\mathcal{I}}$ that gives as much information to each agent (in the Blackwell sense), and gives a strictly Blackwell dominating signal to at least one agent. In the Appendix (Lemma 4), we show that Blackwell-Pareto optimal structures exist, and that, moreover, every private private information structure is weakly dominated by a Blackwell-Pareto optimal one.

As we explain in the introduction, there is some tension between the privacy of an information structure and its informativeness. For example, the most informative structure from the point of view of agent 1 is the one where s_1 completely reveals the state, i.e., $p(s_1) = \delta_\omega$. Likewise, agent 2 would benefit most from a structure where s_2 perfectly reveals the state. But then $p(s_1) = p(s_2)$, and so s_1 and s_2 are not independent. The question is thus: what are the ways to maximally inform the agents, while still maintaining privacy?

3 Blackwell-Pareto Optimality and Conjugate Distributions

The question of Blackwell-Pareto optimality of private private information structures is already non-trivial in the case of two agents and a binary state. For example, is the structure given in Figure 1 Blackwell-Pareto optimal? What about the structure in

Figure 2? In this section, we give a simple description of the Blackwell-Pareto frontier, making it easy to check if a structure is Blackwell-Pareto optimal. In particular, our results imply that the structure in Figure 1 is Blackwell-Pareto optimal while the one in Figure 2 is not.

To state this result, we introduce *conjugate distributions* on $[0, 1]$. Let $F: [0, 1] \rightarrow [0, 1]$ be the cumulative distribution function of a probability measure in $\Delta([0, 1])$. The associated *quantile function*, which we denote by F^{-1} , is given by

$$F^{-1}(x) = \min\{y : F(y) \geq x\}. \quad (1)$$

Since cumulative distribution functions are right-continuous, this minimum indeed exists, and so F^{-1} is well defined. When F is the cumulative distribution function of a full support and nonatomic measure, then F is a bijection and F^{-1} is its inverse. More generally, $F^{-1}(x)$ is the smallest number y such that an x -fraction of the population has values less than or equal to y .

Definition 4. *The conjugate of a cumulative distribution function $F: [0, 1] \rightarrow [0, 1]$ is the function $\hat{F}: [0, 1] \rightarrow [0, 1]$, which is given by*

$$\hat{F}(x) = 1 - F^{-1}(1 - x).$$

Graphically, (x, y) is on the graph of F if and only if $(1 - y, 1 - x)$ is on the graph of \hat{F} : in other words, \hat{F} is the reflection of F around the anti-diagonal of the unit square. An example is depicted in Figure 3.

As we show in the Appendix (Claim 3), \hat{F} is also a cumulative distribution function. Thus, given a measure $\mu \in \Delta([0, 1])$, we can define its conjugate measure $\hat{\mu} \in \Delta([0, 1])$ as the unique measure whose cumulative distribution function is the conjugate of the cumulative distribution function of μ . It is easy to verify that the conjugate of $\hat{\mu}$ is again μ .

The main result of this section is that Blackwell-Pareto optimality can be characterized in terms of conjugates.

Theorem 1. *For a binary state ω and two agents, a private private information structure $\mathcal{I} = (\omega, s_1, s_2)$ is Blackwell-Pareto optimal if and only if the distributions of beliefs $p(s_1)$ and $p(s_2)$ are conjugates.*

Our proof of Theorem 1 combines our more general characterization of Blackwell-Pareto optimality in the n agents case (Theorem 3) together with a classical result of




Figure 3: An example of a cumulative distribution function F and its conjugate \hat{F} . The shapes under the curves are congruent: the transformation that maps one to the other is reflection around the anti-diagonal. Qualitatively, F corresponds to the belief distribution of a more informative signal, and \hat{F} corresponds to that of a less informative signal (because the former assigns less mass to posterior beliefs near 0.5).

Lorentz (1949) about so-called “sets of uniqueness,” which we discuss in detail in §6; these are subsets of $[0, 1]^n$ that are uniquely determined by their projections to each of the n axes.

Figure 3 suggests that on the Blackwell-Pareto frontier, when s_1 is very informative, s_2 must be very uninformative. We formalize this in the following claim, which is proved in the Appendix:

Corollary 1. *Suppose that ω is binary and that (ω, s_1, s_2) and (ω, t_1, t_2) are Blackwell-Pareto optimal private private information structures. If t_1 dominates s_1 , then t_2 is dominated by s_2 .*

Note that for every pair of conjugate distributions μ and $\hat{\mu}$, there exists a private private information structure $\mathcal{I} = (\omega, s_1, s_2)$ where the prior probability of $\omega = 1$ equals the common mean of μ and $\hat{\mu}$, the posterior $p(s_1)$ has the distribution μ , and $p(s_2)$ has the distribution $\hat{\mu}$. By Theorem 1, this structure will be Blackwell-Pareto optimal. To explicitly construct such a structure, calculate the cumulative distribution function F of μ and its conjugate \hat{F} , choose (s_1, s_2) uniformly from the unit square (so that they are independent and each distributed uniformly on $[0, 1]$), and let $\omega = 1$ be the event that $s_2 \geq \hat{F}(1 - s_1)$. A simple calculation shows that




Figure 4: A private private information structure, where the beliefs $p(s_1)$ and $p(s_2)$ are distributed according to the pair of conjugate distributions F and \hat{F} from Figure 3: the signals are uniform on $[0, 1]^2$, and $\omega = 1$ if and only if $s_2 \geq \hat{F}(1 - s_1)$ (black region).

$\hat{F}(1 - s_1)$ is equal to the posterior $p(s_1)$ and has the distribution μ , and $p(s_2)$ has the distribution $\hat{\mu}$. Figure 4 illustrates this construction.

We can use Theorem 1 to understand whether the structures of Figures 1 and 2 are optimal. The uniform distribution on $[0, 1]$ is its own conjugate. Hence, using Theorem 1's belief conjugacy test, we can conclude that Figure 1's information structure is Blackwell-Pareto optimal.

To understand the structure of Figure 2, consider, more generally, a discrete distribution μ on $[0, 1]$ with k atoms. Its conjugate $\hat{\mu}$ is also atomic: each atom of μ with weight w corresponds to an interval of zero mass with length w for $\hat{\mu}$ and, symmetrically, each interval of length l carrying no atoms in μ becomes an atom of weight l in $\hat{\mu}$ (see Figure 5). In particular, $\hat{\mu}$ has either $k - 1$, k or $k + 1$ atoms, corresponding to the cases that (1) μ places positive mass on both 0 and 1, (2) μ places positive mass on exactly one of $\{0, 1\}$, and (3) μ places zero mass on $\{0, 1\}$.

We conclude that the information structure of Figure 2, where both signals induce beliefs $1/4$ or $3/4$ is not Blackwell-Pareto optimal, since two discrete distributions with the same number of atoms can only be conjugates if each of them assigns a non-zero weight to exactly one of $\{0, 1\}$.




Figure 5: The conjugate of a discrete distribution F with three atoms at 0.1, 0.4, and 0.6. Each atom becomes an interval of zero measure with the length equal to the atom’s weight, and vice versa. Since F does not have atoms at the endpoints of $[0, 1]$, the number of intervals of zero measure exceeds the number of atoms by one, so its conjugate \hat{F} has four atoms at 0, 0.5, 0.8, and 1.

4 Optimal Privacy-Preserving Recommendation

The optimal privacy-preserving recommendation is an information-design problem, where an informed party who wishes to disclose to an uninformed party as much information as possible about the state ω using a message s_2 , but must not reveal any information about a given random variable s_1 that is correlated with ω . In this application, we should interpret s_1 not as a signal sent to some agent, but as an additional component of the state that must be kept secret for legal or security reasons, and where the joint distribution of (ω, s_1) is given exogenously.

As a concrete example, suppose an uninformed company wants to learn about a decision-relevant type ω of an applicant (e.g., whether she is a good fit for a job or whether she will pay her rent on time), and an informed party (e.g., a recommender or a credit-rating company) knows both this type and a sensitive or legally protected attribute s_1 of the applicant that correlates with the type: this could be the applicant’s private medical information, gender or race. The informed party faces the problem of optimal privacy-preserving recommendation: convey as much information as possible about the applicant without revealing any information about her protected attribute, so that the company’s downstream decision based solely on the recommendation will be independent of the protected attribute and therefore not cause disparate impact.

Note that even a recommendation that does not explicitly contain the protected attribute may cause disparate impact, if it contains correlates of the attribute.

A less economic (but more colorful) story is that of a government that would like to reveal a piece of intelligence ω , but without revealing any information about the identity of its source s_1 . These could be naturally correlated: for example, if ω is the location of an enemy's weapons facility and the source s_1 is likely to live close to it. So the government's disclosure s_2 should contain as much information as possible about ω , while not revealing any information about s_1 .

Motivated by these examples, we give the following definition.

Definition 5. *Given an information structure (ω, s_1) , a signal s_2 is an optimal privacy-preserving recommendation for (ω, s_1) if there is no s'_2 that is independent of s_1 and Blackwell dominates s_2 .*

Note that in this definition we take an approach that is agnostic to any decision problem faced by the receiver, or indeed any other goal the receiver might have, such as selling this information, using it in a game, etc. Instead, we study signals that are maximal in their information content, and can be used by the receiver towards any of these goals.

The notion of an optimal privacy-preserving recommendation is closely related to Blackwell-Pareto optimality, namely, it is equivalent to requiring that $\mathcal{I} = (\omega, s_1, s_2)$ is a Blackwell-Pareto optimal private information structure. Indeed, if (ω, s_1, s_2) is Blackwell-Pareto optimal, none of the signals can be made more informative under the privacy constraint and so s_2 is an optimal privacy-preserving recommendation. The converse is established in Appendix B.6 (Corollary 5).

When ω and s_1 are correlated, s_2 cannot be a completely revealing signal, as it would provide information about s_1 . A priori, it is not obvious whether there exists a unique solution to the optimal privacy-preserving recommendation problem, or whether there are multiple Blackwell-unordered signals s_2 that are optimal for different decision problems. Our characterization of the Blackwell-Pareto optimal private signals (Theorem 1) implies that the solution is unique and provides a simple recipe to calculate it, when the state is binary.

Theorem 2. *For a binary state ω , there exists an optimal privacy-preserving recommendation s_2^* for every (ω, s_1) . This recommendation is unique up to equivalence: the*

distribution of $p(s_2^*)$ is the conjugate of the distribution of $p(s_1)$. Furthermore, every signal s_2 independent of s_1 is Blackwell dominated by s_2^* .

The last statement in Theorem 2 implies that every decision maker would find the signal s_2^* optimal, regardless of the decision problem at hand. For example, no s_2 that is independent of s_1 can have higher mutual information with ω or lower expected Euclidean distance from ω . This uniqueness of the optimal privacy-preserving recommendation is a rather surprising property as one could expect that, for given (ω, s_1) , there are non-equivalent choices of s_2^* that are both maximal and incomparable in the Blackwell order. In Appendix C, we demonstrate that uniqueness is a feature of the binary state case by considering an example with three states, binary s_1 and a continuum of optimal privacy-preserving recommendations.

Figure 6 shows the optimal privacy-preserving recommendation when the two states are equally likely and s_1 is a symmetric binary signal that matches the state with probability $3/4$. The optimal recommendation s_2^* is trinary: it completely reveals the state with probability $1/2$, and gives no information with the remaining probability. More generally, when the states are equally likely and s_1 is a symmetric binary signal that matches the state with probability $r \in [1/2, 1]$, the optimal recommendation will be trinary. It completely reveals the state with probability $2(1 - r)$, and gives no information with the complementary probability. Thus, as the correlation between s_1 and ω increases, the optimal recommendation becomes less informative.

We provide a simple practical procedure for generating an optimal privacy-preserving recommendation s_2^* , given realizations of (ω, s_1) . We know that s_1 and s_2^* induce conjugate belief distributions, and so we can use the general procedure outlined in Figure 4 to construct s_2^* as follows:

- Calculate $p(s_1)$, the conditional probability of $\omega = 1$ given s_1 .
- If $\omega = 1$, sample s_2^* uniformly from the interval $[1 - p(s_1), 1]$.
- If $\omega = 0$, sample s_2^* uniformly from the interval $[0, 1 - p(s_1)]$.

This procedure yields an s_2^* that, conditioned on s_1 , is distributed uniformly on $[0, 1]$, and hence is independent of s_1 . It is simple to verify that s_2^* is optimal (see the proof of Theorem 2).

This procedure can be simplified if the posterior $p(s_1)$ only takes finitely many values, in which case there exists an optimal privacy-preserving recommendation that




Figure 6: Optimal privacy-preserving recommendation when a $3/4$ -binary signal s_1 must be kept secret. The left panel depicts the cumulative distribution function F of posteriors induced by the symmetric binary signal s_1 matching the state with probability $3/4$. The optimal recommendation s_2^* corresponds to the conjugate distribution \hat{F} depicted in the right panel. We see that s_2^* is trinary: it is completely uninformative with probability $1/2$ and fully reveals the state with the complementary chance, inducing the posteriors 0 or 1 with equal probabilities.

is also finitely valued. Let $[0, 1] = \bigsqcup_{k=0}^K I_k$ be a partition of the unit interval into subintervals using the values of $p(s_1)$. The belief $p(s_2^*)$ is constant when s_2^* ranges within I_k . Hence, the constructed optimal recommendation s_2^* with values in $[0, 1]$ is equivalent to a signal $t_2^* \in \{0, \dots, K\}$ such that $t_2^* = k$ whenever $s_2^* \in I_k$. The signal t_2^* is also an optimal privacy-preserving recommendation and takes at most one more value than the number of values of $p(s_1)$.

Consider the symmetric binary s_1 matching ω with probability $3/4$ from Figure 6. An optimal privacy-preserving recommendation s_2^* of the state can be generated as follows. It takes three values, $\{0, 1, 2\}$. If $s_1 \neq \omega$, then $s_2^* = 2 \cdot \omega$, and if $s_1 = \omega$ then $s_2^* = 1$ with probability $2/3$ and $s_2^* = 2 \cdot \omega$ with probability $1/3$. As a result, the realization $s_2^* = 1$ is completely uninformative and $s_2^* \in \{0, 2\}$ completely reveals ω .

Information Loss in the Optimal Privacy-Preserving Recommendation.

We now discuss how the informativeness of the optimal privacy-preserving recommendation s_2^* depends on the correlation between s_1 and the state ω . As s_1 can represent the collection of all protected attributes, this question is of interest to a policymaker considering the welfare implications of adding an extra protected at-

tribute, thus increasing the correlation.

Corollary 1 implies that informativeness of s_1 and s_2^* move in opposite directions with respect to the Blackwell order. However, this comparative static does not quantify the amount of information lost due to the privacy constraint. Suppose the informativeness of a signal s about binary ω is measured by the mutual information $I(\omega; s) = H(p) - \mathbb{E}[H(p(s))]$, where H is the Shannon entropy defined by $H(q) = q \log_2(q) + (1 - q) \log_2(1 - q)$. The mutual information $I(\omega; s) = 0$ if s is uninformative and $I(\omega; s) = H(p)$ if s pins down the realization of ω . The following inequality provides an upper bound on the informativeness of the optimal privacy-preserving recommendation:

$$I(\omega; s_1) + I(\omega; s_2^*) \leq H(p). \quad (2)$$

The information-theoretic intuition behind (2) is that privacy does not allow s_1 and s_2^* to carry the same “bit of information” about ω , and the “total number of bits” is $H(p)$. The inequality would be less intuitive if we relied on the fairness interpretation of the privacy constraint from §1. A counterintuitive phenomenon is that (2) is always strict whenever $H(p) > I(\omega; s_1) > 0$, i.e., s_1 is informative about ω but does not reveal it completely. This effect is discussed in §7, where we prove a tighter bound in a more general setting and also discuss quadratic loss as an alternative measure of informativeness.

The bound in inequality (2) does not pin down the informativeness level of the optimal privacy-preserving recommendation. We address this question in higher generality. Suppose the optimal recommendation s_2^* is observed by a decision maker, who selects an action $a \in A$ and receives a payoff of $u(\omega, a)$. Such a decision maker may be interested in evaluating how good the signal is in her decision problem. Her expected payoff can be expressed as

$$\mathbb{E} \left[\sup_{\sigma: S_2 \rightarrow A} u(\omega, \sigma(s_2^*)) \right].$$

We assume that u is bounded from above and thus the indirect utility $U(q) = \sup_{a \in A} ((1 - q) \cdot u(a, 0) + q \cdot u(a, 1))$ is a continuous convex function. The following proposition provides an explicit formula for the decision-maker’s expected payoff.

Proposition 1. *Consider the optimal privacy-preserving recommendation s_2^* for (ω, s_1) with binary ω , and let F be the cumulative distribution function of $p(s_1)$. The expected*

payoff of a decision maker with an indirect utility function U observing s_2^* is equal to

$$\int_{[0,1]} U(1 - F(t)) dt. \quad (3)$$

The mutual information discussed above corresponds to a particular decision problem with $U(q) = H(p) - H(q)$.

The proposition follows from Theorem 2. Indeed, the distribution of beliefs induced by the optimal recommendation is the conjugate \hat{F} of F . Thus the expected payoff of the decision maker observing s_2^* is given by

$$\int_{[0,1]} U(q) d\hat{F}(q).$$

Changing the variable $t = \hat{F}(q)$ and using the definition of the conjugate, we get formula (3). Details can be found in Appendix B.7.

Three and More States. Beyond the binary state case, optimal privacy-preserving recommendations are not unique (see Appendix C), and so the choice of the recommendation can depend on the receiver’s decision problem. In this case, we do not have a simple characterization or construction and leave this question as a direction for future research.

5 Welfare-Maximizing Private Private Information Structures

The Blackwell-Pareto frontier characterized in Theorem 1 is agnostic to the designer’s specific objective. Indeed, by maximizing any objective monotone in the informativeness of each agent’s signal, we get a Blackwell-Pareto optimal information structure. Here we refine the characterization of Blackwell-Pareto optimal structures by describing those of them that maximize a welfare objective.

As a stylized example, consider a consulting company advising several clients about the state ω of the market and required—by contract or by law—to avoid giving its customers information about each other. Assuming that the fee paid by a client is proportional to the utility she derives, a revenue-maximizing consulting company faces the problem of choosing a private-private information structure that maximizes clients’ welfare. In contrast to the secret sharing of Shamir (1979), where the goal is to split information in a way that each agent knows as little as possible both about

others and about ω , the goal of the consulting company is to make each individual signal as informative of the ω as possible.

Formally, suppose that each agent $i \in \{1, 2\}$ has to choose an action $a_i \in A_i$ after observing a signal s_i , and receives payoff according to a utility function $u_i(\omega, a_i)$. Our only assumption is that u_i is bounded from above.

For a given binary ω , the social welfare of a given structure (ω, s_1, s_2) is

$$\sum_{i=1,2} \mathbb{E} \left[\sup_{\sigma_i: S_i \rightarrow A_i} u_i(\omega, \sigma_i(s_i)) \right].$$

What are the private private information structures (ω, s_1, s_2) that maximize social welfare? Clearly, every maximizing structure must lie on the Blackwell-Pareto frontier. But while the Blackwell-Pareto frontier contains a rich set of information structures, including some that induce a continuum of beliefs, the ones that maximize social welfare have a simpler form.

Proposition 2. *Given a binary ω , and given u_1 and u_2 , there exists a welfare-maximizing private private information structure (ω, s_1, s_2) such that s_1 takes at most two values, s_2 takes at most three values, and the distributions of beliefs induced by s_1 and s_2 are conjugates.*

By permuting the roles of s_1 and s_2 , we deduce that there is also a welfare-maximizing structure in which s_2 takes two values and s_1 takes three (at most). Curiously, we see that the privacy constraint increases the number of actions in the support from two, as in the single-agent persuasion model of [Kamenica and Gentzkow \(2011\)](#), to possibly three. On the other hand, in multi-agent persuasion there is no limit on the number of actions that may be taken ([Arieli et al., 2021](#)).

The proposition is proved in [Appendix B.8](#) using a combination of an extreme-point argument and the characterization of Blackwell-Pareto optimal structures via conjugate distributions ([Theorem 1](#)).

For an example of a social welfare-maximizing structure, consider the canonical example with two equally likely states, $A_i = \Omega = \{0, 1\}$, where each agent gets utility 1 from matching the state and utility -1 from mismatching it, so that $u_1(\omega, a) = u_2(\omega, a) = 1 - 2|\omega - a|$. If we reveal the state to agent 1 and give agent 2 no information, then the social welfare is 1. Consider instead a private private information structure where each agent has a posterior belief of $\sqrt{1/2}$ with probability $\sqrt{1/2}$




Figure 7: A social welfare-maximizing private private information structure for the decision problem in which $u_1(\omega, a) = u_2(\omega, a) = 1 - 2|\omega - a|$.

and a posterior belief of 0 with the complementary probability. Such a structure exists as this distribution is its own conjugate: see also Figure 7. Then the social welfare is $4 - 2\sqrt{2} \approx 1.17$. Let us check that this is the highest possible social welfare across all private private information structures.

By Proposition 2, we can assume that the distribution of posteriors μ induced by s_1 is supported on two points. It has mean $1/2$ since the average posterior equals the prior, and thus can be represented as $\frac{\alpha}{\alpha+\beta}\delta_{\frac{1}{2}-\beta} + \frac{\beta}{\alpha+\beta}\delta_{\frac{1}{2}+\alpha}$ for some constants $\alpha, \beta \in (0, 1/2]$, where δ_x denotes the point mass at x . The contribution of the first agent to the welfare is, therefore, $\frac{4\alpha\beta}{\alpha+\beta}$.

The conjugate distribution $\hat{\mu}$ takes the form $(\frac{1}{2} - \alpha)\delta_0 + (\alpha + \beta)\delta_{\frac{\beta}{\alpha+\beta}} + (\frac{1}{2} - \beta)\delta_1$. As the problem is state-symmetric, we can assume $\beta \geq \alpha$ without loss of generality and, hence, the middle atom of $\hat{\mu}$ is above $1/2$. Therefore, the second agent contributes $1 - 2\alpha$ to the welfare, and the total welfare equals $\frac{4\alpha\beta}{\alpha+\beta} + 1 - 2\alpha$. A simple calculation shows that this is maximized when $\beta = 1/2$ and $\alpha = \sqrt{1/2} - 1/2$, which yields the structure described above.

6 Blackwell-Pareto Optimality and Sets of Uniqueness

In this section, we study Blackwell-Pareto optimality of private private information in the general setting of n agents and a state ω that takes value in $\Omega = \{0, \dots, m - 1\}$. Our main result shows that Blackwell-Pareto optimality can be characterized using “sets of uniqueness”: subsets of $[0, 1]^n$ that are uniquely determined by their projections to the n axes.

As a first step, we show that it is without loss of generality to focus on information structures that are constructed similarly to the examples we have considered above: each s_i is distributed uniformly on $[0, 1]$, and each value of ω corresponds to some subset of $[0, 1]^n$. That is, ω is a deterministic function of the signals (see Figures 1 and 2, as well as 8 in the Appendix).

More formally, let $\mathcal{A} = (A_0, \dots, A_{m-1})$ be a *partition* of $[0, 1]^n$ into measurable sets. That is, each A_k is a measurable subset of $[0, 1]^n$, the sets in \mathcal{A} are disjoint, and their union is equal to $[0, 1]^n$.

Definition 6. *The private private information structure associated with a partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ is $\mathcal{I} = (\omega, s_1, \dots, s_n)$ where (s_1, \dots, s_n) are distributed uniformly on $[0, 1]^n$ and $\{\omega = k\}$ is the event that $\{(s_1, \dots, s_n) \in A_k\}$.*

Note that if \mathcal{A} and \mathcal{A}' are partitions such that each symmetric difference $A_k \triangle A'_k$ has zero Lebesgue measure, then both are associated with the same information structure, in the strong sense that the joint distributions of $(\omega, s_1, \dots, s_n)$ coincide. Accordingly, we henceforth consider two subsets of $[0, 1]^n$ to be equal if they only differ on a zero-measure set.

Private private information structures associated with partitions have two important properties that make them useful. First, signals are uniform on $[0, 1]$, so that one does not need to consider abstract signal spaces. Second, the collection of signals (s_1, \dots, s_n) reveals the state. The next result shows that—up to equivalence—these assumptions are without loss of generality.

Proposition 3. *For every private private information structure \mathcal{I} , there exists a partition \mathcal{A} whose associated information structure \mathcal{I}' is equivalent to \mathcal{I} .*

While the general proof contained in Appendix B.2 is not constructive, for structures with a finite number of signals and a binary state, we show in Appendix C.1 how to construct a partition with an equivalent associated structure.

The ideas behind the proof of this proposition are the following. Using standard results, one can always reparameterize the signals so that they are uniformly distributed in $[0, 1]$. Thus the main challenge is to ensure that the state is determined by the signals. To this end, given signals that do not determine the state, we add a signal t that resolves the remaining uncertainty, so that ω is a deterministic function of (s_1, \dots, s_n, t) . Then, we use a “secret sharing” technique (similar to Shamir, 1979)

to create a pair of independent random variables t_1, t_2 such that t is determined by the pair (t_1, t_2) , but each t_i is uninformative about the state and the other signals.⁷ We then reveal to agents 1 and 2 the additional signals t_1 and t_2 , respectively. Thus the information structure $(\omega, (s_1, t_1), (s_2, t_2), \dots, s_n)$ is equivalent to $(\omega, s_1, \dots, s_n)$, but now the signals determine the state.

Proposition 3 implies that for the purpose of studying the Blackwell-Pareto optimality of private signals, one can assume without loss of generality that information structures are always associated with partitions. In particular, the question of Blackwell-Pareto optimality can now be phrased as a question about partitions: when does a partition \mathcal{A} correspond to a Blackwell-Pareto optimal structure? Our next result answers this question. We state this result for the case of a binary state, as it involves significantly simpler notation; the result for a general finite state space is given in Appendix B.3. In the case of a binary state, a partition $\mathcal{A} = (A_0, A_1)$ is determined by A_1 , since A_0 is its complement. Hence we will represent \mathcal{A} by a single set $A = A_1$. The information structure associated with A will refer to the structure associated with the partition (A^c, A) .

Given a measurable set $A \subseteq [0, 1]^n$, we define n functions $(\alpha_1^A, \dots, \alpha_n^A)$ that capture the projections of A to the n coordinate axes. Denote by λ the Lebesgue measure on $[0, 1]^{n-1}$. The projection $\alpha_i^A: [0, 1] \rightarrow [0, 1]$ of A to the i th axis is

$$\alpha_i^A(t) = \lambda(\{y_{-i} : (y_i, y_{-i}) \in A, y_i = t\}).$$

If $(\omega, s_1, \dots, s_n)$ is the information structure associated with A , then $\alpha_i^A(t)$ is the posterior of agent i when she observes $s_i = t$.

Definition 7. *A measurable $A \subseteq [0, 1]^n$ is a set of uniqueness if for every measurable $B \subseteq [0, 1]^n$ such that $(\alpha_1^A, \dots, \alpha_n^A) = (\alpha_1^B, \dots, \alpha_n^B)$, it holds that $A = B$.*

Less formally, A is a set of uniqueness if it is determined by the projections $(\alpha_1^A, \dots, \alpha_n^A)$. The main result of this section characterizes Blackwell-Pareto optimality in terms of sets of uniqueness.

Theorem 3. *A private information structure is Blackwell-Pareto optimal if and only if it is equivalent to a structure associated with a set of uniqueness $A \subseteq [0, 1]^n$.*

⁷In the common uses of secret sharing t, t_1, t_2 are binary, while our construction requires that they are distributed uniformly on $[0, 1]$.

To prove that Blackwell-Pareto optimality implies that A is a set of uniqueness, suppose A is not a set of uniqueness, so that $B \neq A$ has the same projections. Hence the structure associated with A is equivalent to the one associated with B . By considering a convex combination of the two structures, we arrive at another equivalent structure, one in which the signals do not always determine the state. We resolve this residual uncertainty via an additional informative signal t , which is independent of (s_1, \dots, s_n) . Now, revealing this new signal to agent 1 results in a private private information structure $(\omega, (s_1, t), s_2, \dots, s_n)$ that dominates the structure associated with A .

Conversely, suppose the information structure associated with A is not Blackwell-Pareto optimal. By considering a Blackwell-Pareto dominating information structure and garbling the signals, we can find a density $f : [0, 1]^n \rightarrow [0, 1]$ that is not an indicator function, but has the same marginals as A . We next apply a result of [Gutmann et al. \(1991\)](#): the set of densities valued in $[0, 1]$ with given marginals is a convex set whose extreme points are indicator functions. Since f is not an indicator function, the corresponding convex set is not a singleton and must have at least two different extreme points. There exists some other set with the same marginals as A , so A is not a set of uniqueness.

Theorem 3 shows an equivalence between the two *a priori* unrelated notions of Blackwell-Pareto optimality and sets of uniqueness; a similar result in Appendix B.3 establishes an analogous equivalence for arbitrary finite state spaces, generalizing sets of uniqueness to *partitions of uniqueness*. This connection lets us use characterization results about sets of uniqueness to study Blackwell-Pareto optimality. Sets of uniqueness have been studied since [Lorentz \(1949\)](#), who gives a simple characterization in the two-dimensional case. A version of his characterization, as we explain below, leads to Theorem 1. Beyond the two-dimensional case, sets of uniqueness have also been more recently studied in the *mathematical tomography* literature (e.g., [Fishburn et al., 1990](#)). We discuss below how these newer results help us understand Blackwell-Pareto optimal structures.

To characterize sets of uniqueness in two dimensions, we will need the following definitions. Say that $A \subseteq [0, 1]^2$ is a *rearrangement* of $B \subseteq [0, 1]^2$ if for $i = 1, 2$ and every $q \in [0, 1]$, the sets $\{t \in [0, 1] : \alpha_i^A(t) \leq q\}$ and $\{t \in [0, 1] : \alpha_i^B(t) \leq q\}$ have the same Lebesgue measure. That is, α_i^A and α_i^B , when viewed as random variables defined on $[0, 1]$, have the same distribution. This has a simple interpretation in

terms of information structures: A is a rearrangement of B if and only if the two associated information structures are Blackwell equivalent. This is immediate, since in the information structure associated with A , $\alpha_i^A(t)$ is the belief of agent i after observing $s_i = t$. Recall that $B \subseteq [0, 1]^n$ is *upward-closed* if $x = (x_1, \dots, x_n) \in B$ implies that $y = (y_1, \dots, y_n) \in B$ for all $y \geq x$.

Theorem 4 (Lorentz (1949)). *A measurable subset $A \subseteq [0, 1]^2$ is a set of uniqueness if and only if it is a rearrangement of an upward-closed set.*

This formulation is different than the one that appears in the paper by Lorentz (1949). We, therefore, show in the Appendix that it is an equivalent characterization. Theorem 1 is a consequence of Theorems 4 and 3.

When $n > 2$, the known characterizations of sets of uniqueness are more involved (Kellerer, 1993). Nevertheless, a simple sufficient condition for uniqueness (Fishburn et al., 1990) is to be an *additive* set: this holds when there are bounded $h_i: [0, 1] \rightarrow \mathbb{R}$ such that $A = \{x \in [0, 1]^n : \sum_{i=1}^n h_i(x_i) \geq 0\}$.

In two dimensions a set is additive if and only if it is a rearrangement of an upward-closed set, and so additivity provides another characterization of the sets of uniqueness (and equivalently, of the Blackwell-Pareto optimal structures). In higher dimensions (i.e., with three or more agents), the sufficiency of additivity implies that every additive set is associated with a Blackwell-Pareto optimal structure. With $n \geq 3$, Kemperman (1991) demonstrated that there are sets of uniqueness that are not additive. However, additivity is “almost necessary”: Kellerer (1993) characterizes the class of sets of uniqueness as the closure, in a certain topology, of the class of additive sets.

In light of these results, progress in understanding Blackwell-Pareto optimal private structures for more than two agents is contingent on new breakthroughs regarding these old questions. A more promising direction for future research is to characterize Blackwell-Pareto optimal private structures for a non-binary state and two agents. As we discuss in Appendix B.3, this will require understanding *partitions of uniqueness* with more than two elements. To the best of our knowledge, this has not been investigated in the literature.

7 Feasibility

In §2, we discussed the fact that a private private information structure $(\omega, s_1, \dots, s_n)$ is equivalent to the “direct revelation” structure $(\omega, p(s_1), \dots, p(s_n))$. Equivalence classes of information structures correspond to n -tuples (μ_1, \dots, μ_n) of measures on $\Delta(\Omega)$, where μ_i is the distribution of $p(s_i)$. In this section, we consider the question of feasibility: which tuples (μ_1, \dots, μ_n) represent some private private information structure?⁸

We use our results from the previous sections to give a simple answer in the two-agent, binary state case. In the general case, we give necessary conditions for feasibility, indicating that we can think of private private structures as a way to divide an “information pie” among the agents. However, there can be negative externalities, so that the sum of the pieces is smaller than the whole pie.

Definition 8. *An n -tuple (μ_1, \dots, μ_n) of probability measures on $\Delta(\Omega)$ is said to be feasible if there exists a private private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ such that μ_i is the distribution of $p(s_i)$.*

For example, Figure 2 shows that for symmetric binary states and two agents, it is feasible for both agents to have binary signals that induce beliefs of either $1/4$ or $3/4$. That is, (μ_1, μ_2) is feasible for

$$\mu_1 = \mu_2 = \frac{1}{2}\delta_{1/4} + \frac{1}{2}\delta_{3/4}.$$

The question of feasibility was studied by [Gutmann et al. \(1991\)](#) and [Arieli et al. \(2021\)](#). The latter investigates feasibility for two agents, symmetric binary states, and focuses on the case of $\mu_1 = \mu_2$. They show that if μ is symmetric to reflection around $1/2$ (i.e., to permuting the states), then (μ, μ) is feasible if and only if μ is a mean-preserving contraction of the uniform distribution on $[0, 1]$. It follows, for example, that $3/4$ is the strongest possible binary signal that two players can have in a symmetric private private information structure with a symmetric binary state. To the best of our knowledge, little is known about feasibility beyond this result.

A necessary condition for feasibility is given by the so-called martingale condition (i.e., by the law of iterated expectations). It implies that if the posterior $p(s_i)$ has

⁸As we explain in [Appendix A](#), the feasibility of belief distributions can be recast as the feasibility of reduced-form mechanisms.

distribution μ_i then the expected posterior $\int q d\mu_i(q)$ must equal to the prior distribution of ω . Thus a necessary condition for feasibility is that $\int q d\mu_i(q) = \int q d\mu_j(q)$ for all agents i and j .

The question of feasibility is closely related to that of Blackwell-Pareto optimality. Indeed, by Lemma 4 from the Appendix, any private private information structure is weakly dominated by Blackwell-Pareto optimal one. Thus (μ_1, \dots, μ_n) is feasible if and only if there exists a Blackwell-Pareto optimal structure represented by some (ν_1, \dots, ν_n) , such that each μ_i is a mean-preserving contraction of ν_i . This holds since mean-preserving contractions of the posterior belief distributions correspond to Blackwell dominance. By Blackwell's Theorem, one can take a structure with posteriors (ν_1, \dots, ν_n) , and apply an independent garbling to each agent's signal to arrive at a structure with posteriors (μ_1, \dots, μ_n) .

This observation, together with Theorem 1, gives the following characterization for the case of a binary state and two agents.

Corollary 2. *For a binary state and two agents, a pair (μ_1, μ_2) of distributions on $[0, 1]$ is feasible if and only if μ_2 is a mean preserving contraction of the conjugate of μ_1 .*

This result generalizes the symmetric case addressed in Proposition 2 of [Arieli et al. \(2021\)](#): in our result, the two states need not be equally likely, the two agents need not have the same belief distribution, and their belief distributions need not be symmetric around $1/2$. It offers a simple tool for checking feasibility. Indeed, by applying a standard characterization of mean-preserving spreads, the pair (μ_1, μ_2) is feasible if and only if they have the same expectation, and the corresponding cumulative distribution functions (F_1, F_2) satisfy

$$\int_y^1 F_2(x) dx \geq \int_y^1 \hat{F}_1(x) dx$$

for every $y \in [0, 1]$.

For $n \geq 2$ agents, the characterization of Blackwell-Pareto optimal private private structures via sets of uniqueness (Theorem 1) implies the following characterization of feasible distributions.

Corollary 3. *For $n \geq 2$ agents, a collection (μ_1, \dots, μ_n) of distributions on $[0, 1]$ is feasible if and only if there exists a set of uniqueness $A \subset [0, 1]^n$ such that each μ_i*

is a mean-preserving contraction of the distribution of agent i 's beliefs in the private information structure associated with A .

This corollary admits a straightforward extension to the general case of $n \geq 2$ agents and $m \geq 2$ states by replacing sets of uniqueness with partitions of uniqueness and relying on Theorem 5 from Appendix B.3 instead of Theorem 1.

Corollary 3, combined with sufficient conditions for a set A to be a set of uniqueness from §3, gives a simple general recipe for constructing feasible distributions. However, it is hard to use Corollary 3 to disprove the feasibility of a given distribution as one needs to check all possible sets of uniqueness A . To this end, we now present a simple necessary condition for feasibility relying on information-theoretic ideas.

We will require two standard definitions. The *Shannon entropy* of a measure $q \in \Delta(\Omega)$ is

$$H(q) = - \sum_{k \in \Omega} q(k) \log_2(q(k)). \quad (4)$$

Given a signal (ω, s_i) , denote the *mutual information* between ω and s_i by

$$I(\omega; s_i) = H(\mathbb{E}[p(s_i)]) - \mathbb{E}[H(p(s_i))]. \quad (5)$$

Note that $I(\omega; s_i)$ can be written in terms of the distribution of posteriors μ_i , and so it is an equivalence invariant:

$$I(\mu_i) = H\left(\int q \, d\mu_i(q)\right) - \int H(q) \, d\mu_i(q). \quad (6)$$

In this expression, the first expectation $\int q \, d\mu_i(q)$ is the prior distribution of ω .

In information theory, entropy is often used to quantify the amount of randomness in a distribution. Mutual information is then the expected reduction in this randomness, and is used as a measure for the amount of information contained in a signal. These notions are also used in economics, e.g., in the rational inattention literature (Sims, 2010; Matějka and McKay, 2015). In our setting, mutual information is useful as it provides the following necessary condition for the feasibility of private information structures.

Proposition 4. *With n agents and m states, the tuple (μ_1, \dots, μ_n) of distributions on $\Delta(\Omega)$ is feasible only if all μ_i have the same expectation $p = \int q \, d\mu_i(q)$ and*

$$\sum_i I(\mu_i) \leq H(p).$$

The fact that $I(\mu_i) \leq H(p)$ follows immediately from the definition of mutual information. For general information structures (e.g., conditionally independent signals), there are no further restrictions on the tuple $(I(\mu_1), \dots, I(\mu_n))$: each can take any value between 0 and $H(p)$. Proposition 4 shows that the situation is different when it comes to private information structures. Here, the sum of mutual information is bounded by the entropy of the prior over ω , so that the entropy of ω behaves like a finite resource that needs to be split between the agents. The proof of this proposition uses standard information-theoretic tools and is contained in Appendix B.10.

Proposition 4 raises a natural question: is this condition tight? That is, does the picture of entropy as a finite resource to be split between the agents tell the whole story, or is there a tighter inequality that relates $\sum_i I(\mu_i)$ and $H(p)$?

Our next proposition shows that Proposition 4 can be strengthened in the case of a binary state.

Proposition 5. *The tuple (μ_1, \dots, μ_n) of distributions on $[0, 1]$ is feasible only if all μ_i have the same expectation $p = \int q d\mu_i(q)$ and*

$$\sum_i I(\mu_i) \leq H(p) - C_p \sum_{i < j} I(\mu_i) I(\mu_j),$$

where $C_p = 2 \ln 2 \cdot p^2(1-p)^2$.

The proposition is proved in Appendix B.11. As far as we know this proposition is a novel information-theoretic inequality, which might have some independent interest. It shows that for a binary state, while entropy is a finite resource, it cannot be fully divided among the agents: the sum of mutual information is strictly less than the entropy of ω (as long as at least two signals are informative). This is a special property of the binary state setting. For example, if ω is uniformly distributed over $\{0, 1\} \times \{0, 1\}$, then the structure in which s_1 is equal to the first coordinate of ω and s_2 is equal to the second satisfies Proposition 4 with equality.

The mutual information I is the expected utility associated with a particular decision problem: one in which the indirect utility is given by a constant minus the entropy. Does an analog of Proposition 4 hold for other decision problems? A positive answer would give additional necessary conditions for feasibility. The next proposition shows that for another natural indirect utility function—a quadratic

one—an analogous statement indeed holds. Curiously, the proof of this proposition is different than that of Proposition 4, and we do not know of a unifying principle that yields both results. Moreover, we do not know of any other indirect utility that yields the same type of result.

In analogy to (4) and (6), for $q \in \Delta(\Omega)$ denote $\bar{H}(q) = \sum_{k \in \Omega} q(k)(1 - q(k))$, and for a measure μ on $\Delta(\Omega)$ define

$$\bar{I}(\mu) = \bar{H} \left(\int q \, d\mu(q) \right) - \int \bar{H}(q) \, d\mu(q).$$

Loosely speaking, for a distribution μ over posterior beliefs, $\bar{I}(\mu)$ is the expected reduction in the variance of the agent’s belief.

Proposition 6. *With n agents and m states, the tuple (μ_1, \dots, μ_n) of distributions on $\Delta(\Omega)$ is feasible only if all μ_i have the same expectation $p = \int q \, d\mu_i(q)$ and $\sum_i \bar{I}(\mu_i) \leq \bar{H}(p)$.*

The proposition is proved in Appendix B.12. While this statement is completely analogous to that of Proposition 4, the proof uses a different technique, exploiting the L^2 orthogonality of independent random variables. Indeed, we do not know of a unifying argument that implies both propositions, and we furthermore do not know of additional decision problems that yield analogous statements. We note that Proposition 6 is a generalization—from the binary state case—of the “concentration of dependence” principle of Mossel et al. (2020). A very similar idea appeared earlier in the economics literature (Al-Najjar and Smorodinsky, 2000) and is standard in the analysis of Boolean functions (see, e.g., Kahn et al., 1988; O’Donnell, 2014).

References

- D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4):218–256, 2022.
- A. Agan and S. Starr. Ban the box, criminal records, and racial discrimination: A field experiment. *The Quarterly Journal of Economics*, 133(1):191–235, 2018.
- N. I. Al-Najjar and R. Smorodinsky. Pivotal players and the characterization of influence. *Journal of Economic Theory*, 92(2):318–342, 2000.

- C. Aliprantis and K. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide (3rd Edition)*. Springer Berlin Heidelberg New York, 2006.
- I. Arieli, Y. Babichenko, F. Sandomirskiy, and O. Tamuz. Feasible joint posterior beliefs. *Journal of Political Economy*, 129(9):2546–2594, 2021.
- A. Aswani and M. Olfat. Optimization hierarchy for fair statistical decision problems. *The Annals of Statistics*, 50(6):3144–3173, 2022.
- S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org, 2019. <http://www.fairmlbook.org>.
- D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- D. Bergemann, B. Brooks, and S. Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017.
- D. Bergemann, A. Bonatti, and T. Gan. The economics of social data. *The RAND Journal of Economics*, 53(2):263–296, 2022.
- D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.
- K. C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica: Journal of the Econometric Society*, pages 1175–1187, 1991.
- B. Brooks and S. Du. Optimal auction design with common values: An informationally robust approach. *Econometrica*, 89(3):1313–1360, 2021.
- K. Burdzy and S. Pal. Contradictory predictions. *arXiv preprint arXiv:1912.00126*, 2019.
- K. Burdzy and J. Pitman. Bounds on the probability of radically different opinions. *Electronic Communications in Probability*, 25, 2020.
- B. Çelen and S. Kariv. Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review*, 94(3):484–498, 2004a.

- B. Çelen and S. Kariv. Observational learning under imperfect information. *Games and Economic Behavior*, 47(1):72–86, 2004b.
- Y.-K. Che, J. Kim, and K. Mierendorff. Generalized reduced-form auctions: A network-flow approach. *Econometrica*, 81(6):2487–2520, 2013.
- Y.-C. Chen, W. He, J. Li, and Y. Sun. Equivalence of stochastic and deterministic mechanisms. *Econometrica*, 87(4):1367–1390, 2019.
- S. Cichomski. Maximal spread of coherent distributions: a geometric and combinatorial perspective. *arXiv preprint arXiv:2007.08022*, 2020.
- S. Cichomski and A. Osekowski. The maximal difference among expert’s opinions. *Electronic Journal of Probability*, 26:1–17, 2021.
- S. Cichomski and F. Petrov. A combinatorial proof of the Burdzy-Pitman conjecture. *arXiv preprint arXiv:2204.07219*, 2022.
- A. Dawid, M. DeGroot, and Mortera. Coherent combination of experts’ opinions. *Test*, 4(2):263–313, 1995.
- C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- R. Eilat, K. Eliaz, and X. Mu. Bayesian privacy. *Theoretical Economics*, 16(4):1557–1603, 2021.
- K. Eliaz, R. Spiegler, and Y. Weiss. Cheating with models. *American Economic Review: Insights*, 2020.
- J. Feigenbaum, L. Fortnow, D. M. Pennock, and R. Sami. Computation in a distributed information market. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 156–165, 2003.

- P. Fishburn, J. Lagarias, J. Reeds, and L. A. Shepp. Sets uniquely determined by projections on axes i. continuous case. *SIAM Journal on Applied Mathematics*, 50(1):288–306, 1990.
- F. Forges. Correlated equilibrium in two-person zero-sum games. *Econometrica*, 58(2):515, 1990.
- D. Gale. What have we learned from social learning? *European Economic Review*, 40(3-5):617–628, 1996.
- R. J. Gardner. *Geometric Tomography*. Cambridge University Press, 1995.
- A. Gershkov, J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi. On the equivalence of Bayesian and dominant strategy implementation. *Econometrica*, 81(1):197–220, 2013.
- J. K. Goeree and A. Kushnir. Reduced form implementation for environments with value interdependencies. *Games and Economic Behavior*, 99:250–256, 2016.
- J. K. Goeree and A. Kushnir. A geometric approach to mechanism design. *Journal of Political Economy Microeconomics*, 1(2):000–000, 2023.
- P. Gopalan, N. Nisan, and T. Roughgarden. Public projects, boolean functions, and the borders of border’s theorem. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):1–21, 2018.
- S. Gutmann, J. Kemperman, J. Reeds, and L. A. Shepp. Existence of probability measures with given marginals. *The Annals of Probability*, pages 1781–1797, 1991.
- S. Hart and P. J. Reny. Implementation of reduced form mechanisms: a simple approach and a new characterization. *Economic Theory Bulletin*, 3(1):1–8, 2015.
- L. Hong and S. Page. Interpreted and generated signals. *Journal of Economic Theory*, 144(5):2174–2196, 2009.
- D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. Quantifying causal influences. *The Annals of Statistics*, 41(5):2324–2358, 2013.

- J. Kahn, G. Kalai, and N. Linial. The influence of variables on boolean functions. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 68–80, 1988.
- E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- H. G. Kellerer. Uniqueness in bounded moment problems. *Transactions of the American Mathematical Society*, 336(2):727–757, 1993.
- J. Kemperman. Sets of uniqueness and systems of inequalities having a unique solution. *Pacific Journal of Mathematics*, 148(2):275–301, 1991.
- J. Kleinberg, S. Mullainathan, and M. Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.
- J. Kleinberg, J. Ludwig, S. Mullainathan, and A. Rambachan. Algorithmic fairness. In *AEA papers and proceedings*, volume 108, pages 22–27, 2018.
- A. Kleiner, B. Moldovanu, and P. Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021.
- X. Lang. Reduced-form budget allocation with multiple public alternatives. *Social Choice and Welfare*, 59(2):335–359, 2022.
- A. Liang and E. Madsen. Data and incentives. *Theoretical Economics*, forthcoming, 2023.
- A. Liang, J. Lu, and X. Mu. Algorithmic design: Fairness versus accuracy. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 58–59, 2022.
- G. Lorentz. A problem of plane measure. *American Journal of Mathematics*, 71(2):417–426, 1949.
- E. Maskin and J. Riley. Optimal auctions with risk averse buyers. *Econometrica: Journal of the Econometric Society*, pages 1473–1518, 1984.
- F. Matějka and A. McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.

- L. Mathevet, J. Perego, and I. Taneva. On information design in games. *Journal of Political Economy*, 2020.
- S. A. Matthews. On the implementability of reduced form auctions. *Econometrica: Journal of the Econometric Society*, pages 1519–1522, 1984.
- N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6): 1–35, 2021.
- E. Mossel, M. Mueller-Frank, A. Sly, and O. Tamuz. Social learning equilibria. *Econometrica*, 88(3):1235–1267, 2020.
- A. Nikzad. Constrained majorization: Applications in mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 330–331, 2022.
- R. O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012.
- J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2009.
- A. Rambachan, J. Kleinberg, S. Mullainathan, and J. Ludwig. An economic approach to regulating algorithms. Technical report, National Bureau of Economic Research, 2020.
- A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- C. Sims. Rational inattention and monetary economics. *Handbook of Monetary Economics*, 3:155–181, 2010.
- I. Taneva. Information design. *American Economic Journal: Microeconomics*, 11(4): 151–85, 2019.
- Q. Valenzuela-Stookey. Greedy allocations and equitable matchings. *arXiv preprint arXiv:2207.11322*, 2022.

- Y. Viossat. Is having a unique equilibrium robust? *Journal of Mathematical Economics*, 44(11):1152–1160, 2008.
- R. V. Vohra. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press, 2011.
- G. Winkler. Extreme points of moment sets. *Mathematics of Operations Research*, 13(4):581–587, 1988.
- C. Zheng. Reduced-form auctions of multiple objects. Technical report, Working paper, 2021.

Appendix

A More Applications of Private Private Signals

Here, we outline three other applications of private private signals. We discuss a connection between the feasibility of reduced-form mechanisms in Bayesian mechanism design and the feasibility of belief distributions, a use of private private signals for information design in zero-sum games, and a link between the informativeness of private private signals and bounds on casual strength in the context of causal inference.

Feasible Reduced-Form Mechanisms as Feasible Belief Distributions. In Bayesian mechanism design, one-agent marginals of a multi-agent mechanism are known as its reduced forms. Pioneered by [Matthews \(1984\)](#), [Maskin and Riley \(1984\)](#), and [Border \(1991\)](#), the reduced-form approach allows the designer, instead of solving a multi-agent problem, to deal with a single-agent problem under additional feasibility constraints.

We consider a general setting where a mechanism specifies a public outcome as a function of agents' independent types. We notice a surprising equivalence between the feasibility of reduced-form mechanisms and the feasibility of private private belief distributions. Thus our results on feasible belief distributions also characterize feasible reduced-form mechanisms. This connection highlights the link between feasible reduced forms and partitions of uniqueness. It also indicates the relevance of information-theory concepts such as Shannon's entropy in the mechanism design context. These insights contribute to a diverse list of known connections of the reduced-form feasibility problem to majorization ([Hart and Reny, 2015](#); [Kleiner et al., 2021](#); [Nikzad, 2022](#)), support functions of convex sets ([Goeree and Kushnir, 2016, 2023](#)), network flows ([Che et al., 2013](#); [Zheng, 2021](#)), and polymatroids ([Vohra, 2011](#); [Valenzuela-Stookey, 2022](#); [Lang, 2022](#)).

Formally, consider a Bayesian mechanism design setting with n agents and a set Ω of m public outcomes. Each agent $i = 1, \dots, n$ has a type $t_i \in T_i$, where the sets of types T_i are arbitrary measurable spaces, and t_i are random draws from $\nu_i \in \Delta(T_i)$, independent across agents. For example, $t_i \in \mathbb{R}_+^m$ may equal i 's value for each of m public projects, and $\omega \in \{1, \dots, m\}$ determines which project gets implemented. Or, $t_i \in \mathbb{R}$ may be the agent's value for a private good, and $\omega \in \{1, \dots, n\}$ specifies who

gets it.

A direct mechanism M selects a randomized outcome for each profile of types,⁹ i.e., $M: T_1 \times \dots \times T_n \rightarrow \Delta(\Omega)$. As each agent i only knows her own type, the relevant information for i about M is captured by the reduced-form mechanism $M_i: T_i \rightarrow \Delta(\Omega)$ obtained by averaging of M over t_{-i} , i.e.,

$$M_i(t_i) = \int_{T_{-i}} M(t_i, t_{-i}) d\nu_{-i}(t_{-i}).$$

The idea dating back to [Maskin and Riley \(1984\)](#) and [Matthews \(1984\)](#) is that constraints of Bayesian incentive compatibility, interim individual rationality, and revenue or welfare objectives can all be reformulated in terms of collection of reduced forms (M_1, \dots, M_n) rather than M . Thus a multi-agent mechanism design problem reduces to a collection of single-agent ones if we know which reduced forms are feasible.

Definition 9. *Single-agent mechanisms (M_1, \dots, M_n) with $M_i: T_i \rightarrow \Delta(\Omega)$ are feasible reduced forms for given type distributions $\nu_i \in \Delta(T_i)$ if there is an n -agent mechanism M such that M_i are its reduced forms.*

The feasibility of reduced-form mechanisms turns out to be tightly related to the feasibility of belief distributions induced by private information structures, as studied in §7. An n -tuple of distributions (μ_1, \dots, μ_n) is *feasible* if there is a private information structure $(\omega, s_1, \dots, s_n)$ such that the belief $p(s_i)$ is distributed according to μ_i for each $i = 1, \dots, n$; see Definition 8.

Claim 1. *Let $M_i: T_i \rightarrow \Delta(\Omega)$, $i = 1, \dots, n$, be a collection of single-agent mechanisms, and each agent's type $t_i \in T_i$ be distributed according to $\nu_i \in \Delta(T_i)$. Denote the distribution of $M_i(t_i)$ by μ_i . Then (M_1, \dots, M_n) are feasible reduced forms if and only if (μ_1, \dots, μ_n) are feasible belief distributions.*

The idea behind the claim is that the agents' types can be interpreted as “signals” about the outcome.

Proof. Each mechanism M defines a joint distribution of the outcome ω and the types (t_1, \dots, t_n) . Interpreting ω as the state and types as independent signals about ω , we

⁹One could augment a mechanism by a vector of transfers, but they are irrelevant to the question of feasibility; see, e.g., [Gopalan et al. \(2018\)](#).

conclude that the posterior belief $p(t_i)$ equals the reduced-form mechanism $M_i(t_i)$. This equivalence completes the proof. \square

This claim allows one to use our feasibility results for belief distributions from §7—criteria from Corollaries 2 and 3 and necessary conditions from Propositions 4, 5, and 6—as results for the feasibility of reduced-form mechanisms.

We note that the general feasibility criteria for public decisions implied by Corollaries 2 and 3 are majorization requirements similar to those obtained by Hart and Reny (2015); Kleiner et al. (2021) for symmetric auctions. The only difference is that, instead of a single dominating distribution, we get a family of such distributions indexed by sets of uniqueness A . The necessity of checking a continuum of dominating distributions may serve as an explanation for the computational intractability of the feasibility problem for public decisions established by Gopalan et al. (2018).

Influencing Competitors in Zero-Sum Games. Suppose agents compete in a zero-sum game, and a designer who knows the state wishes to influence how agents' actions correlate with the state. We show that equilibrium signals must be private. So, our bounds on the informativeness of private signals limit how much the designer can adapt the agents' actions to the state, thus constraining the designer's payoffs.

Consider a zero-sum game played by two players. The action set of player $i \in \{1, 2\}$ is A_i , which we take to be finite, and the utilities are given by $u_1 = -u_2 = u$ for some $u: A_1 \times A_2 \rightarrow \mathbb{R}$. We assume that this game has a unique mixed Nash equilibrium, which holds for generic zero-sum games (Viossat, 2008).

There is a random state ω taking value in Ω . The two players do not know the state and their payoffs do not depend on it. But, there is another agent (the designer) who knows the state and has a utility function $u_d: \Omega \times A_1 \times A_2$ that depends on the state and the actions of the players. This can model a setting where a designer wants to influence the actions of two competitors, with the designer's preference over actions given by her private type ω . The designer commits to a (not necessarily private) information structure (ω, s_1, s_2) . When the state ω is realized, the designer observes it and sends the signal s_1 to player 1 and s_2 to player 2. The players choose their actions after observing the signals.

As a simple example, suppose the game is rock-paper-scissors, so that $A_1 = A_2 =$

$\{R, P, S\}$ and $u(a_1, a_2)$ equals 1 on $\{(P, R), (R, S), (S, P)\}$, zero on the diagonal, and -1 on the remaining action pairs. The state ω takes values in $\{0, 1\}$ and is equal to 1 with probability $1/2$. The designer gets a payoff of 1 for each player who chooses scissors in the high state or chooses rock in the low state.

A pure strategy of player i is a map $f_i: S_i \rightarrow A_i$, and a mixed strategy σ_i is a random pure strategy. An equilibrium consists of an information structure together with a strategy profile (σ_1, σ_2) such that each agent maximizes her expected utility given her signal. That is, for every $s_i \in S_i$ and $a_i \in A_i$

$$\mathbb{E}[u_i(\sigma_i(s_i), \sigma_{-i}(s_{-i}))|s_i] \geq \mathbb{E}[u_i(a_i, \sigma_{-i}(s_{-i}))|s_i].$$

This is just the incentive compatibility condition of a correlated equilibrium, and so, by a direct revelation argument, we can assume that $S_i = A_i$ and that σ_i is always the identity: in equilibrium, the designer recommends an action to each agent, and the agents follow the recommendations. We refer to such equilibria as direct-revelation equilibria.

The next claim shows that private private information structures arise endogenously in this setting.

Claim 2. *In every direct-revelation equilibrium, the information structure (ω, s_1, s_2) is a private private information structure.*

Proof. A zero-sum game with a unique Nash equilibrium has a unique correlated equilibrium which is equal to that Nash equilibrium (Forges, 1990). Thus (s_1, s_2) form a Nash equilibrium, and in particular s_1 must be independent of s_2 . \square

The intuition behind this result is simple: revealing to player i any information about the recommendation given to player $-i$ gives i an advantage that she can exploit to increase her expected utility beyond the value of the game. But player $-i$ can guarantee that i does not get more than the value, and hence s_i cannot contain any information about s_{-i} . Note that Claim 2 applies beyond generic zero-sum games to any game with any number of players, provided that it has a unique correlated equilibrium.¹⁰

¹⁰The set of games with a unique correlated equilibrium is open (Viossat, 2008), so a small enough perturbation of (for example) the rock-paper-scissors game will still have a unique correlated equilibrium, although it will not be zero-sum. As a side note, we are unaware of interesting examples

In the rock-paper-scissors example above, the joint distribution of (s_1, s_2) must be uniform over $\{R, P, S\} \times \{R, P, S\}$, by Claim 2. However, the designer is free to choose the joint distribution between (s_1, s_2) and ω . Thus her problem is to maximize $\mathbb{E}[u_d(\omega, s_1, s_2)]$ over all structures in which (s_1, s_2) is uniform over $\{R, P, S\} \times \{R, P, S\}$. Choosing (s_1, s_2) independently of ω yields a payoff of $6/9$. A straightforward calculation shows that an optimal structure yields her a payoff of $10/9$. By comparison, in a relaxed problem where the designer is allowed to dictate the players' actions without worrying about the privacy constraint, she can achieve utility 2 by revealing the state to both players, telling them to both choose scissors when the state is high and rock when the state is low.

Beyond the specifics of the rock-paper-scissors example, the fact that equilibrium signals are private private means that any bound on the informativeness of private private signals yields a bound on the designer's equilibrium utility: if the designer's recommendations only contain a limited amount of information about the state, then she cannot hope that the players' actions efficiently adapt to the state and yield her high utility. Thus our results, including Theorem 1 and Propositions 4, 5, and 6, constrain what can be achieved by the designer in any such setting.

Bounds on Causal Strength. In causal inference, a *collider* is a causal structure where a number of independent random variables together determine a state (see, e.g., Pearl, 2009), as in the majority signal example above. Our feasibility results from §7 such as Propositions 4, 5, and 6 restrict the informativeness of private private signals. Interpreting causes in collider structures as such signals, one gets bounds on the causal strengths of the various causes (Janzing et al., 2013).

Consider, for example, two “causes” X_1, X_2 —for example, nature (DNA) and nurture (socioeconomic status)—that produce an effect Y —say the winning of an Olympic medal. Assuming that X_1 and X_2 are independent, (Y, X_1, X_2) is a private private structure. Hence it is impossible that Y is very statistically dependent on both X_1 and X_2 . The trade-off between the dependence of Y on X_1 and the dependence of Y on X_2 is captured by the Blackwell-Pareto frontier.

of three-player games with a unique mixed correlated equilibrium. In particular, the following question is open, to the best of our knowledge: does there exist a three-player game with a unique correlated equilibrium in which no player plays a pure strategy?

B Omitted Proofs

Note that we sometimes prove results in a different order than the order that they appear in the main text since some of the results we state earlier are implied by some of the later results.

B.1 Preliminary Lemmas

Let $\mathcal{I} = (\omega, s_1, \dots, s_n)$ be a private private information structure. The signals s_1, \dots, s_n can be combined into a new signal $s = (s_1, \dots, s_n)$. The following lemma gives a lower bound on the informativeness of the combined signal s in terms of the informativeness of the individual signals. It can be seen as superadditivity of mutual information for independent signals. Recall that the mutual information $I(\omega; s)$ is defined in (5).

Lemma 1. *For a private private information structure $(\omega, s_1, \dots, s_n)$ the following inequality holds*

$$\sum_{i=1}^n I(\omega; s_i) \leq I(\omega; (s_1, \dots, s_n)). \quad (7)$$

Proof. The result for $n \geq 3$ follows from the result for $n = 2$ by applying it sequentially to (s_1, \dots, s_k) for $k \leq n$. Consequently, in the rest of the proof we assume $n = 2$.

Our goal is to show that $\Delta = I(\omega; (s_1, s_2)) - I(\omega; s_1) - I(\omega; s_2) \geq 0$. Let $p_1(k) = p(s_1)(k) = \mathbb{P}[\omega = k | s_1]$, define p_2 likewise, and let $p_{12}(k) = p(s_1, s_2)(k) = \mathbb{P}[\omega = k | s_1, s_2]$. Let p denote the prior distribution of ω . By the martingale property $\mathbb{E}[p_{12} | p_i] = p_i$ and $\mathbb{E}[p_i] = p$.

Using this notation and the definition of mutual information, we can write for $i \in \{1, 2\}$

$$I(\omega; s_i) = \mathbb{E} \left[\sum_k p_i(k) \log \frac{p_i(k)}{p(k)} \right] \quad \text{and} \quad I(\omega; s_1, s_2) = \mathbb{E} \left[\sum_k p_{12}(k) \log \frac{p_{12}(k)}{p(k)} \right].$$

By the martingale property we can replace the first p_i by p_{12} :

$$I(\omega; s_i) = \mathbb{E} \left[\sum_k p_{12}(k) \log \frac{p_i(k)}{p(k)} \right].$$

Thus

$$\Delta = -\mathbb{E} \left[\sum_k p_{12}(k) \log \frac{p_1(k)p_2(k)}{p_{12}(k)p(k)} \right].$$

Applying Jensen's inequality to the logarithm, we get that

$$\Delta \geq -\log \mathbb{E} \left[\sum_k p_{12}(k) \frac{p_1(k)p_2(k)}{p_{12}(k)p(k)} \right].$$

By cancelling and rearranging, we get

$$\Delta \geq -\log \sum_k \frac{1}{p(k)} \mathbb{E}[p_1(k)p_2(k)].$$

Since $p_1(k)$ and $p_2(k)$ are independent,

$$\Delta \geq -\log \sum_k \frac{1}{p(k)} \mathbb{E}[p_1(k)] \mathbb{E}[p_2(k)].$$

By the martingale property $\mathbb{E}[p_i(k)] = p(k)$, and so $\Delta \geq -\log \sum_k p(k) = 0$. \square

Note that this proof only used the independence of (s_1, s_2) to the extent that it implies that $p(s_1)$ is uncorrelated with $p(s_2)$.

To show that a given private private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is Blackwell-Pareto dominated, we will often use the following technique: construct an additional informative signal t independent of s_1, \dots, s_n , and reveal it to one of the agents, say, the first one. The new information structure $\mathcal{I}' = (\omega, (s_1, t), s_2, \dots, s_n)$ strictly dominates \mathcal{I} thanks to the following direct corollary of Lemma 1.

Corollary 4. *Fix ω , and consider a pair of signals s and t such that*

- *s and t are independent, and*
- *t is not independent of ω .*

Then the information structure $(\omega, (s, t))$ strictly dominates (ω, s) with respect to the Blackwell order.

Proof. Clearly (ω, s) is weakly dominated by $(\omega, (s, t))$. We show that this domination is strict.

Since t is informative, $I(\omega; t) > 0$. Hence, by Lemma 1, $I(\omega; (s, t)) \geq I(\omega; s) + I(\omega; t) > I(\omega; s)$. Since $I(\omega; s)$ is the value of a particular decision problem (where the indirect utility is given by a constant minus the entropy), it follows that $(\omega, (s, t))$ strictly dominates (ω, s) . \square

The next lemma shows that, without loss of generality, induced posteriors are equal to signals, which can be seen as a version of the revelation principle for private information structures.

Lemma 2. *Any private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is equivalent to $\mathcal{J} = (\omega, t_1, \dots, t_n)$ where each agent's signal t_i is her posterior $p(s_i)$ in the structure \mathcal{I} .*

Proof. By the law of total expectation, $p(t_i) = t_i$. It follows that $p(s_i)$ and $p(t_i)$ have the same distribution, and so are Blackwell equivalent. \square

For a private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$, recall that we denote by $\mu_i \in \Delta(\Delta(\Omega))$ the distribution of the belief $p(s_i)$. Let $\mathcal{M} \subset \Delta(\Delta(\Omega))^n$ be the set of feasible distributions μ_1, \dots, μ_n , i.e., those that correspond to some private information structure \mathcal{I} .

Lemma 3. *The set of feasible distributions \mathcal{M} is compact in the topology of weak convergence.*

Proof. Since the set of probability measures $\Delta(\Delta(\Omega))$ is compact, to prove the compactness of \mathcal{M} , it is enough to check that it is closed. In other words, we need to check that if a sequence of feasible distributions $(\mu_1^l, \dots, \mu_n^l)$ weakly converges to $(\mu_1^\infty, \dots, \mu_n^\infty)$ as $l \rightarrow \infty$, then the limit is also feasible.

Let $\mathcal{I}^l = (\omega, s_1^l, \dots, s_n^l)$ be an information structure inducing $(\mu_1^l, \dots, \mu_n^l)$. By Lemma 2, we can assume without loss of generality that the signals s_i^l are in $\Delta(\Omega)$ and they coincide with the induced beliefs, i.e., $p(s_i^l) = s_i^l$. Let $\psi^l \in \Delta(\Omega \times \Delta(\Omega)^n)$ be the joint distribution of ω and the beliefs s_1^l, \dots, s_n^l . By compactness of the set of probability measures, we can extract a subsequence of ψ^l converging to some ψ^∞ . By definition, the marginal of ψ^∞ on the belief coordinates equals $\mu_1^\infty \times \dots \times \mu_n^\infty$.

Consider a private private information structure $\mathcal{I}^\infty = (\omega, s_1^\infty, \dots, s_n^\infty)$, where signals s_i^∞ belong to $\Delta(\Omega)$ and the joint distribution of the state and signals is given by ψ^∞ . Each signal s_i^∞ has distribution μ_i^∞ . Let us check that the induced beliefs coincide with signals, i.e., $p(s_i^\infty)(k) = s_i^\infty(k)$ almost surely for each $k \in \Omega$. We verify an equivalent integrated version of this identity:

$$\int \left(\sum_k h(k, s_i^\infty) p(s_i^\infty)(k) \right) d\psi^\infty = \int \left(\sum_k h(k, s_i^\infty) s_i^\infty(k) \right) d\psi^\infty \quad (8)$$

for any continuous function h on $\Omega \times \Delta(\Omega)$. Since the left-hand side is just the integral of h , this is equivalent to

$$\int h(\omega, s_i^\infty) d\psi^\infty = \int \left(\sum_k h(k, s_i^\infty) s_i^\infty(k) \right) d\psi^\infty. \quad (9)$$

For each $l < \infty$, the beliefs in \mathcal{I}^l coincide with the signals, i.e.,

$$\int h(\omega, s_i^l) d\psi^l = \int \left(\sum_k h(k, s_i^l) s_i^l(k) \right) d\psi^l.$$

As integration of a continuous function commutes with taking weak limits, letting l go to infinity, we obtain (9).

We conclude that each belief $p(s_i^\infty)$ in \mathcal{I}^∞ coincides with the signal s_i^∞ and the latter is distributed according to μ_i^∞ . Therefore, $(\mu_1^\infty, \dots, \mu_n^\infty)$ is feasible and so the set of feasible distributions is closed and thus compact. \square

The next lemma shows that our order on private private information structures is well-behaved, in the sense that each structure is dominated by a Blackwell-Pareto optimal one: each element of the partially ordered set of private private information structures is upper bounded by a maximal element.

Lemma 4. *For any private private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$, there exists a Blackwell-Pareto optimal structure $\mathcal{I}' = (\omega, s'_1, \dots, s'_n)$ that weakly dominates \mathcal{I} .*

Proof. Recall that $\mathcal{I} \leq \mathcal{J}$ if for any continuous convex $\varphi: \Delta(\Omega) \rightarrow \mathbb{R}$ and any agent $i = 1, \dots, n$,

$$\int \varphi(q) d\mu_i(q) \leq \int \varphi(q) d\nu_i(q) \quad (10)$$

where μ_i and ν_i are the distributions of agent i 's beliefs in \mathcal{I} and \mathcal{J} , respectively.

We say that the collection of distributions (μ_1, \dots, μ_n) is dominated by (ν_1, \dots, ν_n) if (10) holds. Hence, \mathcal{J} dominates \mathcal{I} if and only if the distributions of beliefs in \mathcal{J} dominate those in \mathcal{I} . If $\mu_i^l \rightarrow \mu_i^\infty$ and $\nu_i^l \rightarrow \nu_i^\infty$ weakly as $l \rightarrow \infty$ and $(\nu_1^l, \dots, \nu_n^l)$ dominates $(\mu_1^l, \dots, \mu_n^l)$, then $(\nu_1^\infty, \dots, \nu_n^\infty)$ dominates $(\mu_1^\infty, \dots, \mu_n^\infty)$ as integration of a continuous function φ in (10) is exchangeable with taking weak limits. Thus the dominance order on distributions is continuous in the weak topology.

Let (μ_1, \dots, μ_n) be the distributions of posteriors induced by \mathcal{I} and let \mathcal{M} be the set of feasible distributions endowed with the weak topology. As \mathcal{M} is compact by Lemma 3 and the dominance order is continuous, there is a maximal element $(\nu_1, \dots, \nu_n) \in \mathcal{M}$ dominating (μ_1, \dots, μ_n) . Since (ν_1, \dots, ν_n) is feasible, it is induced by some private private information structure \mathcal{I}' . By the construction, \mathcal{I}' dominates \mathcal{I} and is Blackwell-Pareto optimal. \square

B.2 Proof of Proposition 3

We need to show that, given a private private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ with $\Omega = \{0, \dots, m-1\}$, there is an equivalent structure associated with a partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ of $[0, 1]^n$. The construction relies on two lemmas. Lemma 5 shows that assuming signals s_i are uniform on $[0, 1]$ is without loss of generality. Hence it remains to show that there is an equivalent information structure where signal realizations determine the state. This is done using a secret-sharing scheme from Lemma 6.

Lemma 5. *For any private private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$, there is an equivalent private private information structure $\mathcal{I}' = (\omega, s'_1, \dots, s'_n)$ such that each s'_i is uniformly distributed on $[0, 1]$.*

Proof. Consider the information structure $\mathcal{J} = (\omega, t_1, \dots, t_n)$ where $t_i = (s_i, r_i)$, and each r_i is independent and uniformly distributed on $[0, 1]$. Clearly, \mathcal{I} and \mathcal{J} are equivalent. As t_i is nonatomic, and since all standard nonatomic probability spaces are isomorphic, t_i can be reparametrized to be uniform on $[0, 1]$. \square

We say that a signal t is *split* into r_1 and r_2 if t is a function of r_1 and r_2 , i.e., $t = f(r_1, r_2)$.

Lemma 6. *A signal t distributed uniformly on $[0, 1]$ can be split into r_1 and r_2 such that each r_i is uniformly distributed on $[0, 1]$ and the three random variables t , r_1 , and r_2 are pairwise independent. Furthermore, if t' is an additional signal that is independent of t , then we can take the pair (r_1, r_2) to be independent of t' .*

This lemma extends the classic secret sharing idea from cryptography which applies to discrete random variables. The proof, by construction, is immediate.¹¹

Proof. Denote by $[x]$ the fractional part of $x \in \mathbb{R}$. Take r_1 independent of both t and t' and distributed uniformly on $[0, 1]$, and let $r_2 = [r_1 + t]$. Then $t = [r_2 - r_1]$ and r_1 , r_2 , and t are easily seen to be pairwise independent and also independent of t' altogether. \square

With the help of Lemmas 5 and 6, we are ready to prove Proposition 3.

Proof of Proposition 3. We are given a private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ with sets of signal realizations S_i , $i = 1, \dots, n$. We aim to construct an equivalent one, \mathcal{I}' , where each signal s'_i is uniformly distributed on $[0, 1]$ and the realization of signals (s'_1, \dots, s'_n) determines the state or, equivalently, \mathcal{I}' is associated with some partition $\mathcal{A} = (A_1, \dots, A_{m-1})$ of $[0, 1]^n$.

By Lemma 5, we can find a private information structure $(\omega, t_1, \dots, t_n)$ equivalent to \mathcal{I} where each t_i is uniformly distributed in $[0, 1]$. If the signals (t_1, \dots, t_n) determine the state, then the proof is completed.

Consider the case where (t_1, \dots, t_n) do not determine the state ω . To capture the uncertainty in ω remaining after the signals have been realized, we construct a new signal t as follows.

Let $q : [0, 1]^n \rightarrow \Delta(\Omega)$ be a conditional distribution of ω given all the signals, i.e., $q(x_1, \dots, x_n)(k) = \mathbb{P}[\omega = k | t_1 = x_1, \dots, t_n = x_n]$ for any $k \in \Omega$. With each distribution $q \in \Delta(\Omega)$ we associate a partition of $[0, 1]$ into m intervals

$$B_k(q) = \left[\sum_{l=0}^{k-1} q(l), \sum_{l=0}^k q(l) \right), \quad k = 0, \dots, m-1.$$

The length of $B_k(q)$ equals the mass assigned by q to $\omega = k$. Let t be a random variable uniformly distributed on $[0, 1]$ and independent of (t_1, \dots, t_n) . Consider a new state variable $\omega' \in \Omega$ such that $\omega' = k$ whenever $t \in B_k(q(t_1, \dots, t_n))$. By

¹¹We are thankful to Tristan Tomala for suggesting this construction.

definition, the joint distributions of $(\omega, t_1, \dots, t_n)$ and $(\omega', t_1, \dots, t_n)$ coincide and, therefore, the two structures are equivalent.

The new state ω' is determined by the realizations of t_1, \dots, t_n and the new signal t . Using Lemma 6, we split the signal t into r_1 and r_2 that are independent of each other, and where each r_i is independent of t . Note that, by this lemma, we can take r_1 and r_2 to be independent of (t_1, \dots, t_n) . Since each r_i is uninformative of t , the structure $(\omega', (t_1, r_1), (t_2, r_2), t_3, \dots, t_n)$ where r_1 is revealed to the first agent and r_2 to the second one is a private private structure and is equivalent to \mathcal{I} . Since t is a function of r_1 and r_2 , the signals $(t_1, r_1), (t_2, r_2), t_3, \dots, t_n$ determine the state.

It remains to reparameterize the signals of the first two agents so that, instead of being uniform on $[0, 1]^2$, they become uniform on $[0, 1]$. Consider any bijection $h : [0, 1]^2 \rightarrow [0, 1]$ preserving the Lebesgue measure; such a bijection exists since both are standard nonatomic spaces. and define $s'_1 = h(t_1, r_1)$, $s'_2 = h(t_2, r_2)$, and $s'_i = t_i$ for $i = 3, \dots, n$. The private private information structure $\mathcal{I}' = (\omega', s'_1, s'_2, \dots, s'_n)$ is equivalent to \mathcal{I} , all the signals are uniform on $[0, 1]$, and the realization of signals determines ω' . \square

B.3 Proof of Theorem 3

We formulate and prove an extension of Theorem 3 applicable to non-binary sets of states $\Omega = \{0, 1, \dots, m-1\}$.

Consider a partition of $[0, 1]^n$ into m measurable sets $\mathcal{A} = (A_0, \dots, A_{m-1})$. Recall that the structure $\mathcal{I} = (\omega, s_1, \dots, s_m)$ is said to be associated with a partition \mathcal{A} if all the signals are uniform on $[0, 1]$ and $\omega = k$ whenever $(s_1, \dots, s_n) \in A_k$.

We say that two partitions $\mathcal{A} = (A_0, \dots, A_{m-1})$ and $\mathcal{B} = (B_0, \dots, B_{m-1})$ are *equal* if A_k and B_k differ by a set of zero Lebesgue measure for each k . Recall that the projection of a measurable set $A \subseteq [0, 1]^n$ on the i th coordinate is denoted by α_i^A (see §6). The notion of sets of uniqueness from §6 extends to partitions as follows.

Definition 10. *A partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ is a partition of uniqueness if for any partition $\mathcal{B} = (B_0, \dots, B_{m-1})$ such that $\alpha_i^{A_k} = \alpha_i^{B_k}$ for all i and k , it holds that $\mathcal{A} = \mathcal{B}$.*

Theorem 5 (Extension of Theorem 3 to m states). *A private private information structure \mathcal{I} is Blackwell-Pareto optimal if and only if it is equivalent to a structure associated with a partition of uniqueness \mathcal{A} .*

Note that in the case of $m = 2$ states, a set A_1 in a partition $\mathcal{A} = (A_0, A_1)$ determines $A_0 = [0, 1]^n \setminus A_1$. Hence, $\mathcal{A} = (A_0, A_1)$ is a partition of uniqueness if and only if A_1 is a set of uniqueness. Hence, Theorem 3 is an immediate corollary of its extended version. For an application of the theorem for $m > 2$, see the example contained in Appendix C. This example also demonstrates that partitions of uniqueness are not necessarily composed of sets of uniqueness for $m > 2$ and, hence, the requirement of a partition to be a partition of uniqueness does not boil down to restrictions on its elements unless $m = 2$.

The proof of the theorem is split into a sequence of lemmas. We say that a private information structure is *perfect* if the information received by all the agents together is enough to deduce the realization of ω , i.e., there exists a function $f : S_1 \times \dots \times S_n \rightarrow \Omega$ such that $\omega = f(s_1, \dots, s_n)$. In particular, a structure with signals uniform in $[0, 1]$ is associated with some partition if and only if it is perfect.

The next lemma shows that perfection is necessary for Blackwell-Pareto optimality.

Lemma 7. *If a private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is equivalent to a structure that is not perfect, then \mathcal{I} is not Blackwell-Pareto optimal.*

The construction of the Blackwell-Pareto improvement resembles the proof of Proposition 3 except for the fact that the newly constructed signal is revealed entirely to one of the agents, thus strictly improving her information (in the Blackwell order), by Corollary 4.

Proof. Without loss of generality, \mathcal{I} itself is imperfect. Let $q : S_1 \times \dots \times S_n \rightarrow \Delta(\Omega)$ be the distribution of ω conditional on $s_1 = x_1, \dots, s_n = x_n$, i.e., $q(x_1, \dots, x_n)(k) = \mathbb{P}[\omega = k \mid s_1 = x_1, \dots, s_n = x_n]$, $k = 0, \dots, m - 1$. Since \mathcal{I} is not perfect, we can find a state $k_0 \in \Omega$ such that the event $\{\omega = k_0\}$ is not always determined by the signals. That is, the random variable $q(s_1, \dots, s_n)(k_0)$ does not always take values in $\{0, 1\}$. Without loss of generality, we assume that $k_0 = 0$. With each $q \in \Delta(\Omega)$ we associate a partition of $[0, 1] = \bigsqcup_{k \in \Omega} B_k(q)$, where

$$B_k(q) = [q(\{0, \dots, k - 1\}), q(\{0, \dots, k\})], \quad k = 0, \dots, m - 1$$

so that the length of $B_k(q)$ equals $q(k)$.

We construct a new equivalent structure with an extra signal t as in the proof of Proposition 3. Let t be a random variable uniformly distributed on $[0, 1]$ and independent of s_1, \dots, s_n . Define a new state ω' as a function of these variables in the following way: $\omega' = k$ whenever $t \in B_k(q(s_1, \dots, s_n))$. The joint distribution of $(\omega', s_1, \dots, s_n)$ coincides with that of $(\omega, s_1, \dots, s_n)$ and, hence, the two structures are equivalent.

To get a Blackwell-Pareto improvement, we reveal the realization of t to the first agent and obtain a private private information structure $\mathcal{I}' = (\omega', (s_1, t), s_2, \dots, s_n)$. To argue that \mathcal{I}' is indeed a Blackwell-Pareto improvement we need to show that t itself is an informative signal about ω' , i.e., the posterior $p(t) \in \Delta(\Omega)$ is not equal to the prior p with a positive probability. It is enough to show $p(t)(k_0)$ takes different values for t in $[0, \varepsilon]$ and in $[1 - \varepsilon, 1]$. As we assume without loss of generality that $k_0 = 0$, the interval B_{k_0} is the leftmost one in the partition, and so $p(t)(k_0) = p(t)(0) = \mathbb{P}[t < q(s_1, \dots, s_n)(0) | t]$. That is, if we denote by Q the cumulative distribution function of $q(s_1, \dots, s_n)(0)$, then $p(t)(k_0) = 1 - G(t)$. Since G is a non-constant function on $(0, 1)$ by our assumption on k_0 , the induced belief $p(t)(k_0)$ is not a constant. Thus t is informative. By Corollary 4, this implies that the signal (s_1, t) which the first agent receives in \mathcal{I}' strictly dominates the signal s_1 received in \mathcal{I} . As the signals of all other agents are the same in the two structures, \mathcal{I}' strictly Blackwell-Pareto dominates \mathcal{I} . \square

The next step is to show that only structures corresponding to partitions of uniqueness can be Blackwell-Pareto optimal.

Lemma 8. *If a private private information structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is Blackwell-Pareto optimal, then \mathcal{I} is equivalent to a structure associated with a partition of uniqueness.*

Proof. By Proposition 3, we can find a private private information structure $\mathcal{J} = (\omega, t_1, \dots, t_n)$ equivalent to \mathcal{I} and associated with some partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ of $[0, 1]^n$. Let us demonstrate that \mathcal{A} is a partition of uniqueness. Towards a contradiction, assume that there is another partition $\mathcal{A}' = (A'_0, \dots, A'_{m-1})$ not equal to \mathcal{A} but such that the projections $\alpha_i^{A^k} = \alpha_i^{A'_k}$ for all i and k . So \mathcal{I} is also equivalent to the structure $\mathcal{J}' = (\omega, t'_1, \dots, t'_n)$ associated with \mathcal{A}' .

By Lemma 7, to get a contradiction, it is enough to construct an information structure \mathcal{I}' that is equivalent to \mathcal{I} but not perfect, as this would imply the existence

of a strict Blackwell-Pareto improvement. We define \mathcal{I}' as a structure where the joint distribution of the state and signals is a convex combination of the corresponding distributions in \mathcal{I} and \mathcal{I}' . Formally, let s'_1, \dots, s'_n be independent random variables each uniformly distributed on $[0, 1]$ and let $\theta \in \{0, 1\}$ be a symmetric Bernoulli random variable independent of (s'_1, \dots, s'_n) . Define the state ω' as follows:

$$\omega' = k \quad \text{if} \quad \begin{cases} (s'_1, \dots, s'_n) \in A_k \text{ and } \theta = 0 \\ (s'_1, \dots, s'_n) \in A'_k \text{ and } \theta = 1 \end{cases}.$$

Since elements of the partitions \mathcal{A} and \mathcal{A}' have the same projections, the posterior induced by observing $t_i = x$ in \mathcal{I} is identical to the one induced by observing $t'_i = x$ in \mathcal{I}' . Hence it is again identical to the posterior induced by observing $s'_i = x$ in \mathcal{I}' . As the partitions \mathcal{A} and \mathcal{A}' are not equal, there are $k \neq k'$ such that the intersection $A_k \cap A'_{k'}$ has a non-zero Lebesgue measure. Hence, if $(s'_1, \dots, s'_n) \in A_k \cap A'_{k'}$, whether $\omega = k$ or $\omega = k'$ is determined by θ . We conclude that, with positive probability, the signals (s'_1, \dots, s'_n) do not determine the state, so \mathcal{I}' is not perfect and thus both \mathcal{I}' and \mathcal{I} can be Blackwell-Pareto improved by Lemma 7. \square

We see that Blackwell-Pareto optimal structures are contained in those associated with partitions of uniqueness (up to equivalence of information structures). This shows one direction of Theorem 5. It remains to demonstrate that any partition of uniqueness leads to a Blackwell-Pareto optimal structure, i.e., the structures associated with different partitions cannot dominate each other.

For this purpose, we need two intermediate steps contained in the next two lemmas. Lemma 9 shows that a garbling of an information structure is never perfect and Lemma 10 implies that imperfect structures cannot be equivalent to those associated with partitions of uniqueness. Recall that for a pair of information structures (ω, t) and (ω, s) , the signal t is a garbling of s if t , and ω are independent conditional on s . A structure $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is a garbling of $\mathcal{I}' = (\omega, s'_1, \dots, s'_n)$ if each s_i is a garbling of s'_i and each s_i is independent of $(s'_j)_{j \neq i}$. The last requirement means that each agent's signal is garbled independently. Note that, by Blackwell's Theorem (Blackwell, 1951, Theorem 12), \mathcal{I}' (weakly) dominates \mathcal{I} if and only if \mathcal{I} is equivalent to a garbling of \mathcal{I}' .

Lemma 9. *If \mathcal{I} is a garbling of a private private information structure \mathcal{I}' , then \mathcal{I} is not perfect unless \mathcal{I} and \mathcal{I}' are equivalent.*

Proof. Suppose that \mathcal{I} is perfect, and so $\omega = f(s_1, \dots, s_n)$ for some $f: S_1 \times \dots \times S_n \rightarrow \Omega$. Our goal is to show that \mathcal{I} is equivalent to \mathcal{I}' . For a given realization of s_i , the state ω is a function of the remaining signals s_j , $j \neq i$. Since s'_i and the collection $(s_j)_{i \neq j}$ are independent, we see that ω is independent of s'_i conditional on s_i . In other words, s'_i is also a garbling of s_i . We conclude that both \mathcal{I} is a garbling of \mathcal{I}' and \mathcal{I}' is a garbling of \mathcal{I} , so they are equivalent. \square

The next lemma is used to show that imperfect private information structures cannot correspond to partitions of uniqueness. Before stating it, we will need to introduce the following concept. A *fuzzy partition* is a tuple (g_0, \dots, g_{m-1}) of measurable functions $g_k: [0, 1]^n \rightarrow [0, 1]$ such that $\sum_k g_k = 1$. We can identify this tuple with a single function $g: [0, 1]^n \rightarrow \Delta(\Omega)$. The case of a partition is one in which each g_k is the indicator of a set A_k in a partition of $[0, 1]^n$. As with partitions, we identify two fuzzy partitions if they agree almost everywhere. We denote the collection of fuzzy partitions by G .

We define the projection of g_k to its i th coordinate by $\alpha_i^{g_k}(x_i) = \int_{[0,1]^{n-1}} g_k(x_i, x_{-i}) dx_{-i}$. When g_k is the indicator of a set A_k , the projection $\alpha_i^{A_k}$ as defined in the main text is equal to the projection of g_k . With each partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ of $[0, 1]^n$ we associate the set $G_{\mathcal{A}} = \left\{ g \in G \text{ such that } \forall k, i \alpha_i^{g_k} = \alpha_i^{A_k} \right\}$ of fuzzy partitions that have the same projections as \mathcal{A} .

Lemma 10. *A partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ of $[0, 1]^n$ is a partition of uniqueness if and only if $G_{\mathcal{A}}$ is a singleton.*

Note that $G_{\mathcal{A}}$ always contains at least one element, namely, the indicators of the partition \mathcal{A} , i.e., $(\mathbb{1}_{A_0}, \dots, \mathbb{1}_{A_{m-1}}) \in G_{\mathcal{A}}$. The idea behind the lemma is that all extreme points of $G_{\mathcal{A}}$ are indicators of partitions with the same projections as \mathcal{A} . Hence, if $G_{\mathcal{A}}$ is not a singleton it has at least two distinct extreme points, i.e., there is at least one more partition with the same projections as \mathcal{A} , which is incompatible with the fact that \mathcal{A} is a partition of uniqueness. This identification of extreme points and indicators has appeared before in the context of sets of uniqueness (see [Gutmann et al., 1991](#)).

Proof. First we show that $G_{\mathcal{A}}$ can be treated as a non-empty compact convex subset of a locally convex Hausdorff vector space. Non-emptiness and convexity is straightforward and compactness is to be checked once an appropriate topology is defined.

Let $M([0, 1]^n)$ be the set of all finite signed measures on $[0, 1]^n$ endowed with the topology of weak convergence, making it a locally convex Hausdorff topological vector space. We identify a bounded function $g_k: [0, 1]^n \rightarrow \mathbb{R}$ with a measure μ_k on $[0, 1]^n$ having the density g_k with respect to the Lebesgue measure, i.e., $d\mu_k(x_1, \dots, x_n) = g_k(x)dx_1 \dots dx_n$. Hence, $G_{\mathcal{A}}$ can be identified with a subset of $(M([0, 1]^n))^{\Omega}$. Let $\Delta_{\leq}([0, 1]^n)$ be the set of sub-probability measures, i.e., non-negative measures μ with $\mu([0, 1]^n) \leq 1$. The set $\Delta_{\leq}([0, 1]^n)$ is a compact subset of $M([0, 1]^n)$. As $G_{\mathcal{A}}$ is a subset of the compact set $(\Delta_{\leq}([0, 1]^n))^{\Omega}$, compactness of $G_{\mathcal{A}}$ follows from its closedness. To check closedness, we rewrite the conditions defining $G_{\mathcal{A}}$ in an integrated form using as test functions the continuous functions h on $[0, 1]^n$. The tuple of measures $(\mu_0, \dots, \mu_{m-1}) \in (M([0, 1]^n))^{\Omega}$ belongs to $G_{\mathcal{A}}$ if and only if

$$\int_{[0,1]^n} |h(x_1, \dots, x_n)| d\mu_k \geq 0 \quad (11)$$

$$\sum_k \int_{[0,1]^n} h(x_1, \dots, x_n) d\mu_k = \int_{[0,1]^n} h(x_1, \dots, x_n) dx_1 \dots dx_n \quad (12)$$

$$\int_{[0,1]^n} h(x_i) d\mu_k = \int_{[0,1]} h(x_i) \alpha_i^{A^k}(x_i) dx_i \quad (13)$$

for all $k = 0, \dots, m-1$, $i = 1, \dots, n$, and continuous functions h on $[0, 1]^n$ (in the last condition, h depends on one of the coordinates only). Condition (11) is non-negativity, condition (12) is equivalent to $\sum_k g_k = 1$, and condition (13) corresponds to the equal projections condition $\alpha_i^{g^k} = \alpha_i^{A^k}$. By the definition of the weak topology, integration of a continuous function commutes with taking weak limits. We conclude that $G_{\mathcal{A}}$ contains all its limit points and thus is closed.

By the Krein-Milman theorem, any compact convex subset of a locally convex Hausdorff vector space is the closed convex hull of its extreme points (see [Aliprantis and Border, 2006](#), Theorem 7.68). Thus $G_{\mathcal{A}}$ is the closed convex hull of its extreme points. Consequently, if $G_{\mathcal{A}}$ is not a singleton, it has at least two distinct extreme points. To prove the lemma, it remains to demonstrate that all extreme points of $G_{\mathcal{A}}$ correspond to partitions. Towards a contradiction, assume that $g = (g_0, \dots, g_{m-1})$ is an extreme point of $G_{\mathcal{A}}$ but it is not a partition, i.e., there is a state k_0 such that $g_{k_0}(x) \notin \{0, 1\}$ for $x = (x_1, \dots, x_n)$ in a set of positive Lebesgue measure. Since $\sum_k g_k = 1$, there is $k' \neq k$ such that the set of x where both $g_k(x) > 0$ and $g_{k'}(x) > 0$ has positive measure. Hence, for some $\varepsilon > 0$, the set $D \subseteq [0, 1]^n$ of x such that both

$g_k(x) > \varepsilon$ and $g_{k'}(x) > \varepsilon$ also has positive measure. Without loss of generality, we assume that $k = 0$ and $k' = 1$.

By Corollary 2 of [Gutmann et al. \(1991\)](#), for any D of positive measure, there are two disjoint sets $D_1, D_2 \subseteq D$ also of positive measure having the same projections, i.e., $\alpha_i^{D_1} = \alpha_i^{D_2}$ for any $i = 1, \dots, n$. Hence, the function $a(x) = \varepsilon(\mathbb{1}_{D_1}(x) - \mathbb{1}_{D_2}(x))$ has zero projections, is bounded by ε in absolute value, and is equal to zero outside of the set D . For $\sigma \in \{-1, +1\}$, define

$$g_0^\sigma(x) = g_0(x) + \sigma \cdot a(x), \quad g_1^\sigma(x) = g_1(x) - \sigma \cdot a(x).$$

By definition, g_0^σ and g_1^σ have the same projections as g_0 and g_1 , they are non-negative, and $g_0^\sigma + g_1^\sigma = g_0 + g_1$ (hence, $g_0^\sigma + g_1^\sigma + \sum_{k \geq 2} g_k = 1$).

We conclude that the two tuples $(g_0^\sigma, g_1^\sigma, g_2, g_3, \dots, g_{m-1})$, $\sigma \in \{-1, +1\}$, belong to $G_{\mathcal{A}}$. They are not equal to each other as the sets D_1 and D_2 are disjoint. Since the original collection (g_0, \dots, g_{m-1}) is the average of the two constructed ones, it cannot be an extreme point. This contradiction implies that all the extreme points of $G_{\mathcal{A}}$ correspond to partitions and completes the proof. \square

Relying on the last two lemmas, we can demonstrate that any structure associated with a partition of uniqueness is Blackwell-Pareto optimal.

Lemma 11. *Let \mathcal{I} be a private private information structure equivalent to a structure associated with a partition of uniqueness, then \mathcal{I} is Blackwell-Pareto optimal.*

Proof. Without loss of generality, $\mathcal{I} = (\omega, s_1, \dots, s_n)$ is itself a structure associated with a partition of uniqueness $\mathcal{A} = (A, \dots, A_{m-1})$ of $[0, 1]^n$.

Towards a contradiction, assume that there is a private private information structure \mathcal{J} strictly dominating \mathcal{I} . By Blackwell's theorem, \mathcal{I} is equivalent to some garbling of \mathcal{J} denoted by $\mathcal{I}' = (\omega, s'_1, \dots, s'_n)$. By Lemma 9, \mathcal{I}' is not perfect. Let $t_i = p(s'_i) \in \Delta(\Omega)$ be the posterior belief induced by s'_i and $\mu_i \in \Delta(\Delta(\Omega))$ be its distribution. Consider the structure $\mathcal{I}'' = (\omega, t_i, \dots, t_n)$. It is equivalent to \mathcal{I}' (and hence to \mathcal{I}) by Lemma 2. As t_i is a function of s'_i and \mathcal{I}' is not perfect, \mathcal{I}'' cannot be perfect either (this is also a consequence of the fact that \mathcal{I}'' is a garbling of \mathcal{J}).

Let $f: \Delta(\Omega) \times \dots \times \Delta(\Omega) \rightarrow \Delta(\Omega)$ be the conditional distribution of ω given the realized signals t_1, \dots, t_n . This function is defined μ -everywhere with $\mu = \mu_1 \times \dots \times \mu_n$. As \mathcal{I}'' is not perfect, there is a state $k_0 \in \Omega$ such that $f_{k_0} \notin \{0, 1\}$ on a set of positive μ -measure.

Choose a fuzzy partition $g: [0, 1]^n \rightarrow \Delta(\Omega)$ so that the following identity holds:¹² $g(s_1, \dots, s_n) = f(p(s_1), \dots, p(s_n))$. The distributions of posteriors $(p(s_1), \dots, p(s_n))$ and $(p(t_1), \dots, p(t_n))$ both coincide with μ as the structures \mathcal{I} and \mathcal{I}'' are equivalent. Hence, $g_{k_0} \neq \{0, 1\}$ on a set of positive Lebesgue measure, i.e., g does not correspond to a partition. On the other hand, g has the same projections as the partition \mathcal{A} . Indeed, let us compute $\alpha_i^{g_k}(x)$:

$$\begin{aligned} \alpha_i^{g_k}(x) &= \mathbb{E}[g_k(s_1, \dots, s_n) \mid s_i = x] \\ &= \mathbb{E}[g_k(s_1, \dots, s_{i-1}, x, s_{i+1}, \dots, s_n)] \\ &= \mathbb{E}[f_k(p(s_1), \dots, p(s_{i-1}), q, p(s_{i+1}), \dots, p(s_n))], \end{aligned}$$

where q is the posterior induced by $s_i = x$. Since the distribution of $p(s_j)$ is identical to that of t_j ,

$$\begin{aligned} \alpha_i^{g_k}(x) &= \mathbb{E}[f_k(t_1, \dots, t_{i-1}, q, t_{i+1}, \dots, t_n)] \\ &= \mathbb{E}[f_k(t_1, \dots, t_n) \mid t_i = q] \\ &= q(k), \end{aligned}$$

where in the last equality we rely on the fact that the belief induced by t_i coincides with t_i . Since q is the posterior induced by $s_i = x$, the posterior $q(k)$ is equal to $\alpha_i^{A_k}(x)$.

We thus constructed g not equal to $(\mathbb{1}_{A_0}, \dots, \mathbb{1}_{A_{m-1}})$ but having the same projections. By Lemma 10, the partition $\mathcal{A} = (A_0, \dots, A_{m-1})$ cannot be a partition of uniqueness. This contradiction shows that no structure can dominate the one associated with a partition of uniqueness, i.e., such structures are Blackwell-Pareto optimal. \square

The proof of Theorem 5 is now immediate.

Proof of Theorem 5. By Lemma 8, for each Blackwell-Pareto optimal \mathcal{I} , we can find an equivalent structure associated with a partition of uniqueness $\mathcal{A} = (A_0, \dots, A_{m-1})$. By Lemma 11, any structure admitting such an equivalent representation is Blackwell-Pareto optimal. \square

¹²To construct such a g , define $h_i: [0, 1] \rightarrow \Delta(\Omega)$ by $h_i(x_i)(k) = \mathbb{P}[\omega = k \mid s_i = x_i]$. That is, h_i is the map that assigns to each signal realization the induced posterior, so that $p(s_i) = h_i(s_i)$ holds as an equality of random variables. Then let $g: [0, 1]^n \rightarrow \Delta(\Omega)$ be given by $g(x_1, \dots, x_n) = f(h_1(x_1), \dots, h_n(x_n))$.

B.4 Proof of Theorem 4

Proof. Lorentz (1949)'s characterization of two-dimensional sets of uniqueness uses the idea of a *non-increasing rearrangement* $\hat{\varphi}$ of a function $\varphi : [0, 1] \rightarrow [0, 1]$. The function $\hat{\varphi}$ is defined almost everywhere by the following two properties: it is non-increasing on $[0, 1]$ and, for any $q \in [0, 1]$, the lower-contour sets $\{t \in [0, 1] : \varphi(t) \leq q\}$ and $\{t \in [0, 1] : \hat{\varphi}(t) \leq q\}$ have the same Lebesgue measure. A non-increasing rearrangement exists and moreover is unique (as an element of $L^\infty([0, 1])$).

Lorentz (1949) proved that $A \subseteq [0, 1]^2$ is a set of uniqueness if and only if the non-increasing rearrangements of its two projections are inverses of each other, i.e.,

$$\hat{\alpha}_1^A = (\hat{\alpha}_2^A)^{-1}. \quad (14)$$

Formally, if the inverse $(\hat{\alpha}_2^A)^{-1}(t)$ is not unique for some t , the equality (14) is to be understood as the inclusion: $\hat{\alpha}_2^A(t) \in (\hat{\alpha}_1^A)^{-1}(t)$.

Let us demonstrate that the characterization from Theorem 4 is equivalent to the original characterization of Lorentz (1949). That is, we need to check that a set A is a rearrangement of an upward-closed set if and only if the condition (14) holds. Note that for any downward-closed¹³ set B , its image under the map $x_1 \mapsto 1 - x_1$ and $x_2 \mapsto 1 - x_2$ is upward-closed. Hence, it is enough to check the equivalence between (14) and the existence of a downward-closed rearrangement of A .

Suppose that A is a rearrangement of a downward-closed set B . Towards showing that (14) holds, note that any downward-closed set B can be represented through its projections in two symmetric ways: $B = \{x_2 \leq \alpha_1^B(x_1)\}$ and $B = \{x_1 \leq \alpha_2^B(x_2)\}$ up to a zero-measure set. Hence,

$$\alpha_2^B = (\alpha_1^B)^{-1}. \quad (15)$$

Since B is downward-closed, its projections are non-increasing. Moreover, the sets $\{t \in [0, 1] : \alpha_i^B(t) \leq q\}$ and $\{t \in [0, 1] : \alpha_i^A(t) \leq q\}$ have the same measure for any i and q as B is a rearrangement of A . Thus $\alpha_i^B = \hat{\alpha}_i^A$ and we obtain (14) from (15).

Now assume that the condition (14) is satisfied and construct the downward-closed set B as follows: $B = \{(x_1, x_2) \in [0, 1]^2 : x_2 \leq \hat{\alpha}_1^A(x_1)\}$. By the definition, the projection α_1^B equals $\hat{\alpha}_1^A$. For any downward closed set, the projections satisfy the identity (15) and thus $\alpha_2^B = (\alpha_1^B)^{-1} = (\hat{\alpha}_1^A)^{-1} = \hat{\alpha}_2^A$, where the last equality follows

¹³A set $B \subseteq [0, 1]^2$ is downward-closed if, with each point (x_1, x_2) , it contains all the points $(x'_1, x'_2) \in [0, 1]^2$ such that $x'_1 \leq x_1$ and $x'_2 \leq x_2$.

from (14). Hence, for any i and q , the measure of $\{\alpha_i^B(t) \leq q\}$ coincides with that of $\{\hat{\alpha}_i^A(t) \leq q\}$ and thus with the measure of $\{\alpha_i^A(t) \leq q\}$. We conclude that B is a downward-closed rearrangement of A . \square

B.5 Proof of Theorem 1

First, we show that the conjugate of a cumulative distribution function on $[0, 1]$ is also a cumulative distribution function.

Claim 3. *The conjugate \hat{F} is a cumulative distribution function. Furthermore, it has the same mean: $\int x d\hat{F}(x) = \int x dF(x)$.*

Proof. To show that \hat{F} is a cumulative distribution function it suffices to show that it is weakly increasing, right-continuous, that $\hat{F}(0) \geq 0$, and that $\hat{F}(1) = 1$.

We first note that F^{-1} is weakly increasing, by its definition at x as the minimum of the preimage of $[x, \infty)$ under F . Hence \hat{F} is also weakly increasing.

To see that \hat{F} is right continuous, let $\lim_k x_k = x \in [0, 1]$, with $x_k \leq x$. Then

$$\begin{aligned} \lim_k F^{-1}(x_k) &= \lim_k \min\{y : F(y) \geq x_k\} \\ &= \min\{y : F(y) \geq x\} \\ &= F^{-1}(x) \end{aligned}$$

where the penultimate equality follows from the fact that F is right-continuous. Hence F^{-1} is left-continuous, and so \hat{F} is right-continuous.

It is immediate from the definitions that $\hat{F}(0) \geq 0$ and $\hat{F}(1) = 1$, and thus F is a cumulative distribution function. Finally, the expectations of F and \hat{F} are identical since the shape under F (whose measure is equal to its expectation), given by $\{(x, y) \in [0, 1]^2 : y \leq F(x)\}$ is congruent to the shape under \hat{F} , since one is mapped to the other by the measure-preserving transformation $(x, y) \mapsto (1 - y, 1 - x)$. \square

Now, we prove Theorem 1.

Proof. First, suppose \mathcal{I} is Blackwell-Pareto optimal. By Theorem 3, \mathcal{I} is equivalent to some structure \mathcal{I}' associated with a set of uniqueness A . By Theorem 4, A is a rearrangement of an upward-closed set A' , whose associated structure \mathcal{I}'' must also be equivalent to \mathcal{I} . We show that the two agents' posterior belief distributions induced by \mathcal{I}'' are conjugates of each other.

Define $\tilde{h} : [0, 1] \rightarrow [0, 1]$ by $\tilde{h}(x_1) = \inf\{x_2 : (x_1, x_2) \in A'\}$. We have that \tilde{h} is a decreasing function since A' is upward-closed. Define a left-continuous version of \tilde{h} as $h(x) = \lim_{z \rightarrow x^-} \tilde{h}(z)$. For any $q \in [0, 1]$, in the structure associated with A' , agent 1 has a belief lower or equal to q after observing the signal x_1 if and only if $h(x_1) \geq 1 - q$, so the cumulative distribution function of agent 1's posterior belief is $F_1(q) = \max\{x_1 : h(x_1) \geq 1 - q\}$. For agent 2, note that her posterior belief after any signal lower than $x_2 = h(1 - q)$ is lower or equal to q , while her belief at higher signals are strictly higher than q . So, the cumulative distribution function of agent 2's posterior belief is $F_2(q) = h(1 - q)$. Note that $F_2^{-1}(1 - q) = \min\{y : h(1 - y) \geq 1 - q\} = 1 - \max\{x_1 : h(x_1) \geq 1 - q\} = 1 - F_1(q)$, so F_1 and F_2 are conjugates.

Conversely, suppose the distributions of $p(s_1)$ and $p(s_2)$ in a private private information structure \mathcal{I} are conjugates. Write \tilde{F}_1 and \tilde{F}_2 for the cumulative distribution functions of $p(s_1)$ and $p(s_2)$, and consider the set $A \subseteq [0, 1]^2$ where $(x_1, x_2) \in A$ if and only if $x_2 \geq \tilde{F}_2(1 - x_1)$. We show that the structure associated with A is equivalent to \mathcal{I} ; Figure 4 illustrates the construction. Let $\tilde{h}(x) = \tilde{F}_2(1 - x)$, and define a left-continuous version of \tilde{h} as $h(x) = \lim_{z \rightarrow x^-} \tilde{h}(z)$. For the structure associated with A , by the same argument as above, the distribution function of agent 2's posterior belief is $F_2(q) = h(1 - q) = \tilde{F}_2(q)$. The distribution function of agent 1's posterior belief, by the same argument as above, is $F_1(q) = \max\{x_1 : h(x_1) \geq 1 - q\} = 1 - \min\{x_2 : \tilde{F}_2(x_2) \geq 1 - q\} = 1 - \tilde{F}_2^{-1}(1 - q)$. Using the hypothesis that \tilde{F}_1 and \tilde{F}_2 are conjugates, $1 - \tilde{F}_2^{-1}(1 - q) = \tilde{F}_1(q)$. So, the structure associated with A is equivalent to \mathcal{I} . Because $x_1 \mapsto \tilde{F}_2(1 - x_1)$ is a decreasing function, the set A is upward-closed. Using Theorem 3 and Theorem 4, \mathcal{I} is Blackwell-Pareto optimal. \square

B.6 Proof of Theorem 2

We first demonstrate that the definition of an optimal privacy-preserving recommendation (Definition 4) is equivalent to the requirement of Blackwell-Pareto optimality of the corresponding private private information structure. For this purpose, we need the following result refining Lemma 4 and applicable for any number of states.

Lemma 12. *If $\mathcal{I} = (\omega, s_1, s_2)$ is a private private information structure that is not Blackwell-Pareto optimal, then there exists s'_2 independent of s_1 such that the private private structure (ω, s_1, s'_2) dominates \mathcal{I} and is Blackwell-Pareto optimal.*

Proof. Denote by μ_1 and μ_2 the distributions of beliefs induced by s_1 and s_2 , re-

spectively. Let us verify the existence of a Blackwell-Pareto optimal structure $\mathcal{J} = (\omega, t_1, t_2)$ such that $p(t_1)$ is distributed according to μ_1 and t_2 dominates s_2 . Consider the set $\mathcal{M}(\mu_1)$ of distributions μ'_2 of beliefs such that the pair (μ_1, μ'_2) is feasible, i.e., there is a private private information structure inducing these distributions. In particular, μ_2 belongs to $\mathcal{M}(\mu_1)$. By Lemma 3, $\mathcal{M}(\mu_1)$ is compact in the weak topology as a closed subset of the set of feasible pairs \mathcal{M} . Since the Blackwell order is continuous in the weak topology (see the proof of Lemma 4), there is a maximal element $\mu'_2 \in \mathcal{M}(\mu_1)$ dominating μ_2 . Let $\mathcal{J} = (\omega, t_1, t_2)$ be the private private information structure inducing the pair of distributions (μ_1, μ'_2) . The structure \mathcal{J} must be Blackwell-Pareto optimal. Else, by Lemma 9, there is an equivalent structure $\mathcal{J}' = (\omega, t'_1, t'_2)$ where the signals do not determine the state. Then, by the construction from Lemma 7, there exists an informative signal t independent of (t'_1, t'_2) . By revealing t to the second agent, we obtain a strict Blackwell-Pareto improvement of \mathcal{J}' where the distribution of beliefs induced by the first signal remains fixed, but the distribution of beliefs induced by the second signal is improved to μ''_2 . So we have $\mu''_2 \in \mathcal{M}(\mu_1)$ and μ''_2 strictly dominates μ'_2 , which contradicts the maximality of μ'_2 in $\mathcal{M}(\mu_1)$. This contradiction implies that \mathcal{J} is Blackwell-Pareto optimal.

Without loss of generality, we can assume that $s_1 = p(s_1)$ and $t_1 = p(t_1)$, i.e., signals coincide with the induced posteriors. Thus, by the equivalence of s_1 and t_1 , the joint distributions of (ω, s_1) and (ω, t_1) are the same. We complete the proof by defining s'_2 so that the joint distributions of (ω, s_1, s'_2) and (ω, t_1, t_2) coincide. \square

Corollary 5. *For any number of states, given a one-signal information structure (ω, s_1) , a signal s_2 is an optimal privacy-preserving recommendation if and only if $\mathcal{I} = (\omega, s_1, s_2)$ is a Blackwell-Pareto optimal private private information structure.*

Proof. If \mathcal{I} is Blackwell-Pareto optimal, then it cannot be dominated by any private private structure, in particular, by a structure of the form (ω, s_1, s'_2) . Therefore, s_2 is an optimal privacy-preserving recommendation. Conversely, if \mathcal{I} is not Blackwell-Pareto optimal, then, by Lemma 12, there is a dominating private private structure of the form (ω, s_1, s'_2) . Thus s_2 is dominated by s'_2 , and so it is not an optimal privacy-preserving recommendation. \square

We are now ready to prove Theorem 2.

Proof of Theorem 2. We are given (ω, s_1) and aim to construct a new signal s_2^* independent of s_1 such that any other signal s_2 independent of s_1 is dominated by s_2^* .

As usual, $p(s_1)$ is the belief induced by s_1 . We sample s_2^* uniformly from the interval $[1 - p(s_1), 1]$ if the state is $\omega = 1$ and from $[0, 1 - p(s_1)]$ if $\omega = 0$. Hence, conditioned on s_1 , the constructed signal is distributed uniformly on $[0, 1]$ and so s_2^* is independent of s_1 . Denote by F the cumulative distribution function of $p(s_1)$ and compute the belief induced by s_2^* . The conditional probability of $\omega = 1$ given $s_2^* = t$ is equal to $\mathbb{P}[1 - p(s_1) \leq t]$. Hence, $p(s_2^*) = 1 - F(1 - s_2^*)$. Thus the distribution function F^* of $p(s_2^*)$ is given by

$$F^*(x) = \mathbb{P}[p(s_2^*) \leq x] = \mathbb{P}[1 - F(1 - s_2^*) \leq x] = \mathbb{P}[s_2^* \leq 1 - F^{-1}(1 - x)],$$

where F^{-1} is defined as in (1). Since s_2^* is uniformly distributed on $[0, 1]$, we get $F^*(x) = 1 - F^{-1}(1 - x) = \hat{F}(x)$, where \hat{F} is the conjugate of F .

We conclude that (ω, s_1, s_2^*) is a private private information structure and the distributions of posteriors induced by s_1 and s_2^* are conjugates. Therefore, (ω, s_1, s_2^*) is Blackwell-Pareto optimal by Theorem 1. Corollary 5 implies that s_2^* is an optimal privacy-preserving recommendation.

Now, we show that any s_2 independent of s_1 is weakly dominated by s_2^* . If (ω, s_1, s_2) is itself Blackwell-Pareto optimal, then, by Theorem 1, the cumulative distribution function of beliefs induced by s_2 is \hat{F} and thus s_2 is equivalent to s_2^* . Hence, it suffices to consider the case where $\mathcal{I} = (\omega, s_1, s_2)$ is not Blackwell-Pareto optimal. By Lemma 12, there exists an s_2' dominating s_2 such that (ω, s_1, s_2') is a Blackwell-Pareto optimal private private structure. Hence, by Theorem 1, the distribution of beliefs induced by s_2' is the conjugate of F . Thus s_2' is equivalent to s_2^* and we conclude that s_2^* must also dominate s_2 . \square

B.7 Proof of Proposition 1

The indirect utility of the decision maker $U(q) = \sup_{a \in A} ((1 - q) \cdot u(a, 0) + q \cdot u(a, 1))$ is a continuous convex function as u is assumed to be bounded from above. By a standard argument, the expected payoff of the decision maker observing a signal s is equal to $\mathbb{E}[U(p(s))]$. By Theorem 2, the distribution of $p(s)$ for the optimal privacy-preserving recommendation $s = s_2^*$ is the conjugate to that of $p(s_1)$. Thus Proposition 1 is implied by the following lemma about integrating with respect to conjugate distributions.

Lemma 13. *Let F and \hat{F} be a pair of conjugate distributions, and U be a continuous function. The following identity holds*

$$\int_{[0,1]} U(q) d\hat{F}(q) = \int_{[0,1]} U(1 - F(t)) dt.$$

Proof. Assume first that F is a bijection $[0,1] \rightarrow [0,1]$. By the definition of a conjugate, we obtain

$$\int_{[0,1]} U(q) d\hat{F}(q) = \int_{[0,1]} U(q) d(1 - F^{-1}(1 - q)) = \int_{[0,1]} U(1 - F(t)) dt, \quad (16)$$

where we changed the variable $q = 1 - F(t)$ in the second equality.

Now we show that the identities (16) hold even without the assumption that F is a bijection. Since any continuous function on $[0,1]$ can be approximated by a linear combination of indicators $\mathbb{1}_{[0,a]}$ in the sup-norm, it is enough to prove that

$$\int_{[0,1]} \mathbb{1}_{[0,a]}(q) d\hat{F}(q) = \int_{[0,1]} \mathbb{1}_{[0,a]}(1 - F(t)) dt$$

or, equivalently, that

$$\hat{F}(a) = \lambda(\{t \in [0,1] : 1 - F(t) \leq a\}), \quad (17)$$

where λ stands for the Lebesgue measure. By the monotonicity of F , the set from the right-hand side of (17) is an interval $[t_a, 1]$ where $t_a = \min\{t : F(t) \geq 1 - a\}$, i.e., $t_a = F^{-1}(1 - a)$ as defined in §3. We conclude that (17) holds as it is equivalent to the equality $\hat{F}(a) = 1 - F^{-1}(1 - a)$ defining the conjugate distribution and thus (16) holds as well. \square

B.8 Proof of Proposition 2

Proof. Denote the indirect utility of agent $i \in \{1, 2\}$ by $U_i(q) = \sup_{a_i \in A_i} ((1 - q) \cdot u_i(a_i, 0) + q \cdot u_i(a_i, 1))$. Since each u_i is bounded from above, the indirect utilities are continuous convex functions. The social welfare for a private information structure $\mathcal{I} = (\omega, s_1, s_2)$ can be rewritten as follows: $W(\mathcal{I}) = \mathbb{E}[U_1(p(s_1)) + U_2(p(s_2))]$, where, as usual, $p(s_i)$ is the belief induced by the signal s_i . By the convexity of U_i the welfare is monotone in the Blackwell order, and is therefore always maximized by a Blackwell-Pareto optimal structure.

Let \mathcal{I} be a Blackwell-Pareto optimal structure. By Theorem 1, the distributions of posteriors $p(s_1)$ and $p(s_2)$ are conjugates. Denote the distribution of $p(s_1)$ by μ , which can be an arbitrary measure on $[0, 1]$ with the mean equal to the prior p . Denote the set of all such measures by $\Delta_p([0, 1])$. The choice of $\mu \in \Delta_p([0, 1])$ determines the distribution $\hat{\mu}$ of $p(s_2)$. Thus, to maximize welfare over \mathcal{I} it is enough to find $\mu \in \Delta_p([0, 1])$ maximizing the functional

$$w(\mu) = \int_{[0,1]} U_1(q) d\mu(q) + \int_{[0,1]} U_2(q) d\hat{\mu}(q).$$

Below we check that $w(\mu)$ is convex and continuous in the weak topology. Hence, by Bauer's principle, the optimum is attained at an extreme point of $\Delta_p([0, 1])$. It is well-known that the extreme points of this set are measures with the support of size at most two: see, e.g., [Winkler \(1988\)](#). Since the optimal μ is supported on at most two points, its conjugate $\hat{\mu}$ is supported on at most three points (see the discussion after Theorem 1) and we conclude that there is an optimal structure \mathcal{I} where s_1 takes at most two values and s_2 takes at most three values.

It remains to check that $w(\mu)$ is convex and continuous in the weak topology. For the first integral, this is immediate: it is linear in μ (hence, convex) and continuous thanks to the continuity of the integrand. To show that the two properties hold for the second integral, we rewrite it through the cumulative distribution function F of μ . By Lemma 13,

$$\int_{[0,1]} U_2(q) d\hat{\mu}(q) = \int_{[0,1]} U_2(1 - F(t)) dt.$$

Since U_2 is convex, we conclude that $\int_{[0,1]} U_2(1 - F(t)) dt$ is a convex function of F and, hence, of μ . To show the continuity, note that the weak convergence $\mu_k \rightarrow \mu$ implies the convergence of $F_k(q) \rightarrow F(q)$ for all points q of continuity of F (see [Aliprantis and Border, 2006](#), Theorem 15.3). Since any monotone function is continuous almost everywhere with respect to the Lebesgue measure, the sequence of functions $U_2(1 - F_k)$ converges almost everywhere in $[0, 1]$ and is bounded thanks to the boundedness of U_2 . The Lebesgue dominated convergence theorem implies that $\int_{[0,1]} U_2(1 - F_k(t)) dt$ converges to $\int_{[0,1]} U_2(1 - F(t)) dt$. We conclude that the second integral in $w(\mu)$ is a convex continuous function of μ . Thus the functional $w(\mu)$ is itself continuous and convex. \square

B.9 Proof of Corollary 1

Proof. Since (ω, s_1, s_2) is Blackwell-Pareto optimal, the signal s_2 can be seen as an optimal privacy-preserving recommendation corresponding to (ω, s_1) . By Theorem 2, such s_2 dominates any other signal s'_2 independent of s_1 . Hence, to conclude that s_2 dominates t_2 , it is enough to demonstrate that there is a private private information structure (ω, s'_1, s'_2) such that s'_1 is equivalent to s_1 and s'_2 is equivalent to t_2 . By the assumption, t_1 dominates s_1 and, therefore, the signal s_1 is equivalent to some garbling s'_1 of t_1 . Putting $s'_2 = t_2$, we get the desired private private information structure (ω, s'_1, s'_2) and deduce that t_2 is dominated by s_2 . \square

B.10 Proof of Proposition 4

Proof of Proposition 4. We have $I(\omega; (s_1, \dots, s_n)) \leq H(p)$. By Lemma 1 proved in Appendix B.1, $\sum_{i=1}^n I(\omega; s_i) \leq I(\omega; (s_1, \dots, s_n))$ provided that the signals are independent. \square

B.11 Proof of Proposition 5

Proof. We have $I(\omega; (s_1, \dots, s_n)) \leq H(p)$, so it suffices to show that

$$\sum_i I(\omega; s_i) \leq I(\omega; (s_1, \dots, s_n)) - C_p \sum_{i < j} I(\omega; s_i) I(\omega; s_j). \quad (18)$$

Similarly to the proof of Lemma 1, the result for general n follows from the result for $n = 2$ via an inductive argument. Indeed, assume that the statement holds for $n \leq n_0$ with $n_0 \geq 2$ and show that it holds for $n = n_0 + 1$ as well:

$$\begin{aligned} & I\left(\omega; (s_1, \dots, s_{n_0}, s_{n_0+1})\right) \\ &= I\left(\omega; ((s_1, \dots, s_{n_0}), s_{n_0+1})\right) \\ &\geq I\left(\omega; (s_1, \dots, s_{n_0})\right) + I(\omega; s_{n_0+1}) + C_p \cdot I\left(\omega; (s_1, \dots, s_{n_0})\right) \cdot I(\omega; s_{n_0+1}), \end{aligned}$$

where we applied the two signal version of (18) for the pair of signals (s_1, \dots, s_{n_0}) and s_{n_0+1} . Estimating $I(\omega; s_1, \dots, s_{n_0})$ from below via the n_0 -signal version of (18),

we get

$$\begin{aligned}
& I(\omega; (s_1, \dots, s_{n_0}, s_{n_0+1})) \\
& \geq \sum_{i=1}^{n_0} I(\omega; s_i) + C_p \sum_{1 \leq i < j \leq n_0} I(\omega; s_i) \cdot I(\omega; s_j) + I(\omega; s_{n_0+1}) + \\
& \quad + C_p \cdot I(\omega; s_{n_0+1}) \cdot \left(\sum_{i=1}^{n_0} I(\omega; s_i) + C_p \sum_{1 \leq i < j \leq n_0} I(\omega; s_i) \cdot I(\omega; s_j) \right)
\end{aligned}$$

Eliminating all the cubic terms from the second line can only decrease the right-hand side and leads to inequality (18) for $n = n_0 + 1$:

$$I(\omega; (s_1, \dots, s_{n_0}, s_{n_0+1})) \geq \sum_{i=1}^{n_0+1} I(\omega; s_i) + C_p \sum_{1 \leq i < j \leq n_0+1} I(\omega; s_i) \cdot I(\omega; s_j).$$

It thus remains to prove the result for $n = 2$. We aim to show that

$$I(\omega; s_1) + I(\omega; s_2) - I(\omega; s_1, s_2) \leq -C_p \cdot I(\omega; s_1) \cdot I(\omega; s_2). \quad (19)$$

Denote the left-hand side of (19) by Δ and the posterior probabilities of the high state by $p_i = \mathbb{P}[\omega = 1 | s_i]$ and $p_{12} = \mathbb{P}[\omega = 1 | s_1, s_2]$. By the martingale property, $\mathbb{E}[p_{12} | s_i] = p_i$ and $\mathbb{E}[p_i] = p$. Thanks to the martingale property, we can represent $I(\omega; s_i)$ as follows:

$$I(\omega; s_i) = \mathbb{E} \left[p_{12} \log_2 \left(\frac{p_i}{p} \right) + (1 - p_{12}) \log_2 \left(\frac{1 - p_i}{1 - p} \right) \right],$$

where p_i outside of the logarithm was replaced by p_{12} . Hence,

$$\Delta = \mathbb{E} \left[p_{12} \log_2 \left(\frac{p_1 \cdot p_2}{p_{12} \cdot p} \right) + (1 - p_{12}) \log_2 \left(\frac{(1 - p_1)(1 - p_2)}{(1 - p_{12})(1 - p)} \right) \right].$$

By the concavity of the logarithm, a convex combination of logarithms is at most the logarithm of the convex combination. Therefore,

$$\Delta \leq \mathbb{E} \left[\log_2 \left(\frac{p_1 p_2}{p} + \frac{(1 - p_1)(1 - p_2)}{(1 - p)} \right) \right].$$

Denote the centered posteriors by $\bar{p}_1 = p_1 - p$ and $\bar{p}_2 = p_2 - p$. The right-hand side simplifies to

$$\mathbb{E} \left[\log_2 \left(\frac{p_1 \cdot p_2}{p} + \frac{(1 - p_1)(1 - p_2)}{(1 - p)} \right) \right] = \mathbb{E} \left[\log_2 \left(1 + \frac{\bar{p}_1 \cdot \bar{p}_2}{p(1 - p)} \right) \right].$$

Note that $\frac{\bar{p}_1 \cdot \bar{p}_2}{p(1-p)}$ belongs to the interval $\left[-1, \max\left\{\frac{1-p}{p}, \frac{p}{1-p}\right\}\right]$. Consider the function $f(x) = \log_2(1+x)$. By the Taylor formula, for any $x > -1$,

$$f(x) = f(0) + f'(0) \cdot x + \frac{f''(y)}{2}x^2$$

for some y between 0 and x . Computing the derivatives, we get

$$f(x) = \frac{1}{\ln 2}x + \frac{1}{2\ln 2} \frac{-1}{(1+y)^2}x^2 \leq \frac{1}{\ln 2}x - \frac{\min\{p^2, (1-p)^2\}}{2\ln 2}x^2,$$

where in the last inequality we used the fact that $y \in \left[-1, \max\left\{\frac{1-p}{p}, \frac{p}{1-p}\right\}\right]$. Taking into account that $\min\{p^2, (1-p)^2\} \geq 4p^2(1-p)^2$, we obtain

$$f(x) \leq \frac{1}{\ln 2}x - \frac{2p^2(1-p)^2}{\ln 2}x^2$$

and conclude that

$$\mathbb{E}\left[\log_2\left(1 + \frac{\bar{p}_1 \cdot \bar{p}_2}{p(1-p)}\right)\right] \leq \frac{1}{\ln 2}\mathbb{E}\left[\frac{\bar{p}_1 \cdot \bar{p}_2}{p(1-p)}\right] - \frac{2p^2(1-p)^2}{\ln 2}\mathbb{E}\left[\left(\frac{\bar{p}_1 \cdot \bar{p}_2}{p(1-p)}\right)^2\right].$$

Since the expectation of the product is the product of expectations for the independent random variables \bar{p}_1 and \bar{p}_2 ,

$$\Delta \leq -\frac{2}{\ln 2}\text{Var}[p_1] \cdot \text{Var}[p_2].$$

It remains to lower-bound the variance by the mutual information. The Kullback–Leibler divergence between Bernoulli random variables with success probabilities p and x is defined as follows: $D_{\text{KL}}(x||p) = x \log_2\left(\frac{x}{p}\right) + (1-x) \log_2\left(\frac{1-x}{1-p}\right)$. Then $I(\omega; s_i) = \mathbb{E}[D_{\text{KL}}(p_i||p)]$. Applying the inequality $\ln t \leq t - 1$ to both logarithms and taking into account that $\log_2 t = \frac{1}{\ln 2} \ln t$, we obtain

$$D_{\text{KL}}(x||p) \leq \frac{1}{\ln 2} \left(x \cdot \left(\frac{x}{p} - 1\right) + (1-x) \cdot \left(\frac{1-x}{1-p} - 1\right) \right) = \frac{1}{p(1-p)\ln 2}(x-p)^2$$

for $x \in [0, 1]$. Therefore,

$$\text{Var}[p_i] \geq (p(1-p)\ln 2) \cdot I(\omega; s_i)$$

and we conclude that

$$\Delta \leq -2\ln 2 \cdot p^2(1-p)^2 \cdot I(\omega; s_1) \cdot I(\omega; s_2) = -C_p \cdot I(\omega; s_1) \cdot I(\omega; s_2),$$

which is equivalent to the desired inequality (19). \square

B.12 Proof of Proposition 6

Proof. As in the proof of Proposition 4, we show a stronger statement:

$$\sum_i \bar{I}(\omega; s_i) \leq \bar{I}(\omega; (s_1, \dots, s_n)).$$

This implies the statement of Proposition 6 since \bar{H} is concave, and so, as with mutual information, $\bar{I}(\omega; (s_1, \dots, s_n)) \leq \bar{H}(p)$.

Applying the definition of \bar{I} , and using the martingale property $\mathbb{E}[p(s_i)] = \mathbb{E}[p(s_1, \dots, s_n)] = p$, what we want to prove is that

$$\sum_i \sum_{k \in \Omega} \mathbb{E}[[p(s_i)(k) - p(k)]^2] \leq \sum_k \mathbb{E}[[p(s_1, \dots, s_n)(k) - p(k)]^2].$$

In fact, we prove an even stronger statement, showing that the inequality holds already for each $k \in \Omega$ separately, rather than only when summed over Ω .

To this end, fix k , and denote the centered posteriors by $\bar{p}_i = p(s_i)(k) - p(k)$, so that \bar{p}_i is a zero-mean bounded random variable. Likewise denote $\bar{p} = p(s_1, \dots, s_n)(k) - p(k)$. We want to prove that $\mathbb{E}[\bar{p}^2] \geq \sum_i \mathbb{E}[\bar{p}_i^2]$.

Let V be the vector space of zero-mean random variables spanned by $\{\bar{p}, \bar{p}_1, \dots, \bar{p}_n\}$. As a subspace of L^2 , it is endowed with the inner product given by the expectation of the product.

Since the structure is private private, $\mathbb{E}[\bar{p}_i \cdot \bar{p}_j] = \mathbb{E}[\bar{p}_i] \cdot \mathbb{E}[\bar{p}_j] = 0$ for $i \neq j$. That is, the vectors $\{\bar{p}_1, \dots, \bar{p}_n\}$ are orthogonal. Hence, $V = \text{span}\{q, \bar{p}_1, \dots, \bar{p}_n\}$ for some $q \in V$ that is orthogonal to each \bar{p}_i (note that $q = 0$ is allowed and corresponds to the case where \bar{p} can be represented as a linear combination of \bar{p}_i). Since $\bar{p} \in V$, we can write

$$\bar{p} = \alpha q + \sum_i \alpha_i \bar{p}_i$$

for some scalars $\alpha, \alpha_1, \dots, \alpha_n$. By the martingale property, $\mathbb{E}[\bar{p} | \bar{p}_i] = \bar{p}_i$, and so $\mathbb{E}[(\bar{p} - \bar{p}_i) \cdot \bar{p}_i] = 0$. That is, $\bar{p} - \bar{p}_i$ is orthogonal to \bar{p}_i . Hence $\alpha_i = 1$, and

$$\bar{p} = \alpha q + \sum_i \bar{p}_i.$$

Since $\{q, \bar{p}_1, \dots, \bar{p}_n\}$ are orthogonal,

$$\mathbb{E}[\bar{p}^2] = \alpha^2 \mathbb{E}[q^2] + \sum_i \mathbb{E}[\bar{p}_i^2],$$

and in particular $\mathbb{E}[\bar{p}^2] \geq \sum_i \mathbb{E}[\bar{p}_i^2]$.

□

Note that we used the assumption that the structure is private private only inasmuch as it implies that posteriors of different agents are uncorrelated.

C Non-Uniqueness of Optimal Privacy-Preserving Recommendation for Non-Binary States

Theorem 2 shows that when the state is binary, there is a unique optimal privacy-preserving recommendation s_2^* for each s_1 . In this section, we show that this does not hold for non-binary states.

Consider the case of $\Omega = \{0, 1, 2\}$ where $\omega \in \Omega$ is distributed according to the prior $p = (1/4, 1/2, 1/4)$. The signal s_1 is binary: if $\omega = 2$ then $s_1 = 1$, if $\omega = 0$ then $s_1 = 0$, and if $\omega = 1$ then $s_1 \in \{0, 1\}$ equally likely. The induced beliefs $p(s_1)$ are equal to either $(1/2, 1/2, 0)$ or $(0, 1/2, 1/2)$, each with probability $1/2$.




Figure 8: In the private private information structure associated with the partition (A_0, A_1, A_2) , the signal t_1 induces the same distribution of posteriors as s_1 . For any parameter $\beta \in [0, 1/2]$, this partition is a partition of uniqueness and, hence, we get a one-parametric family of non-equivalent optimal privacy-preserving recommendations given by the signal t_2 .

To construct an optimal privacy-preserving recommendation s_2^* we first build an auxiliary private private information structure (ω, t_1, t_2) , associated with the partition of $[0, 1]^2$ into three sets A_0, A_1 , and A_2 depending on a parameter $\beta \in [0, 1/2]$, as

depicted in Figure 8. The pair of signals (t_1, t_2) is uniformly distributed on $[0, 1]^2$ and the state ω equals k whenever the pair of signals belongs to A_k . Since the area of A_1 is twice the area of A_0 and A_2 , and since the latter two areas are equal, ω has the right distribution $p = (1/4, 1/2, 1/4)$.

Let us check that the signal t_1 is equivalent to s_1 , i.e., it induces the same posterior distribution. Indeed, if the realization of t_1 belongs to $[0, 1/2]$, half of each vertical slice of the square is covered by A_0 and half by A_1 , and so the induced posterior is $p(t_1) = (1/2, 1/2, 0)$ with probability $1/2$. Similarly, for $t_1 \in [1/2, 1]$, we get $p(t_1) = (0, 1/2, 1/2)$ also with probability $1/2$.

Let us check that for different values of β we obtain non-equivalent recommendations. For this purpose, we compute the distribution of posteriors induced by t_2 . Note that t_2 is equivalent to a signal s_2^* taking four different values corresponding to different pairs of sets (A_i, A_j) intersected by the horizontal slice of the square. We get the following distribution of posteriors:

$$p(s_2^*) = \begin{cases} (0, 1/2, 1/2) & \text{with probability } \beta \\ (1/2, 0, 1/2) & \text{with probability } 1/2 - \beta \\ (0, 1, 0) & \text{with probability } 1/2 - \beta \\ (1/2, 1/2, 0) & \text{with probability } \beta \end{cases}$$

For different values of β we get different distributions, i.e., the constructed recommendations are not equivalent.

It remains to show that for any value of β , the signal s_2^* is an optimal privacy-preserving recommendation. To this end, we check that the partition (A_0, A_1, A_2) is a partition of uniqueness (as defined in Appendix B.3). Therefore, by Theorem 5, the information structure (ω, t_1, t_2) is Blackwell-Pareto optimal. Thus t_2 is an optimal privacy-preserving recommendation, and so is s_2^* as it is equivalent to t_2 .

To show that (A_0, A_1, A_2) is a partition of uniqueness, we rely on the following elementary but useful general observation: if in a partition (A_0, \dots, A_{m-1}) of $[0, 1]^n$ all sets except for possibly one are sets of uniqueness, then the partition itself is a partition of uniqueness. In our example, the set A_2 is upward-closed and hence is a set of uniqueness by Theorem 4. The set A_0 is a rearrangement of an upward-closed set (since it can be made upward-closed via a measure-preserving reparametrization of the axes) and so is a set of uniqueness by the same theorem. Thus the partition (A_0, A_1, A_2) is a partition of uniqueness and s_2^* is an optimal privacy-preserving

recommendation for any value of β .

The partition (A_0, A_1, A_2) provides an interesting example of the fact that a partition of uniqueness is not necessarily composed of sets of uniqueness. Indeed, for $\beta \neq 0$, the set A_1 is not a set of uniqueness as it has the same marginals as the set obtained by the reflection of A_1 with respect to the vertical line $t_1 = 1/2$.

C.1 Representing Private Private Signals for Binary ω as Sets

To simplify notation, in this section, we consider the case of $n = 2$ agents and a binary state $\omega \in \Omega = \{0, 1\}$. Nevertheless, the same ideas apply more generally to finitely many agents and possible values of the state. By Proposition 3, any private private information structure \mathcal{I} is equivalent to a structure associated with some set $A \subseteq [0, 1]^2$, which we denote by $\mathcal{I}_A = (\omega, s_1, s_2)$. In this section, we show how to construct \mathcal{I}_A given \mathcal{I} . We begin with the case where \mathcal{I} is uninformative and describe \mathcal{I}_A for any prior $p = \mathbb{P}[\omega = 1]$. Relying on this construction, we then describe how to construct \mathcal{I}_A for any \mathcal{I} with a finite number of possible signal values.

Recall that in \mathcal{I}_A , the signals (s_1, s_2) are uniformly distributed on $[0, 1]^2$, the state is $\omega = \mathbb{1}_A(s_1, s_2)$, and A is some measurable subset of $[0, 1]^2$ with Lebesgue measure $\lambda(A) = p$ so that $p = \mathbb{P}[\omega = 1]$. Recall that the distribution of posteriors induced by \mathcal{I}_A can be computed as follows: The conditional probability of the high state given that agent i receives a signal $s_i = t$ is exactly $\alpha_i^A(t)$, the one-dimensional Lebesgue measure of the cross-section $\{(y_1, y_2) \in A : y_i = t\}$. In other words, $\alpha_i^A(s_i)$ is i 's posterior corresponding to s_i and the induced distribution of posteriors μ_i is the image of the uniform distribution under the map α_i^A , i.e., $\mu_i([0, t])$ equals the Lebesgue measure of $\{x_i \in [0, 1] : \alpha_i^A(x_i) \leq t\}$.

Example 1 (Non-informative signals). Consider a private private information structure \mathcal{I} , where both agents receive completely uninformative signals, i.e., the induced posteriors are equal to the prior p .

To find an equivalent structure \mathcal{I}_A , we need to construct a set $A = A_p \subseteq [0, 1]^2$ such that the Lebesgue measure of all its projections equals p . To this end, let Y be any subset of $[0, 1]$ with measure p (e.g., $[0, p]$), and let

$$A = \{(x_1, x_2) \in [0, 1]^2 : [x_1 + x_2] \in Y\},$$

where $[x]$ is the fractional part of $x \in \mathbb{R}$. It is easy to see that A indeed has the desired property.

It turns out that the construction of an information structure \mathcal{I}_{A_p} representing completely uninformative signals can be used to find a representation for any information structure with a finite number of possible signal values.

Example 2 (Arbitrary finite number of signal values). Let $\mathcal{I} = (\omega, s_1, s_2)$ be a private private information structure with $n = 2$ agents and finite signal spaces S_1 and S_2 . Our goal is to construct a set $A \subseteq [0, 1]^2$ such that the structure \mathcal{I}_A associated with A is equivalent to \mathcal{I} .

For each agent $i \in \{1, 2\}$, consider a disjoint partition of $[0, 1]$ into intervals A_{s_i} , $s_i \in S_i$, so that the length of each A_{s_i} coincides with the probability that the signal $s_i \in S_i$ is sent under \mathcal{I} . Let $q(s_1, s_2) \in [0, 1]$ be the conditional probability of $\{\omega = 1\}$ given signals (s_1, s_2) .

Recall that, in Example 1, we constructed a set $A_p \subseteq [0, 1]^2$ such that its projection to each of the coordinates has a constant density p . Now we construct A by pasting the appropriately rescaled copy of $A_{q(s_1, s_2)}$ into each rectangle $A_{s_1} \times A_{s_2}$. Denote by $T_{[a, b] \times [c, d]}$ an affine map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ that identifies $[0, 1]^2$ with $[a, b] \times [c, d]$:

$$T_{[a, b] \times [c, d]}(x_1, x_2) = (a + (b - a)x_1, c + (d - c)x_2).$$

We define A as the following disjoint union:

$$A = \bigsqcup_{s_1 \in S_1, s_2 \in S_2} T_{A_{s_1} \times A_{s_2}} \left(A_{q(s_1, s_2)} \right).$$

Let $s'_i \in [0, 1]$ be a signal received by an agent i in \mathcal{I}_A . The signal s'_i falls into A_{s_i} with the same probability that i receives the signal s_i in \mathcal{I} . By construction, the conditional probability of $\{\omega = 1\}$ given s'_i is constant over each interval A_{s_i} and coincides with the posterior $p_i(s_i)$ that i gets under \mathcal{I} . We conclude that \mathcal{I} and \mathcal{I}_A induce the same distribution of posteriors and so are equivalent.