

First-Price Auctions with General Information Structures: A Short Introduction

DIRK BERGEMANN

Yale University

and

BENJAMIN BROOKS

University of Chicago

and

STEPHEN MORRIS

Princeton University

We explore the impact of private information in sealed-bid first-price auctions. For a given symmetric and arbitrarily correlated prior distribution over values, we characterize the impact that the structure of private information has on bidding behavior, and the sharing of surplus between the seller and the bidder. Our results provide lower bounds and upper bounds for bids and revenues across all information structures. Our work has implications for the identification of value distributions from data on winning bids and for the informationally robust comparison of alternative bidding mechanisms.

1. INTRODUCTION

In a recent paper, [Bergemann et al. 2017a], we derive results about equilibrium behavior in the first-price auction that hold across *all* common-prior information structures. The purpose of this letter is to give an informal introduction into the results. At the end we offer a brief discussion of related work.

For a given prior distribution over value profiles, we study what can happen for all information structures specifying bidders' information about their own and others' values. For any value distribution, we identify a lower bound on the distribution of winning bids in the sense of first-order stochastic dominance. In other words, no matter what the true information structure is, the distribution of winning bids must first-order stochastically dominate the bound that we describe. In addition, when the prior distribution of values is symmetric, we construct an equilibrium and an information structure in which this lower bound is attained. This minimum winning-bid distribution therefore pins down the minimum amount of revenue that can be generated by the auction in expectation. Moreover, the minimum winning-bid distribution is attained in an efficient equilibrium. As a result, this equilibrium also attains an upper bound on the expected surplus of the bidders, which is equal to the maximum feasible surplus minus minimum revenue.

Let us give a brief intuition for how our bounds are obtained. If the distribution of winning bids places too high of a probability on low bids, then some bidder would find that a modest increase in their bid would result in a relatively large increase in the probability of winning, so that such a deviation would be attractive. For example, it cannot be that all bidders tie with a bid of zero with probability

one, for then some bidder could increase his bid a token positive amount and win the auction outright. This suggests that the relevant constraints for pinning down minimum bidding are those associated with deviating to higher bids. Indeed, we show that the minimum winning-bid distribution is characterized by bidders being indifferent to *all* upward deviations.

To characterize the minimum, it turns out to be sufficient to look at a relatively small class of such deviations: For some bid b , we say that a bidder *uniformly deviates up to b* if he switches to bidding b whenever he would have bid less than b in equilibrium. It is clearly necessary for equilibrium that the bidders should not want to uniformly deviate upward. Moreover, it turns out that the change in a bidder's surplus from a uniform upward deviation depends only on the distribution of winning bids, and not on the distribution of losing bids. This motivates a relaxed program in which we minimize the distribution of winning bids, subject only to the uniform upward incentive constraints. The solution to this relaxed program gives us a lower bound on the winning-bid distribution. We illustrate this proof strategy with an example in which there are two bidders and a uniformly distributed common value.

2. MODEL

We consider the sale of a single unit of a good by a first-price auction. There are N individuals who bid for the good, indexed by $i \in \mathcal{N} = \{1, \dots, N\}$, each of whom has a value which lies in the compact interval $V = [\underline{v}, \bar{v}] \subset \mathbb{R}_+$. The bidders are assumed to be risk neutral and to have quasilinear preferences over the allocation and payments. Values are jointly distributed according to a probability measure $\mu \in \Delta(V^N)$. Each individual $i \in \mathcal{N}$ submits a bid $b_i \in B = [0, \bar{v}]$, and the winner is selected uniformly from among the high bidders. Bidders may receive additional information about the profile of values, beyond knowing the prior distribution. This information comes in the form of signals that are correlated with the profile of values. An *information structure* is a collection $\mathcal{S} = (\{S_i\}_{i=1}^N, \pi)$, where the S_i are measurable spaces and $\pi : V^N \rightarrow \Delta(S)$ is a measurable mapping from profiles of values to probability measures over $S = \times_{i=1}^N S_i$. The interpretation is that S_i is the set of bidder i 's signals and π describes the conditional joint distribution of signals given values. For a fixed information structure \mathcal{S} , the first-price auction is a game of incomplete information, in which bidders' strategies are measurable mappings $\sigma_i : S_i \rightarrow \Delta(B)$ from signals to probability measures over bids. A (*Bayes Nash*) *equilibrium* is a strategy profile $\sigma = (\sigma_1, \dots, \sigma_N)$ such that each bidder's strategy maximizes their ex ante expected surplus, given the strategies of the other bidders.

3. A PURE-COMMON-VALUE EXAMPLE

We will illustrate the main results with a simple example. There are two bidders who share a common value for the good, which is uniformly distributed between 0 and 1. Since we assume there is no reservation price in the auction, the good is always allocated, regardless of the particular information structure and equilibrium. As both bidders have the same value, all equilibria are socially efficient and result in a total surplus of 1/2. There may, however, be variation across information structures and equilibria in how this surplus is split between the bidders and the

seller.

We allow bidders to observe arbitrary and possibly correlated signals about the common value. Indeed, at one extreme the bidders' signals are perfectly correlated, so that they have exactly the same information about the value. For example, the bidders might have no information beyond the prior, so that they both expect the good to be worth $1/2$, or the bidders might both observe the true value of the good, so that they know the good's value exactly. In any such case, the bidders will compete the price up to the interim expected value of the good, which results in zero bidder surplus and expected revenue of $1/2$. These examples illustrate our later general result that, unless we make additional assumptions about what the bidders know, a tight upper bound on revenue is the efficient surplus.

When the bidders have private information, the distribution of ex-ante surplus can be rather different. An important case has been studied by [Engelbrecht-Wiggans et al. 1983]: bidder 1 observes the true value while the bidder 2 is uninformed and observes nothing. In the case of the uniform distribution, the resulting equilibrium involves bidder 1 bidding $v/2$ and the uninformed bidder randomizing uniformly over the interval $[0, 1/2]$. This equilibrium results in a surplus of $1/6$ for bidder 1, a surplus of 0 for bidder 2, and revenue of $1/3$.

What can we say about the outcome of the auction more generally? In the analysis of [Engelbrecht-Wiggans et al. 1983], the informed bidder strictly prefers his equilibrium bid over any other bid. This suggests that it might be possible to construct other information structures in which revenue is even lower. Here is one such construction: The two bidders receive signals $s \in [0, 1]$ that are independent draws from the cumulative distribution $F(s) = \sqrt{s}$, so that the distribution of the maximum signal is standard uniform, the same as the common signal value. Moreover, the signals and the value are correlated so that the maximum signal is exactly equal to the value:

$$v = \max \{s_1, s_2\}. \quad (1)$$

This information structure admits an equilibrium in which the bidders use the following monotonic pure strategy:

$$\sigma(s) = \frac{1}{\sqrt{s}} \int_{x=0}^s x \frac{dx}{2\sqrt{x}} = \frac{s}{3}.$$

Thus, the equilibrium bid is the expectation of other bidder's signal, conditional on it being below s . We will presently verify that these strategies constitute an equilibrium, but let us first note the implied welfare outcomes: the winning bid will always be $\max_i s_i/3 = v/3$, so that revenue is $1/6$. Bidder surplus is therefore $1/3$, which is twice as much as the bidders obtained in the proprietary information model of [Engelbrecht-Wiggans et al. 1983]

Let us now verify that these strategies comprise an equilibrium. It is well known that these strategies are an equilibrium of a slightly different model, in which the bidders receive the same signals drawn from the same distribution, but in which each bidder's signal is their private value. In other words, when there are independent private values (IPV), the equilibrium bid is the expectation of the other bidder's value (i.e., signal) conditional on it being less than one's own signal [Krishna 2002].¹

¹Indeed, this is a necessary consequence of the revenue equivalence between first- and second-price

Now, in our common-value model, a bidder who deviates by bidding $s'/3$ for some $s' < s$ will only win when their own signal was the highest signal, and therefore equal to the common value. Thus, a downward deviator's surplus looks exactly the same as in the as-if IPV setting, and we can immediately conclude that bidders do not want to deviate down. On the other hand, if a bidder deviates up to $s'/3$ with $s' > s$, the bidder continues to win on the event that they had the high signal, and now wins on some events when it was the other bidder who had the high signal, which was the true value. The deviator's surplus is

$$\left(s - \frac{s'}{3}\right)\sqrt{s} + \int_{x=s}^{s'} \left(x - \frac{s'}{3}\right) \frac{1}{2\sqrt{x}} dx = \frac{2}{3}s\sqrt{s}$$

which is independent of s' . In other words, bidders are exactly indifferent to *all* upward deviations!

In fact, no matter how one structures the information or the equilibrium strategies, it is impossible for revenue to fall below the level attained in this example, i.e., $1/6$ is a tight lower bound on revenue when there are two bidders and there is a pure common value that is standard uniform. Moreover, not only is it impossible for revenue to fall below the level of the example, but the distribution of winning bids in any equilibrium under any information structure must first-order stochastically dominate the winning-bid distribution in the equilibrium we just constructed.

We sketch the argument as follows. Any equilibrium under any information structure will induce a distribution of the winning bid, which we denote by $H(b)$. We can further decompose this into a distribution of the winning bid conditional on the value, $H(b|v)$, so that

$$H(b) = \int_{v=0}^1 H(b|v) dv.$$

Let us suppose that the equilibrium is symmetric and that H has no atoms. (These assumptions are dispensed with in the general argument.) Then all bidders are equally likely to win at each winning bid, and hence, each bidder's surplus is simply

$$\frac{1}{2} \int_{v=0}^1 \int_{x=0}^1 (v-x)H(dx|v) dv.$$

Now, fix any bid b , and consider the following deviation for bidder i : whenever the equilibrium strategy says to bid less than b , bid b , and otherwise bid according to the equilibrium strategy. We refer to this as a *uniform deviation up to b* . We claim that the surplus from this deviation is

$$\int_{v=0}^1 \left((v-b)H(b|v) + \frac{1}{2} \int_{x=b}^1 (v-x)H(dx|v) \right) dv.$$

The reason is that the deviating bidder always bids at least b , so that whenever the winning bid would have been less than b , the deviator now wins (and by hypothesis the ex ante probability of a tie at b is zero). On the other hand, if the winning bid under the equilibrium strategies would have been greater than b , then the outcome of the auction is unchanged: if the deviator would have won in equilibrium, then

auctions in the IPV setting.

the deviation does not affect his bid, and if he would not have won, any deviation would be to a bid of b , which is not high enough to win.

Thus, a necessary condition for a winning bid distribution to arise in equilibrium is that no bidder wants to uniformly deviate up, which is equivalent to

$$\int_{v=0}^1 (v-b)H(b|v)dv \leq \frac{1}{2} \int_{v=0}^1 \int_{x=0}^b (v-x)H(dx|v)dv.$$

Now, consider a distribution $H(b)$ that is the expectation of some conditional distribution $H(b|v)$ that satisfies these inequalities for all b . We show that the set of winning bid distributions that can be so induced has a smallest element in the first-order stochastic dominance ordering, which is equivalent to the pointwise ordering on $H(b)$. The conditional distribution that attains this minimum is constructed in a “greedy” manner, whereby lower values are associated with lower winning bids.² In particular, there is a monotonic *winning bid function* $\beta(v)$ so that $H(\cdot|v)$ puts probability one on $\beta(v)$. With this additional structure, the incentive constraint for a uniform upward deviation to $\beta(w)$ can be rewritten as

$$\int_{v=0}^w (v-\beta(w))dv \leq \frac{1}{2} \int_{v=0}^1 \int_{x=0}^w (v-\beta(v))dv,$$

which rearranges to

$$\beta(w) \geq \frac{1}{2w} \int_{v=0}^1 \int_{x=0}^w (v+\beta(v))dv. \quad (2)$$

There is a pointwise smallest solution to this functional inequality, which is attained by a winning bid function $\beta(v) = v/3$.³ This winning bid function induces a distribution of winning bids \underline{H} that is uniform on $[0, 1/3]$, which is the minimum winning bid distribution across all information structures and strategy profiles that satisfy the uniform upward incentive constraints. Since any equilibrium strategy profile must also satisfy these constraints, \underline{H} must be below any equilibrium winning bid distribution. But since this is precisely the winning bid distribution that obtains with the maximum-of-independent-signals information structure, we know that this is also the minimum winning bid distribution across all information structures and equilibria.

Figure 1 shows the possible combinations of bidder surplus (on the x -axis) and revenue (on the y -axis). As the total surplus is $1/2$, and as all equilibrium allocations are efficient in the pure common value environment, the different allocations correspond to the -45 degree line on the right of the picture.

²Suppose that $H(b|v)$ satisfies the uniform upward incentive constraints and induces a distribution $H(b)$. Then we can define $\beta(v) = \min\{b|H(b) \geq v\}$ to be the bid with the same percentile as the value, and let $\tilde{H}(b|v)$ be the conditional distribution that puts probability one on $\beta(v)$. It is readily verified that $\tilde{H}(b|v)$ also deters uniform upward deviations.

³Consider the operator Λ that maps a monotonic winning bid function β to the right-hand side of (2). Then the uniform upward incentive constraint is equivalent to $\beta \geq \Lambda\beta$. Λ is also monotonic, so that iteratively applying Λ to any β that satisfies this inequality generates a decreasing sequence of functions. Finally, Λ is a contraction of modulus $1/2$ (in the sup norm) so that there is a unique fixed point of Λ which must be below any β that satisfies (2). It is easily verified that $\beta(v) = v/3$ is this fixed point.




Fig. 1. Bidder Surplus and Revenue Across Information Structures

Note that at this solution, all of the uniform upward incentive constraints bind, which implies that in equilibrium, bidders must be indifferent between their equilibrium bids and all higher bids. Indeed, this was the case in our construction of an information structure and equilibrium that obtains the bounds. These binding constraints are illustrated in Figure 2, which depicts indirect utility of a bidder who mimics the bid of a bidder with type s' while his true type is s . The mimicked type is on the x -axis, and we plot the indirect utility for three types, $s = 1/4, 1/2, 3/4$. The flat segments to the right of the equilibrium bid (marked the short vertical line) indicate the bidders' upward indifference.

4. MAIN RESULT

In [Bergemann et al. 2017a], we generalize the preceding analysis to provide a lower bound on the winning bid distribution for *any* distribution of values. In particular, this includes value distributions in which the bidders have different values. When the distribution is symmetric, in the sense of exchangeability, then the lower bound is tight and is in fact the minimum winning bid distribution.

We now give a statement of the general result. The bidders only learn the *realized average* of the $N - 1$ lowest valuations:

$$\alpha(v) = \frac{1}{N-1} \left(\sum_{i=1}^N v_i - \max v \right). \quad (3)$$

Let Q denote the distribution of $\alpha(v)$, and write $[\underline{w}, \bar{w}]$ for the convex hull of the support of Q . The bidders receive signals that are independent draws from a distribution $F(s) = (Q(s))^{1/N}$ on the support $S = [\underline{w}, \bar{w}]$. This distribution is

chosen so that the highest signal is distributed according to Q . Indeed, signals will be correlated with values so that:

- (i) the highest signal is equal to the realized average losing value;
- (ii) the bidder with the highest value receives the highest signal.

The associated bidding function is then given by:

$$\underline{\beta}(w) = \frac{1}{Q^{\frac{N-1}{N}}(w)} \int_{x=w}^w x \frac{N-1}{N} \frac{Q(dx)}{Q^{\frac{1}{N}}(x)} \quad (4)$$

which is the expectation of the highest of $N-1$ draws from $(Q(s))^{1/N}$ conditional on it being less than w . The associated winning bid distribution is

$$\underline{H}(b) = Q(\underline{\beta}^{-1}(b)). \quad (5)$$

For a given information structure \mathcal{S} and equilibrium σ , the winning-bid distribution is defined by:

$$H(b; \mathcal{S}, \sigma) = \int_{v \in V^N} \int_{s \in \mathcal{S}} \sigma([0, b]^N | s) \pi(ds|v) \mu(dv), \quad (6)$$

where $\sigma([0, b]^N | s)$ is the conditional probability that all bids are less than b given signal profile s . Our main result is the following:

THEOREM 4.1 MINIMUM WINNING BIDS.

- (i) For any information structure \mathcal{S} and equilibrium σ , the distribution of winning bids $H(\mathcal{S}, \sigma)$ first-order stochastically dominates \underline{H} , i.e., $H(b; \mathcal{S}, \sigma) \leq \underline{H}(b)$ for all b ;
- (ii) There exists an information structure \mathcal{S} and an efficient equilibrium σ such that the distribution of winning bids $H(\mathcal{S}, \sigma)$ is exactly equal to \underline{H} .

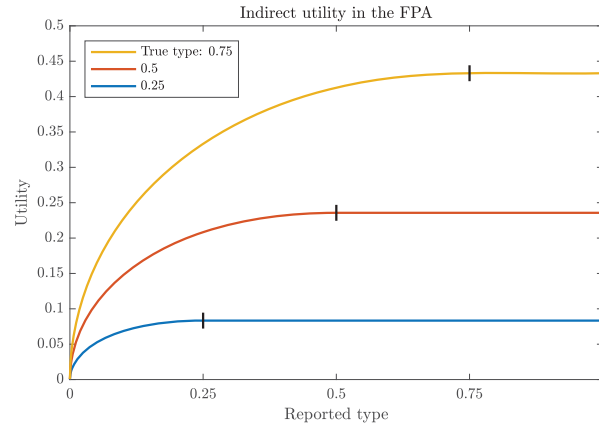


Fig. 2. Uniform Upward Incentive Constraints




Fig. 3. The set of revenue-bidder surplus pairs with independent private values.

With a pure common value model, the realized surplus in the first-price auction is independent of the allocation of the object across the bidder. We now illustrate our results with independent values drawn from the uniform distribution and two bidders. By contrast, in the independent private value environment, the surplus that is generated by the auction now depends on the allocation across the bidders. Figure 3 illustrates results for this example. As the maximum total surplus is $2/3$, the efficient allocations correspond to the -45 degree line on the right of the picture. The worst case for efficiency would be that the object is always sold to the bidder with the lowest value, which would generate a total surplus of $1/3$. Thus, the green trapezoid represents the surplus pairs that satisfy this range restriction on total surplus and also give non-negative surplus to both the bidders and the seller.

We can again consider the whole range of revenue and bidder surplus across all possible information structures and equilibria. This includes, in particular, maximum revenue, minimum bidder surplus, and minimum total surplus. In the analysis thus far, we have considered all information structures, including those in which the bidders' signals do not reveal their own values exactly. We refer to this model as one of *unknown values*, to distinguish it from the case that we consider next. The set of surplus pairs that can arise in the unknown-values model is the area enclosed by the blue curve in Figure 3. Point A corresponds to minimum revenue/maximum bidder surplus characterized by our main theorem. Point B corresponds to maximum revenue/minimum bidder surplus, in which the bidders obtain zero surplus but the allocation is efficient, so that the seller obtains all of the efficient surplus. In the information structure that attains this point, the bidders have relatively precise information about the highest value, but imprecise information about who has the highest value. This induces very aggressive bidding so that the bidders compete away all their rents, although the information is precise enough to facilitate an efficient allocation in which the high-value bidder always wins the good. The point C corresponds to minimum total surplus, which, remarkably, attains the lower bound of $1/3$. This is attained in an information structure and

equilibrium in which the good is always won by the bidder who values it the *least*. The information and equilibrium are easy to describe: each bidder observes a signal equal to the *other* bidder’s value, and then bids their signal.

In the *known-values* model, we assume that the bidders at least know their own values, and their signals may contain additional information about others’ signals. The set of welfare outcomes that can arise under known-values information structures and equilibria is enclosed by the red curve in Figure 3.⁴ In contrast to unknown values, known values implies that each bidder can guarantee himself a strictly positive surplus, so that maximum revenue attained at point D is strictly less than the expected highest value, which is $2/3$. We fully characterize this outcome using methods adapted from [Bergemann et al. 2015], which studies the welfare impact of information in third-degree price discrimination. Point E corresponds to minimum revenue under known values, for which we have an analytical characterization only when values are binary.

The known-values surplus is significantly smaller than the unknown-values surplus set. It is notable that the inefficiencies that can arise in the known-values model are relatively small compared to what can happen with unknown values. This observation is in line with the results of [Syrngkanis and Tardos 2013] and [Syrngkanis 2014] who show that the efficiency loss in the independent private-value auction expressed in terms of the ratio between realized surplus and efficient surplus in the first-price auction is bounded below by $1 - 1/e$.

5. DISCUSSION

A feature of our analysis is that we characterize bidding behavior in all equilibria for all information structures at once. [Bergemann and Morris 2013; 2016] show that the range of such behavior can be described using a certain incomplete-information correlated equilibrium that they term *Bayes correlated equilibrium*. Thus, a contribution of ours is to characterize the Bayes correlated equilibria of the first-price auction.

In the pure-common-value model, the winner’s value and all of the losing bidders’ values are exactly the same, so that the winning-bid-minimizing information structure can be simply described as the bidders having independent and real-valued signals and the value being the maximum of the signals. [Bulow and Klemperer 2002] studied a second-price auction where bidders’ information is of this form and showed that there is an equilibrium in which bidders bid their signals. They showed that the bidder with the highest signal has the lowest marginal revenue as described by the standard virtual value, thus hinting at the low revenue properties of this information structure. In fact, the winner’s curse effect is so strong that revenue would be higher if the seller simply offered the good at the highest posted price such that all bidders are willing to accept.

⁴The fact that the known-values set (in red) is contained within the unknown-values set (in blue) is a reflection of the general observation that providing the bidders with more information decreases the set of outcomes that can be rationalized as an equilibrium with even more information. In other words, the set of Bayes correlated equilibria is decreasing in the minimum information of the players. [Bergemann and Morris 2016] formalize the notion of “more information” and give a precise statement of this result.

In [Bergemann et al. 2017b], we characterize revenue-maximizing auctions for this information structure with pure common values. When the good must be sold, the aforementioned posted price is indeed the optimal mechanism, but more generally the optimal mechanism has a novel form that we refer to as the priority auction. This mechanism biases the allocation towards lower-signal bidders conditional on the good being allocated. In this context, we show that the revenue ranking of [Bulow and Klemperer 1996] between optimal auction with N bidders and standard auction with $N + 1$ bidders is reversed in favor of the optimal auction, even when we allow the standard auction to have $N + K$ bidders for any $K > 0$.

Our approach can be used to compare the set of possible welfare outcomes across mechanisms. We have shown that the first-price auction (even without a reserve price) is guaranteed to generate positive revenue, regardless of the information structure and equilibrium. This is not true of the second-price auction. Even when buyers know their own values, the second-price auction admits weakly-dominated equilibria in which revenue is zero. In [Bergemann et al. 2018], we show more broadly that the first-price auction must have a weakly greater revenue guarantee than any other mechanism that is revenue equivalent to the first-price auction when the environment is one of symmetric and independent private values, which includes second-price auctions, all-pay auctions, and any combinations thereof. When we compute a restricted revenue guarantee across just symmetric affiliated environments and monotonic pure-strategy equilibria, first-price, second-price, and English auctions all become revenue-guarantee equivalent. Finally, in [Bergemann et al. 2016] we identify mechanisms that provide the optimal revenue guarantee when there are two bidders and binary common values, and [Brooks and Du 2018] characterize revenue-guarantee maximizing auctions in more general common value settings.

REFERENCES

- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957.
- (2016): “Informationally Robust Auction Design,” Tech. rep., Cowles Foundation Discussion Paper 2065.
- (2017a): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, 85, 107–143.
- (2017b): “Selling to Intermediaries: Optimal Auction Design in a Common Value Model,” Tech. rep., Yale University, University of Chicago and Princeton University.
- (2018): “Revenue Guarantee Equivalence,” Tech. rep., Yale University, University of Chicago and Princeton University.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BROOKS, B. AND S. DU (2018): “Optimal Auction Design with Common Values: An Informationally Robust Approach,” Tech. rep., University of Chicago and Simon Fraser University.
- BULOW, J. AND P. KLEMPERER (1996): “Auctions vs Negotiations,” *American Economic Review*, 86, 180–194.
- (2002): “Prices and the Winner’s Curse,” *RAND Journal of Economics*, 33, 1–21.
- ENGELBRECHT-WIGGANS, R., P. MILGROM, AND R. WEBER (1983): “Competitive Bidding and Proprietary Information,” *Journal of Mathematical Economics*, 11, 161–169.

- KRISHNA, V. (2002): *Auction Theory*, San Diego: Academic Press.
- SYRGKANIS, V. (2014): “Efficiency of Mechanisms in Complex Markets,” Ph.D. thesis, Cornell University.
- SYRGKANIS, V. AND É. TARDOS (2013): “Composable and Efficient Mechanisms,” in *STOC '13 Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, 211–220.