Optimal Reverse-Pricing Mechanisms

Martin Spann [1]   **Robert Zeithammer** [2]   **Gerald Häubl** [3]

February 19, 2010

Abstract

Reverse pricing is a market mechanism under which a consumer's bid for a product leads to a sale if the bid exceeds a hidden acceptance threshold the seller has set in advance. The seller faces two key decisions in designing such a mechanism: First, he must decide where in the process to collect the revenue—that is, whether to commit to a minimum markup above cost (and thus define the bid-acceptance threshold given cost) and whether to set a fee for the consumer's right to bid. Second, the seller must decide whether to facilitate or hinder consumer learning about the current bid-acceptance threshold. We analyze these decisions for a profit-maximizing small intermediary retailer selling to consumers who can also purchase the product in an outside posted-price market. The optimal revenue model is to charge a fee for the right to bid and then accept all bids above cost, rather than to set a positive minimum markup above cost. Avoiding minimum markups in favor of a bidding fee is more profitable because of increased efficiency arising from more entry by consumers and higher bids by the entrants. When consumers learn about the bid-acceptance threshold before they enter the market, efficiency increases further, and generating revenue through a bidding fee can compensate the seller for his loss of information rent when the competition from the outside posted-price firm is relatively weak.

Keywords: reverse pricing, name-your-own-price, analytical modeling, e-commerce.

---

[1] Martin Spann: Munich School of Management, Ludwig-Maximilians-University (LMU), Geschwister-Scholl-Platz 1, 80539 München, Germany, Phone: +49 (89) 2180-72050, Fax: +49 (89) 2180-72052, E-mail: spann@spann.de

[2] Robert Zeithammer: The Anderson School of Management, UCLA, Los Angeles, USA, Phone: (310) 825-1862, E-mail: robert.zeithammer@anderson.ucla.edu

[3] Gerald Häubl: School of Business, University of Alberta, Edmonton, AB, Canada T6G 2R6, Phone: (780) 492-6886, Fax: (780) 492-3325, E-mail: gerald.haeubl@ualberta.ca

1   Introduction

Reverse pricing (often also referred to as "name-your-own-price selling") is an emerging market mechanism under which a consumer's bid for a specific product leads to a sale if this bid exceeds the seller's hidden acceptance threshold, which in turn depends on the seller's marginal cost (Spann and Tellis 2006). If a bid exceeds the seller's threshold, the consumer receives the product and pays the bid amount. Some current examples of reverse-pricing sellers are Priceline (which specializes in airline tickets, hotel accommodations, and car rentals), some eBay sellers (with eBay's "Best Offer" feature), and discount airlines such as Germanwings.

Reverse pricing is popular as a spot market for handling unexpected demand or supply fluctuations in industries with volatile cost and/or demand, such as the travel industry. Vendors such as Expedia post retail prices based on demand predictions, but the residual uncertainty arising from consumers' actual booking behavior frequently results in excess capacity, which airlines may sell through alternative channels (Koenigsberg, Muller, and Vilcassim 2008). In this paper, we focus on selling excess capacity through a reverse-pricing intermediary retailer, such as Priceline, alongside posted-price retailers. The marginal cost of the excess capacity varies with every seat sold and hotel room booked, and the strength of reverse pricing is its natural ability to accommodate these fluctuating costs, as we discuss next.

The essence of reverse pricing is reversing the order of actions compared to a standard posted-price market, in which the seller first sets the price given his marginal cost and then waits for consumers. A reverse-pricing seller solicits consumer bids first and then evaluates them in light of current marginal costs. This process allows the seller to use a bid-acceptance rule that only accepts profitable bids, that is, bids above the cost of getting the good from a producer at the

moment. In this paper, we analyze the seller's problem of setting such a bid-acceptance rule to maximize profits. The problem is essentially static because our seller is an intermediary and neither the good nor the consumers persist over time: the producer's offer price (the seller's marginal cost) is a temporary exploding offer, so the seller cannot keep the good and sell it later. The prospective customers are those in the market for that particular product or service at that moment, and they are likely to buy the good elsewhere if unsuccessful with reverse pricing. For example, if Priceline does not sell a given seat on a flight to a customer interested in buying it right now, both the seat and the customer may disappear before another opportunity for a trade arises. We thus abstract away from dynamic yield-management issues that would arise if the seller were the producer of the good.

How should our seller design the reverse-pricing mechanism so as to maximize profit? The specification of a seller strategy involves two key decisions. First, where in the interaction with consumers should the seller generate revenue: by relying entirely on the information rent from successful bids (i.e., the difference between the bid amount and cost), by setting a positive minimum markup above marginal costs in specifying the bid-acceptance threshold, by charging consumers a bidding fee (akin to a subscription fee) for the right to submit a bid, or by some combination of these sources? Second, the seller must decide whether to help or hinder the consumers in their learning about the bid-acceptance threshold. Consumers may learn about the current threshold via repeated bidding (Fay 2004) or through communication with other consumers (Hinz and Spann 2008). Given that such learning is typically considered destructive to profits (Fay 2004; Hinz and Spann 2008; Segan 2005), understanding whether the seller is in fact better off preventing consumers from such learning is important. The answer to this design question is intertwined with the answer to the revenue-source question: a seller relying on the

information rent would clearly prefer to hinder learning, whereas a seller relying on an upfront bidding fee may benefit from such learning because resolving the consumers' uncertainty may increase their willingness to pay the fee.

We analyze both seller decisions using a model of a small intermediary seller who faces a population of consumers and a premium posted-price competitor. Regarding the first decision (revenue model), we find that under several different assumptions about consumers, bidding fees dominate minimum markups in terms of expected profit. In other words, the seller should never try to commit to a positive minimum markup. Instead, the seller is always better off lowering the minimum markup slightly and charging a correspondingly higher bidding fee. Bidding fees are a superior source of revenue because they allow more trades to occur via an implied lower bid-acceptance threshold, thereby making the market more efficient. Although similar to the intuition for the superiority of a two-part tariff to simple linear pricing (Tirole 1988, p. 136), the argument for bidding fees is more subtle: unlike in the case of a simple two-part tariff, our reverse-pricing seller is not a monopolist and does not capture the entire consumer surplus because of strategic bid-shading (bidding less than valuation) by the consumers. In addition to their profit advantage, bidding fees are also easier to implement than minimum markups because they require less commitment from the seller: a minimum markup seller needs to credibly promise he will reject some profitable trades in the future in order to increase bids today.

Regarding the second decision, the seller who makes revenue through bidding fees sometimes prefers to facilitate consumers' learning about the bid-acceptance threshold instead of trying to prevent them from discovering it. We find that the seller prefers to facilitate consumer learning about the threshold when the outside price is relatively high, that is, when the outside competition is relatively weak. The additional information assures that only consumers with

4

high-enough valuations enter and submit bids, greatly increasing market efficiency. A high outside price gives the seller market power over more consumers, allowing him to capture more of the efficiency gain. Bidding fees thus partially "immunize" the seller against the low-profit consequences of consumers becoming informed about the threshold: a seller relying only on the information rents for revenue would always prefer to hide threshold information from consumers.

Most prior research on reverse pricing empirically analyzes bidding behavior and focuses on either behavioral aspects (Chernev 2003; Ding, Eliashberg, Huber, and Saini 2005; Spann and Tellis 2006), the inference of certain consumer characteristics from observed bidding behavior (Hann and Terwiesch 2003; Spann, Skiera, and Schäfers 2004), or the impact of information diffusion about the seller's bid-acceptance threshold on bidding behavior (Hinz and Spann 2008). Four recent articles are more closely relevant to the present work in that they examine some aspects of the seller strategy. Terwiesch, Savin, and Hann (2005) discuss the influence of frictional costs on the optimal bid-acceptance threshold. Wang, Gal-Or, and Chatterjee (2009) and Fay (2009) provide two different rationales for the existence of the reverse-pricing channel in addition to a posted-price channel, either used by the same retailer or used by two different retailers, respectively. Amaldoss and Jain (2008) analyze joint bidding for multiple items at a reverse-pricing retailer and find that such joint bids can increase retailer profit.

The most closely related piece of prior work is Fay (2004), who examines whether the seller should limit consumers' ability to learn his bid-acceptance threshold. That work focuses on learning through rebidding as opposed to learning through communication among consumers. Contrary to the common belief that repeat bidding erodes profits by revealing information about the seller's threshold, Fay's findings suggest that repeat bidding may leave profits unaffected as long as the seller anticipates rebidding and adjusts his threshold accordingly. Moreover, they also

indicate that the seller may sometimes wish to encourage rather than discourage repeat bidding. The key assumption of Fay's model is that only two possible thresholds exist, arising from two possible inventory levels. This assumption both simplifies bidding and gives substantial power to the seller through contingent pricing. In contrast, our model allows a continuum of thresholds in equilibrium, increasing realism. Our model is also complementary to Fay (2004) in that his seller is initially uncertain about inventory, whereas our seller is uncertain about marginal cost.

In a nutshell, the contributions of this research relative to prior work on reverse pricing are as follows: We conceptualize the reverse-pricing seller as an intermediary uncertain about his marginal costs. The resulting model allows analysis and optimization of two key parameters of a reverse-pricing mechanism that are under the seller's control: revenue model and information release. We show that the optimal reverse-pricing mechanism charges a positive bidding fee and accepts all bids above the marginal cost rather than charging a minimum markup above the marginal cost akin to a "transaction fee." Moreover, we account for the possibility of consumers discovering the bid-acceptance threshold and show when it is in the seller's interest to facilitate such learning.

The article is organized as follows: section 2.1 lays out our assumptions. In section 2.2, we develop a model of consumer bidding behavior in a reverse-pricing market with an outside option. Section 2.3 then investigates the optimal strategy of a seller who cannot commit to a minimum markup. We then allow the seller to commit to a specific minimum markup in section 2.4, and positive minimum markups are not optimal. Section 2.5 analyzes whether the reverse-pricing seller should facilitate or hinder consumer learning of the bid-acceptance threshold. We conclude in section 3.

2 Optimal Menu-Based Pricing for a Reverse-Pricing Seller

2.1 Assumptions

We analyze the optimization problem of a small perishable-product intermediary seller who cannot set posted prices contingent on his highly variable marginal costs. Capturing the idea of selling excess capacity beyond some outside market for the good, the seller sets his bidding fee and/or minimum markup after learning the prevailing stable posted price in an outside market p but before learning current marginal cost w. Capturing the idea that the seller selling excess capacity is small relative to the outside market, the seller is a price taker of both p and w. This price-taking assumption ensures that our model does not degenerate into Bertrand price competition due to a lack of differentiation between the seller and the outside market. The seller intermediates between producers and consumers of the product. The product (e.g., producer's excess capacity) is perishable in that neither the consumers' nor the producer's offers persist over time. Consequently, the seller can model his profit-maximization problem as a single-shot game between himself and the consumers who happen to be in the market at the moment. In the Priceline plane-ticket setting, the producers are airlines, the product is a seat on a specific flight, and the product is perishable in that each airline makes a short-lived offer of a particular marginal cost for the seat to Priceline. The consumer is short-lived in that he moves on to buy the ticket elsewhere whenever Priceline does not accept his bid. Next, we discuss our assumptions about the outside market, consumers, and marginal costs in more detail, starting with the former.

Outside market. The posted price in the outside market p is variable but set before the game begins. For example, the posted price for a plane ticket usually remains constant for several days and does not react immediately to activity on Priceline. This price p is informative to the consumers, who infer that the reverse-pricing seller's marginal cost is below p.

Consumers. A unit mass of risk-neutral consumers exists. The consumers have private valuations of the product drawn from some continuous distribution H with support on $[0,1]$. We provide as many results as possible for a general H but resort to $H=Uniform[0,1]$ when a general result is not tractable. Given the normalization of utility to the $[0,1]$ interval, we assume the outside posted price p is always below the highest valuation in the market (i.e. $p \leq 1$). Because the highest consumer valuation exceeds the outside price, two distinct segments of consumers emerge: those with a valuation v less than p ("low consumers") can only buy the product from the reverse-pricing seller, whereas others with $v \geq p$ ("high consumers") can also obtain a positive surplus by buying on the outside posted-price market. Note that these segments are endogenous to p, and so the same consumer can be a "low" consumer when p is high and a "high" consumer when p is low. Since some consumers cannot afford the outside price, the outside market can be considered a premium option. All consumers have unit demand in the sense that once they buy the product from the reverse-pricing seller, they do not buy it again on the outside market. Finally, consumers incur no frictional cost for bidding, an assumption that can be relaxed while preserving the key qualitative results (please contact the authors for details).

Marginal Costs. The producers offer the product to the reverse-pricing seller at marginal cost w, which is between 0 and the outside posted price of p. For analytical tractability, we assume marginal costs are distributed uniformly on $[0, p]$. The critical part of this assumption is that p and w are not perfectly correlated, so consumers remain uncertain about w even after observing p.

Timing of the Game. The difference between our model and a model of standard posted-price retailing is that the marginal cost w is not known when the seller sets his strategy. The seller only queries the producers when a bidder appears and makes a request for a specific product.

Thus the seller must maximize profits by specifying his strategy (a bidding fee and a bid-acceptance rule) before he knows w.[1] He has the following two sources of revenue at his disposal (see Figure 1):

Figure 1: Seller Profit from Consumer's Bid

seller profit (Π)= margin + fee (f)
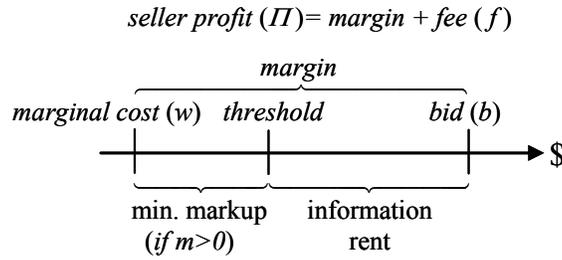
1) The seller can charge each bidder a bidding fee f regardless of whether an individual's eventual bid is successful. This bidding fee enters the bidder's decision-making analogously to an entry cost in an auction (Samuelson 1985). Reverse pricing with a bidding fee is similar to charging consumers for a referral service that matches them with producers but is not equivalent because the seller also keeps the difference between the accepted bid and the marginal cost w.

2) The seller can earn the margin between his marginal cost w and the consumer's bid. This margin consists of the seller's information rent as the difference between his bid-acceptance threshold and the consumer's bid and a potential minimum markup between marginal costs and the bid-acceptance threshold. The seller may or may not be able to commit to the minimum-markup strategy of only accepting bids that are at least a fixed markup $m \geq 0$ above w. Commitment is necessary for the minimum markup strategy to work, because *ex post* (i.e., once w is realized), this strategy foregoes potentially profitable trades (when $w < bid < w+m$)

---

[1] An equivalent assumption to the seller announcing the bid-acceptance rule is that the consumers infer it from recent experience with the same seller and share this information through word-of-mouth (Hinz and Spann 2008).

in order to induce higher bids *ex ante*. Therefore, the seller without commitment is restricted to *m=0*. In this paper, we first solve the game for a seller without this commitment, and then investigate how the seller's strategy would change with the commitment.

Figure 2 illustrates the timing of the game. First, at Stage 0, the posted price *p* is set. After seeing the price *p*, consumers update their prior that the cost is uniformly distributed on [0, 1] to a belief that the cost is uniformly distributed on [0, *p*]. At Stage 1, the seller sets and announces his (m, f) strategy (the seller without commitment is restricted to *m=0* as discussed above). At Stage 2, consumers make their entry decision. Some low consumers do not enter because their expected surplus is negative. The consumers who do enter then submit their bids. Once the seller receives a bid (Stage 3), he queries the producers to determine the current marginal cost and then accepts bids above *w+m*, using the *m* announced in Stage 1. The seller faces competition from the posted-price market because the bidders whose bids are unsuccessful, as well as those who chose not to submit a bid at all, have the opportunity to buy the product on the outside posted-price market at Stage 4 for price *p*.

Figure 2: Timing of the Game

Stage 0	Stage 1	Stage 2	Stage 3	Stage 4
p is set and revealed; consumers update belief about *w* (*w* ≤ *p)*	Seller announces entry fee *f* and minimum markup *m*	One randomly drawn consumer decides whether to enter (for fee *f*) & submits *bid* if enters	Seller learns marginal cost *w* and accepts *bid* if *bid* > *w+m*	Consumer's option to buy at outside posted price *p*

time
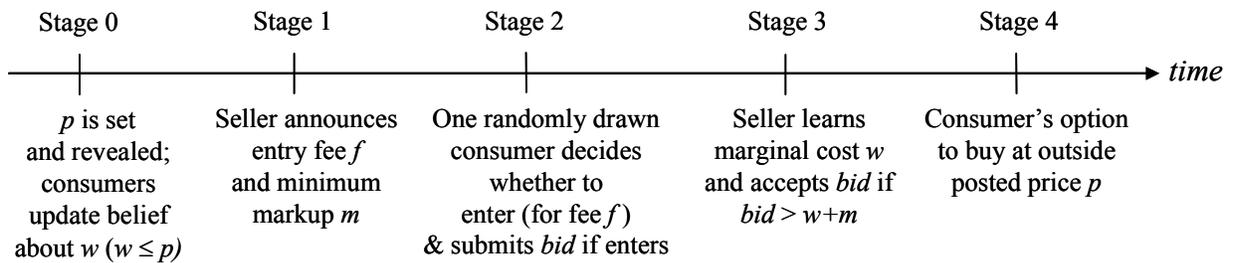
2.2 Consumer strategy: entry decision and subsequent bidding strategy

We solve the consumer's problem in Stage 2 using backward induction, starting with the bidding strategy given entry. We analyze the bidding strategy for the general case of $m \geq 0$, with the consumers facing a non-commitment seller bidding as if $m=0$. Consumers are uncertain about w when they submit their bids in Stage 2. Therefore, they must optimize against a distribution of bid-acceptance thresholds the seller's strategy in Stage 3 implies, namely, Uniform[$m, p +m$]. Note that this distribution is the correct equilibrium belief at this stage of the game, either because of the incentives of the seller without commitment to use $m=0$ or because of the seller's commitment to a particular $m>0$ in Stage 1. We extend the model of Spann and Tellis (2006) to account for the outside option and characterize consumers' bidding behavior at the end of Stage 2 as follows:

Proposition 1: *Consumers who enter bid* $b(v \mid m) = \begin{cases} \dfrac{m+v}{2} & \text{for } m \leq v < p \\ \dfrac{m+p}{2} & \text{for } v \geq p \end{cases}$.

The proofs of Proposition 1 and all subsequent propositions are in the appendix. This bidding problem is analogous to the problem of a bidder in a first-price sealed-bid auction who faces one exogenous competitor. The exogenous competitor bids Uniform[$m, p +m$] . To balance the surplus from winning against the probability of winning, the optimal bidding strategy in such a situation is to "shade" the bid down from one's valuation of the product. As in the first-price sealed-bid auction, winning means paying one's bid, and shading is necessary to leave room for surplus.[2] Since consumers with valuations below m cannot make a positive surplus in this market, it is sufficient to characterize the bidding strategy for $v>m$, for which bid shading is apparent in

---

[2] See Krishna (2002) for an exposition of bidding in a first-price sealed-bid auction.

that $b(v|m) < v$. Also note that the two consumer segments bid differently: because of the outside option, all high consumers with $v > p$ mimic the behavior of a consumer with $v = p$. Finally, the bidding strategy does not depend on the bidding fee f, because f is sunk upon entry. Given the bidding strategy, a comparison of the expected surplus from entering and bidding to the expected surplus from not entering governs the entry decision in the beginning of Stage 2 (Samuelson 1985). Please see Figure 3 for the consumer's decision tree. The expected surplus is derived in the appendix, and it again depends on the consumer segment:

$$(1) \quad S(v|f,m,p) = \begin{cases} \dfrac{(v-m)^2}{4p} - f & \text{for } m \le v < p \\ v - p + \dfrac{(m-p)^2}{4p} - f & \text{for } v \ge p \end{cases}.$$

The functional form of the surplus is instructive with respect to consumers' entry decisions. First, low consumers face quadratic surplus because higher valuations and/or lower margins increase both the probability of bid acceptance and the surplus upon acceptance. These two components multiply to produce the convex function from which the bidding fee is simply subtracted as a direct transfer to the seller. Another insight concerns the high consumers: they receive the outside-option surplus $v - p$ plus a "gambling bonus" equivalent to the surplus of the consumer they mimic in bidding, namely, $v = p$. Since they can obtain the product for the price p without entering, the high consumers will enter only when the gambling bonus exceeds f. Since the gambling bonus is equivalent to the expected surplus of a consumer with $v = p$ and the expected surplus is increasing in v below p, satisfying the consumer with $v = p$ is a prerequisite for the reverse-pricing seller to get any customers at all. When the $v = p$ consumer enters, all high consumers enter as well, and low consumers above a valuation threshold $\underline{v} = m + 2\sqrt{pf}$ enter.

Figure 3: Payoff to Consumers
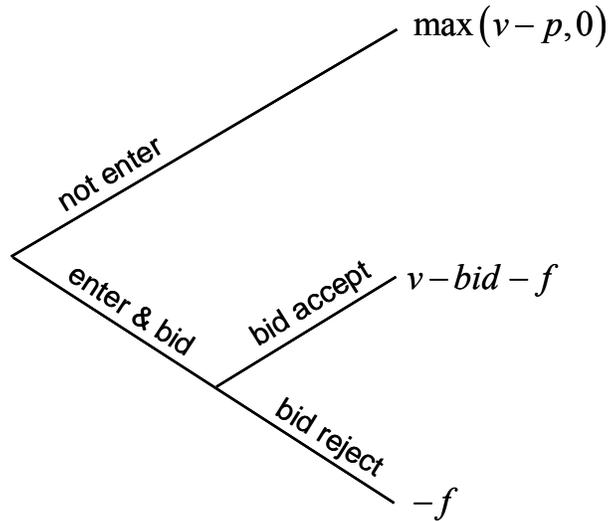
Please note that even if consumers enter, their payoff is uncertain. Thus, some consumers may pay an entry fee and bid but end up paying only the entry fee f when they are unsuccessful. This discussion concludes our analysis of the consumer side in Stage 2. The next section proceeds backwards to Stage 1 to solve the seller's profit-maximization problem.

2.3   Optimal selling without commitment to a minimum markup (*m=0*)

In this section, we derive the optimal strategy of a reverse-pricing seller for the case that he cannot commit to a minimum markup $m > 0$.

At Stage 1, the reverse-pricing seller without commitment ($m = 0$) maximizes his expected profit by selecting the optimal bidding fee f. The profit-maximization is a screening problem because the seller uses the fee to effectively set the entry threshold $\underline{v} = 2\sqrt{pf}$. Increasing f slightly has the marginal benefit of all consumers above $\underline{v}$ paying a higher bidding fee.

The marginal cost of increasing f slightly is the reduction in the number of consumers who enter. Specifically, the threshold consumer with valuation $\underline{v}$ does not enter anymore, which results in a

loss of both the marginal bidding fee the consumer would pay and the loss of the marginal profit

from that consumer not bidding. At the optimal fee level, the marginal benefit equals the

marginal cost, and the following characterization results:

Proposition 2: *When* $\dfrac{3x}{4} - \dfrac{1-H(x)}{h(x)}$ *is an increasing function of x, the optimal consumer*

entry threshold of a non-commitment seller satisfies either $\dfrac{3}{4}\underline{v} = \dfrac{1-H(\underline{v})}{h(\underline{v})}$ *or* $\underline{v} = p$,

whichever is lower. When H(x) = x on [0,p], *the implied optimal fee is*

$$\sqrt{f*(p)} = \min\left(\frac{2}{\sqrt{p}\cdot(p+6)}, \frac{\sqrt{p}}{2}\right) = \begin{cases} \dfrac{\sqrt{p}}{2} & \text{for } p < \sqrt{13}-3 \\ \dfrac{2}{\sqrt{p}\cdot(p+6)} & \text{for } p \geq \sqrt{13}-3 \end{cases}.$$

Note that the special-case assumption *H(x) = x* on [0,p] in Proposition 2 places no

restriction on the distribution of high consumers. The assumption is that there is a mass of *p* low

customers with *v<p* distributed uniformly on [0,p] and a mass of (1-p) high customers with *v > p*

distributed arbitrarily above *p*. There is no need for a specific assumption about the high

consumers' valuations because they all mimic the *v = p* type.

The screening nature of the profit-maximization allows the seller to focus only on the

marginal consumer *v*. The first condition of Proposition 2 is a single-crossing condition

analogous to a monotone hazard-rate condition on *H*. As argued in the previous section, the

threshold must not exceed *p* for the seller to make any money at all, hence the dichotomous

solution: for low *p*, the competition from the outside market is so strong and the relative

proportion of high consumers so high that the seller chooses to serve only the high consumers

and sets *v = p*. Figure 3 outlines the possible payoffs to consumers.

Since the problem of setting the optimal fee is a screening problem, we can easily compare it with other problems of a similar structure, namely, $g(\underline{v}) = InvHazardRate(\underline{v})$. When $g(x) = x$, we get the equation for an optimal reserve of an auctioneer, who can obtain the product for zero marginal cost (Myerson 1981). When, on the other hand, $g(x) = x/2$, we get the equation for an optimal fee of a zero-marginal-cost referral service that merely matches consumers up with producers without collecting any margin. We can therefore easily show that $f*(p)$ is always less than the optimal referral fee but higher than the optimal reserve in a zero-marginal-cost auction.

2.4 Optimal selling with commitment to a minimum markup ($m \geq 0$)

In this section, we show that at least in the uniform case of $H(x) = x/1$ *on* $[0,p]$, the optimal seller strategy constrained (by the lack of commitment) to *m=0* and $f = f*(p)$ of Proposition 2 is in fact optimal even for the unconstrained seller, who can credibly commit to any minimum markup $m \geq 0$. In other words, the seller would rather derive revenue from the bidding fee than from the additional margin arising from the minimum markup.

Suppose the seller can credibly establish commitment to any minimum markup $m \geq 0$, for example, through reputation or transparent billing. The consumers take *m* into account as we show in section 2.2, and the upside for the seller is that all consumers who enter raise their bids to compensate for the higher bid-acceptance threshold. The downside is that fewer consumers now enter at all, as can be seen from the form of the entry threshold $\underline{v} = m + 2\sqrt{pf}$. The resulting profit function of the seller is quite involved, but we can show that for any *m>0,* a corresponding optimal fee $f*(p|m)$ exists, and reducing the *m* slightly in favor of a small increase in *f* from $f*(p|m)$ can help the seller. In other words, the profit of a seller who uses the strategy

$\left(f*(p\,|\,m),m\right)$ is decreasing in *m* for all non-negative *m.* The details of this result form the

proof of Proposition 3—the central result of this paper:

> **Proposition 3**: *Suppose H(x) = x on* [0,*p*]. *Even if the seller can credibly commit to a*
>
> *positive minimum markup*, *the optimal selling strategy uses zero minimum markup and*
>
> *the positive fee* $f*(p)$ *derived in Proposition* 2.

As both the fee and the minimum markup screen out some low consumers, they might appear

equivalent at first glance. However, the minimum markup also changes the entrants' bidding

strategy. Because consumers shade their bids to halfway between their valuation and *m*

(Proposition 1), the minimum markup increases the bids of the entrants less than it increases the

acceptance threshold *w + m*. Therefore, increasing the minimum markup reduces the amount of

trading the entrants achieve on average, thus diminishing market efficiency (i.e., the size of the

pie). This reduction in efficiency is large enough to hurt seller profit as well. Consequently,

generating revenue through a bidding fee, which does not distort bidding, is always better than

generating revenue through a minimum markup, which does. Note that we analyze the extreme

case of perfect and effortless commitment to a margin, and find that it is not useful. Therefore,

intermediate levels of commitment arising from reputation in an alternative dynamic model

would also not be useful to the seller.

As discussed above, the proof of Proposition 3 relies on showing that seller profit is

decreasing in *m* along the locus of $\left(f*(p\,|\,m),m\right)$ strategies. Since the profit is decreasing even

at *m*=0, a seller who could credibly subsidize the consumer bids would do so. We emphasize this

interesting result in the following corollary:

Corollary to Proposition 3: *Suppose H(x) = x on* [0,*p*]. *A seller who charges a bidding fee would find it profitable to credibly subsidize bids and accept at least some bids below cost.*

Illustrative Example. For concreteness, suppose the seller is Priceline and the product is a plane ticket from New York to London, for which p = $1000—the lowest outside posted price competitors such as Expedia and Travelocity offer. Assume the highest valuation that corresponds to v = 1 is $1500, indicating that two thirds of the prospective Priceline customers are "low consumers" in that they value the ticket below p and one third are "high consumers" who value it above p. Table 1 displays the results for (1) the optimal bidding fee given the minimum markup set to zero, (2) the case of an optimal minimum markup given no bidding fee, and (3) for the case of the optimal fee to use given the minimum markup suggested in (2).

Table 1: Illustrative example of a plane ticket from New York to London that costs $1000 on Expedia and that consumers value uniformly between $0 an $1500.

	optimal bidding fee, no minimum markup	*optimal minimum markup, no bidding fee*	*optimal minimum markup and its associated optimal bidding fee*
Bidding fee (f)	$202	$0	$102
Minimum markup (m)	$0	$260	$260
Screening level (v)	$900	$260	$900
Social welfare (gains from trade)	$126	$125	$99
Seller profit	$109	$60	$84

Please see the appendix for details of the calculations that produced Table 1. The results in Table 1 show that the social welfare with a minimum markup and no fee is almost as high as with the optimal fee and no minimum markup, but the seller only makes about half of this welfare in profit, much less than with the optimal fee. Moreover, charging f=0 is not optimal for the seller who uses the minimum markup of $260. Instead, using a positive fee of $102 combined with this

minimum markup of $260 is optimal, which results in lower social welfare, because more consumers are screened out, but higher profit (compare third and fourth column of Table 1). However, profit is still considerably lower than in the case of the optimal bidding fee with no minimum markup (compare the second and fourth columns of Table 1). This comparison illustrates the key point of Proposition 3: the lack of market efficiency when a profit-maximizing seller uses a positive minimum markup as well as a bidding fee hurts profitability.

The uniform assumption does not drive the key qualitative result that the seller prefers fees over minimum markups. We have proven an analogy of Proposition 3 for the exponential distribution (not reported), and we propose that the economic intuition of increased market efficiency is quite general. This intuition is analogous to the intuition behind two-part tariffs outperforming uniform pricing in posted-price markets.

2.5 Should the seller who charges a bidding fee facilitate or hinder consumer learning about the current bid-acceptance threshold?

As discussed in the introduction, consumers may learn the bid-acceptance threshold through communication with recent buyers (Hinz and Spann 2008). Consumers who know the acceptance threshold never bid more than it, erasing the seller's information rent completely. Therefore, a seller who derives profit only from the information rent ($m=f=0$) obviously has a strong incentive to hide his marginal cost w from consumers.[3] This incentive to hide threshold information does not immediately extend to a seller who implements a bidding fee following Proposition 3: facilitating consumer learning may increase participation by high-value bidders, in turn allowing the seller to charge a higher bidding fee and make a higher profit overall. In this subsection, we

---

[3] We believe this to be the basis of Priceline's no-rebidding policy, as well as the rationale for its randomization of the acceptance rule reported in Anderson (2009) and Segan (2005).

sharpen this intuition and qualify its validity with a bound on the intensity of outside market competition captured by p in our model.

To analyze the seller's incentive to hinder consumer learning under bidding fees ($f > 0$), we consider the limiting case of consumers knowing the upcoming bid-acceptance threshold with certainty prior to making their entry decision. That is, we analyze consumers who learn w at the beginning of Stage 2 (see Figure 2). We hope this simple limiting case qualitatively captures what would happen with more realistic learning that would still leave the consumers somewhat uncertain. When consumers know the current acceptance threshold w, they enter and bid $\min(w, p)$ whenever $\min(v, p) - w > f$. The seller accounts for this bidding strategy and charges a different bidding fee than one that would be optimal with uncertain consumers. We denote the bidding fee optimal under the informed-consumer scenario f_I.[4] It is relatively simple to show (please see Appendix for details) that f_I satisfies $2f_I = E\left[\min(p, v) \mid v > f_I\right]$, that is, the seller should charge half of the expected net valuation (minimum of v and p) of consumers who have a positive probability of trading (valuations above f_I). In the uniform low-consumer case of $H(x) = x$ on $[0, p]$, the optimal fee is $f_I^*(p) = \left(2 - \sqrt{4 - 6p + 3p^2}\right)\big/3$, which can be bounded by $p/4 < f_I^*(p) < p/2$. In the plane-ticket example of Table 1, $f_I^* = 2\left(3 - \sqrt{3}\right)\big/9 \approx \$422$. The plane-ticket example suggests the bidding fee is much higher when the potential consumers learn the marginal cost w in advance. This inequality extends to all p as long as the low consumers are uniformly distributed: when $H(x) = x$ on $[0, p]$, $f_I^*(p) > f^*(p)$ for every p because $f_I^*(p) > p/4 > f^*(p)$. The higher bidding fee makes sense because informed consumers do not

---

[4] While we focus on bidding fees ($f > 0$), note that the argument for facilitating consumer learning is equally valid for a seller with commitment to a minimum margin ($m > 0$): with consumers informed about w, bidding fees and minimum margins are interchangeable in that any combination that results in the same $f + m$ yields the same consumer behavior and profit.

have to gamble at entry time and shade their bids afterward. They are thus willing to pay more

for entry. Our final proposition characterizes when facilitating consumer learning about marginal

costs is more profitable for the seller than hindering it:

> **Proposition 4**: *Suppose H(x) = x on* [0,*p*]. *A unique outside price* $\sqrt{13} - 3 < \tilde{p} < 1$ *exists*
> *such that facilitating consumer learning about the seller's current bid-acceptance*
> *threshold is profitable for the seller when* $p > \tilde{p}$, *and vice versa.*

Numerical approximation reveals that $\tilde{p} \approx 0.885$. Recall that when $p < \sqrt{13} - 3 \approx 0.60$, the

reverse-pricing seller only serves high consumers (Proposition 2). Therefore, when the outside

price is so low that the reverse-pricing seller serves only high consumers, Proposition 4 shows

that hindering consumer learning is always optimal.

Figure 4: Seller profit under two consumer-information scenarios
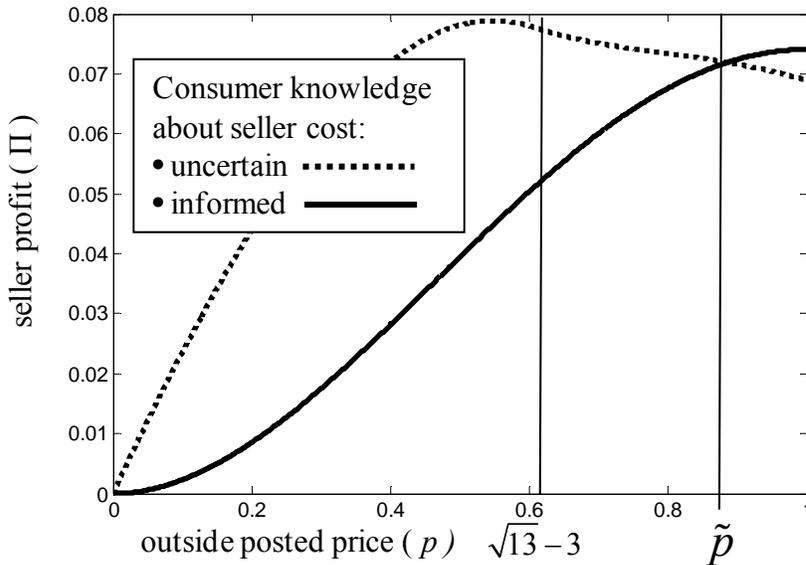
Note to Figure 4: This figure illustrates Proposition 4. The two curves
represent seller profits as a function of the price on the outside posted-price
market. The solid curve involves consumers informed about seller cost
before making their entry decision. The dashed curve involves consumers
uncertain about seller cost at the time of their entry decision.

Not surprisingly, facilitating consumer learning increases the efficiency of the market for all p (proof omitted for parsimony). The intuition for the profitability of facilitating consumer learning only with high p is as follows: a high p means both relatively fewer high consumers and a larger efficiency gain from the added information. The seller can capture enough of the efficiency gain to compensate for the lost information rent provided by the uninformed consumers. As p drops and the proportion of high consumers increases, the efficiency advantage gradually disappears, but the information rent from uninformed consumers remains. Figure 4 illustrates Proposition 4.

3 Discussion

This paper analyzes the design of an optimal reverse-pricing mechanism from the perspective of a profit-maximizing seller. Specification of such a mechanism involves two key decisions: (1) the revenue model, namely, whether to set a minimum markup above cost (and thus set the bid-acceptance threshold) and/or charge a fee for the right to bid, and (2) whether to facilitate or hinder consumer learning about the bid-acceptance threshold. Consumers may learn about the acceptance threshold through repeated bidding or communication with other consumers.

Two key implications for reverse-pricing sellers result from our analyses. First, the optimal revenue model is to charge a bidding fee and then accept all bids above the marginal cost instead of charging a minimum markup above the marginal cost. A bidding fee is preferable because it allows more trades to occur and because it does not require commitment to reject potentially profitable trades. In contrast, a minimum markup increases the bid-acceptance threshold (to cost + minimum markup) without equally increasing bids, thus reducing the amount of trading achieved (i.e., reducing efficiency).

The second major implication of our analyses is that a seller relying on an upfront bidding fee can benefit from consumer learning (e.g., via word of mouth) about the bid-acceptance threshold. This result is in sharp contrast with a seller who relies only on information rents as source of revenue—such a seller always wants to hinder consumer learning. A seller relying on an upfront bidding fee can benefit from such learning due to an increase in market efficiency. The seller can charge a higher bidding fee because informed consumers do not have to gamble at entry time and shade their bids afterward. Consequently, consumers are willing to pay more for entry. Thus, sellers may facilitate consumers' learning of these thresholds, for example, by supporting online communities such as www.biddingfortravel.com or www.betterbidding.com. However, the result of the seller benefiting from such an early revelation of the threshold depends on the distribution of valuations across consumers and the price in the outside market. As the outside price drops and the proportion of high consumers thus increases, the seller may be better off hindering consumer learning about the bid-acceptance threshold, because the information rent from these consumers dominates the efficiency gains.

The results of our analyses also provide interesting insights into the phenomenon of ("live") shopping communities (for example, www.woot.com), in which retailers may charge consumers a fee for the opportunity to enter and buy short-term promotional offers. Furthermore, retailers may even charge consumers a fee for the discovery of the selling price. For instance, on the online shopping site Dubli.com, consumers have to pay .80€ to learn the current price of a product (and simultaneously, they lower the price by .25€). Swoopo.com combines an ascending open auction with bidding fees where consumers have to pay a fee in order to submit a bid.

Bidding fees can be viewed as analogous to subscriptions, which are quite common in discount retailing (BJ's, Sam's Club, etc.). When interpreting the bidding fee as a subscription,

the key simplifying assumption is that the consumer's surplus from bidding at Stage 2 is an accurate characterization of the consumer surplus that arises from her aggregate purchases over the lifetime of the subscription, with these purchases reflecting successful reverse-pricing bids rather than posted-price transactions. By showing that bidding fees are optimal in such a situation, we have provided a model of a subscription-based reverse-pricing discounter.

We derive our results under the simplifying assumption of no frictional costs for ease of exposition. This assumption is not a limitation because, as we can show (not reported), our basic qualitative results would not change if we relaxed it. In contrast, the assumption that the outside market is not a strategic player is necessary for our results, and what would happen if more active competition were allowed is unclear. The analysis of the dynamics of competition between reverse-pricing and posted-price sellers (see Fay 2009) is an important area for future research.

Our analyses are based on the (common) assumption of rational, utility-maximizing consumers. Follow-up work is needed to empirically validate our analytical results. Psychological factors might cause actual bidding behavior in a reverse-pricing context to depart from the normative model we use to characterize it. For instance, examining consumer reactions to bidding fees would be an interesting line of research. We suspect consumers might (irrationally) be more willing to pay a minimum markup than to pay a bidding fee for entry. For example, consumers may use an overriding "rule" to never pay for anything other than the product itself (e.g., information or service), a situation akin to the one Amir and Ariely (2007) describe. A deep understanding of the behavioral aspects of bidding in reverse-pricing markets will require much empirical work, but the analytical findings we report here provide a normative benchmark against which to develop and test behavioral departures and extensions.

The analyses we have presented in this article have important implications for sellers who wish to apply a reverse-pricing mechanism. Our findings provide guidance as to how such firms should determine the optimal reverse-pricing mechanism in terms of how to set a bidding fee or minimum markup, whether to facilitate or hinder consumer learning about the bid-acceptance threshold, and how best to adjust the revenue model when such consumer learning is anticipated.

Appendix: Proofs of Propositions

Proof of Proposition 1: It is sufficient to characterize bidding for valuations that exceed the minimum markup m, because bidders with valuations below m cannot obtain a positive surplus in the market. A bidder with valuation v who has decided to enter and bid an amount x has the expected surplus

$$S(v \mid f, m, p) = \Pr(x \text{ accepted}|m)(v-x) + \mathbf{1}(v > p)\Pr(x \text{ denied}|m)(v-p) - f =$$

$$= \frac{x-m}{p}(v-x) + \mathbf{1}(v > p)\frac{p+m-x}{p}(v-p) - f \text{ . Maximizing this expected surplus yields the bidding}$$

strategy $b(v \mid m) = \begin{cases} \dfrac{m+v}{2} & \text{for } m \le v < p \\ \dfrac{m+p}{2} & \text{for } v \ge p \end{cases}$ *QED.*

Derivation of consumer surplus: Plugging the optimal bidding strategy $b(v \mid m)$ into the surplus

function $S(v \mid f, m, p)$ yields $S(v, f, m, p) = \begin{cases} \dfrac{(v-m)^2}{4p} - f & \text{for } m < v \le p \\ v - p + \dfrac{(m-p)^2}{4p} - f & \text{for } v > p \end{cases}$.

Proof of Proposition 2: The seller maximizes his expected profit $\Pi(f)$ by selecting the optimal bidding fee f. As long as $2\sqrt{pf} \leq p$, that is, as long as at least consumer $v = p$ enters, the seller obtains the following total expected profit:

$$\Pi(f) = \Pr(\underline{v} < v)f + E_{w,v>\underline{v}}\left[\mathbf{1}(b>w)(b-w)\right] =$$

(A1) $$= \left[1 - H(\underline{v})\right]f + \int_{\underline{v}}^{p}\int_{0}^{p}\mathbf{1}\left(\frac{v}{2}>w\right)\left(\frac{v}{2}-w\right)dwdH(v) + \left[1-H(p)\right]\int_{0}^{\frac{p}{2}}\left(\frac{p}{2}-w\right)dw =,$$

$$= \left[1 - H(2\sqrt{pf})\right]f + \int_{2\sqrt{pf}}^{p}\pi(v)dH(v) + \left[1-H(p)\right]\pi(p)$$

where $\pi(v) \equiv \int_{0}^{\frac{v}{2}}\left(\frac{v}{2}-w\right)dw = \frac{v^2}{8}$ is the profit from a bidder of valuation v. The profit $\Pi(f)$ consists of three intuitive parts. The first term is the revenue from bidding fees, the second is the expected revenue from low entrants, and the third is the expected revenue from high entrants.

The second term of the profit function in equation (A1) requires further dissection and explanation. It considers all possible valuations v of consumers who will enter, and for each such v, it computes the seller's ex-ante expected profit from a particular entrant with that valuation. This profit, denoted $\pi(v)$, arises from the seller integrating over all marginal costs that lead to accepted bids $(w$ s.t. $w < b(v))$ while collecting the contribution $b(v) - w$ for each such w. The optimal bidding fee satisfies the following first-order condition:

(A2) $FOC_f : \left(\dfrac{(p+2)\sqrt{pf}}{2}\right)h\left(2\sqrt{pf}\right) = 1 - H\left(2\sqrt{pf}\right).$

The *FOC* captures the following intuition: increasing f slightly has the marginal benefit of all consumers above $\underline{v}$ paying a higher bidding fee (RHS of equation (A2)). The marginal cost of increasing f slightly (LHS of equation (A2)) is the threshold consumer with valuation $\underline{v}$ not entering anymore, which results in a loss of both the marginal bidding fee consumer would pay $f \cdot (d\underline{v}/df) = \sqrt{pf}$ and the loss of the marginal profit from that consumer not bidding $\pi(f|\underline{v}) \cdot (d\underline{v}/df) = \left(p\sqrt{pf}\right)/2$. Both of these components of the marginal cost are weighted by the density $h(\underline{v})$ of the marginal consumer occurring in the first place.

Rewriting the *FOC* in terms of the entry threshold $\underline{v} = 2\sqrt{pf}$ completes the proof of the general case. Plugging in $H(x) = x$ on $[0,p]$ yields $\sqrt{f^*(p)} = 2/(\sqrt{p} \cdot (p+6))$. The implied threshold is less than p when $4 \le p \cdot (p+6)$, above which $\underline{v} = p$ is optimal. *QED*

Proof of Proposition 3: At Stage 1, the reverse-pricing seller maximizes his expected profit $\Pi(m,f)$ by selecting the optimal combination of bidding fee f and minimum markup m. As long as $m + 2\sqrt{pf} \le p$ (i.e., as long as consumer $v=p$ enters), the seller obtains the following total expected profit:

$$\Pi(m,f) = \Pr(v > \underline{v})f + E_{w,v>\underline{v}}\left[\mathbf{1}\big(b(v|m) > w+m\big)b(v|m)(b-w)\right] =$$

(A3)

$$= \left[1 - H\big(m + 2\sqrt{pf}\big)\right]f + \int_{m+2\sqrt{pf}}^{p} \pi(m|v)\,dH(v) + \left[1 - H(p)\right]\pi(m|p)$$

where $\pi(m|v) \equiv \int_{0}^{\frac{v-m}{2}} \left(\frac{v+m}{2} - w\right)dw = \left(\frac{v-m}{2}\right)\left(\frac{v+3m}{4}\right)$. To maximize $\Pi(m,f)$, the seller first finds the optimal bidding fee for every possible level of minimum markup m. This contingent optimal fee satisfies the following first-order condition that equates the marginal cost of raising f slightly with its marginal benefit:

(A4) $FOC_f : \left(mp + \frac{(p+2)\sqrt{pf}}{2}\right)h\big(m + 2\sqrt{pf}\big) = 1 - H\big(m + 2\sqrt{pf}\big).$

Let $H(x) = x$ on $[0,p]$. In other words, assume p low consumers are distributed uniformly on $[0,p]$ and an additional mass $(1-p)$ of arbitrarily distributed high consumers are above p. For this uniform distribution of low consumers, the optimal fee contingent on m implied by FOC_f becomes

(A5) $\sqrt{f^*(p|m)} = \dfrac{2 - 2m(1+p)}{\sqrt{p} \cdot (p+6)},$

which obviously holds for *m* low-enough, specifically $m \leq 1/(1+p)$. When *m* is larger, a positive

fee is not optimal because *m* alone already screens many consumers sub-optimally. The FOC_f

implies an entry threshold of $\underline{v} = \dfrac{4 + 2m - 3mp}{6 + p}$. Since the entry threshold is linear in *m* and the

optimal fee is quadratic in *m*, the profit function $\Pi(m) \equiv \Pi(m, f*(p\,|\,m))$ restricted to the locus

of points that satisfy the FOC_f is only cubic in *m* :

$$\Pi(m) = (1 - \underline{v}) \cdot f + \frac{1}{8} \int_{\underline{v}}^{p} (v - m)(v + 3m)\,dv + \frac{1 - p}{8}(p - m)(p + 3m) =$$

(A6)
$$= -\frac{1}{24(6 + p)^2 p} \cdot \left[\begin{array}{l} -32 + p^3(6 + p)^2(2p - 3) + \underbrace{(2 - 3p)^2(8 + 3p)}_{>0} \cdot m^3 + \\[2mm] +3(3p - 2)(16 + 2p + p^2) \cdot m^2 + 3\left(32(1 + p) + p^2(6 + p)^2(p - 2)\right) \cdot m \end{array} \right]$$

To prove Proposition 3, it is sufficient to show that for all $0 \leq m \leq 1/(1+p)$, $\Pi(m)$ is decreasing

in *m*, that is, $\Pi'(m) < 0$. $\Pi'(m)$ can be evaluated as

$$\Pi'(m) = -\frac{\underbrace{(2 - 3p)^2(8 + 3p)}_{>0} \cdot m^2 + 2(3p - 2)(16 + 2p + p^2) \cdot m + \left(32(1 + p) + (6 + p)^2(p^3 - 2p^2)\right)}{8(6 + p)^2 p}.$$

We start by showing that $\Pi'(m) < 0$ for the two extreme values of *m*: At the upper limit of *m*,

$$\Pi'\left(\frac{1}{1 + p}\right) = -\frac{1 + 4p - 3p^2 + p^4}{8(1 + p)^2} < 0 \Leftrightarrow 1 + p(4 - 3p) + p^4 > 0,$$ which obviously holds for all

$0 \leq p \leq 1$. Second, the profit is also decreasing at the lower limit of *m*:

(A7) $\Pi'(0) = -\dfrac{32(1 + p) + (6 + p)^2 p^2(p - 2)}{8(6 + p)^2 p} < 0 \Leftrightarrow 32 + 32p - 72p^2 + 12p^3 + 10p^4 + p^5 > 0.$

Inequality (A7) clearly holds for *p*=0 and *p*=1. To see that it also holds for all $0 < p < 1$, let

$\lambda(p) = 32 + 32p - 72p^2 + 12p^3 + 10p^4 + p^5$. Evaluating $\lambda'(0) = 32 > 0$ and $\lambda'(1) = -31 < 0$ is

simple. Since λ' is a quadratic with a positive coefficient on *p*, it approaches positive infinity as *p*

grows very large or very small. Therefore, λ' has at least one root above 1 because it is negative

at $p=1$ and continuous. Now suppose a $0 < p* < 1$ exists such that $\lambda(p*) < 0$. Then λ' must actually have all three remaining roots between 0 and 1. In other words, λ' must start out positive around $p=0$, turn negative (λ must decrease below zero toward $p*$), turn positive again (to let λ increase again toward its positive value at 1), and then turn negative near $p=1$. λ' starting out positive, ending negative, and having three roots between 0 and 1 in turn implies that λ'' has exactly two roots between 0 and 1. Since $\lambda''(0) = -144 < 0$ and $\lambda''(1) = 68 > 0$, λ'' having exactly two roots between 0 and 1 is impossible: a cubic that is negative at 0 and positive at 1 must have either one or three roots between 0 and 1. Therefore, the existence of a $p*$ such that $0 < p* < 1$ and $\lambda(p*) < 0$ would lead to a contradiction, and we have shown that $\Pi'(0)<0$ for every $0 \le p \le 1$.

It remains to be shown that $\Pi'(m) < 0$ for all non-extreme m values as well, that is, for $0 < m < 1/(1+p)$. Since $\Pi(m)$ is cubic in m, $\Pi'(m)$ is at most quadratic in m. When $p = \frac{2}{3}$, $\Pi'(m)$ is actually constant at $\Pi'(m) = -\frac{41}{360} < 0$. When $p \ne \frac{2}{3}$, the coefficient on m^2 in $\Pi'(m)$ is unequivocally negative, so $\Pi'(m)$ is concave and has a global maximum at $\Pi''(m*) = 0$.

Simple algebra shows this maximum is $m* = \dfrac{16 + 2p + p^2}{16 - 18p - 9p^2}$. When $p < \frac{2}{3}$, $m* > \dfrac{1}{1+p}$, so $\Pi'(m)$ is monotonically increasing on the $\left[0, 1/(1+p)\right]$ interval and $\Pi'(m) < \Pi'\left(\dfrac{1}{1+p}\right) < 0$ for $0 < m < 1/(1+p)$. When, on the other hand, $p > \frac{2}{3}$, then $m* < 0$, so $\Pi'(m)$ is monotonically decreasing on the $\left[0, 1/(1+p)\right]$ interval and $\Pi'(m) < \Pi'(0) < 0$ for $0 < m < 1/(1+p)$. Since profit $\Pi(m)$ is thus decreasing in m along the locus of possible solutions, the globally optimal solution is captured by the FOC_f given $m = 0$. *QED*

Proof of Proposition 4.

We start by analyzing the seller profits under the two consumer-information scenarios. The profit of the seller facing informed consumers is $\Pi_I\left(f_I^*(p)\right) = f_I^*\left(p - f_I^*\right)\left(\frac{2 - p - f_I^*}{2}\right)$. This can be

expressed as $\Pi_I\left(f_I^*(p)\right) = \frac{1}{27}(2 - D)\left(D - 2 + 3p(2 - p)\right)$ for $D = \sqrt{4 - 3p(2 - p)} \leq 2$, where

the inequality follows from the fact that D is decreasing in p (and hence $f_I^*(p) = \frac{2 - D}{3}$ is

positive and increasing in p as would be expected). Since $\frac{d\Pi_I\left(f_I^*(p)\right)}{dp} = \frac{(1 - p)(2 - D)}{3} > 0$, the

profit is increasing as p increases, that is, as the outside competition weakens. The algebraic expression is complicated by the square root in the optimal fee f_I^*, but the bounds on f_I^* can be used to bound the profit:

$$\frac{p}{4} < f_I^*(p) < \frac{p}{2} \Rightarrow \Pi_I\left(f_I^*(p)\right) < \frac{p}{2}\left(p - \frac{p}{4}\right)\left(2 - p - \frac{p}{4}\right)\frac{1}{2} = \frac{3p^2}{64}(8 - 5p).$$ Another simple-to-

evaluate profit is at $p=1$, when the seller is a monopolist: the fee becomes 1/3, and the profit $2/27 \sim 0.074$.

Now turn to the seller facing uncertain consumers. When p is low enough ($p < \sqrt{13} - 3 \approx 0.60$), the seller facing uncertain consumers only serves the high consumers by

charging $\sqrt{f} = \sqrt{p}/2$ (with $\underline{v} = p$). His profit is $\Pi(f^* | \underline{v} = p) = (1 - p)\left(\frac{p}{4} + \frac{p^2}{8}\right)$, which

exceeds the upper bound on $\Pi_I\left(f_I^*(p)\right)$ when $16 - 32p + 7p^2 > 0$, which holds for most of the

low p because $16 - 32p + 7p^2 > 0 \Leftrightarrow p < \frac{4}{7} \approx 0.57$. Therefore, $p < \frac{4}{7}$ implies

$\Pi_I\left(f_I^*(p)\right) < \Pi\left(f^*(p)\right)$. Note that $\Pi(f^* | \underline{v} = p)$ reaches its maximum at $p = \frac{\sqrt{7} - 1}{3} \approx 0.54$,

after which point it is decreasing. Since $\Pi\left(f^*\left(\sqrt{13} - 3\right)\right) > \Pi_I\left(f_I^*\left(\sqrt{13} - 3\right)\right)$ (as can be shown

by direct computation), $\Pi(f^* | \underline{v} = p)$ must therefore exceed $\Pi_I\left(f_I^*(p)\right)$ also on the

$\left[\frac{4}{7}, \sqrt{13} - 3\right]$ interval. Therefore, a low p implies $\Pi_I\left(f_I^*(p)\right) < \Pi\left(f^*(p)\right)$. When p is high

($p \ge \sqrt{13} - 3 \approx 0.60$), the seller serves some low consumers as well, and his profit is a rather

unwieldy $\Pi\left(f*(p)\right) = \dfrac{32 + p^3\left(3 - 2p\right)\left(6 + p\right)^2}{24p\left(6 + p\right)^2}$. It can be shown that $\Pi\left(f*(p)\right)$ is decreasing

in p. Intuitively, higher p means both a lower competition from the outside posted-price market (positive effect) and a lower chance of very low cost w (negative effect). The negative effect overwhelms the positive effect, as can be seen from the fact that both $f*(p)$ and $\underline{v}(p)$ are decreasing in p: the seller needs to compensate the consumers for the relatively higher costs that lead to less trading. At $p = 1$, the seller charges a fee of $(2/7)^2$, which corresponds to a participation threshold of 4/7. Therefore, the profit

is $\Pi\left(f*(1)\right) = \left[1 - 4/7\right]\left(\dfrac{2}{7}\right)^2 + \displaystyle\int_{4/7}^{1} \dfrac{v^2}{8}\,dv = \dfrac{27}{392} \approx 0.065$. This profit is less than the corresponding

profit of the seller facing informed consumers. Therefore, the seller facing informed consumers can now capture enough of the additional efficiency that he prefers informed consumers, and

$\Pi_I\left(f_I^*(1)\right) > \Pi\left(f*(1)\right)$. From continuity of the profits and the fact that $\Pi\left(f*(p)\right)$ is

decreasing in p for high p while $\Pi_I\left(f_I^*(p)\right)$ is increasing for all p, a unique point $\sqrt{13} - 3 < z < 1$

exists such that $\Pi_I\left(f_I^*(z)\right) = \Pi\left(f*(z)\right)$, and the seller facing informed consumers makes a

higher profit for $p > z$ and vice versa. Finding z algebraically would involve solving a high-order polynomial equation, but it can be easily numerically estimated as $z \approx 0.885$. *QED*

Welfare and Profit Calculation of Illustrative Example in Section 2.4

Let $v=1$ correspond to 1500, and let $p = 2/3$ (which thus corresponds to 1000), and assume consumer valuations are distributed uniformly on [$0,$1500]. The optimal bidding fee to charge

is $f*\left(\dfrac{2}{3}\right) = \dfrac{27}{200} \approx \$202$, which screens at level $\underline{v}(0) = \dfrac{3}{5} \approx \$900$. Thus, most low consumers

do not enter. Those low consumers who do enter bid $v/2$, resulting in bids between $450 and $500. In addition, all high consumers enter, and all bid $500. The expected social welfare W realized through the reverse-pricing seller is the difference $v - w$ when there is a trade, that is, when valuation exceeds $\underline{v}$ and the bid exceeds w:

$$W\left(f*\left(\frac{2}{3}\right)\right)=\int\limits_{\frac{3}{5}}^{\frac{2}{3}}\int\limits_{0}^{\frac{v}{2}}(v-w)\,dwdv+\int\limits_{\frac{2}{3}}^{1}\int\limits_{0}^{\frac{p}{2}}(v-w)\,dwdv=\frac{757}{9000}\approx\$126$$

with an overall probability of trading of about 13 percent. The seller's profit is

$$\Pi\left(f*\left(\frac{2}{3}\right)\right)=\left(1-\frac{3}{5}\right)\frac{27}{200}+\int\limits_{\frac{3}{5}}^{\frac{2}{3}}\frac{v^{2}}{8}\,dv+\left(1-\frac{2}{3}\right)\frac{p^{2}}{8}=\frac{47263}{648000}\approx\$109\,.$$

Now consider the optimal minimum markup to charge contingent on bidding fees being zero.

To derive the optimal markup, let $f=0$ in $\Pi(m,f)$ in the Proof of Proposition 3:

$$\Pi(m\,|\,f=0)=\int\limits_{m}^{p}\pi(m\,|\,v)\,dv+(1-p)\pi(m\,|\,p)\Rightarrow FOC_{m}:m*=\frac{3-\sqrt{9-10p+5p^{2}}}{5}\,.$$

Substituting $p=2/3$ yields $m*=\dfrac{9-\sqrt{41}}{15}\approx\$260$. This markup screens at $260 and results in total

welfare of $W\left(m*\left(\frac{2}{3}\right)\right)=\int\limits_{m*}^{p}\int\limits_{0}^{\frac{v-m*}{2}}(v-w)\,dwdv+\int\limits_{p}^{1}\int\limits_{0}^{\frac{p-m*}{2}}(v-w)\,dwdv=\dfrac{189+59\sqrt{41}}{6750}\approx\$125$. Plugging

the $m*$ into the profit equation yields $\Pi\left(m*\left(\frac{2}{3}\right)\,|\,f=0\right)=\dfrac{61+41\sqrt{41}}{8000}\approx\$60\,.$

Now consider the optimal bidding fee contingent on setting the minimum markup

to $m*=\dfrac{9-\sqrt{41}}{15}\approx\$260$. For $p=2/3$, equation (A5) suggests a fee of

$$f*\left(m\,|\,p=\frac{2}{3}\right)=\frac{3(3-5m)^{2}}{200}\Rightarrow f*\left(\frac{9-\sqrt{41}}{15}\,|\,p=\frac{2}{3}\right)=\frac{41}{600}\approx\$102\,.$$ Since the combination

$(m*,f*)$ is now on the locus of optimal fees given markups, we can simply plug the $m*$ into

equation (A6) to find that the profit is $\dfrac{122+123\sqrt{41}}{16200}\approx\$84$. As expected, this profit is more than

$\Pi\left(m*\left(\frac{2}{3}\right)\,|\,f=0\right)$. Interestingly, the resulting screening level is the same as with the optimal

fee and no markup, namely, $\frac{3}{5} \approx \$900$. The welfare is thus decidedly lower because of bid-

shading: $W\left(m^*, f^*\left(m^*\right)\right) = \int\limits_{\frac{3}{5}}^{\frac{2}{3}} \int\limits_{0}^{\frac{v-m^*}{2}} (v-w)\,dw\,dv + \int\limits_{\frac{2}{3}}^{1} \int\limits_{0}^{\frac{p-m^*}{2}} (v-w)\,dw\,dv = \frac{18 + 277\sqrt{41}}{27000} \approx \$99.$

Claim: *When potential consumers learn the marginal cost before their entry decision, the optimal bidding fee is a solution to* $2f_I = E\left[\min(p,v)\,|\,v > f_I\right].$

Proof of claim: Consumers enter and bid $\min(w, p)$ when $\min(v, p) - f - w > 0$. The seller collects all of his revenue through the fee f, resulting in the following profit function:

$$\Pi(f) = f \Pr\left(entry\,|\,f\right) = f\left[\int\limits_{f}^{p} \Pr\left(w < v - f\right) dH(v) + \int\limits_{p}^{1} \Pr\left(w < p - f\right) dH(v)\right] =$$

$$= f\left[\int\limits_{f}^{p} (v-f)\,dH(v) + (p-f)\left(1 - H(p)\right)\right]$$

where the first term in $\Pr\left(entry\,|\,f\right)$ averages the probability that $w < v - f$ over the valuations of the low consumers. The second term adds the probability that $w < p - f$ times the probability of a high consumer occurring. The first-order conditions are

$$0 = \Pr\left(entry\,|\,f\right) + f\frac{d \Pr\left(entry\,|\,f\right)}{df} = \int\limits_{f}^{p} (v-f)\,dH(v) + (p-f)\left(1 - H(p)\right) - f\left[1 - H(f)\right]$$

$$\Leftrightarrow 2f = \frac{p\left[1 - H(p)\right] + \int\limits_{f}^{p} v\,dH(v)}{1 - H(f)} = E\left[\min(p,v)\,|\,v > f\right] \qquad\qquad QED$$

In the uniform case of $H(x)=x$ on $[0,p]$, the optimal fee is a solution to the equation

$0 = 3f_I^2 - 4f_I + 2p - p^2$, which has a unique root below p of $f_I^*(p) = \frac{2 - \sqrt{4 - 6p + 3p^2}}{3}$. The

fact that $f_I^*(p) < p/2$ follows from $p/2 > \left(2 - \sqrt{4 - 6p + 3p^2}\right)\big/3 \Leftrightarrow p^2 > 0$. QED

References

Anderson, Chris K. 2009. Setting Prices on Priceline. *Interfaces* **39**(4) 307-315.

Amaldoss, Wilfred, Sanjay Jain. 2008. Joint Bidding in the Name-Your-Own-Price Channel: A Strategic Analysis. *Management Science* **54**(10) 1685-1699.

Amir, On, Dan Ariely. 2007. Decisions by Rules: The Case of Unwillingness to Pay for Beneficial Delays. *Journal of Marketing Research* **44**(1) 142-152.

Chernev, Alexander. 2003. Reverse Pricing and Online Price Elicitation Strategies in Consumer Choice. *Journal of Consumer Psychology* **13**(1/2) 51-62.

Ding, Min, Jehoshua Eliashberg, Joel Huber, Ritesh Saini. 2005. Emotional Bidders—An Analytical and Experimental Examination of Consumers' Behavior in Priceline-like Reverse Auctions. *Management Science* **51**(3) 352-364.

Fay, Scott. 2004. Partial Repeat Bidding in the Name-Your-Own-Price Channel. *Marketing Science* **23**(3) 407-418.

Fay, Scott. 2009. Competitive Reasons for the Name-Your-Own-Price Channel. *Marketing Letters* **20**(3) 277-293.

Hann, Il-Horn, Christian Terwiesch. 2003. Measuring the Frictional Costs of Online Transactions: The Case of a Name-Your-Own-Price Channel. *Management Science* **49**(10) 1563-1579.

Hinz, Oliver, Martin Spann. 2008. The Impact of Information Diffusion on Bidding Behavior in Secret Reserve Price Auctions. *Information Systems Research* **19**(3) 351-368.

Koenigsberg, Oded, Eitan Muller, Naufel J. Vilcassim. 2008. easyJet® pricing strategy: Should low-fare airlines offer last-minute deals? *Quantitative Marketing and Economics* **6**(3) 279-297.

Krishna, Vijay. 2002. *Auction Theory*. London: Academic Press of Elsevier Science.

Myerson, Roger B. 1981. Optimal Auction Design. *Mathematics of Operations Research* **6**(1) 58-73.

Samuelson, William F. 1985. Competitive Bidding with Entry Costs. *Economic Letters* **17**(1-2) 53-57.

Segan, Sascha. 2005. *Priceline.com for Dummies*. Hoboken: Wiley.

Spann, Martin, Bernd Skiera, Björn Schäfers. 2004. Measuring Individual Frictional Costs and Willingness-to-Pay via Name-Your-Own-Price Mechanisms. *Journal of Interactive Marketing* **18**(4) 22-36.

Spann, Martin, Gerard J. Tellis. 2006. Does the Internet Promote Better Consumer Decisions? The Case of Name-Your-Own-Price Auctions. *Journal of Marketing* **70**(1) 65-78.

Terwiesch, Christian, Sergei Savin, Il-Horn Hann. 2005. Online-Haggling at a Name-Your-Own-Price Retailer: Theory and Application. *Management Science* **51**(3) 339-351.

Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge MA: MIT Press.

Wang, Tuo, Esther Gal-Or, Rabikar Chatterjee. 2009. The Name-Your-Own-Price Channel in the Travel Industry: An Analytical Exploration. *Management Science* **55**(6) 968-979.