école d'économie d'aix-marseille
aix-marseille school of economics

Working Papers / Documents de travail

I Want to Tell You? Maximizing Revenue in First-Price Two-Stage Auctions

Galit Ashkenazi-Golan
Yevgeny Tsodikovich
Yannick Viossat

WP 2020 - Nr 44

I Want to Tell You?

Maximizing Revenue in First-Price Two-Stage Auctions *

Galit Ashkenazi-Golan[†], Yevgeny Tsodikovich[‡], and Yannick Viossat[§]

October 14, 2020

Abstract

A common practice in many auctions is to offer bidders an opportunity to improve their bids, known as a Best and Final Offer (BAFO) stage. This final bid can depend on new information provided about either the asset or the competitors. This paper examines the effects of new information regarding competitors, seeking to determine what information the auctioneer should provide assuming the set of allowable bids is discrete. The rational strategy profile that maximizes the revenue of the auctioneer is the one where each bidder makes the highest possible bid that is lower than his valuation of the item. This strategy profile is an equilibrium for a large enough number of bidders, regardless of the information released. We compare the number of bidders needed for this profile to be an equilibrium under different information settings. We find that it becomes an equilibrium with fewer bidders when no additional information is made available to the bidders compared to when information regarding the competition is available. As a result, from the auctioneer's revenue perspective, when the number of bidders is unknown, there are some advantages to not revealing information between the stages of the auction.

Journal classification number: D44, D82.

Keywords: Auctions, multistage auctions, BAFO, information utilization.

*The authors wish to thank David Lagziel for his valuable comments. G. Ashkenazi-Golan acknowledge the support of the Israel Science Foundation, grants #217/17 and #722/18, and NSFC-ISF, China Grant #2510/17. Y. Tsodikovich acknowledges the support of the French National Research Agency Grants ANR-17-EURE-0020. The research of G. Ashkenazi-Golan and Y. Tsodikovich was partially supported by the COST Action CA16228 "European Network for Game Theory" (GAMENET).

†The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997800, Israel. e-mail: galit.ashkenazi@gmail.com .

‡Aix Marseille Univ, CNRS, EHESS, Centrale Marseille, AMSE, Marseille, France. e-mail: yevgets@gmail.com

§PSL, Université Paris–Dauphine, CEREMADE UMR7534, Paris, France

1 Introduction

This paper examines information effects in two-stage auctions involving a large number of bidders. It is intuitively clear and theoretically understood that, under rather general assumptions, when large numbers of bidders participate in an auction, the strong competition drives the revenue of the auctioneer up. Wilson [1977], followed by Milgrom [1979], showed that in a first-price sealed-bid auction with affiliated values, as the number of bidders increases, the winning bid converges to the value the winner attributes to winning the auction. This mechanism is not only efficient, but leaves the entire economic surplus in the hands of the auctioneer.[1]

The importance of the number of bidders to the revenue of the auctioneer was highlighted by Bulow and Klemperer [1994], who showed that the additional revenue provided by one extra bidder could exceed the revenue from choosing the optimal allocation mechanism. Another reason for analyzing auctions with many bidders was stated in Wilson [1985] and Neeman [2003]: it allows for lighter assumptions regarding bidders' knowledge of features like the exact number of bidders or the distribution of their valuations. Often, bidders only need to know that there is a large number of participants in the auction. Furthermore, in many settings the exact equilibrium strategies are hard to compute, whereas the asymptotic strategies are more straightforward, as illustrated by Bali and Jackson [2002] among others.

It is common to have two stages in procurement auctions, especially those relating to construction contracts. In the first stage, the bidders submit sealed bids that are evaluated, the highest being chosen. The corresponding bidders move on to the second stage, known as the Best And Final Offer (BAFO) stage, where they are allowed to increase their bids. The winner in the auction is chosen based on the BAFO-stage bids, and typically pays his bid (first-price auction).

This paper was inspired by a consulting job. An auctioneer[2] was seeking advice on what information to provide to the bidders before the BAFO stage in order to maximize her revenue. In general, such information can help bidders learn either their private values or the level of competition. While an extensive literature has examined the former information effect (e.g., Pesendorfer and Swinkels [2000], Swinkels [2001]), little is known about the latter. Our paper fills this gap in the literature by considering which information, if any, should be given to bidders between stages to maximize revenue. To isolate the problem of how to stimulate competition via information, we chose a model of private valuations. Naturally, if there are common unknown components to the valuation, additional information teaches each bidder not only about the competition but also about his private value, which can significantly affect future bids. Hence, in our model, information about the other bids and bidders does not affect valuations. It only affects beliefs regarding the second-stage bids and how competitive the rest of the bidders might be.

[1]More about efficiency in large auctions can be found in Pesendorfer and Swinkels [2000], Swinkels [2001], Bali and Jackson [2002], Fibich and Gavious [2010] and the references therein.

[2]Hereafter, as customary in auctions, the auctioneer will be referred to as a female "seller" and the bidders as male "buyers", although in procurement auctions the roles are reversed.

While the literature suggests that a second-price sealed-bid auction is optimal (in terms of revenue for the auctioneer) for our setting [Myerson, 1981], this type of auction is not widely used. Two-stage auctions, however, are common. The optimal two-stage auction design when the values might be correlated was explored in Perry et al. [2000], that considered a model where the second stage is a second-price sealed-bid auction. The optimal information to be released by the auctioneer in a two-stage auction where the second stage is a second-price sealed-bid auction was investigated by Ganuza [2007]. Our research is restricted to the more common auction structure, where the second stage is a first-price sealed-bid auction. As far as we know, this model, particularly the optimal information structure, has not previously been explored.

In our model, n bidders with private values participate in a sealed-bid two-stage auction with discrete[3] bids. The top two bidders from the first stage proceed to the second stage (the BAFO stage), where they can either increase their bid or maintain the previous one. The highest bidder in the second stage wins the auction, pays this bid, and receives the good. Between the stages, the seller can privately send a message to each of the finalists regarding others' bids, their ranking, or other relevant information.

This information can trigger two opposing effects. On the one hand, information about competitors can drive prices up, as bidders seek to increase their chances of victory. On the other hand, the possibility of future information can reduce bids in the first round. A bidder may choose to submit a lower bid in the first stage and increase it only if he finds himself facing a strong competitor. Otherwise, the smaller initial bid becomes the final one, decreasing the revenue of the seller.

The main goal and contribution of this paper are to provide a theoretical answer to a practical question, and to deduce the conditions under which each of the above-mentioned effects is stronger. More precisely, we seek to determine when the strategy profile that maximizes the revenue of the seller is an equilibrium, and when it is the unique equilibrium for different information structures.

We show that the revenue-maximizing strategy profile is an equilibrium for a large enough number of bidders, regardless of the information released between the stages. We compare the number of bidders required for this strategy profile to be an equilibrium across different information structures. We find that without meaningful information, fewer bidders are required to maintain this strategy profile as an equilibrium. As a result, when the number of bidders is unknown, there are some advantages to conducting an auction without revealing information before the BAFO stage.

We postulate not only that it takes fewer bidders to obtain the revenue-maximizing equilibrium in auctions where information is withheld than in auctions with information, but also that it remains the unique equilibrium under fewer bidders when no information is provided. This strengthens the claim that providing additional information between stages is not to the seller's advantage, as it gives bidders an opportunity to play an equilibrium where they bid less. As is frequently the case (see, for example, Quint and Hendricks [2018]), the uniqueness

[3] We discuss the continuous case in Appendix C.

of this equilibrium can only be proven in special cases. We provide such a proof in Appendix A for a simple scenario and test this hypothesis for more complicated scenarios using a computer simulation.

Our analysis differs from common practice in the auction literature. Instead of finding equilibrium strategies in a given auction or finding the optimal mechanism for specific settings, we solve an implementation problem. In our work, a particular strategy profile, the revenue-maximizing profile, is desired and we construct an information structure in which this profile is an equilibrium in a two-stage auction with discrete bids and many bidders. For a small number of bidders, multiple equilibria can exist, including equilibria with mixed actions or relatively low bids. The example in Section 3.1 shows that the optimal information structure in an auction with a small number of bidders depends on the fine details of the value distribution.

The implication of our results is that information provided before the BAFO stage might decrease the revenue of the seller. Hence, when the sole purpose of the BAFO stage is to encourage competition, providing no information tends to be better for the seller. It also renders the BAFO meaningless. Since the bidders learn nothing between stages, the optimal strategy is to bid the same amount in both stages. In such situations, there is an argument for avoiding the BAFO stage altogether, which can reduce costs and prevent delays [Ahadzi and Bowles, 2001, Dudkin and Välilä, 2006].

The remainder of the paper is organized as follows. This introduction is followed by a short survey of related literature and findings regarding multi-stage auctions. Section 2 formally presents the model and the various information structures. Section 3 demonstrates the model through two examples and sets the ground for the main results, which are presented in Section 4. Section 5 includes the concluding remarks and a discussion of the remaining open questions. A partial answer regarding the uniqueness of the equilibrium is given in Appendix A. Finally, in Appendix B we present a more general but less intuitive assumption regarding the joint distribution of the private values that remains consistent with our results and in Appendix C we discuss auctions with a continuum of private values and possible bids. We show that in the continuous case, as in the discrete, there are good reasons for not revealing information.

1.1 Motivation and Related Literature

While commonly used, two-stage auctions with a BAFO stage have received little research attention. For example, the Federal Transit Administration recommends a BAFO stage in its manual[4] as a means to conclude the auction and receive improved offers. A BAFO stage is also commonly used in the UK and France as the final stage of procurement auctions [Noumba and Dinghem, 2005]. It appears in large-scale infrastructure projects around the world, such as the water distillation factory in Ashkelon, Israel [Sauvet-Goichon, 2007] and in prison procurement and operations in South Africa [Merrifield et al., 2002]. It was even used as a tie-breaking stage when the auctioneer could not decide between two offers in terms of highest value for

[4]Best Practices Procurement Manual, Federal Transit Administration, November 2001.

money, and decided to implement an additional (unplanned) BAFO stage [Rintala et al., 2008]. Furthermore, it is the US government's preferred method of buying goods, as suggested in Chelekis [1992] and illustrated in Roll [2000], Finley [2001] and others.

The World Bank, in a procurement guide entitled "Negotiation and Best and Final Offer (BAFO)"[5] states that:

> "BAFO is appropriate when the procurement process may benefit from Bidders/Proposers having a final opportunity to improve their Bid/Proposal, including by reducing prices, clarifying or modifying their Bid/Proposal, or providing additional information. It is normally particularly effective when markets are known to be highly competitive and there is strong competitive tension between Bidders/Proposers."

Then the guide lists the objectives of the BAFO – the two main aims being to increase understanding by bidders of the auctioneer's requirements and to enhance competition among bidders who have made a proposal. The results presented here mainly relate to the second aim, addressing the following question: does the BAFO enhance competition? We also investigate its effectiveness when markets are highly competitive, as suggested by the World Bank's guide.

Another possible reason to employ two-stage auctions is for screening purposes: the auctioneer screens out bidders that do not meet some minimum quality or price standards, and in the second stage, she runs an English auction among the remaining bidders. Thus, the purpose of the first stage is to ensure that the good provided meets the minimum requirements, and the only remaining issue, the price, is settled in the second stage.

Screening can also serve the purpose of limiting competition to encourage bidders to participate in an auction with entry costs. For example, in indicative bidding [Ye, 2007, Quint and Hendricks, 2018] the seller asks bidders to submit non-binding first-stage bids to evaluate their interest in the good. She then chooses the best bidders to proceed to the second stage, in which they prepare their "real" bid. Since preparing the second-stage bid is costly (requires time, effort, learning about the asset, and so on), it is important to reassure the bidders that their chances of winning are high enough to justify the entry cost. This cost also discourages bidders from bidding high just to qualify, as they are not sure that they want to participate in the second stage.

Multiple stages can also help bidders learn their rivals private value. This happens when the bidder does not know for certain his own valuation (but rather has an estimation of this value given his signal) and the values are correlated (as in Milgrom and Weber [1982]). In this case, the information about bids of others changes the bidder's own evaluation of the gain in winning the auction. In such settings, a sealed-bid first-price auction is significantly different from an English ascending auction. In the former, bids are placed according to initial evaluations of the values, there is no updating, and the chances of overbidding and suffering from the "winner's curse" are higher. In contrast, the constant updating of a bidder's estimation of the value

[5] http://pubdocs.worldbank.org/en/663621519334519385/Procurement-Guidance-Negotiation-and-Best-Final-Offer.pdf

during an English auction, based on the current price and the number of buyers that quit at each level, can reduce bids and prevent such problems [Perry et al., 2000].

Since we are studying the effect of information on bidding competition, we assume that private values are known from the outset of the auction. Bids are therefore updated in the second round for the sole purpose of increasing the chances of winning. This can occur when the bidder learns the other bids and realizes that he is about to lose. However, in practice, it also occurs without any additional information. For example, the following question appears in the FAQ section of the website of the US Federal Transit Administration[6]:

> Q: Is it permissible for an offeror to lower its price for the Best and Final Offer (BAFO) without any basis for the change, other than trying to beat out the competition, or does the price reduction have to be based on changes or other clarifications discussed during the presentation? We are participating as an offeror and have been asked for a BAFO.
>
> A: Yes, the basis for change may be solely the desire to increase your chances of winning the contract award by lowering your price. If the contract you are competing for is a cost-reimbursement type contract, the procuring agency may well ask you for your rationale in lowering the original cost estimates to do the work. [...]

In Lemma 1 we show that this kind of bidding is sub-optimal. If it is profitable for the bidder to submit a better offer in the second stage without additional information, he might just as well submit it in the first stage and increase his chances of winning, as nothing changes between the stages.

2 The Model

We consider a two-stage auction with private valuations. In the first stage, all bidders submit their bids. The two highest bidders proceed to the second stage, called Best-And-Final-Offer (BAFO) stage. The first-stage bids are binding in the sense that each participant in the BAFO stage must bid at least as high as his first-stage bid. The highest bidder at the BAFO stage wins the auction and pays his second-stage bid. In both stages, we use a symmetric tie-breaking rule.

Formally, let $N = \{1, \ldots, n\}$ be the set of $n \geqslant 3$ bidders, $V = \left\{v^1, \ldots, v^K\right\}$ the set of possible private valuations (types) in ascending order and F_n the joint distribution assigning types to bidders. In Assumption 1 we explain how F_n changes with n. This assumption is slightly generalized in Appendix B.

After observing his own valuation, bidder i places his first-stage bid, b_i^1, chosen from a given set of allowable bids $B = \left\{b^1, \ldots, b^M\right\}$. Presetting a set of allowable bids is a common practice, as the auctioneer prefers bids to be "rounded" or in fixed increments; moreover, this is a natural

[6] https://www.transit.dot.gov/funding/procurement/third-party-procurement/best-and-final-offer

limitation as the currency is not continuous. Appendix C provides several results when the sets of bids and private valuations are continuous. We assume that for any $v^j \in B$, there is an allowable bid no higher than v_j, and if $j > 1$, a bid strictly between v^{j-1} and v^j. The quantities $\beta^j = \max\{b \in B | b < v^j\}$ are thus well defined, and for $j > 1$, $\beta^j > v^{j-1}$.[7]

The two[8] bidders who bid highest in the first stage move on to the BAFO stage. These two bidders receive information from the auctioneer regarding the results of the first stage before the second stage starts. The information is a function from the set of first-stage bids to a set of vectors of messages (one for each bidder). Let Θ_i be the set of possible messages bidder i may receive. The information structure is $\Theta : B^n \to \Delta(\Theta_1 \times \ldots \times \Theta_n)$. The realized vector of messages is denoted by $(\theta_1, \ldots, \theta_n) \in \Theta_1 \times \ldots \times \Theta_n$.

In the second stage, the two bidders place a second bid (once again, within B), the only restriction being that their second bid must be at least as high as their first bid. Hence, a (mixed) strategy of bidder i is a function $\sigma_i = (\sigma_i^1, \sigma_i^2)$ such that $\sigma_i^1 : V \to \Delta(B)$ and $\sigma_i^2 : V \times B \times \Theta_i \to \Delta(B)$. The strategy profile of all the bidders is denoted by $\Sigma = \bigtimes_{i \in N} \sigma_i$.

We wish to compare information structures. Specifically, we want to compare an *informative* information structure to a *non-informative* one.[9] In the context of this paper, the bidders may use their information when deciding on their second-stage bids. In equilibrium, a bidder optimizes his gain at the second stage given the joint distribution of values, his value, his first-stage bid, his opponents' strategies, and the signal he observes between the stages. Given these, the bidder rationally forms his beliefs regarding his opponent's bid at the second stage based on Bayes' Rule, and then he best responds to these beliefs. An information structure is non-informative if all signals conveyed with positive probability do not affect these beliefs, whatever the strategies of the players.

Definition 1. *Let B_j be the random variable that represents the second-stage bid of bidder j, given F_n and Σ. An information structure Θ is* non-informative *if for every bidder $i \in N$ who proceeded to the BAFO stage with a private value $v_i \in V$, first-stage bid $b_i^1 \in B$ while the others use the strategy profile Σ_{-i}:*

$$\forall F_n, \forall \Sigma_{-i}, \forall \theta_i \in \Theta_i, \forall j \neq i, b \in B : \Pr(B_j = b | F_n, v_i, b_i^1, \Sigma_{-i}, \theta_i) = \Pr(B_j = b | F_n, v_i, b_i^1, \Sigma_{-i}), \tag{1}$$

where the $\forall \theta_i \in \Theta_i$ part refers only to messages with non-zero probability given the first-stage bids. Otherwise, the information structure is informative.

A non-informative information structure is such that the message from the auctioneer does not change the belief of bidder i regarding the distribution of the second stage bids of the other

[7]In fact, if we remove the richness of allowable bids assumption, the only thing lost is the efficiency of the auction – the item might no longer go to the bidder who values it highest, since the bids sometimes cannot differentiate between bidders with different private values.

[8]Our results can be generalized to cases where more than two bidders proceed to the BAFO stage.

[9]Often referred to as "with information" and "without information", respectively.

finalist. The simplest non-informative information structure is that where for all i, Θ_i is a singleton and the same message is conveyed to the bidders regardless of the actual bids.

In an informative information structure, there is a *possibility* that the information changes this belief. Intuitively, the "most informative" information structure in this sense is to reveal all the first-stage bids to the finalists. Given the bid of bidder j, bidder i can update his belief about j's private value and, combined with Σ_{-i}, his belief regarding j's second-stage bid.

The highest bidder at the second stage wins the auction (ties are broken randomly with equal probabilities). The utility for the winner is the difference between his valuation and the amount he pays – his second-stage bid.

What we are examining here is the effect of the number of bidders and the information structure on equilibrium strategies and equilibrium payoff. We note that in general, the solution concept we study is symmetric equilibrium, i.e. an equilibrium where all bidders with the same private value use the same strategy. Hence, the strategies are well-defined and self-consistent when the number of bidders changes. Thus, we also should specify how the joint distribution of valuations changes when the number of bidders change, that is, how F_n changes with n. Let F_n be some distribution and for every $k \in \{1, \ldots, K\}$, we denote by D_k the event such that the types of all bidders are at most v^k. Given D_k, and given that the valuation of a specific bidder i is v^k, we denote by $\delta^k(F_n)$ the minimal probability of each other bidder having the type v^k, independently of all others' types and $\delta(F_n) := \min_k \delta^k(F_n)$.

Here, *independently* does not necessarily mean that the bidders have independent valuations. Instead, the idea is that given D_k and given the valuations of every possible subset of bidders that does not include bidder j, the latter has a probability of at least $\delta(F_n)$ of having the type v^k. Valuations can be correlated, but not *too correlated* in the sense that type v^k cannot be ruled out for bidder j, regardless of the types of all the other bidders. For example, when the valuations are i.i.d. (with full support over V), $\delta^k(F_n)$ is simply $\frac{p^k}{\sum_{j \leqslant k} p^j}$, where p^j is the probability of each bidder having type v^j, and $\delta(F_n) = \min_k \frac{p^k}{\sum_{j \leqslant k} p^j}$. Typically, when the number of bidder is large or when the bids are increasing with the valuation, in equilibrium, a bidder with high private value always outbids a bidder with low private value, so a bidder with a private value v^k (for $k < K$) has a positive probability to win only conditioning on D_k. We formalize this reasoning in the Results section.

We assume that as n increases, these lower bounds become strictly positive and independent of n:

Assumption 1. *The family of distributions F_n have* uniform full support*, i.e.*

$$\delta := \liminf_{n \to \infty} \delta(F_n) > 0. \tag{2}$$

The intuition behind the uniform full support assumption can be taken from the following pathological example. Assume F_n is such that it assigns to one of the bidders the highest type and to all the others the lowest type. The high-type bidder knows that he is the only one, and can bid a very low bid that is still higher than the value of all the other bidders. This remains true regardless of n or the design of the auction. To ensure that this high-type bidder feels a competitive pressure to increase his bid as the number of bidders increases, the expected number of other high-type bidders should not go to 0 with n. This is the case under the uniform full support assumption, since the expected number of type k bidders is at least $n\delta$.

Note that the condition on F_n can be slightly weaker (but also less elegant). These conditions can be found in Appendix B. This concludes the description of the model at the bidders' part.

As for the seller, she utilizes her role as a mechanism designer that chooses an information structure to maximize her expected payoff in the resulting auctions. Since we assume that the bidders adopt equilibrium strategies in the auction, she wishes each bidder to bid his highest possible bid which is not strictly dominated, i.e. the corresponding β^i to him. Such a strategy will be referred to as the *revenue-maximizing strategy* and the strategy profile in which all bidders use the revenue maximizing strategies is called the *revenue-maximizing profile*. [10] Denote by σ^* the strategy in which a bidder with valuation v^j bids in both stages β^j and by Σ^* the strategy profile in which all bidders use the revenue-maximizing strategies. The main questions that we explore in this paper is when Σ^* is indeed an equilibrium and when it is the unique one.

3 Examples

We present two simple examples to illustrate the model and which provide key insights into our results. Specifically, the examples demonstrate how the revenue-maximizing information structure depends on the number of bidders.

The first example (Section 3.1) demonstrates that when different information structures are compared, the results vary with the number of bidders. We show that with a small number of bidders, the optimal information structure for the auctioneer depends on the exact distribution of valuations. In particular, an informative information structure can be better. However, with a larger number of bidders the non-informative information structure becomes optimal. The next question we address is how large the number of bidders needs to be to ensure that the non-informative structure is optimal. The second example (Section 3.2) suggests that this number is not necessarily very large. Moreover, we show that for a large enough number of bidders, both information structures are equivalent: under both structures, the only equilibrium is the revenue-maximizing strategy profile. All the results are formally presented and proved in Section 4.

[10] There might be other strategy profiles with the same expected payoff for the seller, but they all are weakly dominated by this one.

3.1 Small Number of Bidders

Suppose there are $n = 3$ bidders, two valuations $\{v^L = 0.01, v^H = 1\}$, and four possible bids $B = \{0, 0.33, 0.66, 0.99\}$. The probability of each bidder having a high valuation is $p = 0.01$, regardless of the other bidders. Clearly, a bidder with a low valuation bids $b^1 = 0$ in all equilibria.

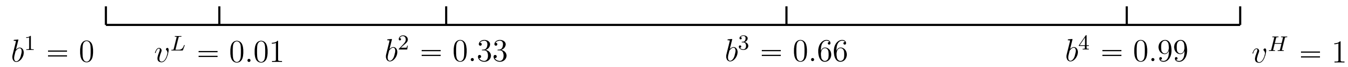

We compare the two extreme information structures: the non-informative information structure and an information structure where the bids of the bidders in the first stage are revealed. Since the last structure reveals all the information available to the auctioneer, we refer to this structure as *the fully-informative* information structure. Claim 1 shows that the non-informative structure is not revenue-maximizing: there is no symmetric equilibrium where a bidder with a high valuation bids higher than $b^2 = 0.33$ with a positive probability. The unique symmetric equilibrium, in this case, is the one where a high valuation bidder bids $b^2 = 0.33$ in both stages.

Claim 1. *In a non-informative information structure, the only equilibrium is that where a high valuation bidder submits the bid b^2 in both stages.*

Proof. Following Lemma 1, we may assume, without loss of generality, that the bids are the same in both stages. Consider a bidder with a high valuation. When bidding b^4, his profit cannot exceed 0.01; when bidding b^3, his profit cannot exceed 0.34; when bidding b^1, his profit cannot exceed $\frac{1}{3}$. However, when bidding $b^2 = 0.33$, both opponents have low valuations with probability 0.99^2, so the profit from bidding b^2 is at least $0.99^2 \cdot 0.67$, which is much higher. Hence, bidding b^2 is the dominant strategy and the only equilibrium is bidding b^2. ■

In the full information case, all equilibria are such that the high type never bids b_1 (because this would imply a gain less than $\frac{1}{3}$ in that case), and bids 0.66 or more with positive probability in the second round (if the second round bid was always b^2, then deviating and bidding b^3 in the second round when the other bidder has the valuation v^H is a profitable deviation). Hence, any equilibrium of the full information model (there are several) is better for the seller than the best equilibrium without information. Note also that the best equilibrium for the seller in the full information case is the one where, when there are at least two bidders with high type, they both bid b^4 in the second round.

In this example, revealing information benefits the seller. If there are only low valuation bidders, the information structures are equivalent. However, given there is at least one high-type bidder, the profit of the seller in equilibrium without information is 0.33, while the profit of the seller in the fully informative structure is higher. For example, if in equilibrium a high-typed bidder bids 0.33 and increases his bid to 0.66 only if there is another high-typed bidder, the payoff to the seller would be $0.33 \cdot 0.99^2 + 0.66(1 - 0.99^2) > 0.33$.

Note that changes to the joint distribution of the values affect this observation. For example, if the dependency is such that a bidder with a high valuation induces a very high probability of an opponent having a high valuation as well, the information structures become equivalent.

Direct computation reveals that the reasoning in the full-information model remains true as long as $n \leqslant 184$ (for $n > 184$ the two information structures are revenue-equivalent), while Claim 1 remains true only as long as $n \leqslant 155$. When the number of bidders is between 155 and 184, the only equilibrium without information is to bid b^3 in both stages, while the equilibrium in the fully-informed model is to bid b^2 in the first stage and to raise the bid only if necessary. Hence, for the auctioneer, not releasing information becomes more profitable for a large number of bidders.[11] An auction with more than 100 bidders is atypical. However, the next example demonstrates that revealing no information may prove advantageous even with a much smaller number of bidders, depending on the other auction parameters.

3.2 Large Number of Bidders

Suppose that there are $n > 3$ bidders, two private values, $V = \{v^L = 0.3, v^H = 1\}$, and five possible equally-spaced bids $B = \{0, 0.2, \dots, 0.8\}$. Each bidder has a probability p of having the private value v^H, independent of the other bidders. In Σ^*, a bidder with private value v^L bids 0.2 and a bidder with private value v^H bids 0.8. If everyone is bidding according to Σ^*, there is no profitable deviation for a bidder with private value v^L regardless of the information structure, so we consider only a high-type bidder.

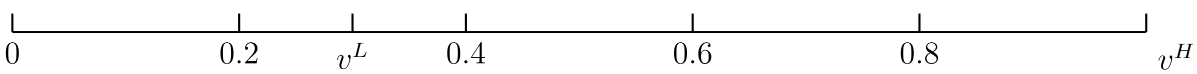

In the non-informative information structure, Σ^* is an equilibrium iff a high-type bidder has no profitable deviation. Of all possible deviations, the most profitable is where he bids the lowest bid that is above v^L, so the condition for Σ^* to be an equilibrium is

$$(1-0.8)\tfrac{1-(1-p)^n}{np} \geqslant (1-0.4)(1-p)^{n-1}. \tag{3}$$

For each p, there exists a minimal n for which this inequality holds, denoted by N^{NI}, and it remains true for every $n > N^{NI}$. This N^{NI} is the minimum number of bidders required for Σ^* to be an equilibrium.

Naturally, N^{NI} decreases with p. When p is small, a high-type bidder can bid the lower bid because there is a high probability that he is the only high-type bidder and no one will outbid him. When p is large, there is a high probability that other high-type bidders exist and bid 0.8, and a lower bid is bound to lose.

[11]For a much larger number of bidders (~ 600), the only equilibrium is the revenue-maximizing equilibrium, as Lemma 2 suggests.

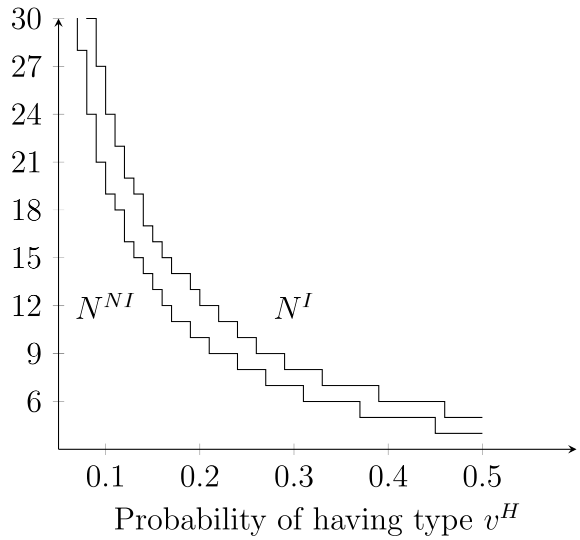

Figure 1: The minimum number of bidders for Σ^* to become an equilibrium when no information is given before the BAFO stage (N^{NI}) and when all the bids are revealed before the BAFO stage (N^I) for the auction presented in Section 3.2. For larger probabilities, both functions continue to decrease towards the minimal possible number of bidders in this model, $n = 3$.
.

A similar computation can be done for any information structure. For example, in the fully-informative structure, the most profitable deviation from Σ^* for a high-type bidder is to bid 0.4 in the first stage and raise his bid in the BAFO stage only if the other bidder outbid him. The condition for Σ^* to be an equilibrium is

$$(1-0.8)\tfrac{1-(1-p)^n}{np} \geqslant (1-0.4)(1-p)^{n-1} + \tfrac{1-0.8}{2}(n-1)p(1-p)^{n-2}. \tag{4}$$

Observe that this inequality is harder to satisfy than the former one. Again, this inequality can be solved for each p to obtain N^I, the minimum number of bidders required for Σ^* to be an equilibrium in this information structure.

The general result is proved in Lemma 2 – for n large enough, Σ^* is an equilibrium (and in fact, the only equilibrium), regardless of the information structure. The corresponding solutions of N^{NI} and N^I as a function of p are summarized in Figure 1. As Corollary 1 suggests, N^I is always larger than N^{NI}. When the probability of being a high-type bidder is not extremely small, the required number of bidders is around 10 for both types of information structures, and when the probability of being a high-type bidder is around 0.5, it is around 5.

4 Results

We start our analysis with a useful lemma which states that when the information structure is non-informative, the auction is equivalent to a first-price sealed-bid auction in the sense that the bidders, if they are behaving rationally, do not change their bids in the second stage. More precisely, every strategy that changes the bid in the second stage is dominated by one that does not. The reason is that each bidder could already have bid his second-stage bid in the first round, increasing the probability of reaching the second stage (and ultimately winning) and without affecting his payoff. This significantly simplifies the analysis of auctions without information, as was shown in Claim 1.

Lemma 1. *If Θ is non-informative, then each strategy in which the second bid differs from the first is strictly dominated by a strategy in which the second bid equals the first.*

Proof. When the information structure is non-informative, there is no advantage to making the second-stage bid conditional on that information. Hence, all decisions regarding the second-stage bid can be made before the first-stage bids are placed. Increasing the bid between stages (even with some positive probability) offers no advantage over bidding higher in the first stage (or conducting randomization before the first stage rather than between the stages).

Bidding twice b_i^2 increases the chances of proceeding to the second stage, and thus increases the chances of winning without reducing the gain from victory. Thus, with a non-informative Θ, the only strategies worth considering are those where $b_i^1 = b_i^2$. Under this information structure, the auction is equivalent to a first-price (single-stage) sealed-bid auction. ■

Our next result, which lays the path to the main theorem, is that the revenue-maximizing equilibrium is the unique equilibrium for a large enough number of bidders. This result is similar to established findings from other types of auctions and follows the same logic. As the number of bidders increases, the probability of being chosen for the second stage, and ultimately winning, decreases. Hence, it is optimal to bid as high as possible to increase the chances of being selected.

Lemma 2. *For any set of private values V, information structure Θ, and family of distributions over V which satisfy the uniform full support assumption $(F_n)_{n\in\mathbb{N}}$ (with a lower bound $\delta > 0$), there exists N_0 such that if the number of bidders is larger than N_0, then the unique symmetric equilibrium is Σ^*.*

Proof. Assume, by contradiction, that the result does not hold. Then the set of $n \in \mathbb{N}$ for which Σ^* is not the unique symmetric equilibrium is infinite. For each n in this set (actually, a sequence), there exists an equilibrium in which there exists a type v^j that bids β^j in the first stage with probability less than 1. Since the number of types is finite and the number of bids is finite, there exists a sub-sequence of these ns in which in equilibrium the highest type that does not play his respective β in the first stage with probability 1 is v^k for some k, and the lowest bid he bids in equilibrium with nonzero probability is the same, $b_{\min} < \beta^k$. W.l.o.g.

we can assume that $k = K$ and this bidder is the bidder with the highest possible valuations. Otherwise, in equilibrium, all bidders with types $v^j > v^k$ bid their respective β^j and outbid him. His only chance for a positive profit is if there are no bidders with a higher valuation. We can denote this event by D and the rest of the proof is identical even if $k < K$, when everything is conditional on D.

Denote by $q_2(n)$ the probability that a bidder with valuation v^K will bid $b_{\min}$ and let $q_2^{lim} = \lim_{n\to\infty} q_2(n)$ (the limit exists upto a subsequence). Also denote $q_1(n) = 1 - q_2(n)$. Let A be the expected payoff of the bidder with valuation v^K in the equilibrium in which he is bidding $b_{\min}$ at stage 1 and let B be the payoff when he deviates to bid β^K. We divide the remaining discussion into two cases: $q_2^{lim} < 1$ and $q_2^{lim} = 1$.

Case 1 : $q_2^{lim} < 1$**.**

In that case, for n large enough, there is a probability bounded away from zero that a bidder with valuation v^k will bid higher than b_{min}. Formally, there exists $\bar{q} \in (0, 1)$ and $N_0 \in \mathbb{N}$ such that for all $n \geqslant N_0$: $q_1(n) \geqslant \bar{q} > 0$. In both stages, no bidder bids strictly more than β^K. Thus, the probability of a bidder with valuation v^K winning the auction when using strategy σ^* can be (very loosely) bounded from below by $\frac{1}{n}$. To conclude, $B \geqslant (v^K - \beta^K)\frac{1}{n}$.

When a bidder bids b_{min} in the first stage, he will move to the second stage (and thus have some positive probability of winning) only if at most one bidder bids β^K in the first stage.

$$A \leqslant \left[(1 - \delta\bar{q})^{n-1} + (n - 1)(1 - \delta\bar{q})^{n-2}\delta\bar{q}\right](v^K - b_{min}). \tag{5}$$

For large enough n, this bound converges to 0 exponentially while the lower bound on B converges to 0 more slowly. Thus, for large enough n, $A < B$.

Case 2 : $q_2^{lim} = 1$**.**

In this case, at the limit, no bidder bids β^K in the first stage. Formally, for every $\varepsilon > 0$ there exists $N_0 \in \mathbb{N}$ such that the probability that a bidder with valuation v^K bids β^K in the first stage is smaller than ε. If a bidder with valuation v^K bids according to σ^*, then he wins with the probability of being chosen among all those bidding the same.

The number of bidders who have high valuation and bid β^K can be bounded from above using a binomial random variable with the parameters n and ε. It is well-known that for a binomial random variable $X \sim Bin(n, p)$,

$$E\left[\tfrac{1}{X+1}\right] = \tfrac{1}{np}\left[1 - (1 - p)^n\right], \tag{6}$$

therefore, the following bound is obtained: For ε small enough and n large enough $B \geqslant \frac{v^K - \beta^K}{n\varepsilon}(1 - (1 - \varepsilon)^n) = f_\varepsilon(n)$.

A bidder who bids b_{min} can win only if there is at most one bidder who bids above b_{min} in the first stage. Denote this event by E. Given E, the bidder needs to be chosen among all the other bidders who bid b_{min}.

Conditioning on E, the probability to win is lower than if all other bidders were (independently) bidding b_{min} with probability $\delta(1 - \varepsilon)$ and strictly lower with the remaining probability.

From Eq. (6) again, we have

$$A \leqslant \Pr(E)\frac{v^K - b_{min}}{n\delta(1-\varepsilon)} \leqslant 2\frac{v^K - b_{min}}{n\delta} = g_\varepsilon(n). \tag{7}$$

Since $\frac{f_\varepsilon(n)}{g_\varepsilon(n)} \to O(\frac{1}{\epsilon})$ as $n \to \infty$, then for ε small enough there exists N_0 such that for every $n > N_0$, $A < B$. To conclude, b_{min} cannot be played in the first stage in equilibrium. ■

Our next goal is to compare information structures. More precisely, we wish to compare the non-informative information structure to the informative ones in terms of expected revenue to the seller. Clearly, the revenue-maximizing equilibrium provides the best expected equilibrium payoff for the seller. He therefore strives to design an information model under which the revenue-maximizing equilibrium will be the unique one. Our next theorem states that whenever the revenue-maximizing equilibrium appears for some informative Θ, it also appears for the non-informative Θs.

Theorem 1. *When Σ^* is an equilibrium in an auction with some informative information structure, it is also an equilibrium in an auction with non-informative information structures.*

Proof. Following Lemma 1, the only strategies to consider in non-informative information structures are strategies in which the bid in both stages is the same. These strategies (which disregard the information) are also available in the informative information structures, so any possible deviation from Σ^* in the non-informed model is also a possible deviation in the informed model. As there is no profitable deviation in the informed model, such deviation obviously does not exist in the non-informed model either. ■

This result can easily be generalized. A strategy profile that is an equilibrium when information is available but not used, is still an equilibrium when this information is not available, simply because the set of available deviations is smaller without information. For example, if bidding b^1 regardless of type is an equilibrium in some informative information structure it is also an equilibrium in non-informative information structures, as the information about the bids does not change the prior regarding the bidders' types.

A direct corollary of this result is that a smaller number of bidders is needed to generate the revenue-maximizing equilibrium when no information is conveyed to the bidders. Any information provided by the seller might serve as a basis for profitable deviations.

Corollary 1. *Let N^{NI} be the minimum number of bidders needed for Σ^* to be an equilibrium when the information structure is non-informative. Fix Θ as some informative information model and let N^I be the minimum number of bidders needed for Σ^* to be an equilibrium when the information structure is Θ. Then $N^{NI} \leqslant N^I$.*

Proof. For N^I bidders, Σ^* is an equilibrium when the information is Θ. According to Theorem 1, Σ^* is also an equilibrium in the non-informed model with N^I bidders. Since N^{NI} is the minimum number of bidders for which Σ^* is an equilibrium in the non-informed model, it follows that $N^{NI} \leqslant N^I$. ■

The required number of bidders can be computed in each auction based on the auction parameters (bids, private values, and F_n). We demonstrate this calculation in Section 3.2. It is clear from an examination of the mechanism behind the example that for reasonable parameter values, ten or even fewer bidders could be sufficient for the revenue-maximizing equilibrium to emerge.

5 Concluding Remarks

This paper sought to determine the optimal information about bids that a seller should disclose to the bidders so as to increase her revenue in two-stage auctions. The answer to this question strongly depends on the number of bidders. When the number of bidders is small, all the parameters of the auction need to be considered; there are examples where it is better to reveal information and other cases where it is better to withhold it. For a very large number of bidders, the answer is that the information structure is irrelevant, as the only equilibrium is the one where the bidders submit their maximum bid, regardless of the information. Moreover, this strategy profile also becomes an equilibrium for fewer bidders when no information is revealed, unlike under any informative information structure.

We showed that the revenue-maximizing strategy profile requires fewer bidders to be an equilibrium when the information model is non-informative. We conjecture that this also applies to uniqueness – fewer bidders are required for the revenue-maximizing equilibrium to be the *unique* equilibrium when the information structure is non-informative. Any information revealed by the seller can create additional equilibria. This strengthens the idea that the best tactic on information is to reveal nothing – not only does the revenue-maximizing profile is an equilibrium first in this case, it also becomes the unique equilibrium first.

As elsewhere (e.g. Quint and Hendricks [2018]), uniqueness can be shown only in special cases. We were able to show that this holds for a simple example with pure strategies, which can be generalized to a large class of two-stage auctions. This proof is given in Appendix A. In addition, our conjecture regarding uniqueness was verified by computer simulation on another class of auctions for a large set of parameters.

We conclude that when the number of bidders is not too small[12] or is unknown, there is some advantage to not revealing information. This might allow the revenue-maximizing strategy profile to become an equilibrium and possibly the unique equilibrium.

If the sole purpose of the second stage is to encourage competition and the valuations are private, then it is advisable to consider simplifying the process and conducting a single-stage sealed-bid auction. However, the main purpose of the second stage may be to discuss

[12] As a rule of thumb, the number of bidders required is such that for every private valuation v^j, the expected number of bidders that have this private value given it is the maximal private value anyone has, is at least 3. For example, if valuations are independent and each bidder has equal probability to have any of the private valuation, there should be roughly $n = 3K$ bidders for a bidder with the valuation v^K to feel enough competitive pressure to raise his bids.

design (in design-build auctions) or to enable the bidders to learn more about the item being auctioned, either from the auctioneer or from the bids of the other bidders (when valuations have a common-value component). In such cases, the optimal information policy regarding first-stage bids remains to be determined by future research.

References

Marcus Ahadzi and Graeme Bowles. The private finance initiative: the procurement process in perspective. In *17th annual ARCOM Conference*, pages 5–7, 2001.

Valentina Bali and Matthew Jackson. Asymptotic revenue equivalence in auctions. *Journal of Economic Theory*, 106(1):161–176, 2002.

Jeremy Bulow and Paul Klemperer. Auctions vs. negotiations. Technical report, National Bureau of Economic Research, 1994.

George C Chelekis. *The Official Government Auction Guide*. Crown, 1992.

Gerti Dudkin and Timo Välilä. Transaction costs in public-private partnerships: a first look at the evidence. *Competition and regulation in network industries*, 1(2):307–330, 2006.

Gadi Fibich and Arieh Gavious. Large auctions with risk-averse bidders. *International Journal of Game Theory*, 39(3):359–390, 2010.

Kenneth Finley. Dynamic pricing possibilities in the purchase of bulk fuel for the department of defense. Technical report, Naval Postgraduate School Monterey CA, 2001.

Juan-José Ganuza. Competition and cost overruns in procurement. *The Journal of Industrial Economics*, 55(4):633–660, 2007.

Andrew Merrifield, Tjiamogale Eric Manchidi, and Stephen Allen. The asset procurement and operating partnership system (apops) for prisons in south africa. *International Journal of Project Management*, 20(8):575–582, 2002.

Paul R. Milgrom. A convergence theorem for competitive bidding with differential information. *Econometrica*, 47(3):679–688, 1979.

Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

Zvika Neeman. The effectiveness of english auctions. *Games and Economic Behavior*, 43(2): 214–238, 2003.

Paul Noumba and Severine Dinghem. Private participation in infrastructure projects in the republic of Korea. *World Bank Policy Research Working Paper*, (3689), 2005.

Motty Perry, Elmar Wolfstetter, and Shmuel Zamir. A sealed-bid auction that matches the english auction. *Games and Economic Behavior*, 33(2):265–273, 2000.

Wolfgang Pesendorfer and Jeroen M Swinkels. Efficiency and information aggregation in auctions. *American Economic Review*, 90(3):499–525, 2000.

Daniel Quint and Kenneth Hendricks. A theory of indicative bidding. *American Economic Journal: Microeconomics*, 10(2):118–51, 2018.

Kai Rintala, David Root, Graham Ive, and Paul Bowen. Organizing a bidding competition for a toll road concession in South Africa: the case of Chapma's peak drive. *Journal of Management in Engineering*, 24(3):146–155, 2008.

Bruce A Roll. Online reverse auctions: a pricing tool for government contracting. Technical report, Naval Postgraduate School Monterey CA, 2000.

Bruno Sauvet-Goichon. Ashkelon desalination plant – A successful challenge. *Desalination*, 203 (1):75 – 81, 2007.

Jeroen M Swinkels. Efficiency of large private value auctions. *Econometrica*, 69(1):37–68, 2001.

Robert Wilson. A bidding model of perfect competition. *The Review of Economic Studies*, 44 (3):511–518, 1977.

Robert Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1115, 1985.

Lixin Ye. Indicative bidding and a theory of two-stage auctions. *Games and Economic Behavior*, 58(1):181–207, 2007.

A Uniqueness of the Revenue-Maximizing Equilibrium

We show that Σ^* is the unique equilibrium even for fewer bidders when the information structure is non-informative, taking the case of two possible valuations, three possible bids and considering pure strategies only. We conjecture that this is the case in general, regardless of the specific parameters of the auction. Partial results in the general case, based on computer simulations, are also reported.

The following proposition compares two extreme information structures: the non-informative information structure and the fully-informative information structure (where the first-stage bids are revealed to the bidders before the second stage).

Proposition 1. *Consider an auction with n bidders, two possible types $V = \{v^1, v^2\}$, and three possible bids $B = \{b^1, b^2, b^3\}$ ordered in the following way: $b^1 < v^1 < b^2 < b^3 < v^2$. Assume that the valuations are independent and identically distributed, with probability p of each bidder having valuation v^2. If Σ^* is the unique equilibrium in pure strategies when the information structure is fully informative, then it is also the unique equilibrium in pure strategies when the information structure is non-informative.*

Proof. Assume w.l.o.g. that $b^1 = 0$ and $v^2 = 1$. Since $b^1 < v^1 < b^2 < b^3 < v^2$, bidders of type v^1 always bid in equilibrium b^1. Bidders of type v^2 can choose any of the three possible bids. In the rest of the discussion, we focus on one such bidder. We denote by X the number of other high-type bidders, i.e. $X \sim Bin(n-1, p)$. The possible bids, private values, and revenues are illustrated in this figure:

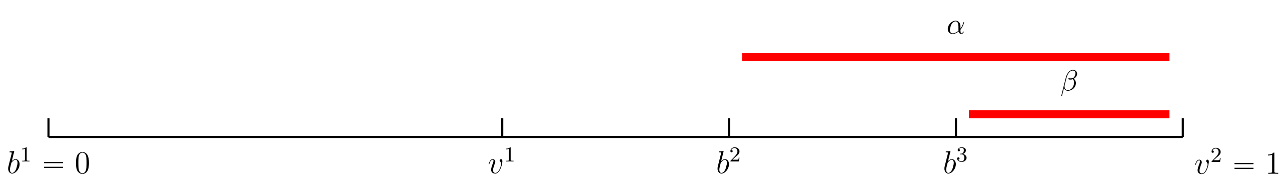

The strategy of the proof is to construct a counter-example – a choice of parameters for which Σ^* is the unique equilibrium in the informed model but not in the non-informed model. We will show that such a counter-example is impossible.

The non-informed model

For a counter-example, we need an equilibrium in pure strategies additional to Σ^* in the non-informed model, but no additional equilibria in the fully-informed model.

Clearly, if the strategy profile in which everyone bids b^1 is an equilibrium in the non-informed model, then it is also an equilibrium in the fully-informed model. This is because they are essentially the same – no new information is revealed before the BAFO stage, as all bids are known and are the same regardless of private value. Therefore in a counter-example, it cannot be an equilibrium, which means that there exists a profitable deviation from this strategy profile. The profitable deviation can be to bid b^2 or to bid b^3. However, if all other bidders bid b^1, bidding b^3 has no strategic advantage over bidding b^2: both guarantee a win, and the revenue is higher when b^2 is bid. Thus, bidding b^2 is a profitable deviation:

$$U_i(b^1|v^2) = \frac{1}{n} < \alpha = U_i(b^2, b^1_{-i}|v^2) \tag{A.1}$$

The only other pure strategy profile is to bid b^2 (denoted by Σ), so it is assumed to be the additional equilibrium. Hence, bidding b^3 is not a profitable deviation:

$$U_i(\Sigma|v^2) = \alpha\mathbb{E}(\frac{1}{X+1}) \geqslant U_i(b^3, \Sigma_{-i}|v^2) = \beta \tag{A.2}$$

In addition, Σ^* is equilibrium in the non-informed model so bidding b^2 is not a profitable deviation:

$$U_i(\Sigma^*|v^2) = \beta\mathbb{E}(\frac{1}{X+1}) \geqslant U_i(b^2, \Sigma^*_{-i}|v^2) = \alpha\Pr(X=0) \tag{A.3}$$

The fully-informed model

We construct an example where Σ^* is the unique equilibrium with information. Thus, there is no profitable deviation from Σ^* and no other equilibrium. More precisely:

Bidding b^2 in both stages regardless of information is not an equilibrium in the fully-informed model. The only profitable deviation can be to employ a strategy that uses information, which must be either σ_0="play 0 at the first bid, then best respond to information" or σ_α="play b^2 at the first bid, then best respond to information".

Is σ_α a profitable deviation? The best response is to continue bidding b^2 if the other bid is 0 (there are no v^2 bidders) and bid b^3 if the other bid is b^2 (otherwise, bidding b^2 again would result in the equilibrium strategy). To conclude, the best response is getting β with probability 1 rather than splitting α between the two. Thus, with some implicit conditions on the bids, we get:

$$0 < \tfrac{\alpha}{2} < \beta < \alpha < 1. \tag{A.4}$$

In addition, σ_α cannot be an equilibrium in the fully-informed model, and the only possible deviation is bidding b^3 in both stages:

$$U_i(\Sigma_\alpha|v^2) = (\alpha-\beta)\Pr(X=0) + \beta\mathbb{E}(\frac{1}{X+1}) < U_i(b^3, \Sigma_{\alpha-i}|v^2) = \frac{\beta}{2}(1+\Pr(X=0)). \tag{A.5}$$

We can rearrange these inequalities as inequalities between α and β and combine them all:

$$0 < \alpha\max\{\frac{1}{2}, \frac{\Pr(X=0)}{\mathbb{E}(\frac{1}{X+1})}, \frac{\Pr(X=0)}{\frac{1}{2}+\frac{3\Pr(X=0)}{2}-\mathbb{E}(\frac{1}{X+1})}\} < \beta < \alpha\mathbb{E}(\frac{1}{X+1}) < 1 \tag{A.6}$$

In addition, from the properties of expectation, it is clear that $\mathbb{E}(\frac{1}{X+1}) < \frac{1}{2} + \frac{\Pr(X=0)}{2}$. Set $x = \mathbb{E}(\frac{1}{X+1}), y = \Pr(X=0)$. A necessary condition for (A.6) to hold is

$$\frac{1}{2} < x < \frac{1}{2} + \frac{y}{2} \tag{A.7}$$

$$y < x^2 \tag{A.8}$$

$$y < (\frac{1}{2} + \frac{3}{2}y - x)x \tag{A.9}$$

Summing (A.8) and (A.9) yields $y < \frac{x}{4-3x}$ which combined with (A.7) results in $\frac{1}{2} < x < \frac{2}{3}$. On the other hand, (A.9) is equivalent to $y(1 - \frac{3}{2}x) < \frac{1}{2}x - x^2$. The right-hand side is negative in the region $\frac{1}{2} < x < \frac{2}{3}$ whereas the left-hand side is positive. A contradiction – σ_α is not a profitable deviation and $\beta \leqslant \frac{\alpha}{2}$. △

Can σ_0 be a profitable deviation? The profit with the deviation is:

$$U_i(\sigma_0, \Sigma_{\alpha_{-i}}|v^2) = \Pr(X = 0)\frac{2}{n}\max\{\alpha, \frac{1}{2}\} + \Pr(X = 1)\frac{1}{n-2}\max\{\beta, \frac{\alpha}{2}\} \tag{A.10}$$

Our previous case resulted in $\beta \leqslant \frac{\alpha}{2}$. In addition, $\frac{1}{n} \leqslant \alpha$ and $\frac{2\alpha}{n} \leqslant \alpha$ so $\frac{2}{n}\max\{\alpha, \frac{1}{2}\} \leqslant \alpha$:

$$\begin{aligned} U_i(\sigma_0, \Sigma_{\alpha_{-i}}|v^2) &\leqslant \Pr(X = 0)\alpha + \Pr(X = 1)\frac{1}{n-2}\frac{\alpha}{2} \leqslant \\ &\leqslant \Pr(X = 0)\alpha + \Pr(X = 1)\frac{\alpha}{2} = \alpha(\Pr(X = 0) + \frac{1}{2}\Pr(X = 1)) < \\ &< \alpha\mathbb{E}(\frac{1}{X+1}) = U_i(\Sigma_\alpha|v^2) \end{aligned} \tag{A.11}$$

which means that σ_0 is not a profitable deviation either. △

To conclude, it is impossible to satisfy all the conditions in these settings, and if Σ^* is not unique without information, it cannot be unique with information. ■

The key component of the proof is that only bidders with private valuation v^2 need to be considered and there is a relatively small number of parameters and inequalities that lead to a contradiction. This proof can be applied to more general cases, with more possible private values and bids, as long as these features remain.

Changing the order to $b^1 < b^2 < v^1 < b^3 < v^2$ complicates the problem significantly. After removing dominated strategies, there are 13 (pure) strategy profiles in the informed model which need to be excluded from being an equilibrium. This results in a large number of inequalities and "cases" to verify, that do not converge to such a neat result as in the Proposition and, more importantly, cannot be easily generalized.

We, therefore, wrote a computer simulation to test our hypothesis in slightly more general settings for different values of p, n, v^1 and M, the number of possible bids (assuming that they are equally spaced, $B = \{0, \frac{1}{M}, \ldots, \frac{M-1}{M}\}$). For each $M \in \{4, \ldots, 15\}$ and $n \in \{4, \ldots, 15\}$, we constructed a grid with 100 points for (p, v^1) in the domain $[0.1, 0.4] \times [0.09, 0.5]$ and calculated all possible pure-strategy equilibria for the chosen parameters. Despite our efforts, we did not manage to find a counter-example for our conjecture, i.e. a set of parameters for which Σ^* is the unique equilibrium with information but not the unique equilibrium without information.

B Alternative Assumption Regarding the Joint Distribution of the Private Valuations

A close examination of the proof of Lemma 2 reveals that the uniform full support distribution can be slightly relaxed. We chose the uniform full support assumption since it is both general and easy to verify. Alternative assumptions are more general but harder to verify; moreover, they rely on the bidding strategies and not the distributions. Here we present one such generalization.

Instead of the uniform full support assumption, assume that there exists $\delta > 0$ such that for every $k \in \{1, ..., K\}$, when a bidder has a valuation v^k, conditional on D_k[13] and given that:

- If each bidder with valuation v^k bids within some range of values with probability $\bar{q}$ at most (independently from the other bidders having valuation v^k), then the probability that at most one other bidder will bid within that range is bounded from above by

$$1 - \bar{q}\delta)^{n-1} + (n-1)\bar{q}\delta(1-\bar{q}\delta)^{n-2}. \tag{B.1}$$

  This term comes from the best-case scenario for the bidder. If all the other bidders have the private valuation v^k with the lowest possible probability (δ), then the number of v^k-type bidders who bid in this range is a random variable X with binomial distribution with parameters $n-1$ and $\delta\bar{q}$, and the right-hand side is $\Pr(X \leqslant 1)$.

- If we denote by X the number of other bidders with valuation v^k, then $E[\frac{1}{X+1}] \leqslant \frac{[1-(1-\delta\bar{q})^n]}{n\delta\bar{q}}$. The right-hand side is the expectation when X has the binomial distribution with parameters $n-1$ and $\bar{q}\delta$, i.e. when the bidders are independent (see Eq.(6)).

Note that this assumption is satisfied when valuations are independent with a lower bound of δ on the probabilities of the values in the support. However, some positive correlation typically makes this condition easier to satisfy – given a bidder's valuation, the probability of bidders with similar valuations should increase. Negative correlation is also possible. However in extreme cases, if a bidder's valuation lowers the probability of the other bidders having a similar valuation, then the assumption does not hold.

C Continuous Auctions

The main text presented a two-stage auction model with discrete sets of valuations and possible bids. This mirrors common practice, where currency is discrete and in many cases, auctioneers require "rounded" bids (for example, bids may have to be in increments of $10,000\$$). Moreover, this assumption simplifies the analysis, as the existence of equilibria under all information structures is ensured in this model and the revenue-maximizing strategy profile, Σ^*, is well defined. In this Appendix, we provide several results regarding the continuous case. Although not a full analysis of the problem, our findings point to additional drawbacks of revealing information and strengthen the claim made in the main text.

Our model is identical to the model in Section 2, except for the continuity of the bids and valuations. Hence, we assume that $B = [0, \infty)$ and $V = [v_{min}, v_{max}]$ with $0 \leqslant v_{min} < v_{max}$.

[13]Recall, D_k is when the types of all bidders are at most v^k.

We only consider the full information case, and assume that the first-round bids of the two finalists are revealed before the second round. The tie-breaking rule in the second round is not necessarily symmetric. For example, it can favor the bidder who made the highest first-round bid.

In this model, the β^js are not well defined and therefore neither is the revenue-maximizing strategy profile (or, if defined as the strategy profile in which all bidders bid their valuations in both rounds, it is obviously not an equilibrium). Instead, we study the existence of an efficient equilibrium, i.e. an equilibrium in which the bidder with the highest valuations wins the item, which in many cases also provides the highest expected payoff to the seller. We show that under mild additional assumptions, there are no symmetric efficient equilibria in pure strategies. This leads us to the conclusion that, as in the discrete model, there are no advantages to a BAFO stage whose sole purpose is to increase auctioneer's revenue, without providing additional information about the item being auctioned. We note that it is not even clear that an equilibrium exists in this model. We therefore prove in this Appendix that *if* an equilibrium exists, it is not efficient.

Consider a symmetric efficient equilibrium in pure strategies. In such an equilibrium there exists a function $f : V \to B$ which represents the first-round bid. Since the equilibrium is efficient, f must be strictly increasing (otherwise, the highest bidder might not be selected for the second round). Hence, on the equilibrium path, valuations are revealed by the first-round bids. The second-round bids can be described by the two functions $g, h : V \times V \to B$, where $g(x, y)$ is the second-round bid of the bidder who bid the highest value $f(x)$ in the first round and $h(x, y)$ is the bid of the runner-up, who bid $f(y)$ in the first round. The equilibrium is efficient, so $g(x, y) \geqslant h(x, y)$.

For simplicity, we assume that the valuations are i.i.d. and drawn from V with a probability distribution $p(v)$ that is continuous and positive on V. This ensures that the probability that a valuation falls within a particular interval of (small) length ϵ is of order ϵ. The following Lemma provides additional properties of the equilibrium.

Lemma 3. *Let x and y be possible valuations, with $x > y$. Almost surely:*

1. $y \leqslant g(x, y) \leqslant x$.
2. $g(x, y) = \max(f(x), h(x, y))$.
3. *If* $f(x) < y$*, then* $h(x, y) = g(x, y) \geqslant y$.
4. *The assumption that* $h(x, y) \geqslant y$ *is made without loss of generality.*
5. *If* $x > v_{min}$*, then* $f(x) < x$.

Proof.

1. If $g(x, y) > x$, the player with valuation x is bidding above his value and should slightly decrease his second-stage bid; if $g(x, y) < y$, the player with valuation y should deviate and bid slightly above $g(x, y)$.

2. The auction rules prevent bidding below the first-round bid, so $g(x, y) \geqslant f(x)$. Clearly, $g(x, y) \geqslant h(x, y)$ because the auction is assumed to be efficient, so $g(x, y) \geqslant \max(f(x), h(x, y))$. The reverse inequality holds because otherwise the winner should decrease his second-round bid.

3. Immediate from 1 and 2: $g(x, y) \geqslant y > f(x)$ so $g(x, y) \neq f(x)$ and the only remaining possibility is $g(x, y) = h(x, y)$.

4. If we replace h by $\tilde{h}$ defined by $\tilde{h}(x, y) = \max(y, h(x, y))$, then we still have an equilibrium. Indeed, this does not change anything when $f(x) < y$ by point 3), and when $f(x) > y$, this does not change the equilibrium payoffs nor the possible deviations and their payoffs. Therefore, we can assume that the equilibrium we have in hand is the one where $h(x, y) \geqslant y$.

5. If $f(x) = x$ and $x > v_{min}$, then the bidder obtains a zero payoff though he could obtain a positive payoff by playing as if he had a slightly lower valuation.

■

For our proof, we still need an extra assumption. Several natural assumptions are possible. We choose to assume that the runner-up never bids above his value:

Assumption 2. *For all $x, y \in V$ with $x > y$, we assume $h(x, y) \leqslant y$.*

Combined with Lemma 3 (part 4), on equilibrium path and without loss of generality, $h(x, y) = y$ and by part 2, $g(x, y) = \max(f(x), y)$. It follows that g is non-decreasing in each of its arguments and strictly increasing if both arguments increase: $x < x'$ and $y < y'$ imply $g(x, y) < g(x', y')$. We are now ready to show that in this model, an efficient equilibrium does not exist, i.e. that the above properties cannot all hold together.

Proposition 2. *In a two-stage BAFO auction with a continuous set of bids and valuations, there are no symmetric efficient equilibria in pure strategies when all bids are revealed to the bidders before the BAFO stage.*

Proof. Assume by contradiction that such an equilibrium exists and consider bidder n with valuation $z > v_{min}$. Let $x \geqslant y$ be the two highest valuations among the other $n - 1$ bidders. Almost surely, $x > y$, which we now assume. Let $\epsilon > 0$ and consider a deviation by bidder n that consists of bidding $f(z - \epsilon)$ in the first round instead of $f(z)$ and then playing optimally in the second round. We study the gain from this deviation in all possible cases (with non-zero probability) and show that it is positive for ϵ small enough. The main two cases are the case where bidder n does not have the highest valuation ("case A": $z < x$) and where bidder n has the highest valuation ("case B": $z > x$). The full analysis of all sub-cases that occur with positive probability is provided here and summarized in Table C.

Case A: $z < x$

Case A.1: $f(x) < z - \epsilon < z < x$.

As ϵ goes to 0, this case occurs with probability of order 1 (bounded away from zero). Bidder n's equilibrium payoff is 0, as the highest bidder wins the auction. If bidder n deviates, he is still selected and his opponent's second-round bid is $g(x, z-\epsilon) = h(x, z-\epsilon)$, so the deviation payoff is arbitrarily close to $\max(0, z - h(x, z - \epsilon))$, which is equal to ϵ under Assumption 2. Thus, the gain is of order ϵ.

Case A.2 Otherwise.

The payoff from the equilibrium strategy is 0 and the payoff when deviating is non-negative, so the deviation gain is non-negative.

Case B: $z > x$

Case B.1 $z - \epsilon > x$.

This occurs with probability of order 1. The equilibrium payoff is $z - g(z, x)$ and the payoff when deviating is $z - g(z - \epsilon, x)$. The gain from deviation is therefore

$$g(z, x) - g(z - \epsilon, x) = \max(f(z), h(z, x)) - \max(f(z - \epsilon), h(z - \epsilon, x))$$

under Assumption 2 this results in $\max(f(z), x) - \max(f(z - \epsilon), x)) \geqslant 0$.

Case B.2 $y < z - \epsilon < x < z$.

This occurs with probability of order ϵ. The equilibrium payoff is $z - g(z, x) \leqslant z - x$, whereas when deviating, the payoff is arbitrarily close to $z - g(x, z - \epsilon) \geqslant z - x$. Thus, there is at most an arbitrarily small loss from deviating. Under Assumption 2, the gain is actually strictly positive:

$$g(z, x) - g(x, z - \epsilon) = \max(f(z), x) - \max(f(x), z - \epsilon) \geqslant x - \max(f(x), z - \epsilon) > 0.$$

Case B.3 $z - \epsilon < y < x < z$.

This occurs with probability of order ϵ^2. The payoff when deviating is zero while the payoff in equilibrium is $z - g(z, x) \leqslant z - x \leqslant \epsilon$. Hence, by deviating, the bidder loses at most ϵ.

To summarize, under Assumption 2, by bidding $f(z - \epsilon)$ instead of $f(z)$, the bidder can obtain a positive profit of order ϵ and a loss of order ϵ^3. Hence, for ϵ small enough, this is a profitable deviation and the functions f, g, h do not form a symmetric efficient equilibrium. ■

This proposition shows that with information, equilibria, if they exist, cannot be efficient. Contrastingly, without information the auction is equivalent to a first-price (single-stage) sealed-bid auction (the proof of Lemma 1 is also valid in this model), so there exists an efficient equilibrium. Although efficiency does not directly imply best revenue for the seller, it is often

Case	Order	Deviation gain
A1. $f(x) < z - \epsilon < z < x$	1	ε
A2. Other cases with $z < x$	1	$\geqslant 0$
B1. $z - \epsilon > x$	1	$g(z, x) - g(z - \epsilon, x) \geqslant 0$
B2. $y < z - \epsilon < x < z$	ϵ	$g(z, x) - g(x, z - \epsilon) > 0$
B3. $z - \epsilon < y < x < z$	ϵ^2	$g(z, x) - z \geqslant x - z \geqslant -\epsilon$

Table 1: Summary of the 5 cases under Assumption 2.

the case and the lack of efficiency suggests that there might be disadvantages to revealing information. We conclude that there are good reasons not to reveal information between the stages in the continuous case as well, which reinforces the paper's findings under the discrete model.