

Privacy-Constrained Signals

Zhang Xu*

Wei Zhao[†]

November 27, 2025

Abstract

This paper provides a unified approach to characterize the set of all feasible signals subject to privacy constraints. The Blackwell frontier of feasible signals can be decomposed into minimum informative signals achieving the Blackwell frontier of privacy variables, and conditionally privacy-preserving signals. A complete characterization of the minimum informative signals is then provided. We apply the framework to ex-post privacy (including differential and inferential privacy) and to constraints on posterior means of arbitrary statistics.

*School of Economics, Renmin University of China. *Email:* xuzhang@ruc.edu.cn

[†]School of Economics and Management, Tsinghua University. *Email:* wei.zhao@outlook.fr

Contents

1	Introduction	3
2	Model	4
2.1	Privacy-Constrained Signals	6
3	Characterization of Privacy-Constrained Signals	7
4	Blackwell Frontier of Privacy-Permissible Set	12
4.1	Ex-Post Privacy	12
4.2	Posterior-Mean Privacy	14
5	Discussion and Future Work	15
A	Appendix	17
A.1	Proofs for Section 3	17
A.2	Proofs for Section 4	20

1 Introduction

The big data plays a critical role in economic decisions, promoting efficiency in allocation. At the same time, growing concern on privacy has been drawn. The abuse of sensitive (personal) data, leading to statistical and price discrimination, imposes negative externality on the economy as a whole. A natural question is, what is the set of feasible datasets, subject to privacy constraints on sensitive information? The past literature on this question mainly focus on perfectly privacy-preserving constraints. However, privacy-preserving constraints may be too demanding in practical operations. In this spirit, various orders, both complete and partial, have been proposed to measure the degree of information leakage.

This paper develops a general framework for privacy constraints. We model information disclosure through signals defined on an abstract state of the world $\tilde{\omega} \in \Omega$, and represent the sensitive information as a random variable $\tilde{\theta} : \Omega \rightarrow \Theta$ defined on the same state space. For instance, if Ω is an n -dimensional space, the sensitive component Θ may correspond to its first m dimensions with $m < n$. Following Blackwell (1953), each signal (experiment) can be represented by the distribution of posteriors it induces. We therefore model privacy constraints as a subset of distributions over posteriors about the sensitive variable $\tilde{\theta}$, which we refer to as the *privacy-permissible set*. A signal is privacy-constrained if the posterior distribution it induces, when marginalized over Θ , belongs to this permissible set. For the analysis to be well behaved and natural, we assume that the privacy-permissible set is a lower set with respect to the Blackwell order, that is, whenever a signal is permissible, any less informative (in Blackwell sense) signal is also permissible.

To characterize the set of privacy-constrained signals, it is equivalent to describe its Blackwell frontier, that is, the set of all privacy-constrained signals that are Blackwell-undominated. Theorem 1 reduces this task to characterizing the Blackwell frontier of the privacy-permissible set itself. Given a distribution over posteriors about the sensitive variable, $\gamma \in \Delta(\Delta(\Theta))$, we first construct a minimum-informative extension τ_γ , which preserves the marginal distribution over $\Delta(\Theta)$ while revealing as little information as possible about the state $\tilde{\omega}$. Afterward, we disclose τ_γ , and then disclose a Blackwell-undominated signal among those that are conditionally privacy-preserving given τ_γ . The latter class of signals is characterized in Strack and Yang (2024). This sequential procedure yields an element of the Blackwell frontier of privacy-constrained signals.

Theorem 2 describes how to generate all minimum-informative extensions for a given $\gamma \in \Delta(\Delta(\Theta))$. Under a minimum-informative extension, each posterior over $\tilde{\theta}$ in the support of γ is extended to exactly one posterior over the full state $\tilde{\omega}$. Thus, computing a minimum-

informative extension amounts to assigning a conditional distribution of $\tilde{\omega}$ given $\tilde{\theta}$ to every element in the support of γ . In the discrete setting, this assignment can be expressed through a finite collection of linear constraints.

However, characterizing the Blackwell frontier of an abstract privacy-permissible set remains challenging. In Section 4, we focus on two important classes of privacy constraints and show how to characterize their Blackwell frontiers. The first class is *ex-post privacy*, which encompasses differential privacy Dwork et al. (2006), inferential privacy Ghosh and Kleinberg (2016); Wang et al. (2025), and Bayesian privacy Eilat et al. (2021). When the regulator is concerned with ex post privacy loss, the constraint is imposed directly on the posterior beliefs about the sensitive variable $\tilde{\theta}$. Ex-post privacy specifies a subset of posteriors over $\tilde{\theta}$, and a signal is ex post privacy-constrained if every realized posterior about $\tilde{\theta}$ lies in this subset. Proposition 1 shows that under ex-post privacy, characterizing the Blackwell frontier of the privacy-permissible set reduces to identifying the extreme points of the permissible posteriors over $\tilde{\theta}$. Proposition 2 further provides an explicit characterization of the Blackwell frontier under inferential privacy. Related results for the discrete cases of inferential privacy and differential privacy appear in Xu and Zhao (2025). The second class is *posterior-mean privacy*, where the regulator is concerned only with the information revealed about the posterior mean of a statistic defined on the sensitive variable. In this case, Proposition 3 provides a clean characterization of the corresponding Blackwell frontier.

The remainder of the paper is organized as follows. Section 2 introduces the formal setting. Section 3 presents the general characterization of privacy-constrained signals. Section 4 characterizes the Blackwell frontier of the privacy-permissible set for several privacy constraints. Section 5 offers further discussion. All proofs are collected in the Appendix.

2 Model

Let $(\Omega, \mathcal{B}(\Omega), \mu_0)$ be a probability space, where $\mathcal{B}(\cdot)$ is the Borel σ -algebra generator, $(\Omega, \mathcal{B}(\Omega))$ is a standard Borel space and $\mu_0 \in \Delta(\Omega)$ is an interior prior.¹ The *state* is a random variable $\tilde{\omega} \sim \mu_0$ and $\omega \in \Omega$ is a realization of state. The *privacy* is formalized as a random variable $\tilde{\theta} : \Omega \rightarrow \Theta$, where $(\Theta, \mathcal{B}(\Theta))$ is a standard Borel space. $\theta \in \Theta$ is a realization of privacy. For example, ω is the vector containing consumer characteristics such as gender, race, willingness to pay, and history records, and θ is the sensitive subvector consisting of gender and race.

¹A Borel space $(E, \mathcal{B}(E))$ is standard if there is an isomorphism $\psi : E \rightarrow F$ for some $F \in \mathcal{B}(\mathbb{R})$. An isomorphism is a bijection ψ such that both ψ and ψ^{-1} are measurable. (Çınlar (2011), p.11.)

A *signal* is a random variable $\pi : \Omega \times [0, 1] \rightarrow S$, where $(S, \mathcal{B}(S))$ is a standard Borel space. An element $s \in S$ is a *signal realization*. Let $\tilde{r}$ be an auxiliary random variable uniformly distributed on $[0, 1]$. For each realization (ω, r) , the signal realization is given by $\pi(\omega, r)$. Let λ be the Lebesgue measure and $\mathbb{P} := \mu_0 \times \lambda$ be the product measure induced by μ_0 over Ω and λ over $[0, 1]$. Then the distribution of π is $p^\pi(\cdot) := \mathbb{P}(\{(\omega, r) : \pi(\omega, r) \in \cdot\})$ (i.e., $p^\pi(B) = \mathbb{P}(\{(\omega, r) : \pi(\omega, r) \in B\})$, $\forall B \in \mathcal{B}(S)$, and similarly hereinafter).² The conditional probability of π given ω is $p_\omega^\pi(\cdot) := \mathbb{P}(\pi \in \cdot | \omega)$.³

Observing signal realization s induces the posterior $\mu_s(\cdot) := \mathbb{P}(\tilde{\omega} \in \cdot | s)$. Given signal π , let $\tilde{\mu}_\pi$ denote the associated belief-valued random variable of posterior belief μ_s . Moreover, let $\langle \pi \rangle(\cdot) := p^\pi(\{s \in S : \mu_s \in \cdot\})$ denote the distribution of $\tilde{\mu}_\pi$. We write $(\cdot)^\theta$ for the operator mapping a distribution over Ω or $\Delta(\Omega)$ to its marginal over Θ or $\Delta(\Theta)$. For $\mu \in \Delta(\Omega)$, $\mu^\theta(\cdot) = \mu(\{\omega : \tilde{\theta}(\omega) \in \cdot\})$, and for $\tau \in \Delta(\Delta(\Omega))$, $\tau^\theta(\cdot) = \tau(\{\mu : \mu^\theta \in \cdot\})$. Therefore, μ_s^θ represents the posterior about privacy induced by signal realization s , and $\langle \pi \rangle^\theta$ represents the distribution of posteriors about privacy.

To keep the notation clear, we use $\mu \in \Delta(\Omega)$ to denote an arbitrary posterior over the state $\tilde{\omega}$, and $\nu \in \Delta(\Theta)$ to denote an arbitrary posterior over the privacy variable $\tilde{\theta}$. Likewise, we use $\tau \in \Delta(\Delta(\Omega))$ for an arbitrary distribution over posteriors about $\tilde{\omega}$, and $\gamma \in \Delta(\Delta(\Theta))$ for an arbitrary distribution over posteriors about $\tilde{\theta}$.

Denote Π by the set of all signals. For two signals $\pi, \pi' \in \Pi$, we say that π *Blackwell dominates* π' and write $\pi \succeq \pi'$ provided $\langle \pi \rangle$ is a *mean-preserving spread* of $\langle \pi' \rangle$ which is also denoted as $\langle \pi \rangle \succeq \langle \pi' \rangle$ (Blackwell 1951, 1953; Strassen 1965).⁴ Given $\pi \succeq \pi'$, if $\pi \preceq \pi'$ also holds, then π and π' are *Blackwell equivalent*, written $\pi \sim \pi'$; otherwise, π *strictly* Blackwell dominates π' , written $\pi \succ \pi'$. We say that π Blackwell dominates π' *in terms of* $\tilde{\theta}$ and write $\pi \succeq_\theta \pi'$ if $\langle \pi \rangle^\theta$ is a mean-preserving spread of $\langle \pi' \rangle^\theta$. The corresponding relations $\sim_\theta$ and $\succ_\theta$ are defined analogously. Whenever we refer to Blackwell dominance without the qualifier “in terms of $\tilde{\theta}$,” we mean dominance with respect to $\tilde{\omega}$.

² p^π is a probability measure on $(S, \mathcal{B}(S))$; it is called *distribution* of π . If $S \subseteq \mathbb{R}$, then $F^\pi : \mathbb{R} \rightarrow [0, 1]$, s.t. $F^\pi(s) = p^\pi(\pi \leq s)$ for all $s \in \mathbb{R}$ is called the *distribution function* or CDF of π .

³Since $(\Omega \times [0, 1], \mathcal{B}(\Omega) \otimes \mathcal{B}([0, 1]))$ and $(S, \mathcal{B}(S))$ are standard Borel spaces, there exist regular versions of $\mathbb{P}(\pi \in \cdot | \omega)$ and $\mathbb{P}(\tilde{\omega} \in \cdot | s)$ respectively (Çinlar (2011), Theorem 2.19, p.154).

⁴For $\tau, \tau' \in \Delta(\Delta(\Omega))$, τ is a mean-preserving spread of τ' if there is a *dilation* $K : \Delta(\Omega) \rightarrow \Delta(\Delta(\Omega))$ [i.e., a Markov kernel such that $\forall \mu' \in \text{supp}(\tau')$, $\mu' = \int_{\Delta(\Omega)} \mu K(d\mu | \mu')$ (mean-preservation)], such that $\tau(\cdot) = \int_{\Delta(\Omega)} K(\cdot | \mu') \tau'(d\mu')$ (spread). The similar definition applies to any $\gamma, \gamma' \in \Delta(\Delta(\Theta))$.

2.1 Privacy-Constrained Signals

From an informational perspective, each signal can be identified with the posterior distribution it induces. Thus, analyzing signals is equivalent to working directly with Bayesian-plausible distributions over posteriors (cf. Blackwell (1953); Kamenica and Gentzkow (2011)).

Let $\Gamma := \{\gamma \in \Delta(\Delta(\Theta)) : \mathbb{E}_\gamma[\nu] = \mu_0^\theta\}$ denote the Bayesian-plausible distributions of posteriors about privacy. If there is no constraint on how much information about privacy may be disclosed, then any posterior distribution in Γ can be induced. Hence, in general, a privacy constraint can be imposed directly on the set Γ .

In particular, let

$$\mathcal{P} \subseteq \Gamma$$

be a nonempty subset containing the posterior distributions about $\tilde{\theta}$ that are allowed to be disclosed. We refer to $\mathcal{P}$ as the *privacy-permissible set*. To ensure that the privacy constraint is well behaved, we impose the following assumptions on $\mathcal{P}$.

Assumption 1. $\mathcal{P}$ is a lower set with respect to the Blackwell order; that is, if $\gamma \in \mathcal{P}$ and $\gamma' \preceq \gamma$, then $\gamma' \in \mathcal{P}$.

Assumption 1 is natural and is required by any reasonable notion of a privacy constraint. It states that if a certain amount of information about privacy may be disclosed, then revealing any less informative disclosure must also be permissible.

Assumption 2. $\mathcal{P}$ is a closed set; that is, if $\{\gamma_t\}_{t \in \mathbb{N}_+}$ satisfies $\gamma_t \in \mathcal{P}$ for all $t \in \mathbb{N}_+$ and $\gamma^* := \lim_{t \rightarrow \infty} \gamma_t$ exists, then $\gamma^* \in \mathcal{P}$.

Assumption 2 is a technical requirement that simplifies our characterization of privacy-constrained signals. When a decision-maker selects a signal in $\mathcal{P}$ to maximize an objective function, taking the supremum is equivalent to optimizing over the closure of $\mathcal{P}$. Let

$$\overline{\mathcal{P}} := \{\gamma \in \mathcal{P} : \nexists \gamma' \in \mathcal{P} \text{ such that } \gamma' \succ \gamma\}$$

be the *Blackwell frontier* of $\mathcal{P}$. Under Assumption 1 and 2, $\overline{\mathcal{P}} \neq \emptyset$ and $\mathcal{P} = \{\gamma \in \Gamma : \exists \bar{\gamma} \in \overline{\mathcal{P}} \text{ such that } \gamma \preceq \bar{\gamma}\}$.⁵ $\overline{\mathcal{P}}$ gives several upper bounds of the amount of information

⁵Let $\{\gamma_t\}_{t \in \mathbb{N}_+}$ be a sequence in $\mathcal{P}$ such that $\gamma_t \prec \gamma_{t+1}$ for all t . For each t , let $\tilde{\nu}_t$ be a belief-valued random variable satisfying $\tilde{\nu}_t \sim \gamma_t$. Then $\{\tilde{\nu}_t\}_{t \in \mathbb{N}_+}$ is a martingale (rigorously, it requires γ_{t+1} is sufficient for γ_t , see footnote 9). By the martingale convergence theorem (Doob 1951), the limit $\nu^* = \lim_{t \rightarrow \infty} \nu_t$ exists. Let γ^* be the distribution of ν^* . Then, $\lim_{t \rightarrow \infty} \gamma_t = \gamma^*$. By Assumption 2, $\gamma^* \in \mathcal{P}$. Then, $\gamma^* \in \overline{\mathcal{P}}$.

can be disclosed about privacy. Moreover, when $\mathcal{P}$ is not closed, our characterization of privacy-constrained signals can be made correct by removing some signals whose induced distribution over posterior about privacy is in the Blackwell frontier $\overline{\mathcal{P}}$.

Definition 1. *A signal π is a $\mathcal{P}$ -privacy-constrained signal if $\langle \pi \rangle^\theta \in \mathcal{P}$.*

Example 1 (Privacy-Preserving Signals). *When $\mathcal{P} = \{\delta_{\mu_0^\theta}\}$, $\mathcal{P}$ -privacy-constrained signals reduce to the privacy-preserving signals introduced by [He et al. \(2021\)](#) and [Strack and Yang \(2024\)](#). [Strack and Yang](#) show that all privacy-preserving signals can be generated by garbling and reordering of a conditionally revealing quantile signal.*

Example 2 (Single-Bound Privacy). *A natural generalization of privacy-preserving signals is to consider $\mathcal{P} = \{\gamma \in \Gamma : \gamma \preceq \bar{\gamma}\}$ for some $\bar{\gamma} \in \Gamma$. The element $\bar{\gamma}$ provides a single upper bound on the amount of information that may be disclosed about privacy. A signal π is said to be $\bar{\gamma}$ -privacy-constrained if $\langle \pi \rangle^\theta \preceq \bar{\gamma}$.*

This framework also includes differential privacy ([Dwork et al. 2006](#)), inferential privacy ([Ghosh and Kleinberg 2016](#); [Wang et al. 2025](#)), and privacy constraints defined through the posterior mean of a statistic. These concepts are discussed in greater detail in [Section 4](#).

3 Characterization of Privacy-Constrained Signals

Compare with privacy-preserving constraint, our $\mathcal{P}$ -privacy constraint allows us to disclose some information about privacy. Based on this, we can intuitively construct the following two-stage disclosure rule:

Stage 1. Construct a signal $\pi_1 : \Omega \times [0, 1] \rightarrow S_1$ to release permissible information about privacy $\tilde{\theta}$.

Stage 2. Construct another signal $\pi_2 : \Omega \times [0, 1] \rightarrow S_2$, which is conditionally privacy-preserving given π_1 , i.e., for almost every $B \in \mathcal{B}(\Theta)$ and $s_1 \in S_1, s_2 \in S_2$,

$$\mathbb{P}(\tilde{\theta} \in B | s_2, s_1) := \mu_{(s_1, s_2)}^\theta(B) = \mu_{s_1}^\theta(B).$$

Let $\pi_1 \vee \pi_2 : \Omega \times [0, 1] \rightarrow S_1 \times S_2$ denote the *join* of two signals π_1 and π_2 . When the state-randomizer pair (ω, r) is realized, the joint signal is $\pi_1 \vee \pi_2(\omega, r) = (\pi_1(\omega, r), \pi_2(\omega, r))$. Observing $\pi_1 \vee \pi_2$ is equivalent to observing π_1 and then π_2 sequentially, with beliefs updated

⁵ δ_a is the Dirac delta function, i.e., $\delta_a(x) = 0$ if $x \neq a$ and $\int \delta_a(x) dx = 1$.

at each stage according to Bayes' rule. It is immediate that if π_1 is $\mathcal{P}$ -privacy-preserving and π_2 is conditionally privacy-preserving given π_1 , then the joint signal $\pi_1 \vee \pi_2$ is $\mathcal{P}$ -privacy-preserving as well. Conditional privacy preservation corresponds to designing π_2 separately for each realization of π_1 ; see Remark 2 in [Strack and Yang \(2024\)](#). In other words, constructing π_2 amounts to constructing a privacy-preserving signal on each posterior belief induced by π_1 .

Denote by $\Pi_{\mathcal{P}}$ the set of $\mathcal{P}$ -privacy-constrained signals, and define its Blackwell frontier as $\bar{\Pi}_{\mathcal{P}} := \{\pi \in \Pi_{\mathcal{P}} : \nexists \pi' \in \Pi_{\mathcal{P}} \text{ such that } \pi' \succ \pi\}$. Any $\pi \in \bar{\Pi}_{\mathcal{P}}$ is called a *Blackwell-undominated* $\mathcal{P}$ -privacy-constrained signal. In what follows, we present our main result, Theorem 1, which characterizes the Blackwell frontier $\bar{\Pi}_{\mathcal{P}}$ via two-stage disclosure rules. Any $\mathcal{P}$ -privacy-constrained signal can then be obtained by garbling elements of this frontier.

Lemma 1. *A signal is $\mathcal{P}$ -privacy-constrained if and only if it is Blackwell dominated by a signal π' such that $\langle \pi' \rangle^\theta \in \bar{\mathcal{P}}$.*

Lemma 1 states that, without loss of generality, constructing the Blackwell frontier of $\mathcal{P}$ -privacy-constrained signals reduces to considering only those signals whose induced distribution over posteriors about privacy lies on the Blackwell frontier of $\mathcal{P}$. Then, by garbling these signals in the frontier, we obtain all $\mathcal{P}$ -privacy-constrained signals. Therefore, in what follows, we focus on constructing the Blackwell frontier of $\mathcal{P}$ -privacy-constrained signals from the Blackwell frontier of $\mathcal{P}$.

For each $\gamma \in \bar{\mathcal{P}}$, we construct its *minimum-informative extensions*, namely every posterior about state $\tau_\gamma \in \mathcal{T}$ satisfying

$$\begin{aligned} \tau_\gamma^\theta &= \gamma \text{ almost surely} && \text{(Extension),} \\ \nexists \tau \in \mathcal{T} \text{ such that } \tau^\theta &= \gamma \text{ almost surely, and } \tau \prec \tau_\gamma && \text{(Minimum informative).} \end{aligned}$$

That is, τ_γ extends γ from a distribution over posteriors about privacy $\tilde{\theta}$ to a distribution over posteriors about the full state $\tilde{\omega}$, while revealing the least additional information about the state needed to sustain the same distribution of posteriors about privacy. Since a posterior distribution can itself be viewed as a signal, the object τ_γ provides exactly the first-stage disclosure we seek. To simplify notation, we use τ_γ to denote any signal π such that $\langle \pi \rangle = \tau_\gamma$.⁶

⁶Formally, τ can induce a joint distribution over $\Omega \times \Delta(\Omega)$, $p^{(\tilde{\omega}, \tilde{\mu})}$ such that for any $B_\Omega \times B_{\Delta(\Omega)} \in \mathcal{B}(\Omega) \otimes \mathcal{B}(\Delta(\Omega))$, $p^{(\tilde{\omega}, \tilde{\mu})}(B_\Omega \times B_{\Delta(\Omega)}) = \int_{B_{\Delta(\Omega)}} \mu(B_\Omega) \tau(d\mu)$. Since $\Delta(\Omega)$ with the Borel σ -algebra induced by the weak-* topology, is standard Borel, there exists a regular version of conditional distribution of $\tilde{\mu}$ given ω , denoted by $p_\omega^{\tilde{\mu}}$ ([Çinlar \(2011\)](#), Theorem 2.18, p.154). Then there exists a measurable function $\pi_\tau : \Omega \times [0, 1] \rightarrow \Delta(\Omega)$, such that $\pi_\tau(\omega, \tilde{r})$ has distribution $p_\omega^{\tilde{\mu}}$ ([Kallenberg \(1997\)](#), Lemma 2.22, p.34). Therefore, π_τ is the signal with $\langle \pi_\tau \rangle = \tau$.




Figure 1: Construction of Blackwell-Undominated Privacy-Constrained Signals

For a $\gamma \in \overline{\mathcal{P}}$, the minimum-informative extension need not be unique. Let $\mathcal{T}_\gamma$ denote the set of all minimum-informative extensions of γ .

Theorem 1 (Characterization of Privacy-Constrained Signals). *A signal π belongs to the Blackwell frontier of $\mathcal{P}$ -privacy-constrained signals, $\overline{\Pi}_{\mathcal{P}}$, if and only if there exist $\tau_\gamma \in \mathcal{T}_\gamma$ for some $\gamma \in \overline{\mathcal{P}}$ and $\tilde{q}_{\tau_\gamma}$ that is Blackwell-undominated conditionally privacy-preserving given τ_γ such that π is Blackwell equivalent to the joint signal $\tau_\gamma \vee \tilde{q}_{\tau_\gamma}$.*

Briefly, Theorem 1 shows that, once τ_γ is made public, the remaining task of constructing the Blackwell frontier of $\mathcal{P}$ -privacy-constrained signals reduces to constructing the Blackwell frontier of privacy-preserving signals, as characterized in [Strack and Yang \(2024\)](#).

The key insight of our approach is that, in order to construct the most informative signal whose distribution over posteriors about privacy is γ , we begin by identifying the least informative one. Once this baseline signal is obtained, every signal with the same γ can be generated by disclosing additional conditionally privacy-preserving information. By contrast, if one starts with a signal that already contains unnecessary information beyond what is required to sustain γ , then taking its join with any conditionally privacy-preserving signal inevitably preserves this extra information, and thus cannot characterize the full set of feasible signals sustaining γ . Figure 1 summarizes the structure of Blackwell-undominated privacy-constrained signals described in Theorem 1.

Additionally, the characterization of the minimum-informative extension of a given $\gamma \in \overline{\mathcal{P}}$ is straightforward.

Theorem 2 (Characterization of Minimum-Informative Extensions). $\underline{\tau} \in \mathcal{T}$ is a minimum-informative extension of a $\gamma \in \overline{\mathcal{P}}$ if and only if

- (1) $\underline{\tau}^\theta = \gamma$ almost surely;
- (2) For almost every $\mu, \hat{\mu} \in \text{supp}(\underline{\tau})$, if $\mu^\theta = \hat{\mu}^\theta$ almost surely, then $\mu = \hat{\mu}$ almost surely.

The key condition in Theorem 2 is (2), which requires that, in the minimum-informative extension, each $\nu \in \text{supp}(\gamma)$ extends to exactly one $\mu \in \Delta(\Omega)$. Otherwise, any two such μ 's with the same marginal μ^θ can be merged into a single posterior, implying that the original extension was not minimal.

Theorem 2 is particularly useful in the discrete setting. Consider finite sets $\Omega = \{\omega_i\}_{i=1}^I$, $\Theta = \{\theta_j\}_{j=1}^J$ with $J < I$ and let $\text{supp}(\gamma) = \{\nu_n\}_{n=1}^N$. By Theorem 2, characterizing the minimum-informative extensions of γ amounts to allocating, for each $\nu \in \text{supp}(\gamma)$, a conditional distribution over Ω given Θ . For any $\mu \in \Delta(\Omega)$, denote by $\mu(\omega_i|\theta_j)$ the conditional probability of ω_i given θ_j . Any sequence of posterior distributions $\{\mu_n\}_{n=1}^N$ constitutes a minimum-informative extension of γ if and only if it satisfies the following conditions:

$$\mu_n(\omega_i) = \sum_{j=1}^J \nu_n(\theta_j) \mu_n(\omega_i|\theta_j), \quad \text{for all } n, \quad (1)$$

$$\mu_0(\omega_i|\theta_j) = \sum_{n=1}^N \frac{\gamma(\nu_n) \nu_n(\theta_j) \mu_n(\omega_i|\theta_j)}{\sum_{n=1}^N \gamma(\nu_n) \nu_n(\theta_j)}, \quad \text{for all } i, j, \quad (2)$$

$$\sum_{i=1}^I \mu_n(\omega_i|\theta_j) = 1, \quad \text{for all } j, n, \quad (3)$$

$$\mu_n(\omega_i|\theta_j) \geq 0, \quad \text{for all } i, j, n. \quad (4)$$

Given any collection $\{\mu_n\}_{n=1}^N$ satisfying these constraints, the associated minimum-informative extension is the distribution τ_γ over $\{\mu_1\}_{n=1}^N$ defined by $\tau_\gamma(\mu_n) = \gamma(\mu_n^\theta) = \gamma(\nu_n)$ for all n . Equation (1) ensures that, for each ν_n , there exists a unique posterior μ_n such that $\mu_n^\theta = \nu_n$. By the construction of τ_γ , it follows immediately that $\tau_\gamma^\theta = \gamma$. Equation (2) guarantees that the induced distribution τ_γ satisfies $\mathbb{E}_{\tau_\gamma}[\mu] = \mu_0$. Finally, Equations (3) and (4) ensure that $\mu_n(\omega_i|\theta_j)$ is a well-defined probability distribution.

⁷Since $\tilde{\theta}$ is a measurable function of Ω , the realization of ω uniquely determines θ . Hence $\mu_n(\omega_i) = \mu_n(\omega_i, \theta_j)$, and $\mu_n(\omega_i, \theta_j) > 0$ only if $\theta_j = \tilde{\theta}(\omega_i)$. Thus, in Equation (1), at most one term in the summation can be positive.

To illustrate how our characterization can be used to describe all privacy-constrained signals, we present a simple example.

Example 3. Consider the following 2×2 setting. The state space $\Omega = X \times \Theta$ consists of a variable of interest $X = \{x_1, x_2\}$ and a privacy $\Theta = \{\theta_1, \theta_2\}$. Define $\bar{\gamma} \in \Gamma \setminus \{\delta_{\mu_0}\}$ by $\bar{\gamma} = \alpha\delta_{\nu_1} + (1 - \alpha)\delta_{\nu_2}$. Now we construct all Blackwell-undominated $\bar{\gamma}$ -privacy-constrained signals (Example 2).

First, by Theorem 2, any minimum-informative extension τ_γ of $\bar{\gamma}$ takes the form $\tau_\gamma = \alpha\delta_{\mu_1} + (1 - \alpha)\delta_{\mu_2}$, where μ_1 and μ_2 satisfy

$$\mu_1((x_i, \theta_j)) = \nu_1(\theta_j)\mu_1(x_i|\theta_j), \quad \mu_2((x_i, \theta_j)) = \nu_2(\theta_j)\mu_2(x_i|\theta_j), \quad \text{for all } i, j.$$

with the pairs $\{\mu_n(x_i|\theta_j)\}_{i,j,n=1}^2$ that

$$\begin{aligned} \mu_2(x_1|\theta_j) &= \mu_0(x_1|\theta_j) + \frac{\alpha\nu_1(\theta_j)}{(1 - \alpha)\nu_2(\theta_j)}[\mu_0(x_1|\theta_j) - \mu_1(x_1|\theta_j)], \quad \text{for all } j, \\ \mu_n(x_1|\theta_j) + \mu_n(x_2|\theta_j) &= 1, \quad \mu_n(x_i|\theta_j) \geq 0, \quad \text{for all } i, j, n. \end{aligned}$$

Second, applying Theorem 1, we construct the Blackwell-undominated conditionally privacy-preserving signals associated with each τ_γ . For each realization $\mu_n \in \tau_\gamma$, Theorem 1 in Strack and Yang (2024) implies that there exists a quantile signal $\tilde{q}_{\mu_n}^*$ such that

$$\tilde{q}_{\mu_n}^*(x_1, \theta_j) \sim U[0, \mu_n(x_1|\theta_j)], \quad \tilde{q}_{\mu_n}^*(x_2, \theta_j) \sim U[\mu_n(x_1|\theta_j), 1],$$

which is Blackwell-undominated since revealing the realization of $\tilde{\theta}$ would allow the receiver to infer the state $\tilde{\omega}$ exactly. By reordering the quantile signal, we can generate all signals that are Blackwell-undominated and conditionally privacy-preserving given μ_n . Specifically, any such reordered quantile signal, denoted by $\tilde{q}_{\mu_n}$, satisfies

$$\tilde{q}_{\mu_n}(x_1, \theta_j) \sim U(I_{\mu_n(x_1|\theta_j)}), \quad \tilde{q}_{\mu_n}(x_2, \theta_j) \sim U([0, 1] \setminus I_{\mu_n(x_1|\theta_j)}),$$

where $I_a \subseteq [0, 1]$ is any measurable subset with Lebesgue measure $\lambda(I_a) = a$.

Let $\tilde{q}_{\tau_\gamma}$ denote the collection of these reordered quantile signals across realizations of τ_γ . The set of all $\tau_\gamma \vee \tilde{q}_{\tau_\gamma}$ then constitutes the Blackwell frontier of $\bar{\gamma}$ -privacy-constrained signals.

Now consider a decision-making problem. For a decision maker who seeks to maximize an objective subject to the γ -privacy constraint, the resulting optimization problem can

be solved in two steps. First, under each minimum-informative extension, one solves an optimal transport problem (Strack and Yang 2024). Second, one optimizes over the set of minimum-informative extensions, which reduces to a linear programming problem. For a general $\mathcal{P}$ -privacy constraint, however, an additional third step is required: one must also optimize over the Blackwell frontier of $\mathcal{P}$.

4 Blackwell Frontier of Privacy-Permissible Set

For a general privacy-permissible set $\mathcal{P}$, characterizing the $\mathcal{P}$ -privacy-constrained signals first requires characterizing the Blackwell frontier of $\mathcal{P}$. In some cases, the Blackwell frontier is explicitly given, as in Example 2. In other cases, such as differential privacy, the set $\mathcal{P}$ is defined by a collection of constraints, and its Blackwell frontier is not directly evident. Unfortunately, there is no uniform method for characterizing the Blackwell frontier of an arbitrary or abstract $\mathcal{P}$. In this section, we focus on two important classes for which the Blackwell frontier can be analyzed: ex-post privacy and posterior-mean privacy.

4.1 Ex-Post Privacy

When the regulator cares about the *ex post* cost of disclosing information about privacy, the constraint applies not to the distribution over posteriors about privacy, but directly to the posteriors themselves. Formally, let

$$\mathcal{M} \subseteq \Delta(\Theta)$$

denote the set of permissible posteriors about privacy. In line with the spirit of Blackwell-closeness in Assumption 1, we assume that $\mathcal{M}$ is a compact convex subset of $\Delta(\Theta)$. If a signal π is permissible, meaning that every posterior $\nu \in \text{supp}(\langle \pi \rangle^\theta)$ lies in $\mathcal{M}$, then any less informative signal $\pi' \preceq \pi$ is also permissible. It follows that every convex combination of posteriors in $\text{supp}(\langle \pi \rangle^\theta)$ must also lie in $\mathcal{M}$. To avoid triviality, we assume that $\mu_0^\theta \in \mathcal{M}$.

Definition 2. A signal π is a $\mathcal{M}$ -ex-post-privacy-constrained signal if $\langle \pi \rangle^\theta(\mathcal{M}) = 1$.

Given $\mathcal{M}$, the induced privacy-permissible set is $\mathcal{P}_\mathcal{M} := \{\gamma \in \mathcal{T} : \gamma(\mathcal{M}) = 1\}$. Since $\mu_0^\theta \in \mathcal{M}$, the set $\mathcal{P}_\mathcal{M}$ is nonempty. Because $\mathcal{M}$ is compact and convex, the set $\mathcal{P}_\mathcal{M}$ satisfies Assumptions 1 and 2. Therefore, $\mathcal{M}$ -ex-post privacy is a special case of $\mathcal{P}$ -privacy.

Let $\text{ext } \mathcal{M}$ as the the set of extreme points of $\mathcal{M}$, i.e,

$$\text{ext } \mathcal{M} := \{\nu \in \mathcal{M} : \nexists \nu' \neq \nu'' \in \mathcal{M}, \alpha \in (0, 1) \text{ such that } \nu = \alpha \nu' + (1 - \alpha) \nu''\}.$$

Let $\overline{\mathcal{P}}_{\mathcal{M}}$ denote the Blackwell frontier of $\mathcal{P}_{\mathcal{M}}$ as defined previously.

Proposition 1 (Characterization of Blackwell Frontier of Ex-Post Privacy-Permissible Set). *A distribution of posteriors about privacy $\gamma \in \Gamma$ belongs to $\overline{\mathcal{P}}_{\mathcal{M}}$ if and only if $\gamma(\text{ext } \mathcal{M}) = 1$.*

Remark 1. *When $\text{ext } \mathcal{M}$ is finite and we restrict attention to distributions γ with finite support, Proposition 1 becomes immediate. The general case requires substantially more work. For the “only if” direction, one must construct a dilation (see footnote 4) to show that any distribution γ' with $\gamma'(\text{ext } \mathcal{M}) < 1$ can be mean-preserving spread to a distribution γ satisfying $\gamma(\text{ext } \mathcal{M}) = 1$. This requires two ingredients. First, for each $\nu \in \mathcal{M}$ there must exist a probability distribution P_{ν} over $\text{ext } \mathcal{M}$ such that $\mathbb{E}_{P_{\nu}}[\nu'] = \nu$, which follows from Choquet’s Theorem (Theorem 10.7, p.168, [Simon \(2011\)](#)). Second, the mapping $\nu \mapsto P_{\nu}$ from $\mathcal{M}$ to $\Delta(\text{ext } \mathcal{M})$ must be measurable. This measurability is ensured by Theorem 9.1 in [Simon \(2011\)](#) (p.136), together with the Kuratowski–Ryll–Nardzewski measurable selection theorem (Theorem 6.9.3, p.36, Vol.II, [Bogachev \(2007\)](#)).*

By Proposition 1, under $\mathcal{M}$ -ex-post privacy the Blackwell frontier of permissible distributions over posteriors about privacy, $\overline{\mathcal{P}}_{\mathcal{M}}$, can be generated by first computing the extreme points of the permissible posterior set $\mathcal{M}$. Any distribution in $\overline{\mathcal{P}}_{\mathcal{M}}$ is then obtained by convex combinations of these extreme points with respect to the prior μ_0^{θ} .

Ex-post privacy encompasses all privacy notions that impose constraints directly on posterior beliefs about privacy and that satisfy Blackwell-closeness. When privacy realizations are finite and the constraints imposed on posteriors are finite and linear, the set $\mathcal{M}$ becomes a convex polytope whose extreme points are finite and can be computed explicitly. For example, the well-known concept of differential privacy introduced by [Dwork et al. \(2006\)](#) is defined by a finite collection of linear constraints on posterior beliefs about privacy. This concept has been adopted by major institutions such as Google, Microsoft, and the U.S. Census Bureau ([Abowd 2018](#)). The characterization of Blackwell frontier of differential privacy has been conducted by [Xu and Zhao \(2025\)](#).

Similar with differential privacy, [Ghosh and Kleinberg \(2016\)](#) introduces the notion of inferential privacy in the context of vector-valued datasets. [Wang et al. \(2025\)](#) adapts this concept to the framework of [He et al. \(2021\)](#), which corresponds to the discrete case of our setting. Our Definition 3 generalizes inferential privacy to arbitrary standard Borel spaces.

Definition 3. For any $\varepsilon \in [0, +\infty)$. A signal π is ε -inferential-privacy-constrained if for almost every $\nu \in \text{supp}(\langle \pi \rangle^\theta)$ and $B', B'' \in \mathcal{B}(\Theta)$ such that $\mu_0(B') > 0$, $\mu_0(B'') > 0$,

$$\frac{\nu(B')}{\nu(B'')} \leq e^\varepsilon \cdot \frac{\nu(B')}{\nu(B'')}. \quad (5)$$

Let $\mathcal{I}$ denote the set of posteriors that satisfy the ε -inferential privacy constraint (5), $\mathcal{P}_{\mathcal{I}}$ the privacy-permissible set induced by $\mathcal{I}$, and $\overline{\mathcal{P}}_{\mathcal{I}}$ its Blackwell frontier.

Proposition 2. A distribution of posteriors about privacy $\gamma \in \Gamma$ belongs to $\overline{\mathcal{P}}_{\mathcal{I}}$ if and only if for almost every $\nu \in \text{supp}(\gamma)$, there is a subset $E_\nu \in \mathcal{B}(\Theta)$ such that $\mu_0^\theta(E_\nu) \in (0, 1)$, and

$$\nu(B) = \frac{e^\varepsilon \mu_0^\theta(B \cap E_\nu) + \mu_0^\theta(B \setminus E_\nu)}{e^\varepsilon \mu_0^\theta(E_\nu) + (1 - \mu_0^\theta(E_\nu))}, \quad (6)$$

for almost every $B \in \mathcal{B}(\Theta)$.

Proposition 2 provides a clean characterization of the extreme points under ε -inferential privacy, a setting with infinitely many privacy realizations and linear constraints. An extreme point of $\mathcal{I}$ partitions the privacy realization space into two sets of positive measure, E and $\Theta \setminus E$. Relative to the prior, the posterior probability assigned to E is increased uniformly, while the posterior probability assigned to $\Theta \setminus E$ is decreased uniformly. Using Proposition 2, we can simplify the main characterizations of ε -inferential-private private information structure presented in Wang et al. (2025).

4.2 Posterior-Mean Privacy

In some settings, an individual's privacy loss depends only on the posterior expected type rather than on the full posterior distribution. In other words, the privacy cost associated with each posterior belief is a linear function of its realization. In this case, a natural privacy constraint is to impose an upper bound on the distribution of posterior means of a statistic of the privacy.

Specially, let $\tilde{f} : \Theta \rightarrow \mathbb{R}$ be a one-dimensional statistic of privacy. Denote by ν_0^f the prior distribution of $\tilde{f}$, and suppose that $\bar{\kappa}$ is a mean-preserving contraction of ν_0^f . A signal π is f -posterior-mean-privacy-constrained (with respect to $\bar{\kappa}$) if the distribution over posterior means of $\tilde{f}$ is a mean-preserving contraction of $\bar{\kappa}$. For any $\gamma \in \Gamma$, denote the induced

distribution of posterior mean about $\tilde{f}$ by

$$\kappa_\gamma(\cdot) := \gamma(\{\nu \in \text{supp}(\gamma) : \mathbb{E}_\gamma[\tilde{f}(\theta)] \in \cdot\}).$$

When $\Theta \subseteq \mathbb{R}$, the statistic $\tilde{f}$ captures all moments of $\tilde{\theta}$. Since the distribution of posterior means of $\tilde{f}$ depends on the distribution of posteriors about privacy, posterior-mean privacy is not an ex-post privacy.

Definition 4. A signal π is a f -posterior-mean-privacy-constrained signal if $\kappa_{\langle \pi \rangle^\theta} \preceq \bar{\kappa}$.

Let $\mathcal{E}$ denote the set of posteriors that satisfy the f -posterior-mean privacy, $\mathcal{P}_\mathcal{E}$ the privacy-permissible set induced by $\mathcal{E}$, and $\overline{\mathcal{P}}_\mathcal{E}$ its Blackwell frontier.

Proposition 3. A distribution of posteriors about privacy $\gamma \in \Gamma$ belongs to $\overline{\mathcal{P}}_\mathcal{E}$ if and only if (1) for almost every $\nu \in \text{supp}(\gamma)$, there exists $y_1, y_2 \in \tilde{f}(\Theta)$ and $\alpha \in (0, 1]$, ν puts α on a point $\theta_1 \in \tilde{f}^{-1}(y_1)$ and $(1 - \alpha)$ on another point $\theta_2 \in \tilde{f}^{-1}(y_2)$ and (2) $\kappa_\gamma = \bar{\kappa}$.

Proposition 3 characterizes the Blackwell frontier of the privacy-permissible set induced by posterior-mean privacy. When $\gamma \in \overline{\mathcal{P}}_\mathcal{E}$, every realized posterior about privacy is a two-point distribution, and the induced distribution over posterior mean about $\tilde{f}$ attains the upper bound.

5 Discussion and Future Work

We provide a characterization of signals that, in the Blackwell sense, do not reveal more private information than those in a given privacy-permissible set $\mathcal{P}$, where $\mathcal{P}$ is a subset of distributions over posterior beliefs about privacy. Specifically, we show that the most informative $\mathcal{P}$ -privacy-constrained signals can be constructed as the join of two components: (i) a minimum-informative extension of a distribution on the Blackwell frontier of $\mathcal{P}$, and (ii) a Blackwell-undominated, conditionally privacy-preserving signal as characterized by [Strack and Yang \(2024\)](#). We then characterize the Blackwell frontier of the privacy-permissible set, with particular attention to ex-post privacy and posterior-mean privacy.

This paper has several limitations. First, we do not provide a general method for identifying the Blackwell frontier of an arbitrary privacy-permissible set. Our current results apply only to two specific formulations of privacy. Extending these results to more general notions of privacy remains an important direction for future research.

Second, for the general privacy constraint, when there is a unique bound on privacy information, our approach reduces the decision-making problem to a first-stage optimal transport problem followed by a second-stage linear program. Although this method is conceptually clean, it is computationally demanding. Developing a more tractable solution approach remains an important direction for future research.

References

- Abowd, J. M. (2018). The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2867–2867.
- Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, Volume 1, pp. 26.
- Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics*, 265–272.
- Bogachev, V. I. (2007). *Measure theory*. Springer.
- Çinlar, E. (2011). *Probability and stochastics*. Springer.
- Doob, J. L. (1951). Continuous parameter martingales. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Volume 2, pp. 267–276. University of California Press.
- Dwork, C., F. McSherry, K. Nissim, and A. Smith (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284.
- Eilat, R., K. Eliaz, and X. Mu (2021). Bayesian privacy. *Theoretical Economics* 16(4), 1557–1603.
- Ghosh, A. and R. Kleinberg (2016). Inferential privacy guarantees for differentially private mechanisms. *arXiv preprint arXiv:1603.01508*.
- He, K., F. Sandomirskiy, and O. Tamuz (2021). Private private information. *arXiv preprint arXiv:2112.14356*.
- Kallenberg, O. (1997). *Foundations of modern probability*, Volume 2. Springer.

- Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review* 101(6), 2590–2615.
- Royden, H. and P. Fitzpatrick (2010). *Real Analysis*. Prentice Hall.
- Simon, B. (2011). *Convexity: an analytic viewpoint*, Volume 187. Cambridge University Press.
- Strack, P. and K. H. Yang (2024). Privacy-preserving signals. *Econometrica* 92(6), 1907–1938.
- Strassen, V. (1965). The existence of probability measures with given marginals. *The Annals of Mathematical Statistics* 36(2), 423–439.
- Wang, S., S. Zheng, Z. Lin, G. Fanti, and Z. S. Wu (2025). Inferentially-private private information. In *Proceedings of the ACM on Web Conference 2025*, pp. 2579–2595.
- Xu, Z. and W. Zhao (2025). Privacy structure and blackwell frontier. *arXiv preprint arXiv:2511.10226*.

A Appendix

A.1 Proofs for Section 3

Lemma 2. *For two signals π, π' , if $\pi \succeq \pi'$, then $\pi \succeq_\theta \pi'$.*

Proof. Since $\pi \succeq \pi'$, then there exists a dilation $K : \Delta(\Omega) \rightarrow \Delta(\Delta(\Omega))$, such that for almost every $\mu_{s'} \in \text{supp}(\langle \pi' \rangle)$, $\mu_{s'} = \int_{\Delta(\Omega)} \mu_s dK(\mu_s | \mu_{s'})$, and $\langle \pi \rangle(\cdot) = \int_{\Delta(\Omega)} K(\cdot | \mu_{s'}) d\langle \pi' \rangle(\mu_{s'})$ (see footnote 4). Define $Q : \Delta(\Theta) \rightarrow \Delta(\Delta(\Theta))$, such that for $\langle \pi' \rangle^\theta$ -almost every $\nu' \in \Delta(\Theta)$ and all $B \in \mathcal{B}(\Delta(\Theta))$,

$$Q(B | \nu') = \mathbb{E}_{\langle \pi' \rangle} [K(\{\mu \in \Delta(\Omega) : \mu^\theta \in B\} | \mu_{s'}) | \mu_{s'}^\theta = \nu'] .$$

Then, for almost every $\nu' \in \text{supp}(\langle \pi' \rangle^\theta)$,

$$\begin{aligned} \nu' &= \mathbb{E}_{\langle \pi' \rangle} [\mu_{s'}^\theta | \mu_{s'}^\theta = \nu'] \\ &= \mathbb{E}_{\langle \pi' \rangle} \left[\int_{\Delta(\Omega)} \mu_s^\theta dK(\mu_s | \mu_{s'}) | \mu_{s'}^\theta = \nu' \right] \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}_{\langle \pi' \rangle} \left[\int_{\Delta(\Theta)} \nu dK(\{\mu \in \Delta(\Omega) : \mu^\theta = \nu\} | \mu_{s'}) | \mu_{s'}^\theta = \nu' \right] \\
&= \int_{\Delta(\Theta)} \nu \mathbb{E}_{\langle \pi' \rangle} [dK(\{\mu \in \Delta(\Omega) : \mu^\theta = \nu\} | \mu_{s'}) | \mu_{s'}^\theta = \nu'] \\
&= \int_{\Delta(\Theta)} \nu dQ(\nu | \nu').
\end{aligned}$$

For almost every $B \in \mathcal{B}(\Delta(\Theta))$,

$$\begin{aligned}
\langle \pi \rangle^\theta(B) &= \langle \pi \rangle(\{\mu \in \Delta(\Omega) : \mu^\theta \in B\}) \\
&= \mathbb{E}_{\langle \pi' \rangle} [K(\{\mu \in \Delta(\Omega) : \mu^\theta \in B\} | \mu_{s'})] \\
&= \int_{\Delta(\Theta)} \mathbb{E}_{\langle \pi' \rangle} [K(\{\mu \in \Delta(\Omega) : \mu^\theta \in B\} | \mu_{s'}) | \mu_{s'}^\theta = \nu'] d\langle \pi' \rangle^\theta(\nu') \\
&= \int_{\Delta(\Theta)} Q(B | \nu') d\langle \pi' \rangle(\nu').
\end{aligned}$$

Therefore, $\pi \succeq_\theta \pi'$. □

For two signals π and π' , we say that π is *sufficient* for π' if $\pi \sim \pi \vee \pi'$. It is equal to say that given π , π' is independent with $\tilde{\omega}$, i.e., $\mathbb{P}(\pi' \in \cdot | \omega, s) = \mathbb{P}(\pi' \in \cdot | s)$ almost surely.⁸ Furthermore, if $\pi \succeq \pi'$, there exists $\pi'' \sim \pi$ such that π'' is sufficient for π' (Blackwell (1951), Theorem 6).⁹ If $\pi \succeq_\theta \pi'$, then there exists $\pi'' \sim_\theta \pi$ such that π'' is sufficient for π' in terms of $\tilde{\theta}$, i.e., $\pi'' \sim_\theta \pi'' \vee \pi'$.¹⁰

⁸By the definition, $\pi \sim \pi \vee \pi'$ says that $\mathbb{P}(\tilde{\omega} \in \cdot | s, s') = \mathbb{P}(\tilde{\omega} \in \cdot | s)$ almost surely. According to Theorem 7 in Blackwell (1951), it is equivalent to $\mathbb{P}(\pi' \in \cdot | \omega, s) = \mathbb{P}(\pi' \in \cdot | s)$ almost surely.

⁹The key distinction between Blackwell domination and sufficiency is that sufficiency requires the process from $\tilde{\mu}_{\pi'}$ to $\tilde{\mu}_\pi$ to form a martingale. Blackwell domination, on the other hand, does not impose this martingale requirement, allowing for more general transformations between signals. To understand it, consider the following example:

Under signal π , the posterior belief about $\tilde{\omega}$ takes the value 0 (i.e, a distribution putting 1 at $\omega = 0$) with prob 1/4, 1/2 with prob 1/2, and 1 with prob 1/4. Under signal π' , $\tilde{\omega} = 1/4$ with prob 1/2 and 3/4 with prob 1/2. Sufficiency implies that, conditional on the belief $\tilde{\omega} = 1/4$ being realized after observing π' , the updated belief after observing π should be $\tilde{\omega} = 0$ with prob 1/2 and $\tilde{\omega} = 1/2$ with prob 1/2. In contrast, Blackwell domination allows for different mappings. For example, conditional on $\tilde{\omega} = 1/4$ after observing π' , the belief after π could be $\tilde{\omega} = 1/2$ with prob 1/2 and $\tilde{\omega} = 1$ with prob 1/2.

However, we can construct another signal, π'' , by applying the dilation K (as defined in footnote 4) to split π' . By construction, π'' is sufficient for π' , and $\pi'' \sim \pi$.

¹⁰We can extend a Markov kernel $Q : \Delta(\Theta) \rightarrow \Delta(\Delta(\Theta))$ to a Markov kernel $K : \Delta(\Omega) \rightarrow \Delta(\Delta(\Omega))$ by the following construction: for any $\mu' \in \Delta(\Omega)$ with marginal $\mu'^\theta = \nu'$, define $K(\nu'' \otimes \mu'(\tilde{\omega} | \tilde{\theta}) | \mu') = Q(\nu'' | \nu')$, where $\mu'(\tilde{\omega} | \tilde{\theta})$ is the conditionally distribution of $\tilde{\omega}$ given $\tilde{\theta}$ induced by μ' . In other words, $K(\cdot | \mu')$ is defined as the pushforward of $Q(\cdot | \nu')$ under the map $\nu'' \mapsto \nu'' \otimes \mu'(\cdot | \tilde{\theta})$, which reconstructs a measure on Ω from its marginal ν'' on Θ and the conditionals $\mu'(\cdot | \tilde{\theta})$.

Proof of Lemma 1. “If”. Suppose $\pi \preceq \pi'$ where $\langle \pi' \rangle^\theta \in \overline{\mathcal{P}}$. According to Lemma 2, $\langle \pi \rangle^\theta \preceq \langle \pi' \rangle^\theta$. Therefore by Assumption 1, $\pi \in \Pi_{\mathcal{P}}$.

“Only if”. Suppose $\langle \pi \rangle^\theta \notin \overline{\mathcal{P}}$; otherwise the claim is trivial since $\pi \preceq \pi$. Then, there exists a signal π' with $\langle \pi' \rangle^\theta \in \overline{\mathcal{P}}$ such that π' is sufficient for π in terms of $\tilde{\theta}$. Then, since $\langle \pi \vee \pi' \rangle^\theta = \langle \pi' \rangle^\theta \in \mathcal{P}$, $\pi \vee \pi'$ is a $\mathcal{P}$ -privacy-constrained signal. We get $\pi \preceq \pi \vee \pi'$. $\square$

Lemma 3. *For any signal π with $\langle \pi \rangle^\theta \in \mathcal{P}$, $\pi \in \overline{\Pi}_{\mathcal{P}}$ if and only if there does not exist a conditionally (given π) privacy-preserving signal $\tilde{q}$ such that $\pi \prec \pi \vee \tilde{q}$.*

Proof. “Only if”. Suppose there exists such signal $\tilde{q}$. Since $\tilde{q}$ is conditionally privacy-preserving given π , $\pi \vee \tilde{q}$ is $\mathcal{P}$ -privacy-constrained. $\pi \prec \pi \vee \tilde{q}$ contradicts $\pi \in \overline{\Pi}_{\mathcal{P}}$.

“If”. Suppose there exists a $\mathcal{P}$ -privacy-constrained signal π' such that $\pi \prec \pi'$. Since $\langle \pi \rangle^\theta \in \mathcal{P}$, $\langle \pi' \rangle^\theta = \langle \pi \rangle^\theta$. W.l.o.g, assume π' is sufficient for π , i.e., $\pi \vee \pi' \sim \pi'$. Since $\langle \pi \vee \pi' \rangle^\theta = \langle \pi' \rangle^\theta = \langle \pi \rangle^\theta$, π' is a conditionally (given π) privacy-preserving signal such that $\pi \prec \pi \vee \pi'$. $\square$

Proof of Theorem 1. “If”. Let $\tau_\gamma \in \underline{\mathcal{T}}_\gamma$ for some $\gamma \in \overline{\mathcal{P}}$ and $\tilde{q}$ denote a signal which is Blackwell-undominated conditionally privacy-preserving given τ_γ . Suppose $\tau_\gamma \vee \tilde{q} \notin \overline{\Pi}_{\mathcal{P}}$. Following Lemma 3, there exists another conditionally (given $\tau_\gamma \vee \tilde{q}$) privacy-preserving signal $\tilde{q}'$ such that $\tau_\gamma \vee \tilde{q} \prec \tau_\gamma \vee \tilde{q} \vee \tilde{q}'$. Since $\langle \tau_\gamma \vee \tilde{q} \vee \tilde{q}' \rangle^\theta = \langle \tau_\gamma \vee \tilde{q} \rangle^\theta = \tau_\gamma^\theta$, $\tilde{q} \vee \tilde{q}'$ is also a conditionally (given τ_γ) privacy-preserving signal. Therefore, $\tau_\gamma \vee \tilde{q} \prec \tau_\gamma \vee (\tilde{q} \vee \tilde{q}')$ contradicts that $\tilde{q}$ is Blackwell-undominated among all conditionally privacy-preserving signals given τ_γ .

“Only if”. For any $\pi \in \overline{\Pi}_{\mathcal{P}}$, if $\pi \notin \underline{\mathcal{T}}_{\langle \pi \rangle^\theta}$, then there exists another signal $\pi_1 \in \Pi_{\mathcal{P}}$ such that $\langle \pi_1 \rangle^\theta = \langle \pi \rangle^\theta$ and π is sufficient of π_1 . If $\pi_1 \notin \underline{\mathcal{T}}_{\langle \pi \rangle^\theta}$, we can find another signal $\pi_2 \in \Pi_{\mathcal{P}}$ such that $\langle \pi_2 \rangle^\theta = \langle \pi \rangle^\theta$ and π_1 is sufficient of π_2 . Continuing this process, if it terminates in a finite number of steps, we will eventually obtain a signal $\pi_N \in \underline{\mathcal{T}}_{\langle \pi \rangle^\theta}$. Otherwise, we construct a sequence $\{\tilde{\mu}_{\pi_t}\}_{t \in \mathbb{N}_+}$, where $\tilde{\mu}_{\pi_t}$ is random variable about posterior belief induced by π_t . According to Section 4 in Blackwell (1953), the sequence $\{\tilde{\mu}_{\pi_t}\}_{t \in \mathbb{N}}$ forms a reverse martingale. By the martingale convergence theorem (Doob 1951), $\tilde{\mu}_{\pi_t} \rightarrow \tilde{\mu}^*$ as $t \rightarrow \infty$. Let π_∞ be one of the signals that induces $\tilde{\mu}^*$. Then, we have $\pi_\infty \in \underline{\mathcal{T}}_{\langle \pi \rangle^\theta}$. Define π^* as a unified notation that refers to either π_N or π_∞ . Since π is conditionally privacy-preserving signal given π^* . Therefore, there is a $\langle \pi \rangle^\theta \in \overline{\mathcal{P}}$ (Lemma 1), $\pi^* \in \mathcal{T}$ and π is a corresponding Blackwell-undominated conditional privacy-preserving signal such that $\pi \sim \pi^* \vee \pi$. $\square$

Proof of Theorem 2. “Only if”. If condition (1) is not satisfied, by definition $\underline{\tau} \notin \underline{\mathcal{T}}_\gamma$. Now, suppose condition (2) is not satisfied. Then, there exists a positive measure set $B \in \text{supp}(\underline{\tau}^\theta)$

such that for any $\mu \in \text{supp}(\underline{\tau})$ with $\mu^\theta \in B$, $Q_{\underline{\tau}}(\mu|\mu^\theta) < 1$, where $Q_{\underline{\tau}}$ denotes the conditional distribution over $\Delta(\Omega)$ given $\Delta(\Theta)$ induced by $\underline{\tau}$. By construction, $Q_{\underline{\tau}}$ is non-degenerate. We can then construct a new distribution over posteriors, τ' , such that $\tau'^\theta = \underline{\tau}^\theta$ and for all $\mu^\theta \in \text{supp}(\underline{\tau}^\theta)$, in the new distribution τ' , the conditional probability on $\hat{\mu} = \mathbb{E}_{\underline{\tau}}(\{\mu \in \Delta(\Omega) : \mu^\theta = \nu\}|\nu)$ is 1 given $\nu \in \text{supp}(\underline{\tau}^\theta)$. $\underline{\tau}$ is a strict mean-preserving spread of τ' , since $\underline{\tau}(\cdot) = \int_{\Delta(\Omega)} Q_{\underline{\tau}}(\cdot|\mu'^\theta) d\tau'(\mu')$, and $Q_{\underline{\tau}}(\cdot|\mu'^\theta)$ is non-degenerate. $\tau' \prec \underline{\tau}$ contradicts $\underline{\tau} \in \mathcal{T}_\gamma$.

“If”. Suppose there exists a signal τ' such that $\tau'^\theta \in \overline{\mathcal{P}}$ and $\tau' \prec \underline{\tau}$. There is a non-degenerate dilatation $K : \Delta(\Omega) \rightarrow \Delta(\Delta(\Omega))$ spreads τ' to $\underline{\tau}$. Moreover, since $\underline{\tau}$ satisfies condition 2, there is a one-to-one mapping from $\nu \in \text{supp}(\underline{\tau}^\theta)$ to $\mu \in \text{supp}(\underline{\tau})$. Therefore, the kernel $Q : \Delta(\Theta) \rightarrow \Delta(\Delta(\Theta))$ defined in Lemma 2 is non-degenerate if K is non-degenerate. As a result, $\tau'^\theta \prec \underline{\tau}^\theta$, which contradicts $\tau'^\theta \in \overline{\mathcal{P}}$. $\square$

A.2 Proofs for Section 4

Proof of Proposition 1. “If”. Suppose there exists another $\gamma' \in \mathcal{P}_\mathcal{M}$ such that $\gamma \prec \gamma'$, then there exists a nondegenerate dilation $K : \Delta(\Theta) \rightarrow \Delta(\Delta(\Theta))$ such that for almost every $\nu \in \text{supp}(\gamma)$,

$$\nu = \int_{\Delta(\Theta)} \nu' K(d\nu'|\nu).$$

This means that for almost every $\nu \in \text{supp}(\gamma)$, it can be expressed by linear combination of $\nu' \in \text{supp}(\gamma')$. Since K is nondegenerate, there exists a positive measure subset $A \subseteq \text{supp}(\gamma)$ such that $K(\nu|\nu) < 1$ for any $\nu \in A$. Hence, $\nu \notin \text{ext } \mathcal{M}$ for any $\nu \in A$, which is contradicted with $\gamma(\text{ext } \mathcal{M}) = 1$.

“Only if”. We construct a dilation from $\mathcal{M}$ to $\Delta(\text{ext } \mathcal{M})$. Since $(\Theta, \mathcal{B}(\Theta))$ is a standard Borel space, $\Delta(\Theta)$ embeds into a locally convex space and endowed with the topology of weak convergence is metrizable. $\mathcal{M}$ is a compact convex subset of $\Delta(\Theta)$, which is also metrizable. By Choquet’s Theorem (Theorem 10.7, p.168, Simon (2011)), for any $\nu \in \mathcal{M}$, the set

$$\Phi(\nu) := \left\{ P_\nu \in \Delta(\text{ext } \mathcal{M}) : \nu = \int \nu' dP_\nu(\nu') \right\}$$

is nonempty. Moreover, $\Phi(\mu)$ is closed in the weak-* topology. Define the barycenter map $B : \Delta(\mathcal{M}) \rightarrow \mathcal{M}$ by $B(P_\nu) = \int \nu' dP_\nu = \nu$. By Simon (2011), Theorem 9.1 (p.136), the map B is continuous. Consequently, for any open set $U \subseteq \Delta(\text{ext } \mathcal{M})$, $\Phi^{-1}(U) = \{\nu \in \mathcal{M} : B^{-1}(\nu) \cap U \neq \emptyset\} = B(U)$, which is an open set. Therefore, by the Kuratowski–Ryll–Nardzewski measurable selection theorem (Theorem 6.9.3, p.36, Vol. II, Bogachev (2007)),

there exists a measurable selection $P_\nu^* : \mathcal{M} \rightarrow \Delta(\text{ext } \mathcal{M})$ such that $\nu = \int \nu' dP_\nu^*(\nu')$. Hence, the map $K : \nu \mapsto P_\nu^*$ defines a dilation. If $\gamma(\text{ext } \mathcal{M}) \neq 1$, then K is nondegenerate, which implies $\gamma \notin \overline{\mathcal{P}}_{\mathcal{M}}$. $\square$

Lemma 4. *Suppose $\gamma \in \mathcal{P}_{\mathcal{I}}$. For almost every $\nu \in \text{supp}(\gamma)$, there is a nonnegative measurable function g_ν defined for which*

$$\nu(B) = \int_B g_\nu d\mu_0^\theta, \quad \text{for all } B \in \mathcal{B}(\Theta). \quad (7)$$

and g_ν is essential bounded. Let $e^{\bar{\varepsilon}_\nu}$ as the essential supremum of g_ν , i.e.,

$$\begin{aligned} \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) > e^{\bar{\varepsilon}_\nu}\}) &= 0, \\ \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) > e^{\bar{\varepsilon}_\nu} - \delta\}) &> 0, \quad \text{for all } \delta > 0. \end{aligned} \quad (8)$$

then, $\bar{\varepsilon}_\nu < \varepsilon$ and

$$\mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [e^{\bar{\varepsilon}_\nu - \varepsilon}, e^{\bar{\varepsilon}_\nu}]\}) = 1. \quad (9)$$

Moreover, if $g_\nu \neq 1$ μ_0^θ -almost surely, then $\bar{\varepsilon}_\nu > 0$ and

$$\begin{aligned} \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [e^{\bar{\varepsilon}_\nu - \varepsilon}, 1]\}) &> 0, \\ \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in (1, e^{\bar{\varepsilon}_\nu}]\}) &> 0. \end{aligned} \quad (10)$$

Proof. Suppose $\gamma \in \mathcal{P}_{\mathcal{I}}$. Radon-Nikodym Theorem ([Royden and Fitzpatrick \(2010\)](#), p.382) shows that for almost every $\nu \in \text{supp}(\gamma)$, g_ν defined by (7) exists. By the inferential-privacy constraint (5), for all $B \in \mathcal{B}(\Theta)$ with $\mu_0^\theta(B) > 0$,

$$e^{-\varepsilon} \frac{\mu_0^\theta(B)}{\mu_0^\theta(\Theta)} \leq \frac{\nu(B)}{\nu(\Theta)} \leq e^\varepsilon \frac{\mu_0^\theta(B)}{\mu_0^\theta(\Theta)} \Rightarrow e^{-\varepsilon} \mu_0^\theta(B) \leq \nu(B) \leq e^\varepsilon \mu_0^\theta(B),$$

which is due to the fact that $\mu_0^\theta(\Theta) = \nu(\Theta) = 1$. Therefore,

$$\mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [e^{-\varepsilon}, e^\varepsilon]\}) = 1.$$

g_ν is essentially bounded below and above. Since the completeness of $\mathbb{R}$, g_ν has essential supremum and infimum, denoted as $e^{\bar{\varepsilon}_\nu}$ and $e^{\underline{\varepsilon}_\nu}$, respectively.

Since the fact that $\nu(\Theta) = \nu_0^\theta(\Theta)$, if

$$\begin{aligned} \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [e^{\underline{\varepsilon}_\nu}, 1]\}) &> 0, \text{ then} \\ \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in (1, e^{\bar{\varepsilon}_\nu}]\}) &> 0. \end{aligned}$$

Hence, if $g_\nu \neq 1$ μ_0^θ -almost surely, (10) holds and $\underline{\varepsilon}_\nu < 0 < \bar{\varepsilon}_\nu$. We need only to show that $\bar{\varepsilon}_\nu < \varepsilon$ and $\bar{\varepsilon}_\nu - \underline{\varepsilon}_\nu \leq \varepsilon$, which induce (9).

Suppose $\bar{\varepsilon}_\nu - \underline{\varepsilon}_\nu > \varepsilon$, then for a constant number $0 < \Delta\varepsilon < \frac{1}{2}(\bar{\varepsilon}_\nu - \underline{\varepsilon}_\nu - \varepsilon)$, we have $(\bar{\varepsilon}_\nu - \Delta\varepsilon) - (\underline{\varepsilon}_\nu + \Delta\varepsilon) > \varepsilon$. Since the definition of essential supremum and infimum, define two measurable sets

$$B_1 := \{\theta \in \Theta : g_\nu(\theta) < e^{\underline{\varepsilon}_\nu + \Delta\varepsilon}\},$$

$$B_2 := \{\theta \in \Theta : g_\nu(\theta) > e^{\bar{\varepsilon}_\nu - \Delta\varepsilon}\},$$

$\mu_0^\theta(B_1) > 0$ and $\mu_0^\theta(B_2) > 0$. Obviously,

$$\frac{\nu(B_2)}{\nu(B_1)} > e^{\bar{\varepsilon}_\nu - \Delta\varepsilon - (\underline{\varepsilon}_\nu + \Delta\varepsilon)} \frac{\mu_0^\theta(B_2)}{\mu_0^\theta(B_1)} > e^\varepsilon \frac{\mu_0^\theta(B_2)}{\mu_0^\theta(B_1)},$$

contradicts with (5).

If $g_\nu = 1$ μ_0^θ -almost surely, then $\bar{\varepsilon}_\nu = 0 < \varepsilon$. Otherwise, since above we show that $\underline{\varepsilon}_\nu < 0 < \bar{\varepsilon}_\nu$ and $\bar{\varepsilon}_\nu - \underline{\varepsilon}_\nu \leq \varepsilon$, then $\bar{\varepsilon}_\nu < \varepsilon$. $\square$

Proof of Proposition 2. “If”. We will show that for almost every $\nu \in \text{supp}(\gamma)$, there does not exist non-degenerated $K_\nu \in \Delta(\Delta(\Theta))$ such that $\nu = \int_{\Delta(\Theta)} \nu' K_\nu(d\nu')$ and inferential-privacy constraint (5) holds almost everywhere on $\text{supp}(K_\nu)$. This statement indicates that $\gamma \in \overline{\mathcal{P}}_{\mathcal{I}}$.

Suppose there is a non-degenerated K_ν . Since K_ν is non-degenerated, there is a positive measurable subset $M \in \text{supp}(K_\nu)$ and, w.l.o.g, a subset $B \subseteq E_\nu$ such that for all $\nu' \in M$, $\nu'(B) > \nu(B)$. Because of the inferential-privacy constraint (5), $\frac{\nu'(B)}{\nu'(E_\nu^c)} \leq \frac{\nu(B)}{\nu(E_\nu^c)} = e^\varepsilon \frac{\mu_0^\theta(B)}{\mu_0^\theta(E_\nu^c)}$, then $\nu'(E_\nu^c) > \nu(E_\nu^c)$. Since the mean-preserving condition, there is another positive measurable subset $M' \in \text{supp}(K_\mu)$ such that $\nu''(E_\nu^c) < \nu(E_\nu^c)$ holds for all $\nu'' \in M'$. Again, due to the constraint (5), $\frac{\nu''(E_\nu)}{\nu''(E_\nu^c)} \leq \frac{\nu(E_\nu)}{\nu(E_\nu^c)} = e^\varepsilon \frac{\mu_0^\theta(E_\nu)}{\mu_0^\theta(E_\nu^c)}$, then $\nu''(E_\nu) < \nu(E_\nu)$. Therefore $\nu''(\Theta) = \nu''(E_\nu) + \nu''(E_\nu^c) < \nu(E_\nu) + \nu(E_\nu^c) = 1$ which contradicts with the fact that ν'' is a probability.

“Only if”. Suppose a γ' for which there is a positive measurable subset $M \in \text{supp}(\gamma')$ and for each $\nu \in M$, there is a positive measurable subset $F \subseteq \Theta$ and some $\varepsilon' \in (0, \varepsilon)$ such that for any $B_1 \in \mathcal{B}(F)$ and $B_2 \in \mathcal{B}(\Theta)$ with $\mu_0^\theta(B_1) > 0$ and $\mu_0^\theta(B_2) > 0$,

$$e^{-\varepsilon'} \frac{\mu_0^\theta(B_1)}{\mu_0^\theta(B_2)} \leq \frac{\nu(B_1)}{\nu(B_2)} \leq e^{\varepsilon'} \frac{\mu_0^\theta(B_1)}{\mu_0^\theta(B_2)} \quad \mu_0^\theta\text{-almost surely.} \quad (11)$$

Then, for a positive constant $\delta < \min\{e^{\varepsilon - \varepsilon'} - 1, 1/\nu(F)\}$, ν can split into ν_1 with probability

$\frac{1}{2}(1 + \delta\mu(F))$ and ν_2 with probability $\frac{1}{2}(1 - \delta\mu(F))$, where $\nu_1(\cdot) := \frac{(1+\delta)\nu(\cdot \cap F)}{1+\delta\nu(F)} + \frac{\nu(\cdot \setminus F)}{1+\delta\nu(F)}$, $\nu_2(\cdot) := \frac{(1-\delta)\nu(\cdot \cap F)}{1-\delta\nu(F)} + \frac{\nu(\cdot \setminus F)}{1-\delta\nu(F)}$. Since for $B_3, B_4 \in \mathcal{B}(\Theta)$ with $\mu_0^\theta(B_4) > 0$, almost surely,

$$\begin{aligned} \frac{\nu_1(B_3)}{\nu_1(B_4)} &= \frac{\nu_1(B_3 \cap F)}{\nu_1(B_4)} + \frac{\nu_1(B_3 \setminus F)}{\nu_1(B_4)} \leq \frac{\nu(B_3 \cap F)(1 + \delta)}{\nu(B_4)} + \frac{\nu(B_3 \setminus F)}{\nu(B_4)} \\ &\leq (1 + \delta)e^{\varepsilon'} \frac{\nu_0^\theta(B_3 \cap F)}{\nu_0^\theta(B_4)} + e^\varepsilon \frac{\nu_0^\theta(B_3 \setminus F)}{\nu_0^\theta(B_4)} \leq e^\varepsilon \frac{\nu_0^\theta(B_3)}{\mu_0^\theta(B_4)}, \\ \frac{\nu_2(B_3)}{\nu_2(B_4)} &= \frac{\nu_2(B_3 \cap F)}{\nu_2(B_4)} + \frac{\nu_2(B_3 \setminus F)}{\nu_2(B_4)} \geq \frac{\nu(B_3 \cap F)(1 - \delta)}{\nu(B_4)} + \frac{\nu(B_3 \setminus F)}{\nu(B_4)} \\ &\geq (1 - \delta)e^{-\varepsilon'} \frac{\mu_0^\theta(B_3 \cap F)}{\mu_0^\theta(B_4)} + e^{-\varepsilon} \frac{\mu_0^\theta(B_3 \setminus F)}{\mu_0^\theta(B_4)} \geq e^{-\varepsilon} \frac{\mu_0^\theta(B_3)}{\mu_0^\theta(B_4)}, \end{aligned}$$

ν_1 and ν_2 satisfies ε -inferentially private constraint (5). Thus, $\gamma' \notin \overline{\mathcal{P}}_{\mathcal{I}}$.

Next, we only need to show that if γ' does not satisfies (6), γ' has a positive measurable $M \in \text{supp}(\gamma')$ and for each $\mu \in M$, there is a subset $F \in \mathcal{B}(\Theta)$ with $\mu_0^\theta(F) > 0$ and some $\varepsilon' \in (0, \varepsilon)$ such that (11) holds.

Let M be the set of $\nu \in \text{supp}(\gamma')$ such that (6) does not hold and there is a measurable function g_ν defined by (7). Suppose $g_\nu \neq 1$ μ_0^θ -almost surely, otherwise $\nu = \mu_0^\theta$ almost surely which is a trivial case. Then, $\bar{\varepsilon}_\nu$ defined as (8) is contained in $(0, \varepsilon)$. Define two measurable sets

$$\begin{aligned} \overline{B} &:= \{\theta \in \Theta : g_\nu(\theta) = e^{\bar{\varepsilon}_\nu}\}, \\ \underline{B} &:= \{\theta \in \Theta : g_\nu(\theta) = e^{\bar{\varepsilon}_\nu - \varepsilon}\}. \end{aligned}$$

(1) Suppose $\mu_0^\theta(\overline{B}) = 0$ or $\mu_0^\theta(\underline{B}) = 0$, and w.l.o.g. assume $\mu_0^\theta(\overline{B}) = 0$, that is,

$$\mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [e^{\bar{\varepsilon}_\nu - \varepsilon}, e^{\bar{\varepsilon}_\nu}]\}) = 1.$$

Similarly as (10), we can show that

$$\mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [1, e^{\bar{\varepsilon}_\nu}]\}) > 0. \quad (12)$$

(2) Suppose $\mu_0^\theta(\overline{B}) > 0$ and $\mu_0^\theta(\underline{B}) > 0$. If $\mu_0^\theta(\overline{B} \cup \underline{B}) = 1$, then ν satisfies (6). Hence, under the assumption that ν does not satisfy (6),

$$\mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in (e^{\bar{\varepsilon}_\nu - \varepsilon}, e^{\bar{\varepsilon}_\nu})\}) > 0.$$

W.l.o.g., we can assume that (12) holds. Since

$$\begin{aligned} 0 &= \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [e^{\bar{\varepsilon}_\nu}, e^{\bar{\varepsilon}_\nu}]\}) \\ &= \mu_0^\theta\left(\bigcap_{n=1}^{\infty} \left\{\theta \in \Theta : g_\nu(\theta) \in (e^{\frac{n-1}{n}\bar{\varepsilon}_\nu}, e^{\bar{\varepsilon}_\nu}]\right\}\right), \end{aligned}$$

where the first equality is due to $\{\theta \in \Theta : g_\nu(\theta) \in [e^{\bar{\varepsilon}_\nu}, e^{\bar{\varepsilon}_\nu}]\} = \emptyset$ and the second equality holds since for any $\theta \in \Theta$ for which $g_\nu(\theta) < e^{\bar{\varepsilon}_\nu}$, there is $N > 0$ such that $g_\nu(\theta) < \frac{N-1}{N}e^{\bar{\varepsilon}_\nu}$. Therefore, since the continuity of (finite) measure, there exists $N > 0$ such that

$$\mu_0^\theta\left(\bigcap_{n=1}^N \left\{\theta \in \Theta : g_\nu(\theta) \in (e^{\frac{n-1}{n}\bar{\varepsilon}_\nu}, e^{\bar{\varepsilon}_\nu}]\right\}\right) < \mu_0^\theta(\{\theta \in \Theta : g_\nu(\theta) \in [1, e^{\bar{\varepsilon}_\nu}]\}),$$

then

$$\mu_0^\theta\left(\left\{\theta \in \Theta : g_\nu(\theta) \in [1, e^{\frac{N-1}{N}\bar{\varepsilon}_\nu}]\right\}\right) > 0. \quad (13)$$

Denote $F := \left\{\theta \in \Theta : g_\nu(\theta) \in [1, e^{\frac{N-1}{N}\bar{\varepsilon}_\nu}]\right\}$, then since (13) and (9), for any $B_1 \in \mathcal{B}(F)$ and $B_2 \in \mathcal{B}(\Theta)$ with $\mu_0(B_1) > 0$ and $\mu_0(B_2) > 0$,

$$e^{-\bar{\varepsilon}_\nu} \frac{\mu_0^\theta(B_1)}{\mu_0^\theta(B_2)} \leq \frac{\nu(B_1)}{\nu(B_2)} \leq e^{\varepsilon - \frac{1}{N}\bar{\varepsilon}_\nu} \frac{\mu_0^\theta(B_1)}{\mu_0^\theta(B_2)}.$$

Since $\bar{\varepsilon}_\nu \in (0, \varepsilon)$, we construct a set F satisfying (11). □

Proof of Proposition 3. “Only if”. We first show the part (1).

- (i) Suppose there is a subset $Y \in \tilde{f}(\Theta)$ which contains more than two points such that ν puts a positive probability on $\tilde{f}^{-1}(y)$ for all $y \in Y$. Since $\mathbb{E}_\nu[\tilde{f}(\theta)]$ is a combination of $\tilde{f}(\theta)$ for $\theta \in \text{supp}(\nu)$, there exists $\underline{y} < \mathbb{E}_\nu[\tilde{f}(\theta)] < \bar{y}$ such that $\underline{p} := \nu(\{\theta \in \Theta | \tilde{f}(\theta) \leq \underline{y}\}) > 0$, $\bar{p} := \nu(\{\theta \in \Theta | \tilde{f}(\theta) \geq \bar{y}\}) > 0$, $\underline{p} + \bar{p} < 1$ and there is $\alpha \in (0, 1)$ such that $\alpha \mathbb{E}_\nu[\tilde{f}(\theta) | \tilde{f}(\theta) \leq \underline{y}] + (1 - \alpha) \mathbb{E}_\nu[\tilde{f}(\theta) | \tilde{f}(\theta) \geq \bar{y}] = \mathbb{E}_\nu[\tilde{f}(\theta)]$. Denote $\lambda := \min\{\frac{\underline{p}}{\alpha}, \frac{\bar{p}}{1-\alpha}\} < 1$. Then ν can be split into $\nu_1(\theta) := \mathbf{1}\{\tilde{f}(\theta) \leq \underline{y}\} \alpha \frac{\nu(\theta)}{\underline{p}} + \mathbf{1}\{\tilde{f}(\theta) \geq \bar{y}\} (1 - \alpha) \frac{\nu(\theta)}{\bar{p}}$ and $\nu_2 = \frac{1}{1-\lambda}(\nu - \lambda \nu_1)$ with probability λ and $(1 - \lambda)$ respectively. Since $\mathbb{E}_{\nu_1}[\tilde{f}(\theta)] = \mathbb{E}_\nu[\tilde{f}(\theta)] = \mathbb{E}_{\nu_2}[\tilde{f}(\theta)]$, this split does not change the distribution of posterior mean of $\tilde{f}$.
- (ii) Suppose there a positive measure subset of belief such that ν puts α on the set $\tilde{f}^{-1}(y_1)$ and $(1 - \alpha)$ on the set $\tilde{f}^{-1}(y_2)$ but $\text{supp}(\nu) \cap \tilde{f}^{-1}(y_1)$ is not a singleton. Denote $\nu_\theta(\theta') := \alpha \mathbf{1}\{\theta' = \theta\} + \mathbf{1}\{\theta' \in \tilde{f}^{-1}(y_2)\} \nu(\theta')$ for $\theta \in \text{supp}(\nu) \cap \tilde{f}^{-1}(y_1)$, then $\nu =$

$\int_{\text{supp}(\nu) \cap \tilde{f}^{-1}(y_1)} \nu_\theta d\nu(\theta)$. Since such μ has a split, and notice that $\mathbb{E}_{\nu_\theta}[\tilde{f}(\theta)] = \mathbb{E}_\nu[\tilde{f}(\theta)]$ for all $\theta \in \text{supp}(\nu) \cap \tilde{f}^{-1}(y_1)$, γ has a strictly mean-preserving spreading that cannot change the induced distribution of posterior mean of $\tilde{f}$.

Suppose the part (2) does not hold, then κ_γ is a strictly mean-preserving contraction of $\bar{\kappa}$, i.e., there exists a non-degenerated dilation $K : \mathbb{R} \rightarrow \Delta(\mathbb{R})$, that is, $y = \int_{\mathbb{R}} y' K(dy'|y)$ for all $y \in \text{supp}(\kappa_\gamma)$, such that

$$\bar{\kappa}(B) = \int_{\mathbb{R}} K(B|y) d\kappa_\gamma, \quad \forall B \in \mathcal{B}(\mathbb{R}).$$

That means, there is a positive measure subset $Y \subseteq \text{supp}(\kappa_\gamma)$ such that any $y \in Y$ is split into $\text{supp}(K(\cdot|y))$ according to non-degenerated distribution measure $K(\cdot|y)$. We now construct a dilation from $\Delta(\Theta)$ to $\Delta(\Delta(\Theta))$ based on K .

W.l.o.g., suppose the part (1) holds. Denote E as the set of $y \in \text{supp}(\kappa_\gamma)$ such that almost surely, all $\nu \in \text{supp}(\gamma)$ with $\mathbb{E}_\nu[\tilde{f}(\theta)] = y$ satisfies $\text{supp}(\nu) \subseteq \tilde{f}^{-1}(y)$. $\text{supp}(\kappa_\gamma) \setminus E$ has a positive measure. Otherwise, by the definition of E , κ_γ disclose full information about posterior mean of $\tilde{f}$. Hence $\kappa_\gamma = \bar{\kappa}$, contradiction. Assume $Y \subseteq \text{supp}(\kappa_\gamma) \setminus E$. For $y \in Y$, it is comprised by some $\nu \in \text{supp}(\gamma)$ such that its support consists of two points $\theta_1 \in \tilde{f}^{-1}(y_1)$ and $\theta_2 \in \tilde{f}^{-1}(y_2)$, where $y_1 < y_2$. For each $y = \alpha y_1 + (1 - \alpha)y_2$, with $\alpha \in [0, 1]$, define $\nu_y := \alpha \delta_{\theta_1} + (1 - \alpha) \delta_{\theta_2}$. If $\text{supp}(K(\cdot|y)) \subseteq [y_1, y_2]$, then $\nu = \int_{\text{supp}(K(\cdot|y))} \nu_{y'} K(y'|y)$. Thus, γ has a strictly mean-preserving spread and satisfies the constraint. If $\text{supp}(K(\cdot|x)) \not\subseteq [y_1, y_2]$, then there exists a non-degenerated $Q(\cdot|y)$ which is the mean-preserving contraction of $K(\cdot|y)$ and $\text{supp}(Q(\cdot|y)) \subseteq [y_1, y_2]$. Following the same argument above, replacing by $Q(\cdot|y)$.

“If”. Suppose there exists another $\gamma' \in \mathcal{P}_{\mathcal{E}}$ that strictly Blackwell-dominates γ , then there is a non-degenerated dilation $K : \Delta(\Theta) \rightarrow \Delta(\Delta(\Theta))$, that is, $\nu = \int_{\Delta(\Theta)} \nu' K(d\nu'|\nu)$, for all $\nu \in \text{supp}(\gamma)$, such that $\gamma'(B) = \int_{\Delta(\Theta)} K(B|\nu) d\gamma$, for all $B \in \mathcal{B}(\Delta(\Theta))$. Define $Q(Y|y) := \mathbb{E}_\gamma[K(\{\nu' : \mathbb{E}_{\nu'}[\tilde{f}(\theta)] \in Y\}|\nu_y) | \{\nu_y : \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = y\}]$, for almost every $Y \in \mathcal{B}(\mathbb{R})$. Because of the part (2) that support set of all $\nu \in \text{supp}(\gamma)$ contains at most two points, for all ν_y such that $K(\nu_y|\nu_y) < 1$, it must have $K(\{\nu : \mathbb{E}_\nu[\tilde{f}(\theta)] \neq y\}|\nu_y) > 0$. Because K is non-degenerated, $\{\nu : K(\nu|\nu) < 1\}$ has a positive measure. Hence, Q is non-degenerated.

The following proof is similar with proof of Lemma 2.

$$\begin{aligned} \kappa_{\gamma'}(Y) &= \gamma'(\{\nu : \mathbb{E}_\nu[\tilde{f}(\theta)] \in Y\}) = \int_{\Delta(\Theta)} K(\{\nu' : \mathbb{E}_{\nu'}[\tilde{f}(\theta)] \in Y\}|\nu) d\gamma(\nu) \\ &= \int_{\mathbb{R}} \int_{\{\nu_y : \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = y\}} K(\{\nu' : \mathbb{E}_{\nu'}[\tilde{f}(\theta)] \in Y\}|\nu_y) d\gamma(\nu_y) dy = \int_{\mathbb{R}} Q(Y|y) d\kappa_\gamma(y), \end{aligned}$$

$$\begin{aligned}
yd\kappa_\gamma(y) &= \int_{\{\nu_y: \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = x\}} yd\gamma(\nu_y) = \int_{\{\nu_y: \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = y\}} \left[\int_{\Theta} \tilde{f}(\theta) d \left(\int_{\Delta(\Theta)} \nu' K(d\nu'|\nu_y) \right) \right] d\gamma(\nu_y) \\
&= \int_{\{\nu_y: \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = y\}} \int_{\Delta(\Theta)} \int_{\Theta} \tilde{f}(\theta) d\nu' K(d\nu'|\nu_y) d\gamma(\nu_y) \\
&= \int_{\{\nu_y: \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = y\}} \int_{\Delta(\Theta)} \mathbb{E}_{\nu'}[\tilde{f}(\theta)] K(d\nu'|\nu_y) d\gamma(\nu_y) \\
&= \int_{\{\nu_y: \mathbb{E}_{\nu_y}[\tilde{f}(\theta)] = y\}} \int_{\mathbb{R}} y' dK(\{\nu_{y'} : \mathbb{E}_{\nu_{y'}}[\tilde{f}(\theta)] = y'\} | \mu_y) d\gamma(\nu_y) \\
&= \int_{\mathbb{R}} y' Q(dy'|y) d\kappa_\gamma(y) \Rightarrow y = \int_{\mathbb{R}} y' Q(dy'|y) \text{ almost surely.}
\end{aligned}$$

Thus, $\kappa_{\gamma'}$ is a strictly mean-preserving spread of $\kappa_\gamma = \bar{\kappa}$, which contradict with $\gamma' \in \mathcal{P}_\mathcal{E}$. $\square$