Computing Profit-Maximizing Bid Shading Factors in First-Price Sealed-Bid Auctions*

Paulo Fagandini† Ingemar Dierickx‡

2022

Abstract

Computational methods are used to determine a profit-maximizing shading factor by which rational bidders shade their bid in first price sealed bid auctions for a broad range of realistic scenarios when the prior is diffuse. Bidders' valuations may have both common value and firm-specific components, and the accuracy of their estimates of the common value component may differ. In addition, we allow for a subset of "naive" rivals, defined as bidders who do not account for the Winners' Curse. Our computations show that profit-maximizing shading is greatly impacted by asymmetries in the bidding population and, in particular, by the presence of naive bidders. Failing to account for the presence of naive bidders results in underbidding only in one case, when facing a single rival who is naive, and in overbidding in all other cases. Overbidding is particularly severe when the population of naive competitors is large.

Keywords: Winner's CurseAuctionsBiddingAsymmetric AgentsNaive Bidders.

1 Introduction

In this paper we address the problem that bidders face in first price sealed bid (FPSB) auctions for common value goods: How much to shade their signals about the value of the object being sold. Countless firms are grappling with this problem every day. Two different disciplines, game theory and decision sciences, have taken totally different approaches to this problem. Excellent reviews of the different approaches to auctions can be found in Rothkopf (2007) and Lorentziadis (2016).

---

*We are indebted to Steffen Hoernig for his insightful advice in this paper. This paper also profited from valuable interaction with Guido Maretto, Susana Peralta, Fernando Anjos, Pedro Vicente, Patrick Rey, and Robert Wilson, as well as the referees who helped to significantly improve this paper. The usual disclaimer applies. Paulo Fagandini thanks the support of the Fundação para a Ciência e a Tecnologia FCT under the grant SFRH/BD/105669/2015. This is partially based on a dissertation for a PhD in Economics at Nova School of Business and Economics.

†CEFAGE-ISCAL, and Nova School of Business and Economics

‡I.D. Consulting Ltd.

We build upon the pioneering work of Rothkopf (1969) and Wilson (1984), who use non-Bayes Nash equilibrium models. Like them, we work on the classical mineral rights model, in which an indivisible good is auctioned in a first price sealed bid fashion among competitive bidders, who produce an unbiased estimate about its value, unknown to all of them at the moment of the auction. Like Hubbard et al. (2013); Hubbard and Paarsch (2014), we propose a computational method to address a broad range of realistic valuation and information scenarios. However, our paper differs in several aspects: 1) we focus on a common value component instead of a private valuations, 2) we consider an unbounded support for the signals, 3) we allow for naive bidders, and 4) we use a different methodology that enables firms to submit bids that maximizes ex-ante expected profits.

Specifically, we derive first order conditions and compute a constant shading factor (SF) ex-ante of receiving the signal that (i) allows for a common value component as well as a firm-specific component in valuations, (ii) allows for differences in the accuracy of bidder signals, and (iii) allows for the introduction of non-rational bidders.[1] In real life, constant shading rules such as constant absolute markup and constant percentage markup are commonly used. Shachat and Wei (2012) find that bids and prices in laboratory experiments agree with game-theoretic predictions in English auctions, but not in First Price Sealed Bid auctions, where constant shading strategies are used more often. Rothkopf (1980), and Compte and Postlewaite (2012) provide discussions on why constant strategies, and shading before observing the signal, should approximate Bayesian equilibrium strategies when the prior is diffuse.[2]

To test whether the SF results in Bayes Nash equilibrium bids, we use as benchmark the analytical solution to the symmetric problem with normally distributed noise Levin and Smith (1991); Hoernig and Fagandini (2018),[3] and find that the SF exactly replicates those results when the prior is diffuse. In addition, to verify the SF in a broader range of scenarios, we compared it with results from "brute force" Monte Carlo simulations. In all cases, we found that results from the SF and "brute force" Monte Carlo simulations coincide.

We also generalize Robert Wilson's bias factor (BF) to obtain a measure of the Winner's Curse. In short, the BF shades the bidder's signal by the expected error of the signal conditional on winning. Therefore, this correction allows the bidder to obtain zero expected winning profits, avoiding the Winner's Curse. To obtain positive expected winning profits, shading must exceed the BF by some margin. This margin depends on the number of rivals that a bidder is facing. The more bidders in an auction, the more aggressive a bidder must be in order to have a chance at winning. However, the more bidders in an auction, the more severe the Winner's Curse, and therefore the bids should be more conservative. Optimal shading should take into account both the Winners' Curse and the competitive effect (Thaler, 1988, p.192). To disentangle these two opposite effects we consider the BF as that part of the SF that takes care of avoiding the Winner's Curse, while the remainder ($SF - BF$) accounts for the competitive effect.[4]

[1] We use rational and sophisticated as equivalent terms. A naive bidder is a type of irrational bidder whose behavior does not follow any profit maximization given all the available information. These can use basic rules of thumb or other strategies. We assign a strategy to our naive bidder later in the text.

[2] As Heumann (2019) notes, because of its tractability the model studied by Wilson (1998) is often used in empirical work. Diffuse priors are also considered in recent theoretical and empirical papers, *e.g.* Vives (2011); Hong et al. (2013).

[3] Wilson (1969) derived the first explicit bidding function; however it was limited to only two symmetric bidders.

[4] Hoernig and Fagandini (2018) decomposes the optimal shading as the expected value plus the dis-

Finally, we allow for a subset of naive bidders, who follow a simple rule of thumb and shade their signals by an arbitrary fixed amount. Dyer et al. (1989) posit that even experienced bidders follow simple strategies when facing an auction that would work only in invariant environments. While our model allows us to set any fixed shading for these naive bidders we assume they are naive only in a limited sense, *viz.* that they do not account for the Winner's Curse. Specifically, we assume that naive bidders shade their bids by $(SF - BF)$. That is, they do not account for the Winner's Curse, but do properly account for the competitive effect and analyze their impact on the optimal bids of sophisticated bidders.[5] The presence of naive bidders in real life bidding problems cannot be denied. The notion of the Winner's Curse was first discussed by three Atlantic Richfield engineers in a study of field data in the oil industry Capen et al. (1971). The Winner's Curse cannot occur when all bidders act rationally Cox and Isaac (1984). However, as Thaler (1988) stresses, bidding in a common value auction can be very challenging. Occurrence of the Winner's Curse in common value auctions has been acknowledged for more than a half a century Kagel and Levin (2002), providing strong evidence for the presence of naive bidders. Dyer et al. (1989) document that in laboratory experiments even experienced executives in the construction industry, who are successful in their jobs, suffer from the Winner's Curse. They suggest that industry specific learning and situation-specific rules of thumb, which could not be applied in a laboratory setting, may help them avoid overbidding in the field. Furthermore, experienced contractors do suffer unanticipated losses when bidding on a type of project they are not familiar with.[6] These findings indicate that in real life auctions it is not unlikely that a subset of bidders may be naive. We find that failing to account for the presence of naive bidders results in underbidding only in one case, when facing a single rival who is naive, and in overbidding in all other cases. Losses due to overbidding are particularly severe when the population of naive competitors is large.

The paper is organized as follows: Starting with Section 2 we present a brief literature review for the main works related to this article. In Section 3 we present the model and introduce the shading factor. In Section 4, we present the bias factor. In Section 5 we show how shading factors react to specific asymmetries in the bidding population. Section 6 summarizes key conclusions and implications of our work and suggests avenues for future research.

2 Literature Review

The game theory literature models auction outcomes as Bayes Nash equilibria. While it offers valuable advice for auctioneers on auction design, it does not offer bidders a practical methodology to compute optimal bids for a wide range of realistic scenarios. Some papers, most notably Wilson (1967, 1969, 1977, 1998), Milgrom and Weber (1982a); Levin and Smith (1991), and more recently Hoernig and Fagandini (2018) actually pin down a bidding function for the bidders, but elegant analytical solutions are derived at the cost of a good amount of simplifying assumptions that are rarely satisfied in real life.

persion index of the maximum error. While the first deals with the Winner's Curse, the second accounts for the competitive effect.

[5] The term naive is used in the same sense by Kagel and Levin (2002), and in a similar way in Lorentziadis (2012).

[6] Thaler (1988) provides a very good review of laboratory and field studies that document the existence of the Winner's Curse.

Another limitation of this literature is that theoretical results are usually derived in models with uniformly additive noise, an assumption that is required to solve those models analytically.[7] In the applied literature, normal and log-normal distributions are usually thought to be best suited to model real life problems. For example, electricity markets are usually modeled with additive normal noise, while hydrocarbon reservoirs in the petroleum industry are treated with a log-normal distribution.[8]

Although some game theoretical papers include asymmetries in the distributions of the signal noise — e.g. Wilson (1998) — most of the literature assumes some symmetry in at least one dimension, as well as rationality of all the bidders.[9] These are important limitations. While these assumptions are often necessary to ensure tractability of the models, they are, as Armantier and Sbai (2006); Hubbard and Paarsch (2009); Hubbard et al. (2013); Hubbard and Paarsch (2014) have pointed out, restrictive and not often found in the real world. Furthermore, as we show in this paper, even small deviations from these assumptions result in sharply different optimal bids.

Recently game theoretic papers have begun to model different degrees of sophistication among the population of bidders. In particular, Eyster and Rabin (2005), with their concept of *Cursed Equilibrium*, provide a new explanation for the Winner's Curse. In their model, some agents fail to take into account how information impacts the other players' strategies. The cursedness hypothesis indeed improves fitting laboratory data to the models. However, this approach also suffers from simplifying assumptions common to the theoretical literature: the use of uniform distributions (and bounded domains), and the necessity to impose other symmetry assumptions, such as the degree of cursedness.

Crawford and Iriberri (2007) also study out-of-equilibrium models to explain bidders' behavior that is inconsistent with the traditional Bayes Nash solution. In particular, they apply the concept of *Level-K Thinking* (LK) introduced by Stahl and Wilson (1994, 1995). That means that bidders are assumed to go through a limited number of K iterations to reach a certain level of 'best response'. Say, if there are two bidders, both bidders assume a particular strategy by the other bidder. They best respond to that strategy. Later, realizing that the other player could have followed the same reasoning, they decide to best respond to the previous strategy. How many times they follow the reasoning of best responding to the previous strategy is the value of "K." The higher the number of iterations K, the closer is the equilibrium to a Nash Equilibrium. *Level-K Thinking* provides an explanation for overbidding in laboratory data in first price auctions, both in common value and in independent-private-value scenarios. They estimate from the data that the distribution of the Ks has the highest weights on $L1$ and $L2$.

Other avenues taken by the recent game theoretic literature on auction theory are focused on complex problems, such as auctions on divisible goods Wang and Zender (2002), multi-unit auctions Hortaçsu and Puller (2008), ascending auctions Heumann (2019), or uniform versus discriminatory auctions Fabra et al. (2006); Hortaçsu and McAdams (2010). While these contributions do indeed study more intricate and more realistic scenarios, they do not provide a practical methodology enabling bidders to compute optimal bids, which is the central aim of this paper.

[7] Notable literature reviews can be found in Klemperer (2004); Krishna (2010); Milgrom (2004); Salant (2014).

[8] For example Crawford (1970), Smiley (1979).

[9] Following the literature, we use the term "symmetry" when referring to the distribution from where the signals are drawn. To avoid confusion, we will use "heterogeneity" when referring to differences in the the bidders' level of sophistication.

Decision theory has paid closer attention to the bidders' problem. Since it is not, in general, feasible to analytically derive the Bayes Nash equilibrium bidding strategies in realistic scenarios, bidding strategies are optimized against a given distribution of competitors' bids relying mostly on Monte Carlo simulations.[10] There are some early decision theoretic and experimental contributions from the 50's and the 60's,[11] there are papers that address the bidders' problem within the context of a specific industry,[12] and there is a body of papers that propose computational methods. One computational approach uses data from previous auctions and runs Monte Carlo simulations to determine optimal bids. Key papers include David (1993), Wen and David (2001), Ma et al. (2005), among others. In a second computational approach, data from earlier auctions are used to estimate the moments of the distribution from which bidders' signals are drawn, assuming that the other bidders play Bayes Nash equilibrium strategies, and optimal bids are then computed using classical symmetric models. Key papers include Bajari (1998), Bajari and Hortaçsu (2005), Campo et al. (2003), among others.

Recent contributions to the literature on computational economics that look to approximate bidding functions in settings where analytical solutions are not feasible include Hubbard and Paarsch (2009); Hubbard et al. (2013); Hubbard and Paarsch (2014). Hubbard and Paarsch (2009) use computational methods to characterize the bidding behavior of asymmetric bidders in a setting where the auctioneer has a preference for some of those bidders. Among other interesting results they are able to show how non-preferred bidders bid more aggressively to remain competitive. Hubbard et al. (2013) use numerical methods to approximate the inverse-bid functions in first price auctions with asymmetric bidders, which allows the bidder to estimate the probability of winning conditional on their bid. They find that low order polynomials perform poorly approximating the bidding functions and use theoretical results about the crossing of the cumulative distribution of valuations to assess the quality of this numerical approximation. Hubbard and Paarsch (2014) expand this framework, including a wide set of extensions such as risk aversion, collusion, among others, and analyze different algorithms such as the shooting method (see Marshall et al. (1994)) and the projection methods (which encompass the polynomials approach).

3 The Model

In this section, we introduce the model for the bidders' behavior. We find the bidder's condition that maximizes profits, and use this to compute her best response. Next, we iterate over all the bidders to reach the equilibrium, i.e. when every bidders' best response coincides, within an arbitrary degree of tolerance, with its response in the previous iteration.

Assume that there are n bidders, competing in a First Price Sealed Bid auction for an item with unknown value $\mu_i = \mu + \Delta_i$, where μ is the unknown common value part for everyone, and Δ_i represents a private value component. All bidders receive an unbiased signal of their valuation for the auctioned object $s_i = \mu_i + \epsilon_i$, with ϵ_i independently distributed $N(0, \sigma_i^2)$.

[10] R. Wilson and M. Rothkopf are probably among the most prolific authors that actually looked at both sides, game theory and decision sciences.

[11] For example Friedman (1956), Ortega Reichert (1968), among others.

[12] *E.g.* Capen et al. (1971).

This specification allows us to consider bidder-specific valuation differences (Δ_i) and differences in the accuracy of bidders' signals (σ_i^2). Specifically, we allow for two groups of bidders,[13] whose members receive a signal drawn from the same family of distributions but with different moments.[14] These may stem from a variety of factors such as different economies of scale, experience, technological factors, and so on. If $\Delta_i = 0$ for every bidder we have a pure common value scenario.

We also allow for the presence of naive bidders as well as rational bidders. Rational bidders choose ex-ante equilibrium bids against the population of - rational and/or naive - opponents they are facing. The number of bidders, the distributions from which their signals are drawn, and whether or not bidders account for the Winner's Curse are common knowledge. For now, consider that naive bidders shade by a given constant, and do not optimize in the same way as the rational bidders do. Later on, we will allow for some degree of sophistication of these naive bidders, and pin down this constant for comparative statics and simulations.

Our model applies to two main cases: one with additive noise $s_i = \mu + \Delta_i + \epsilon_i$ with $\epsilon_i \sim N(0, \sigma_i^2)$, the other with multiplicative noise $s_i = \mu\Delta_i\eta_i$ with $\eta_i \sim LN(0, \sigma_i^2)$. We focus on the additive model, as the multiplicative model can be transformed into the additive model by taking the natural logarithm.[15]

We look for ex-ante equilibrium bids among all the rational bidders. As in Wilson (1998), we consider the case in which the prior is diffuse. In practical terms this means that the signal and the standard deviation of the error capture all the information the bidder has about the true value of the object when making the decision of how much to bid.

Instead of simulating many scenarios we obtain equilibrium constant shading from the first order conditions.[16] This produces accurate estimates quickly and significantly reduces demands on computing power.

Let c_i be the shading applied by bidder i, for $i = 1, 2, ..., n$. That is, she submits a bid $b_i = s_i - c_i$, and consequently, conditional on μ, her bids are distributed $N(\mu + \Delta_i - c_i, \sigma_i^2)$.

[13] The analysis can be extended to three or more groups. However, 2 groups keeps the comparative statistics parsimonious.

[14] Note however that, while the moments may differ, the nature (Normal or Log Normal) of the underlying distribution is the same for all bidders. This is consistent with the auction literature and reflects the fact that the distribution of the error term is related to the characteristics of the auctioned item and not to the characteristics of the bidders.

[15] An objection to a constant additive shading model is that it may appear unrealistic to apply the same additive shading factor, say 2, in cases where the signals are as different as, say 5 and 100. Furthermore, this specification does not guarantee a natural lower bound — zero — for bids. However, while we focus on additive model, it is important to bear in mind that this specification includes the log transformation of the multiplicative model, which does not suffer from any of those problems.

[16] i.e. instead of using Monte Carlo simulations. This process is briefly described in Appendix C

The ex-ante expected profits for bidder i are:

$$
\begin{aligned}
E_{\mu,\epsilon_i,\epsilon_{-i}}[\pi_i] &= E_{\mu,\epsilon_i,\epsilon_{-i}}\left[(\mu_i - b_i)\prod_{j\neq i}\mathbb{1}_{b_i>b_j}\right] = \\
&= E_{\mu,\epsilon_i,\epsilon_{-i}}\left[(\mu+\Delta_i-(\mu+\Delta_i+\epsilon_i-c_i))\prod_{j\neq i}\mathbb{1}_{\mu+\Delta_i+\epsilon_i-c_i>\mu+\Delta_j+\epsilon_j-c_j}\right] = \\
&= E_{\epsilon_i}\left[(c_i-\epsilon_i)\prod_{j\neq i}F_j(\Delta_i-c_i-(\Delta_j-c_j)+\epsilon_i)\right] \qquad (1)
\end{aligned}
$$

where $f_k(\cdot)$ and $F_k(\cdot)$ are the density and cumulative distribution functions of ϵ_k for $k = 1, ..., n$ respectively. It is clear how from the ex-ante perspective, the common component μ cancels out, and therefore the ex-ante expected profits do not depend on μ nor on its prior distribution.[17] The decision variable of bidder i is c_i, the amount she will shade her bid.

The shading factor for bidder i then is the c_i that maximizes (1). A Nash Equilibrium in pure strategies in shading factors is defined as:

Definition 1. *Let there be M bidders, from which $N \leq M$ are rational bidders. Let c_i be the shading applied by each bidder. Fix the naive bidders' shading at $(c_{N+1}, ..., c_M)$. The vector $(c_1, c_2, ..., c_N)$ is an equilibrium if, for a given $(c_{N+1}, ..., c_M)$ and for all i from 1 to N, it holds that:*

$$
c_i \in \arg\max_{\hat{c}} \int_{\mathbb{R}} f_i(\epsilon_i)(\hat{c}-\epsilon_i)\left(\prod_{j\neq i}F_j(\Delta_i-\hat{c}-(\Delta_j-c_j)+\epsilon_i)\right)d\epsilon_i \qquad (2)
$$

The shadings applied by the rational bidders in equilibrium $\{c_i\}_1^N$ are called the Shading Factors (SF).

Lemma 1. *Considering the optimization problem of bidder i, its first order condition is:*

$$
\int_{\mathbb{R}} f_i(\epsilon_i)\left(\prod_{j\neq i}F_j(z_{ij}+\epsilon_i)\right)\left(1-(c_i-\epsilon_i)\sum_{j\neq i}\frac{f_j(z_{ij}+\epsilon_i)}{F_j(z_{ij}+\epsilon_i)}\right)d\epsilon_i = 0 \qquad (3)
$$

where $z_{ij} = \Delta_i - c_i - (\Delta_j - c_j)$.

Proof. The proof can be found in Appendix B. □

Computational methods are used to solve this problem. The steps for the algorithm are as follows:

1. Assume that the initial shading for all the rational bidders is zero.
2. Find the optimal shading from the first order condition (3) for all the rational bidders, assuming all the other rational bidders' are shading according to the previously assumed shading.

[17] Note the importance of the diffuse prior assumption for this result.

3. Update the optimal shading for every rational bidder, using the values found in the previous step.
4. Assess how much has the SF changed compared to the previous value.[18] If some bidder's SF has changed more than a tolerance, iterate from (ii) using the updated shadings found for the rational bidders. If not, these are the chosen Shading Factors.

We implemented this algorithm in both Octave and in Java, with equivalent results.

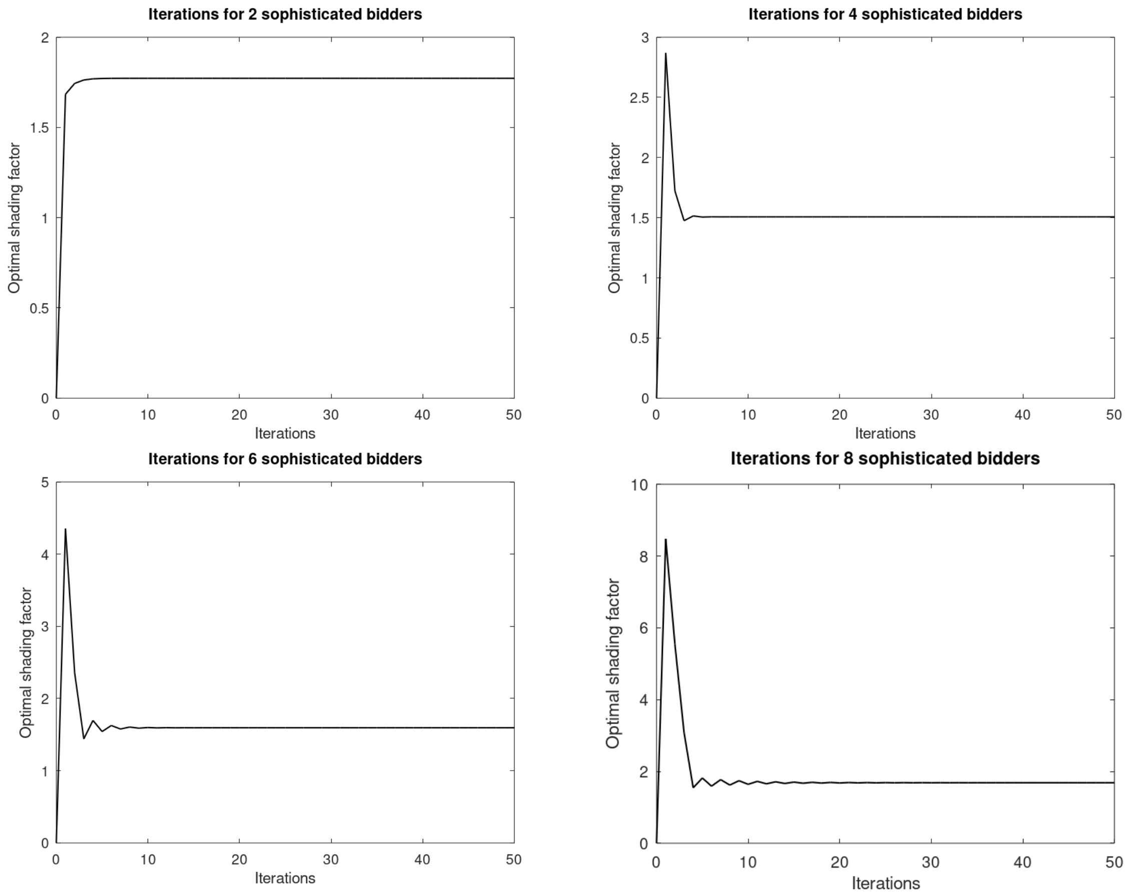

Figure 1: Convergence of optimal shading for all rational and symmetric bidders drawing signals from a standard normal distribution. On the horizontal axis the number of iterations it took, while on the vertical axis, the shading factor chosen in each iteration.

For cases where the standard deviation or the number of bidders is not too large, the algorithm converges relatively fast, in the order of a couple of seconds, or 6 iterations for 2 rational symmetric bidders. Convergence plots are shown in Figure 1.[19]

In Figure 2 we plot the best response functions of two sophisticated bidders. We also plot the 45 degree line to identify the equilibrium. Both are best responding at the Shading Factor of approximately 1.77.

We benchmarked our results to the equilibrium bidding schedule as a function of the signal for the symmetric problem with a diffuse prior. Wilson (1969) was the first

[18] We used as tolerance 10^{-4}. This choice was arbitrary. We saw no further gains in precision to the bias factors and shading factors for the distributions employed in our simulations, but this parameter can be adjusted as required.

[19] We plot convergence on iterations instead of reporting time to convergence, which depends on the particular hardware that is available to us.

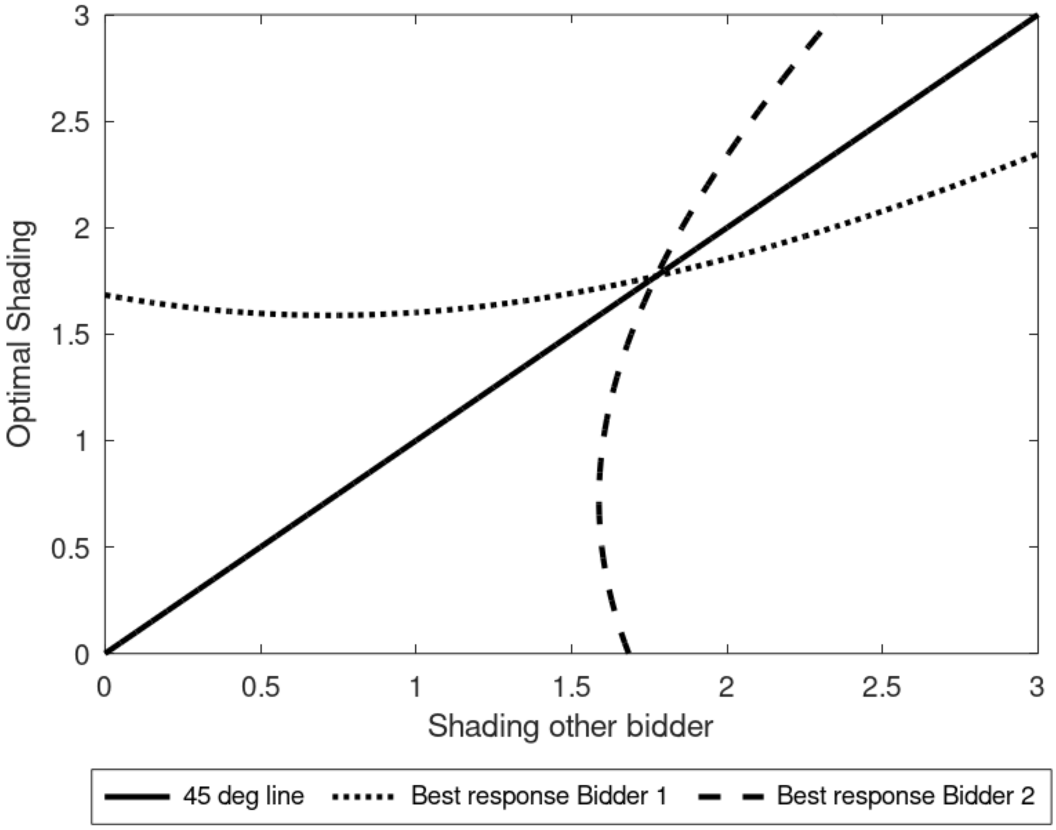

Figure 2: Best response of two symmetric and rational competitors, playing the equilibrium $SF \approx 1.77$.

to provide a solution to the Bayes Nash equilibrium problem with a diffuse prior for 2 symmetric bidders receiving normal additive noise. Wilson (1998), in a footnote, claims to have proof that constant shading is also optimal with more than 2 bidders when the prior is diffuse. Unfortunately, those proofs were not made public.[20] Levin and Smith (1991), and later Hoernig and Fagandini (2018), derived the Bayes Nash equilibrium bidding function for the symmetric problem with n bidders and a diffuse Gaussian prior.[21] Their solution is linear and corresponds to:

$$b(s) = s - \sigma \times \frac{(n(n-1))^{-1} + \int_{\mathbb{R}} w f(w)^2 F(w)^{n-2} dw}{\int_{\mathbb{R}} f(w)^2 F(w)^{n-2} dw} \tag{4}$$

where σ is the standard deviation of each bidder's signal error. The shading factor (SF) coincides exactly with this analytical solution.[22] This encouraging result establishes that the SF is indeed optimal for the symmetric problem where all bidders are rational. It does not, however, indicate whether the SF also produces optimal bids for asymmetric auctions. To verify the SF in asymmetric scenarios as well as scenarios where some of the bidders are naive — for which analytical solutions are not available — we verified our equilibrium with "brute force" Monte Carlo simulations. We did that by generating signals for the all the bidders; then, we applied the SF to all the rivals, and averaged profits for the bidder. Finally, we looked for the shading that gave bidder i the highest expected profits, when considering the others' shadings fixed. This confirmed that the SF was an equilibrium for a wide range of scenarios that may include valuation and information asymmetries, as well as the presence of naive bidders. For the symmetric case, we started with an arbitrary shading (different from the SF) for each bidder, and then iterating, as we did for the FOC, allowed us to obtain the equilibrium SF. The estimates confirmed the SF obtained from the first order condition in (3).[23]

It is important to remark that we do not claim uniqueness of the equilibrium. We do not prove convergence to the true solution, a common problem in research concerning asymmetric first-price auctions in other computational approaches, as pointed out by Hubbard et al. (2013). However, for all cases we tried, the algorithm converges, and that suffices to obtain comparative statics, some of those will be described in Section 5.

4 The Bias Factor

A bidder in a first price sealed bid common value auction should take into account two considerations that point in opposite directions: competition and the Winner's Curse. Increased competition implies that you must bid more aggressively to win the auction. On the other hand, accounting for the Winner's Curse requires bidders to bid more cautiously as the number of bidders increases.

[20] In a private communication with the author we were informed that he does not have those proofs anymore.

[21] Hoernig and Fagandini (2018) also find the solution for the non symmetric, but all rational, equilibrium as a linear function.

[22] Tried for several symmetric and all-rational bidders drawing signals from a standard normal distribution.

[23] On Appendix C we describe the Monte Carlo simulations and show some histograms for the estimates for a couple of symmetric bidder scenarios and also for a couple of scenarios with asymmetric bidders.

The bias factor (BF) introduced by Wilson (1984) — and revisited by Cramton (1995) — enables us to disentangle these two effects.

The BF indicates by how many standard deviations the signal received by the winner exceeds her true value of the auctioned item.

$$BF_i = \frac{E[\varepsilon_i|win]}{\sigma_i} \tag{5}$$

For the symmetric case considered by Wilson, the BF corresponds to the expected signal error conditional on winning divided by the standard deviation.[24] Applying the (additive) bias factor to the signal ensures that the adjusted signal $\hat{s}_i$ is unbiased conditional on winning:

$$\hat{s}_i = s_i - BF_i \times \sigma_i$$
$$E[\hat{s}_i|win] = \mu + \Delta_i$$

Shading signals by the bias factor, i.e. $b_i = \hat{s}_i$, results in zero expected winning profits. Thus, the BF given c_{-i} can be computed by solving

$$\int_{\mathbb{R}} f_i(\epsilon_i)(c_i - \epsilon_i) \left(\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right) d\epsilon_i = 0 \tag{6}$$

Bidders	BF	Prob. Win
2	0.56	0.500
3	0.85	0.333
4	1.03	0.250
5	1.16	0.200
6	1.27	0.167
7	1.35	0.143
8	1.42	0.125

Table 1: Bias factors as found by Wilson (1984) for symmetric bidders. $\epsilon_i \sim N(0, 1)$.

With asymmetric bidders, the bias factor is no longer the highest order statistic of the signal error. As asymmetries impact the severity of the Winner's Curse, the bias factor is now different for each bidder.

We find that the classical results in the literature for the Winner's Curse Capen et al. (1971); Rothkopf (1969); Wilson (1967); Kagel and Levin (2002) hold. The Winner's Curse effect is stronger — and the bias factor is correspondingly larger — when the

[24] Wilson divided by σ just to have a result that was scale-independent.

number of rivals is greater (Table 1), when rivals have an intrinsic firm-specific valuation advantage, and when the bidders' signals are less accurate. We also find that the bias factor increases sharply when some rivals are naive. Intuitively, it is clear that winning against a naive bidder, who does not account for the Winner's Curse, would be "worse news" than winning against a rational bidder; therefore one would expect that a larger adjustment is required to avoid the Winner's curse. However, the sheer magnitude of this effect is unexpected — particularly when there are two or more naive rivals (See Figure 8 in the Appendix).

While the bias factor is by no means a tool to generate optimal bids, it is useful in three very important ways. First, it provides a lower bound for optimal shading. Second, it enables us to isolate the Winner's Curse effect, which helps to explain why bids do not change monotonically in the number of bidders, a fact that might at first seem surprising. We do that by computing the BF for a bidder facing competitors who use the correct SF. The SF, as we mentioned earlier, encompasses two effects, the Winner's Curse and competition. By isolating the Winner's Curse with the BF, we can obtain the competitive effect ($CF = SF - BF$), which describes the trade-off between the gains in case of winning, and the probability of beating the other bidders. This is shown later in Table 2. Third, the bias factor enables us to suggest a plausible candidate for bids submitted by naive bidders; in turn this makes it possible to assess the impact of naive bidders on the bids — and expected profits — of rational bidders. These issues are taken up in Section 5 below.

5 Simulations and Predictions

In this section we first examine the shading factor in a few simple cases to verify whether our methodology produces results that are in accord with the standard results in the literature. Subsequently, we analyze the impact of valuation asymmetries (firm-specific valuation differences), information asymmetries (some bidders receiving a noisier signal than others), and the presence of naive bidders on equilibrium bidding strategies.

5.1 Rational Symmetric Bidders

In this subsection, we consider the symmetric all-rational model to verify consistency with predictions in the literature. Examining the SF, winning profits and the probability of winning for different numbers of bidders (N) and different levels of noise in their signals, confirms that our model reproduces standard results for a population of rational symmetric bidders.

For example, our results corroborate the intuition that the BF is a lower bound to the SF, and that the SF approaches the BF as the number of bidders increases, implying zero expected profits when bidders approach infinity, in accordance with the literature.[25]

Table 2 shows an interesting pattern in the SF for a symmetric all-rational auction when signals are drawn from a standard normal distribution.

[25] For example, the analytical solution Hoernig and Fagandini (2018), coinciding with all our simulations, allows to compute the symmetric values for 50 ($BF = 2.25$, $CF = 0.10$, $SF = 2.34$), 100 ($BF = 2.51$, $CF = 0.07$, $SF = 2.58$), and even 10,000 ($BF = 3.85$, $CF = 0.02$, $SF = 3.88$) bidders.

Bidders	BF	CF	SF	Exp. Winning Profits	Prob. Win
2	0.17	1.60	1.77	1.21	0.50
3	0.43	1.08	1.51	0.66	0.33
4	0.62	0.89	1.51	0.48	0.25
5	0.76	0.79	1.55	0.38	0.20
6	0.87	0.73	1.60	0.33	0.17
7	0.96	0.68	1.64	0.29	0.14
8	1.04	0.65	1.69	0.26	0.13
9	1.49	0.24	1.73	0.24	0.11
10	1.54	0.24	1.73	0.24	0.10

Table 2: Correction for symmetric bidders with $\epsilon_i \sim N(0,1)$.

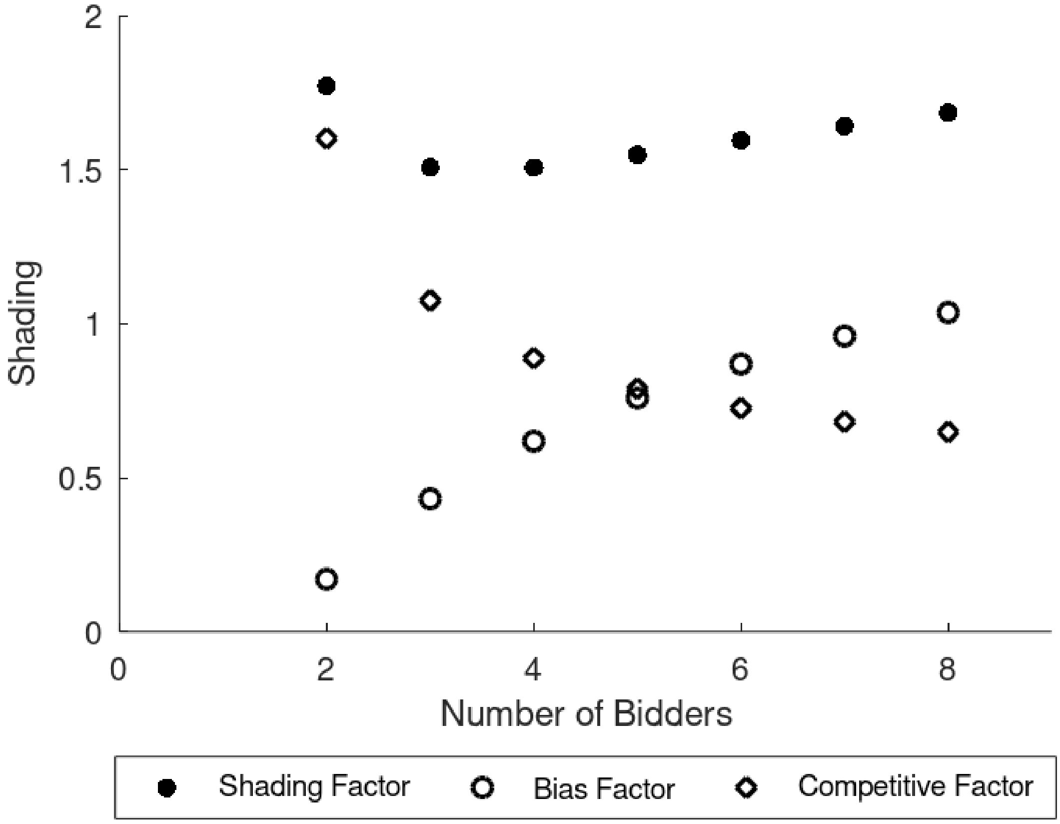

Figure 3: Decomposition of the shading factor in the Winner's Curse and the competitive effect. Symmetric bidders with $\epsilon_i \sim N(0,1)$.

As the number of bidders increases, the probability of winning and expected winning profits monotonically both decline, as they should. However, there appears to be an anomaly in the SF: the amount by which bidders shade their bids first decreases when

N goes from 2 to 4, and then steadily increases from there on.[26] As suggested earlier in this paper, this non-monotonic pattern is the result of two effects working in opposite directions. As a bidder faces more rivals, (i) the Winner's Curse looms larger, which requires a larger adjustment of her signal, and (ii) competition increases, which requires her to bid closer to her signal, *i.e.* to sacrifice margin for a higher probability of winning. Initially, the competition effect dominates. With 2 bidders, low rivalry results in bidders shading their bids substantially; as the number of bidders gets larger, increasing rivalry leads them to bid closer to their valuation. With more than 4 bidders, however, the need to compensate for the Winner's Curse becomes the dominant effect, driving lower bids as the number of bidders increases.

Table 3 shows the impact of all bidders having less accurate signals, for a fixed number of bidders. As expected, the adjustment implied by the shading factor increases. It is interesting to note that in equilibrium profits increase when bidders have less accurate signals.

σ	SF	Winning Profits
0.5	0.75	0.24
0.7	1.05	0.33
0.9	1.36	0.43
1.1	1.66	0.53
1.3	1.96	0.62
1.5	2.26	0.72

Table 3: SF for 4 rational bidders. $\epsilon_i \sim N(0, \sigma)$.

This result is consistent with a general finding that the efficiency of auctions increases as bidders' information (symmetrically) improves. For example, in a very different model Milgrom and Weber (1982b) analyze various information policies and show that auctioneer revenues increase when bidders are provided better information.

5.2 Valuation Asymmetries

We consider two groups of bidders. The valuation of Group I is held constant, whereas the valuation advantage parameter (Δ_{II}) of Group II is varied. Throughout this subsection, both groups receive signals with errors generated $N(0,1)$.

First we examine the case with two sophisticated bidders. Starting from the symmetric case where $\Delta_{II} = 0$, we progressively increase the mean of the distribution for Group II (i.e. Δ_{II}). Remember that Δ_{II} is assumed to be common knowledge.

When there is no valuation asymmetry, both bidders obviously shade the same. As the valuation asymmetry increases we observe that the bidder with the lower valuation is shading less while the bidder with the higher valuation is shading more. The valuation advantage provides space for the high-valuation bidder to increase profits by shading

[26] This finding was confirmed also with Monte Carlo simulations to rule out potential coding issues.

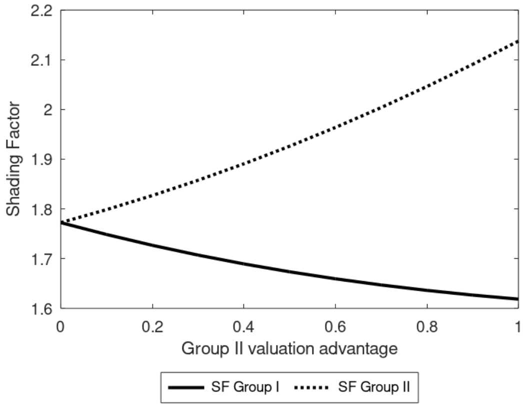

(a) 1 vs. 1

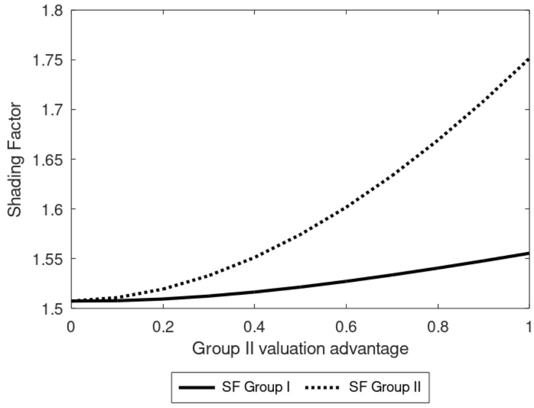

(b) 2 vs. 1

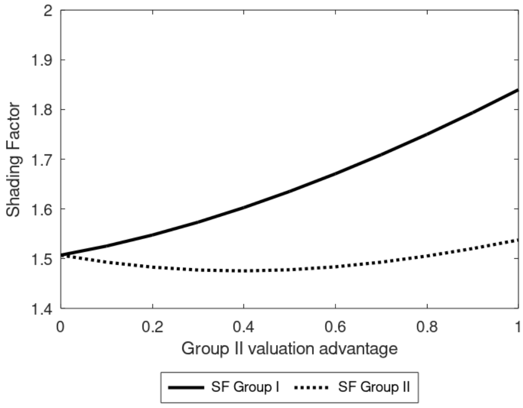

(c) 2 vs. 2

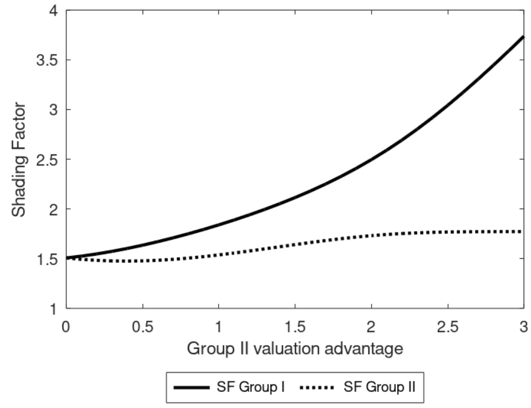

(d) 2 vs. 2, larger valuation gaps

Figure 4: Valuation asymmetry between two groups, Group I (low valuation) and Group II (high valuation)), of sophisticated bidders.

more without reducing too much her probability of winning. The low-valuation bidder, who knows that his rival is shading more, thus reducing the risk of the Winner's Curse, is able to bid more aggresively, as can be seen from the decreasing SF.

The picture changes dramatically when we have multiple high valuation bidders. In Figure 4c and 4d we observe that Group I always shades more (i.e. bids less) than Group II, and increasingly so as the valuation asymmetry increases. This makes sense: competition amongst high-valuation bidders in Group II requires low valuation bidders in Group I to bid prudently to avoid the Winner's Curse. We also note a very interesting phenomenon: the pattern for Group II is non-monotonic in its valuation advantage. Initially, for small valuation asymmetries the shading factor declines (reaching a minimum at about $0.4 \times \sigma$) and then increases. As can be seen in Figure 4d, when valuation differences become very large the Shading Factor converges to the value for an auction with only two high-valuation bidders (about 1.77). The increasing valuation gap combined with vigorous competition within Group II requires Group I bidders to bid more cautiously to avoid the Winner's Curse. Put simply, when valuation differences become large, bidders in Group I become largely irrelevant, approaching a symmetric scenario with only high-valuation bidders.

Hubbard and Paarsch (2009) study a related but different problem, viz. bidding behavior in procurement auctions with private valuations when the auctioneer has a preference for qualified firms and follows a policy of scaling down their bids for the purpose of bid evaluation. This policy has three effects: first, preferred firms may inflate their bids yet still win the auction; second, non-preferred firms may bid more aggressively, and third, participation may be affected as well. These interesting results are consistent with auction theory. Since asymmetries lower the auctioneer's revenue, a policy that artificially amplifies a bidder's advantage is counterproductive. In horse races, the favorite horse gets a handicap, not a head start. Similarly in auctions, an optimal policy discriminates against favored bidders to neutralize symmetric advantages.

5.3 Signal Quality

Information asymmetries have been widely studied in the context of bidding for oil and gas leases on the Outer Continental Shelf, among others. As Hendricks and Porter (2007) put it:

> *Oil and gas leases are classified into two categories. Wildcat tracts are located in previously unexplored areas. Prior to a wildcat auction, firms are allowed to conduct seismic studies, but they are not permitted to drill any exploratory wells. The seismic studies provide noisy, but roughly equally informative signals about the amount of oil and gas on a lease. We argue that wildcat auctions are likely to satisfy the symmetry assumption on the signal distribution. Drainage leases are adjacent to wildcat tracts where oil and gas deposits have been discovered previously. Firms that own adjacent tracts possess drilling information that makes them better informed about the value of the drainage tract than other firms, who are likely to have access only to seismic information. We argue that these auctions can be modeled by assuming one bidder has a private, informative signal and all other bidders have no private information.*

We assume that while bidders in Group I receive estimates with error $N(0, 1)$, bidders on Group II receive signals with error $N(0, \sigma)$. To focus on the effect of the quality of the signal only, we assume that $\Delta_I = \Delta_{II} = 0$. Starting with $\sigma = 1$, we progressively improve Group II's signal quality — decreasing the standard deviation on their signal's error — to study the impact of an increasing information asymmetry. Note in the horizontal axis, that a decreasing value implies a better signal quality for Group II.

Figure 5a shows that for the case of one less-informed bidder competing against one better-informed rival, both bidding schedules roughly coincide. In equilibrium, the better-informed bidder shades her bid just enough to get close to her less-informed rival. Hendricks et al. (1994) obtain a similar result in a somewhat different common value auction where one uninformed bidder, who only observes a public signal, competes with one informed bidder who also observes a private signal of the true value.[27]

However, this result only holds in this one particular case, where one less-informed bidder is bidding against one better-informed rival. In all other cases, where there are more than two bidders, the bidding schedules no longer coincide, as can be seen in Figures 5b and 5c. Hendricks et al. (1994) argue that in the case studied — bidding for OCS oil and gas leases — informed bidders collude so that they can effectively be thought

[27] See also Porter (1995), and Hendricks and Porter (2007).

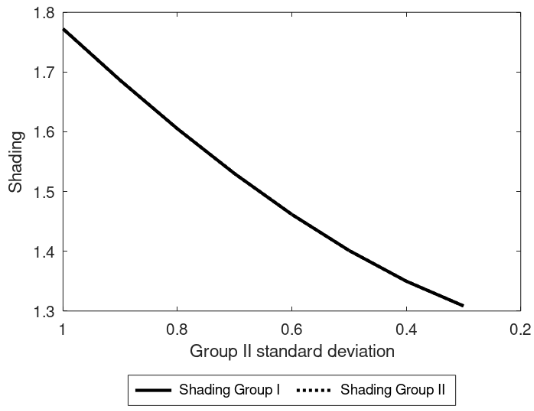

(a) 1 vs. 1

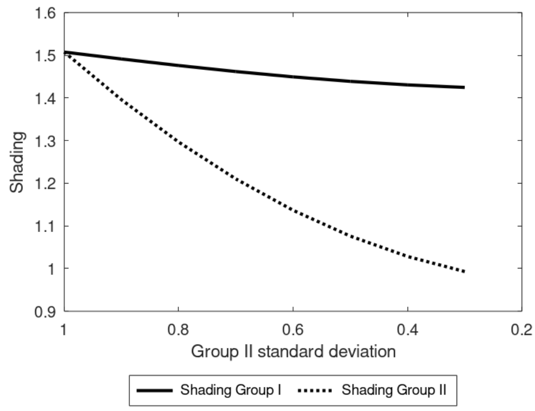

(b) 2 vs. 1

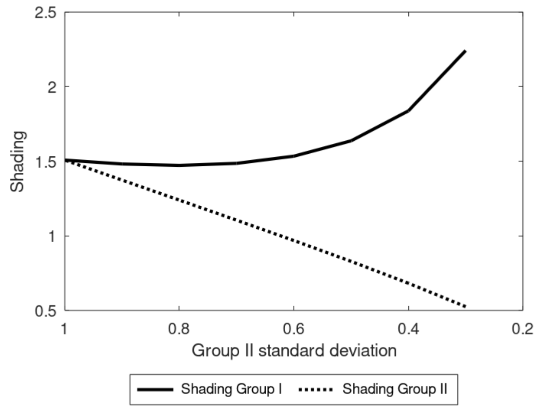

(c) 1 vs. 2

Figure 5: Information asymmetry between two groups, Group I (noisier estimate) and Group II (better estimate), of sophisticated bidders.

of as a single informed bidder. As for uninformed bidders, Porter (1995) points out that modelling a single uninformed bidder is irrelevant, as the equilibrium distribution function will be concerned only with the distribution of the highest bid among of the uninformed bidders. The author correctly identifies that the strategy of the informed bidders depend on the number of uninformed bidders.

In our setup, the number of uninformed bidders is unimportant *only if they have no private information at all*, which is a strong assumption. In our model, less-informed bidders do have a private signal of the true value, albeit a less accurate one, and therefore the number of less-informed bidders does matter, as can be seen from the different bidding schedules in Figure 5b.

In Figure 5c we see the impact of multiple better-informed bidders: increasingly aggressive bids from better-informed rivals force the poorly informed bidder to bid very cautiously.

Figure 6 shows that expected profits of a less-informed bidder who competes with better-informed rivals decline precipitously as the latter's information advantage increases. Note that they are not zero, however. This because in our model, the less-informed bidder does have a private signal of the true value, albeit a less accurate one.

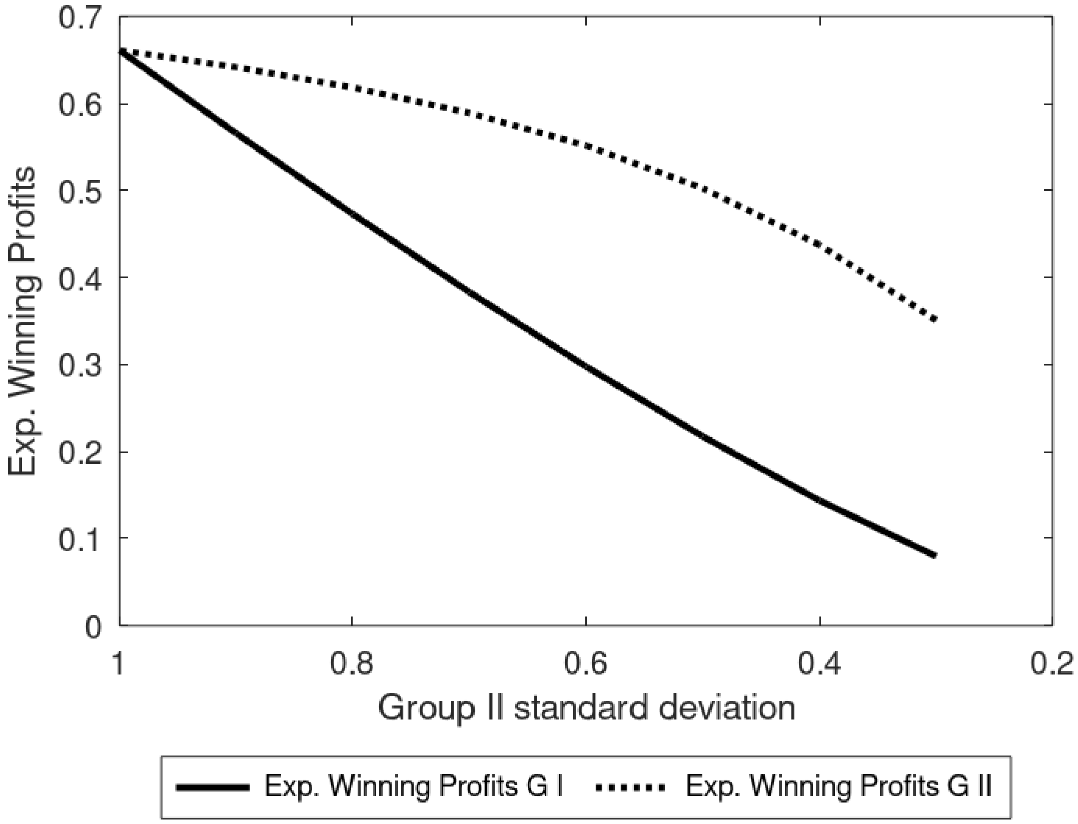

Figure 6: 1 bidder in G I (noisier estimate) vs. 2 bidders in G II (better estimate)

5.4 Naive Bidders

In this section we illustrate the effects of naive bidders on equilibrium bidding and expected profits of a sophisticated bidder. To isolate the impact of naive bidders, we assume no valuation or information advantages for any bidder.

As defined in this paper, naive bidders are aware that they need to shade somehow their bids in order to obtain positive profits. Evidence Kagel and Levin (2002) suggests that even experienced bidders use various rules of thumb. While our model allows us to use any shading rule for naive bidders we decided to endow naive bidders with a fair amount of sophistication. Specifically, we assume that they correctly account for the competitive effect, they are naive only in the sense that they do not adjust for the Winner's Curse. That is, naive bidders shade their bids by the Competitive Factor $CF \equiv SF - BF$.

	Number of Naive Rivals					
Number of Rivals	0	1	2	3	4	5
1	0.61	0.52				
2	0.22	0.15	0.10			
3	0.12	0.08	0.04	0.02		
4	0.08	0.05	0.02	0.01	0.00	
5	0.06	0.03	0.02	0.01	0.00	0.00

Table 4: Expected profits for sophisticated bidders. $\epsilon_i \sim N(0, 1)$.

Number of Rivals	CF
1	1.60
2	1.08
3	0.89
4	0.79
5	0.73

Table 5: Shading applied by the naive bidders, i.e. the Competitive Factor CF, in the scenarios of Table 4.

For a sophisticated bidder, facing naive rivals is bad news. Table 4 shows the different dimensions in which naive rivals can impact the profits of sophisticated bidders, and the naive rivals' margin.

In each column, the number of naive rivals is kept constant while the total number of competitors — hence the number of sophisticated bidders — varies. Moving down in any column shows the competitive pressure of rational competitors.

Moving along the diagonal shows competitive pressure of naive competitors: profits vanish *extremely* fast when their number increases. The devastating impact of naive rivals is due to the fact that they ignore the Winner's Curse and therefore shade *less* when they are facing more competitors, as can be seen in table 5.

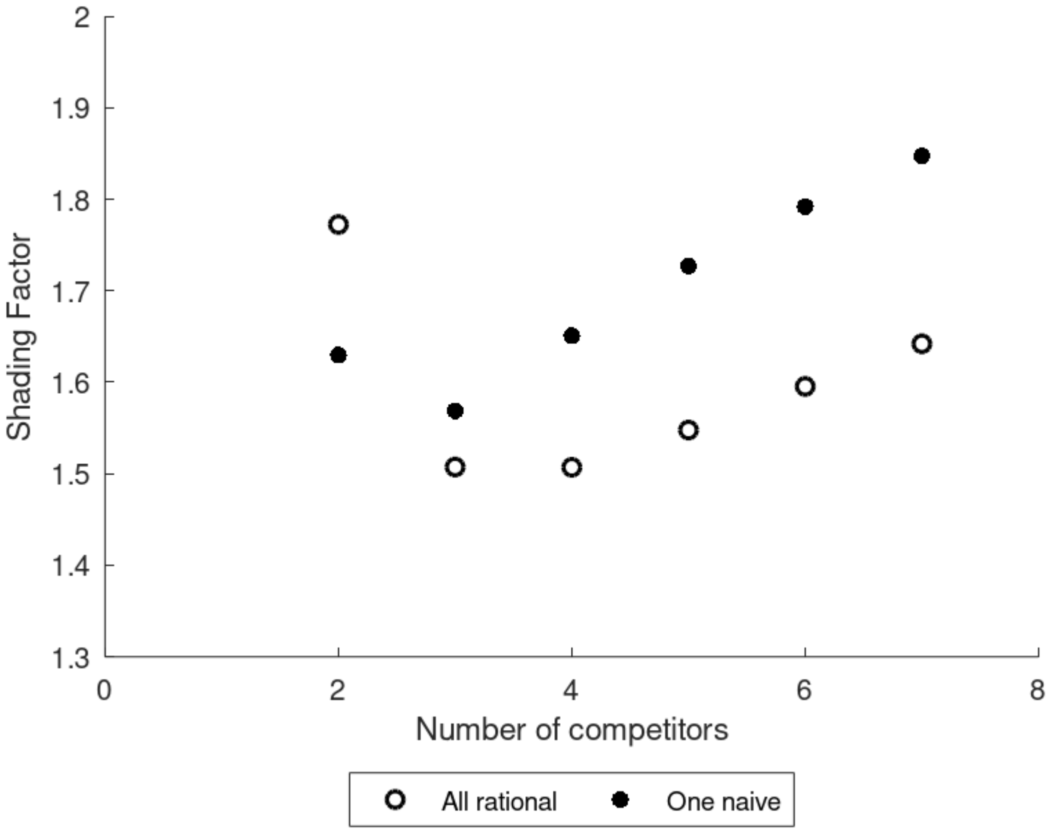

Figure 7: Optimal shading when all competitors are rational *vs.* one is naive.

Figure 7 shows Shading Factors for two cases: when all rivals are rational (square marks) and when one of the rivals is naive (round marks). When facing only one rival, one should bid more aggressively (*i.e.* shade less) when this rival is naive instead of rational. When facing only one rival, concerns about the Winner's Curse — which in this case is relatively weak — are dominated by the more aggressive bid from a naive rival. With more rivals, the opposite holds: concerns about the Winner's Curse dominate the impact of increased competition. This is because a naive bidder, as shown in Table 5, bids more aggressively when facing more competitors; in turn, this aggressive bidding behavior aggravates the Winner's Curse, requiring a more cautious bid to guard against it. Thus, as Figure 7 shows, failing to account for the presence of a naive bidder results in underbidding only in one case, when facing one opponent, and overbidding in all the other cases.

6 Conclusions and Implications

In this paper, we compute a shading factor (SF) to obtain optimal bids in first price sealed bid common value auctions. The SF is computed ex-ante of receiving a signal, does not require a bounded support of either signals or bids, allow for differences in the accuracy of bidders' estimates, as well as firm-specific valuation differences. Furthermore, the SF also allow for (a subset of) naive bidders, who fail to account for the Winner's Curse.

We find that the SF generates the same bid as the Bayes Nash equilibrium with n symmetric bidders and a diffuse Gaussian prior found in Hoernig and Fagandini (2018). For asymmetric scenarios, the values generated by the SF are confirmed by "brute force" Monte Carlo simulations. These results validate the proposed approach and confirm that it does indeed reliably compute optimal bids for a wide range of scenarios, that may include valuation and information asymmetries, as well as the presence of naive bidders.

We also generalize Wilson's (1984) Bias Factor (BF) to obtain a measure of the Winner's Curse effect, allowing us to disentangle the SF's shading in two parts: the Winner's Curse effect, and the competitive effect. Even though the interplay of different dimensions of bidder heterogeneity may lead to surprising shadings that are not monotonic in the number of bidders, all results can be understood intuitively by analyzing how those two effects which work in opposite directions, affect bids.

The BF also enables us to suggest a plausible candidate for bids submitted by naive bidders and to assess their impact on the expected profits of sophisticated bidders. Our results show that this impact is devastating. Hence, a critical task in real life bidding problems is to correctly gauge the level of sophistication of one's competitors. Absent specific information, it may be better to underestimate their capacities — and shade one's bid accordingly — rather than overestimate them and fall victim to the Winner's Curse. This, of course, opens a Pandora's box of tactical opportunities, since it would be advantageous for a sophisticated bidder to be *perceived* as naive. Given that bidders in real life do not know each other's type (rational or naive) with certainty, this opens up an interesting line of inquiry into the incentives for sophisticated bidders to pose as naive. Savvy rational bidders may wish to cultivate an image of being unsophisticated!

Declarations

Funding Paulo Fagandini: Fundação para a Ciência e a Tecnologia FCT, grant SFRH/BD/105669/2015
Conflict of Interest/Competing interests The authors declare to have no conflict of interest.
Availability of data and material Octave code available on `https://tinyurl.com/wrve87dd` and `https://tinyurl.com/yc39rbzx`

References

Armantier, O. and E. Sbai (2006). Estimation and comparison of treasury auction formats when bidders are asymmetric. *Journal of Applied Econometrics 21*(6), 745–779.

Bajari, P. (1998). Econometrics of sealed bid auctions. In *Proceedings of the Business and Economic statistics section of the American Statistical Association*, pp. 41–49.

Bajari, P. and A. Hortaçsu (2005). Are structural estimates of auctions models reasonable? evidence from experimental data. *Journal of Political Economy 113*(4), 703–741.

Campo, S., I. Perrigne, and Q. Vuong (2003). Asymmetry in first-price auctions with affiliated private values. *Journal of Applied Econometrics 18*(2), 179–207.

Capen, E. C., R. V. Clapp, W. M. Campbell, et al. (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology 23*(6), 641–653.

Compte, O. and A. Postlewaite (2012). Simple auctions. Working paper.

Cox, J. C. and R. M. Isaac (1984). In search of the winners' curse. *Economic Inquiry 22*(4), 579–592.

Cramton, P. C. (1995). Money out of thin air: The nationwide narrowband pcs auction. *Journal of Economics & Management Strategy 4*, 267–343.

Crawford, P. B. (1970). Texas offshore bidding patterns. *Journal of Petroleum Technology 22*(3), 283–289.

Crawford, V. P. and N. Iriberri (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica 75*, 1721–1770.

David, A. K. (1993). A bidding model of perfect competition. In *IEE proceedings-C*, pp. 421–426.

Dyer, D., J. H. Kagel, and D. Levin (1989). A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysys. *The Economic Journal 99*(394), 108–115.

Eyster, E. and M. Rabin (2005). Cursed equilibrium. *Econometrica 73*, 1623–1672.

Fabra, N., N.-H. von der Fehr, and D. Harbord (2006). Designing electricity auctions. *The RAND Journal of Economics 37*(1), 23–46.

Friedman, L. (1956). A competitive-bidding strategy. *Operations Research 4*(1), 104–112.

Hendricks, K. and R. H. Porter (2007). Chapter 32 an empirical perspective on auctions. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization*, Volume 3 of *Handbook of Industrial Organization*, pp. 2073 – 2143. Elsevier.

Hendricks, K., R. H. Porter, and C. A. Wilson (1994). Auctions for Oil and Gas Leases with an Informed Bidder and a Random Reservation Price. *Econometrica 62*(6), 1415–1444.

Heumann, T. (2019). An ascending auction with multi-dimensional signals. *Journal of Economic Theory 184*, 104938.

Hoernig, S. and P. Fagandini (2018). An exact solution to the multiple-bidder pure common value first-price sealed bid auction with gaussian noise and diffuse prior. Nova SBE mimeo.

Hong, H., H. J. Paarsch, and P. Xu (2013). On the asymptotic distribution of the transaction price in a clock model of a multi-unit, oral, ascending-price auction within the common-value paradigm. *The RAND Journal of Economics 44*(4), 664–685.

Hortaçsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *Journal of Political Economy 118*(5), 833–865.

Hortaçsu, A. and S. L. Puller (2008). Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics 39*(1), 86–114.

Hubbard, T. P., R. Kirkegaard, and H. J. Paarsch (2013). Using economic theory to guide numerical analysis: Solving for equilibria in models of symmetric first-price auctions. *Computational Economics 42*, 241–266.

Hubbard, T. P. and H. J. Paarsch (2009). Investigating bid preferences at low-price, sealed-bid auctions with endogenous participation. *International Journal of Industrial Organization 27*(1), 1–14.

Hubbard, T. P. and H. J. Paarsch (2014). Chapter 2 - on the numerical solution of equilibria in auction models with asymmetries within the private-values paradigm. In K. Schmedders and K. L. Judd (Eds.), *Handbook of Computational Economics Vol. 3*, Volume 3 of *Handbook of Computational Economics*, pp. 37–115. Elsevier.

Kagel, J. H. and D. Levin (2002). *Common value auctions and the winner's curse.* Princeton University Press.

Klemperer, P. (2004). *Auctions: Theory and Practice.* Princeton University Press.

Krishna, V. (2010). *Auction Theory* (Second Edition ed.). San Diego: Academic Press.

Levin, D. and J. L. Smith (1991). Some evidence on the winner's curse: Comment. *The American Economic Review 81*, 370–375.

Lorentziadis, P. (2012). Optimal bidding in auctions of mixed populations of bidders. *European Journal of Operational Research 217*(3), 653–663.

Lorentziadis, P. (2016). Optimal bidding in auctions from a game theory perspective. *European Journal of Operational Research 248*, 347–371.

Ma, X., F. Wen, Y. Ni, and J. Liu (2005). Towards the development of risk-constrained optimal bidding strategies for generation companies in electricity markets. *Electric Power Systems Research 73*(3), 305–312.

Marshall, R. C., M. J. Meurer, J.-F. Richard, and W. Stromquist (1994). Numerical analysis of asymmetric first price auctions. *Games and Economic Behavior 7*(2), 193–220.

Milgrom, P. and R. J. Weber (1982a). The value of information in a sealed-bid auction. *Journal of Mathematical Economics 10*(1), 105–114.

Milgrom, P. R. (2004). *Putting auction theory to work.* Cambridge University Press.

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

Ortega Reichert, A. (1968). *Models for Competitive Bidding Under Uncertainty.* Ph. D. thesis, Department of Operations Research and Department of Statistics, Standford University.

Porter, R. H. (1995). The Role of Information in U.S. Offshore Oil and Gas Lease Auction. *Econometrica 63*(1), 1–27.

Rothkopf, M. H. (1969). A model of rational competitive bidding. *Management Science 15*(7), 362–373.

Rothkopf, M. H. (1980). Equilibrium linear bidding strategies. *Operations Research 28*(3), 576–583.

Rothkopf, M. H. (2007). Decision analysis: The right tool for auctions. *Decision Analysis 4*, 167–172.

Salant, D. J. (2014). *A Primer on Auction Design, Management, and Strategy.* MIT Press.

Shachat, J. and L. Wei (2012). Procuring commodities: First-price sealed-bid or english auctions. *Marketing Science 31*(2), 317–333.

Smiley, A. K. (1979). *Competitive Bidding Under Uncertainty: The Case of Offshore Oil.* Ballinger Pub. Co.

Stahl, D. O. and P. W. Wilson (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization 25*, 309–327.

Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior 10*, 218–254.

Thaler, R. H. (1988). Anomalies: The winner's curse. *Journal of Economic Perspectives 2*, 191–202.

Vives, X. (2011). Strategic supply function competition with private information. *Econometrica 79*(6), 1919–1966.

Wang, J. J. and J. F. Zender (2002). Auctioning divisible goods. *Economic theory 19*(4), 673–705.

Wen, F. and A. K. David (2001). Optimal bidding strategies and modeling of imperfect information among competitive generators. *Power Systems, IEEE Transactions on 16*(1), 15–21.

Wilson, R. B. (1967). Competitive bidding with asymmetric information. *Management Science 13*(11), 816–820.

Wilson, R. B. (1969). Competitive bidding with disparate information. *Management Science 15*(7), 446–448.

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

Wilson, R. B. (1984). Lecture notes for course 310b analytical foundations of pricing strategy, stanford business school.

Wilson, R. B. (1998). Sequential equilibria of asymmetric ascending auctions: The case of log-normal distributions. *Economic Theory 12*(2), 433–440.

Appendix A Effect of naive competitors on the Winner's Curse

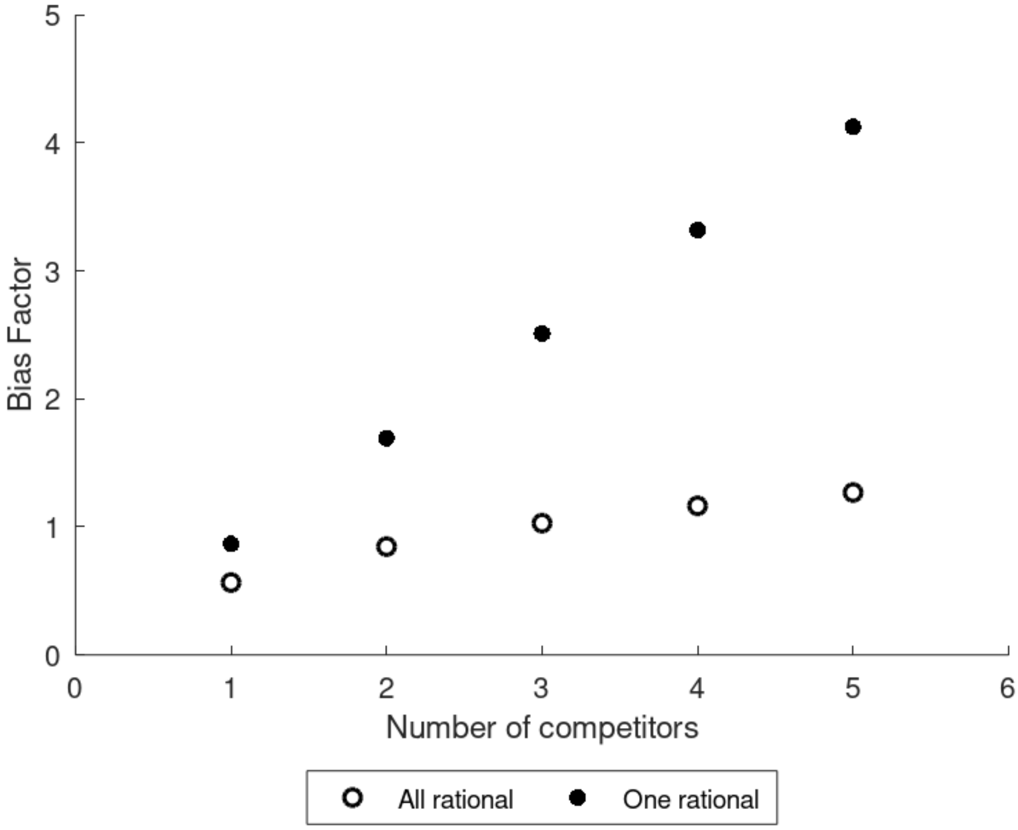

Figure 8: Bias factor for an increasing number of bidders. In one group only one of the bidders is rational (solid) and in the other, all bidders are rational (dashed).

Appendix B Proof of Lemma 1

As stated in the main text, the Shading Factors satisfy the following:

$$c_i \in \arg\max_{\hat{c}} \int_{\mathbb{R}} f_i(\epsilon_i)(\hat{c} - \epsilon_i) \left(\prod_{j \neq i} F_j(\Delta_i - \hat{c} - (\Delta_j - c_j) + \epsilon_i) \right) d\epsilon_i \tag{7}$$

Take the derivative with respect to c_i and set it equal to zero:

$$\frac{\partial}{\partial c_i} \left[\int_{\mathbb{R}} f_i(\epsilon_i)(c_i - \epsilon_i) \left(\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right) d\epsilon_i \right] = 0 \tag{8}$$

Noting that $f_i(\epsilon_i)$ does not depend on c_i, solve:

$$\int_{\mathbb{R}} f_i(\epsilon_i) \frac{\partial}{\partial c_i} \left[(c_i - \epsilon_i) \left(\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right) \right] d\epsilon_i = 0 \tag{9}$$

The derivative corresponds to the addition of two terms:

$$\frac{\partial}{\partial c_i} [(c_i - \epsilon_i)] \left(\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right) + \dots \tag{10}$$

$$(c_i - \epsilon_i) \frac{\partial}{\partial c_i} \left[\left(\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right) \right] \tag{11}$$

It is clear that (10) is equal to:

$$\left(\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right)$$

Finally, to solve (11) let's focus for now on the term:

$$\frac{\partial}{\partial c_i} \left[\prod_{j \neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \right] =$$

$$- \sum_{j \neq i} \left(f_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i) \prod_{k \neq j,i} F_k(\Delta_i - c_i - (\Delta_k - c_k) + \epsilon_i) \right)$$

Multiplying and dividing by $F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)$ we obtain

$$- \sum_{j \neq i} \left(\frac{f_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)}{F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)} \prod_{k \neq i} F_k(\Delta_i - c_i - (\Delta_k - c_k) + \epsilon_i) \right)$$

Note that the product does not depend on j, so it can go out of the summation:

$$-\left(\prod_{k\neq i} F_k(\Delta_i - c_i - (\Delta_k - c_k) + \epsilon_i)\right)\left(\sum_{j\neq i}\frac{f_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)}{F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)}\right)$$

So the third term, replacing k with j, corresponds to:

$$-(c_i - \epsilon_i)\left(\prod_{j\neq i} F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)\right)\left(\sum_{j\neq i}\frac{f_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)}{F_j(\Delta_i - c_i - (\Delta_j - c_j) + \epsilon_i)}\right)$$

Finally, to obtain the full expression, apply the change of variables $z_i = \Delta_i - c_i - (\Delta_j - c_j)$, to obtain:

$$\left(\prod_{j\neq i} F_j(z_{ij} + \epsilon_i)\right) - (c_i - \epsilon_i)\left(\prod_{j\neq i} F_j(z_{ij} + \epsilon_i)\right)\left(\sum_{j\neq i}\frac{f_j(z_{ij} + \epsilon_i)}{F_j(z_{ij} + \epsilon_i)}\right)$$

Replacing back into the integral in (9) we have:

$$\int_{\mathbb{R}} f_i(\epsilon_i)\left(\prod_{j\neq i} F_j(z_{ij} + \epsilon_i)\right)\left(1 - (c_i - \epsilon_i)\sum_{j\neq i}\frac{f_j(z_{ij} + \epsilon_i)}{F_j(z_{ij} + \epsilon_i)}\right) d\epsilon_i = 0 \qquad (12)$$

This is the first order condition.

Appendix C Monte Carlo method

We benchmarked the *Shading Factor* in some cases using Monte Carlo simulations, as set out below:

Generate signals randomly for the number of bidders, in a matrix of size $n \times m$. Each column represents a bidder, and each row a draw. Define an interval in which to look for the equilibrium. Apply the shading to the competitors and test shading within that interval and up to 2 decimal places, and pick the one with the highest expected profit.

This procedure has limitations. First, we are taking a fixed number of shadings within a bounded and fixed interval; however, given the results obtained using the first order condition, we can use intervals easily wide enough to give reasonable assurances that it includes the optimum. Second, this procedure is inefficient, as it tries every single possibility within that set. However, this procedure also has important advantages. It can deal with any trouble embedded functions might have and it is agnostic about the curvature of the expected profits functions (does not assume differentiability or a unique maximum within the interval).

We plot the histograms of the estimates of specific simulations.[28] Each SF in the sample, consists on the equilibrium shading that maximizes expected profits using 100.000 hypothetical draws for each bidder. We do this 500 times to obtain the sample of optimal factors we include in the histograms shown below. We use a tolerance of 0.01 for each optimal factor and a maximum of iterations (over best responses) of 50 times. In all these samples, the algorithm converged (within tolerance) every time.

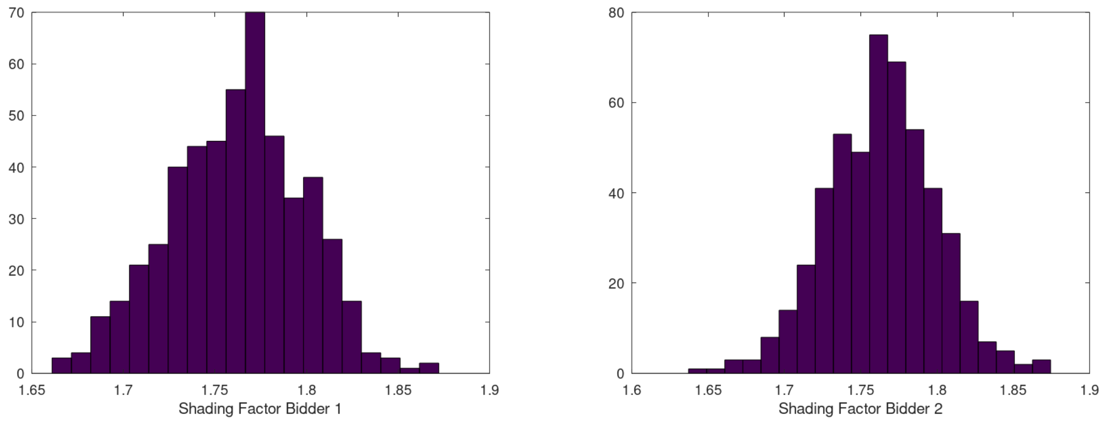

(a) Bidder 1, mean $\widehat{SF}_1 = 1.7621$.

(b) Bidder 2, mean $\widehat{SF}_2 = 1.7632$.

Figure 9: Two sophisticated bidders, both with signals $N(0, 1)$. The SF is 1.77 for both.

[28]Code available at `https://tinyurl.com/yc39rbzx`.

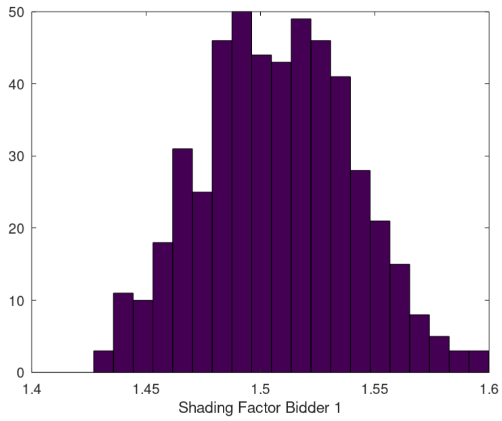

(a) Bidder 1, mean $\widehat{SF}_1 = 1.5075$

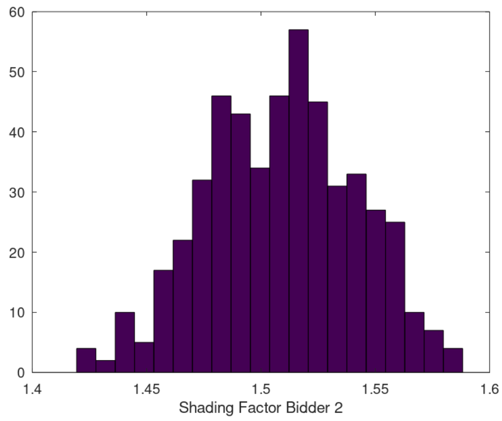

(b) Bidder 2, mean $\widehat{SF}_2 = 1.5091$

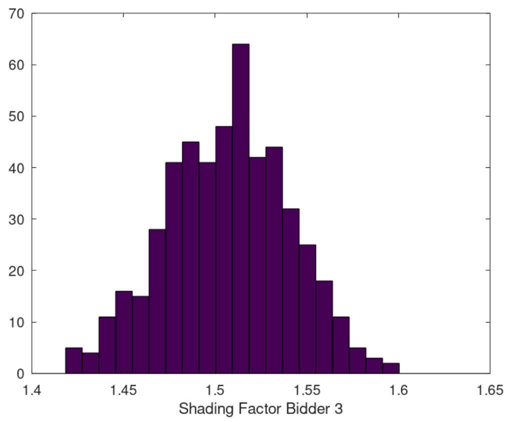

(c) Bidder 3, mean $\widehat{SF}_3 = 1.5069$

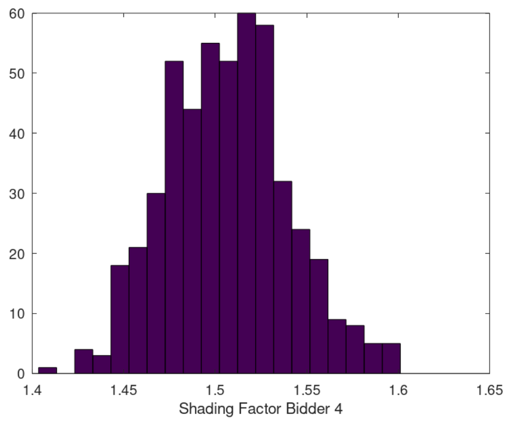

(d) Bidder 4, mean $\widehat{SF}_4 = 1.5066$

Figure 10: Four sophisticated bidders, with signals $N(0,1)$. The $SF = 1.51$ for all.

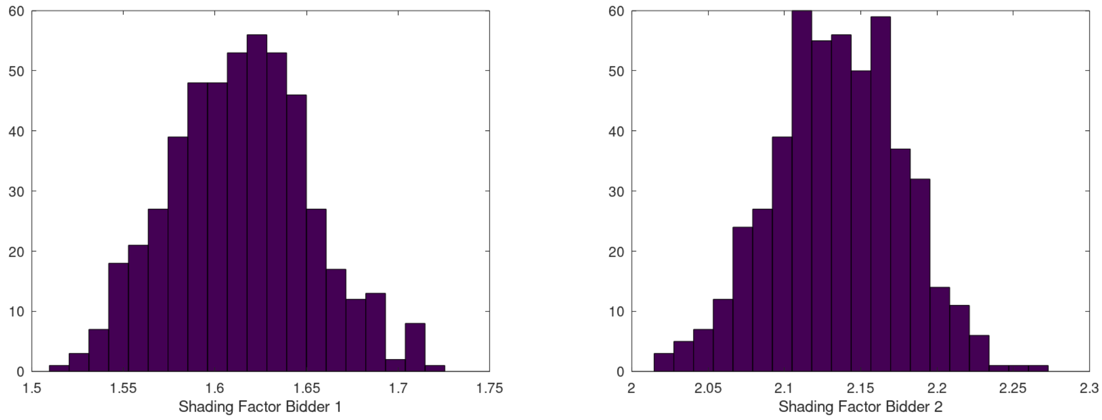

(a) Bidder 1, mean $\widehat{SF}_1 = 1.6141$ (b) Bidder 2, mean $\widehat{SF}_2 = 2.1347$

Figure 11: Two sophisticated bidders, one with signals $N(0,1)$, the other with signals $N(1,1)$. The true corrections are $SF = 1.62$ and $SF = 2.14$ respectively.

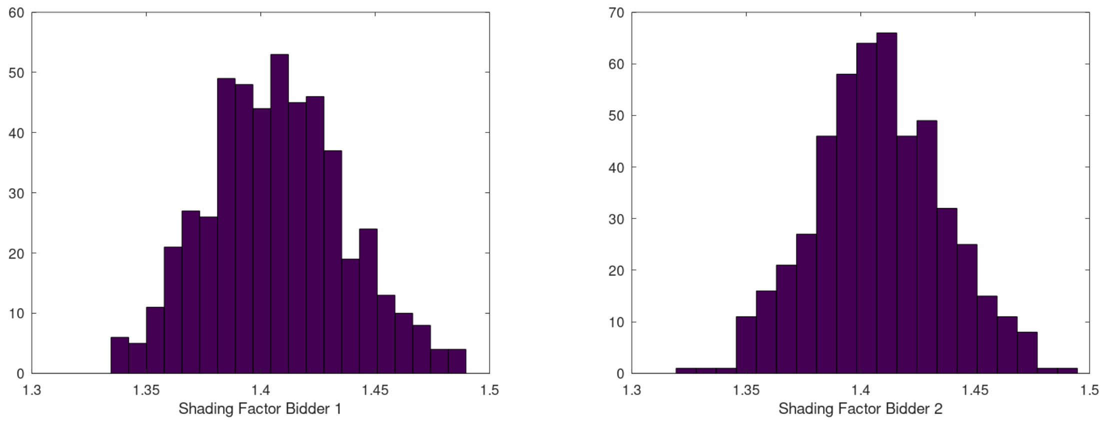

(a) Bidder 1, mean $\widehat{SF}_1 = 1.4069$ (b) Bidder 2, mean $\widehat{SF}_2 = 1.4084$

Figure 12: Two sophisticated bidders, one with signals $N(0,1)$, the other with signals $N(0,0.5)$. The true correction is 1.40 for both.