

PoA of Simple Auctions with Interdependent Values *

Alon Eden ¹, Michal Feldman ², Inbal Talgam-Cohen ³, Ori Zviran ²

¹ Harvard University

² Tel Aviv University

³ Technion, Israel Institute of Technology

alonarden@gmail.com, michal.feldman@cs.tau.ac.il, italgam@cs.technion.ac.il, ori.zviran@gmail.com

Abstract

We expand the literature on the price of anarchy (PoA) of simultaneous item auctions by considering settings with correlated values; we do this via the fundamental economic model of interdependent values (IDV). It is well-known that in multi-item settings with private values, correlated values can lead to bad PoA, which can be polynomially large in the number of agents n . In the more general model of IDV, we show that the PoA can be polynomially large even in single-item settings. On the positive side, we identify a natural condition on information dispersion in the market, which enables good PoA guarantees. Under this condition, we show that for single-item settings, the PoA of standard mechanisms degrades gracefully. For settings with multiple items we show a separation between two domains: If there are more buyers, we devise a new simultaneous item auction with good PoA, under limited information asymmetry. To the best of our knowledge, this is the first positive PoA result for correlated values in multi-item settings. The main technical difficulty in establishing this result is that the standard tool for establishing PoA results — the smoothness framework — is unsuitable for IDV settings, and so we must introduce new techniques to address the unique challenges imposed by such settings. In the domain of more items, we establish impossibility results even for surprisingly simple scenarios.

1 Introduction

In this work, we study simple and practical mechanisms for selling heterogeneous items to unit-demand agents,¹ with the goal of maximizing social welfare. While much of the literature focuses on truthful mechanisms (for which it is in the agents' best interest to report their true values for the items), such mechanisms are often complicated, computationally and cognitively demanding, and must be run in a centralized manner (Dobzinski 2011; Li 2017; Ausubel, Milgrom et al. 2006). In many real-life settings like e-commerce, the mechanisms used in practice are simple and run in a distributed

manner, but are non-truthful. A prime example is the auctions run by e-Bay, where each item is sold separately in a second-price auction, so that an agent interested in winning at most one auction may be better off reporting lower than her true values, to avoid multiple wins.

Motivated by such settings, Christodoulou et al. [2016] pioneered the study of *simultaneous item auctions*, where a single-item auction is run for each item separately. Since such auctions are non-truthful, their performance is measured by their *Bayesian Price of Anarchy* (B-POA) — the ratio between the optimal welfare and the welfare guarantee of the worst equilibrium. Christodoulou et al. and many follow-ups show that this simple format achieves near-optimal welfare in combinatorial settings when the auction in use is the first- or second-price auction [Roughgarden et al. 2017]. Importantly, all prior works assume independence of different buyers' valuations, since otherwise the PoA might be polynomial in the number of agents (Bhawalkar and Roughgarden 2011; Feldman et al. 2013; Roughgarden 2016).

There are many settings in which the independence assumption is unrealistic. For instance, if an item being auctioned has the potential of being resold in the future, this potential factors into agents' values for the item, creating dependence on how others value it. Such dependence is formally captured by the *interdependent values* (IDV) model (Milgrom and Weber 1982). In this model, the correlation of values stems directly from one of the most fundamental aspects of a market — the way in which information is dispersed among agents. In more detail, each agent has a privately-known *signal*, which captures her partial knowledge about the items for sale. Her value for each item is a function of all the information on the market related to this item; that is, of her own signal as well as the signals of all other agents (which are unknown to her). Since the valuations of different agents depend on the same signals, values are correlated.

IDV settings have been studied extensively in the economic literature since Milgrom and Weber (1982), and in the computer science literature (Ito and Parkes 2006; Constantin, Ito, and Parkes 2007; Robu et al. 2013; Chawla, Fu, and Karlin 2014; Roughgarden and Talgam-Cohen 2016; Eden et al. 2019); see Krishna (2009) for an overview and the full version for additional related work. Realistic scenarios captured by this model include common value auc-

*For the full version, see <https://arxiv.org/abs/2011.00498> (Eden et al. 2020). Due to space limitation, all missing statements and proofs are deferred to the full version.
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Where an agent has a value for each item, and her value for a bundle is the maximal value for a bundle's item.

tions (Klemperer 1998), mineral rights auctions (Wilson 1969) and resale (Myerson 1981; Roughgarden and Talmag-Cohen 2016). We use the last as a running example throughout the paper:

Example 1.1 (Resale model). *Let s_i be the private signal of agent $i \in [n]$, distributed uniformly between 0 and 1. This signal captures the appreciation of agent i for the item for sale. Agent i 's valuation function is $v_i(\mathbf{s}) = s_i + \beta \sum_{j \neq i} s_j$ for some parameter $\beta \in (0, 1)$. That is, the agent also takes into account the others' assessment of the item (possibly since she plans to resell the item eventually).*

Main Question. *Our main question in this paper is: What is the B-PoA of simultaneous item auctions with IDV?*

The bounds we establish partially answer an open question of Roughgarden et al. [2017] regarding natural, economically-meaningful forms of correlation for which the B-PoA of simultaneous item auctions is bounded.

Challenges. Several challenges arise when approaching the task of bounding the B-PoA in IDV. As intuition suggests, the domain of problems that can be modeled by interdependent values is very large, and includes the mostly-studied setting of private values, correlated or uncorrelated. Due to the wide scope, it is unsurprising that in full generality there is no hope to achieve a good bound for the B-PoA. Our lower bounds formalize the need for assumptions in order to get a bound independent of the number of agents.

A standard assumption in the interdependent model is *single-crossing* (SC). Intuitively it means that the information possessed by an agent has more influence on her own value than on others' values. SC enables many positive results in settings with interdependence (Ausubel et al. 2000; Maskin 1992; Chawla, Fu, and Karlin 2014). For example, in a single item setting, the truthful (generalized) Vickrey auction achieves optimal welfare only with SC. However, SC alone is insufficient for ensuring good B-PoA results of simple auctions, even with a single item. The following proposition shows that any mechanism that allocates the item to the agent with the highest value is prone to a large degree of social inefficiency even with SC.

Proposition 1.2. *There exists a single-item, n -bidder setting satisfying SC such that the PoA of every auction that allocates the item to the highest-valued bidder is $\Omega(n)$.*²

The intuition behind this negative result is instructive for identifying necessary assumptions for positive results: Consider n bidders whose signals belong to subsets $S_1 = \{1\}$, $S_2 = \{1\}$, and $S_i = [0, 1]$ for $i \geq 3$, respectively. Consider the valuation profile $v_1 = \sum_{i \in [n]} s_i + \epsilon$, $v_2 = 2(s_2 + \epsilon)$, and $v_i = s_i$ for $i \geq 3$, where $\epsilon > 0$ is arbitrarily small. In this setting, the signals of bidders $i \geq 3$ have a significant effect on bidder 1's value but have no effect on bidder 2's value. In a scenario where these bidders have high signals, bidder 1's value is significantly higher than bidder 2's value. However, bidders $i \geq 3$ never win the item, since

²This holds even under the standard no-overbidding assumption (see Section 2) and under additional conditions such as submodularity over signals (Eden et al. 2019).

bidder 1 out-values them for every signal profile. Thus, they may as well report low signals in equilibrium, resulting in an outcome where bidder 2 wins the item. A similar result to Proposition 1.2 holds for every deterministic truthful (ex-post incentive compatible) mechanism. For further details see the full version.

Our Results. We study simple mechanisms in which the agents report their signals to the auction(s). All our positive results hold with respect to Bayesian equilibrium, which is the strongest guarantee and propagates to other equilibrium notions. Moreover, we show that pure equilibria need not exist even in a single-item setting, which further motivates our choice of benchmark.³ Since we are tackling new uncharted territory, our starting point is a single item, where we already face new challenges that do not exist in private values settings.

Single item. We consider natural auction formats like first-price and second-price. As discussed above, the PoA of such auctions inherently depends on the number of agents. In the proofs of Propositions 1.2, a single bidder is highly influenced by the signals of others, while all other bidders have private values. Thus, the effect of bidder signals on different bidders may vary greatly. We parameterize the extent of this variation by γ (Definition 3.2) and establish positive results that depend on γ .

Our main positive result for single item applies to generalized Vickrey auction (GVA) — a natural generalization of the Vickrey auction to interdependent values (Maskin 1992; Ausubel et al. 2000). We show that the B-POA of GVA under SC is γ (see Theorem 3.3), and this is tight. When considering a relaxed notion of SC, c -SC, we get a bound of $\max\{\gamma, c\} + 1$ for both GVA and second-price auctions, which is almost tight. A non-trivial implication of this result is that in every IDV setting where $\gamma = c = 1$, it holds that every equilibrium is fully efficient. For example, this is the case in our running example of resale model (Example 1.1).

Multiple items. We consider a combinatorial setting with heterogeneous items and unit-demand valuations. We allow for signals to be multi-dimensional, a notoriously hard setting in the IDV literature (Dasgupta and Maskin 2000; Jehiel and Moldovanu 2001). Our results for this case are more nuanced. First of all, we show that in general, one cannot approximate the optimal welfare for a very natural class of mechanisms which includes the ones we reason about in this paper. In the example showing this, the value of all bidders is originated from a single bidder's signal; that is, this setting suffers from extreme information asymmetry. Therefore, we provide results for settings with limited-information asymmetry (for a formal definition of the condition, see Section 4.2), like our running example.

Moreover, we show a separation between two domains: the case where there are more bidders than items ($n \geq m$), and the case where there are many more items than bidders. For the case $n \geq m$, we consider a mechanism that runs a second-price auction for each item, when considering bid-

³In contrast, for simultaneous auctions with private values, it is known that a PNE exists even for general classes of valuations (e.g. (Christodoulou, Kovács, and Schapira 2016)).

der’s valuation valuated at her bid, zeroing out other bidders’ bids. We refer to this mechanism as “simultaneous privatized second-price auction”. One subtlety that arises in our setting is the need to enable agents to explicitly express willingness to participate in auctions for individual items. Our main positive result for multiple items is that for γ -heterogeneous unit demand bidders, and under sufficiently limited information asymmetry, the B-POA of simultaneous privatized second-price auctions is $O(\gamma^2)$ (under c -SC, we get a bound of $O(\max\{\gamma, c\}^2)$, see Theorem 4.7 for details).

Our proof diverges from the smoothness framework, as the typical proof requires reasoning about a deviation of bidders who know their value for an item, which is not the case with interdependence. Our proof decomposes the welfare into two terms and bounds the performance of our auction for each one separately (for more details, see Section 4).

For the case $m \gg n$, we provide a strong negative result. We consider the simplest setting one can imagine: there are n^2 items, and bidders have common values for each item, which is the sum of signals of bidders for the item ($v_{i\ell}(\mathbf{s}) = 1 + \sum_j s_{j\ell}$ for every item ℓ). All signals for any bidder/item are sampled i.i.d. This example demonstrates the difficulty of coordination in our setting. For every item, the bidders’ values are equal, and thus the bidder’s difficulty is in identifying the valuable items. The fact that the information is distributed among bidders poses a major difficulty. A priori, all items are the same, but the signals’ distribution is set up in a way that gives rise to a ball-and-bins type phenomenon, where the expected value of each item is a constant, but the maximal value of n items is $\Theta(\ln n / \ln \ln n)$. It follows that $\text{OPT} = \Theta(n \ln n / \ln \ln n)$.

We then use the standard no-overbidding assumption (which is crucial for our positive results in the $n \leq m$ regime) to bound the number of items a bidder bids on, which in turn upper bounds the expected welfare of the bidder. As a result, we show that for the settings discussed, for every simultaneous item auction and every equilibrium σ , $\text{OPT}/\text{EQ}(\sigma) = \tilde{\Omega}(\log n)^4$ (see the full version for the full details). Note that this implies that the Price of Stability is $\tilde{\Omega}(\log n)$, which is a stronger negative result.

2 Preliminaries

Notation. Let \mathbf{x}, \mathbf{y} be vectors, i an index and A a subset of indices. Then \mathbf{x}_A is the vector obtained by taking indices A of \mathbf{x} . Let $\mathbf{x}_{-A} = \mathbf{x}_{\bar{A}}$ (in particular \mathbf{x}_{-i} is vector \mathbf{x} with the i th entry removed). By $\mathbf{x} \leq \mathbf{y}$ we mean $x_i \leq y_i$ for all i . Vectors $\mathbf{1}$ and $\mathbf{0}$ are the all-ones and all-zeros vectors.

Basic problem setting. The following is a standard interdependence setting for selling a single item (Milgrom and Weber 1982): There are n bidders (agents), each bidder i with a privately-known signal s_i from signal space S_i (a continuous interval in $\mathbb{R}_{\geq 0}$). Let $\mathbf{s} = (s_1, \dots, s_n) \in S \subseteq S_1 \times \dots \times S_n$ denote the signal profile of all bidders. Every bidder i also has a publicly-known valuation $v_i : S \rightarrow \mathbb{R}_{\geq 0}$, which is a function of the signal profile \mathbf{s} . This dependence of the values on other bidders’ signals is the defin-

ing property of interdependence. Function v_i is weakly increasing in each coordinate and strictly increasing in s_i . Let $\mathbf{v} = (v_1, \dots, v_n)$ denote the valuation profile of all bidders.

We also consider combinatorial settings in which there are m items for sale: For multiple items, \mathbf{s}_i is multi-dimensional and there is a signal $s_{i\ell}$ for every item ℓ . Value $v_{i\ell}$ of the ℓ th item is a function of $\mathbf{s}_\ell = (s_{1\ell}, \dots, s_{n\ell})$. We focus on *unit-demand* valuations, for which the value for a subset of items $T \subseteq [m]$ is $v_i(T; \mathbf{s}) = \max_{\ell \in T} \{v_{i\ell}(\mathbf{s}_\ell)\}$.

As part of the classic interdependence model, valuation profiles are assumed to be *single-crossing (SC)* (e.g. Milgrom and Weber 1982; Maskin 1992; Ausubel et al. 2000).

Definition 2.1 (SC). *Valuation profile \mathbf{v} is SC if for every bidder pair i, i' , every signal profile \mathbf{s} and every $\delta \geq 0$, $v_i(s_i + \delta, \mathbf{s}_{-i}) - v_i(\mathbf{s}) \geq v_{i'}(s_i + \delta, \mathbf{s}_{-i}) - v_{i'}(\mathbf{s})$.*

Intuitively SC means that a bidder’s own value is influenced the most from a change in her signal. Eden et al. (2018) introduced a natural relaxation called c -SC for $c \geq 1$, which requires $c(v_i(s_i + \delta, \mathbf{s}_{-i}) - v_i(\mathbf{s})) \geq v_{i'}(s_i + \delta, \mathbf{s}_{-i}) - v_{i'}(\mathbf{s})$. Our results generalize to this relaxation.

Who knows what. We address both full-information and Bayesian settings. In either model the signal spaces S_1, \dots, S_n and the valuation functions v_1, \dots, v_n are assumed to be public knowledge. In full information settings, all bidders know all signals (but the mechanism does not). In Bayesian settings (our main focus), bidder i has private knowledge of s_i , which is drawn from a publicly-known distribution F_i with density f_i . We denote by F the publicly-known joint distribution of signal profile \mathbf{s} .

Correlation. Even if the signals are sampled independently, the values are correlated since they depend on the same signals. This generalizes various types of correlation found in other papers (Bateni et al. 2015; Chawla, Malec, and Sivan 2015; Immorlica, Singla, and Waggoner 2020). When dealing with a single item we further allow the signals themselves to be correlated, i.e., F is not necessarily a product distribution. We denote by $F_{|s_i}$ the distribution of \mathbf{s}_{-i} given signal s_i .

Submodularity. Submodularity of the valuation functions is assumed in Section 4; it means that as a bidder’s own signal increases, an increase in the others’ signals has diminishing marginal influence on her value:⁵

Definition 2.2 ((Chawla, Fu, and Karlin 2014)). *A valuation v_i is submodular if for every signal s_i , every two profiles of the other bidders $\mathbf{s}_{-i} \leq \mathbf{s}'_{-i}$ and every $\delta \geq 0$, $v_i(s_i + \delta, \mathbf{s}_{-i}) - v_i(s_i, \mathbf{s}_{-i}) \geq v_i(s_i + \delta, \mathbf{s}'_{-i}) - v_i(s_i, \mathbf{s}'_{-i})$.*

For multiple items, the same definition applies to $v_{i\ell}$ (rather than v_i) with s_i, \mathbf{s}_{-i} replaced by $s_{i\ell}, (\mathbf{s}_\ell)_{-i}$.

2.1 Mechanisms

Objective. Our goal is to maximize welfare, that is, for a single item to allocate it to the bidder with the highest value $v_i(\mathbf{s})$, and for multiple items and unit-demand bidders to find a matching of items to bidders with the maximum total value. We denote the optimal welfare for a given setting

⁴ $\tilde{\Omega}$ hides $o(\log n)$ terms.

⁵(Eden et al. 2019) define a stronger version of valuation submodularity called submodularity over signals (SoS).

with signal profile \mathbf{s} by $\text{OPT}(\mathbf{s})$, and the expected optimal welfare by $\text{OPT} = \mathbb{E}_{\mathbf{s} \sim F}[\text{OPT}(\mathbf{s})]$.

Simple mechanisms. A mechanism consists of a pair (x, p) of (deterministic) allocation rule x and payment rule p . The mechanism solicits a signal report (*bid*) b_i from each bidder i . Let $\mathbf{b} = (b_1, \dots, b_n)$ denote the bid profile. The mechanism outputs for every bidder i an indicator $x_i(\mathbf{b}) \in \{0, 1\}$ of whether she wins the item, and her payment $p_i(\mathbf{b})$. The allocation rule x guarantees feasibility, i.e., $\sum_i x_i(\mathbf{b}) \leq 1$ for every \mathbf{b} . Bidder i 's expected utility given bid profile \mathbf{b} and true signal profile \mathbf{s} is $u_i(\mathbf{b}, \mathbf{s}) = x_i(\mathbf{b})v_i(\mathbf{s}) - p_i(\mathbf{b})$.

No-overbidding (NOB). As standard in the literature on PoA (e.g. Bhawalkar and Roughgarden 2011) we often assume NOB, defined as follows:

Definition 2.3 (Single-item NOB for interdependence). *Bid profile \mathbf{b} satisfies NOB if $b_i \leq s_i$ for every bidder i . Strategy profile σ satisfies NOB if for all $\mathbf{b} \sim \sigma(\mathbf{s})$, \mathbf{b} satisfies NOB.*

NOB assumptions reflect bidders' reluctance to expose themselves to negative utility from overbidding and help explain the prevalence of certain auction formats like the second-price auction in practice. It is well-known that without NOB the second-price auction has unbounded PoA.⁶

Standard mechanisms. In mechanisms for IDV, the *highest-valued bidder* is computed according to solicited bids \mathbf{b} and the public-knowledge valuation profile \mathbf{v} .

Definition 2.4 (Critical bid). *Given a partial bid profile \mathbf{b}_{-i} , bidder i 's critical bid b_i^* is the lowest report for which she wins the item, i.e., $b_i^* = \min\{b_i \mid x_i(\mathbf{b}) = 1\}$.*

Definition 2.5. *The generalized Vickrey auction (GVA) solicits signal bids, allocates the item to the highest-valued bidder i , and charges her critical bid value $v_i(b_i^*, \mathbf{b}_{-i})$.*

There are also natural generalization of the first- and second-price auctions to interdependence (throughout we refer to these and not the independent private values versions). The following is the generalization of second-price; the generalization of first-price is defined in the full version.

Definition 2.6 (Second-price (2PA) with interdependence). *The second-price auction solicits signal bids, allocates the item to the highest-valued bidder, and charges the second-highest value as payment.*

2.2 Price of Anarchy Background

Let $\sigma_i(\cdot)$ be bidder i 's bidding *strategy* as a function of her true signal. A bidding strategy can be *pure* (mapping to a single reported signal) or *mixed* (mapping to a distribution over reported signals).

Fix a mechanism. Given a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$, let $\sigma(\mathbf{s})$ be a mapping of signals \mathbf{s} to bids, and $\sigma_{-i}(\mathbf{s}_{-i})$ be the mapping excluding agent i 's bid. We focus on the following equilibrium concepts:

⁶Consider two bidders with (independent private) values $\epsilon, 1$ where $\epsilon \ll 1$. Bidding 1, 0 is an equilibrium of the second-price auction and its welfare is ϵ , but the optimal welfare is 1. In this equilibrium, the first bidder is bidding much higher than her value.

1. **Ex-post equilibrium (EPE):** A bidding strategy profile σ constitutes an EPE if σ is deterministic,⁷ and for every bidder i with signal s_i , and every signal profile \mathbf{s}_{-i} , $u_i(\sigma(\mathbf{s}); \mathbf{s}) \geq u_i((b_i, \sigma(\mathbf{s}_{-i})); \mathbf{s})$ for every b_i .

2. **Bayes-Nash equilibrium (BNE):** A (possibly randomized) bidding strategy profile σ constitutes a BNE of a Bayesian setting with signal distribution F if for every bidder i and signal s_i , and every b_i , $\mathbb{E}_{\mathbf{s}_{-i} \sim F|s_i} [u_i(\sigma(\mathbf{s}); \mathbf{s})] \geq \mathbb{E}_{\mathbf{s}_{-i} \sim F|s_i} [u_i((b_i, \sigma(\mathbf{s}_{-i})); \mathbf{s})]$

Price of Anarchy (PoA). The expected welfare achieved by a single-item mechanism (x, p) for signal profile \mathbf{s} and bid profile \mathbf{b} is $\sum_i x_i(\mathbf{b})v_i(\mathbf{s})$. Where the mechanism (in particular, the allocation rule) is evident from the context, we denote this by $\text{SW}(\mathbf{b}, \mathbf{s})$. For an equilibrium strategy σ , let $\text{EQ}(\sigma, \mathbf{s}) = \mathbb{E}[\text{SW}(\sigma(\mathbf{s}), \mathbf{s})]$ be the expected welfare of σ at \mathbf{s} , and $\text{EQ}(\sigma) = \mathbb{E}_{\mathbf{s} \sim F}[\text{EQ}(\sigma, \mathbf{s})]$. The ex-post PoA and Bayesian PoA of an auction, respectively, are:

$$\text{EP-POA} = \sup_{\substack{\mathbf{s}, \mathbf{v}, \sigma: \\ \sigma \text{ is EPE}}} \frac{\text{OPT}(\mathbf{s})}{\text{EQ}(\sigma, \mathbf{s})}; \quad \text{B-POA} = \sup_{\substack{F, \mathbf{v}, \sigma: \\ \sigma \text{ is BNE}}} \frac{\text{OPT}}{\text{EQ}(\sigma)}.$$

Pure Nash equilibrium (PNE) and Nash PoA (N-PoA) are defined in the full version.

In the above, the supremum is taken over *all* Bayesian settings. It is also possible to restrict the class of settings, for example, to those with valuation profiles that satisfy the SC property, or bidding strategies that satisfy a NOB property. Such restriction will, in general, improve the PoA, and we leverage this method to achieve our results.

Fixing the mechanism and agents' signals, it is easy to observe that every *EPE* is a *PNE*, and every *PNE* is a *BNE* for every signal distribution. This directly implies that $\text{EP-POA} \leq \text{NE-POA} \leq \text{B-POA}$. Thus, lower bounds for EP-POA and upper bounds for B-POA propagate to the other notions. For further discussion, see the full version.

3 Single Item: A Positive Result

In this section, we focus on single-item settings. Our main result is a (parameterized) property which, along with SC, leads to good PoA bounds.

Recall the setting described above, where the signal of, say, agent 3 affects only the value of herself and agent 1, and no other agent. In the opposite extreme, a change in an agent's signal affects the values of all other agents equally. That is: $\forall i, \forall j, j' \neq i :$

$$v_j(s_i + \delta, \mathbf{s}_{-i}) - v_j(\mathbf{s}) = v_{j'}(s_i + \delta, \mathbf{s}_{-i}) - v_{j'}(\mathbf{s}).$$

We call this condition homogeneous influence. One can verify that homogeneous influence holds in our running example – the resale model, as well as in other prominent settings like Klemperer's Wallet-game (Klemperer 1998), Common-values (Wilson 1969), private values, and private/common value interpolation (Bergemann and Morris 2013).

We find that homogeneous influence, together with SC, ensures full efficiency.

⁷Ex-post equilibria are not necessarily deterministic, but as we use these equilibria in the context of lower bounds, the restriction to deterministic only strengthens our results.

Proposition 3.1. *In single item settings with SC and homogeneous influence, every equilibrium of the GVA is fully efficient.*

The following question arises: how robust is this result? In other words, how would the PoA deteriorates as we move away from homogeneous influence? To this end, we introduce the following parameterized property, which measures how far a valuation profile is from homogeneous influence.

Definition 3.2 (γ -heterogeneity in signal-value impact). *A valuation profile is γ -heterogeneous in signal-value impact, or in short γ -heterogeneous, if for every agent i , two other agents j, j' , signal profile \mathbf{s} and $\delta > 0$,*

$$\gamma(v_j(s_i + \delta, \mathbf{s}_{-i}) - v_j(\mathbf{s})) \geq v_{j'}(s_i + \delta, \mathbf{s}_{-i}) - v_{j'}(\mathbf{s}).$$

By definition, γ is always at least 1; the special case of $\gamma = 1$ is homogeneous influence.

Our main result in this section is a tight bound for the B-POA for γ -heterogeneous, SC valuations. Our result extends beyond SC, to c -SC profiles (Eden et al. 2019), where the PoA degrades gracefully with the parameter c (SC is c -SC for $c = 1$).

Theorem 3.3 (Main positive result for single-item). *Consider a single-item setting with γ -heterogeneous, c -SC, continuous valuations. The B-POA of GVA under NOB is bounded by $1 + \max\{\gamma, c\}$. For SC valuations the B-POA is tightly γ . These results hold even for correlated signals.*

Theorem 3.3 is essentially tight, as illustrated in the full version. We remark that the upper bound of $1 + \max\{\gamma, c\}$ and the lower bound apply also with respect to 2PA with interdependence. Before proving Theorem 3.3 we introduce the following useful technical Lemma:

Lemma 3.4. *Let \mathbf{v} be a valuation profile satisfying γ -heterogeneity and c -SC. For every two agents i, j , signal profile \mathbf{s} and coordinate-wise non-negative vector $\delta = (\delta_1, \dots, \delta_{i-1}, 0, \delta_{i+1}, \dots, \delta_n)$, if $v_j(\mathbf{s}) \geq v_i(\mathbf{s})/d$ for $d \geq \max\{\gamma, c\}$ then $v_j(\mathbf{s} + \delta) \geq v_i(\mathbf{s} + \delta)/d$.*

Proof of Theorem 3.3. Let \mathbf{s} be a signal profile and \mathbf{b} be a bid profile such that $\mathbf{b} \leq \mathbf{s}$. Let $i \in \operatorname{argmax}_j v_j(\mathbf{s})$, and $w(\mathbf{b})$ be the winner under \mathbf{b} . We prove that $u_i((s_i, \mathbf{b}_{-i}); \mathbf{s}) \geq v_i(\mathbf{s}) - \max\{\gamma, c\}v_{w(\mathbf{b})}(\mathbf{s})$. (1)

We distinguish between the following two cases:

Case 1: Bidder i does not win the item under bidding profile (s_i, \mathbf{b}_{-i}) . Thus, her utility is 0. Moreover, there exists a bidder $j \neq i$ such that $v_j(s_i, \mathbf{b}_{-i}) \geq v_i(s_i, \mathbf{b}_{-i})$. Since bidder $w(\mathbf{b})$ wins at \mathbf{b} , $v_{w(\mathbf{b})}(\mathbf{b}) \geq v_j(\mathbf{b})$. By Lemma 3.4 (using $d = \max\{\gamma, c\}$), $\max\{\gamma, c\}v_{w(\mathbf{b})}(s_i, \mathbf{b}_{-i}) \geq v_j(s_i, \mathbf{b}_{-i}) \geq v_i(s_i, \mathbf{b}_{-i})$. Applying Lemma 3.4 again, we get $\max\{\gamma, c\}v_{w(\mathbf{b})}(\mathbf{s}) \geq v_i(\mathbf{s})$. Therefore, $u_i((s_i, \mathbf{b}_{-i}); \mathbf{s}) = 0 \geq v_i(\mathbf{s}) - \max\{\gamma, c\}v_{w(\mathbf{b})}(\mathbf{s})$.

Case 2: Bidder i wins the item under bidding profile (s_i, \mathbf{b}_{-i}) . Let $b_i^* \leq s_i$ be i 's critical bid given \mathbf{b}_{-i} . $u_i((s_i, \mathbf{b}_{-i}); \mathbf{s}) = v_i(\mathbf{s}) - v_i(b_i^*, \mathbf{b}_{-i})$. If $w(\mathbf{b}) = i$, $u_i((s_i, \mathbf{b}_{-i}); \mathbf{s}) \geq 0 \geq v_i(\mathbf{s}) - \max\{\gamma, c\}v_{w(\mathbf{b})}(\mathbf{s})$. Otherwise, $w(\mathbf{b}) \neq i$. Let $b'_i \geq b_i$ be the lowest signal such that $v_i(b'_i, \mathbf{b}_{-i}) \in \operatorname{argmax}_k v_k(b'_i, \mathbf{b}_{-i})$. There exists such $b'_i \leq s_i$ since i does not win under bidding profile \mathbf{b} but wins under bidding profile (s_i, \mathbf{b}_{-i}) . By the same argument

and the continuity of the valuations, there exists $j \neq i$ such that $v_j(b'_i, \mathbf{b}_{-i}) \in \operatorname{argmax}_k v_k(b'_i, \mathbf{b}_{-i})$. By Corollary 3.4, $\max\{\gamma, c\}v_{w(\mathbf{b})}(b'_i, \mathbf{b}_{-i}) \geq v_j(b'_i, \mathbf{b}_{-i}) = v_i(b'_i, \mathbf{b}_{-i}) \geq v_i(b_i^*, \mathbf{b}_{-i})$. Therefore, $u_i((s_i, \mathbf{b}_{-i}); \mathbf{s}) \geq v_i(\mathbf{s}) - \max\{\gamma, c\}v_{w(\mathbf{b})}(b'_i, \mathbf{b}_{-i}) \geq v_i(\mathbf{s}) - \max\{\gamma, c\}v_{w(\mathbf{b})}(\mathbf{s})$, where the last inequality follows by $\mathbf{b} \leq \mathbf{s}$ and the monotonicity of the valuations. This concludes the proof of (1).

We can now establish the bound on the B-POA: Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a BNE satisfying NOB; Let $\mathbb{I}_i(\mathbf{s})$ be the indicator variable of the event $i = \operatorname{argmax}_j v_j(\mathbf{s})$ (breaking ties arbitrarily). We get:

$$\text{EQ}(\sigma) \geq \mathbb{E}_{\mathbf{s}} \left[\sum_i u_i(\sigma(\mathbf{s}); \mathbf{s}) \right] \quad (2)$$

$$\geq \sum_i \mathbb{E}_{\mathbf{s}} [u_i((s_i, \sigma_{-i}(\mathbf{s}_{-i})); \mathbf{s})] \quad (3)$$

$$\geq \sum_i \mathbb{E}_{\mathbf{s}} [\mathbb{I}_i(\mathbf{s}) \cdot (v_i(\mathbf{s}) - \max\{\gamma, c\}v_{w(\sigma(\mathbf{s}))}(\mathbf{s}))] \quad (4)$$

$$= \mathbb{E}_{\mathbf{s}} [\max_i v_i(\mathbf{s})] - \max\{\gamma, c\} \mathbb{E}_{\mathbf{s}} [v_{w(\sigma(\mathbf{s}))}(\mathbf{s})] \\ = \text{OPT} - \max\{\gamma, c\} \text{EQ}(\sigma), \quad (5)$$

where (2) holds since the sum of utilities is dominated by the welfare. (3) holds by the equilibrium hypothesis (and linearity of expectation). (4) holds by (1). With exact SC, the bound further improves to γ . \square

4 Multiple Items: A Positive Result

In this section, we study settings with multiple items and unit-demand agents. After some multi-item preliminaries, we identify a condition limiting knowledge asymmetry among agents, which can be used to accommodate positive PoA results. We then use it to prove a PoA guarantee for a simple simultaneous auction in the $n \geq m$ regime (Theorem 4.7). (For the $n \ll m$ regime see the full version.)

4.1 Multiple Items Preliminaries

Participation. Similarly to (Syrgekianis and Tardos 2013), we assume agents can communicate to the item auctions whether or not they wish to participate and compete for the item being sold (in addition to their signal report). The overall report of agent i is (b_i, a_i) , where b_i is the agent's vector of reports (bids) of her m signals, and a_i is a participation vector with $a_{i\ell} = 1$ if the agent participates in the auction for item ℓ and $a_{i\ell} = 0$ otherwise. The agent's (mixed) strategy $\sigma_i(s_i)$ given her vector of signals is then a distribution over possible reports (b_i, a_i) . In the full version we show that a bad PoA may occur when agents have to participate in each of the auctions.

Notation and assumptions. Let \mathbf{b} (respectively, \mathbf{a}) be a profile of n signal reports b_i (respectively, a_i), and denote by $\sigma(\mathbf{s})$ a strategy profile given signal profile \mathbf{s} , that is, a distribution over (\mathbf{b}, \mathbf{a}) pairs. We use $X_i(\mathbf{b}, \mathbf{a})$ to denote the item subset allocated in total by the separate auctions to agent i , given the bid and participation profiles \mathbf{b}, \mathbf{a} .

For simplicity, in this section, we state our results for valuation profiles satisfying SC and γ -heterogeneity. The proofs

are deferred to the full version where we prove stronger results for the relaxed notion of c-SC. Standard NOB assumptions for multiple items essentially restrict the sum of every agent's bids for items in a subset X to at most the agent's value for X . We now formalize this notion under IDV:

Definition 4.1 (Multi-item NOB for IDV). *With multiple items, a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ satisfies NOB if for every agent i with signal s_i it holds that*

$$\mathbb{E}_{\substack{\mathbf{s}_{-i} \sim F_{-i}, \\ (\mathbf{b}, \mathbf{a}) \sim \sigma(\mathbf{s})}} \left[\sum_{\ell \in X_i(\mathbf{b}, \mathbf{a})} v_{i\ell}(b_{i\ell}, \mathbf{s}_{-i\ell}) \right] \leq \mathbb{E}_{\substack{\mathbf{s}_{-i} \sim F_{-i}, \\ (\mathbf{b}, \mathbf{a}) \sim \sigma(\mathbf{s})}} [v_i(X_i(\mathbf{b}, \mathbf{a}); \mathbf{s})]$$

Note that this coincides with the standard NOB assumption if valuations are private.

4.2 Limited Knowledge Asymmetry

We formulate a condition that enables our positive PoA result. We use the following definition:

Definition 4.2 (Truncated values and welfare). *For every agent i and item ℓ , the truncated value $\tilde{v}_{i\ell}$ given signal profile \mathbf{s} is $\tilde{v}_{i\ell}(\mathbf{s}) = \tilde{v}_{i\ell}(s_\ell) = \min_{j \neq i} v_{i\ell}(\mathbf{s}_{-j\ell}, 0_{j\ell})$. The truncated optimal matching $\tilde{m}(\mathbf{s})$ is a matching in $\operatorname{argmax}_{\text{matching } \mu} \sum_{(i,\ell) \in \mu} \tilde{v}_{i\ell}(\mathbf{s})$, and the expected truncated welfare is $\widetilde{\text{OPT}} = \mathbb{E}_{\mathbf{s} \sim F} \left[\sum_{(i,\ell) \in \tilde{m}(\mathbf{s})} \tilde{v}_{i\ell}(\mathbf{s}) \right]$.*

In words, truncated value $\tilde{v}_{i\ell}(\mathbf{s})$ is agent i 's value for item ℓ after the most significant signal except for i 's own has been zeroed out; $\tilde{m}(\mathbf{s})$ is an allocation that maximizes the social welfare with respect to the truncated values, given \mathbf{s} ; and $\widetilde{\text{OPT}}$ is the optimal social welfare in expectation over \mathbf{s} with respect to truncated values. Observe that by the monotonicity of values in signals, $\widetilde{\text{OPT}} \leq \text{OPT}$.

Given an interdependent setting, the *limited knowledge asymmetry* condition is the following: there exists a constant d such that $d \cdot \widetilde{\text{OPT}} \geq \text{OPT}$. (6)

Intuitively, if a truncated value $\tilde{v}_{i\ell}(\mathbf{s})$ is far from the true value $v_{i\ell}(\mathbf{s})$, this means that agent i 's value for item ℓ is largely determined by the information that is held by a single agent who is not i herself. This means there is information asymmetry among i and the informed agent. The condition in (6) rules out settings in which nearly all the welfare stems from such information asymmetry. In such extreme settings, all information shaping the values can be traced to a single agent – as in the classic and well-studied “drainage tract” model of Wilson (Wilson 1969; Milgrom 2004). We now give two opposite examples: a natural setting in which the condition in (6) holds (Example 4.3), and an extreme setting in which one agent has all the information (Example 4.5). For the second example we show a lower bound of $\Omega(m)$ on the PoA of a natural family of mechanisms.

Example 4.3 (Weighted-sum valuations revisited). *Recall the Resale model, introduced in Example 1.1. We extend this example to multiple items: Assume for simplicity that $n = m$. For every unit-demand agent i and item ℓ , let $v_{i\ell}(\mathbf{s}) = s_{i\ell} + \beta \sum_{j \neq i} s_{j\ell}$ for $\beta \leq 1$ (where the restriction on β maintains the SC property). We assume all signals are drawn i.i.d. from a monotone hazard rate (MHR) distribution (such as a uniform, normal, or exponential).*

Proposition 4.4. *In Example 4.3, for a sufficiently large n , $(1 + \epsilon)\widetilde{\text{OPT}} \geq \text{OPT}$.*

The proof of Proposition 4.4 follows from the symmetric nature of both the valuations themselves and the informativeness of the signals.

Example 4.5 (Agent 1 holds almost all information). *Consider n unit-demand agents with single-dimensional signals s_1, \dots, s_n , and signal spaces $S_1 = [0, 1]$ for agent 1 and $S_i = \{0\}$ for every agent $i \neq 1$. The agents have the following values for every item $\ell \in [m]$: Agent 1's value is $v_{1\ell} = s_1$, and for every agent $i \neq 1$, $v_{i\ell} = (1 - \epsilon)s_1$. Note these valuations are SC and γ -homogeneous.*

In Example 4.5, the truncated values are $\tilde{v}_{1\ell} = s_1$ and $\tilde{v}_{i\ell} = 0$ for $i \neq 1$. We get that $\widetilde{\text{OPT}} = s_1$ while $\text{OPT} > m(1 - \epsilon)s_1$, so there is no constant d for which (6) holds.

Proposition 4.6. *For each bidder let $h_{i\ell}(\mathbf{s})$ be a monotone function such that $\forall \mathbf{s}, h_{i\ell}(\mathbf{s}) \leq v_{i\ell}(\mathbf{s})$, and $h_{i\ell}(s_{i\ell}, \mathbf{0}_{-i\ell}) = v_{i\ell}(s_{i\ell}, \mathbf{0}_{-i\ell})$. Every mechanism that allocates each item separately to the highest-valued agent according to $h_{i\ell}(\mathbf{s})$, and charges an agent zero payment if no-one else participates, has EP-POA of $\Omega(m)$ even under NOB. In particular, by $h_{i\ell}(\mathbf{s}) = v_{i\ell}(\mathbf{s})$ we get that this holds for simultaneous 2PA or GVA auctions.*

The proof of Proposition 4.6 uses Example 4.5. As shown in the next section, the above proposition holds for the mechanism we develop and use to get good PoA guarantees. We leave open the question of whether there exists a simple mechanism with good PoA for settings such as Example 4.5 that do not satisfy (6).

4.3 Positive Result: $n \geq m$ Regime

We present a simple simultaneous item-bidding mechanism and show that every BNE of it achieves an $O(\gamma^2)$ -approximation to $\widetilde{\text{OPT}}$. For valuations satisfying Equation (6), it implies an $O(d\gamma^2)$ -approximation to OPT .

Intuition. Before introducing the mechanism, we provide some intuition that helps shed light on the choice of our mechanism and its analysis. Specifically, we emphasize how our analysis is different from the smoothness framework, which has become the standard technique for establishing PoA results for simple auctions.

A key step in the smoothness paradigm is to find an appropriate hypothetical deviation for each player such that for any bids of the other bidders, the utility achieved by the deviation is lower bounded by some fraction of the player's contribution to the optimal welfare, less some error term. For the case of independent private values and unit-demand bidders⁸, Christodoulou et al. (2016) show that the hypothetical deviation of going *all-in* for item ℓ gives the following guarantee:⁹ $u_i((v_{i\ell}, \mathbf{0}_{i-\ell}), \mathbf{b}_{-i}) \geq v_{i\ell} - \max_{j \neq i} b_{j\ell}$ (7).

Inequality (7) suffices to give PoA bounds in full information settings. Moreover, this bound also applies to settings with incomplete information via an extension theorem.

⁸For simultaneous second price auctions.

⁹Bidding her true value $v_{i\ell}$ for item ℓ and 0 for all other items.

Even though the agent does not know the realization of other agents' valuations, she can sample from the value distribution of others, and go "all-in" for the item ℓ she receives in the optimal allocation in the sampled market, thus effectively simulating the contribution of this agent to the optimal assignment (the first term in the right-hand side of (7)). This is termed the 'doppelgänger' technique. The error term is then bounded by tying it to an allocation for sampled equilibrium bids and using a NOB assumption.

Unfortunately, there are some major obstacles to generalizing this type of analysis to IDV. Whereas one can derive a smoothness-type inequality, its right-hand side expression would take the form of $\gamma \tilde{v}_{i\ell}(\mathbf{s}) - \max_{j \neq i} v_{j\ell}(b_{j\ell}, s_{-j\ell})$. Using this inequality, one can derive PoA results for full information games¹⁰, but these bounds would not generalize to incomplete information settings for the following reasons:

- The doppelgänger technique cannot be directly applied. Agent i is only aware of s_i and is uninformed of s_{-i} , which affects her valuations. Therefore, sampling signals of other bidders t_{-i} , and going all-in accordingly, will simulate a different value distribution than the desired one, as the bid takes into account only s_i and not s_{-i} .
- The error term in Christodoulou et al. (2016) crucially depends only on the bids of other bidders, but in IDV the error term depends also on the agents' *real* signals, which prevents from their analysis to go through.
- The benchmark in our setting should be the truncated optimal allocation and not the true optimal allocation.

The following observation drives our analysis: the doppelgänger technique can be successfully applied when considering what we refer to as *privatized valuations*; namely, the valuation of agent i under i 's private signal, while zeroing out the contribution of others' signals. We can then use the γ -heterogeneous condition to claim that the effect of other bidders' signals on the valuation of two different agents is roughly the same.

The discussion above motivates the design of the mechanism we propose, which is a simultaneous second-price item auction with respect to the privatized valuations (see Fig. 1). For the sake of analysis, we split the benchmark into two terms. The first term (SELF) accounts for the privatized valuations, and the second term (OTHER) accounts for the contribution of the signals of other bidders to the benchmark (see Equation (8)). We proceed by providing smoothness-type inequalities for each of these terms separately. Specifically, for each term, we identify an appropriate hypothetical deviation, such that the expected welfare in equilibrium "covers" it. We formalize the intuition given above.

Privatized values. Given a (reported) signal profile \mathbf{b}_ℓ for item ℓ , agent j 's *privatized value* $\hat{v}_{j\ell}$ for ℓ is her value when other agents' signals are set to zero: $\hat{v}_{j\ell}(\mathbf{b}_\ell) = v_{j\ell}(b_{j\ell}, \mathbf{0}_{-j\ell})$. We use $\hat{v}_{j\ell}(\mathbf{b}_\ell)$ and $v_{j\ell}(b_{j\ell}, \mathbf{0}_{-j\ell})$ interchangeably. Observe that the privatized value is upper-bounded by the truncated value $\hat{v}_{j\ell}(\mathbf{b}_\ell) \leq \tilde{v}_{j\ell}(\mathbf{b}_\ell)$.

¹⁰We do not include this derivation in the manuscript since it does not help us in proving a bound on the B-PoA.

Theorem statement. We focus on the simple mechanism described in Fig. 1, which consists of simultaneous *privatized* second-price auctions. This mechanism allocates every item separately by running a second-price auction over the privatized values of agents participating for this item. Our main theorem in this section is the following:

Simultaneous privatized second-price auctions
Input: Bid and participation profiles \mathbf{b} and \mathbf{a} .
Output: An allocation and payments.
For every item ℓ (simultaneously):

- For every agent j , compute the privatized value $\hat{v}_{j\ell}(\mathbf{b}_\ell) = v_{j\ell}(b_{j\ell}, \mathbf{0}_{-j\ell})$
- Allocate item ℓ to agent $i \in \operatorname{argmax}\{\hat{v}_{j\ell}(\mathbf{b}_\ell) | a_{j\ell} = 1\}$
- Charge agent i a payment $p_{i\ell}(\mathbf{b}_\ell, \mathbf{a}_\ell) = \max\{0, \max_{j \neq i} \{\hat{v}_{j\ell}(\mathbf{b}_\ell) | a_{j\ell} = 1\}\}$

Figure 1: Simultaneous privatized second-price auctions.

Theorem 4.7. Consider a setting with m items, $n \geq m$ unit-demand agents, and γ -heterogeneous, SC valuations. If $d \cdot \widetilde{\text{OPT}} \geq \text{OPT}$ for some constant d , then the B-POA of simultaneous privatized second-price auctions (See Fig. 1) under no-overbidding is $O(d\gamma^2)$.

The condition that $d \cdot \widetilde{\text{OPT}} \geq \text{OPT}$ in Theorem 4.7 is necessary, as by $h_{i\ell}(\mathbf{s}) = v_{i\ell}(s_{i\ell}, \mathbf{0}_{-i\ell})$, Proposition 4.6 holds for the simultaneous privatized second-price auction as well. The remainder of the section is dedicated to proving Theorem 4.7. We use the following notation: Let $u_i((\mathbf{b}, \mathbf{a}); \mathbf{s})$ denote agent i 's utility under bidding profile (\mathbf{b}, \mathbf{a}) and signal profile \mathbf{s} . Let $p_{i\ell}(\mathbf{b}_\ell, \mathbf{a}_\ell)$ be as defined in Fig. 1. We first decompose the welfare as stated above. Recall from Definition 4.2 that $\tilde{m}(\mathbf{s})$ is the matching that maximizes the social welfare with respect to the truncated values at signal profile \mathbf{s} . We decompose $\widetilde{\text{OPT}}$ as follows:

$$\begin{aligned} \widetilde{\text{OPT}} = & \mathbb{E}_{\mathbf{s}} \left[\underbrace{\sum_{(i,\ell) \in \tilde{m}(\mathbf{s})} v_{i\ell}(s_{i\ell}, \mathbf{0}_{-i\ell})}_{\text{SELF}} \right] \\ & + \mathbb{E}_{\mathbf{s}} \left[\underbrace{\sum_{(i,\ell) \in \tilde{m}(\mathbf{s})} \tilde{v}_{i\ell}(\mathbf{s}) - v_{i\ell}(s_{i\ell}, \mathbf{0}_{-i\ell})}_{\text{OTHER}} \right]. \end{aligned} \quad (8)$$

In order to prove Theorem 4.7, we prove that the equilibrium 'covers' the two terms in the decomposition above.

Lemma 4.8. For every BNE σ , $2\text{EQ}(\sigma) \geq \text{SELF}$.

Lemma 4.9. For every BNE σ , $O(\gamma^2) \cdot \text{EQ}(\sigma) \geq \text{OTHER}$.

These lemmas are proven using two deviations which make use of smoothness-type inequalities. For Lemma 4.8, we use a deviation based on the doppelgänger technique described above. For Lemma 4.9, we show that an arbitrary deviation can be used to bound the OTHER term.

Proof of Theorem 4.7. Combining Lemmas 4.8 and 4.9 with Equations (8) and (6), for every BNE σ , we have $(2 + O(\gamma^2)) \text{EQ}(\sigma) \geq \text{SELF} + \text{OTHER} = \widetilde{\text{OPT}} \geq \text{OPT}/d$, concluding the proof. \square

Acknowledgments

The work of A. Eden, M. Feldman and O. Zviran was partially supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866132), and by the Israel Science Foundation (grant number 317/17). The work of I. Talgam-Cohen was supported by the Israel Science Foundation (grant No. 336/18) and by the Taub Family Foundation. We deeply thank David Parkes for helpful discussions.

References

- Ausubel, L. M.; Milgrom, P.; et al. 2006. The lovely but lonely Vickrey auction. *Combinatorial auctions* 17: 22–26.
- Ausubel, L. M.; et al. 2000. A generalized Vickrey auction. In *Econometric Society World Congress*. Citeseer.
- Batani, M.; Dehghani, S.; Hajiaghayi, M.; and Seddighin, S. 2015. Revenue Maximization for Selling Multiple Correlated Items. In Bansal, N.; and Finocchi, I., eds., *Algorithms - ESA 2015 - 23rd Annual European Symposium*, volume 9294 of *Lecture Notes in Computer Science*, 95–105. Springer.
- Bergemann, D.; and Morris, S. 2013. Robust predictions in games with incomplete information. *Econometrica* 81(4): 1251–1308.
- Bhawalkar, K.; and Roughgarden, T. 2011. Welfare Guarantees for Combinatorial Auctions with Item Bidding. In Randall, D., ed., *Proceedings of the 22 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, 700–709. SIAM.
- Chawla, S.; Fu, H.; and Karlin, A. R. 2014. Approximate revenue maximization in interdependent value settings. In *ACM Conference on Economics and Computation, EC ’14*, 277–294. ACM.
- Chawla, S.; Malec, D. L.; and Sivan, B. 2015. The power of randomness in Bayesian optimal mechanism design. *Games Econ. Behav.* 91: 297–317.
- Christodoulou, G.; Kovács, A.; and Schapira, M. 2016. Bayesian Combinatorial Auctions. *J. ACM* 63(2): 11:1–11:19.
- Christodoulou, G.; Kovács, A.; Sgouritsa, A.; and Tang, B. 2016. Tight Bounds for the Price of Anarchy of Simultaneous First-Price Auctions. *ACM Trans. Economics and Comput.*, *EC ’16* 4(2): 9:1–9:33.
- Constantin, F.; Ito, T.; and Parkes, D. C. 2007. Online auctions for bidders with interdependent values. In Durfee, E. H.; Yokoo, M.; Huhns, M. N.; and Shehory, O., eds., *6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, 110. IFAAMAS.
- Dasgupta, P.; and Maskin, E. 2000. Efficient Auctions. *The Quarterly Journal of Economics* 115(2): 341–388.
- Dobzinski, S. 2011. An impossibility result for truthful combinatorial auctions with submodular valuations. In Fortnow, L.; and Vadhan, S. P., eds., *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011*, 139–148. ACM.
- Eden, A.; Feldman, M.; Fiat, A.; and Goldner, K. 2018. Interdependent Values without Single-Crossing. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 369. ACM.
- Eden, A.; Feldman, M.; Fiat, A.; Goldner, K.; and Karlin, A. R. 2019. Combinatorial Auctions with Interdependent Valuations: SOS to the Rescue. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019*, 19–20. ACM.
- Eden, A.; Feldman, M.; Talgam-Cohen, I.; and Zviran, O. 2020. Price of Anarchy of Simple Auctions with Interdependent Values. *arXiv* URL <https://arxiv.org/abs/2011.00498>.
- Feldman, M.; Fu, H.; Gravin, N.; and Lucier, B. 2013. Simultaneous auctions are (almost) efficient. In Boneh, D.; Roughgarden, T.; and Feigenbaum, J., eds., *Symposium on Theory of Computing Conference, STOC’13*, 201–210. ACM.
- Immorlica, N.; Singla, S.; and Waggoner, B. 2020. Prophet Inequalities with Linear Correlations and Augmentations. *CoRR* abs/2001.10600.
- Ito, T.; and Parkes, D. C. 2006. Instantiating the contingent bids model of truthful interdependent value auctions. In Nakashima, H.; Wellman, M. P.; Weiss, G.; and Stone, P., eds., *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, 1151–1158. ACM.
- Jehiel, P.; and Moldovanu, B. 2001. Efficient design with interdependent valuations. *Econometrica* 69(5): 1237–1259.
- Klemperer, P. 1998. Auctions with almost common values: The Wallet Game and its applications. *European Economic Review* 42(3-5): 757–769.
- Krishna, V. 2009. *Auction theory*. Academic press.
- Li, S. 2017. Obviously strategy-proof mechanisms. *American Economic Review* 107(11): 3257–87.
- Maskin, E. 1992. *Auctions and Privatization*, 115–136. J.C.B. Mohr Publisher.
- Milgrom, P. R. 2004. *Putting auction theory to work*. Cambridge University Press.
- Milgrom, P. R.; and Weber, R. J. 1982. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* 50(5): 1089–1122.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6(1): 58–73.
- Robu, V.; Parkes, D. C.; Ito, T.; and Jennings, N. R. 2013. Efficient Interdependent Value Combinatorial Auctions with Single Minded Bidders. In Rossi, F., ed., *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 339–345. IJCAI/AAAI.
- Roughgarden, T. 2016. *Twenty lectures on algorithmic game theory*. Cambridge University Press.

Roughgarden, T.; Syrgkanis, V.; and Tardos, É. 2017. The Price of Anarchy in Auctions. *J. Artif. Intell. Res.* 59: 59–101.

Roughgarden, T.; and Talgam-Cohen, I. 2016. Optimal and Robust Mechanism Design with Interdependent Values. *ACM Trans. Economics and Comput.* 4(3): 18:1–18:34.

Syrgkanis, V.; and Tardos, É. 2013. Composable and efficient mechanisms. In *Symposium on Theory of Computing Conference, STOC'13*, 211–220. ACM.

Wilson, R. 1969. Competitive Bidding with Disparate Information. *Management Science* 15(7): 446–452.