http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–18
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Soft Floors in Auctions

Robert Zeithammer[a]

[a] UCLA Anderson School of Management, University of California, Los Angeles, Los Angeles, CA 90095

Contact: rzeitham@ucla.edu, http://orcid.org/0000-0002-8388-247X (RZ)

Received: July 6, 2017
Revised: March 16, 2018; May 28, 2018
Accepted: June 30, 2018
Published Online in Articles in Advance: June 11, 2019

https://doi.org/101287/mnsc.2018.3164

Abstract. Several of the auction-driven exchanges that facilitate programmatic buying of internet display advertising have recently introduced "soft floors" in addition to standard reserve prices (called "hard floors" in the industry). A soft floor is a bid level below which a winning bidder pays his own bid instead of paying the second-highest bid as in a second-price auction most ad exchanges use by default. This paper characterizes soft floors' revenue-generating potential as a function of the distribution of bidder independent private values. When bidders are symmetric (identically distributed), soft floors have no effect on revenue, because a symmetric equilibrium always exists in strictly monotonic bidding strategies, and standard revenue-equivalence arguments thus apply. The industry often motivates soft floors as tools for extracting additional expected revenue from an occasional high bidder, for example a bidder retargeting the consumer making the impression. Such asymmetries in the distribution of bidder preferences do not automatically make soft floors profitable. This paper presents two examples of tractable modeling assumptions about such occasional high bidders, with one example implying low soft floors always hurt revenues because of strategic bid-shading by the regular bidders, and the other example implying high soft floors can increase revenues by making the regular bidders bid more aggressively.

History: Accepted by Juanjuan Zhang, marketing.

Keywords: advertising and media • marketing: pricing • microeconomics: market structure and pricing

1. Introduction

Since the world's first banner ad in 1994 (Singel 2010), advertising dollars have followed the shift of consumer attention to digital media, reaching more than one-third of total U.S. advertising spending by 2016. Despite starting with display banner ads, the lion's share of digital advertising dollars was initially spent on search ads, because they offered an unparalleled level of targeting (Goldfarb 2014). However, for the first time in the more recent history of digital advertising, spending on display ads surpassed spending on search ads in 2016 (emarketer 2016). An improved targeting ability is one of the key forces behind the resurgence of banners: unlike the banners from the 1990s, today's banner ads are targeted to the individual viewer one impression at a time by computer algorithms—a practice called "programmatic buying." A dominant method of allocating and pricing the display advertising space sold programmatically is real-time bidding (RTB), whereby each available impression is sold to interested advertisers by a sealed-bid auction that lasts a fraction of a second. Experts estimate that more than $20 billion in advertising is sold by RTB per year in the United States (emarketer 2016) in more than 30 trillion unique transactions (Friedman 2015).

What are the rules of these trillions of auctions? The vast majority of the "ad exchange" auctioneers use second-price sealed-bid "Vickrey" auctions—a dramatic shift from the obscurity of the Vickrey pricing rule in past auction-driven marketplaces documented by Rothkopf et al. (1990). However, several important players in the RTB industry have recently partially reversed this shift by introducing "soft floors"—bid levels below which the auction's pricing rule switches from second-price to first-price, sometimes also called "high-bid."[1] The "soft" part of "soft floor" contrasts with a "hard floor"—a bid level below which the auctioneer will not sell the impression, also known as "reserve price" in the auction literature (Myerson 1981). This paper provides the first theoretical treatment of soft floors and shows that their usefulness depends on the distribution of bidder preferences. Throughout this paper, bidders are assumed to have independent private values. When the bidders are symmetric (i.e., when their valuations are drawn from the same distribution), I show that the use of soft floors is misguided because they complicate bidding and have no effect on expected revenue. When the bidders are asymmetric in an RTB-relevant fashion (i.e., when high-valuation bidders occasionally join the auction), I show by two examples that soft floors can both hurt and benefit the auctioneer, depending both on the magnitude of the soft floor and on the valuation overlap between the regular bidders and the high bidders. The next few paragraphs introduce the three different modeling assumptions used in this paper and preview the results they imply. Please see Figure 1 for a representation of

Figure 1. (Color online) Effects of Soft Floor on Auction Revenue, and Assumptions They Rely on

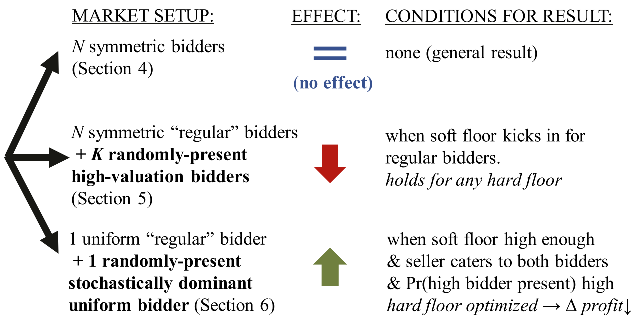

all the modeling variants covered in this paper, by section.

As long as the bidders are symmetric, I show that soft floors have no impact on auction revenue. In other words, when the different advertisers' valuations of each impression are drawn from the same distribution, soft-floor auctions are revenue-equivalent with standard auctions that have the same hard floor. The revenue-equivalence result is *not* a trivial extension of the well-known equivalence between first- and second-price auctions: just because first- and second-price auctions yield the same expected revenue (under bidder symmetry), it does not immediately follow that their hybrid arising from the presence of a soft floor will also be revenue-equivalent with the simple second-price auction: strategic bidders may react to the introduction of a soft floor by playing mixed strategies or by pooling, thus changing the relationship between valuations and the chance of winning. The first main result of this paper (discussed in Section 4) is a general proof that whereas bidders indeed react to the introduction of a soft floor by adjusting their bids, the resulting equilibrium is in pure monotonic strategies. The monotonicity of the bidding equilibrium in the soft-floor auction guarantees the soft-floor auction does not change any bidder's chance of winning relative the second-price auction with the same hard floor, which in turn keeps the expected revenue of the auctioneer unaffected according to the revenue-equivalence result of Myerson (1981). In the extensions, analogous arguments are then used to also show that revenue equivalence continues to hold even when bidders participate randomly, and also when the soft floor is hidden, as it tends to be on some exchanges.

Given the robust revenue equivalence in the symmetric model, the rest of this paper explores the obvious possibility that a rationalization of the soft-floor industry practice can arise from asymmetries among bidders. Inspired by the industry analysts who originally motivated the use of soft floors (e.g., Weatherman 2013), I consider the possibility that high-value bidders may occasionally enter the auction. For example, a "retargeting advertiser" (whose website the customer has just visited before arriving to the publisher auctioning off the customer's impression) likely values the impression much more than other advertisers who bid only on demographics. If such a high-value advertiser were always present, there would be little benefit to soft floors—the seller could simply increase the hard floor; but such a high-value advertiser may not participate in every RTB auction, so a soft floor might seem to be a clever adaptive mechanism that automatically activates a higher reserve price only when the advertiser does appear (Weatherman 2013). In contrast to this industry intuition, I show that adding randomly appearing asymmetrically high bidders always makes *low-enough* soft floors *sub*optimal for the auctioneer. On the other hand, very high soft floors can increase revenue under some assumptions. I now discuss both of these contrasting examples in turn.

In the second main result of the paper (covered in Section 5), I show that when the high bidders are guaranteed to have valuations above the regular bidders, soft floors low enough that all the high bidders face second-price pricing reduce expected revenues. The reason is that the strategic bid-shading by regular bidders always more than offsets the additional pricing pressure on the high bidders generated by the soft floor. The revenue loss can be derived in closed form, and it is bounded by the amount of revenue an auctioneer running a second-price auction would lose from losing one of the regular bidders. The bound echoes the classic result of Bulow and Klemperer (1996), who show that a revenue gain from setting the hard floor optimally is less than the gain from finding one more bidder.

If low soft floors hurt revenue, might higher ones help? The analysis of higher soft floors that "kick in" for the high bidders is not tractable, but tractability resumes in a slightly modified asymmetric model when the soft floor is so high that the auction effectively becomes a first-price sealed-bid auction. The third main result of this paper (covered by Section 6) provides a lower bound of the profitability of a soft-floor auction in an RTB-relevant asymmetric setting by analyzing the revenue potential of a first-price auction when one of two uniformly distributed potential bidders stochastically dominates the other, but only participates occasionally. This analysis extends the analysis of equilibrium bidding by Maskin and Riley (2000) and Kaplan and Zamir (2012) to random participation by one of the two bidders. The bidding equilibrium is in closed form, but the expected revenue calculation, and hence the ultimate revenue comparison with a second-price auction, involves an intractable integral. Approximating the integral numerically, I find that first-price auctions can revenue-dominate second-price auctions as long as the high bidder's chance of participation is high enough to induce aggressive bidding by the regular bidder. Because soft floor auctions weakly dominate first-price auctions, this result is an example of a situation in which soft floor can strictly increase the auctioneer's revenue. The exact conditions for the revenue dominance, as well as the magnitude of the revenue difference, are sensitive to whether the auctioneer can optimize the hard floor. Once the hard floor is optimized for the demand situation, only a small amount of additional revenue is available from also optimizing the pricing rule.

2. Literature Review

The literature on online display advertising (for thorough literature reviews, see Hoban and Bucklin 2015, Choi et al. 2017, or Johnson et al. 2017) contains very few papers about soft floors: Yuan et al. (2013) examine data from a large ad exchange that uses soft floors and estimate that more than half of the exchange's revenue is transacted using the first-price rule instead of the second-price rule. They conclude soft floors are an economically important phenomenon in the RTB marketplace. In contrast to the predictions of this paper, Försch et al. (2017) analyze the profitability of soft floors using a large-scale field experiment and conclude that even relatively low soft floors can increase the auctioneer's revenue. One possible explanation for the divergence between this paper's predictions and the field experiment's results is that the bidders in the experiment did not have enough time to adjust their strategies to the novel mechanism: unlike the rational bidders assumed herein, the bidders in Försch et al. (2017) do not reduce their bids when soft floors are introduced.

This paper also contributes to the broader literature on mechanism design in the online display advertising industry. Most work in that literature focuses on optimizing the pricing and allocation rules of RTB exchanges to address specific ways the bidders for online display impressions differ from bidders assumed in canonical models. For example, Abraham et al. (2016) focus on the informational asymmetry arising from informative "cookies" available to only some bidders in a pure common-value model, and they compare the two dominant auction pricing rules in terms of revenue. Arnosti et al. (2016) also study the impact of bidder asymmetries in a common-value model, with the asymmetry arising from the difference between "performance advertisers" who know their valuations of each impression and "brand advertisers" who do not. They focus mainly on market efficiency and propose a new theoretical mechanism that is nearly efficient. In contrast to Abraham et al. (2016) and Arnosti et al. (2016), this paper assumes that bidders have and know their private values, focuses on the ex ante asymmetry between "regular" bidders who tend to bid low and often and "high" retargeting bidders who bid high but rarely, and restricts attention to a prominent design implemented in the industry—the soft-floor auction. Celis et. al (2014) take a different approach to analyzing the competition between regular and retargeting bidders: instead of considering an ex ante asymmetric set of bidders, they assume an interim-asymmetric model of bidders drawn from a mixture of regular bidders poorly matched with the customer making the impression and high-valuation bidders who do match with the customer well. They note that such a mixture distribution is irregular in the sense of Myerson (1981), so standard auctions may not perform well, and they propose a novel mechanism called "buy-it-now or take-a-chance" which does better.

All of the above mechanism-design papers—including this paper—focus on the sale of a single impression. There is also an emerging stream of work that addresses the multiunit nature of RTB marketplaces. For example, Balseiro et al. (2015) use a modern fluid mean-field equilibrium notion to simplify an otherwise intractable model of budget-constrained bidders participating in a sequence of second-price sealed-bid auctions. Having outlined the contribution of this paper to the literature, I now describe the main mechanism of interest—the soft-floor auction.

3. Soft-Floor Auction Definition and Other Supply-Side Assumptions

One object (e.g., an ad impression in the RTB context) is for sale. The auctioneer values the object at zero and sets two reserves: a hard floor $h \geq 0$ and a soft floor $s > h$. The soft-floor sealed-bid auction collects bids, sorts them such that $b_{(1)} \geq b_{(2)} \geq b_{(3)} \ldots$, and determines the auction winner and the price paid as follows:

(1) When $b_{(1)} > s$, the bidder who submitted $b_{(1)}$ wins and pays $\max\{s, b_{(2)}\}$.

(2) When $s \geq b_{(1)} \geq h$, the bidder who submitted $b_{(1)}$ wins and pays $b_{(1)}$.

(3) When $h > b_{(1)}$, the auctioneer keeps the object.

In words, the soft floor functions as a reserve price in a second-price sealed-bid auction (2PSB) as long as at least one bid exceeds it (case 1). When no bids exceed s, the auction becomes a first-price sealed-bid auction (1PSB) with a reserve price equal to h (cases 2 and 3).

Throughout this paper, I assume h is common knowledge; that is, the auctioneer announces the reserve price before the auction. Regarding the bidders' information about s, I first assume the auctioneer also preannounces s (or that, equivalently, the bidders figure out both values through experimentation) and then address the possibility of keeping s hidden from bidders whenever tractable.

An analysis of the revenue implications of soft floors requires a demand-side model of bidders. This paper considers independent private-valuation (IPV) bidders—a standard assumption in auction theory. IPV is a reasonable model of bidders in the RTB context that motivates this paper: "valuation" of an impression is the increase in the advertiser's profit from winning the impression, "private" means no advertiser can learn about his own valuation of the impression from how much another advertiser values it, and "independent" means the values are statistically independent of each other in the population of bidders. Given the IPV assumption, a population distribution of valuations completes the model. This paper makes three partially nested assumptions about the distribution, summarized in Figure 1 and introduced in the previous section. I turn to the symmetric case next.

4. Symmetric Bidders: Soft Floors Have No Impact on the Auctioneer's Revenue

Suppose N bidders indexed by $i = 1,2,\ldots,N$ have private valuations v_i drawn independently from a continuous distribution F with full support on $[0,M]$. Following Krishna's (2009) notation, let G be the distribution of the maximum from $N-1$ *independent and identically distributed (iid)* draws from F, itself denoted Y_1: $G(Y_1) \equiv F^{N-1}(Y_1)$, and let X_1 be the highest of N *iid* draws from F, distributed $F^N(X_1)$.

This section demonstrates that when bidders are symmetric, soft floors have no impact on the auctioneer's expected revenue. The proof proceeds in two steps. First, for any $s > h \geq 0$, I construct a monotonically increasing equilibrium bidding strategy $\beta(v)$ that best responds to s and h. Second, the fact that the bidding strategy is monotonic means the soft-floor auction allocates the object to the same bidder as a standard auction with a hard floor of h would, and so the revenue-equivalence theorem of Myerson (1981) implies the soft-floor auction also produces the same expected revenue to the auctioneer. The exact form of the bidding equilibrium depends on the bidders' information about the soft floor. The following subsection (Section 4.1) analyzes the case of the soft floor being common knowledge among a fixed set of participating bidders. Section 4.2 then generalizes the bidding strategies to bidders participating randomly, and Section 4.3 takes up the case of bidders uncertain about the soft floor at the time of bidding. The main revenue-equivalence result is outlined and discussed in Section 4.4.

4.1. Bidding Equilibrium When the Soft Floor Is Common Knowledge and All Bidders Participate

I begin the exposition of bidding in a soft-floor auction under the canonical assumption that all N bidders know s and participate in the auction for sure. Let $\beta_I(v)$ denote the standard symmetric bidding equilibrium in a 1PSB with N bidders and a public reserve h (for a detailed derivation, see Krishna 2009):

$$\beta_I(v) = h\frac{G(h)}{G(v)} + \frac{1}{G(v)}\int_h^v xg(x)dx = E\left[\max\{Y_1, h\}|Y_1 < v\right], \tag{1}$$

where the roman numeral subscript on β indicates the first-price pricing rule. Then the bidding equilibrium $\beta(v)$ in the soft-floor auction can be characterized in terms of $\beta_I(v)$ as follows:

Proposition 1. *When $s < \beta_I(M)$, the following is a unique symmetric monotonic pure-strategy equilibrium of the soft-floor auction:*

$$\beta(v) = \begin{cases} v \leq \beta_I^{-1}(s) : \beta_I(v) \\ v > \beta_I^{-1}(s) : v \end{cases}.$$

When $s \geq \beta_I(M)$, the soft-floor auction becomes a 1PSB auction, and $\beta(v) = \beta_I(v)$.

Please see the appendix for detailed proofs of all propositions in this paper. The intuition for the result is as follows: when $s < \beta_I(M)$ (i.e., when the highest bidder would bid above it in a 1PSB), the equilibrium β involves a jump discontinuity at a valuation v^* such that $\beta_I(v^*) = s$. Bidders with $v < v^*$ shade their bids as if they were in a 1PSB auction. They effectively ignore the higher-valuation bidders because they cannot win against them. Bidders with $v > v^*$ bid their valuations as if they were in a second-price auction. Their bids are unaffected by the behavior of lower-valuation bidders, because bidding one's valuation is a dominant strategy under the 2PSB incentives. The jump discontinuity's location in the space of valuations and the magnitude of the jump ensure no bidder wants to unilaterally deviate from the pricing rule "assigned" to him by his valuation.

Example (F = Uniform[0,1]). Illustrating Proposition 1 on a concrete distributional example is useful. A uniform distribution of valuations implies $G(x) = x^{N-1}$, so the 1PSB bidding strategy is

$$\beta_I(v) = \frac{N-1}{N}v + \frac{h^N}{Nv^{N-1}}. \tag{2}$$

Figure 2. Equilibrium Bidding Strategy with a Known Soft Floor and Guaranteed Bidder Participation

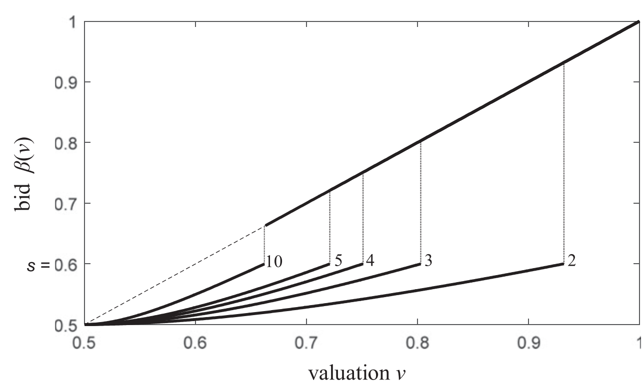

Notes. F = uniform [0,1], $s = 0.6$, and $h = 0.5$ (the h is optimal given the F). The dashed line indicates the 45° line; the dotted vertical lines indicate the jump discontinuities at v^* for the given numbers of bidders indicated by the numbers next to the lines.

The indifference equation $\beta_I(v^*) = s$ becomes $(N-1)\cdot(v^*)^N - sN(v^*)^{N-1} + h^N = 0$, which does not have a closed-form solution for a general N, but does for $N = 2$: $v^*(s; N = 2) = s + \sqrt{s^2 - h^2} < 1 \Leftrightarrow s < \frac{1+h^2}{2}$. Figure 2 illustrates the bidding function $\beta(v)$ for $N = 2$, 3, 4, and 10. I now turn to the possibility that the bidders participate randomly.

4.2. Bidding Equilibrium When the Soft Floor Is Common Knowledge and Participation Is Random

One of the apparent benefits of a soft floor is its ability to put pricing pressure on a single high-valuation bidder, who only pays the hard floor h under 2PSB rules. When such a bidder's presence is assured, the auctioneer can simply increase the hard floor; but when the number of bidders is uncertain, the soft floor $s > h$ can "kick in" precisely when there *happens to be* just one bidder. I will show this intuition is incomplete because it does not consider the associated revenue decrease when there *happen to be* multiple bidders.

Assume N symmetric potential bidders exist, and each of them enters independently of the other with probability $0 \le \alpha \le 1$. Because an entrant might face fewer opponents, he should bid less aggressively compared with facing all potential opponents for sure. Indeed, Harstad et al. (1990) show the existence of a $\beta_\alpha(v) < \beta_I(v)$, which can be expressed as a weighted average of the *contingent* 1PSB bidding functions that would apply for a fixed number of present bidders between 1 and N. Please see Equation (A.7) in the appendix for the $\beta_\alpha(v)$ for general F when $N = 2$. A concrete example with a uniform distribution is again helpful:

Example (F = Uniform[0,1], N = 2, and h = 0). This example makes clear how random participation reduces 1PSB bids relative to certain participation:

$$\beta_\alpha(v) = \frac{\alpha v^2}{2[1-\alpha(1-v)]} = \beta_I(v) - \frac{(1-\alpha)v}{2[1-\alpha(1-v)]}. \quad (3)$$

Not surprisingly, Proposition 1 generalizes to the situation with random participation:

Corollary 1. *For any $s \le \beta_\alpha(M)$, the following is a unique symmetric monotonic pure-strategy equilibrium of the soft-floor auction:*

$$\beta(v) = \begin{cases} v \le \beta_\alpha^{-1}(s) : \beta_\alpha(v) \\ v > \beta_\alpha^{-1}(s) : v \end{cases}.$$

When $s \ge \beta_\alpha(M)$, the soft-floor auction becomes a 1PSB auction, and $\beta(v) = \beta_\alpha(v)$.

The proof is analogous to the proof of Proposition 1, with $G(v) \equiv (1-\alpha) + \alpha F(v)$ as the equilibrium probability of winning. Figure 3 illustrates the bidding strategies for a range of α's in the uniform example.

Figure 3. Equilibrium Bidding Strategy with a Known Soft Floor and Two Potential Bidders Who Participate Randomly.

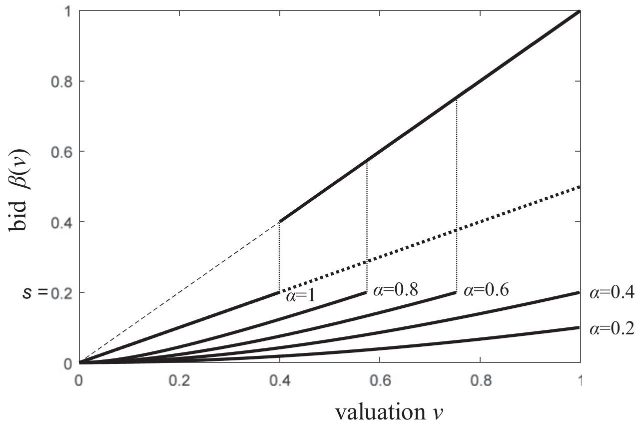

Notes. F = uniform [0,1], $h = 0$, $s = 0.2$, $N = 2$. The dashed line indicates the 45° line; the thicker dotted line indicates the 1PSB bidding strategy with two bidders, and the vertical lines indicate the jump discontinuities at v^* for the participation probabilities shown next to the jumps, if any.

4.3. Bidding Equilibrium When the Soft Floor Is Hidden at the Time of Bidding

Suppose the bidders are uncertain about the soft floor at the time of bidding, and they all summarize their beliefs about it by some distribution Ω on $[h,M]$. Optimal bidding must now average over the possibility that the soft floor happens to be low (and 2PSB rules will thus apply) and the possibility that the soft floor happens to be high (and the price paid will be equal to the winning bid). Unlike in the previous two subsections, characterizing the equilibrium in closed form is not possible even in the uniform example. However, the following proposition provides weak sufficient conditions for a symmetric monotonic equilibrium to exist and bounds the resulting bidding function below with $\beta_I(v)$:

Proposition 2. *When $f(v)\cdot F^{N-2}(v)$ and $\Omega(v)$ are continuous on $[h,M]$, the sealed-bid auction with a hard floor h and a hidden soft floor drawn from Ω on $[h,M]$ has a symmetric pure-strategy equilibrium characterized by an increasing bidding function $\beta(v) > \beta_I(v)$ that satisfies*

$$\beta'(v) = \frac{g(v)(v-\beta(v))}{G(v)[1-\Omega(\beta(v))]}. \quad (4)$$

The proof uses the Peano existence theorem to assure us Equation (4) has a solution. Compared with the 1PSB differential equation that gives rise to $\beta_I(v)$, Equation (4) adds the term in the square bracket. Because $\beta(v)$ is thus steeper everywhere, the relative ranking of the two bidding functions follows.

Figure 4. Bidding Strategy When Soft floor Is Hidden and Bidder Participation Is Guaranteed

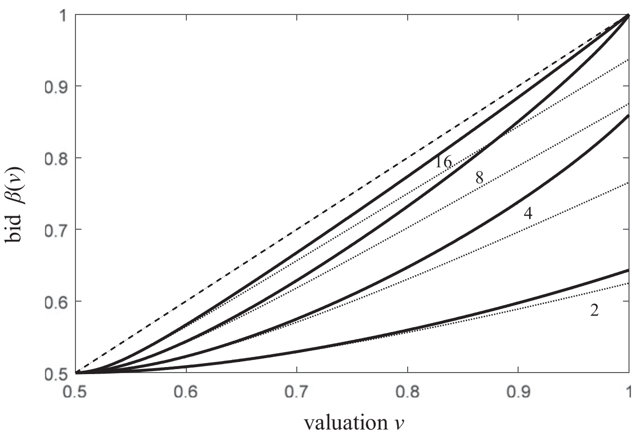

Notes. F = uniform [0,1], and $h = 0.5$ (the h is optimal given the F), and $s \sim$ Uniform[h,1]. The dashed line indicates the 45° line. The dotted lines indicate 1PSB bidding strategies without a soft floor, and the solid lines indicate bidding strategies with a hidden soft floor. Several levels of the number of bidders are indicated by the numbers next to the lines.

Intuitively, relative to 1PSB, random soft floors partially mitigate the increased payment associated with higher bids by switching the pricing rule to 2PSB. The resulting "random discount" gives the bidders an incentive to raise bid levels, and so the $\beta(v)$ exceeds equilibrium bidding in a 1PSB with the same number of bidders everywhere above h . See Figure 4 for a concrete uniform-distribution example.

4.4. Revenue Equivalence Under Bidder Symmetry

All of the previous subsections (Sections 4.1–4.3) find a monotonically increasing symmetric pure-strategy equilibrium of the soft-floor auction game. Under all three potential assumptions regarding the bidders' information about the soft floor and the bidders' participation, the introduction of a soft floor therefore does not affect any bidder's probability of winning. The introduction of a soft floor also does not affect the payoff of the bidder with the lowest trading valuation $v = h$, who makes zero surplus both with and without the soft floor. Therefore, the revenue-equivalence result of Myerson (1981) implies the soft floor does not affect the auctioneer's revenue and the bidders' surpluses—a fact I summarize in the next proposition.

Proposition 3. *Suppose bidders are symmetric in their valuations, their participation behavior, and their beliefs about the soft floor. Then, for every hard floor h, the introduction of a soft floor s > h has no effect on the auctioneer's revenue or any bidder's expected surplus.*

The "magic" of revenue equivalence stems from the fact that we only need to consider the allocation probability (the chance of winning) for every bidder type—the revenues are then implied by incentive compatibility. Please see Myerson (1981) for the original result and Krishna (2009) for the straightforward extension to the case of randomly participating bidders.

The case of randomly participating bidders is especially interesting to analyze deeper because it seems to agree with a common argument in favor of soft floors. Let $N = 2$ to simplify the combinatorics. Bidders with $v < h$ have no impact on revenue. When only one bidder with $v > h$ happens to participate, he wins for sure regardless of the pricing rule. Whereas he would pay only h in the 2PSB, the soft-floor auction charges him more, namely $\min(\beta_\alpha(v), s) > h$. It seems that this revenue advantage of the soft-floor auction over 2PSB might dominate its associated revenue disadvantage when both bidders happen to participate and they both have $v > h$ because that scenario happens much less often. For example, when $\alpha = h = 1/2$, the chance of only one $v > h$ bidder participating is $3/8$—three times greater than the chance that both bidders participate and have $v > h$. Proposition 3 shows that $\beta_\alpha(v)$ is calibrated such that the single-bidder advantage of the soft-floor auction is exactly offset by its two-bidder disadvantage. In the appendix, I demonstrate the revenue equivalence under random participation explicitly (i.e., without relying on the mechanism-design results used in the quick proof of Proposition 3) to illustrate how the advantage and the disadvantage cancel each other out in the expected revenue calculation.

Reflecting on the predictions of this section for bidding behavior is also useful empirically. Looking at Figures 2–4, one can make the following observations: keeping the hard floor constant, adding a soft floor should lead to bid-shading by low-valuation bidders, and so the distribution of the observed bids should become more skewed to the right. If the soft floor is common knowledge, the distribution of observed bids should also have a hole just above the soft floor. The data collected by Försch et al. (2017) do not have either of these features, suggesting the bidders in the experiment did not rationally adjust their bidding strategies to the presence of the soft floor.

5. Randomly Appearing High-Valuation Bidders: Low Soft Floors Reduce Auctioneer's Revenue, and High Soft Floors Break the Monotonicity of Bidding Strategies

In a prominent explanation of soft floors, Kevin Weatherman of the MoPub platform used a stylized example of a seller who sets a hard floor of \$1 and faces occasional bids around \$2 in addition to regular bidding activity in the \$0.75–\$0.90 range (Weatherman 2013). Weatherman's argument for why such a seller would benefit from soft floors is that the seller can lower his hard floor toward the bidding range of the regular bidders while introducing a soft floor above \$1: such an arrangement seems to preserve the price pressure on the high bidder (when exactly one such bidder happens to participate—multiple high bidders put price pressure on each other) while also collecting more revenue from low bidders (when no high bidders happen to be present). As Weatherman puts it, "the goal is to 'harvest' higher bids while not compromising on lower bid opportunities" (Weatherman 2013). Diksha Sahni of AppLift eloquently makes the same argument by pointing out that "when the gap between a bid and the second bid is significant, it may create a gap between potential revenues and actual revenues" (Sahni 2016). Motivated by these industry experts, this section analyzes the possibility of randomly present high-valuation bidders somehow making soft floors profitable for the auctioneer. I focus on the following asymmetric case:[2]

Definition. Let a *market with randomly appearing high-valuation bidders* always contain N "regular" bidders drawn from some F on [0,1] and K potential "high" bidders with valuations drawn from some Φ on $[L, M]$ with $L \geq 1$. The high bidders participate randomly and independently of each other with probability α. Participation by competing bidders is not observable by anyone before bidding.

An analysis of 1PSB equilibrium bidding in the above-defined market is not tractable even with $\alpha = 1$ (Maskin and Riley 2000), so assessing the profitability of a soft floor so high that nobody pays it is difficult. When the soft floor "kicks in," the analysis remains intractable when $s > \beta_I(1)$: no globally monotonic equilibrium strategy exists that would produce an incentive-compatible separation of the high bidders above s from the rest of the bidders. On the flipside, when $s \leq \beta_I(1)$ (i.e., the soft floor is low enough that at least some regular bidders bid above it), the analysis of equilibrium bidding is simple in that the soft-floor auction has the same equilibrium (outlined in Proposition 1) as without the high bidders:

Lemma 1. *When the soft floor is small enough that at least some regular bidders would bid above it in a 1PSB, i.e., when $s \leq \beta_I(1)$, the soft-floor auction in a market with randomly appearing high-valuation bidders has the same bidding equilibrium as the soft-floor auction without the high-valuation bidders characterized in Proposition 1, namely*

$$\beta(v) = \begin{cases} v \leq \beta_I^{-1}(s) : \beta_I(v) \\ v > \beta_I^{-1}(s) : v \end{cases}$$

for both bidder types.

The argument behind Lemma 1 is straightforward: if the high bidders indeed bid their valuations, the regular bidders assume they can only win when no high bidder is present, and so they behave the same as in an auction without high bidders. The high bidders, on the other hand, do not want to deviate from bidding their valuations to bidding the soft floor, because the associated lower price comes with leaving too many potential wins on the table. Intuitively, the low-enough soft floor keeps bidding tractable because it guarantees the bidders facing 1PSB rules are symmetric (they are just the regular bidders with valuations up to v^*) while all the remaining bidders face 2PSB incentives that are unaffected by asymmetries (bidding one's true valuation remains an equilibrium strategy in a 2PSB even with asymmetries because it is a dominant strategy).

So can soft floors increase revenue in this market? Because the bidding equilibrium is the same as in Proposition 1 when $s \leq \beta_I(1)$, the auctioneer's revenue from a soft-floor auction is easy to compare with that in a 2PSB auction with the same hard floor. When no high bidder enters, Proposition 3 proves that the soft floor does not impact revenue. When two or more high bidders enter, 2PSB rules with bids above $L > s$ determine the price in the soft-floor auction, so the soft floor does not impact revenue. Adding the soft floor only makes a difference when exactly one high bidder enters. The lone high bidder wins and pays either s or the highest regular valuation X_1, whichever is greater: $F^N(v^*)s + (1 - F^N(v^*))E(X_1 | X_1 > v^*)$ instead of paying $E(\max\{X_1, h\})$ without the soft floor. The revenue impact of soft floor is therefore

$$\Pi^A(h,s) - \Pi^A_{2PSB}(h) = \Pr(1\,\text{high})\left[F^N(v^*)s + (1 - F^N(v^*))E(X_1 | X_1 > v^*) - E(\max\{X_1, h\})\right], \tag{5}$$

where the probability of exactly one high bidder entering is $\Pr(1\,\text{high}) = K\alpha(1-\alpha)^{K-1}$. Now recall that s can be expressed as a conditional expectation of order statistics: $s = \beta_I(v^*) = E[\max\{Y_1, h\} | Y_1 < v^*]$. Substituting for s in Equation (5) yields

$$\begin{aligned} &\frac{\Pi^A(h,s) - \Pi^A_{2PSB}(h)}{K\alpha(1-\alpha)^{K-1}} \\ &= \underbrace{F^N(v^*)}_{\Pr(X_1<s)}\, \underbrace{E(\max\{Y_1, h\} | Y_1 < v^*)}_{s} \\ &\quad + (1 - F^N(v^*))E(X_1 | X_1 > v^*) \\ &\quad - F^N(v^*)E(\max\{X_1, h\} | X_1 < v^*) \\ &\quad - (1 - F^N(v^*))\,\overbrace{E(\max\{X_1, h\} | X_1 > v^*)}^{=E(X_1|X_1>v^*)\ \text{because}\ v^*>h} \\ &= F^N(v^*)[E(\max\{Y_1, h\} | Y_1 < v^*) \\ &\quad - E(\max\{X_1, h\} | X_1 < v^*)] < 0. \end{aligned} \tag{6}$$

I have just derived the key ingredient of the following result:

Proposition 4. *For every hard floor h, adding a soft floor small enough that the highest regular bidder would bid above it in a 1PSB reduces the expected revenue compared with a 2PSB with the same hard floor. The expected revenue reduction increases in the soft floor magnitude, and it is the same as if the auction followed 2PSB pricing, but lost one regular bidder whenever exactly one high-valuation bidder entered and all the regular bidders happened to have valuations low enough to bid below the soft floor.*

Note that soft floors reduce expected auctioneer revenue precisely when they "kick in" to put pricing pressure on a randomly appearing high bidder, that is, precisely in the situation discussed by soft-floor advocates (e.g., Weatherman 2013, Sahni 2016). The advocates are correct in noting the soft floor adds pricing pressure on the high bidder whenever he is present; but Proposition 4 shows the coincident bid-shading by low-valuation regular bidders more than erodes the benefits of the added pricing pressure.

Because the revenue reduction is increasing in s and weighted by the probability of exactly one high bidder entering, the revenue reduction is bounded by the revenue loss from losing one regular bidder. Proposition 4 thus echoes the classic result of Bulow and Klemperer (1996), who show that a revenue gain from setting the hard floor optimally in a symmetric model is less than the gain from finding one more bidder. Unlike in the Bulow and Klemperer (1996) case in which reserves obviously discourage bidder entry, it is not clear whether having a soft floor might encourage bidder entry; but if it did, Proposition 4 shows that the auctioneer would always prefer adding a small soft floor to losing one regular bidder in a 2PSB.

6. Example of Profitable Soft Floors: First-Price Auctions Can Revenue-Dominate Second-Price Auctions in a Market with a Randomly Appearing Stochastically Dominant Bidder

The previous section argues that low-enough soft floors hurt expected revenue of the auctioneer. Unfortunately, the analysis of medium soft floors in the market with randomly appearing high-valuation bidders defined above is not tractable. However, tractability resumes for very high soft floors in a slightly modified but still RTB-relevant market with randomly appearing *stochastically dominant* bidders, where a "very high" soft floor is one that never kicks in. Such a soft floor effectively implements 1PSB in a marketplace originally designed around the 2PSB rule. The soft-floor auction obviously weakly dominates both the 1PSB auction and the 2PSB auction. So a situation in which 1PSB strictly revenue-dominates 2PSB is an example of a market in which soft floors strictly increase the auctioneer's expected revenue. This section provides such an example.

Throughout this section, I assume that there is only one regular bidder (bidder 1) and one randomly appearing bidder (bidder 2), both have uniformly distributed valuations, and the randomly appearing bidder who appears with probability α is stochastically dominant in that $v_1 \sim U[0,1]$ while $v_2 \sim U[0, M]$ for some $M \geq 1$. The common lower bound of the two valuation supports simplifies analysis, and the stochastic dominance captures the idea of a "high" bidder (differently from Section 5, which assumed guaranteed dominance). To provide an example of a market in which soft floors increase revenue, this section exhibits a range of (α, M) parameter values for which the 1PSB pricing rule (and hence also the soft-floor auction) revenue-dominates the 2PSB pricing rule even with rule-optimized hard floors.

6.1. The Optimal Mechanism Favors the Regular Bidder

Before comparing the revenues of the two standard pricing rules, it is useful to consider the optimal mechanism of Myerson (1981), which would allocate the impression to the bidder with the highest *positive* virtual value $\psi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. The random participation of the high bidder does not change his virtual value, and the uniform F_i then imply $\psi_1(v_1) = 2v_1 - 1 > 2v_2 - M = \psi_2(v_2)$. Therefore, the optimal mechanism would impose a higher hard floor of $h_2 = M/2$ on the high bidder than the optimal hard floor of $h_1 = 1/2$ for the regular bidder, and it would level the playing field further by favoring the regular bidder in picking the auction's winner. Specifically, when $\psi_i(v_i) > 0$, it would award the impression to the regular bidder whenever $v_1 > v_2 - \frac{M-1}{2}$. The problem with such a scheme is obvious: the high bidder would try to obscure its identity—a behavior called "false-name bidding" by Arnosti et al. (2016). Nevertheless, the optimal mechanism suggests when 1PSB is likely to have a revenue advantage over 2PSB, namely whenever the regular bidder bids aggressively in the 1PSB, and thus wins the auction despite having a smaller valuation than the high bidder.

6.2. Equilibrium Bidding Under the First-Price Rule

Bidding in the 2PSB is in dominant strategies, and the derivation of the expected revenue π_{II} and its associated optimal hard floor h_{II}^* are relegated to the appendix. In contrast, the characterization of 1PSB bidding in the market with a stochastically dominant randomly appearing bidder is new to the literature, extending the analysis of Kaplan and Zamir (2012) who derived the $\alpha = 1$ special case for an arbitrary h, and themselves extended the result of Griesmer et al. (1967), who derived the $\alpha = 1$ & $h = 0$ special case. The next proposition characterizes the equilibrium inverse bidding functions in closed form:

Proposition 5. *For any $h, \alpha \in [0,1]$ and $M \geq 1$, equilibrium bidding involves both bidders submitting bids in the $[h, \bar{b}]$ interval, where*

$$\bar{b} = \frac{hM + \alpha(h^2 + M - hM)}{\alpha + M}.$$

The inverse bidding function of the regular bidder is

$$v_1(b) = \frac{h(A+h)}{A+b+[h(A+h)-(A+\bar{b})]\left(\frac{b-h}{\bar{b}-h}\right)^{\theta}\left(\frac{b+h+A}{\bar{b}+h+A}\right)^{1-\theta}},$$

where

$$A = \frac{M(1-\alpha)}{\alpha}$$

and

$$\theta = \frac{h}{A+2h}.$$

The inverse bidding function of the high bidder is

$$v_2(b) = \frac{b[v_1(b)+A] - h(h+A)}{v_1(b) - b}.$$

The proof generally follows the approach of Kaplan and Zamir (2012), adapted for the randomly missing bidder. Figure A.2 in the appendix illustrates the following intuitive comparative statics of equilibrium bidding: as α decreases, the support of the bids shrinks down to h ($\bar{b}$ approaches h), and most regular bidders bid just above h, effectively banking on the high bidder not showing up. In response, the high bidder becomes less aggressive (bids less for a given value of v) because the regular bidder's behavior presents an opportunity to win very often. As M increases, the support of the bids expands to accommodate the greater gains of trade on the table, the high bidder becomes more aggressive in his attempt to win the gains from trade, and the regular bidder becomes more aggressive in response. No closed-form solution exists for the bidding functions shown in Figure A.2, but bidding functions are not necessary for the computation of the expected 1PSB revenue π_I, because the expected revenue follows directly from the inverse bidding functions $v_i(b)$:

Lemma 2. $E(\pi_I | h < 1)$

$$= h[1 - h(1-\alpha+\alpha h)] + \int_h^{\bar{b}} 1 - v_1(x)[1-\alpha+\alpha v_2(x)]dx.$$

The integral in Lemma 2 does not have a closed form but can be trivially approximated numerically as a sum on a fine grid. Optimizing the hard floor is then straightforward. I now turn to comparing the revenue-generating potential of the two pricing rules.

6.3. Comparing the Expected Revenue of the First-Price Rule (1PSB) with the Second-Price Rule (2PSB)

Let the *percentage revenue lift* from 1PSB versus 2PSB be defined as follows:

$$\%\text{ revenue lift from 1PSB versus 2PSB} \equiv \frac{E(\pi_I | h = h_I^*) - E(\pi_{II} | h = h_{II}^*)}{E(\pi_{II} | h = h_{II}^*)}. \quad (7)$$

The top panel of Figure 5 plots the percentage revenue lift as a function of α, for three qualitatively different examples of M: 1, 2, and 3. When $M = 1$, the bidders' valuations are symmetric, but one bidder participates randomly. When $M \leq 2$, the optimizing seller caters to both bidders regardless of pricing regime by setting both hard floors below 1. When $M = 3$, the seller caters to both bidders when α it low, and only the high bidder otherwise. To aid in understanding the forces underlying the revenue

Figure 5. Revenue Comparison of the Two Pricing Rules for Different M, as a Function of α

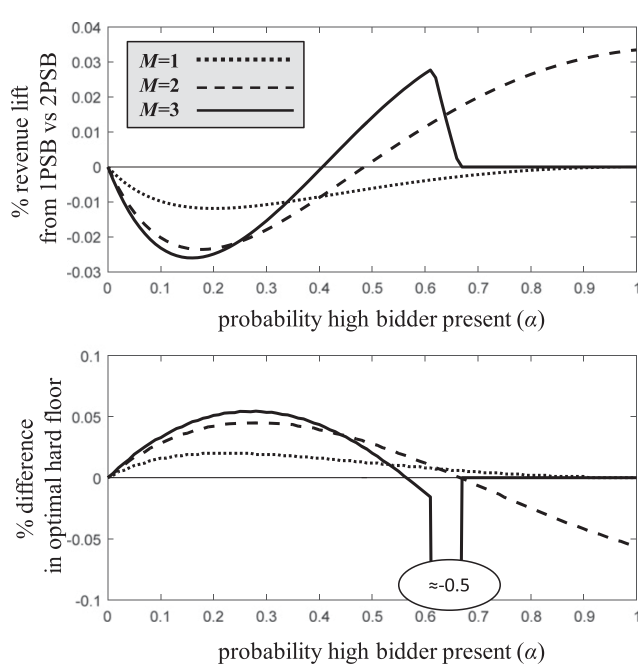

results, the bottom panel plots the percentage difference in the optimal hard floor, defined as % hard floor difference $\equiv \frac{h_I^* - h_{II}^*}{h_{II}^*}$ on the same α axis as the top panel.

Slowly unpacking the different effects illustrated in Figure 5 is helpful, starting with the high bidder always being present ($\alpha = 1$): when $M = 1$, the two bidders are symmetric, and standard revenue equivalence thus implies no difference in revenue. When $M = 2$, we have the "stretch case" of Maskin and Riley (2000)—a situation they show to favor the 1PSB as long as $h = 0$. I find 1PSB continues to dominate 2PSB in the stretch case even with optimal rule-specific hard floors: in addition to the aggressive bidding by the regular bidder, the optimal 1PSB is also more efficient (has lower reserve; see bottom panel of Figure 5). Once M increases to 3, revenue equivalence under $\alpha = 1$ is restored because both pricing rules only cater to the high bidder and effectively turn from auctioning to posted pricing.

Now consider intermediate values of α: Figure 5 indicates 1PSB can revenue dominate 2PSB when enough asymmetry exists in valuations ($M > 1$), the high bidder's participation is likely enough to make the regular bidder bid aggressively (α high enough), and the optimal hard floor is low enough to cater to both bidders (α and M small enough). A goldilocks (α, M) is thus required for 1PSB to dominate 2PSB. One way to understand the revenue-dominance of 1PSB is via its better approximation of the optimal mechanism whenever it induces the regular bidders to bid aggressively and effectively favors them in allocation.

Finally, consider very low α: when $\alpha = 0$, the two pricing rules are revenue-equivalent because they again reduce to posted pricing with a single bidder. When α is small but positive, 2PSB outperforms 1PSB for all M because the 1PSB's $\underline{b}$ approaches h, whereas 2PSB involves a broader range of prices (up to 1). The optimal hard floor for 1PSB rises above that of 2PSB to try and compensate for the low bids, but it apparently cannot compensate fully.

To generalize from the three particular levels of M considered in Figure 5, as well as to examine the role of hard-floor optimization, please see Figure 6, which displays a contour plot of percentage revenue lift from 1PSB. One way to summarize the patterns in Figures 5 and 6 is as follows:

Summary of Revenue Comparisons. *The first-price rule revenue dominates the second-price rule when α is high enough for the regular bidder to bid aggressively. When the hard floor is optimized, a second condition for the revenue dominance of the first-price rule is the asymmetry being small enough for the seller to cater to both bidders instead of just making a take-it-or-leave-it offer to the high bidder.*

The result illustrated by the right-hand panel of Figure 6 extends the findings of Maskin and Riley (2000), who separately analyze the two asymmetries involved in the construction of the RTB-relevant "high" bidder considered here and find that (1) stretching the distribution of bidder 2 relative to the bidder 1's distribution favors the first-price rule, and (2) random participation of bidder 2 (which Maskin and Riley call "shifting of probability weight to the bottom of support") favors the second-price rule. The right-hand panel of Figure 6 replicates both individual effects as special cases ($\alpha = 1$, $M > 1$) and ($\alpha < 1$, $M = 1$), respectively, and shows the effect of a *combination of both asymmetries* on the relative profitability of the two price rules. One way to interpret the joint effect is to conclude the effect of random participation is stronger than the effect of stochastic dominance by one bidder because the 2PSB rule revenue dominates the first-price rule for all $\alpha < 0.75$ regardless of M, and the revenue loss from picking the wrong pricing rule is much bigger when the 1PSB rule is chosen erroneously than when the 2PSB rule is chosen erroneously.

Figure 6. Percentage Revenue Lift of 1PSB vs. 2PSB in (α, M) Space

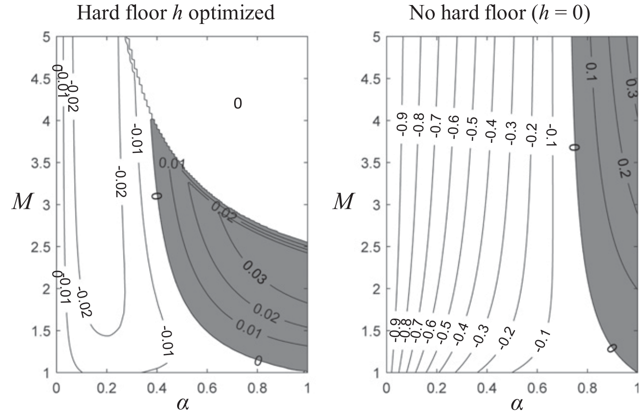

Notes. The shaded area indicates the region of the parameter space for which 1PSB generates higher expected revenue than 2PSB. The contours in the left plot are in 0.01 intervals, and the contours in the right plot are in 0.1 intervals.

The left-hand panel of Figure 6 then examines how the joint effect of "stretching" and random participation changes when the auctioneer optimizes the hard floor. Two qualitative differences relative to the right-hand panel emerge, each due to the auctioneer effectively catering to only one bidder: when the high bidder is unlikely to show up (α low), the auctioneer caters only to the regular bidder. When the high bidder is likely to show up and likely to have a much higher valuation than the regular bidder (α and M high), the auctioneer caters only to the high bidder. In either situation, no revenue difference exists between the two pricing rules.

One important difference between the left and the right panels in Figure 6 is the magnitude of the revenue lift: when the hard floor is optimized, the absolute difference between the two pricing rules is only a few percentage points. By contrast, 1PSB can increase profit by more than 30% when no binding hard floor exists and α is very high, but it can also lose the entire potential revenue when α is very low. Once the hard floor is optimized for the (α, M) situation, only a small amount of revenue seems to be available from also optimizing the pricing rule.

7. Discussion

Soft floors have emerged in the RTB digital display advertising industry as a potential tool for increasing publisher revenues. This paper shows soft floors are not likely to deliver on this promise in the long run when the bidders are ex ante symmetric, even if the auctioneer keeps the exact soft-floor level hidden from the bidders, or when bidders participate in the auction randomly. Adding randomly appearing "high" bidders (e.g., retargeting advertisers in the RTB context) to the auction does not automatically make soft floors profitable either: the profitability of soft floors depends on their magnitude and on the amount of valuation overlap between regular and high bidders. To illustrate the nuanced profitability of soft floors in RTB-relevant asymmetric markets, this paper provides both an example in which soft floors reduce revenue and an example in which they increase it. Overall, the paper provides three main results shown in Figure 1. I now discuss the three results in turn, with one paragraph devoted to each of them.

When the bidder valuations are all drawn from the same distribution (and the bidders are thus "symmetric" in the auction-theory parlance), low-valuation bidders shade their bids down in response to a soft floor, but the fact that low-bidding winners pay their bids exactly compensates for the seeming loss of revenue. Soft floors are revenue neutral because the equilibrium bidding function remains monotonically increasing as in the benchmark second-price auction with the same hard floor, and the classic revenue-equivalence result of Myerson (1981) thus applies. The monotonicity of equilibrium bidding (and hence the revenue equivalence) continues to hold even when the auctioneer hides the exact level of the soft floor before bidding, or when bidders participate randomly.

Soft-floor advocates often point to a gap between the winning bid and the second-highest bid in RTB auctions and argue the auctioneer can capture some of this gap as extra revenue using a soft floor. The symmetric case discussed in the previous paragraph shows an occasional *random* realization of a large gap by a set of otherwise ex ante similar bidders is not a good argument for soft floors. However, the possibility remains that the gap is *systematic*, arising from the presence of bidders whose valuation is known to be relatively high. For example, one of the bidders bidding on a particular impression may be a retargeting advertiser whose website the customer has just visited. A soft floor might seem to put pricing pressure on such an "asymmetrically high" bidder while preserving revenues when he happens not to show up at the auction. The second main result of this paper shows that this intuition is incomplete when the soft floor is low: the soft floor does indeed add pricing pressure on the asymmetrically high bidders, but the coincident bid-shading by low-valuation bidders always more than erodes the benefits of the added pricing pressure. The phenomenon of bid-shading shows that even in a soft-floor auction, the pricing pressure on the asymmetrically high bidder ultimately stems from the presence of lower-valuation bidders, more of whom shade their bids down when the soft floor increases. As a result, soft floors can actually reduce expected auctioneer revenue precisely in the asymmetric situation that motivates their use in the industry.

The third main result of this paper is an example of an asymmetric market with randomly appearing "high" bidders, in which the auctioneer can profit from a soft floor. Unlike in the second result, which considered the effect of a soft floor that "kicks in" for at least some bidders, the third result only considers soft floors high enough to turn the auction into a 1PSB auction. Maskin and Riley (2000) provide an example of an asymmetric market with one stochastically dominant bidder in which 1PSB revenue dominates 2PSB. This paper extends Maskin and Riley's (2000) revenue-dominance result to the RTB-relevant situation of the stochastically dominant bidder present only occasionally and also shows the revenue dominance can survive hard-floor optimization by the auctioneer under some conditions: 1PSB rules can revenue-dominate 2PSB rules as long as the high bidder's participation probability is high enough and the asymmetry is small enough for the auctioneer to cater to both bidders. However, the relative revenue advantage of the 1PSB rule is much smaller once the auctioneer optimizes hard floors. Thus, changing the pricing rule in RTB auctions seems likely to only lead to large revenue effects if the hard floors are difficult to optimize for some reason. This analysis is orthogonal to other arguments for switching to the 1PSB rule, such as its increased transparency as argued by Chen (2017) and Moesman (2017), who echo the earlier analysis of Rothkopf et al. (1990).

Several directions of future research remain. One potentially fruitful avenue would be to consider the direct impact of soft floors in the model of Section 6: as the soft floor rises toward the point when the auction turns into a 1PSB auction for all the bidders, what happens to the revenue advantage of the soft-floor auction over the underlying simple second-price auction? If the revenue advantage is decreasing, a case could be made for optimizing soft floors; but if it increases toward the point when soft floors do not kick in at all, the auctioneer would have a much simpler choice between the two standard pricing rules.

Soft floors are not the only proposed novel mechanisms in the RTB space: for example, Celis et. al (2014) propose a "buy-it-now or take-a-chance" mechanism to address an irregularity in the distribution of bidders arising from random matching with advertisers. It would be interesting to analyze how soft floors would perform under their assumptions because the idea of random matching is similar in spirit to the randomly present high bidders analyzed here. Unlike in this paper, the bidders in Celis et. al (2014) are ex ante symmetric, albeit coming from a mixture distribution. So one can conjecture on the basis of Proposition 3 that soft floors would not impact revenues under random matching as long the mixture had a full support and implied a monotonic 1PSB equilibrium bidding strategy.

Another extension could address the multiunit reality of RTB marketplaces: throughout the paper, I focused on a single auction attended by several independent private-value bidders. However, advertisers looking to purchase impressions on ad exchanges face a sequence of opportunities to show their ad, and they often view these opportunities as substitutes because they are budget constrained. Some advocates of soft floors correctly point out that bidding one's full private valuation in a sequential auction for substitutes is not optimal (e.g., Nolet 2010, Strong 2012). Instead, one needs to bid the valuation of winning net of the opportunity cost of trying again later, and the opportunity-cost calculation needs to take into account equilibrium considerations, because the opportunity cost depends on the behavior and types of competing bidders (see, e.g., Engelbrecht-Wiggans 1994, Milgrom and Weber 2000). Balseiro et al. (2015) provide a new solution concept called "Fluid Mean Field Equilibrium" and stationarity assumptions that together make accounting for multiple bidding opportunities practical. However, existing theory does not support the advocates' leap of faith that soft-floor or 1PSB auctions somehow resolve this issue, but rather continues to find 1PSB and 2PSB are revenue-equivalent even in sequential settings under the symmetric model (e.g., Reiß and Schöndube 2010, Chattopadhyay and Chatterjee 2012). Given these revenue-equivalence results, one can conjecture that an analysis analogous to Section 4 of this paper would show soft floors have no effect on auction revenue even in a sequential-auction model under bidder symmetry. An interesting future direction of inquiry would examine soft floors in sequential auctions with asymmetric bidders.

The managerial recommendations of the results presented in this paper are clear. First, soft floors should be eliminated from ad exchanges that do not seem to have asymmetrically high bidders. Second, low soft floors (i.e., soft floors that kick in for regular bidders) should be avoided, but markets with randomly appearing asymmetrically high bidders might benefit from high soft floors or even soft floors that never "kick in" as reserves. Third, managers should focus on setting their hard-floor levels correctly given the demand they face, instead of worrying about switching pricing rules. Finally, managers should not worry about the "revenue gap" between the top two bids in the second-price auction identified by Sahni (2016). They should resist the temptation to somehow monetize the gap and rest easy knowing the winner needs to capture the entire gap as his surplus to continue bidding truthfully in dominant strategies, that is, to preserve the clear bidding incentives that make the second-price rules desirable. One of the most powerful implications of Myerson's (1981) revenue equivalence is that this strategic simplicity for bidders does not come at a cost to the auctioneer as long as the bidders are symmetric—the second-price auction with a correctly chosen reserve is at least as profitable as any other auction format the manager may wish to implement.

Appendix. Proofs Not Covered in the Main Text

Proof of Proposition 1. The proof proceeds in three steps, establishing the following claims in turn.

Claim 1. Any monotonically increasing equilibrium bidding function $\beta(v)$ with $\beta(M) > s$ must have a jump discontinuity at the valuation level v^* such that $\beta(v^*) = s$, such that no bids in the (s, v^*) interval are submitted.

Claim 2. The proposed bidding strategy is a Nash equilibrium strategy.

Claim 3. The proposed bidding strategy is a unique Nash equilibrium strategy.

Figure A.1. Equilibrium and Deviation Expected Surpluses of the Focal Bidder ($N = 3$, F=U[0,1])

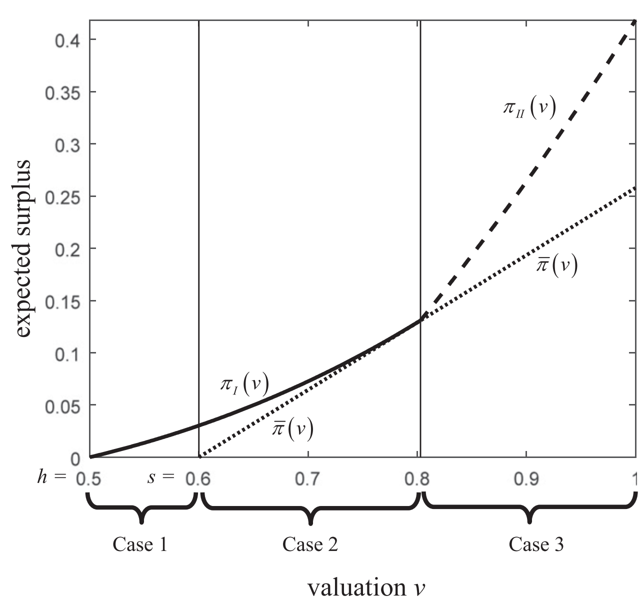

Proof of Claim 1 (Jump Discontinuity). Suppose a symmetric monotonically increasing bidding equilibrium $\tilde{\beta}(v)$ exists such that $\tilde{\beta}(M) > s$. Let v^* be the valuation level such that $\tilde{\beta}(v^*) = s$. By construction, the v^* bidder receives a positive surplus of $\Pr(win)(v^* - s) > 0$, so $v^* > s$. Now consider bidders with $v > v^*$, who also bid above s by monotonicity and hence face 2PSB pricing with a reserve of s. By the dominant-strategy properties of 2PSB, these bidders bid their valuations, that is, $\tilde{\beta}(v) = v$ for all $v > v^*$. Therefore, the bidding function $\tilde{\beta}(v)$ must approach v^* as v approaches v^* from above, resulting in the jump discontinuity:$\lim_{v \to v^{*+}} \tilde{\beta}(v) = v^* > s = \beta(v^*)$. By monotonicity, bidders with $v > v^*$ bid above v^* and bidders with $v < v^*$ bid below s, so no bids in the (s, v^*) are submitted.

Proof of Claim 2 (The Proposed Bidding Strategy Is a Nash-Equilibrium Strategy). Suppose all N-1 competitors follow the bidding strategy $\beta(v)$ outlined in the proposition, and consider a focal bidder with valuation v. It is enough to show that bidding according to $\beta(v)$ is the focal bidder's best response to those competitors. Three cases depend on the magnitude of v; please see Figure A.1 for an illustration of the three cases and all the relevant expected surpluses for the $N = 3$ and $F = \text{Uniform}[0,1]$ example with $h = 0.5$ and $s = 0.6$.

Case 1 ($v \le s$). The bidder can bid below s and guarantee 1PSB pricing should he win. Winning is only possible if all competitors also bid below s, and such competitors follow β_I by assumption. Because β_I is an equilibrium bidding function in 1PSB, bidding $\beta_I(v)$ is the focal bidder's best response to the relevant competition. Hence, he can make a positive expected surplus of $\pi_I(v) = G(v)(v - \beta_I(v))$ by bidding below s. There is only one nonlocal deviation to consider: bidding more than s and triggering 2PSB pricing. But triggering 2PSB pricing also triggers the soft floor as the reserve, so the focal bidder will pay at least s upon winning, which is in turn weakly more than his valuation. Therefore, the focal bidder cannot make a positive payoff by bidding above s, and his overall best response to the soft-floor auction incentives is to follow the proposed $\beta(v)$ and bid below s.

Case 2 ($s < v < v^*$). The bidder can bid below s and guarantee 1PSB pricing should he win. The argument presented in Case 1 shows that doing so will yield an expected surplus of $\pi_I(v) = G(v)(v - \beta_I(v))$. Also as in Case 1, this strategy dominates bidding over v^* because $v < v^*$. Alternatively, the bidder can bid in the (s, v^*) interval under 2PSB rules. Because no competitors bid in the (s, v^*) interval (they follow the proposed $\beta(v)$ by assumption), any bid there by the focal bidder wins whenever all the competitors bid weakly below s, that is, with probability $G(v^*)$, and thus triggers the soft floor as the reserve price. The alternative payoff from bidding in the (s, v^*) interval is thus $\overline{\pi}(v) = G(v^*)(v - s)$.

I now prove $\pi_I(v) \ge \overline{\pi}(v)$ for all $s < v < v^*$; that is, $\beta(v) = \beta_I(v)$ is the best response of the focal bidder. By construction, the two payoffs are increasing in v and coincide at v^*: $\pi_I(v^*) = \overline{\pi}(v^*)$. However, $\pi_I(v)$ dominates $\overline{\pi}(v)$ on (s, v^*) at lower valuations because $\pi_I(v)$ has a lower slope: $\pi_I'(v) = G(v) < G(v^*) = \overline{\pi}'(v)$. The slope of $\overline{\pi}(v)$ is trivial from its form above, and the slope of $\pi_I(v)$ can either be derived from Equation (1) or obtained from the standard mechanism-design result that the slope of equilibrium expected surplus in a standard auction is the probability of winning.

Case 3 ($v \ge v^*$). The bidder has three options for obtaining a positive expected surplus:

(a) Bidding over v^* triggers 2PSB pricing, and so the best bid above v^* is one's true valuation. The expected surplus is $\pi_{II}(v) = G(v^*)(v - s) + \int_{v^*}^{v} (v - Y_1) dG(Y_1)$.

(b) As in Case 2, bidding in the (s, v^*) interval results in an expected payoff of $\overline{\pi}(v)$. Bidding his true valuation as in (a) dominates this option because $\pi_{II}(v) - \overline{\pi}(v) = \int_{v^*}^{v} (v - Y_1) dG(Y_1) > 0$.

(c) Bidding below s triggers the 1PSB pricing rule, but the bidder cannot bid the same amount as he would in a 1PSB, because $\beta_I(v) > s$ by construction. Instead, the bidder solves a constrained optimization problem, finding the best valuation w below v^* to mimic: $\max_{w \le v^*} G(w)(v - \beta_I(w))$. It is easy to see the objective function is increasing on $[0, v^*]$: its derivative in the mimicked type w is $D(w; v) \equiv \frac{d}{dw} G(w)(v - \beta_I(w)) = g(w)(v - \beta_I(w)) - G(w)\beta_I'(w)$. Because the w type is optimizing, the first-order condition $D(w; w) = 0$ must hold, so $g(w)(w - \beta_I(w)) = G(w)\beta_I'(w)$. Substituting the first-order condition back into $D(w; v)$ yields $D(w; v) \equiv g(w)(v - w) > 0$. Therefore, the optimal type to mimic is v^*; that is, the optimal bid weakly below s is s, and the expected surplus from bidding it is the same as $\overline{\pi}(v)$. As shown above in (b), this surplus is also dominated by bidding true valuation.

In summary, the best response of the focal bidder with $v \ge v^*$ is to bid his true valuation, that is, to follow the proposed $\beta(v)$.

Proof of Claim 3 (Uniqueness). Now consider another monotonic equilibrium bidding strategy $\tilde{\beta}(v)$. From Claim 1 we know $\tilde{\beta}(v)$ must have a jump discontinuity at $\tilde{v}^*$ such that $\tilde{\beta}(\tilde{v}^*) = s$. Consider bidders with valuations below $\tilde{v}^*$ first: from monotonicity of $\tilde{\beta}(v)$, the bidding strategy of any bidder with $v < \tilde{v}^*$ depends only on the bidding strategy of bidders with valuations below v and hence also below $\tilde{v}^*$. Therefore, all bidders with $v < \tilde{v}^*$ effectively face 1PSB incentives. It is well known that $\beta_I(v)$ is the unique equilibrium bidding function of 1PSB, so $\tilde{\beta}(v)$ must coincide with $\beta_I(v)$ below $\tilde{v}^*$ unless the bidders can profitably deviate nonlocally to bid above s. However, the above derivation of the deviation payoff $\overline{\pi}(v)$ in Case 2 holds for any monotonic bidding function such that no bids in the (s, v^*) are submitted, so it also holds for the present $\tilde{\beta}(v)$, and the above argument in Case 2 implies the nonlocal deviation to bidding above s by bidders with $v < \tilde{v}^*$ cannot pay off. It follows that $\tilde{\beta}(v) = \beta_I(v)$ for $v <$ s, and so $\tilde{v}^* = v^* = \beta_I^{-1}(s)$. Now consider the bidders with valuations above $\tilde{v}^* = v^*$, who either bid their valuation (as implied by the properties of 2PSB) or deviate nonlocally below v^*. Because we have established that $\tilde{\beta}(v)$ must coincide with $\beta_I(v)$ below v^*, the above argument in Case 3 implies $\tilde{\beta}(v) = v$ for v>v*, and so $\tilde{\beta}(v)$ must coincide with $\beta(v)$ on the entire support of F. In other words, $\beta(v)$ is the unique symmetric Nash-equilibrium bidding strategy. Q.E.D. Proposition 1.

Proof of Proposition 2. Consider one bidder with valuation v and suppose all N-1 of his competitors bid according to some increasing bidding function $\beta(v)$. The focal bidder solves

$$\max_b \underbrace{G(\beta^{-1}(b))}_{\Pr(win)}\Bigg[\underbrace{\Omega(b)(v - E[\max\{s,\beta(Y_1)\}|s<b\,\&\,Y_1<\beta^{-1}(b)])}_{b>s\rightarrow 2PSB} + \underbrace{(1-\Omega(b))(v-b)}_{s>b\rightarrow 1PSB}\Bigg]. \tag{A.1}$$

The $E[\max\{s,\beta(Y_1)\}|s<b\,\&\,Y_1<\beta^{-1}(b)]$ term, which captures price paid whenever the bid exceeds the soft floor, seems rather complex at first but simplifies to

$$E\big[\max\{s,\beta(Y_1)\}|s<b\,\&\,Y_1<\beta^{-1}(b)\big] = b - \int_h^b\int_s^b \frac{G(\beta^{-1}(x))}{G(\beta^{-1}(b))}dxd\frac{\Omega(s)}{\Omega(b)}. \tag{A.2}$$

I now prove the above simplification: write the expected payment as a double integral:

$$E[\max\{s,\beta(Y_1)\}|s<b\&Y_1<\beta^{-1}(b)] = \int_0^b \frac{1}{G(\beta^{-1}(b))}\bigg[\int_0^{\beta^{-1}(s)} sdG(Y_1)+\int_{\beta^{-1}(s)}^{\beta^{-1}(b)}\beta(Y_1)dG(Y_1)\bigg]d\frac{\Omega(s)}{\Omega(b)}. \tag{A.3}$$

The material in the square bracket simplifies as follows:

$$\begin{aligned}&\int_0^{\beta^{-1}(s)} sdG(Y_1) + \int_{\beta^{-1}(s)}^{\beta^{-1}(b)}\beta(Y_1)g(Y_1)dY_1\\ &= G(\beta^{-1}(s))s + G(\beta^{-1}(b))b - G(\beta^{-1}(s))s - \int_{\beta^{-1}(s)}^{\beta^{-1}(b)} G(Y_1)\beta'(Y_1)dY_1\\ &= G(\beta^{-1}(b))b - \int_s^b G(\beta^{-1}(x))dx,\end{aligned}$$

where the second line follows from the first line using integration by parts and a subsequent change of variables $x=\beta(Y_1)$. Plugging the simplified material into Equation (A.3) yields Equation (A.2). In words, Equation (A.2) shows the expected payment of a winner who randomly faces a soft floor below his bid b involves a discount below b (the inside integral), which in turn depends on the realized s: Given Equation (A.2), Equation (A.1) simplifies to

$$\max_b \underbrace{G(\beta^{-1}(b))(v-b)}_{\text{1PSB surplus}} + \underbrace{\int_h^b\int_s^b G(\beta^{-1}(x))dxd\Omega(s)}_{\text{2PSB discount when } s<b}. \tag{A.4}$$

The first-order condition of the bidding problem is

$$\underbrace{g(\beta^{-1}(b))\frac{1}{\beta'(\beta^{-1}(b))}(v-b)}_{\text{winning more often}} - \underbrace{G(\beta^{-1}(b))}_{\substack{\text{paying more in}\\ \text{a simple 1PSB}}} + \underbrace{G(\beta^{-1}(b))\Omega(b)}_{\text{larger discount when } s<b} = 0, \tag{A.5}$$

where the first two terms are the same as in a textbook solution of a 1PSB problem, and the third term arises from the hidden soft floor. In a symmetric equilibrium, it must be that $b=\beta(v)$, and so the equilibrium bidding function must satisfy the differential equation in Equation (4). The differential equation does not have a closed-form solution, but the Peano existence theorem implies a solution exists whenever the right-hand side (RHS) of Equation (4) is continuous in (v,β), for which a sufficient condition is that $g(y)=(N-1)f(y)F^{N-2}(y)$ and $\Omega(v)$ are continuous. Q.E.D. Proposition 2.

Proof of Lemma 1. First consider the incentives of a regular bidder: if the high bidders indeed bid their valuations as suggested by Proposition 1, the regular bidders assume they can only win when no high bidder is present, and so they behave the same as in an auction without high bidders. A deviation by the highest regular bidder $v=1$ to a bid above L that would compete with the high bidders is not profitable, because it only changes the outcome of the game when it actually beats a high bidder and so results in a payment above L that must exceed the regular bidder's valuation by construction. Because the highest regular bidder does not deviate, neither do other regular bidders.

Second, consider the incentives of a high bidder who happens to participate and who should bid his valuation under the proposed equilibrium. The only nontrivial deviation I need to analyze is bidding s or below to trigger 1PSB pricing, resulting in winning much less often but also paying less. Given the high bidder's valuation level, bidding exactly s is the best such deviation from all possible bids weakly below s (see Case 3(c) in the Proof of Proposition 1 for a mathematical argument for why s is the best deviation from all possible bids weakly below s). This deviation is not profitable, because it foregoes both the positive surplus available by possibly beating the other high bidders should they also participate and the positive surplus available by beating high-valuation regular bidders should the other high bidder stay out:

$$\begin{aligned}E[surplus\,|\,bid=v] &= \underbrace{(1-\alpha)^{K-1}}_{\substack{\text{other high}\\ \text{bidders out}}}\underbrace{[v-E(price\,|\,s)]}_{\substack{\text{win for sure, pay either}\\ X_1 \text{ if } X_1>v^*,\ s \text{ otherwise}}}\\ &\quad+\overbrace{\underbrace{[1-(1-\alpha)^{K-1}]}_{\substack{\text{at least 1 other}\\ \text{high bidder in}}}\underbrace{E\big[surplus\,|\,other(s)in\big]}_{\substack{\rightarrow\text{win if present high bidder(s)}\\ \text{have valuations below } v}}}^{>0 \text{ by properties of 2PSB}}\\ &> \underbrace{(1-\alpha)^{K-1}}_{\substack{\text{other high}\\ \text{bidders out}}}\Bigg[\underbrace{F^N(v^*)(v-s)}_{X_1<v^*\rightarrow s \text{ kicks in } 1} + \overbrace{\underbrace{(1-F^N(v^*))[v-E(X_1|X_1>v^*)]}_{X_1>v^*\rightarrow \text{pay } X_1}}^{>0}\Bigg]\\ &> (1-\alpha)^{K-1}F^N(v^*)(v-s) = E\big[surplus|bid=s\big].\end{aligned} \tag{A.6}$$

In words, bidding v yields an additional expected surplus compared with bidding (and paying) s. The additional expected surplus arises from winning more often, namely, when the other high bidder(s) is (are) out and the highest regular competitor is above v^*, and also from the 2PSB auction that results when the other high competitor(s) is (are) in but his (their) valuation is below v. Q.E.D. Lemma 1.

Verification of Revenue Equivalence When Soft Floor Is Common Knowledge and Participation Is Random

Let $N = 2$ in Section 4.2. The 1PSB equilibrium bidding function is the following, easily derived by substituting $G(v) = (1-\alpha) + \alpha F(v)$ in Equation (1) to account for the increased probability of winning:

$$\beta_\alpha(v) = h\frac{(1-\alpha)+\alpha F(h)}{(1-\alpha)+\alpha F(v)} + \frac{\alpha}{(1-\alpha)+\alpha F(v)} \int_h^v z dF(z). \tag{A.7}$$

The auctioneer's expected revenue depends on how many, if any, bidders are present:

$$\begin{aligned}\Pi(h,s) = \underbrace{\alpha^2}_{\Pr(2\text{ bidders})} & \left\{ \begin{aligned} & \overbrace{\Pr(h \le X_1 < v^*)E_{X_1}[\beta_\alpha(X_1)|h \le X_1 < v^*]}^{\text{highest bid below soft floor} \to \text{1PSB with reserve of } h} + \\ & \underbrace{\begin{aligned}\Pr(X_1 > v^*)E_{X_2}[s\Pr(\beta_\alpha(X_2) < s) \\ +X_2\Pr(\beta_\alpha(X_2) > s)|X_1 > v^*]\end{aligned}}_{\text{highest bid above soft floor} \to \text{2PSB with reserve of } s} \end{aligned} \right\} \\ + \underbrace{2\alpha(1-\alpha)}_{\Pr(\text{only 1 bidder})} & \left\{ \begin{aligned} & \overbrace{\Pr(h \le X < v^*)E_X[\beta_\alpha(X)|h \le X < v^*]}^{\text{the only bid below soft floor} \to \text{1PSB with reserve of } h} \\ & + \underbrace{\Pr(X > v^*)s}_{\substack{\text{the only bid above soft floor} \\ \to \text{price is the reserve } s}} \end{aligned} \right\},\end{aligned} \tag{A.8}$$

where X_1 is the highest of two valuations distributed $\sim F^2$, and where X_2 is the second-highest of two valuations, and its distribution conditional on $X_1 = x$ is just $\frac{F(\cdot)}{F(x)}$: the highest of N-1 = 1 draws from F conditional on the draws being below x. Plugging these distributions of the order statistics into Equation (A.8), the auctioneer's expected revenue becomes

$$\begin{aligned}\Pi(h,s) = {} & \alpha^2 \int_h^{v^*} \beta_\alpha(x) dF^2(x) \\ & + \alpha^2 \int_{v^*}^1 \left[\frac{sF(v^*)}{F(x)} + \frac{1}{F(x)} \int_{v^*}^x z dF(z)\right] dF^2(x) \\ & + 2\alpha(1-\alpha) \int_h^{v^*} \beta_\alpha(x) dF(x) \\ & + 2\alpha(1-\alpha)s[1-F(v^*)].\end{aligned} \tag{A.9}$$

Collecting terms yields

$$\begin{aligned}\frac{\Pi(h,s)}{2\alpha} = {} & \int_h^{v^*} \beta_\alpha(x)[(1-\alpha)+\alpha F(x)] dF(x) \\ & + \int_{v^*}^1 \left[s[(1-\alpha)+\alpha F(v^*)] + \alpha \int_{v^*}^x z dF(z)\right] dF(x).\end{aligned} \tag{A.10}$$

Now note that $\beta_\alpha(v^*) = s$, so one can substitute for s using Equation (A.7) as follows:

$$s[(1-\alpha)+\alpha F(v^*)] = h[(1-\alpha)+\alpha F(h)] + \alpha \int_h^{v^*} z dF(z).$$

Figure A.2. 1PSB Bidding Strategies with One Regular and One Randomly Present High Bidder

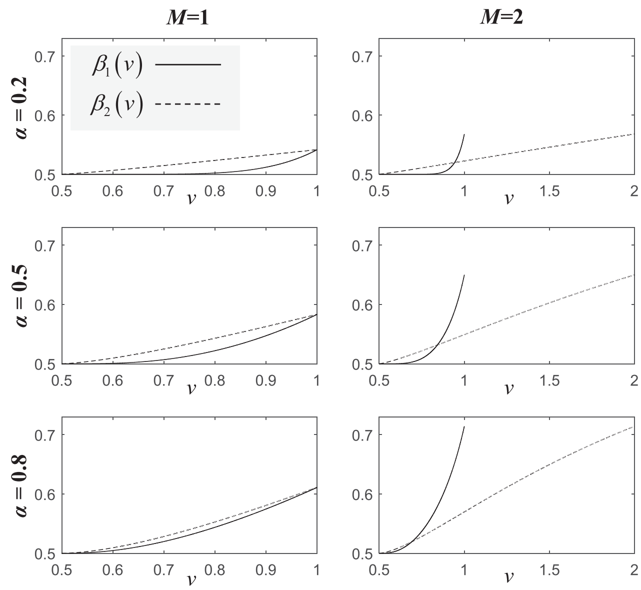

Note. Bidding functions implied by Proposition 5 for two values of M and three values of α, with the x-axis corresponding to the support of the high bidder's valuations, $h = 0.5$ throughout, and the y-axis identical in all six subplots.

After this substitution, note the integrand in the second integral (from v^* to 1) is exactly the same as the integrand in the first (h to v^*), and so Equation (A.10) simplifies to

$$\begin{aligned}\Pi(h,s) &= 2\alpha \int_h^1 \left\{h[(1-\alpha)+\alpha F(h)] + \alpha \int_h^x z dF(z)\right\} dF(x) \\ &= \Pi(h),\end{aligned} \tag{A.11}$$

where the last equality emphasizes that the s has no impact on Π, because neither it nor v^* are present. To see why the expected revenue $\Pi(h)$ is exactly the same as in the second-price auction with the same reserve, rearrange Equation (A.11) as follows:

$$\begin{aligned}\Pi(h) = h & \left\{\underbrace{2\alpha(1-\alpha)[1-F(h)]}_{\Pr(1\text{ bidder with } v>h)} + \underbrace{\alpha^2 2F(h)[1-F(h)]}_{\Pr(2\text{ bidders with } v_i<h<v_{-i})}\right\} \\ & + \underbrace{\alpha^2}_{\Pr(2\text{ bidders})} \underbrace{\int_h^1 2z[1-F(z)]dF(z)}_{\Pr(X_2>h)E(X_2|X_2>h)},\end{aligned}$$

where the

$$\int_h^1 2z[1-F(z)]dF(z) = \Pr(X_2 > h)E(X_2|X_2 > h)$$

because the *pdf* of a minimum of two draws is $2z[1-F(z)]$. This concludes the direct proof of revenue equivalence when two randomly participating symmetric bidders are present.

Proof of Proposition 4. Equation (6) derives the formula for the difference in profits $\Delta\Pi(v^*) \equiv \frac{\Pi^A_{2PSB}(h) - \Pi^A(h,s)}{K\alpha(1-\alpha)^{K-1}}$. It is enough to show $\Delta\Pi$ is positive for all $v^* > h$ and increasing in v^*. Omitting the asterisk on v for clarity, plug in the distributions of X_1 and Y_1 in terms of F:

$$\Delta\Pi(v) = F^N(v)\left[h\frac{F^N(h)}{F^N(v)} + \frac{1}{F^N(v)}\int_h^v x dF^N(x) - h\frac{F^{N-1}(h)}{F^{N-1}(v)} - \frac{1}{F^{N-1}(v)}\int_h^v x dF^{N-1}(x)\right]$$
$$= \int_h^v x dF^N(x) - F(v)\int_h^v x dF^{N-1}(x) - hF^{N-1}(h)\,(F(v) - F(h)).$$

Therefore, $\Delta\Pi(h) = 0$. To show $\Delta\Pi(v) > 0$ and $\Delta\Pi'(v) > 0$ for $v > h$, it is enough to show $\Delta\Pi'(v) > 0$ for all $v > h$:

$$\Delta\Pi'(v) = f(v)\left[vNF^{N-1}(v) - \int_h^v x dF^{N-1}(x) - F(v)v(N-1)F^{N-2}(v) - hF^{N-1}(h)\right]$$
$$= f(v)\left[vF^{N-1}(v) - (F^{N-1}(v) - F^{N-1}(h))E(Y_1|h < Y_1 < v) - hF^{N-1}(h)\right]$$
$$= f(v)\left\{[F^{N-1}(v) - F^{N-1}(h)]\underbrace{[v - E(Y_1|h < Y_1 < v)]}_{>0} + F^{N-1}(h)(v-h)\right\} > 0,$$

where the last line follows from the previous expression by adding and subtracting $f(v)vF^{N-1}(h)$. Q.E.D. Proposition 4.

Details of Section 6: Expected Revenue Under the Second-Price Rule (2PSB) in the Market with One Regular Uniform Bidder and One Stochastically Dominant Randomly Appearing Bidder

The expected revenue π_{II} of a 2PSB auction is straightforward to derive because the bidders bid their valuations as a dominant strategy. As a function of h, the expected revenue function is not necessarily globally concave, because any $h > 1$ excludes the regular low bidder, effectively acting as a posted price for the high bidder. The expected revenue from such a high hard floor is obviously $E(\pi_{II}\,|\,h > 1) = \frac{\alpha h(M-h)}{M}$, so the optimal high hard floor is $h_{II}^H \equiv \frac{M}{2}$, which exceeds 1 whenever $M \geq 2$ and yields the expected revenue of $E(\pi_{II}\,|\,h_{II}^H) = \frac{\alpha M}{4}$. Alternatively, the seller can select an $h \leq 1$, which engages the regular bidder. The resulting expected revenue is as follows:

Claim 4.

$$E(\pi_{II}|h \leq 1) = \frac{\alpha}{6M}(3M - 1 + 3h^2(1+M) - 8h^3) + (1-\alpha)h(1-h). \tag{A.12}$$

Proof of Claim 4. When the high bidder does not enter, the revenue is $h(1-h)$ because the low bidder wins and pays the reserve price whenever his valuation is above the reserve price—a probability of $1 - h$. When the high bidder does enter, four distinct revenue regions of the (v_1, v_2) space exist:

(1) With probability $\frac{h^2}{M}$, $\max v_i < h$, and so the revenue is zero.

(2) With probability $\frac{h(M-h)+h(1-h)}{M}$, $v_i < h < v_{-i}$, and the revenue is thus h.

(3) With probability $\frac{(M-1)(1-h)}{M}$, $h < v_1 < 1 < v_2$, so the high bidder wins for sure and pays the expected conditional valuation of the low bidder $E(v_1|v_1 > h) = \frac{1+h}{2}$.

(4) With probability $\frac{(1-h)^2}{M}$, $h < v_i < 1$, so the bidding competition is as if two *iid* bidders from Uniform[h,1] exist, and the relevant revenue is thus $E(\min v_i|h < v_i > 1) = \frac{1+2h}{3}$.

Combining the above four cases with the possibility of the high bidder staying out yields Equation (A.12). This concludes the proof of the Claim.

The optimal hard floor below 1, denoted h_{II}^L, optimizes the expected revenue in Equation (A.12). The solution is in closed form because the first-order condition of $\max_{h\leq 1} E(\pi_{II}|h \leq 1)$ is quadratic. It is possible to show that $h_{II}^L < 1 \Leftrightarrow \alpha < \frac{M}{2M-3}$. Therefore, $\alpha > \frac{M}{2M-3}$ is a sufficient condition for $E(\pi_{II}|h_{II}^H) > E(\pi_{II}|h_{II}^L)$ because it implies $E(\pi_{II}|h < 1)$ is increasing in h on the [0,1] interval. However, it is clearly possible for $E(\pi_{II}|h_{II}^H) > E(\pi_{II}|h_{II}^L)$ even when $h_{II}^L < 1$—the h_{II}^L may only be a local minimum.

The globally optimal hard floor h_{II}^* is clearly h_{II}^L or h_{II}^H, depending on which leads to higher expected revenue. I now show $h_{II}^* = h_{II}^L$ for all α as long as M is small, namely as long as $M^3 - 21M^2 + 51M - 15 > 0 \Leftrightarrow M <\approx 2.4$, above which point h_{II}^H becomes globally optimal for high-enough α. To see the condition, note first that $h_{II}^L > 1/2 =$ the hard floor optimal without the high bidder. For any $\frac{1}{2} < h < 1$, the difference $E(\pi_{II}|h < 1) - \frac{\alpha M}{4}$ decreases in α. Therefore, for a given M, the optimal hard floor is below 1 for all α as long as $E(\pi_{II}|h_{II}^L, \alpha = 1) > \frac{M}{4}$. When $\alpha = 1$, $h_{II}^L = \frac{1+M}{4}$, the associated revenue is $\frac{M^3+3M^2+51M-15}{96M}$, which exceeds $\frac{M}{4}$ whenever $M^3 - 21M^2 + 51M - 15 > 0$.

Proof of Proposition 5. Following Kaplan and Zamir (2012), let $v_i(b)$ be the inverse bidding function of bidder i and look for a pure-strategy equilibrium whence both bidders bid inside an interval $[h, \bar{b}]$, with initial conditions $v_i(h) = h$. The two bidders solve the following optimization problems, respectively:

$$\max_{b\geq h}\left[\alpha\left(\frac{v_2(b)}{M}\right) + 1 - \alpha\right](v_1 - b)$$
$$\max_{b\geq h}[v_1(b)](v_2 - b), \tag{A.13}$$

where the square brackets are the probabilities of winning with a bid of b. For $v_i > h$, the $b \geq h$ constraints are not binding, and so first-order conditions are necessary for optimality. The first-order conditions imply the two inverse bidding functions must satisfy the following system of differential equations:

$$\alpha v_2'(b)(v_1(b) - b) = \alpha v_2(b) + M(1-\alpha)$$
$$v_1'(b)(v_2(b) - b) = v_1(b). \tag{A.14}$$

Multiplying both sides of the second equation by α and adding to the first equation yields

$$\underbrace{v_2'(b)v_1(b) + v_1'(b)v_2(b)}_{[v_2(b)v_1(b)]'} = \underbrace{bv_1'(b) + bv_2'(b) + v_1(b) + v_2(b)}_{[b(v_1(b)+v_2(b))]'} + A, \text{ where } A = \frac{M(1-\alpha)}{\alpha}. \tag{A.15}$$

Integrating both sides yields the following relationship between the two functions, up to a constant B:

$$v_2(b)v_1(b) = b[v_1(b) + v_2(b)] + bA + B. \tag{A.16}$$

Because $v_1(h) = h$, the constant B must satisfy $B = -h(h + A)$. Plugging this expression for B into Equation (A.16) and rearranging yields $v_2(b)$ in terms of $v_1(b)$:

$$v_2(b) = \frac{b[v_1(b) + A] - h(h + A)}{v_1(b) - b}. \tag{A.17}$$

Finally, plug the $v_2(b)$ in Equation (A.17) into the first-order condition of the high bidder (the second equation in 21), and obtain a differential equation that includes only the $v_1(b)$:

$$v_1'(b)(b-h)(b+h+A) = v_1(b)[v_1(b) - b]. \tag{A.18}$$

To solve Equation (A.18), divide both sides by $v_1^2(b)(b-h)^{1+\theta}(b+h+A)^{2-\theta}$ and rearrange to obtain

$$\frac{v_1'(b)}{v_1^2(b)(b-h)^{\theta}(b+h+A)^{1-\theta}} + \frac{b}{v_1(b)(b-h)^{1+\theta}(b+h+A)^{2-\theta}} = \frac{1}{(b-h)^{1+\theta}(b+h+A)^{2-\theta}}. \tag{A.19}$$

Observe that when we set $\theta = \frac{h}{A+2h}$, the LHS of Equation (A.19) can be expressed as a derivative:

$$\frac{v_1'(b)}{v_1^2(b)(b-h)^{\theta}(b+h+A)^{1-\theta}} + \frac{b + \overbrace{[A\theta - h(1-2\theta)]}^{0}}{v^1(b)(b-h)^{1+\theta}(b+h+A)^{2-\theta}} = \left(\frac{-1}{v_1(b)(b-h)^{\theta}(b+h+A)^{1-\theta}}\right)'. \tag{A.20}$$

The RHS of Equation (A.19) integrates, up to a constant, as

$$\int \frac{1}{(b-h)^{1+\theta}(b+h+A)^{2-\theta}}db = \frac{-1}{(b-h)^{\theta}(b+h+A)^{1-\theta}}\left(\frac{A+b}{h(A+h)}\right) + const. \tag{A.21}$$

Combining Equations (A.20) and (A.21) thus yields the following solution of the differential Equation (A.18) up to some constant C:

$$v_1(b) = \left[\frac{(A+b)}{h(A+h)} + C(b-h)^{\theta}(b+h+A)^{1-\theta}\right]^{-1}. \tag{A.22}$$

The fact that $v_1(\bar{b}) = 1$ fixes the constant to $C = \frac{1-\frac{(A+\bar{b})}{h(A+h)}}{(\bar{b}-h)^{\theta}(\bar{b}+h+A)^{1-\theta}}$, so the solution becomes

$$v_1(b) = \frac{h(A+h)}{A + b + [h(A+h) - (A+\bar{b})]\left(\frac{b-h}{\bar{b}-h}\right)^{\theta}\left(\frac{b+h+A}{\bar{b}+h+A}\right)^{1-\theta}}. \tag{A.23}$$

Because $v_1(\bar{b}) = 1$ and $v_2(\bar{b}) = M$, Equation (A.17) fixes the upper support of bidding to $\bar{b} = \frac{M+h(h+A)}{1+M+A} = \frac{hM+\alpha(h^2+M-hM)}{\alpha+M}$. Observe that $\lim_{\alpha\to 0}\bar{b} = h$, so $\lim_{\alpha\to 0} E(\pi_I) = h(1-h)$. Q.E.D. Proposition 5.

Proof of Lemma 2. When h < 1, the distribution G of seller revenue π_I is

$$\begin{aligned} G(x) \equiv \Pr(\pi_I < x) &= \Pr(\beta_1(v_1) < x)[1 - \alpha + \alpha \Pr(\beta_2(v_2) < x)] \\ &= v_1(x)[1 - \alpha + \alpha v_2(x)]. \end{aligned} \tag{A.24}$$

The expected revenue then follows from G by a single integration over possible revenue levels:

$$\begin{aligned} E(\pi_I | h < 1) &= \int_h^{\bar{b}} x dG(x) = -\int_h^{\bar{b}} x(1 - G(x))' dx \\ &= -\bar{b}\underbrace{(1 - G(\bar{b}))}_{0} + h\underbrace{(1 - G(h))}_{\Pr(v_i < h \text{ or high out})} + \int_h^{\bar{b}} 1 - G(x)dx \\ &= h[1 - h(1 - \alpha + \alpha h)] + \int_h^{\bar{b}} 1 - v_1(x)[1 - \alpha + \alpha v_2(x)]dx. \end{aligned} \tag{A.25}$$

Q.E.D. Lemma 2.

Endnotes

[1] The use of soft floors is widespread. At least AdX, OpenX, and AppNexus exchanges have used them in the U.S. market. When used, they tend to affect many transactions: Yuan et al (2013) analyze a bidding platform on which more than half of the transactions involve a price equal to the winner's bid, that is, an active soft floor.

[2] Note that randomly present high-valuation bidders are just one possible asymmetry. This paper does not attempt to characterize the impact of soft floors in all possible asymmetric markets, but instead focuses on a particular asymmetry used in the industry to justify the soft-floor practice.

References

Abraham I, Athey S, Babaioff M, Grubb MD (2016) Peaches, lemons, and cookies: Designing auction markets with dispersed information. Working paper, Stanford University, Stanford, CA.

Arnosti N, Beck M, Milgrom P (2016) Adverse selection and auction design for Internet display advertising. *Amer. Econom. Rev.* 106(10): 2852–2866.

Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* 61(4):864–884.

Bulow J, Klemperer P (1996) Auctions vs. negotiations. *Amer. Econom. Rev.* 86(1):180–194.

Celis LE, Lewis G, Mobius M, Nazerzadeh H (2014) Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Sci.* 60(12):2927–2948.

Chattopadhyay S, Chatterjee R (2012) Revenue equivalence in sequential auctions. *Econom. Bull.* 32(2):1272–1281.

Chen Y (2017) Programmatic advertising is preparing for the first-price auction era. Accessed December 1, 2017, https://digiday.com/marketing/programmatic-advertising-readying-first-price-auction-era/.

Choi H, Mela CF, Balseiro S, Leary A (2017) Online display advertising markets: A literature review and future directions. Research paper, Columbia Business School, New York.

emarketer (2016) More than two-thirds of US digital display ad spending is programmatic. Accessed December 1, 2017, https://www.emarketer.com/.

Engelbrecht-Wiggans R (1994) Sequential auctions of stochastically equivalent objects. *Econom. Lett.* 44:87–90.

Försch S, Heise M, Skiera B (2017) The impact of floor prices in real-time online display advertising auctions on publisher's profit. Working paper, University of Frankfurt, Frankfurt, Germany.

Friedman J (2015) I study billions of online ad impressions to know the truth. Accessed December 1, 2017, https://www.quora.com/How-do-I-estimate-the-total-number-of-online-ad-impressions-that-might-be-seen-by-internet-users-over-a-30-day-period-for-a-given-U-S-DMA.

Goldfarb A (2014) What is different about online advertising? *Rev. Indust. Organ*. 44(2):115–129.

Griesmer J, Levitan R, Shubik M (1967) Toward a study of bidding processes, part four: Games with unknown costs. *Naval Res. Quart.* 14(4):415–443.

Harstad R, Kagel J, Levin D (1990) Equilibrium bid functions for auctions with an uncertain number of bidders. *Econom. Lett*. 33(1):35–40.

Hoban PR, Bucklin RE (2015) Effects of internet display advertising in the purchase funnel. *J. Marketing Res.* 52(3):375–393.

Johnson GA, Lewis RA, Nubbemeyer EI (2017) Ghost ads: A revolution in measuring ad effectiveness. *J. Marketing Res*. 54(6):867–884.

Kaplan T, Zamir S (2012) Asymmetric first-price auctions with uniform distributions: Analytic solutions to the general case. *Econom. Theory* 50(2):269–302.

Krishna V (2009) *Auction Theory*, 2nd ed. (Academic Press, San Diego).

Maskin E, Riley J (2000) Asymmetric auctions. *Rev. Econom. Stud.* 67(3):413–438.

Milgrom PR, Weber R (2000) A theory of auctions and competitive bidding II. Klemperer P, ed. *The Economic Theory of Auctions* (Edward Elgar, Cheltenham, UK), 1–16.

Moesman S (2017) Three ways first-price auctions will expand programmatic transparency and revenue. Accessed December 1, 2017, http://360.advertisingweek.com/3-ways-first-price-auctions-will-expand-programmatic-transparency-revenue/.

Myerson RB (1981) Optimal auction design. *Math. Oper. Res*. 6(1):58–73.

Nolet M (2010) Price floors, second price auctions and market dynamics. Accessed December 1, 2017, http://www.mikeonads.com.

Reiß JP, Schöndube JR (2010) First-price equilibrium and revenue equivalence in a sequential procurement auction model. *Econom. Theory* 43(1):99–141.

Rothkopf MH, Teisberg TJ, Kahn EP (1990) Why are Vickrey auctions rare? *J. Political Econom.* 98(1):94–109.

Sahni D (2016) How does the mobile RTB auction process work? Accessed December 1, 2017, https://www.applift.com/blog/real-time-bidding-auction.

Singel R (2010) Oct. 27, 1994: Web gives birth to banner ads. *Wired* (October 27), https://www.wired.com/2010/10/1027hotwired-banner-ads/.

Strong E (2012) Second-guessing the second-price auction model. Sell-sider column on Ad Exchanger website. Accessed December 1, 2017, https://adexchanger.com/data-driven-thinking/second-guessing-the-second-price-auction-model/.

Weatherman K (2013) Introducing soft price floors. MoPub platform announcement. Accessed December 1, 2017, http://www.mopub.com/2013/04/12/introducing-soft-price-floors/.

Yuan S, Wang J, Zhao X (2013) Real-time bidding for online advertising: Measurement and analysis. *Proc. 7th Internat. Workshop Data Mining Online Advertising* (ACM, New York).