Supplemental Appendix: Contextually Private Mechanisms

By Andreas Haupt and Zoë Hitzig

This supplemental appendix contains additional impossibility results in non-auction domains, and discusses the overdescending-join auction and the relative-informativeness order.

Impossibility Results in Non-Auction Domains

Many social choice rules exhibit collective but not individual pivotality on some subset of the type space. We give some examples of such choice rules commonly studied in the literature and used in practice which, by Corollary 1, do not admit a fully contextually private protocol. That is, any protocol that computes these choice rules must produce contextual privacy violations.

SA.1. Housing Assignment

Consider first the house assignment problem Shapley and Scarf (1974). All agents are initially endowed with an object from C. The type space is the universal domain of all preferences over C. Denote the initial assignment by an injective function $e\colon N \to C$, where $e(i) \in C$ refers to agent i's initial endowment. For our result it will be irrelevant whether the endowments are private information or known to the designer. We call a choice rule ϕ *individually rational* if for all $i \in N$ and all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$,

$$\phi_i(\boldsymbol{\theta}) \succeq_i e(i).$$

Proposition SA.1. *Assume $|N| \geq 2$. Every partition P for an individually rational and efficient housing assignment choice rule ϕ produces contextual privacy violations.*

Proof. The proof shows an instance of collective but not individual pivotality. Consider two agents i and j and two possible preference profiles for each agent. For agent i, consider a type θ_i which contains $e(i) \succ_i e(j)$, and a type θ_i' which contains $e(j) \succ_i e(i)$. For agent j, consider θ_j which contains $e(j) \succ_j e(i)$, and a type θ_j' which contains $e(i) \succ_j e(j)$. Hold fixed all other types $\boldsymbol{\theta}_{-ij}$, to be such that they prefer their own endowment to all other objects, i.e. $\boldsymbol{\theta}_{-ij} = (e(k) \succeq_k c$ for all $c \in C \setminus \{e(k)\})_{k \in N \setminus \{i,j\}}$.

When the type profile is $(\theta_i, \theta_j, \boldsymbol{\theta}_{-ij})$, the agents both prefer their own endowment to the other's. When the profile is $(\theta_i', \theta_j, \boldsymbol{\theta}_{-ij})$, or $(\theta_i, \theta_j', \boldsymbol{\theta}_{-ij})$, one of the agents prefers their own endowment while the other agent prefers the other's. When $(\theta_i', \theta_j', \boldsymbol{\theta}_{-ij})$, they each prefer the other's endowment to their

own. Let x be the outcome in which both agents retain their endowment, i.e. $x = (i, e(i)), (j, e(j))$. Let y be the outcome in which each agent gets each other's endowment $y = (i, e(j)), (j, e(i))$. Then, individual rationality makes the requirements shown on the mid-left in Figure SA.1: $\phi(\theta_i, \theta_j, \boldsymbol{\theta}_{-ij}) = \phi(\theta_i, \theta'_j, \boldsymbol{\theta}_{-ij}) = \phi(\theta'_i, \theta_j, \boldsymbol{\theta}_{-ij}) = x$. Meanwhile, efficiency requires $\phi(\theta'_i, \theta'_j, \boldsymbol{\theta}_{-ij}) = y$ (shown on the mid-right in Figure SA.1). Hence, by Corollary 2, every partition for an individually rational and efficient choice rule produces contextual privacy violations. □

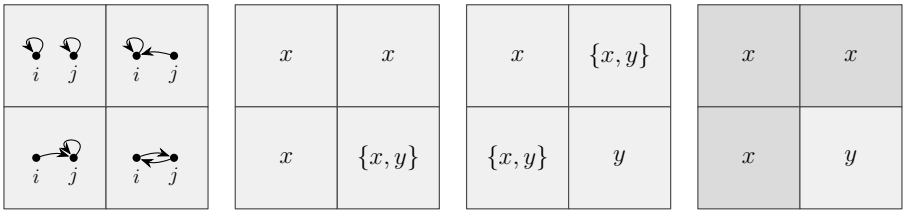

FIGURE SA.1. CONSTRUCTING COLLECTIVE BUT NOT INDIVIDUAL PIVOTALITY FOR HOUSE ASSIGNMENT.

Note: Type profiles $\{\theta_i, \theta'_i\} \times \{\theta_j, \theta'_j\}$ used in the proof (left, arrows denote whether agent prefers own or other's endowed object); required outcomes for each type profile under individual rationality (mid-left), under efficiency (mid-right), and under both efficiency and individual rationality (right).

The failure of contextual privacy in the house assignment problem is illuminating. There will be a violation at any type profile $\boldsymbol{\theta}$ in which there is a pair of agents who each prefer their endowment to the other's. Note that in a setting with many agents and many objects, there are many such type profiles $\boldsymbol{\theta}$ in $\boldsymbol{\Theta}$. In such cases, the conjunction of efficiency and individual rationality produce an instance where the agents are collectively but not individually pivotal.

SA.2. Voting

A natural case of collective but not individual pivotality arises in voting. If three agents vote on a binary issue, then a single agent cannot certainly change the outcome, but two agents together can. This issue holds more generally for generalized median voting rules. These are defined on an ordered outcome space $(X, \leq)$. The type space is the set of single-peaked preferences with respect to $\leq$. A preference $\preceq$ on $(X, \leq)$ is single-peaked if for each agent $i = 1, 2, \ldots, n$ and type θ_i, there is a *peak* $p(\theta_i)$ such that $x < x' \leq p(\theta_i) \implies x' \succ_i x$ and $x > x' \geq p(\theta_i) \implies x' \succ_i x$.

We consider a commonly studied class of voting rules. Namely, we study *generalized median voting rules*. A generalized median voting rule considers the peaks $p(\theta_1), p(\theta_2), \ldots p(\theta_n) \in X$ of agent types as well as *phantom ballots* $k_1, k_2, \ldots, k_{n-1} \in X \cup \{-\infty, \infty\}$, where for all $x \in X$, $-\infty < x < \infty$. The output

is the median of the submitted votes and phantom votes

$$\phi_{(k_1,k_2,\ldots,k_{n-1})}(\boldsymbol{\theta}) = \text{median}(p(\theta_1), p(\theta_2), \ldots, p(\theta_n), k_1, k_2, \ldots, k_{n-1}).$$

As shown in Moulin (1980), generalized median voting rules are exactly the anonymous, strategy-proof and Pareto-efficient voting rules on the class of single-peaked preferences. (A social choice rule ϕ is *anonymous* if $\phi(\boldsymbol{\theta}) = \phi((\boldsymbol{\theta}_{\pi(i)})_{i=1,2,\ldots,n})$ for all permutations $\pi\colon \{1,2,\ldots,n\} \to \{1,2,\ldots,n\}$.) Generalized median voting rules do not admit fully contextually private protocols.

Proposition SA.2. *Assume $|N| \geq 2$ and $|X| \geq \max\{2, {}^{|N|-1}/_{2}\}$ and the preference domain of single-peaked preferences. Any iterative partition P for a generalized median voting rule that is neither the maximum nor the minimum rule produces contextual privacy violations.*

In such voting rules, classical examples of collective but not individual pivotality arise from voting thresholds, e.g., a qualified majority. Consider a voting rule with at least two phantom ballots k_i on an alternative x that would win under a type profile $\boldsymbol{\theta}$. If there are two agents on the right of this alternative, and *both* of them having a type on the left of x, it would change the outcome—they are *collectively pivotal*—and yet it is still possible that neither of them alone would change the outcome.

Proof. Consider two types $\theta, \theta' \in \Theta$ with adjacent peaks and x and x', and a type profile $\boldsymbol{\theta}_{-ij} \in \boldsymbol{\Theta}_{-ij}$ such that

$$x = \text{median}(p(\boldsymbol{\theta}_{-ij}), k_1, k_2, \ldots, k_{n-1}).$$

We can find θ, θ', and $\boldsymbol{\theta}_{-ij}$ such that exactly one outcome in $p(\boldsymbol{\theta}_{-ij}), k_1, k_2, \ldots, k_{n-1}$ is x by $|X| \geq \max\{2, {}^{|N|-1}/_{2}\}$. The type profiles $(\theta, \theta, \boldsymbol{\theta}_{-ij})$, $(\theta', \theta, \boldsymbol{\theta}_{-ij})$ and $(\theta, \theta', \boldsymbol{\theta}_{-ij})$ will all result in outcome x. However, $(\theta', \theta', \boldsymbol{\theta}_{-ij})$ will result in x'. This hence produces an instance of collective pivotality without individual pivotality and shows that there must be a contextual privacy violation for at least one of the agents i, j. □

The Overdescending-Join Protocol

In Section IV, we defined the ascending-join protocol and showed that it is a maximally contextually private protocol for the k-item Vickrey auction. In this appendix, we consider a descending protocol $P_{\text{odesc-j}}$, and use a result from Section IV to show that it is maximally contextually private. In the *overdescending-join protocol*, for every type profile $\tilde{\boldsymbol{\Theta}}$, the designer asks agents $i = n, n-1, \ldots, 1$ whether they can rule out $\max \phi(\tilde{\boldsymbol{\Theta}})$, where, as in the main text, the outcomes are ordered first by price then in the strong set order on the set of winners. We choose the next outcome as the maximal outcome that is not ruled out by an agent.

For the k-item Vickrey auction rule $\phi_{V\text{-}k}$, $k = 1, 2, \ldots, n-1$ we obtain the following two results.

Proposition SB.1. *The overdescending-join protocol is maximally contextually private for the k-item Vickrey auction $\phi_{V\text{-}k}$.*

Proof. Let P be a protocol for $\phi_{V\text{-}k}$ such that $\Gamma(P, \phi_{V\text{-}k}) \subseteq \Gamma(P_{\text{odesc-j}}, \phi_{V\text{-}k})$. We construct the *reflected* social choice rule $\phi^r_{V\text{-}k} \colon \boldsymbol{\Theta}^r \mapsto (W^r, t^r)$, which, for any $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ has $i \in W^r(\boldsymbol{\theta})$ if and only if $i \notin W(\boldsymbol{\theta})$ and i is not the $(k{+}1)^{\text{st}}$ highest agent (with ties broken in the order $i = 1, 2, \ldots, n$), and $t^r(\boldsymbol{\theta}) = t(\boldsymbol{\theta})$. Define the type space $\Theta' = (\Theta, \leq')$ with the reflected order, where $\theta \leq' \theta' \iff \theta \geq \theta'$. Also define the reflected agent order $\boldsymbol{\theta}^r = (\theta_n, \theta_{n-1}, \ldots, \theta_1)$. Observe that $\phi^r_{V\text{-}k} \colon \boldsymbol{\Theta'} \to 2^N \times \boldsymbol{\Theta}$ is a $(n-k-1)$-item Vickrey auction,

$$\phi^r_{V\text{-}k}(\boldsymbol{\theta}^r) = \phi_{V\text{-}(n-k-1)}(\boldsymbol{\theta}). \tag{SB.1}$$

We also define the reflected protocols P^r and $P^r_{\text{odesc-j}}$, which, whenever P resp. $P_{\text{odesc-j}}$ queries agent i, it queries agent $n-i+1$, and agent action $\tilde{\Theta} \subseteq \Theta$ is translated to agent action $\{\theta^r | \theta \in \tilde{\Theta}\}$. Observe that

$$P^r_{\text{odesc-j}} = P_{\text{asc-j}} \text{ and } P^r_{\text{asc-j}} = P_{\text{odesc-j}} \tag{SB.2}$$

and that reflection preserves the contextual privacy order. i.e. for any two protocols P_1, P_2

$$\Gamma(P_1, \phi) \subseteq \Gamma(P_2, \phi) \implies \Gamma(P^r_1, \phi^r) \subseteq \Gamma(P^r_2, \phi^r). \tag{SB.3}$$

Hence, as we assume $\Gamma(P, \phi_{V\text{-}k}) \subseteq \Gamma(P_{\text{odesc-j}}, \phi_{V\text{-}k})$, we have by (SB.3) that

$$\Gamma(P^r, \phi^r_{V\text{-}k}) = \Gamma(P^r, \phi_{V\text{-}(n-k-1)}) \subseteq \Gamma(P_{\text{asc-j}}, \phi_{V\text{-}(n-k-1)}) = \Gamma(P^r_{\text{odesc-j}}, \phi^r_{V\text{-}k}).$$

Here, the last equality is by (SB.1). By Lemma 5, it must be that

$$\Gamma(P^r, \phi_{V\text{-}(n-k-1)}) = \Gamma(P_{\text{asc-j}}, \phi_{V\text{-}(n-k-1)}),$$

hence by (SB.3) $\Gamma(P^{rr}, \phi^r_{V\text{-}(n-k-1)}) = \Gamma(P^r_{\text{asc-j}}, \phi^r_{V\text{-}(n-k-1)})$ and hence, as $P^{rr} = P$, $\phi^r_{V\text{-}(n-k-1)} = \phi_{V\text{-}k}$ and $P^r_{\text{asc-j}} = P_{\text{odesc-j}}$

$$\Gamma(P, \phi_{V\text{-}k}) = \Gamma(P_{\text{odesc-j}}, \phi_{V\text{-}k}),$$

concluding the proof. □

We can express which agents' privacy is protected with the overdescending protocol. Define $d'(\boldsymbol{\theta})$ as the maximal index i such that the $(k+1)^{\text{st}}$ highest

value of $\theta_i, \theta_{i+1}, \ldots, \theta_n$ is the $(k+1)^{\text{st}}$ highest value of $\boldsymbol{\theta}$, and let

$$\begin{aligned} \mathrm{L}(\boldsymbol{\theta}) &= \{i \mid i \text{ is neither a winner nor determines the price at } \boldsymbol{\theta}\} \\ \mathrm{LQW}(\boldsymbol{\theta}) &= \mathrm{W}(\boldsymbol{\theta}) \cap \{i \mid i < d'(\boldsymbol{\theta})\}. \end{aligned}$$

Proposition SB.2. *Let $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. The overdescending-join protocol protects losing and late-queried winning bidders. That is, for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$,*

$$i \in \mathrm{L}(\boldsymbol{\theta}) \cup \mathrm{LQW}(\boldsymbol{\theta}) \Rightarrow (\boldsymbol{\theta}, i) \notin \Gamma(P_{\text{odesc-j}}, \phi_{V\text{-}k}).$$

A naive overdescending protocol which would ask all agents from the top until the $(k+1)^{\text{st}}$ highest bid is found can protect losers, but does not protect late-queried winners. Delay in querying protects these winners in addition to losers. For simplicity, we provide a direct proof of this result, not using the reflection approach in the proof of Proposition SB.1.

Proof. The protocol stops when the $(k+1)^{\text{st}}$ highest bid is found, hence only revealing the fact that losers have a type at most $\boldsymbol{\theta}_{[k+1]}$. As the queries are going from $i = n$ down to 1, the first time that late-queried winners are asked, this means that the type $\boldsymbol{\theta}_{[k+1]}$ has been found, and they need only confirm that their type is at least $\boldsymbol{\theta}_{[k+1]}$. This shows the claim. □

Relative Informativeness

Our analysis has emphasized contextual privacy, a notion explicitly linked to the social choice rule ϕ and sensitive to the privacy concerns of individual agents. Other privacy criteria offer insights about information revelation in mechanisms that are largely complementary to those revealed through a contextual privacy analysis. In particular, as noted in Section II.C, the notion of *relative informativeness* (Mackenzie and Zhou, 2022; Segal, 2007) can allow a designer to further narrow down their selection of a mechanism within a set of mechanisms that are equivalent in terms of contextual privacy.

Definition. An iterative partition P is *less informative* than an iterative partition P', denoted $P \preceq_{\text{RI}} P'$, if any type profiles $\boldsymbol{\theta}$, $\boldsymbol{\theta}'$ that P distinguishes are also distinguished by P'.

We note that every distinction made by the relative informativeness partial order is also made by the contextual privacy partial order.

Proposition SC.1. *Let P, P' be iterative partitions for ϕ. Then*

(i) for all P, P' for ϕ,

$$P \preceq_{\text{RI}} P' \implies \Gamma(P, \phi) \subseteq \Gamma(P', \phi).$$

(ii) there exist P, P' for ϕ such that $P \not\preceq_{\text{RI}} P', P' \not\preceq_{\text{RI}} P$ but $\Gamma(P,\phi) \subset \Gamma(P',\phi)$.

Proof. We begin with (i) and prove the contrapositive. Assume that $P \preceq_{\text{RI}} P'$ but $\Gamma(P,\phi) \not\subseteq \Gamma(P',\phi)$. Then, there is $(\boldsymbol{\theta}, i) \in \Gamma(P,\phi) \setminus \Gamma(P',\phi)$. That means that there is a $\boldsymbol{\theta'} = (\theta_i', \boldsymbol{\theta}_{-i})$ that are distinguished by P but not P' (these satisfy $\phi(\boldsymbol{\theta}) = \phi(\boldsymbol{\theta'})$, but this is insubstantial for this claim). This means that $P \not\preceq_{\text{RI}} P'$.

For claim (ii), let $n = 2$, $\Theta_1 = \{0,1\}$, $\Theta_2 = \{0,1,2\}$ and let the choice rule be the simple indicator function $\phi(\theta_1,\theta_2) = \mathbb{1}_{\theta_1=1}$, i.e. the outcome is 1 if and only if $\theta_1 = 1$, otherwise it is zero.

We consider two iterative partitions. P first asks agent 1 "Is $\theta_1 = 1$?" On an affirmative answer, ask agent 2 "Is $\theta_2 = 0$?", otherwise stop. This leads to the final partition

$$P = \Big\{\underbrace{\{(0,0),(0,1),(0,2)\}}_{\text{outcome 0}},\ \underbrace{\{(1,0)\}}_{\text{outcome 1}},\ \underbrace{\{(1,1),(1,2)\}}_{\text{outcome 1}}\Big\}.$$

A second partition P' again asks agent 1 "Is $\theta_1 = 1$?" On an affirmative answer, ask agent 2 "Is $\theta_2 = 2$?" On a negative answer, ask agent 2 "Is $\theta_2 = 0$?" The final partition for P' is

$$P' = \Big\{\underbrace{\{(0,0)\}}_{\text{outcome 0}},\underbrace{\{(0,1),(0,2)\}}_{\text{outcome 0}},\ \underbrace{\{(1,0),(1,1)\}}_{\text{outcome 1}},\ \underbrace{\{(1,2)\}}_{\text{outcome 1}}\Big\}.$$

P distinguishes $(1,0)$ from $(1,1)$ whereas P' does not; conversely, P' distinguishes $(0,0)$ from $(0,1)$ while P does not. Thus $P \not\preceq_{\text{RI}} P'$ and $P' \not\preceq_{\text{RI}} P$.

Because ϕ ignores agent 2's type, *any* difference in θ_2 is economically irrelevant. In P the designer learns θ_2 *only* when $\theta_1 = 1$, giving

$$\Gamma(P,\phi) = \{((1,0),2),\ ((1,1),2),\ ((1,2),2)\}.$$

In P' she obtains information about θ_2 for both $\theta_1 = 1$ and $\theta_1 = 0$, hence

$$\Gamma(P',\phi) = \Gamma(P,\phi) \cup \{((0,0),2),\ ((0,1),2),\ ((0,2),2)\}.$$

Hence $\Gamma(P,\phi) \subset \Gamma(P',\phi)$ and therefore $P \preceq_{CP} P'$. We have found two protocols that are *incomparable* in the relative-informativeness but are strictly ranked by contextual privacy, completing the proof of part (ii). □

Relative informativeness can be used to refine and select elements that are contextual privacy equivalent.[1]

Example. Equivalence under $\preceq_{\text{CP}}$ but comparability under $\preceq_{\text{RI}}$. Consider $n = 1$, $\Theta = \{1,2,\ldots,k\}$ and $\phi(\theta) = \mathbb{1}_{\theta=k}$. In this case, there is a unique maximally contextually private protocol $P_{\theta=k}$ which asks the agent "is your value equal to

[1] We thank an anonymous reviewer for offering a version of this example.

k?" Any other two protocols P, P' for ϕ that are not $P_{\theta=k}$ are contextual privacy equivalent to each other. (Formally, all protocols except for the one with the final partition $\{\{k\}, \{1, 2, \ldots, k-1\}\}$ are contextual privacy equivalent.)

However, it is clear that some protocols yield less information overall, and thus, for a privacy-conscious designer, they may be more desirable. For example, consider the protocol $P_{\theta=1,\theta=k}$ which asks the agent "is your value equal to 1?" and then "is your value equal to k?" This yields a partition $\{\{1\}, \{k\}, \{2, \ldots, k-1\}\}$. And then consider the protocol P_{reveal} which asks the agent to reveal her type exactly. This yields a partition $\{\{1\}, \{2\}, \{3\}, \ldots, \{k\}\}$. $P_{\theta=1,\theta=k}$ and P_{reveal} are contextual privacy equivalent. However, there is clearly a sense in which P_{reveal} reveals more information—indeed, this is captured by relative informativeness, $P_{\theta=1,\theta=k} \prec_{\text{RI}} P_{\text{reveal}}$.

To summarize the discussion thus far, relative informativeness and contextual privacy are complementary. Among a set of protocols that are *equivalent* in contextual privacy order, relative informativeness *can sometimes make further comparisons.* Among a set of protocols that are *incomparable* in the relative informativeness order, contextual privacy *can sometimes compare them.*

This allows us to evaluate the relative informativeness of protocols considered in this paper. Consider first the protocol for the first-price auction rule that asks agents in decreasing order the value of their bid function $b_i(\theta)$ whether their type is above some threshold type $\tilde{\theta}$. Because of lexicographic tiebreaking, ask in the order $i = 1, 2, \ldots, n$. We call this protocol P_{desc}.

Proposition SC.2. *P_{desc} is minimally relatively informative among protocols computing ϕ_{FPA}.*

Proof. The final partition of P_{desc} consists of sets with the same values:

$$\left(\min \arg\max_{i=1,2\ldots,n} b_i(\theta_i), \max_{i=1,2,\ldots,n} b_i(\theta_i) \right).$$

Hence, in particular, it distinguishes only type profiles that lead to different outcomes. Any coarser partition of $\boldsymbol{\Theta}$ (not only final partitions of protocols) can therefore not compute ϕ_{FPA}. □

The result for the first-price auction rule shows a strong kind of optimality from a privacy perspective. The protocol only distinguishes type profiles that have different outcomes—it is fully contextually private and also a minimal element in the relative informativeness order.

We can show a similar result for the ascending-join protocol. To show that the ascending-join protocol $P_{\text{asc-j}}$ is minimally relatively informative, we first observe that it is without loss to restrict to *binary* protocols, where each internal node $\tilde{\boldsymbol{\Theta}}$ has two children.

Consider an arbitrary protocol P. Let L be the total number of children of internal nodes in excess of 2. In a topological order, consider nodes $\tilde{\boldsymbol{\Theta}}$, and

repeatedly replace its children $\tilde{\boldsymbol{\Theta}}_1, \tilde{\boldsymbol{\Theta}}_2, \ldots, \tilde{\boldsymbol{\Theta}}_l$ by one child $\tilde{\boldsymbol{\Theta}}_1$ and another child $\boldsymbol{\Theta}_2 \cup \cdots \cup \tilde{\boldsymbol{\Theta}}_l$, which is partitioned as $\boldsymbol{\Theta}_2, \ldots, \boldsymbol{\Theta}_l$, with no change to the rest of the iterative partition. This operation reduces L by 1, hence the process terminates. It leads to an iterative partition P' with the same final partition, and hence $P \sim_{\text{RI}} P'$.

To state our sufficient condition, we need two definitions. We say that a binary protocol P is *flat* if there does not exist a node $\tilde{\boldsymbol{\Theta}}$, agent i, types $\theta_i, \theta_i' \in \tilde{\Theta}_i$ such that for all partial type profiles $\boldsymbol{\theta}_{-i} \in \tilde{\boldsymbol{\Theta}}_{-i}$, $(\theta_i, \boldsymbol{\theta}_{-i})$ and $(\theta_i', \boldsymbol{\theta}_{-i})$ are distinguished by P, but not at $\tilde{\boldsymbol{\Theta}}$. In flat protocols, all distinctions are made as soon as they can be made, and not further down in the iterative partition (which, recall, is a rooted tree). We say a protocol is *nonredundant* if for all internal nodes $\tilde{\boldsymbol{\Theta}}$, $\phi(\tilde{\boldsymbol{\Theta}})$ is non-constant.

Lemma SC.1. *Let P be a flat and nonredundant protocol for ϕ. Then P is minimally relatively informative for ϕ.*

Proof. For contradiction assume that there was a less relatively informative protocol $P' \preceq_{\text{RI}} P$. We will show that the protocols lead to identical final partitions. If not, there must be an earliest node $\tilde{\boldsymbol{\Theta}}$ (with respect to the precedence order of P' and P) such that the children of $\tilde{\boldsymbol{\Theta}}$ are not identical in P and P'.

There are two cases, which are both ruled out by our assumption of nonredundancy and flatness.

Terminal node for P' First, P' might terminate at $\tilde{\boldsymbol{\Theta}}$. In this case, it must be that P is redundant, contradicting the assumption of nonredundancy.

Non-Terminal node for P' Second, P' might ask a different query than P. As P' is weakly less relatively informative, it must be the case that all type profiles that are distinguished at $\tilde{\boldsymbol{\Theta}}$ in P' must be distinguished for all children of $\tilde{\boldsymbol{\Theta}}$ in P, but not at $\tilde{\boldsymbol{\Theta}}$. This is impossible by the assumed flatness of P.

Therefore, such a node cannot exist and $P' = P$. □

Proposition SC.3. *$P_{\text{asc-j}}$ is minimally relatively informative among protocols computing $\phi_{V\text{-}k}$.*

Proof. We first show that $P_{\text{asc-j}}$ is flat. For a bimonotonic protocol like $P_{\text{asc-j}}$ it is sufficient to show that there is no node $\tilde{\boldsymbol{\Theta}}$ such that there is a threshold query, say (j, t') (to agent j at threshold t'), that is asked for all $\boldsymbol{\theta}_{-i} \in \tilde{\boldsymbol{\Theta}}_{-i}$. As $P_{\text{asc-j}}$ is bimonotonic, we can identify the node $\tilde{\boldsymbol{\Theta}}$ of $P_{\text{asc-j}}$ with a threshold query (i, t). There are two cases for queries (j, t') and (i, t), based on the fact that (i, t) comes before (j, t').

The first one is where (j, t') is a threshold query for a higher threshold, (j, t'), $j = 1, 2, \ldots, n$, $t' > t$. These are not queried if the protocol terminates at running

price t, which is possible if all but k agents $i = 1, 2, \ldots, n$ have type $\theta_i \leq t$. Hence, flatness holds for such nodes $\tilde{\mathbf{\Theta}}$ and future queries (j, t').

A second case is a query to another agent at the same running price, (j, t), $j \in N \setminus \{0, i\}$. We distinguish two cases. If $j < i$, then by (A2) if all agents $j' \geq i$ drop out at t, $P_{\text{asc-j}}$ does not query (t, j), but terminates. If $j > i$, then if all agents $j' \leq i$ stay in at cutoff t, then (j, t) is not queried, but the price is raised "before it's j's turn." Hence, we also establish flatness in this case.

Second, we show that $P_{\text{asc-j}}$ is nonredundant. This is readily seen from the definition of $P_{\text{asc-j}}$: since the ascending join protocol must rule out an outcome for every query, it cannot ask redundant queries.

Since $P_{\text{asc-j}}$ is both flat and nonredundant, it is minimal in the relative informativeness order. □

Using symmetry like in the proof of Proposition SB.1 it can be shown that the overdescending-join protocol presented in Appendix SB is minimally relatively informative.

*

REFERENCES

Mackenzie, Andrew and Yu Zhou, "Menu mechanisms," *Journal of Economic Theory*, 2022, *204*, 105511.

Moulin, Hervé, "On strategy-proofness and single peakedness," *Public Choice*, 1980, *35* (4), 437–455.

Segal, Ilya, "The communication requirements of social choice rules and supporting budget sets," *Journal of Economic Theory*, 2007, *136* (1), 341–378.

Shapley, Lloyd and Herbert Scarf, "On cores and indivisibility," *Journal of Mathematical Economics*, 1974, *1* (1), 23–37.