Mechanism Design with Information Leakage*

Samuel Häfner
University of St.Gallen
samuel.haefner@unisg.ch

Marek Pycia
University of Zurich
marek.pycia@econ.uzh.ch

Haoyuan Zeng
University of Zurich
haoyuan.zeng@econ.uzh.ch

October 2025
Latest version here.

Abstract

We study the design of mechanisms—e.g., auctions—when the designer does not control information flows between mechanism participants. A mechanism equilibrium is leakage-proof if no player conditions their actions on leaked information; a property distinct from ex-post incentive compatibility. Only leakage-proof mechanisms can implement social choice functions in environments with leakage. Efficient auctions need to be leakage-proof, while revenue-maximizing ones not necessarily so. Second-price and ascending auctions are leakage-proof; first-price auctions are not; while whether descending auctions are leakage-proof depends on tie-breaking.

Keywords: Mechanism Design, Auction, Information Leakage, Eavesdropping

JEL Codes: D44, D47, D82, D83

*We thank Piotr Dworczak, Paul Klemperer, Nick Netzer, Axel Ockenfels, Armin Schmutzler, and seminar participants in Zürich and St. Gallen for their comments.

A small leak will sink a great ship. (Benjamin Franklin)

1 Introduction

Standard mechanism design focuses on environments in which the mechanism designer has complete control over the extensive-form game the mechanism participants play, including complete control over who observes the moves they make. This complete control assumption underlies the revelation principle (Myerson, 1981) and many other insights in mechanism design. In many modern market environments, however, the designer cannot fully control what players observe. In this paper, we study mechanism design in such environments.

For example, many online auctions are prone to so-called eavesdropping attacks, in which some participants can listen to the network for upcoming bids and use this information to their advantage when submitting their own bids; in the presence of eavesdropping, the designer does not have complete control over what the bidders observe.[1] In financial markets, participants with fast access to trading venues can be the first to assess the limit order book and execute their trade orders more quickly than competitors, thereby capturing transient arbitrage opportunities. These market participants have more flexibility over what they observe than the market designer may want to grant them.[2] In peer-to-peer networks, such as blockchains, faster participants also observe slower participants' actions. Exploiting this information—known as front-running—frequently occurs on so-called decentralized exchanges and is regarded as one of the most pressing problems of blockchain technology.[3]

To see the implications of information leakage (or eavesdropping), consider the sale of an object to one of two bidders who have private values drawn independently from the same distribution. In absence of information leakage, the first-price auction with no reserve is efficient while the first-price auction with optimal reserve is revenue

[1] For institutional details on eavesdropping attacks, see, e.g., Franklin and Reiter (1996).

[2] The speed of access is so vital that some large traders place their servers as close to the exchange as possible and connect the two via dedicated optical fiber connections, resulting in what is sometimes referred to as a high-frequency trading arms race (Budish, Cramton, and Shim, 2015).

[3] See, e.g., Eskandari, Moosavi, and Clark (2020). Blockchains are ledgers that record transactions. New transactions are added to blockchains in batches, called blocks. The production of new blocks happens in discrete time intervals. Between block production, the recent transactions are stored in the network. By design, the network is open, and any participant can inspect upcoming transactions.

maximizing. Both of these properties of first-price auctions break in the presence of information leakage. Indeed, suppose the the bid of the first bidder is seen by the second bidder before the second bidder submits their bid. The second bidder may then condition their bid on the bid of the first bidder, which breaks both efficiency and revenue-maximization of first-price auctions. Indeed, efficiency fails as, in any equilibrium, all but one types of the first bidder bid below their own value thus allowing the second bidder to win the auction when their value is lower than that of the first bidder but still above the bid of the first bidder. Similarly, revenue-maximization fails because the information leakage makes it impossible to always assign the object to the bidder with higher Myersonian virtual value.[4]

We study a designer who can choose any finite extensive-form game with perfect recall but does not control what history of the game the players see. Even when two or more players are to move simultaneously in the game chosen by the designer, some of them ("faster" players) might see the moves made by others ("slower" players) before deciding on their own move.[5] The information about players' moves might hence leak to other players and the designer cannot prevent it.

In this setting we first address the question of what extensive-form game can implement an arbitrarily fixed (implementable) social choice function. We show that a game implements the social choice function if and only if it admits an equilibrium in which all players pursue strategies that are independent of the leaked information and this equilibrium implements the social choice function. We call such equilibria—in which leaked information is effectively ignored by the players—leakage-proof. Leakage-proofness is hence a necessary condition for implementability in environments in which the designer does not control information flows between the players.

We then address the more applied problems of designing efficient auctions and revenue-maximizing auctions. We show that an auction is efficient only if it admits a leakage-proof ε-equilibrium that always results in efficient outcome. Second-price (and ascending) auctions remain efficient in the presence of leakage, whether we impose a

[4]We provide more details on this example in Section 2.

[5]For expositional reasons, we focus on games that allow simultaneous moves but otherwise have perfect information. All our insights remain true for general imperfect-information games. Not only all of our arguments extend to the general case, also our results directly imply the analogous results for the general case. The reason is simple: we are interested in mechanisms that perform well in the presence of leakage, including leakage that reveals to every player the history of past play.

common prior assumption or whether we allow heterogeneous priors. In contrast, first-price auctions are no longer efficient in the presence of leakage. When bidders share a common prior, then second-price (and ascending) auctions with optimally set reserve price are revenue maximizing. In contrast, neither the second-price nor the first-price auctions maximize revenue independent of bidders' beliefs about leakage. Indeed, the truthful-bidding equilibria in these mechanisms are leakage-proof and hence attain the Myersonian revenue upper bound irrespective of leakage. In light of the above-discussed example, for some leakage priors second-price and ascending auctions raised noticeably more revenue than many other mechanisms, including first-price auctions. Beyond the common prior environment, the second-price and ascending auctions may however fail to maximize revenue. For instance, if no bidder can eavesdrop on other bidders but all bidders are concerned about the possibility that their bid leaks to others, then the first-price auction with optimal reserve achieves higher revenue than the second-price (and ascending) auctions.

We further show that pure-strategy ex-post equilibria are always leakage-proof, but mixed-strategy ex-post equilibria are not necessarily so. Also, static (one simultaneous move) leakage-proof equilibria are also ex-post equilibria, but this property does not extend to leakage-proof equilibria of dynamic mechanisms.

Finally, as discussed above, second-price (and ascending) auctions are leakage-proof while first-price auctions are not. Interestingly, whether descending (Dutch) auctions are leakage-proof depends on the tie-breaking rule employed in them. Under the standard equal-probability of winning tie-break rules, Dutch auctions are not leakage-proof. However, these auctions become leakage-proof if the good is not allocated in case of a tie-break.

1.1 Related Literature

While we believe we are the first to systematically study information leakage in a mechanism design framework, information flows have been studied in games and applied economic contexts. In particular, Solan and Yariv (2004) consider two-player normal-form games, in which one player can noisily observe the strategy of the other player at some cost, and fully characterize the distributions over the players' payoffs that can obtain in equilibrium. Penta and Zuazo-Garin (2022) characterize the predictions of rationality and common belief in rationality that do not depend on players'

infinite order beliefs over whether their actions are observable to their opponents.[6]

Ex-post equilibrium (Hurwicz, 1972; Dasgupta, Hammond, and Maskin, 1979a) may be motivated by the concern about leakage of information about players' types, as opposed to actions.[7] As we show, these two concerns are distinct. Madarász (2011) and Madarász (2012) study the impact of beliefs about leakage of type information on Coasian dynamics, while Daley and Green (2012) study the impact of the exogenous arrival of news on types in Coasian dynamics.[8]

By studying leakage, we contribute to the recent wave of literature that differentiate between auction mechanism on the ground of their robustness to attacks such as (the lack of) credibility (Akbarpour and Li, 2020; Banchio, Skrzypacz, and Yang, 2025), auditability (Woodward, 2020), or shill-bidding (e.g., Komo, Kominers, and Roughgarden, 2024; Zeng, 2025). While this literature focuses on the seller's inability to fully commit, we study the seller's inability to control information flows.

We also contribute to the literature on misspecified beliefs in mechanisms and markets (Ledyard, 1978; Bergemann and Morris, 2005; Chu and Shen, 2006; Chassang, 2013; Carroll, 2015; Wolitzky, 2016; Madarász and Prat, 2017; Li, 2017; Börgers and Li, 2019; Pycia and Troyan, 2023; Li ⓡ and Dworczak, 2024). While the main thrust of this literature is that robustness to misspecification requires the mechanism to be simple, we show that achieving efficiency in a way that is robust to misspecification requires the mechanism to be leakage-proof.[9]

Our analysis of leakage is relevant for many applied problems. In financial markets, traders often use faster access to information about upcoming trades to exploit arbitrage opportunities (see, e.g., Budish, Cramton, and Shim, 2015; Baldauf and Mollner, 2022, for extensive evidence). Budish, Cramton, and Shim (2015) propose batch auctions to limit the adverse effects of such front-running, while we focus on establishing the link between robust implementation and leakage-freeness.[10]

[6]For analysis of repeated games with mediation, in which mediators recommendations leak after each period, see Ewerhart and Zeng (2025).

[7]For a recent use of such motivated ex-post equilibrium concepts in mechanism design, see, e.g., Zeng (2025).

[8]Madarász and Pycia (2022) allow the parties to control the news arrival and focus on privacy considerations, while Madarász and Pycia (2025) study endogenous arrival of news in a wide class of trading games.

[9]See, e.g., Harrison and Kreps (1979); Eyster and Piccione (2013); Heidhues, Kőszegi, and Strack (2018); Jantschgi, Nax, Pradelski, and Pycia (2024) for the analysis of the complementary problem of how traders with misspecified beliefs behave.

[10]Information flows between participants matter also in treasury auctions and over-the-counter

Leakage is also a major issue in online markets and there is an extensive computer science literature on cryptographic approaches to make auctions robust to information leakage (Kudo, 1998; Abe and Suzuki, 2002; Parkes, Rabin, Shieber, and Thorpe, 2006).[11] Cryptographic methods that allow to shield bids during the auction and to verify them (or at least their ranking) after the auction have received renewed interest in decentralized blockchain settings, lately (Blass and Kerschbaum, 2018, 2020; Galal and Youssef, 2019). We view our analysis as complementary to these approaches, because we focus on incentives, rather than cryptography, to prevent bidders from using leaked information.

Another previously studied solution to leakage are candle auctions, in which the acceptance of bids is uncertain. Gehrlein, Häfner, and Oechssler (2025) show that a candle auction may implement the efficient and optimal outcome despite leakage, while Häfner and Stewart (2025) show that candle auctions admit an approximate ex-post equilibrium in quickly escalating bids.

Finally, auctions are sometimes designed with explicit leakage component in order to privilege a specific bidder, who has the right-of-first refusal. Such rights are commonly seen in the landlord-tenant and competitive procurement settings, see Riley and Samuelson (1981), Burguet and Perry (2009), Bikhchandani, Lippman, and Ryan (2002), Choi (2009), and Doran (2018).

2 Example

We give a brief example of the implications of leakage in auctions. Consider a first-price auction with one object for sale and two bidders $i \in \{1, 2\}$. Bidders have independent private values, θ_i, drawn from the uniform distribution on $[0, 1]$ and quasi-linear utility.

Without information leakage, we have a standard static auction. The unique equilibrium, where bidder i follows a bidding strategy $\beta_i(\theta_i)$ is well known to be symmetric and given by

$$\beta_i(\theta_i) = \frac{\theta_i}{2}, \quad i = 1, 2.$$

markets; for the discussion of evidence see, e.g., Hortaçsu and Sareen (2004), Hortaçsu and Kastl (2012), and Garratt, Lee, Martin, and Townsend (2019).

[11] Leakage of information about actions is related to sniping that is submitting one's bid as close as possible to the end-time of the auction (Roth and Ockenfels, 2002). Our analysis of leakage-proofness might be viewed as addressing the inefficiencies associated with sniping.

Now, suppose we have information leakage and bidder 2 observes bidder 1's bid, which is common knowledge among the bidders. The auctioneer, however, cannot condition the allocation on a bidder's identity. To guarantee the existence of best responses, we assume that the object goes to bidder 2 in case of a tie. This can be justified because bidder 2 can marginally overbid bidder 1 and win. The equilibrium strategies can be obtained by backward induction and are as follows.

$$\beta_1(\theta_1) = \frac{\theta_1}{2} \text{ and } \beta_2(\theta_2, b_2) = \begin{cases} b_1 & \text{if } \theta_2 \geq b_1 \\ 0 & \text{otherwise.} \end{cases}$$

Two observations follow from this example. First, efficiency is violated, because bidder 2 wins the auction when $\theta_1/2 < \theta_2 < \theta_1$. Second, compared to the case without information leakage, the expected revenue is reduced. This is straightforward to see because the equilibrium strategy for bidder 1 is the same as in an environment without information leakage, but bidder 2 bids strictly lower conditional on winning. Indeed, direct computation gives that revenue decreases from $\frac{1}{3}$ to $\frac{1}{6}$.

3 Model

3.1 Environment

Players, Outcomes, and Payoffs Throughout, we consider finite games. The set of players is $N = \{1, \ldots, n\}$. Player i's payoff type is $\theta_i \in \theta_i$, where θ_i is a finite set. We write $\theta \in \times_{i \in N} \theta_i \equiv \Theta$ for a payoff type profile. All players have a common prior $\rho(\cdot)$ on Θ. We assume that for each player $i \in N$, the payoff type θ_i is independently drawn from Θ_i. Hence, $\rho(\theta) = \rho_1(\theta_1)\rho_2(\theta_2)...\rho_n(\theta_n)$, where ρ_i is the marginal distribution over θ_i. The finite set of outcomes is X. Each player has a von Neumann-Morgenstern utility function $u_i : X \times \Theta \to \mathbb{R}$. The payoff structure is fixed, and it is common knowledge.

Extensive Form The designer can choose an extensive form. We restrict attention to extensive-form games with perfect recall, which allow simultaneous moves but otherwise have perfect information. While this restriction simplifies our exposition and terminology, it can be fully relaxed: as mentioned in the introduction, all our

insights remain true for general imperfect-information games.[12] Formally, we consider multistage games with observed actions, G, where

$$G = \left\{H, \subset, P, (A_i)_{i \in N}, g\right\}.$$

The elements of G are defined as follows and summarized in Table 1.

1. H is a finite set of sequences, consisting of the public histories of moves before each stage $k = 1, 2, \ldots$.
   (a) The history before first stage, h^0, is equal to the empty sequence $h_\emptyset \in H$.
   (b) The history before stage $k \geq 2$ is a sequence of length $k-1$, denoted as $h^{k-1} = \left(a^1, a^2, \ldots, a^{k-1}\right) \in H$. Each vector $a^t = (a_1^t, ..., a_n^t)$ consists of the actions a_i^t taken by player i at stage t, where some, but not all, actions a_i^t may be $\emptyset$.
   (c) If $h^k = \left(a^1, a^2, \ldots, a^k\right) \in H$, then $h^l = \left(a^1, a^2, \ldots, a^\ell\right) \in H$ for any $\ell < k$. We also say that h^k is a successor of h^ℓ, which is denoted by $h^\ell \subset h^k$.
2. The game terminates at the end of a stage k with $h^k = \left(a^1, a^2, \ldots, a^k\right) \in H$ if there is no a^{k+1} such that $h^k \| a^{k+1} \in H$, where $\|$ is the concatenation operator. The set of terminal histories is denoted by Z.
3. The player function $P : H \setminus Z \to 2^N$ assigns to each nonterminal history $h^{k-1} \in H \setminus Z$, a set of players $P\left(h^{k-1}\right)$ who simultaneously take actions at stage k.
4. H and P jointly satisfy the condition that for each nonterminal history $h^{k-1} \in H \setminus Z$, there is a set of feasible actions $A_i\left(h^{k-1}\right)$ for each player $i \in N$ at stage k such that
   (a) $A_i\left(h^{k-1}\right) = \emptyset$ for $i \notin P\left(h^{k-1}\right)$ and $|A_i\left(h^{k-1}\right)| \geq 2$ for $i \in P\left(h^{k-1}\right)$.
   (b) $\left\{a^k : h^{k-1} \| a^k \in H\right\} = \times_{i \in N} A_i\left(h^{k-1}\right)$.
5. The outcome resulting from any terminal history $z \in Z$ is denoted by $g(z) \in X$.

[12] All of our arguments extend to the general case, and, furthermore, our results directly imply the analogous results for the general imperfect-information case. Indeed, we study mechanisms that perform well in the presence of leakage and in general imperfect information games leakage might reveal to every player the history of past play.

Table 1: Notation

Name	Notation
Histories before each stage $k = 1, 2, \cdots$	$h^{k-1} \in H$
Precedence relation over histories	$\subset$
Set of players whose actions are considered at h^{k-1}	$P\left(h^{k-1}\right)$
Actions available at h^{k-1} for player i	$A_i\left(h^{k-1}\right)$
Terminal histories	$z \in Z$
Outcomes resulting from z	$g\left(z\right)$

In short, the game evolves over a finite number of stages $k = 1, 2, ...$ At the start of each stage $k = 1, 2, \ldots,$ given the history $h^{k-1} \in H$, the game G selects a set of players $P\left(h^{k-1}\right)$ who move simultaneously at stage k. Each active player $i \in P\left(h^{k-1}\right)$ takes an action $a_i^k \in A_i\left(h^{k-1}\right)$. The resulting action profile $a^k = (a_1^k, ..., a_n^k)$, where $a_i^k = \emptyset$ for $i \notin P\left(h^{k-1}\right)$, determines the updated history at the start of stage $k+1$, i.e. $h^k = h^{k-1} \| a^k$. This process continues until a terminal history is reached, determining the outcome of the game. At the start of each stage $k = 1, 2, ...,$ the histories $h^{k-1} \in H$ are publicly disclosed. Consequently, every player is perfectly informed of all actions previously taken by others, and we refer to H as the set of *public* histories.

Information Leakage In addition to the standard setup above, we assume uncontractible information leakage between players about the simultaneous actions taken within a stage. The main idea, loosely speaking, is that some players "see" the concurrent actions of the other active players. Yet, the designer cannot discriminate players based on who sees whom, because leakages are not verifiable.

Formally, we define a binary leakage order $\precsim$ on the set of players N. The order is complete and transitive. We write $i \precsim j$ iff $(i, j) \in \precsim$, we write $i \sim j$ iff $i \precsim j$ and $j \precsim i$, and we write $i \prec j$ iff $i \precsim j$ is true but $i \sim j$ is not. The interpretation of this order is that if, for two players i and j, it holds $i \prec j$, then j can observe i's concurrent actions whenever they are both asked to move at a stage, but not the other way round. We refer to player j as being faster than i (or, equivalently, i being slower than j). That is, information leaks from slow players to fast players. If for two players i and j it holds that $i \sim j$, then we say they are equally fast. In that case, players i and j cannot see each other's actions. For further reference, we let $\mathcal{L}$

denote the set of all leakage orders on the set N.

With information leakage, the players' information during the game differs. At stage k given the public history $h^{k-1} \in H$, player i under the leakage order $\precsim$ has a *private* history,

$$h_i^{k-1} = h^{k-1} \| \left\{ a_j^k \right\}_{j \prec i},$$

where $a_j^k \in A_j\left(h^{k-1}\right)$. The private history consists of both the public history and the leaked information.

We extend the precedence relation over public histories to private histories. At each stage k, player i's private history succeeds the public history of the current stage and precedes the public history of the next stage, i.e. $h^{k-1} \subseteq h_i^{k-1} \subseteq h^k$. In particular, when player i is (weakly) slower than any other player, their private history coincides with the public history, i.e., $h^{k-1} = h_i^{k-1}$.

Leakage-Order Beliefs To finish, we need to specify the players' leakage order beliefs. We require beliefs to be consistent with the actual leakage order, but we allow for beliefs that are inconsistent with those of other players.

To formalize this, let $\precsim$ be the actual leakage order and define $E_i(\precsim) \subseteq \mathcal{L}$ as the set of leakage orders that are consistent with $\precsim$ from the viewpoint of player i. Consistency here means that $E_i(\precsim)$ contains precisely those leakage orders under which i observes the same players' actions as under $\precsim$ (but not necessarily in the same order). Formally,

Definition 1 (The Set of Consistent Leakage Orders, $E_i(\precsim)$)**.** *For a given leakage order $\precsim$ and a player i, the set of consistent leakage orders, $E_i(\precsim)$, is given by*

$$E_i(\precsim) = \{ \precsim' \in \mathcal{L} : j \prec i \iff j \prec' i, \ \forall j \}.$$

We say that a player i is of leakage type t_i from some finite set of possible leakage types T_i. We write $T = T_1 \times ... \times T_n$ for the set of all possible type profiles $t = (t_1, ..., t_n)$. We work with leakage beliefs that can be captured by the classical notion of a type space, $\mathcal{T} = \{T, (\tau_1, ..., \tau_n)\}$ with $\tau_i : T_i \to \Delta(\mathcal{L}, T_{-i})$ for all i, where τ_i can be iteratively used to construct a player's first-order belief about $\mathcal{L}$ and their respective higher-order beliefs (e.g. Siniscalchi, 2008). Each player i of type t_i has a prior $\gamma_{-i}(.) = \text{marg}_{T_{-i}} \tau_i(t_i)(.)$ on the types of others, T_{-i}, where marg_Z denotes the marginal distribution over Z.

We make two substantive assumptions on any feasible type space $\mathcal{T}$. The first formalizes the idea that first-order beliefs must be consistent. Writing supp(.) for the support of the distribution given in the argument, we have:

Assumption 1 (Consistency of First-Order Beliefs)**.** *Suppose the leakage order is* $\precsim \in \mathcal{L}$*. For every* $t_i \in T_i$*, it holds that* $supp\,(marg_{\mathcal{L}}\,\tau_i(t_i)) \subseteq E_i(\precsim)$*.*

While all types in a feasible leakage type space must hold consistent first-order beliefs, we allow players to hold wrong higher-order beliefs. The following example illustrates what we have in mind.

Example 1. *Let* $\precsim$ *be the actual order, let* $L_1 = \{i \in N : i \precsim j, \forall j \in N\}$ *be the set of the slowest players and define the set of the k-th slowest players recursively as* $L_k = \{i \in N : i \precsim j, \forall j \in N \setminus (L_{k-1} \cup ... \cup L_1)\}$*. Let* $\bar{k}$ *be the fastest group of players in N; i.e.* $\bar{k}$ *is the lowest k for which* $\cup_{\kappa \leq k} L_\kappa = N$*. Suppose there are* $\precsim_1, ..., \precsim_{\bar{k}}$ *with* $\precsim_{\bar{k}} = \precsim$ *such that:*

(B_1*) Players in* L_1 *believe it is common knowledge among all players in* $L_{\bar{k}} \cup ... \cup L_1$ *that* $\precsim_1$ *is true, where* $\precsim_1 \in E_j(\precsim)$*,* $\forall j \in S_1$*.*

(B_2*) Players in* L_2 *believe it is common knowledge among all players in* $L_{\bar{k}} \cup ... \cup L_2$ *that both (*B_1*) and* $\precsim_2$ *are true, where* $\precsim_2 \in E_j(\precsim)$*,* $\forall j \in S_2$*.*

(B_3*) Players in* L_3 *believe it is common knowledge among all players in* $L_{\bar{k}} \cup ... \cup L_3$ *that both (*B_1*)–(*B_2*) and* $\precsim_3$ *are true, where* $\precsim_3 \in E_j(\precsim)$*,* $\forall j \in S_3$*.*

...

($B_{\bar{k}}$*) Players in* $L_{\bar{k}}$ *believe it is common knowledge among all players in* $L_{\bar{k}}$ *that both (*B_1*)–(*$B_{\bar{k}-1}$*) and* $\precsim_{\bar{k}}$ *are true.*

Under the leakage types described in Example 1, the players in the fastest group know the true leakage order, $\precsim$. Players in any group are aware of the leakage types of the slower players. Players in slower groups may have a common belief that differs from the beliefs of players in a faster group and wrongly attribute this belief to the faster players. When $\precsim_1 = ... = \precsim_{\bar{k}}$, then we have common knowledge about the leakage order among all players.

In general, beliefs about the slower players' beliefs need not be correct, nor need they be the same for players that are equally fast for our results to hold. Furthermore,

beliefs can be probabilistic. We assume, however, that the leakage type space contains leakage type profiles of at least two variants that we call the zero-profile and one-profile, respectively.

Definition 2 (Zero-Profile and One-Profile)**.**

1. *The zero-profile* $t^0 = (t_1^0, ..., t_n^0)$ *is a leakage type profile, where* t_i^0 *corresponds to the belief: It is common knowledge that everyone is equally fast.*

2. *A one-profile is any of the* n *permutations of the leakage type profile*

$$(t_1^0, ..., t_{i-1}^0, t_i^1, t_{i+1}^0 ...., t_n^0),$$

where t_i^0 *is as above and* t_i^1 *corresponds to the belief: I am faster than everyone else, who all believe it is common knowledge that everyone is equally fast.*

Assumption 2 (Minimally Rich Type Space)**.** *Any feasible type space* $\mathcal{T}$ *contains the zero-profile and all permutations of the one-profile.*

The leakage order in which everyone is equally fast is the unique order that is consistent with the zero-profile. Moreover, leakage orders in which exactly one player is faster than all other players (who are equally fast) are uniquely consistent with a one-profile. In other words, we consider situations with at least $n + 1$ leakage orders: everyone is equally fast, and one of the n players is faster than the other $n - 1$ players. Assumption 2 is crucial for establishing the necessity of a mechanism's leakage-proofness in various contexts, as we discuss after the results below.[13]

3.2 Equilibrium

The primary object of our interest is the leakage environment Γ, comprising the multistage game G (which the designer can control) together with a leakage type space $\mathcal{T}$ (which the designer cannot control),

$$\Gamma = (G, \mathcal{T}) .$$

A strategy $S_i(\theta_i, t_i)(h_i^{k-1})$ for player i maps, for every private history h_i^{k-1}, the type (θ_i, t_i) into a probability distribution over the set of available actions $A_i\left(h^{k-1}\right)$,

[13] We further discuss how our setup relates to the special case of common knowledge about the leakage order in Section B.

which depends on the public history h^{k-1} preceding the private history h_i^{k-1}. That is, $S_i(\theta_i, t_i)(h_i^{k-1}) \in \Delta A_i\left(h^{k-1}\right)$, where $h^{k-1} \subseteq h_i^{k-1}$. For a fixed leakage type t_i, we denote the set of available strategies $S_i(., t_i)$ by $\sum_i(t_i)$. The set of private histories is denoted by $H_i(t_i)$.

The belief of player i after private history h_i^{k-1} is denoted by

$$\mu_i\left(h_i^{k-1}, t_i\right) \in \Delta\left(\Theta_{-i}, T_{-i}\right),$$

which is a probability distribution over the value types of the opponents, Θ_{-i}, and the leakage types of the opponents, T_{-i}. Beliefs depend on the private history h_i^{k-1} and the particular leakage type t_i determining how the player interprets the observed history. The belief at the beginning of the game is $\mu_i\left(h_\emptyset, t_i\right) = \prod_{j \neq i} \rho_j(\theta_j) \times \gamma_{-i}(T_{-i})$. We collect the profile of conditional beliefs in $\mu = (\mu_i)_{i \in N}$.

To define utilities, let $\sigma = (\sigma_i)_{i \in N}$ denote a profile of contingent action plans for each player, $\sigma_i(h_i^{k-1}) \in \Delta A_i(h^{k-1})$, prescribing the distribution over available actions after any private history h_i^{k-1}. Given such a profile σ, let $\zeta(\sigma)$ denote the probability distribution over the terminal histories Z in the game G. The expected utility of player i under such a profile σ is

$$U_i\left(\sigma, \theta\right) \equiv \sum_{z \in Z} u_i\left(z, \theta\right) \zeta(\sigma)(z).$$

Our equilibrium concept is standard PBE (Fudenberg and Tirole, 1991), requiring that player i responds optimally to what they expect the other players to play, given their beliefs about payoff and leakage types. Whenever possible (in particular, if t_i is such that h_i^{k-1} appears on the path of play given an opponent strategy profile $(S_j(.,.))_{j \neq i}$ and the corresponding beliefs about payoff and leakage types), player i updates their beliefs according to Bayes' rule. Off the path of play, we allow for arbitrary beliefs.

Definition 3 (Equilibrium)**.** *A strategy-belief profile (S, μ) is a perfect Bayesian equilibrium (PBE) in the environment $\Gamma = (G, \mathcal{T})$, if the following conditions are satisfied.*

1. *Sequential rationality with information leakage: For all players $i \in N$, all types $(\theta_i, t_i) \in \Theta_i \times T_i$, all stages k, and all histories $h_i^{k-1} \in H_i\left(t_i\right)$,*

$$S_i(\theta_i, t_i) \in \arg \max_{\sigma_i \in \Sigma_i(t_i)} \mathbb{E}_{\mu_i}\left[U_i\left(\sigma_i, (S_j(\theta_j, t_j))_{j \neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right]$$

2. *Bayesian updating whenever possible: If possible for player i having leakage type t_i $\mu_i\left(h_i^{k-1}, t_i\right)(\theta_{-i}, t_{-i})$ is updated according to Bayes' rule at stage k and private history h_i^{k-1}.*

3.3 Leakage-Proof Equilibrium

We say that the game G is the *default game*. G is a standard finite extensive form game, for which we know a PBE to always exist (Selten, 1975). We call this PBE the default PBE and denote it by (S^0, μ^0). Further, we refer to (G, S^0, μ^0) as the default mechanism.

In all our results we assume that the game G is pruned (Akbarpour and Li, 2020) with respect to the default strategy profile. Pruning the game in such a way removes all histories that cannot be reached under any payoff type profile in the equilibrium S^0 and all stages in which just one bidder can move.[14] Formally,

Assumption 3 (Pruned Game G)**.** *The game G is pruned with respect to the default strategy profile S^0, if, for each history $h \in H$, there exists some payoff type profile θ such that h is on the path of play of some realization of $\left(S_j^0(\theta_j)\right)_{j \in N}$.*

Working with pruned games is standard, as it allows to focus on relevant histories. Indeed, when the game is pruned in accordance with Assumption 3, then the only off-path histories that remain relevant are those in which some players have mimicked a value type with different equilibrium actions in the past. These histories are off-path only for the deviating player; for the other players, they appear on path.

Our goal is to characterize mechanisms that are robust to information leakage — ensuring that no player has an incentive to deviate from the default strategy profile S^0 based on leaked information about others' actions.

Definition 4 (Leakage-Proof Equilibrium)**.** *A strategy-belief profile (S^0, μ^0) is a* leakage-proof *equilibrium in the game G if, for each admissable leakage type space $\mathcal{T}$, the environment $\Gamma = (G, \mathcal{T})$ admits a PBE (S, μ) such that for all payoff type profiles $\theta \in \Theta$ and leakage type profiles $t \in T$, all players $i \in N$, and all private histories $h_i^{k-1} \in H_i(t_i)$, where h_i^{k-1} is on the path of play of S^0,*

$$S_i(\theta_i, t_i)(h_i^{k-1}) = S_i^0(\theta_i)(h^{k-1}).$$

[14] We discuss unpruned games in Section B.

We also say that (G, S^0, μ^0) *is a leakage-proof mechanism.*

Leakage-proofness entails that, for all beliefs (and higher order beliefs) regarding leakage orders, it is optimal for players to continue following the default equilibrium strategies if they have done so in the past.[15] In other words, players' best responses disregard any leaked information, treating the game as if there were none.

4 Leakage Proofness and Implementation

We begin by showing that leakage-proofness is an essential requirement for implementation of a mechanism in the presence of leakages. Suppose that the designer's goal is to implement a social choice function f, which maps from payoff type profiles to outcomes, i.e. $f : \Theta \to X$.

Definition 5 (Partial Implementation under Leakages)**.** *The game* G *implements a social choice function* f*, if for each feasible leakage type space* $\mathcal{T}$*, the environment* $\Gamma = (G, \mathcal{T})$ *admits a PBE* (S, μ) *yielding outcome* $f(\theta)$ *for all payoff type profiles* $\theta \in \Theta$ *and all leakage type profiles* $t \in T$*.*

Definition 5 adapts the standard definition of partial implementation (Dasgupta et al., 1979b) to our environment with information leakage. It is equivalent to partial implementation in the absence of leakages, i.e., when the type space $\mathcal{T}$ only consists of the zero-profile (as in Definition 2). In the presence of information leakage, we require partial implementation to hold across all possible leakage type spaces and particular leakage type profiles.

The following is the main result of this section. It establishes that if the game G implements a social choice function despite information leakage, then G is part of a leakage-proof mechanism, and vice versa.

Theorem 1 (Implementability under Leakages)**.** *The game* G *implements a social choice function* f *if and only if the game* G *admits a leakage-proof equilibrium that implements* f*.*

[15] By imposing leakage-proofness only on the default equilibrium path-of-play, we follow the tradition in the literature on the extensive-form games of imposing additional equilibrium restrictions only on path, e.g. Pearce (1984), Shimoji and Watson (1998), Li (2017), and Pycia and Troyan (2023).

The "if" direction follows immediately from the definition of a leakage-proof equilibrium (Definition 4). To prove the converse, we invoke Assumption 2 of a minimally rich type space, together with the premise that the choice function $f(\theta)$ is implemented for any feasible type space $\mathcal{T}$ and leakage type profile $t \in T$.

The key observation is that these two ingredients imply the following: when all other players follow their default strategies, even a player who is uniquely the fastest cannot obtain a higher payoff than by playing the uninformed player's strategy. Since this uninformed strategy coincides with the default strategy, it follows that the default strategy is a best response when everyone else plays it.

5 Information Leakage in Auctions

We now apply the methods developed for general mechanisms to auctions.

5.1 Setup

Throughout, we consider auctions with a single, indivisible good for sale. We will refer to the players as bidders. The seller is denoted as bidder 0. An outcome $x = (q, m) \in X$ now consists of a vector of allocations $q = (q_1, \ldots, q_n) \in [0,1]^n$, where q_i denotes the probability that the good is allocated to bidder i and we have $\sum_{i=1}^{n} q_i \in [0,1]$, and a profile of payments $m = (m_1, \ldots, m_n) \in \mathbb{R}^n$, where m_i denotes the payment of bidder i to the seller.

The payoff type of bidder $i \in N$, θ_i, corresponds to a nonnegative real number, representing her valuation for the good on sale. The payoff of bidder i is given by

$$u_i(x, \theta) = q_i \theta_i - m_i.$$

For simplicity, we assume that the seller derives no value from the object, $\theta_0 = 0$. Hence, the seller's payoff is given by

$$u_0(x, \theta) = \sum_{i \in N} m_i.$$

For further reference, we write the value type sets as $\Theta_i = \{\theta_{i1}, ..., \theta_{i\bar{m}_i}\}$ where $\theta_{im+1} > \theta_{im}$ for all $m = 1, ..., \bar{m}_i - 1$ and $\bar{m}_i$ is the number of value types for bidder

i, $\bar{m}_i = |\Theta_i|$. We assume equally spaced value types; i.e., there is $\delta > 0$ such that

$$\theta_{im+1} - \theta_{im} = \delta$$

for all players $i \in N$ and all $m = 1, ..., \bar{m}_i - 1$.

In the environment $\Gamma = (G, \mathcal{T})$, a strategy-belief profile (S, μ) induces, for every value and leakage type profile (θ, t), a potentially random allocation and a potentially random payment, whose realizations we denote by $q(\theta, t) = (q_i(\theta, t))_{i \in N}$ and $m(\theta, t) = (m_i(\theta, t))_{i \in N}$ respectively. For any stage k and history h_i^{k-1} that is on the path of play of S, we then denote the subjective interim expected allocation and payment for bidder i of leakage type t_i choosing the continuation action plan of a type $(\hat{\theta}_i, \hat{t}_i)$ as

$$\begin{aligned} Q_i\left(\hat{\theta}_i, \hat{t}_i, t_i, h_i^k\right) &= \mathbb{E}_{\mu_i}\left[q_i\left((\hat{\theta}_i, \theta_{-i}), (\hat{t}_i, t_{-i})\right) \middle| h_i^{k-1}, t_i\right] \\ M_i\left(\hat{\theta}_i, \hat{t}_i, t_i, h_i^k\right) &= \mathbb{E}_{\mu_i}\left[m_i\left((\hat{\theta}_i, \theta_{-i}), (\hat{t}_i, t_{-i})\right) \middle| h_i^{k-1}, t_i\right]. \end{aligned}$$

Together, these quantities allow us to succinctly express the bidders' interim expected utilities. Again, fix a strategy profile S. Then, in every stage k and for every on-path history h_i^{k-1}, bidder i of type (θ_i, t_i) has an expected payoff equal to

$$\begin{aligned} &\mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i, t_i), (S_j(\theta_j, t_j))_{j \neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] \\ &\qquad = \theta_i Q_i(\theta_i, t_i, t_i, h_i^k) - M_i(\theta_i, t_i, t_i, h_i^k). \end{aligned}$$

The discrete type space necessitates that we work with approximate, rather than strict equilibria. Specifically, we define a Perfect Bayesian ϵ-equilibrium as follows.

Definition 6 (ϵ-Equilibrium)**.** *A strategy-belief profile (S, μ) is a perfect Bayesian ϵ-equilibrium (ϵ-PBE) in the environment $\Gamma = (G, \mathcal{T})$, if the following conditions are satisfied.*

1. *Sequential rationality with information leakage: For all players $i \in N$, all types $(\theta_i, t_i) \in \Theta_i \times T_i$, all stages k, and all histories $h_i^{k-1} \in H_i(t_i)$,*

$$\begin{aligned} &\mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i, t_i), (S_j(\theta_j, t_j))_{j \neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] + \epsilon \\ &\qquad \geq \mathbb{E}_{\mu_i}\left[U_i\left(\sigma_i, (S_j(\theta_j, t_j))_{j \neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right], \forall \sigma_i \in \Sigma_i(t_i). \end{aligned}$$

2. *Bayesian updating whenever possible: If possible for player i having leakage type t_i $\mu_i\left(h_i^{k-1}, t_i\right)(\theta_{-i}, t_{-i})$ is updated according to Bayes' rule at stage k and private history h_i^{k-1}.*

For the rest of this section, we make two assumptions about the auction mechanisms that we consider. First, we assume that the auctions are anonymous, meaning that the allocations and payments must be invariant under permutations.

Assumption 4 (Anonymous Auctions)**.** *For every permutation $\varphi : N \to N$, it holds*

$$q_i(\theta, t) = q_{\varphi(i)}((\theta_{\varphi(1)}, ..., \theta_{\varphi(n)}), ((t_{\varphi(1)}, ..., t_{\varphi(n)}))$$
$$m_i(\theta, t) = m_{\varphi(i)}((\theta_{\varphi(1)}, ..., \theta_{\varphi(n)}), ((t_{\varphi(1)}, ..., t_{\varphi(n)})).$$

Second, at every stage and feasible history, the lowest type of any player that is still in the auction dissipates all potential rent from obtaining the good, irrespective of the actual and pretended leakage type. Formally,

Assumption 5 (No Rent to the Lowest Type)**.** *For any player i and leakage type t_i, stage k and on-path history h_i^k, it holds*

$$\mathbb{E}_{\mu_i}\left[U_i\left(S_i(\underline{\theta}_i, \hat{t}_i), (S_j(\theta_j, t_j))_{j\neq i}, \underline{\theta}_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] = 0, \forall \hat{t}_i,$$

where $\underline{\theta}_i$ is the lowest value type $\theta_i \in \Theta_i$ for which $Q_i(\theta_i, t_i, t_i, h_i^{k-1}) > 0$.

5.2 Efficient Auctions under Leakages

Because we assume the seller's valuation is zero, efficiency requires that the good be allocated to the bidder with the highest valuation. We say that an auction is efficient under leakages if it has a PBE that yields an efficient allocation for every leakage-type profile in any leakage-type space. Formally,

Definition 7 (Efficiency under Leakage)**.** *The auction G is efficient under leakage if for each leakage type space $\mathcal{T}$, the environment $\Gamma = (G, \mathcal{T})$ admits a PBE (S, μ) that is efficient for all payoff type profiles $\theta \in \Theta$ and all leakage type profiles $t \in T$.*

From our anonymity assumption (Assumption 4), it follows that ties between equal types need to be resolved uniformly and, hence, that any efficient auction satisfies a property that we call allocation invariance under leakages.

Definition 8 (Allocation Invariance under Leakage)**.** *The auction mechanism* (Γ, S, μ) *satisfies allocation invariance under leakage if there is an allocation* $(q_1(\theta), \ldots, q_n(\theta))$ *such that, for all type spaces* $\mathcal{T}$*, the mechanism implements* $(q_1(\theta), \ldots, q_n(\theta))$ *for all* $t \in T$*.*

The following lemma outlines two properties of mechanisms that satisfy allocation invariance under leakage. These two properties will be instrumental for the main result of this section.

Lemma 1 (Increasing Allocation and Payoff Bounds)**.** *Take an auction mechanism* (Γ, S, μ) *and suppose it satisfies the allocation-invariance-under-leakage property. Fix a player* $i \in N$*, a stage* k*, and a private on-path history* h_i^{k-1}*.*

1. *Then,* $Q_i(\theta_{is}, t_i, t_i, h_i^{k-1})$ *is non-decreasing in* θ_{is} *on the set of value types* $\{\theta_i \in \Theta_i : Q_i(\theta_i, t_i, t_i, h_i^{k-1}) > 0\}$*.*

2. *Further, let* m *be such that* $\theta_{im} = \theta_i$ *and let* $\underline{m}$ *be such that* $\theta_{i\underline{m}} = \underline{\theta}_i$*, where* $\underline{\theta}_i$ *is the lowest value type* $\hat{\theta}_i \in \Theta_i$ *for which* $Q_i(\hat{\theta}_i, t_i, t_i, h_i^{k-1}) > 0$*. Then, it holds, for all feasible types* $\hat{t}_i$ *at* h_i^{k-1} *who believe to be weakly slower, that*

$$\begin{aligned}\sum_{s=\underline{m}+1}^{m} (\theta_{is} - \theta_{is-1}) Q_i(\theta_{is-1}, t_i, t_i, h_i^{k-1}) \\ \leq \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i, \hat{t}_i), (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] \\ \leq \sum_{s=\underline{m}+1}^{m} (\theta_{is} - \theta_{is-1}) Q_i(\theta_{is}, t_i, t_i, h_i^{k-1}). \quad (1)\end{aligned}$$

The proof of Lemma 1 first follows standard arguments for discrete value types to show that the interim allocation is increasing at every stage and derive bounds on the payoffs (see Lovejoy, 2006; Bergemann and Pesendorfer, 2007, for static games). The specific payoff bounds for bidders mimicking a slower bidder in (1) are then obtained by using the allocation-invariance-under-leakage property.

The following result is the main result of this section. To establish it, we use the payoff bounds from Lemma 1 together with our assumption of equally spaced value type spaces Θ_i, which allows us to tightly bound the deviation payoffs for a player following their default strategies when the others do so as well. The strategy of the proof is otherwise similar to that in Theorem 1.

Theorem 2 (Efficiency under Leakages)**.** *Suppose the auction G is efficient under leakages. Then, G admits a leakage-proof ϵ-PBE with $\epsilon = 2\delta$ that always results in an efficient outcome.*

5.3 Revenue-Maximizing Auctions under Leakages

We now turn to revenue-maximizing auctions. Throughout, we assume that virtual values are increasing.[16] Recall $\rho_i(\theta_i)$ is the common prior belief about bidder i's value type θ_i. Then,

Assumption 6 (Increasing Virtual Values)**.** *For every bidder $i \in N$, it holds that the virtual value,*

$$v_{is} = \theta_{is} - (\theta_{is+1} - \theta_{is}) \frac{1 - \sum_{m=1}^{s} \rho_i(\theta_{im})}{\rho_i(\theta_{is})},$$

increases in s.

For an auction (G, S, μ), the seller's revenue for a specific type profile t is

$$\Pi(G, S, \mu, t) = \mathbb{E}_\rho \left[\sum_{i \in N} m_i(\theta, t) \right].$$

It is tempting to require, akin to efficiency, that revenue maximization holds across all feasible leakage type profiles; i.e., an auction is revenue-maximizing for all possible leakage types t. However, other than for efficiency, such a robustness requirement is potentially problematic.

To illustrate this, consider the following example of a first-price auction with an optimal reserve price and two bidders who have leakage types that we call *paranoid*. Both believe that they are slower than the other bidder, which, under the particular type distribution assumed, leads them to bid more aggressively than in the default auction, thus raising revenue above the Myersonian upper bound for regular auctions.

Example 2 (Paranoid Bidders in a First-Price Auction)**.** *Consider a first-price auction with a single object for sale and two bidders $i \in \{1, 2\}$. Bidders have independent private values, θ_i, drawn from the Pareto distribution $F(\theta_i) = 1 - (\theta_i + 1/2)^{-2}$ on*

[16] For an analysis of auctions without leakage that allows general distributions and does not require monotonic virtual values, see Jeong and Pycia (2025).

$[1/2, +\infty)$ *and quasi-linear utility. The Myersonian optimal reserve price is* $r^* = \frac{1}{2}$. *Without information leakage, the symmetric equilibrium bidding strategy is given by*

$$\beta_i(\theta_i) = \theta_i - \frac{\int_{r^*}^{\theta_i} F(x)dx}{F(\theta_i)} = \theta_i - \frac{(\theta_i + 1/2)(\theta_i - 1/2)}{\theta_i + 3/2} = \frac{6\theta_i + 1}{4\theta_i + 6}.$$

Now, suppose both bidders are paranoid; i.e., they both believe the other bidder to be faster. The optimization problem for paranoid bidder i *is*

$$\max_b U_i(\theta_i, b) = (\theta_i - b)F(b).$$

We denote the maximizer of this problem as $b_i^*(\theta_i)$. *The first derivative of* U_i *with respect to* b *is given by*

$$\frac{\partial U_i(\theta_i, b)}{\partial b} = \frac{2(\theta_i - b)}{(1/2 + b)^3} - \left(1 - \frac{1}{(1/2 + b)^2}\right) \equiv g(b),$$

and the second derivative is readily verified to be smaller than 0 *when* $b < \theta_i$.

The maximizer $b^*(\theta_i)$ *satisfies the FOC, i.e.,* $g(b_i^*(\theta_i)) = 0$. *To show that* $b_i^*(\theta_i) > \beta_i(\theta_i)$ *for all* $\theta_i > 1/2$, *we only need to check that* $g(\beta_i(\theta_i)) > 0$ *for all* $\theta_i > 1/2$. *Indeed,*

$$\begin{aligned}
&(1/2 + \beta_i(\theta_i))^3 \times g(\beta_i(\theta_i)) \\
=&2(\theta_i - \beta_i(\theta_i)) - (\beta_i(\theta_i) - 1/2)(\beta_i(\theta_i) + 3/2)(\beta_i(\theta_i) + 1/2) \\
=&\frac{(2\theta_i + 1)(2\theta_i - 1)}{2\theta_i + 3} - \frac{2\theta_i - 1}{2\theta_i + 3}\frac{6\theta_i + 5}{2\theta_i + 3}\frac{4\theta_i + 2}{2\theta_i + 3} \\
=&\frac{(4\theta_i^2 - 1)^3}{(2\theta_i + 3)^2} > 0
\end{aligned}$$

Hence, paranoid bidders bid strictly higher than they would bid in the absence of information leakage. This first-price auction with paranoid bidders generates revenue strictly higher than the Myersonian optimal revenue.[17]

On the other hand, if one player believes both players to be equally fast, while the other believes to be faster than the opponent and the opponent to believe everyone is

[17] Having paranoid bidders could also backfire in a first-price auction. For example, if θ_i is drawn from the uniform distribution $F(\theta_i) = \theta_i$ on $[0, 1]$ then, with paranoid bidders, the optimal reserve price is $\sqrt{3}/3$ and the expected revenue is $2\sqrt{3}/9$, which is strictly lower than the Myersonian no-leakage optimal revenue of $5/12$.

equally fast (the one-profile), revenue is below the Myersonian upper bound. But this implies that we cannot rank the first-price auction against the second-price auction, which is leakage-proof and thus always yields the Myersonian upper bound.

Nevertheless, we conclude this section with a positive result: For common priors regarding the leakage order, the static second-price auction is revenue-maximizing. Specifically, let $\gamma : T \to [0,1]$ be the common prior distribution over types in T for a given type $\mathcal{T}$ (shared by the bidders and the seller).

Definition 9 (Revenue Maximization under a Common Prior)**.** *Fix a leakage type space $\mathcal{T}$. Auction (G, S, μ) maximizes revenue under leakages if*

$$\mathbb{E}_\gamma \left[\Pi(G, S, \mu, t)\right] \geq \mathbb{E}_\gamma \left[\Pi((\Gamma', \mathcal{T}), S', \mu', t)\right],$$

for all auctions $((\Gamma', \mathcal{T}), S', \mu')$.

Lemma 2 (Allocation in the Revenue-Maximizing Auction under a Common Prior)**.** *In the revenue-maximizing auction under a common prior, the allocation satisfies invariance under leakage and is as follows: Suppose bidder i with value $\theta_{is} \in \Theta_i$ has the highest virtual value v_{is} among all bidders. If $v_{is} \geq 0$, allocate the good to bidder i; if $v_{is} < 0$, do not allocate the good to anyone. If multiple bidders have the highest virtual value, randomize uniformly among them.*

The proof leverages standard arguments from Myerson (1981) to our setup. The allocation described in Lemma 2 corresponds to that in the second-price auction with an optimal reserve price when leakages are absent. Because the corresponding equilibrium in the second-price auction is in dominant strategies, the second-prize auction is automatically leakage-proof: a dominant strategy remains optimal even after observing any actions of other bidders. We may thus conclude:

Proposition 1 (Revenue-Maximization under Common Priors)**.** *With a common prior, the second-prize auction with an optimal reserve prize maximizes revenue under leakages.*

6 Discussion

In this section, we first discuss the leakage-proofness (or lack thereof) of standard auction formats. Then, we explore the relationship between our leakage-proofness

concept and ex-post incentive compatibility.

6.1 Leakage-Proofness and Standard Auctions

As observed above, if an auction admits a dominant-strategy equilibrium, then it is automatically leakage-proof: a dominant strategy remains optimal even after observing any actions of other bidders. This not only applies to the static second-price auction but also to the English (button) auction, where everyone remains active until the price reaches one's value.

Strategy-proofness clearly fails in the static first-price auction, and so does leakage-proofness. Under the standard assumption of continuously distributed valuations, the unique equilibrium without leakage features strictly increasing continuous bidding strategies (Lebrun, 1996; Maskin and Riley, 2003). Now suppose bidder i has the highest valuation, and all other bidders play their equilibrium strategies. Consider a one-profile where i observes the highest bid and knows they are the fastest (while others believe speeds are symmetric). Bidder i can then profitably bid just above the observed maximum bid, strictly improving upon their no-leakage equilibrium payoff. In other words, player i has a profitable deviation, showing that leakage-proofness fails.

The Dutch auction is likewise not leakage-proof. Under our anonymity assumption, tie-breaking is uniform. This creates incentives for fast bidders to wait slightly longer, as the risk of delaying is partly offset by information leakage. Conversely, slow bidders would optimally accept the price earlier to avoid revealing their bids to faster bidders.

A further observation is that, unlike the second-price and English auctions—which remain strategically equivalent under information leakage—the static first-price and Dutch auctions do not. In the first-price auction, a bidder who can uniquely exploit leaked information will bid just above the highest observed bid if their valuation is higher, guaranteeing a win. In the Dutch auction, such a strategy is impossible: once bidder i observes another bidder accept the price, their best response is to accept as well and hope to win the tie. Although the conditional payment is essentially the same as in the first-price auction, bidder i now wins with probability at most one half.[18]

[18]The only tie-breaking rule under which the two formats would be equivalent is one that always awards the good to the fastest tying bidder. This is excluded by our anonymity assumption, which

Finally, note that the Dutch auction can be made leakage-proof by adjusting the tie-breaking rule. Fast bidders can benefit from leaked information only by tying with slower bidders. Thus, the simplest way to neutralize such leakages is to ensure that fast bidders never win in a tie. Under anonymity, the only way to achieve this is to not allocate the good at all in the event of a tie.[19]

6.2 Leakage-Proofness and Ex-Post Incentive Compatibility

Leakage-proofness requires that a player's incentives are independent of the simultaneous actions of others, making it reminiscent of ex-post incentive compatibility (EPIC). However, as we will see in this section, neither concept implies the other in general. In particular, an EPIC mechanism is leakage-proof if equilibrium play is in pure strategies, but not necessarily under mixed strategies. Conversely, a leakage-proof mechanism is EPIC if it is static, but not if it is dynamic.

We begin with the standard definition of an ex-post incentive-compatible mechanism. The definition is for a mechanism (Γ, S, μ) without information leakage; i.e., for strategies $S_i(\theta_i)$ that only depend on the value type θ_i.

Definition 10 (*EPIC* mechanism)**.** *A mechanism (Γ, S, μ) is an ex-post incentive compatible (EPIC) mechanism if the strategy profile S is ex-post incentive compatible: For each player $i \in N$ and each payoff type profile $\theta \in \Theta$, we have*

$$U_i\left(S_i\left(\theta_i\right), \left(S_j\left(\theta_j\right)\right)_{j\neq i}, \theta\right) \geq U_i\left(\sigma_i, \left(S_j\left(\theta_j\right)\right)_{j\neq i}, \theta\right). \tag{2}$$

for any $\sigma_i \in \Sigma_i(t_i^0)$. If S is a pure-strategy profile, we call (Γ, S, μ) a pure-strategy EPIC mechanism.

Our first result establishes that any pure strategy EPIC mechanism is leakage-proof. The central insight used in the proof is that the ex-post property, stated above for the beginning of the game, holds at any stage and after any on-path history, which allows leakage-proofness to be inferred immediately. For mixed strategies, the implication does not hold, as we show with an example following the statement.

Proposition 2. *Any pure-strategy EPIC mechanism is leakage-proof. Mixed-strategy EPIC mechanisms are not, in general, leakage-proof.*

prevents the mechanism from conditioning on the leakage order.

[19] Gans and Holden (2022) propose a closely related idea in a different context.

The reason mixed-strategy EPIC mechanisms may not be leakage-proof is that observing the realized actions in mixed strategies provides more information than knowing the player's payoff type. For example, consider the classic zero-sum game, matching pennies.[20]

	Heads	Tails
Heads	+1,-1	-1,+1
Tails	-1,+1	+1,-1

The matching pennies game has a unique mixed-strategy Nash equilibrium in which each player plays each action with probability 1/2. The expected payoff for both players is zero. The equilibrium is ex post incentive-compatible because of the degenerate type space. However, it is not leakage-proof. A player who can observe the action of the other one will always win and have a payoff of one.

Regarding the other direction, we also have that leakage-proofness implies EPIC in specific mechanisms but not in others. Here, the relevant distinction is between static and dynamic mechanisms, where a mechanism is static if it has only one stage.

Proposition 3. *Any static leakage-proof mechanism is an EPIC mechanism. Dynamic leakage-proof mechanisms are not, in general, EPIC.*

The proof for the positive claim is in the appendix. To see the negative claim in Proposition 3, recall the Dutch auction with no allocation in case of a tie, which we discussed in Section 6.1 above. This auction is leakage-proof, but it is clearly not EPIC. If you know the other players' types, then you know at what price they would take the good. So, if you are the highest-valuing bidder, you can always increase your payoff by waiting and snatching the good from the second-highest bidder just before they take it.

7 Concluding Remarks

There are many modern market environments in which a mechanism designer has limited control over information leakage between participants — that is, over which actions of others a player may observe when making a decision. Examples include

[20] Although we present the argument with a static game of complete information, the general idea extends to dynamic games of incomplete information.

financial markets, online platforms, and blockchain-based markets. In this paper, we study the implications of such leakages for mechanism design.

Our contributions are threefold. First, we introduce an analytical framework for studying information leakage. The central objective is to design mechanisms that are leakage-proof, meaning that equilibrium outcomes remain invariant to any (consistent) beliefs about the leakage structure. Second, we apply this robustness notion to a general implementation problem and to the specific cases of efficient and revenue-maximizing auctions. In both settings, we show that leakage-proofness is the critical property: if a mechanism is to remain implementable — or an auction efficient or optimal — under potential leakage, then it must be leakage-proof (and conversely in the case of implementation). Finally, we construct a leakage-proof Dutch auction and show that leakage-proofness is an independent property of mechanisms. We further clarify in what sense it is orthogonal to the seemingly related concept of ex-post incentive compatibility.

A Proofs

A.1 Proofs for Section 4

Proof of Theorem 1. The if part from our definition of a leakage-proof equilibrium (Definition 3). For the only-if part, let $t^0 = (t_1^0, ..., t_n^0)$ denote a profile of leakage types where everyone believes it is common knowledge that everyone is equally fast (the zero profile; cf. Definition 2) and let $t^1 = (t_1^0, ..., t_{i-1}^0, t_i^1, t_{i+1}^0, ..., t_n^0)$ denote a generic one profile, corresponding to the leakage order where one player is faster than all the others who are equally fast (again, cf. Definition 2). We will be explicit about which player is fastest below.

Fix any type space $\mathcal{T}$. Because the game G implements the social choice function f, the environment $\Gamma = (G, \mathcal{T})$ admits a PBE (S^*, μ^*) with outcome $f(\theta)$ for all $\theta \in \Theta$ and $t \in T$. The equilibrium (S^*, μ^*) allows us to determine equilibrium play in the default game G, which we call the default strategies. Observe that the private histories of the any player i under the zero-leakage type profile t^0 are identical to public histories, i.e. $h^{k-1} = h_i^{k-1}$. Hence, for all players $i \in N$, all payoff types $\theta_i \in \theta_i$,

and all histories $h \in H$, we may define

$$S_i^0(\theta_i)(h) = S_i^*(\theta_i, t_i^0)(h).$$

For the following, let $\mu_i^d(h_i^{k-1}, t_i)$ be the updated belief of player i with a leakage type t_i when the others follow their default strategies in S^0. Further, let $H_i^{k-1}(t_i)$ be player i's set of all private histories up to (and including) stage $k-1$ when having leakage type t_i, and let $H_i^{k-1}(t_i, \theta_i) \subseteq H^{k-1}(t_i)$ be all the private histories of player i that can be rationalized when player i is of value type θ_i and follows her default strategy, $S_i^0(\theta_i)$, that is,

$$H_i^{k-1}(t_i, \theta_i) = \\ \left\{ h_i^{k-1} \in H^{k-1}(t_i) : \exists \theta_{-i} \in \Theta_{-i} \text{ such that } h_i^{k-1} \subset \text{supp}\left(\zeta((S_j^0(\theta_j))_{j \in N})\right) \right\}.$$

On the other hand, let $\bar{H}_i^{k-1}(t_i, \theta_i) \subseteq H^{k-1}(t_i) \setminus H_i^{k-1}(t_i, \theta_i)$ be the set of histories that cannot be thus rationalized. Because the game G is pruned, the set $\bar{H}_i^{k-1}(t_i, \theta_i)$ consists exactly of those histories in which player i has pretended to be of different value type (and taken an action that is different from her equilibrium default action) for at least one round before, and including, $k-1$.

Now, fix player i with value type θ_i and leakage type t_i, and take a history $h_i^{k-1} \in \bar{H}_i^{k-1}(t_i, \theta_i)$. Suppose, for the time being, that the other players $j \neq i$ all continue following their default strategies, S_j^0, in the all future rounds and that player i has updated her belief to $\mu_i^d(h_i^{k-1}, t_i)$. Because the game G is finite, player i has an optimal continuation strategy in such a situation, which we denote by $\hat{S}_i(\theta_i, t_i)$.

Now, we combine the default strategies with the off-path best response $\hat{S}_i(\theta_i, t_i)$. defined above and consider the following strategy-belief profile (S, μ). For all players $i \in N$, all payoff types $\theta_i \in \theta_i$, all leakage types $t_i \in t_i$ and all private histories $h_i^{k-1} \in H_i^{k-1}(t_i)$, let

$$S_i(\theta_i, t_i)(h_i^{k-1}) = \begin{cases} S_i^0(\theta_i)(h^{k-1}) & \text{if } h_i^{k-1} \in H_i^{k-1}(t_i, \theta_i) \\ \hat{S}_i(\theta_i, t_i)(h_i^{k-1}) & \text{otherwise} \end{cases}$$

and

$$\mu_i(h_i^{k-1}, t_i) = \mu_i^d(h_i^{k-1}, t_i).$$

We want to show that (S, μ) is an equilibrium in the environment $(G, \mathcal{T})$; i.e. that there is no player i, value type θ_i, and leakage type t_i with a deviation σ_i from $S_i(\theta_i, t_i)$ such that

$$\mathbb{E}_{\mu_i}\left[U_i\left(\sigma_i, (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] > \\ \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i, t_i), (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] \quad (3)$$

for some stage k and history h_i^{k-1}. Observe that we have mutual sequential optimality off the path of play by construction. Consequently, we only need to verify on-path histories.

So, fix a type θ_i and any k and $h_i^{k-1} \in H_i^{k-1}(t_i, \theta_i)$. First, consider any type profile t in which player i knows that she is the uniquely fastest player, observing everyone else's actions. We want to argue that, for any such leakage type profile, if all other players follow their default strategies, it is optimal for player i to do so, too.

For the argument, consider the one-profile t^1 in which player i believes to be the uniquely fastest player and the others to be equally slow. Because the leakage type of the slow players in the one profile t^1 is equal to their leakage type in the zero profile t^0 (they believe that everyone is equally fast and that this is common knowledge), and we used t^0 to construct the default strategy profile, the equilibrium strategy S_i^* of the slow players in t^1 is equal to their default strategy.

Consequently, when the others follow their default strategies, then the payoff to player i *under any type profile in which she knows to be the uniquely fastest player* when choosing a strategy σ_i is equal to that of the fastest player in t_1 (as the fast player observes all other actions in either case and draws the same conclusions),

$$\mathbb{E}_{\mu_i}\left[U_i\left(\sigma, (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i^1\right], \quad (4)$$

for any k and h_i^{k-1} on the path of play of the strategy profile S.

In particular, when player i follows the equilibrium strategy of the fastest player in the one-profile, $S_i^*(\theta_i, t_i^1)$, then the payoff is

$$\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i, t_i^1), (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i^1\right] \\ = \mathbb{E}_{\mu_i}\left[u_i(f(\theta_i, \theta_{-i}), (\theta_i, \theta_{-i})) \middle| h_i^{k-1}, t_i^1\right]. \quad (5)$$

At the same time, because all opponents of i follow the default strategy on the path of play, if player i follows the default action, then by construction $f(\theta)$ obtains as an outcome for any given value profile θ, giving us

$$\mathbb{E}_{\mu_i}\left[U_i\left(S_i^0(\theta_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right] = \mathbb{E}_{\mu_i}\left[u_i(f(\theta_i,\theta_{-i}),(\theta_i,\theta_{-i}))\middle|h_i^{k-1},t_i^1\right]. \quad (6)$$

Consequently, $S_i^0(\theta_i)$ yields the upper bound on possible payoffs. As this holds for all value types $\theta_i \in \Theta_i$, following the default strategy is a best response.

Next, consider any type profile in which player i believes to be faster than all other players but one (if there is such a profile in $\mathcal{T}$). From the arguments above, we know that, even if player i were to observe the action of the unobservable player, and no matter what that player does, player i will not want deviate from $S_i^0(\theta_i)$ for any value type θ_i. So, again, that player best responds by playing the default strategy. In fact, we can repeat the above argument for any type profile in which player i believes to be faster than all but some $m \geq 2$ other players, giving us that, no matter how fast a player, choosing the default strategy is a best response. We have, thus, established that (S,μ) is an equilibrium.

To finish the proof, observe that, because in the equilibrium (S,μ), everyone behaves according to S^0 on the path of play, we have shown that (S^0,μ^0), where $\mu_i^0(\theta_i)(h) = \mu_i^*(\theta_i,t_i^0)(h)$ for all $h \in H$, is a leakage-proof equilibrium. □

A.2 Proofs for Section 5

Proof of Lemma 1. For leakage type t_i and the history h_i^{k-1} under consideration, we say that a type $(\hat{\theta}_i,t_i)$ is feasible when $Q_i(\hat{\theta},t_i,t_i,h_i^{k-1}) > 0$. The value types θ_i thus identified correspond to the value types that may still be in the auction at the history under consideration. Incentive compatibility for a (feasible) bidder (θ_i,t_i) mimicking a feasible bidder $\left(\tilde{\theta}_i,t_i\right)$ requires

$$\theta_i Q_i\left(\theta_i,t_i,t_i,h_i^{k-1}\right) - M_i\left(\theta_i,t_i,t_i,h_i^{k-1}\right) \geq \theta_i Q_i\left(\tilde{\theta}_i,t_i,t_i,h_i^{k-1}\right) - M_i\left(\tilde{\theta}_i,t_i,t_i,h_i^{k-1}\right). \quad (7)$$

Similarly, the incentive compatibility for a bidder $\left(\tilde{\theta}_i, t_i\right)$ mimicking a bidder (θ_i, t_i) yields

$$\tilde{\theta}_i Q_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right) - M_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right) \geq \tilde{\theta}_i Q_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) - M_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right). \quad (8)$$

Combining inequality (7) with inequality (8) and rearranging terms, we have

$$\begin{aligned}\tilde{\theta}_i &\left(Q_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) - Q_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right)\right) \\ &\leq M_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) - M_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right) \\ &\leq \theta_i\left(Q_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) - Q_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right)\right). \quad (9)\end{aligned}$$

The outer inequality in (9) requires

$$\left(\theta_i - \tilde{\theta}_i\right)\left(Q_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) - Q_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right)\right) \geq 0,$$

which implies that $Q_i(\cdot)$ is non-decreasing in θ_i, giving us claim (i).

To continue, we we define

$$\begin{aligned}U_i &= \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i, t_i), (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] \\ \tilde{U}_i &= \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\tilde{\theta}_i, t_i), (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right]\end{aligned}$$

and use the facts

$$\begin{aligned}U_i &= \theta_i Q_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) - M_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) \\ \tilde{U}_i &= \theta_i Q_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right) - M_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right)\end{aligned}$$

to re-express the inequalities in (9) as

$$(\tilde{\theta}_i - \theta_i) Q_i\left(\theta_i, t_i, t_i, h_i^{k-1}\right) \leq \tilde{U}_i - U_i \leq (\tilde{\theta}_i - \theta_i) Q_i\left(\tilde{\theta}_i, t_i, t_i, h_i^{k-1}\right). \quad (10)$$

In particular, for $\theta_{is+1} = \tilde{\theta}_i$ and $\theta_{is} = \theta_i$, we obtain

$$(\theta_{is+1} - \theta_{is}) Q_i\left(\theta_{is}, t_i, t_i, h_i^{k-1}\right)$$

$$\leq \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_{is+1},t_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_{is+1},\theta_{-i}\right)\middle|h_i^{k-1},t_i\right]-$$
$$\mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_{is},t_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_{is},\theta_{-i}\right)\middle|h_i^{k-1},t_i\right]$$
$$\leq (\theta_{is+1}-\theta_{is})Q_i\left(\theta_{is+1},t_i,t_i,h_i^{k-1}\right). \quad (11)$$

Adding up these inequalities from $s=\underline{m}$ to $s=m-1$ and using Assumption 5 then yields

$$\sum_{s=\underline{m}+1}^{m}(\theta_{is}-\theta_{is-1})Q_i(\theta_{is-1},t_i,t_i,h_i^{k-1})$$
$$\leq \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i,t_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i\right]$$
$$\leq \sum_{s=\underline{m}+1}^{m}(\theta_{is}-\theta_{is-1})Q_i(\theta_{is},t_i,t_i,h_i^{k-1}).$$

But then, it also holds

$$\sum_{s=\underline{m}+1}^{m}(\theta_{is}-\theta_{is-1})Q_i(\theta_{is-1},\hat{t}_i,t_i,h_i^{k-1})$$
$$\leq \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i,\hat{t}_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i\right]$$
$$\leq \sum_{s=\underline{m}+1}^{m}(\theta_{is}-\theta_{is-1})Q_i(\theta_{is},\hat{t}_i,t_i,h_i^{k-1}),$$

for feasible $\hat{t}_i$ that believe to be slower than t_i, because the outer expressions in the inequalities above correspond to the bounds on the expected utility of type $(\theta_i,\hat{t}_i)$, conditional on the information of the faster leakage type t_i. To finish the proof, we then recall that the allocation-invariance under leakage property implies $Q_i(\theta_i,t_i,t_i,h_i^{k-1})=Q_i(\theta_i,\hat{t}_i,t_i,h_i^{k-1})$ for all feasible t_i and $\hat{t}_i$, giving us claim (ii). □

Proof of Theorem 2. As in the proof to Theorem 1, let $t^0=(t_1^0,...,t_n^0)$ denote a profile of leakage types where everyone believes it is common knowledge that everyone is equally fast (the zero profile; cf. Definition 2) and let $t^1=(t_1^0,...,t_{i-1}^0,t_i^1,t_{i+1}^0....,t_n^0)$ denote a generic one profile, corresponding to the leakage order where one player is faster than all the others who are equally fast (again, cf. Definition 2). We will be explicit about which player is fastest below.

Fix any type space $\mathcal{T}$. By assumption, the environment $\Gamma=(G,\mathcal{T})$ admits a

PBE (S^*, μ^*) that is efficient for all $\theta \in \Theta$ and $t \in T$. The equilibrium (S^*, μ^*) allows us to determine equilibrium play in the default game G, which we call the default strategies. Observe that the private histories of the any player i under the zero-leakage type profile t^0 are identical to public histories, i.e. $h^{k-1} = h_i^{k-1}$. Hence, for all players $i \in N$, all payoff types $\theta_i \in \theta_i$, and all histories $h \in H$, we may define

$$S_i^0(\theta_i)(h) = S_i^*(\theta_i, t_i^0)(h).$$

For the following, let $\mu_i^d(h_i^{k-1}, t_i)$ be the updated belief of player i with a leakage type t_i when the others follow their default strategies in S^0. Further, let $H_i^{k-1}(t_i)$ be player i's set of all private histories up to (and including) stage $k-1$ when having leakage type t_i, and let $H_i^{k-1}(t_i, \theta_i) \subseteq H^{k-1}(t_i)$ be all the private histories of player i that can be rationalized when player i is of value type θ_i and follows her default strategy, $S_i^0(\theta_i)$, that is,

$$H_i^{k-1}(t_i, \theta_i) = \\ \left\{ h_i^{k-1} \in H^{k-1}(t_i) : \exists \theta_{-i} \in \Theta_{-i} \text{ such that } h_i^{k-1} \subset \text{supp}\left(\zeta((S_j^0(\theta_j))_{j \in N})\right) \right\}.$$

On the other hand, let $\bar{H}_i^{k-1}(t_i, \theta_i) \subseteq H^{k-1}(t_i) \setminus H_i^{k-1}(t_i, \theta_i)$ be the set of histories that cannot be thus rationalized. Because the game G is pruned, the set $\bar{H}_i^{k-1}(t_i, \theta_i)$ consists exactly of those histories in which player i has pretended to be of different value type (and taken an action that is different from her equilibrium default action) for at least one round before, and including, $k-1$.

Now, fix player i with value type θ_i and leakage type t_i, and take a history $h_i^{k-1} \in \bar{H}_i^{k-1}(t_i, \theta_i)$. Suppose, for the time being, that the other players $j \neq i$ all continue following their default strategies, S_j^0, in all future rounds and that player i has updated her belief to $\mu_i^d(h_i^{k-1}, t_i)$. Because the game G is finite, player i has an optimal continuation strategy in such a situation, which we denote by $\hat{S}_i(\theta_i, t_i)$.

Now, we combine the default strategies with the off-path best response $\hat{S}_i(\theta_i, t_i)$. defined above and consider the following strategy-belief profile (S, μ). For all players $i \in N$, all payoff types $\theta_i \in \theta_i$, all leakage types $t_i \in t_i$ and all private histories

$h_i^{k-1} \in H_i^{k-1}(t_i)$, let

$$S_i(\theta_i, t_i)(h_i^{k-1}) = \begin{cases} S_i^0(\theta_i)(h^{k-1}) & \text{if } h_i^{k-1} \in H_i^{k-1}(t_i, \theta_i) \\ \hat{S}_i(\theta_i, t_i)(h_i^{k-1}) & \text{otherwise} \end{cases}$$

and

$$\mu_i(h_i^{k-1}, t_i) = \mu_i^d(h_i^{k-1}, t_i).$$

We want to show that (S, μ) is an ϵ-PBE in the environment $(G, \mathcal{T})$; i.e. that there is an $\epsilon > 0$ such that there is no player i, value type θ_i, and leakage type t_i with a deviation σ_i from $S_i(\theta_i, t_i)$ such that

$$\mathbb{E}_{\mu_i}\left[U_i\left(\sigma_i, (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] > \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_i, t_i), (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i\right] + \epsilon \quad (12)$$

for any stage k and history h_i^{k-1}. Observe that we have mutual sequential optimality off the path of play by construction. Consequently, we only need to verify on-path histories.

So, fix a type θ_i and any k and $h_i^{k-1} \in H_i^{k-1}(t_i, \theta_i)$. First, consider any type profile t in which player i knows that she is the uniquely fastest player, observing everyone else's actions. We want to argue that, for any such leakage type profile, if all other players follow their default strategies, it is optimal for player i to do so, too.

For the argument, consider the one-profile t^1 in which player i believes to be the uniquely fastest player and the others to be equally slow. Because the leakage type of the slow players in the one profile t^1 is equal to their leakage type in the zero profile t^0 (they believe that everyone is equally fast and that this is common knowledge), and we used t^0 to construct the default strategy profile, the equilibrium strategy S_i^* of the slow players in t^1 is equal to their default strategy.

Consequently, when the others follow their default strategies, then the payoff to player i *under any type profile in which she knows to be the uniquely fastest player* when choosing a strategy σ_i is equal to that of the fastest player in t_1 (as the fast player observes all other actions in either case and draws the same conclusions),

$$\mathbb{E}_{\mu_i}\left[U_i\left(\sigma, (S_j(\theta_j, t_j))_{j\neq i}, \theta_i, \theta_{-i}\right) \middle| h_i^{k-1}, t_i^1\right], \quad (13)$$

for any k and h_i^{k-1} on the path of play of the strategy profile S.

In particular, when player i follows the equilibrium strategy of the fastest player in the one-profile, $S_i^*(\theta_i, t_i^1)$, then the payoff is

$$\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_i^1),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right].$$

On the other hand, because all opponents of i follow the default strategy on the path of play, if player i follows the default action, then the payoff is

$$\begin{aligned}&\mathbb{E}_{\mu_i}\left[U_i\left(S_i^0(\theta_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right]\\&\qquad=\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_0),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right].\end{aligned}$$

We want to argue that

$$\begin{aligned}&\mathbb{E}_{\mu_i}\left[U_i\left(S_i^0(\theta_i),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right]+2\delta\geq\\&\qquad\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_i^1),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right].\end{aligned}$$

Indeed, the allocation-invariance-under-leakage property together with the payoff bounds from Claim (ii) in Lemma 1 imply

$$\begin{aligned}&\sum_{s=\underline{m}+1}^{m}(\theta_{is}-\theta_{is-1})\left[Q_i(\theta_{is-1},t_i^1,t_i^1,h_i^{k-1})-Q_i(\theta_{is},t_i^1,t_i^1,h_i^{k-1})\right]\\&\quad\leq\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_i^0),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right]\\&\qquad-\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_i^1),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right]\\&\qquad\leq\sum_{s=\underline{m}+1}^{m}(\theta_{is}-\theta_{is-1})\left[Q_i(\theta_{is},t_i^1,t_i^1,h_i^{k-1})-Q_i(\theta_{is-1},t_i^1,t_i^1,h_i^{k-1})\right],\end{aligned}$$

where the definitions of m and $\underline{m}$ are as in Lemma 1.

Now, recall that $|\theta_i-\theta_i'|\leq\delta$ for all $\theta_i,\theta_i'\in\Theta_i$. Then, because Q_i is non-decreasing in θ_i on the set of value types $\{\theta_i\in\Theta_i: Q_i(\theta_i,t_i,t_i,h_i^{k-1})>0\}$ (cf. the first claim in Lemma 1) and its value lies between zero and one, above inequalities imply

$$\begin{aligned}-\delta\leq{}&\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_i^0),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right]\\&-\mathbb{E}_{\mu_i}\left[U_i\left(S_i^*(\theta_i,t_i^1),(S_j(\theta_j,t_j))_{j\neq i},\theta_i,\theta_{-i}\right)\middle|h_i^{k-1},t_i^1\right]\leq\delta,\quad(14)\end{aligned}$$

as desired.

Next, consider any type profile in which player i believes to be faster than all other players but one (if there is such a profile in $\mathcal{T}$). From the arguments above, we know that, even if player i were to observe the action of the unobservable player, and no matter what that player does, player i cannot gain more than 2δ by deviating from $S_i^0(\theta_i)$ for any value type θ_i. So, again, that player approximately best responds by playing the default strategy. In fact, we can repeat the above argument for any type profile in which player i believes to be faster than all but some $m \geq 2$ other players, giving us that, no matter how fast a player, choosing the default strategy is an approximate best response. We have, thus, established that (S, μ) is an $\epsilon-$PBE with $\epsilon = 2\delta$.

To finish the proof, observe that, because in the equilibrium (S, μ), everyone behaves according to S^0 on the path of play, we have shown that (S^0, μ^0), where $\mu_i^0(\theta_i)(h) = \mu_i^*(\theta_i, t_i^0)(h)$ for all $h \in H$, is a leakage-proof ϵ-PBE with $\epsilon = 2\delta$ that is efficient. □

Proof of Lemma 2. From the proof of Lemma 1, we know the maximum expected payment from a bidder i having value θ_{is} is equal to

$$\begin{aligned} M_i\left(\theta_{is}, t_i\right) &= \theta_{is} Q_i(\theta_{is}, t_i) - \mathbb{E}_{\mu_i}\left[U_i\left(S_i(\theta_{is}, t_i), \left(S_j(\theta_j, t_j)\right)_{j \neq i}, \theta_i, \theta_{-i}\right)\right] \\ &= \theta_{is} Q_i(\theta_{is}, t_i) - \sum_{m=1}^{s-1} \left(\theta_{im+1} - \theta_{im}\right) Q_i(\theta_{is}, t_i), \end{aligned}$$

where $Q_i(\theta_{is}, t_i) = Q_i\left(\theta_{is}, t_i, t_i, h_\emptyset\right)$ is the expected allocation of type (θ_{is}, t_i) at the onset of the auction. Letting $m_i = |\Theta_i|$, the expected revenue is

$$\begin{aligned} &\mathbb{E}_\gamma\left[\mathbb{E}_\rho\left[\sum_{i \in N} M_i\left(\theta_i, t_i\right)\right]\right] \\ &= \mathbb{E}_\gamma\left[\sum_{i \in N}\left[\sum_{m=1}^{m_i} \rho_i\left(\theta_{im}\right) M_i\left(\theta_{im}, t_i\right)\right]\right] \\ &= \mathbb{E}_\gamma\left[\sum_{i \in N}\left[\sum_{m=1}^{m_i} \rho_i\left(\theta_{im}\right)\left[\theta_{im} Q_i\left(\theta_{im}, t_i\right) - \sum_{\hat{m}=1}^{m-1}\left(\theta_{i\hat{m}+1} - \theta_{i\hat{m}}\right) Q_i\left(\theta_{i\hat{m}}, t_i\right)\right]\right]\right] \end{aligned}$$

$$
\begin{aligned}
&= \mathbb{E}_\gamma \left[\sum_{i \in N} \left[\sum_{m=1}^{m_i} \rho_i(\theta_{im}) \theta_{im} Q_i(\theta_{im}, t_i) - \sum_{m=1}^{m_i} \rho_i(\theta_{im}) \sum_{\hat{m}=1}^{m-1} (\theta_{i\hat{m}+1} - \theta_{i\hat{m}}) Q_i(\theta_{i\hat{m}}, t_i) \right] \right] \\
&= \mathbb{E}_\gamma \left[\sum_{i \in N} \left[\sum_{m=1}^{m_i} \rho_i(\theta_{im}) \theta_{im} Q_i(\theta_{im}, t_i) - \sum_{m=1}^{m_i} (\theta_{im+1} - \theta_{im}) Q_i(\theta_{im}, t_i) \sum_{\hat{m}=m+1}^{m_i} \rho_i(\theta_{i\hat{m}}) \right] \right] \\
&= \mathbb{E}_\gamma \left[\sum_{i \in N} \left[\sum_{m=1}^{m_i} \rho_i(\theta_{im}) Q_i(\theta_{im}, t_i) \left[\theta_{im} - (\theta_{im+1} - \theta_{im}) \frac{\sum_{\hat{m}=m+1}^{m_i} \rho_i(\theta_{i\hat{m}})}{\rho_i(\theta_{im})} \right] \right] \right] \\
&= \mathbb{E}_\gamma \left[\sum_{i \in N} \left[\sum_{m=1}^{m_i} \rho_i(\theta_{im}) Q_i(\theta_{im}, t_i) \left[\theta_{im} - (\theta_{im+1} - \theta_{im}) \frac{1 - \sum_{\hat{m}=1}^{m} \rho_i(\theta_{i\hat{m}})}{\rho_i(\theta_{im})} \right] \right] \right] \\
&= \mathbb{E}_\gamma \left[\sum_{i \in N} \left[\sum_{m=1}^{m_i} \rho_i(\theta_{im}) \sum_{\theta_{-i}} \rho_{-i}(\theta_{-i}) \mathbb{E}_{\gamma_{-i}} \left[q_i((\theta_{im}, \theta_{-i}), (t_i, t_{-i})) \right] \times \right. \right. \\
&\qquad \left. \left. \left[\theta_{im} - (\theta_{im+1} - \theta_{im}) \frac{1 - \sum_{\hat{m}=1}^{m} \rho_i(\theta_{i\hat{m}})}{\rho_i(\theta_{im})} \right] \right] \right] \\
&= \mathbb{E}_\gamma \left[\sum_{\theta \in \Theta} \rho(\theta) \sum_{i \in N} \mathbb{E}_{\gamma_{-i}} \left[q_i(\theta, (t_i, t_{-i})) \right] \left[\theta_{im} - (\theta_{im+1} - \theta_{im}) \frac{1 - \sum_{\hat{m}=1}^{m} \rho_i(\theta_{i\hat{m}})}{\rho_i(\theta_{im})} \right] \right] \\
&= \mathbb{E}_\gamma \left[\sum_{\theta \in \Theta} \rho(\theta) \sum_{i \in N} q_i(\theta, t)) \left[\theta_{im} - (\theta_{im+1} - \theta_{im}) \frac{1 - \sum_{\hat{m}=1}^{m} \rho_i(\theta_{i\hat{m}})}{\rho_i(\theta_{im})} \right] \right].
\end{aligned}
$$

The third-to-last equality follows because value and leakage types are independent (recall γ_{-i} is i's prior about T_{-i}). The second-to-last equality follows because value types are independent. And the last equality follows from the law of iterated expectations. In the final expression, the term in the square brackets corresponds to the virtual valuation v_{im}, which is independent of t, thus giving us the claim. □

A.3 Proofs for Section 6

For the proofs in this section, we make use of following useful observation.

Observation 1. Any EPIC mechanism (Γ, S, μ) is sequentially *ex-post* incentive compatible on the equilibrium path: For each payoff type profile $\theta \in \Theta$, each stage history $h \subset z$ where z is in the support of $\zeta((S_j(\theta_j))_{j \in N})$ and each player $i \in P(h)$, we have

$$
U_i^h(S_i(\theta_i), S_{-i}(\theta_{-i}), \theta) \geq U_i^h(\sigma_i, S_{-i}(\theta_{-i}), \theta)
$$

for any σ_i, where $U_i^h(.)$ denotes player i's continuation utility under the respective

strategy and type profiles.

Proof of Observation 1. Suppose not. Then, there is a player i, a type θ_i, a stage history $h \subset z$ where z is in the support of $\zeta\left(\left(S_j(\theta_j)\right)_{j\in N}\right)$ and a deviation σ_i, such that

$$U_i^h\left(S_i\left(\theta_i\right), S_{-i}\left(\theta_{-i}\right), \theta\right) < U_i^k\left(\sigma_i, S_{-i}\left(\theta_{-i}\right), \theta\right).$$

We construct a strategy $S_i'\left(\theta_i\right)$ that is a profitable deviation from $S_i\left(\theta_i\right)$. Let i play $S_i\left(\theta_i\right)$ unless they encounter h, in which case they play σ_i from that point on. Formally,

$$S_i'\left(\theta_i\right) = \begin{cases} \sigma_i & \text{at } h' \text{ if } h \subseteq h', \\ S_i\left(\theta_i\right) & \text{otherwise.} \end{cases}$$

Because h is on the equilibrium path, the path-of-play passes through h with positive probability. Then, we have

$$U_i\left(S_i\left(\theta_i\right), S_{-i}\left(\theta_{-i}\right), \theta\right) < U_i\left(\sigma_i, S_{-i}\left(\theta_{-i}\right), \theta\right).$$

But this says that S is not EPIC, giving us the desired contradiction. □

Proof of Proposition 2. Fix a pure strategy EPIC mechanism (Γ, S, μ). We want to show that S is a leakage-proof equilibrium in G. Because S is EPIC, Observation 1 gives that, for each payoff type profile $\theta \in \Theta$, each stage history $h \subset z$ where z is in the support of $\zeta\left(\left(S_j(\theta_j)\right)_{j\in N}\right)$ and each player $i \in P\left(h\right)$, we have

$$U_i^h\left(S_i\left(\theta_i\right), S_{-i}\left(\theta_{-i}\right), \theta\right) \geq U_i^h\left(\sigma_i, S_{-i}\left(\theta_{-i}\right), \theta\right)$$

for any σ_i. But then, because we are dealing with pure strategies, $S_i\left(\theta_i\right)$ being a best response irrespective of the payoff types of other players implies that player i does not want to deviate even if i observes the actions of other players, as long as they are playing according to the pure strategy profile S_{-i}. In other words, leakage will not lead to profitable deviations, thereby providing leakage-proofness as desired. □

Proof of Proposition 3. By contraposition. Let (Γ, S, μ) be a static mechanism that is not EPIC. By Observation 1, there is a player $i \in N$, a type profile $\theta \in \Theta$, and a

deviation σ_i such that

$$U_i^{h_\emptyset}\left(S_i\left(\theta_i\right), S_{-i}\left(\theta_{-i}\right), \theta\right) < U_i^{\emptyset}\left(\sigma_i, S_{-i}\left(\theta_{-i}\right), \theta\right)$$

In particular, there exists at least one action a_{-i} in the support of $\left(S_j\left(\theta_j\right)\left(h_\emptyset\right)\right)_{j\neq i}$ for which player i optimally chooses an action that is not in the support of $S_i\left(\theta_i\right)\left(h_\emptyset\right)$. Now, by Assumption 2, there is a type profile in which player i is the uniquely fastest player and all other players believe that everyone is equally slow. Under such a profile, all players except i play according to $\left(S_j\left(\theta_j\right)\left(h_\emptyset\right)\right)_{j\neq i}$. So, for the private history $h_i = \{a_{-i}\}$, player i would like to choose an action that is different from what S_i prescribes, violating leakage-proofness and, thus, giving us the claim. □

B The Role of Pruning and a Minimally Rich Leakage Type Space

In this section, we provide an example showing that if we drop Assumptions 2 (Minimally Rich Type Space) and 3 (Pruning), then Theorem 1 fails: Even if the same social choice function f is implemented under any leakage order, there may not exist a leakage-proof equilibrium in the game G which implements f.

Example 3. *There are two players. Player 1 has two possible types:* $\Theta_1 = \{\theta_H, \theta_L\}$ *with equal probability, and player 2 has one type:* $\Theta_2 = \{\theta_2\}$. *There are 5 possible outcomes:* $X = \{x, y, z, m, n\}$.

The following table depicts, for each outcome in X *and type profile* $\theta \in \Theta_1 \times \Theta_2$, *the payoffs* (u_1, u_2) *of the two players (*u_1 *corresponds to player* 1*'s payoff and* u_2 *corresponds to player* 2*'s payoff):*

x	θ_2
θ_H	(2,1)
θ_L	(0,1)

y	θ_2
θ_H	(0,1)
θ_L	(2,1)

z	θ_2
θ_H	(-2,-2)
θ_L	(-2,-2)

m	θ_2
θ_H	(-2,2)
θ_L	(-2,2)

n	θ_2
θ_H	(2,-2)
θ_L	(2,-2)

The social choice function f *is given by:*

f	θ_2
θ_H	x
θ_L	y

The game G is a static game as follows (each box specifies the outcome resulting from each pair of actions):

G	a^*	a^*_{2H}	a^*_{2L}	a'_{2H}	a'_{2L}
a_{1H}	x	z	z	m	z
a_{1L}	y	z	z	z	m
a^*_{1H}	z	x	z	n	z
a^*_{1L}	z	z	y	z	n

For the sake of the argument, we consider a setup in which leakage orders are common knowledge. Let us list all the possible leakage orders and the corresponding equilibria:[21]

$\precsim_0 \hat{=} 1 \sim 2$**:** This corresponds to the leakage order under which the two players move simultaneously without observing each other's action. Clearly,

$$S_1(\theta_H, \precsim_0) = a_{1H}, \quad S_1(\theta_L, \precsim_0) = a_{1L}, \quad S_2(\theta_2, \precsim_0) = a^*$$

is a Bayesian Nash equilibrium strategy profile, which implements f.

$\precsim_1 \hat{=} 1 \succ 2$**:** Here, player 1 can observe player 2's action. The following strategy profile is a perfect Bayesian equilibrium which implements f:

$$S_1(\theta_1, \precsim_1)(a_2) = \begin{cases} a_{1H} & \text{if } a_2 = a^*, \theta_1 = \theta_H \\ a_{1L} & \text{if } a_2 = a^*, \theta_1 = \theta_L \\ a^*_{1H} & \text{if } a_2 \in \{a^*_{2H}, a'_{2H}\} \\ a^*_{1L} & \text{if } a_2 \in \{a^*_{2L}, a'_{2L}\}, \end{cases} \qquad S_2(\theta_2, \precsim_1) = a^*.$$

$\precsim_2 \hat{=} 1 \prec 2$**:** Here, player 2 can observe player 1's action. The following strategy profile is a perfect Bayesian equilibrium which implements f:

$$S_1(\theta_1, \precsim_2) = \begin{cases} a^*_{1H} & \text{if } \theta_1 = \theta_H \\ a^*_{1L} & \text{if } \theta_1 = \theta_L, \end{cases} \quad S_2(\theta_2, \precsim_2)(a_1) = \begin{cases} a'_{2H} & \text{if } a_1 = a_{1H} \\ a'_{2L} & \text{if } a_1 = a_{1L} \\ a^*_{2H} & \text{if } a_1 = a^*_{1H} \\ a^*_{2L} & \text{if } a_1 = a^*_{1L} \end{cases}.$$

[21] For simplicity, we omit the belief system, which can be easily derived from the strategy profile.

Consequently, the game G implements the same social function f under any leakage order. Nonetheless, it lacks a leakage-proof equilibrium. The intuition is that player 1 has costly precautionary actions, a_{1H}^* and a_{1L}^*. These actions are worthwhile if information actually leaks to player 2, but may have negative consequences if no leakage occurs. As a result, the equilibrium fails to be leakage-proof.

If we prune the game, this issue disappears. Specifically, pruning to the equilibrium under no leakage, $\precsim_0$, restricts player 1 to $\{a_{1H}, a_{1L}\}$ and player 2 to a^*. In this reduced game, the equilibrium outcome is $f(\theta)$ regardless of the leakage order ($\precsim_0$, $\precsim_1$, or $\precsim_2$). Pruning thus simplifies the strategic environment under leakages.

In contrast, assuming a minimally rich type space prevents $f(\theta)$ from being implemented under leakages. For example, consider a type profile where player 1 believes they are as slow as player 2, while player 2 knows both that they are faster and what player 1 believes. In this case, the outcome will be m under any possible value type profile θ. The minimally rich type space assumption, therefore, restricts the strategic settings to which our analysis applies.

References

Abe, M. and K. Suzuki (2002): "M+ 1-st price auction using homomorphic encryption," in *International workshop on public key cryptography*, Springer, 115–124.

Akbarpour, M. and S. Li (2020): "Credible Auctions: A Trilemma," *Econometrica*, 88, 425–467.

Baldauf, M. and J. Mollner (2022): "Fast traders make a quick buck: The role of speed in liquidity provision," *Journal of Financial Markets*, 58, 100621.

Banchio, M., A. Skrzypacz, and F. Yang (2025): "Dynamic Threats to Credible Auctions," in *Proceedings of the 26th ACM Conference on Economics and Computation*, 186–186.

Bergemann, D. and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

Bergemann, D. and M. Pesendorfer (2007): "Information Structures in Optimal Auctions," *Journal of Economic Theory*, 137, 580–609.

Bikhchandani, S., S. A. Lippman, and R. Ryan (2002): "On the Right-of-first-refusal," *Available at SSRN 621181*.

Blass, E.-O. and F. Kerschbaum (2018): "Strain: A Secure Auction for Blockchains," in *Computer Security*, ed. by J. Lopez, J. Zhou, and M. Soriano, Cham: Springer International Publishing, Lecture Notes in Computer Science, 87–110.

——— (2020): "BOREALIS: Building Block for Sealed Bid Auctions on Blockchains," in *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, New York, NY, USA: Association for Computing Machinery, ASIA CCS '20, 558–571.

Börgers, T. and J. Li (2019): "Strategically simple mechanisms," *Econometrica*, 87, 2003–2035.

Budish, E., P. Cramton, and J. Shim (2015): "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *The Quarterly Journal of Economics*, 130, 1547–1621.

Burguet, R. and M. K. Perry (2009): "Preferred suppliers in auction markets," *The Rand journal of economics*, 40, 283–295.

Carroll, G. (2015): "Robustness and linear contracts," *American Economic Review*, 105, 536–563.

Chassang, S. (2013): "Calibrated incentive contracts," *Econometrica*, 81, 1935–1971.

Choi, A. H. (2009): "A rent extraction theory of right of first refusal," *The Journal of Industrial Economics*, 57, 252–262.

Chu, L. Y. and Z.-J. M. Shen (2006): "Agent competition double-auction mechanism," *Management Science*, 52, 1215–1222.

Daley, B. and B. Green (2012): "Waiting for News in the Market for Lemons," *Econometrica*, 80, 1433–1504.

Dasgupta, P., P. Hammond, and E. Maskin (1979a): "The implementation of social choice rules: Some general results on incentive compatibility," *The Review of Economic Studies*, 46, 185–216.

——— (1979b): "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *The Review of Economic Studies*, 46, 185–216.

Doran, N. (2018): "Asymmetric, Right of First Refusal Auctions," Working Paper.

Eskandari, S., S. Moosavi, and J. Clark (2020): "SoK: Transparent Dishonesty: Front-Running Attacks on Blockchain," in *Financial Cryptography and Data Security*, ed. by A. Bracciali, J. Clark, F. Pintore, P. B. Rønne, and M. Sala, Cham: Springer International Publishing, Lecture Notes in Computer Science, 170–189.

Ewerhart, C. and H. Zeng (2025): "Mediated Subgame Perfect Equilibrium," Working Paper.

Eyster, E. and M. Piccione (2013): "An approach to asset pricing under incomplete and diverse perceptions," *Econometrica*, 81, 1483–1506.

Franklin, M. and M. Reiter (1996): "The Design and Implementation of a Secure Auction Service," *IEEE Transactions on Software Engineering*, 22, 302–312.

Fudenberg, D. and J. Tirole (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236–260.

Galal, H. S. and A. M. Youssef (2019): "Verifiable Sealed-Bid Auction on the Ethereum Blockchain," in *Financial Cryptography and Data Security*, ed. by A. Zohar, I. Eyal, V. Teague, J. Clark, A. Bracciali, F. Pintore, and M. Sala, Berlin, Heidelberg: Springer, Lecture Notes in Computer Science, 265–278.

Gans, J. S. and R. Holden (2022): "A Solomonic Solution to Ownership Disputes: Theory and Applications," .

Garratt, R. J., M. J. Lee, A. Martin, and R. M. Townsend (2019): "Who Sees the Trades? The Effect of Information on Liquidity in Inter-dealer Markets," Tech. rep., Federal Reserve Bank of New York, staff Report No. 892.

Gehrlein, J., S. Häfner, and J. Oechssler (2025): "The Candle Auction in the Field and the Lab," Working Paper.

Häfner, S. and A. Stewart (2025): "Front-Running and Candle Auctions," .

Harrison, J. M. and D. Kreps (1979): "Martingales and arbitrage in multiperiod securities markets," *Journal of Economic Theory*, 20, 381–408.

Heidhues, P., B. Kőszegi, and P. Strack (2018): "Unrealistic expectations and misguided learning," *Econometrica*, 86, 1159–1214.

Hortaçsu, A. and J. Kastl (2012): "Valuing dealers' informational advantage: A study of Canadian treasury auctions," *Econometrica*, 80, 2511–2542.

Hortaçsu, A. and S. Sareen (2004): "Order Flow and the Formation of Dealer Bids in Treasury Auctions," 2004 Meeting Papers 50, Society for Economic Dynamics.

Hurwicz, L. (1972): "On informationally decentralized systems," *Decision and organization: A volume in Honor of J. Marschak*.

Jantschgi, S., H. H. Nax, B. Pradelski, and M. Pycia (2024): "Double auctions and transaction costs," *CEPR Discussion Papers*.

Jeong, B.-h. J. and M. Pycia (2025): "First-Price Principle and the Failure of Revenue Equivalence," Working Paper.

Komo, A., S. D. Kominers, and T. Roughgarden (2024): "Shill-proof auctions," *arXiv preprint arXiv:2404.00475*.

Kudo, M. (1998): "Secure electronic sealed-bid auction protocol with public key cryptography," *IEICE transactions on fundamentals of electronics, communications and computer sciences*, 81, 20–27.

Lebrun, B. (1996): "Existence of an Equilibrium in First Price Auctions," *Economic Theory*, 7, 421–443.

Ledyard, J. (1978): "Incentive compatibility and incomplete information," *Journal of Economic Theory*, 18, 171–189.

Li, S. (2017): "Obviously Strategy-Proof Mechanisms," *American Economic Review*, 107, 3257–3287.

Li ⓡ, J. and P. Dworczak (2024): "Are Simple Mechanisms Optimal when Agents are Unsophisticated?" .

Lovejoy, W. S. (2006): "Optimal Mechanisms with Finite Agent Types," *Management Science*, 52, 788–803.

Madarász, K. (2011): "Information Projection: Model and Applications," *The Review of Economic Studies*, 79, 961–985.

——— (2012): "Information projection: Model and applications," *The Review of Economic Studies*, 79, 961–985.

Madarász, K. and A. Prat (2017): "Sellers with misspecified models," *Review of Economic Studies*, 84, 790–815.

Madarász, K. and M. Pycia (2022): "Towards a resolution of the privacy paradox," CEPR Discussion Paper No. DP16873.

——— (2025): "Cost over Content: Information Choice in Trade," Working Paper.

Maskin, E. and J. Riley (2003): "Uniqueness of Equilibrium in Sealed High-Bid Auctions," *Games and Economic Behavior*, 45, 395–409.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Parkes, D. C., M. O. Rabin, S. M. Shieber, and C. A. Thorpe (2006): "Practical secrecy-preserving, verifiably correct and trustworthy auctions," in *Proceedings of the 8th international conference on Electronic commerce: The new e-commerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet*, 70–81.

Pearce, D. G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.

Penta, A. and P. Zuazo-Garin (2022): "Rationalizability, Observability, and Common Knowledge," *The Review of Economic Studies*, 89, 948–975.

Pycia, M. and P. Troyan (2023): "A Theory of Simplicity in Games and Mechanism Design," *Econometrica*, 91, 1495–1526.

Riley, J. G. and W. F. Samuelson (1981): "Optimal auctions," *The American Economic Review*, 71, 381–392.

Roth, A. E. and A. Ockenfels (2002): "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92, 1093–1103.

Selten, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.

Shimoji, M. and J. Watson (1998): "Conditional Dominance, Rationalizability, and Game Forms," *Journal of Economic Theory*, 83, 161–195.

Siniscalchi, M. (2008): "Epistemic game theory: Beliefs and types," in *The New Palgrave Dictionary of Economics*, Springer, 1–7.

Solan, E. and L. Yariv (2004): "Games with Espionage," *Games and Economic Behavior*, 47, 172–199.

Wolitzky, A. (2016): "Mechanism design with maxmin agents: Theory and an application to bilateral trade," *Theoretical Economics*, 11, 971–1004.

Woodward, K. (2020): "Self-auditable auctions," Working Paper.

Zeng, H. (2025): "Identity-Compatible Auctions," Working Paper.