Quant Mark Econ (2015) 13:135–171
DOI 10.1007/s11129-015-9157-y

Optimal selling strategies when buyers name their own prices

Robert Zeithammer[1]

Received: 14 August 2014 / Accepted: 25 May 2015 / Published online: 28 June 2015

Abstract This paper models a name-your-own-price (NYOP) retailer who allows buyers to initiate their retail interactions by describing a product and submitting a binding bid for it. The buyers have an outside option to buy the same good for a commonly known posted price that also acts as an informative upper bound on the cost the NYOP retailer faces. We conceptualize a selling strategy of such an NYOP retailer to be the probability that a buyer's bid gets accepted. The selling strategy is a function of only the bid level; it does not depend on the particular realization of the retailer's procurement cost. Using mechanism-design techniques, we characterize the optimal selling strategy and the equilibrium bidding function that best responds to it. We show that the optimal strategy implements the first-best ex-post optimal mechanism: for every cost realization, the retailer can make as much profit as he would if he could learn his cost first and use the optimal mechanism contingent on it. The complexity involved in credibly communicating an entire bid-acceptance function to buyers can make the first-best strategy impractical in some real-world markets, so we also analyze several simpler NYOP strategies: setting a minimum bid, charging a participation fee, and accepting all bids above cost. We find that under many scenarios, the minimum-bid strategy dominates the other simpler strategies and achieves a majority of the maximal profit improvement available from the first best strategy. However, NYOP retailers in thin markets can do better by charging participation fees than by setting minimum bids.

Keywords Name-your-own-price selling · Pricing · Mechanism design · Ex-post implementation · Bidding · Two-part tariff · Minimum bid · Reserve price · Auction · Revenue equivalence

✉ Robert Zeithammer
rzeitham@ucla.edu

[1] Anderson School of Management, University of California, Los Angeles, CA, USA

1 Introduction

Name-your-own-price (NYOP) selling allows buyers to initiate their interactions with a retailer by describing a product and submitting a binding bid for it. The retailer accepts the bid whenever it exceeds his secret acceptance threshold, and he keeps the difference between the bid and his procurement cost as profit. NYOP selling was pioneered by Priceline.com in 1998 in the travel industry, and has since been adopted globally by a range of online retailers in different product categories such as white goods (prisminister.dk), restaurant meals (chiching.com), and designer fashion accessories (nyopoly.com). This paper considers several strategies an NYOP retailer can use to increase his profits.

Most existing models of NYOP selling assume the retailer is essentially passive and accepts all bids above his procurement cost.[1] Priceline does not officially disclose the details of its bid-acceptance algorithm, but researchers knowledgeable about Priceline's suppliers report a policy very similar to accepting all bids above cost (Anderson 2009; Anderson and Wilson 2011). Chiching.com keeps a 15 % commission, effectively reducing the consumer bids and delegating the NYOP retailer role to the restaurant. The restaurant may in turn be tempted to accept bids above its opportunity cost. Although passively accepting all bids above cost is easy to explain and guarantees positive profits for the NYOP retailer, it cannot be profit maximizing because it does not exercise any market power. Can an NYOP retailer do anything to increase his profits from the passive benchmark? How should we even conceptualize the space of possible NYOP selling strategies? This paper makes two contributions: First, it defines a large set of such strategies, and uses mechanism-design principles to characterize the profit-maximizing one under general assumptions about the distributions of buyer valuations and retailer cost. Second, it compares the first-best strategy with several second-best simpler strategies: a minimum-bid strategy whereby low bids get rejected, a participation-fee strategy akin to a two-part tariff, and a non-NYOP fixed posted price strategy.

Our model of the market in which the NYOP retailer operates assumes the object the NYOP retailer sells can also be obtained in an outside spot market for some commonly known price. For example, the object sold can be a particular designer handbag offered by Nyopoly.com (the NYOP retailer), and the outside market price is the lower of the price posted in the designer's own store for the same handbag and that posted by third-party retailers (if any). Alternatively, the object sold can be a seat on a particular flight sold by Priceline (the NYOP retailer), and the outside market price is the lowest price for that seat posted by other travel retailers with levels of product opacity similar to that of Priceline[2] (e.g., Hotwire). We make the natural assumption that the outside market price is an informative upper bound on the procurement cost the NYOP retailer faces at that moment.

We define an *NYOP selling strategy* to be the probability the NYOP retailer accepts a bid as a function of the bid level and the commonly known outside spot-market price. For example, the selling strategy implied by accepting all bids over cost (the commonly assumed policy described above) is just the cumulative distribution function of the

[1] Amaldoss and Jain 2008; Fay 2009; Wang et al. 2009, and others. Please see Anderson and Wilson (2011) for a review of analytical approaches to modeling NYOP selling.

[2] Priceline makes its offering opaque (Fay 2008) by hiding the airline name and exact time of departure. Other retailers, e.g., byopoly.com, prisminister.dk, or chiching.com, do not make the products opaque. We abstract away from opacity in this paper because it is an orthogonal issue. Please contact the authors for the optimal strategy when the retailer's offering is opaque but the outside market is transparent.

procurement cost. We propose an NYOP retailer should actively manage his selling strategy and strive to communicate it credibly to the buyers.

Our first contribution is a characterization of the optimal NYOP selling strategy and the buyers' bidding function that best responds to it. The NYOP retailer we consider needs to set a single selling strategy for a range of possible procurement costs, but learns the exact cost realization before making his bid-acceptance decision. We have two mutually related reasons for this modeling choice: First, NYOP retailers sometimes only query their suppliers for current cost quotes after receiving a bid.[3] For example, chiching.com approaches participating suppliers only after the bid is finalized. Second, even if the retailer can costlessly look up his current cost for any product offered and post the optimal selling price,[4] the buyers he faces may prefer to bid instead of purchasing at a posted price. We assume that NYOP selling emerged as a natural response to such buyer activism, which was in turn enabled by the Internet. In other words, we take the existence of NYOP selling as exogenously given. For modeling clarity, we adopt the assumption that the retailer needs to set his selling strategy before learning the exact cost realization, and base it only on the distribution of costs. However, all our results will also hold for a retailer who privately knows his cost before a buyer bid arrives, but cannot feasibly post a fixed selling price contingent on it.

To derive the optimal NYOP strategy, we adapt mechanism-design techniques to accommodate the retailer's ex-ante lack of cost information. We first derive the optimal direct-revelation mechanism and then characterize its NYOP implementation. In the direct-revelation mechanism, the optimal allocation rule *conditional* on a cost realization is not surprising: the retailer should accept buyer valuations above the monopoly price implied by the realized cost. However, NYOP selling is not a direct-revelation mechanism, because successful NYOP buyers pay their bids and hence bid strictly below their valuations. Interestingly, we find that under standard regularity conditions of Myerson (1981), the optimal allocation rule from the direct-revelation mechanism can be implemented even when buyers pay their bids and bid strategically.

The optimal NYOP selling strategy (i.e. the optimal bid-acceptance probability function) involves a minimum bid above which higher bids result in higher probabilities of acceptance. The optimal strategy is unique on an interval of low bids, but consists only of an upper bound on an interval of higher bids. The outside spot market influences both the buyer's bidding strategy (it involves a jump-discontinuity and pooling at a particular bid-level) and the retailer's bid-acceptance probability (it accepts high-enough bids with certainty, regardless of the cost realization), but we show that it does not hinder NYOP implementation of the optimal allocation.

The nature of our optimal solution suggests a theoretically interesting corollary: it implies that NYOP *can be* an ex-post optimal selling strategy. Riley and Zeckhauser (1983) show that a monopolist who knows his procurement cost and is selling one indivisible object to a single risk-neutral buyer should set a posted price and make a take-it-or-leave-it offer. Because optimal NYOP selling results in the same allocation as such cost-contingent monopoly pricing, the NYOP retailer makes as much profit as he

3 Several existing models make an analogous assumption, e.g. Amaldoss and Jain (2008) or Spann et al. (2010).

4 Posted pricing would be the optimal cost-contingent mechanism for such a retailer (Riley and Zeckhauser 1983).

would if he could learn his cost first and use the optimal mechanism contingent on it. The key intuition behind the ex-post optimality of NYOP is that although the retailer does not know his cost at the time of setting his strategy, he learns it before making the bid-acceptance decision, and hence eventually has all the pieces of information needed for implementation of the optimal allocation. In other words, NYOP selling can accommodate buyer activism in the form of a desire to submit bids before the procurement cost is realized, without any loss of retailer profit.

The NYOP retailer needs to credibly communicate the strategy to prospective buyers. Credibility requires commitment to a particular acceptance probability for every possible bid level, but is does *not* require the retailer to credibly communicate his cost realization or commit to any action contingent on a cost realization. We devote a special subsection (5.5) to a discussion of two mechanisms that can facilitate credibility in the real world: reputation and a third-party auditor.

The complexity involved in credibly communicating an entire bid-acceptance function to buyers can make the first-best strategy impractical in some real-world markets. Our second main contribution is a thorough exploration of second-best NYOP selling strategies that are simpler to credibly communicate. Specifically, we analyze two kinds of retailers who passively accept all considered bids above their cost, but can commit not to consider certain bids. First, we analyze a seller who can commit to only consider bids above some minimum level—a strict simplification of the first-best strategy. Communicating the minimum bid is simpler because the retailer only needs to post a single number. Credibility is easier to achieve because the minimum bid is analogous to a public reserve in an auction—a standard feature in real-world markets. Second, we consider a seller who can commit to only consider bids by buyers who paid a participation fee—a two-part-tariff strategy proposed by Spann et al. (2010, 2015). We also compare these two second-best NYOP strategies with the non-NYOP strategy of posting a fixed selling price before learning the cost realization. We are able to compare the profitability of these alternative strategies analytically when the distributions of valuations and costs are both uniform. To achieve the profitability comparison outside the uniform-uniform setting, we rely on numerical approximations of the achievable profits under a variety of assumptions about the shape of both key distributions (increasing, decreasing, concave, convex, U, inverted-U) in a full-factorial design.

Not surprisingly, we document that all "active" strategies strictly outperform the passive strategy of accepting all bids above cost. More surprisingly, we find the minimum-bid strategy achieves much of the maximum theoretical profit (i.e., the profit of first best strategy) in most of our simulation scenarios. Specifically, the minimum-bid strategy does particularly well when the distribution of valuations involves a lot of consumers who can afford the outside option, and when the density of valuations around the optimal minimum bid level is not downward sloping. Regarding the comparison of simpler strategies with each other, the minimum-bid strategy obviously weakly dominates the fixed posted-price benchmark. When the distributions of valuations and costs are both uniform, the minimum-bid strategy also strictly dominates the participation-fee strategy for all levels of the outside market price. However, we also find situations in which this ordering reverses: when the distributions of valuations and costs are such that the gains from trade are small, participation fees outperform minimum bids. We conclude that much of the "heavy lifting" of the first-best strategy is often accomplished by the much simpler minimum-bid strategy, but NYOP retailers in thin markets can do better by charging participation fees.

2 Related literature

Although NYOP selling has generated a lot of academic interest as a novel environment for studying consumer decision-making (e.g., Chernev 2003; Ding et al. 2005; Spann and Tellis 2006; Spann et al. 2012, and others), relatively less is known about the selling strategy NYOP retailers should use to maximize their profit (we review the relevant contributions in detail below). This paper contributes to the small but growing theoretical literature about optimal NYOP selling.

As noted in the Introduction, most existing models of NYOP selling assume the retailer accepts all bids above his procurement cost. All other existing papers restrict attention to specific variants of possible strategies, for example, the extent of randomization in the acceptance rule (Shapiro 2011) or minimum markups and participation fees (Spann et al. 2012, 2015). By contrast, this paper is the first to apply mechanism-design techniques to the problem of NYOP selling, thereby considering a much larger set of possible strategies.

The most closely related paper is by Spann et al. (2010, 2015), who show that an NYOP retailer analogous to the one assumed herein profits more from charging a participation fee than from charging a minimum markup (or from some combination of a fee and a minimum markup). Whereas they assume both distributions that parameterize the model are uniform, we prove our results in full generality. We show that no participation-fee strategy can implement the same allocation as the optimal direct-revelation mechanism, and so employing participation fees is strictly less profitable than employing our optimal mechanism.

We take the existence of NYOP selling as given, but our results are relevant to the literature that tries to rationalize the existence of NYOP. One theoretical justification for NYOP selling is as a second-best solution to the optimal selling problem that accommodates buyer activism enabled by the Internet. We show that under standard regularity assumptions about the distribution of buyer valuations, NYOP selling can accommodate buyer activism without compromising retailer profits: despite having to set his selling strategy (i.e., his schedule of bid-acceptance probabilities) before learning his cost of production, the NYOP retailer can achieve first-best ex-post profits. Therefore, NYOP selling does not actually give any more market power to the buyers. In other words, the first-best optimal strategy allows the retailer to recapture his market power despite buyers actively bidding. We propose that even if our strategy is not always the most practical, it thus serves as an important theoretical benchmark.

In another closely related work, Shapiro (2011) shows buyer risk aversion is a way to rationalize the existence of NYOP selling. Specifically, he shows the NYOP monopoly profit is higher than the posted-price monopoly profit when buyers are risk averse, because such buyers bid more than risk-neutral buyers to avoid the risk of not winning at all. Shapiro's (2011) model makes several of the same assumptions we do: his retailer can commit to a probabilistic bid acceptance, and he sometimes faces a non-strategic premium posted-price retailer akin to the one we assume. In contrast to Shapiro (2011), our buyers are risk neutral and our retailer is ex-ante uncertain about his procurement cost. Extending our model with ex-ante retailer uncertainty to the case of risk-averse buyers is not tractable to us, but Shapiro's logic suggests an optimal NYOP retailer facing risk-averse buyers would strictly outperform his posted-price counterpart instead of merely getting the same profit as he does under risk neutrality.

Several other papers consider NYOP retailers who accept all bids above their cost, and show that even such a passive NYOP strategy may dominate posted pricing in a competitive setting (Fay 2009) or as a price-discrimination tool (Wang et al. 2009; Shapiro and Zillante 2009). Such models rely on buyer heterogeneity in an inherent preference for buying via NYOP, perhaps because of varying frictional costs (Hann and Terwiesch 2003) or the varying impact of the risk arising from opacity of the NYOP offering (Fay 2009). The buyers in this paper instead care only about their surplus, and do not heterogeneously favor posted-price buying over NYOP, or vice versa. However, our results can be used as a building block in a competitive model because we characterize the best response of an NYOP retailer to any regular demand function.

Chernev (2003) and Spann et al. (2012) show that consumer behavior is different when the NYOP retailer allows any price to be named (a true "name your own price" in Chernev's nomenclature) compared to when the NYOP retailer presents a menu of prices (called "select your price" by Chernev). The mechanism proposed here is more in line with "select your price." Moreover, we suggest the retailer should present not only a menu of prices, but also the associated acceptance probabilities. Such an institution facilitates credibility, simplifies bidding, and enables the retailer to better learn consumer preferences by fixing bidder beliefs.

NYOP selling is also related to auction theory in that it corresponds to a first-price sealed-bid auction with a single bidder (the buyer) and a stochastic secret reserve (the retailer's bid-acceptance function). The implementability of the optimal mechanism via NYOP selling can thus be interpreted as an extension of the revenue equivalence between first-price and second-price auctions, itself a well-known result in auction theory (Vickrey 1961; Myerson 1981): the direct-revelation mechanism we find is equivalent to a set of second-price sealed-bid auctions, each with an optimal reserve price (different for different costs) and only one bidder. Optimal NYOP selling can be thought of as a set of first-price sealed-bid auctions, each of which is constructed to accept a bid exactly often enough to maintain revenue equivalence with the corresponding second-price auction.

We analyze a static model in which each buyer is restricted to a single bid, but real-world bidders may manage to bid multiple times. Fay (2004) uses a stylized model to show that repeat bidding does not necessarily erode retailer profits, as long as the retailer is aware of the behavior and uses the right dynamic thresholds. In a more recent paper on repeat bidding, Chen (2012) suggests Priceline's "lockout period restriction, a design alleged to protect sellers, can actually benefit customers," and links the issue to bargaining theory. Extending our setup to a dynamic environment is beyond the scope of this paper, but we propose that something akin to Fay's (2004) finding would replicate: as long as the retailer anticipates repeat bidding and conditions his acceptance-probability strategy on the number of bids the bidder has submitted to date, the retailer should be able to extract first-best profits from the sequence of interactions.

3 Model

Our notation follows Krishna (2002) whenever possible, and it is summarized for easy reference in Table 1. A risk-neutral buyer is interested in buying one particular indivisible object. The buyer's valuation x of the object is drawn from a continuous distribution $F(x)$ with density $f(x)$ and support on $[\underline{x}, \overline{x}]$. Assume the *virtual value* $\psi(x) \equiv x - \frac{1-F(x)}{f(x)}$ is increasing, that is, that the distribution F is *regular* in the sense of Myerson (1981). The virtual value function is a central concept in the theory of mechanism design, and it represents the marginal revenue a seller can extract from a buyer of type x in a direct-revelation mechanism (see Krishna 2002 for a more detailed exposition).

The object is readily available in an outside posted-price market for a commonly known price $\psi^{-1}(0) < p \leq \overline{x}$, where $\psi^{-1}(0)$ is the price a monopolist with zero marginal cost would charge for the object. When the buyer does not buy the object from either retailer, her payoff is zero. Following Spann et al. (2010) and Shapiro (2011), we assume the NYOP retailer is small in that he takes the posted price as fixed, and the outside spot market does not adjust its posted price in response to the NYOP retailer's strategy.[5]

An NYOP retailer can procure the object for a procurement cost $c \sim H(c)$, where the distribution H has full support on $[0,p]$. In other words, the outside posted price is a public upper bound on the NYOP retailer's procurement cost. For example, the NYOP retailer can be Priceline selling excess capacity on a particular flight, whereas p is the price of a seat on the same flight posted by Hotwire (a posted-price site with similar detail of product description, also known as "opacity"). Alternatively, the NYOP retailer Nyopoly could be reaching price-conscious consumers of designer handbags, whereas p is the price of the same handbag at the designer'store.[6] The entire model is thus parameterized by the two distributions F and H, where the support of the latter depends on p. A lower p is both bad news (tougher competition) and good news (lower expected cost) for the NYOP retailer.

Note that in our baseline model, we abstract from the fact that some NYOP products are opaque—a feature pioneered by Priceline, and a potential source of differentiation between the two retailers (Fay 2008; Shapiro and Shi 2008). Opacity is not intertwined with NYOP selling in the real world: none of the other three retailers we mention in the Introduction (chiching.com, prisminister.dk, nyopoly.com) are opaque. It can be shown that opacity does not change our main result qualitatively. Please contact the authors for the optimal strategy when the retailer's offering is opaque but the outside market is transparent.

Timing of the game is as follows (please see the bottom timeline in Fig. 1): in the beginning of the game, the NYOP retailer announces his bid-acceptance strategy $A(b) \equiv \Pr(accept\ b)$ for all possible levels of bid b submitted by the buyer for the object. The buyer then submits a binding bid. After receiving a bid, the retailer queries his suppliers for a cost quote to learn his actual cost c and decides whether to accept the bid. At any time during the game, the buyers can choose to buy from the posted-price outside market and pay the price p. Figure 1 highlights the key contrast with posted pricing in the timing of the cost information.

[5] Because the two competitors are selling the same object, Bertrand competition would result if the outside competitor responded. To prevent a complete collapse of profits, we could introduce horizontal differentiation arising from heterogeneity in buyer inherent preference for NYOP over posted pricing similar to Hann and Terwiesch (2003) or Fay (2009). Within such a larger model, our paper characterizes what the NYOP best response would look like.

[6] The underlying assumption is that the NYOP retailer does not have a special technology for producing the object, but rather obtains the object from the same supplier as his posted-price competitors. Even after learning the posted price, uncertainty about c remains because p is a relatively stable price, set to reflect long-run revenue-management considerations and quite possibly a larger set of customers (as in Spann et al. 2010).

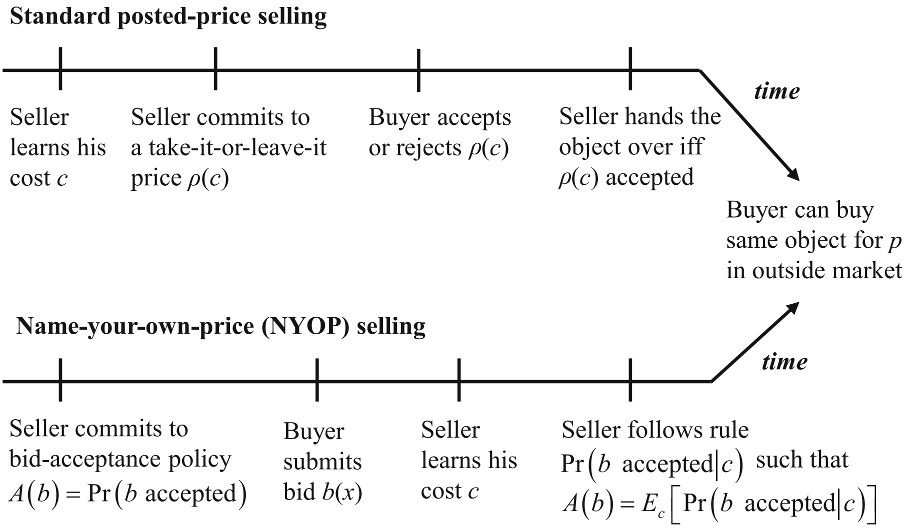

Fig. 1 Timing of the NYOP game, as compared to standard posted pricing

To derive the optimal mechanism, we follow Myerson (1981) and the rest of the mechanism-design literature, and assume the NYOP retailer can commit to any bid-acceptance strategy $A(b) \equiv \Pr(accept\ b)$. In the second half of the paper, we relax this assumption in several ways and explore the impact on retailer profits. When the retailer has no pre-commitment ability, he obviously accepts all bids above c, so everyone knows $A(b) = \Pr(b > c) = H(c)$, and the retailer's announcement in the beginning of the game contains no new information.

Before solving for the optimal bid-acceptance strategy, we summarize the impact of the outside spot market on the demand the NYOP retailer faces. For any bid-acceptance strategy, all buyers with $x > p$ mimic the type $x = p$ because they have a real option to buy in the outside market when the price they name is rejected. In other words, the NYOP retailer faces buyers with a distribution of *net* valuations F on $[0,p)$ and $[1-F(p)]$ mass at p. To see this fact, note the expected surplus U of an $x > p$ buyer who bids b is

$$U(b,x) = A(b)(x-b) + [1-A(b)](x-p) = (x-p) + U(b,p), \tag{1}$$

where $U(b,p)$ is the utility of the buyer with $x = p$. It is immediate that the same b that maximizes $U(b,p)$ also maximizes $U(b,x)$. Intuitively, the buyer thus receives all of his valuation in excess of p as surplus, and his participation with the NYOP retailer is akin to free gambling in hopes of randomly getting a price below p.

4 Optimal direct-revelation mechanism

We use the revelation principle (Myerson 1981) to restrict attention to direct-revelation mechanisms whereby the buyer reports her valuation truthfully. Much of the material in this section is standard, so the details are relegated to the Appendix. The only deviation from a textbook treatment (e.g., Krishna 2002) is the fact that the retailer does not know his cost c when he sets his strategy, but he does learn c before making his bid-acceptance decision.

Knowing the cost before the acceptance decision needs to be made allows us to first optimize the contingent bid-acceptance rule $\pi(x,c)$=Pr(allocate object to x when cost is c). Not knowing the cost at the outset restricts the retailer to probabilistic assignments that are only a function of x. Let $q(x)$ be the probability that a buyer with x receives the object. A retailer who is planning to use a given $\pi(x,c)$ can only commit to $q(x) = \int_0^p \pi(x,c)dH(c)$. The aforementioned atom at p also provides an interesting wrinkle relative to the textbook in that the optimal q is discontinuous at p. The optimal contingent bid-acceptance rule is as follows (please see the Appendix for all proofs):

Proposition 1 *The optimal contingent bid-acceptance rule in a direct-revelation mechanism is*

$$\pi(x,c) = \begin{cases} 1 \ when(x < p \ and \ c < \psi(x)) \ or \ when \ x = p \\ 0 \ otherwise \end{cases}.$$

For every cost c, this rule achieves the same profit as a posted-price monopolist who knows c before setting his price r, namely,* $[1-F(r^*(c))][r^*(c)-c]$ where $r^*(c)$= $\min[\psi^{-1}(c),p]$.

The allocation rule for $x < p$ is familiar from the mechanism-design literature: Myerson's (1981) optimal reserve price applies for low buyers—for every $x < p$, the retailer should sell (i.e., set $\pi(x,c)=1$) iff $\psi(x)>c \Leftrightarrow x>$monopoly price($c$). In addition to sometimes serving some of the low buyers, the retailer should also always sell to all high-value buyers ($x \geq p$) because $p \geq c$ holds for all c by construction.

The intuition for $\pi(x,c)$ goes back to a simple posted-price monopolist with a marginal cost $c < p$ who faces demand $[1-F(z)]$ for all prices $z < p$, and a point mass of $[1-F(p)]$ customers willing to pay exactly p. Such a monopolist charges precisely $\min(\psi^{-1}(c),p)$ to maximize his profit. Because the retailer can condition his $\pi(x,c)$ on c, he can effectively get the ex-post monopoly profits. In other words, for every cost c, he can get the same profit as a posted-price monopolist who knows c before setting his price. Note that the optimal allocation applies regardless of the distribution of retailer cost c, but the implied bid-acceptance strategy will depend on H.

4.1 Example: F and H uniform

Throughout the paper, we will use the example of F = Uniform[0,1] and H = Uniform[0,p] to illustrate the findings in closed form. When F is uniform on [0,1], $\psi(x)$=2x−1, so the complete optimal allocation is $sell \Leftrightarrow x > \frac{1+c}{2}$ or $x = p$. Note that only buyers with $x > \frac{1}{2}$ have any chance of winning, so the retailer effectively sets a minimum bid of $\frac{1}{2}$. The retailer makes a profit of $\Pi_A(p) = E_c\left[\int_{\min\left[\frac{1+c}{2},p\right]}^{1} (2x-1-c)dx\right]$, where the A subscript denotes the first best strategy by referring to the bid-acceptance strategy. When H is also uniform, the retailer's expected profit is $\Pi_A(p) = \int_0^{2p-1}\left(\int_{\frac{1+c}{2}}^{1}\frac{2x-1-c}{p}dx\right)dc + \int_{2p-1}^{p}\left(\int_p^1\frac{2x-1-c}{p}dx\right)dc = \frac{2p^3+6p(1-p)-1}{12p}$

Note that regularity of F in the sense of Myerson (1981) is not required for posted pricing to be the optimal strategy contingent on c: Riley and Zeckhauser (1983) show that the optimal $\pi(x,c)$ is a step function with a single step even when $\psi(x)$ is not increasing in x. Had we not assumed regularity in this section, Proposition 1 would be modified to $\pi(x,c)=1$ when either $x>x^*$ for some x^* that satisfies $c=\psi(x^*)$, or when $x = p$ (see Proposition 1 of Riley and Zeckhauser 1983). Exposition is easier with a regular F, and the next section will show that regularity is actually a necessary condition for an implementation of the optimal allocation rule through NYOP selling.

The form of the optimal allocation rule in Proposition 1 is not surprising. The key question of this paper is how to implement it within the NYOP institution. The retailer could simply promise to charge a price of $\min(\psi^{-1}(c),p)$ to all bidders with bids that exceed it, preserving the incentive to bid truthfully. In other words, the retailer could run a Becker et al. (1964) procedure with a carefully selected distribution of prices. However, an NYOP retailer promises to charge buyers their bids whenever a sale occurs. In response to paying their bids, buyers shade their bids below their private valuations. In the next section, we derive the buyer's bidding function and the retailer's bid-acceptance rule that implements the optimal mechanism.

5 Implementation of the optimal mechanism through NYOP

An NYOP retailer promises to charge buyers their bids whenever a sale occurs. One advantage of accepted buyers paying their bids (vs. the optimal monopoly price conditional on the realized c) is that the retailer does not need to credibly communicate his c to the buyers. However, how buyers will respond is not a priori clear. Will they bid according to an increasing (and hence invertible) bidding function that allows the retailer to implement his optimal allocation, or will they somehow obfuscate their type to avoid being exploited? In this section, we show constructively that the buyers' best-response bidding function is invertible—and the NYOP retailer can thus implement the optimal allocation rule described in Proposition 1—if and only if the virtual value $\psi(x)$ is increasing in x, irrespective of H. In other words, regularity of F in the sense of Myerson (1981) is necessary and sufficient for an NYOP implementation of the optimal mechanism. The proof proceeds in three steps, each captured in a separate lemma:

1) Derivation of the bidding strategy for low-valuation bidders ($x < p$)
2) Derivation of the bidding strategy for high-valuation bidders ($x = p$)
3) Specification of the optimal bid-acceptance probability for intermediate bid levels

5.1 Bidding strategy for low-valuation bidders ($x < p$)

First, consider a buyer with $\psi^{-1}(0)<x<p$ who follows a bidding strategy $\beta(x)$. Proposition 1 shows his probability of winning in the optimal mechanism is $q(x) = \int_0^p \pi(x,c)dH(c) = H(\psi(x))$. The proof of Proposition 1 shows his expected utility is $U(x) = \int_{\underline{x}}^{x} q(t)dt = \int_{\psi^{-1}(0)}^{x} H(\psi(t))dt = \int_{\psi^{-1}(0)}^{x} (x-t)dH(\psi(t))$, where the last

equality follows from integration by parts. For NYOP selling to be revenue equivalent with the optimal direct-revelation mechanism, the bidding strategy $\beta(x)$ must satisfy:

$$U(x) = q(x)(x - \beta(x)) \Rightarrow \beta(x) = x - \frac{U(x)}{q(x)} \tag{2}$$

The structure of the candidate bidding function in Eq. (2) yields the following result:

Lemma 1 *For bidders with $x < p$, the NYOP bidding strategy implied by the optimal direct-revelation mechanism of Proposition 1 is increasing iff ψ is increasing, and it can be characterized by $\beta(x) = E_c[\psi^{-1}(c)|\psi^{-1}(c) < x]$.*

In words, Lemma 1 shows the optimal bid by a buyer with valuation $x < p$ is the average (over c) of monopoly prices the buyer would be willing to pay, namely, prices below x. Given the regularity of F we assume throughout, the bidding function is increasing, and hence invertible, so the retailer can infer each bidder's x and apply the allocation rule of Proposition 1. Interestingly, regularity of F is also required for a bidding function that satisfies Eq. (2) to be increasing. Given the bidding function, we can also solve for the cost-contingent bid-acceptance rule implied by the $\pi(x,c)$ in Proposition 1:

$$c < \psi\left(\beta^{-1}(b)\right) \Leftrightarrow b > \beta\left(\psi^{-1}(c)\right) \Leftrightarrow b > E_{\text{cost}}\left[\psi^{-1}(\text{cost})\middle|\text{cost} < c\right] \tag{3}$$

That is, the retailer accepts bids over the average monopoly price he would charge for all costs below his actual cost realization.

5.2 Bidding strategy for high-valuation bidders ($x = p$)

Now consider a buyer with $x = p$, and recall that the optimal allocation rule is to sell to this buyer for all levels of c. Because $\psi(x) < x$ for all $x < 1$, the limit of $\pi(x,c)$ as x approaches p from below involves an acceptance probability below one for all x other than $x = \bar{x}$(special because $\psi(\bar{x}) = \bar{x}$). In other words, $A\left(\beta_p^-\right) = \Pr(c < \psi(p)) < 1$, where $\beta_p^- = \lim_{x \to p-} \beta(x)$. To implement the needed $\pi(p,c) = 1$, a bid level $\beta(p) > \beta_p^-$ must exist such that buyers with $x = p$ prefer to bid $\beta(p)$, but buyers with $x < p$ do not deviate from bidding according to the above $\beta(x)$. Our next lemma shows that a unique such $\beta(p)$ does exist:

Lemma 2 *For bidders with a net valuation of the NYOP offering equal to the outside price p, the NYOP bidding strategy implied by the optimal direct-revelation mechanism of Proposition 1 is $\beta(p) = p - \int_{\psi^{-1}(0)}^{p} H(\psi(z))dz > \beta_p^-$.*

The intuition behind Lemma 2 is that the surplus of a "high-valuation" buyer with $x = p$ is a sure gain of $p - \beta(p)$. The high-valuation buyer is willing to bid $\beta(p)$ as long as $\beta(p)$ is low-enough that the resulting surplus is at least as much as the expected surplus of the highest low-valuation buyer (with x just under p) and Proposition 1 pins the latter to be exactly the cumulated probability of winning represented by the integral expression in Lemma 2. The highest low-valuation buyer is in turn willing to follow his Lemma 1 strategy as long as the bid $\beta(p)$ is high-enough that deviating up to the sure payout of $p-\beta(p)$ is not profitable for him. Therefore the incentives of the two key buyers ($x=p$ and x just under p) put exactly opposing pressures on the level of $\beta(p)$, resulting in a unique feasible magnitude. Despite having their bids accepted with certainty, the high-valuation buyers make the same "gambling payout" (see Eq. 1) from the NYOP channel as the highest low-valuation buyer whose strictly lower bid is often rejected.

5.3 Optimal bid-acceptance probability for intermediate bid levels

The optimal mechanism determines the acceptance probability of all bids in $[0,\bar{\beta_p}]$ and $[\beta(p),p]$. No bidders should submit bids in the "intermediate" region of $[\bar{\beta_p},\beta(p)]$, but the NYOP retailer still needs to specify a bid-acceptance strategy function for such bids, ensuring they indeed remain off equilibrium. One simple rule is to simply reject such bids. Another simple rule that keeps $A(b)$ non-decreasing is to simply let $\Pr(accept\ \ b)=H(\psi(p))$ for all $b\in[\bar{\beta_p},\beta(p))$. Our next Lemma derives the upper bound on the probability acceptance of every intermediate b such that nobody wants to bid in the $[\bar{\beta_p},\beta(p)]$ interval:

Lemma 3 *When* $A(b)\leq\left(\frac{1}{p-b}\right)\int_{\psi^{-1}(0)}^{p} H(\psi(z))dz$ *for* $b\in[\bar{\beta_p},\beta(p))$, *no buyer submits a bid* $b\in[\bar{\beta_p},\beta(p))$.

Plugging $b=\beta(p)$ into the bound in Lemma 3 shows the upper bound approaches 1 as b approaches $\beta(p)$. Therefore, the full $A(b)$ can actually be continuous and increasing on $[0,\beta(p)]$. Lemmas 1–3 complete the proof of the first main result of this paper:

Proposition 2 *For every regular continuous F on $[\underline{x},\overline{x}]$ and every outside posted price $p>\psi^{-1}(0)$, the following bid-acceptance probability function implements the ex-post optimal mechanism*:

$$A(b)=\begin{cases} b<\lim\limits_{x\to p-}\beta(x):H\left(\psi\left(\beta^{-1}(b)\right)\right) \\ \lim\limits_{x\to p-}\beta(x)\leq b<\beta(p):anything\leq\int_{\psi^{-1}(0)}^{p}\dfrac{H(\psi(z))}{p-b}dz \\ b\geq\beta(p):1 \end{cases}$$

where the bidding function

$$\beta(x)=\begin{cases} x\in\left[\psi^{-1}(0),p\right):\int_{0}^{\psi(x)}\psi^{-1}(c)\frac{dH(c)}{H(\psi(x))} \\ x=p:p-\int_{\psi^{-1}(0)}^{p}H(\psi(z))dz \end{cases}$$

is the unique best response to $A(b)$. After learning his production cost c, the retailer accepts a bid b whenever $b>\beta(\psi^{-1}(c))$ for $b<\lim_{x\to p-}\beta(x)$, and he accepts $\beta(p)$ with certainty.

Note the NYOP implementation of the optimal mechanism consists of a pair of a bid-acceptance probability $A(b)$ and a bidding function $\beta(x)$ that best respond to each other. The minimum bid implied by $\beta(x)$ and $A(b)$ is $\psi^{-1}(0)$. It is useful to consider a closed-form example:

5.4 Example: F and H Uniform

It is easy to show that an $F = Uniform[0,1]$ implies $\psi^{-1}(0) = \frac{1}{2}$, and hence for $\frac{1}{2} < x < p$,

$$\beta(x) = \int_{1/2}^{x} 2z \frac{h(2z-1)}{H(2x-1)} dz = \int_{0}^{2x-1} \left(\frac{c+1}{2}\right) \frac{h(c)}{H(2x-1)} dc = E\left(\frac{c+1}{2} \middle| x > \frac{c+1}{2}\right) \tag{4}$$

The bid of the $x = p$ bidders is

$$\beta(p) = p\left[1 - \Pr\left(p > \frac{c+1}{2}\right)\right] + \Pr\left(p > \frac{c+1}{2}\right) E_c\left(\frac{c+1}{2} \middle| p > \frac{c+1}{2}\right) \tag{5}$$

When H is also uniform on $[0,p]$, $\beta(x) = \frac{1+2x}{4}$ for all $x < p$, and $\beta(p) = 1 - \frac{1}{4p}$. Please see Fig. 2a for this bidding function.

Given this bidding strategy, we can also solve for the bid-acceptance rule in closed form:

$$\pi(b,c) = 1 \Leftrightarrow \begin{cases} b \leq \dfrac{1+2p}{4} : \text{ accept when } c < 2\beta^{-1}(b) \Leftrightarrow b > \dfrac{1}{2} + \dfrac{c}{4} \\[2ex] \dfrac{1+2p}{4} < b < 1 - \dfrac{1}{4p} : \text{accept with } \Pr \leq \dfrac{(2p-1)^2}{4p(p-b)} \\[2ex] b \geq 1 - \dfrac{1}{4p} : \text{accept always} \end{cases} \tag{6}$$

Therefore, the retailer only accepts bids over 1/2, accepts bids over $1 - \frac{1}{4p}$ with certainty, and is willing to sell below cost whenever $\frac{1}{2} + \frac{c}{4} < c \Leftrightarrow c > \frac{2}{3}$. The implied bid-acceptance strategy is

$$A(b) = \begin{cases} 0 \text{ for } b < \dfrac{1}{2} \\ \quad \dfrac{4b-2}{p} \text{ for } \dfrac{1}{2} \leq b \leq \dfrac{1+2p}{4} \\[2ex] \text{below } \dfrac{(2p-1)^2}{4p(p-b)} \text{ for } \dfrac{1+2p}{4} < b < 1 - \dfrac{1}{4p} \\[2ex] 1 \text{ for } b \geq 1 - \dfrac{1}{4p} \end{cases} \tag{7}$$

5.5 Two real-world mechanisms for making the first-best strategy credible

The optimal strategy outlined above requires the retailer to communicate a particular bid-acceptance function and commit to it. One way to facilitate such credible communication would be for the NYOP retailer to post a table with the probabilities of different bid-amounts being accepted, and allow buyers to easily report posted

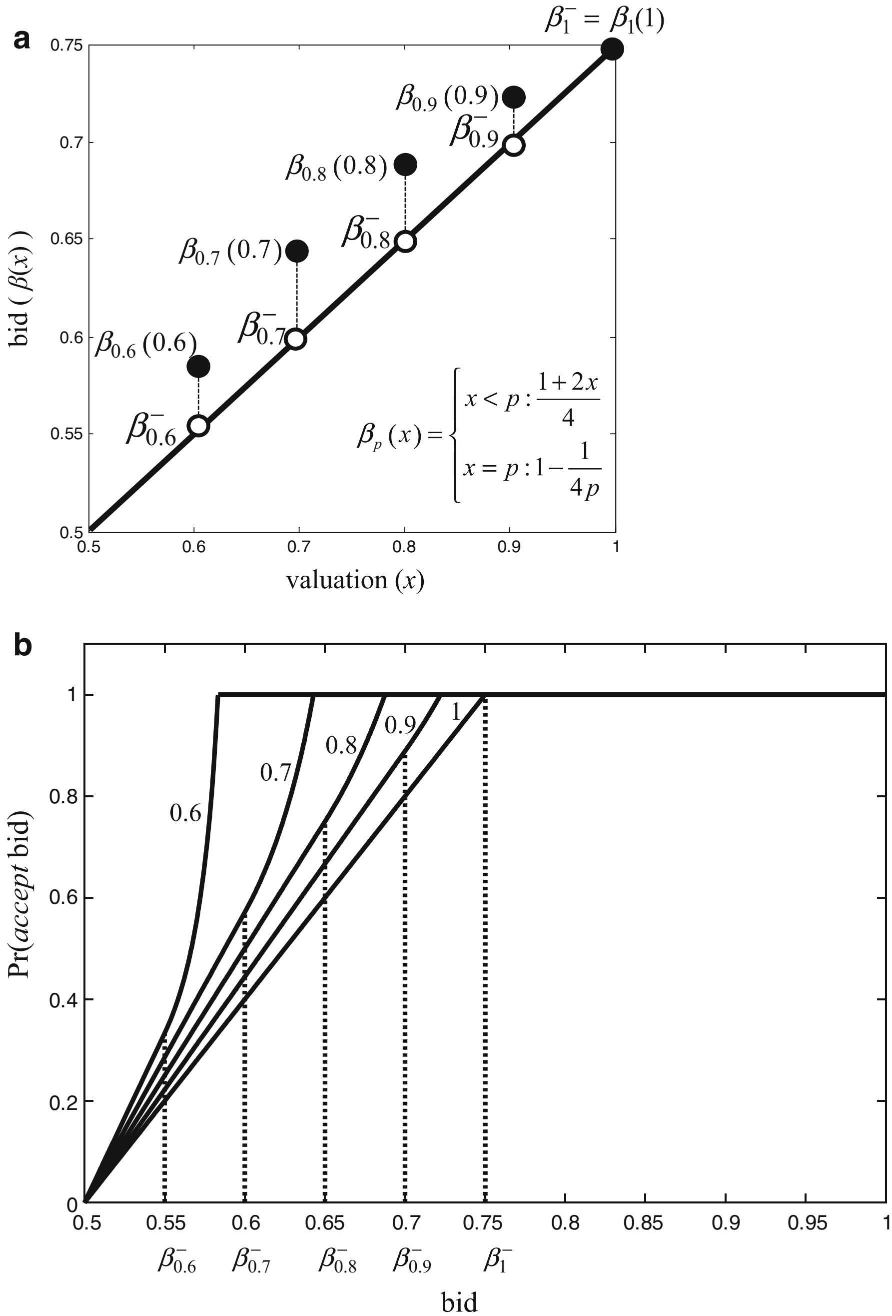

Fig. 2 **a** Bidding functions that implement the optimal mechanism, for five levels of the outside market price p (F=*Uniform*[0,1],H=*Uniform*[0,p]). Note to Figure: The subscript of the bidding function indicates p. The empty circles indicate the jump discontinuities from β_p^- to $\beta_p(p)$. For $q < p$, $\beta_q(z)=\beta_p(z)$ for all $z < q$. **b**: Optimal bid-acceptance strategy, for five different levels of outside price p. Note to Figure: An illustration of Eq. (7). The numbers next to the lines indicate the five levels of p. Each line reaches 1 at the respective $\beta(p)$. For bids between β_p^- and $\beta(p)$, the lines indicate the maximum acceptance probability such that no bidder bids at those levels

probabilities and acceptance outcomes to a third-party auditor. The auditor could then post the summary statistics, ensuring the empirical probabilities match the promised ones. For example, the auditor could simply report a single graph, relating the stated probabilities to recency-weighted empirical estimates of the actual probabilities. As long as that graph is a 45-degree line, the retailer is not under- or over-stating the buyers' chances. Note that the auditor only needs to know the stated probability of success and the outcome (accept or reject), not the retailer's eventual cost realization or anything specific about the product at hand or the bid level. We propose the NYOP retailers interested in using our optimal selling strategy should facilitate the information flow to such a third-party auditor and perhaps even support its operations with a flat subsidy.

Note that an auditor of a multi-product retailer can pool information not only across independent cost draws for one product, but also across all of his different products within the same period. Specifically, the auditor needs to only report a single graph that pools across all recent submitted bids on all products. The retailer can thus be more convincing about his stated probabilities of acceptance when he carries more products, so the auditor technique for commitment exhibits an economy of scope.

Without an official auditor, consumer word of mouth on websites such as biddingfortravel.com can serve as an unofficial auditor and facilitate credibility via reputation. This scenario is already happening in the case of Priceline, at least to some degree: when a buyer selects a hotel region on which to bid, the website flashes a message saying that some person "recently won a hotel in the selected area" for a certain price – presumably a price with a high chance of acceptance. When the buyer then types in a very low bid, Priceline automatically flashes a message saying, "Based on recent data, your price has almost no chance of being accepted." Thus, the website communicates at least the support of prices that have a positive but still uncertain chance of being accepted. User forums such as betterbidding.com and biddingfortravel.com contain many reports of how accurate this information is, making sure Priceline does not exaggerate the range of suggested prices it reports.

6 Simpler second-best selling strategies: four alternatives

The complexity of the first-best NYOP selling strategy characterized in the previous section may be impractical for some real-world markets. In this section of the paper, we consider three simpler NYOP selling strategies and one obvious non-NYOP benchmark: first, retailers may not be able to credibly communicate a full bid-acceptance strategy, but commit to only considering bids above some minimum level akin to a public reserve price in an auction. We contrast this strategy with an important non-NYOP benchmark that is offered by a seller who sets a fixed posted price despite his ex-ante uncertainty about cost. Second, the retailers may be able to charge a participation fee, that is, ask to be paid for considering bids, as proposed by Spann et al. (2010). Third, we analyze a "passive" retailer who can only credibly reject unprofitable bids (bids below cost). The profits available to such a retailer are the relevant baseline from which all the above retailers improve. We thus begin by defining this benchmark next.

6.1 Passive retailer

A passive retailer cannot commit to anything and cannot charge a participation fee, and hence accepts all bids above cost: $A(b)=H(b)$. We assume the bidder's belief about H is correct; either the retailer communicates his A or the bidders can learn it from experience or word of mouth. The optimal bidding strategy is

$$\beta_0(x) = \underset{b \geq 0}{\operatorname{argmax}}\, H(b)(x-b) + [1-H(b)]\max(0, x-p) \tag{8}$$

where the second term represents the option value of buying in the outside posted-price market, and where the zero subscript will henceforth denote the passive-retailer situation. The expected surplus in Eq. 8 has the same form as that in Eq. (1), so $\beta_0(x)=\beta_0(p)$ for all $x>p$. Consider $x\leq p$, and let b^* solve the first-order condition (FOC) $b^* = x-\frac{H(b^*)}{h(b^*)}$. Under standard regularity assumptions about H, $\beta_0(x|x\leq p)=b^*$.[7] Given the bidding function, the expected retailer profit is: $\Pi_0(p)=E_x[\pi_0(\min(x,p))]$, where $\pi_0(x) = \int_0^{\beta_0(x)} (\beta_0(x)-c)dH(c)$ is the expected profit contribution of a buyer with valuation x.

Uniform-uniform example: When F = Uniform[0,1] and H = Uniform[0,p], $\beta_0(x) = \frac{x}{2}$, $\pi_0(x) = \frac{x^2}{8p}$, and $\Pi_0(p) = \frac{p(3-2p)}{24}$. Note that $\Pi_0(p)$ is hill shaped, and maximized at $p = \frac{3}{4}$. This non-monotonicity demonstrates how a lower outside market price is both bad news (tougher competition) and good news (lower expected cost) for the NYOP retailer. The competitive effect is stronger for low p, the cost-reduction effect for high p.

6.2 Minimum-bid and fixed-price strategies

Suppose the NYOP retailer cannot commit to a probability schedule $A(b)$, but can credibly refuse to consider bids below a certain minimum level m. Once he considers a bid, he accepts all bids above his cost just like the passive retailer of the previous section.

Commitment to a minimum bid is easier than credibly communicating an arbitrary $A(b)$, because minimum bid is a pure strategy, and hence can be verified on a case-by-case basis: even a single instance of the retailer accepting a bid below his stated minimum can be paraded in public as proof that the retailer is not credible. Therefore, the retailer can commit to a minimum bid by putting his reputation for trustworthiness on the line. Moreover, a minimum bid can also be made credible via common knowledge that getting cost quotes from suppliers is costly (a low bid is then not worth considering, because it does not cover the cost of a supplier quote in expectation).

[7] Note the FOC characterization of passive selling requires additional assumptions about H compared to the first best mechanism. One standard regularity assumption equivalent to the Myerson regularity discussed earlier in this paper is that $c+H(c)/h(c)$ is monotonically increasing in c.

Setting a minimum bid is clearly a realistic strategy because it is equivalent to a public reserve in auctions—a commonly observed feature in most auction markets. Most existing NYOP retailers already use a minimum bid, but none of them list it explicitly: for example, Priceline warns a buyer submitting a low bid that “Your price has almost no chance of being accepted,” and almost never accepts the bid after that warning. Chiching.com rejects a low bid, and responds, “Be a risk taker! Try some higher numbers.” Although minimum bids are already in use, we suspect their levels are generally too low: in private discussions, NYOP managers revealed to us that they set minimum bids low because they are concerned about turning down potential trades in the (mistaken) belief that sales volume and market efficiency go hand in hand with profits. In this section, we hope to provide clear guidance about how to better set the levels.

Another reason to analyze the minimum-bid strategy is the fact that minimum bids are a salient feature of the first best policy. It is interesting to investigate how much of the profit-increasing potential of the first best policy is achieved via rejecting low bids alone, and how the added profit from the first best policy depends on the outside market price p. We characterize the optimal minimum-bid strategy by first solving for the optimal bidding behavior, and then considering the best response of the retailer to the bidder.

Bidding strategy Assume $m \le p$ to ensure the possibility of trade. Buyers with $x < m$ do not enter the NYOP market, because they cannot earn a non-negative surplus. Buyers with $x \ge m$ solve a constrained version of Eq. 9:

$$\beta_m(x) = \underset{b \ge m}{\operatorname{argmax}}\, H(b)(x-b) + [1-H(b)]\max(0, x-p) \tag{9}$$

It is immediate that as long as $H(b)(x-b)$ is decreasing in b on $[m,x]$,[8] a mass of buyers with valuations just above m all pool on bidding m : $\beta_m(x)=\max(\beta_0(x),m)$. We now turn to describing how the retailer should optimally set the minimum bid m.

Optimal minimum bid The FOC of the bidding problem against a passive retailer implies that the mass of bidders bidding exactly m can be expressed in terms of F and H:

$\beta_0(x) > m \Leftrightarrow x > m + \frac{H(m)}{h(m)} \equiv x_m$, so bidders with $x \in [m, x_m]$ bid exactly m. The retailer profit thus involves two cases depending on whether m is so high that $x_m > p$ (and hence all bidders bid m) or m is low enough that $x_m < p$ (and hence some high-valuation bidders bid more than m):

$$\Pi_m(m,p) = \begin{cases} x_m < p : [F(x_m) - F(m)] \displaystyle\int_0^m H(c)dc + \int_{x_m}^p \pi_0(\min(x,p))dF(x) \\ x_m \ge p : [1 - F(m)] \displaystyle\int_0^m H(c)dc \end{cases} \tag{10}$$

[8] Concavity of $H(b)(x-b)$on the same interval is sufficient but not necessary for this.

The following proposition characterizes the optimal minimum bid the retailer should set:

Proposition 3 *Let* $D(m) = \left[F\left(m + \frac{H(m)}{h(m)}\right) - F(m)\right]H(m) - f(m)\int_0^m H(c)dc$.
When $D(m)>0$ *for all m such that* $x_m<p$, *the candidate* m^* *for the optimal minimum bid is implicitly characterized by* $\psi(m^*)=E(c|c<m^*)$, *and the candidate* m^* *exceeds the optimal bid in the first best mechanism. Otherwise, the candidate satisfies* $D(m^*)=0$. *Given the candidate* m^*, *the optimal minimum bid is* $min(m^*,p)$.

The proof of Proposition 3 is straightforward: first of all, the optimal minimum bid obviously cannot exceed p. D is the first derivative of the first case of Π_m shown in Eq. (10). When D is increasing for all low-enough m, the m^* must be in case two (Eq. 10). The relatively simple expression $\psi(m^*)=E(c|c<m^*)$ is the first-order condition of the second case of Eq. (10). From the same simple expression, it is immediate that the optimal minimum bid exceeds the minimum bid in the first best mechanism of $\psi^{-1}(0)$.

Example of the first case: When F = Uniform[0,1] and H = Uniform[0,p], $D(m) = \frac{m^2}{p} > 0$, so we are in the second case and all bidders with $x>m$ bid exactly m. The H = Uniform implies $E(c|c<z) = \frac{z}{2}$, so $m^* = \frac{2}{3} > \frac{1}{2} = \psi^{-1}(0)$ and the optimal minimum bid is $\min\left(\frac{2}{3}, p\right)$. The resulting retailer profit is

$$\Pi_m(p) \equiv \max_m \Pi_m(m,p) = \begin{cases} p > \dfrac{2}{3} : \dfrac{2}{27p} \\ p < \dfrac{2}{3} : (1-p)\dfrac{p}{2} \end{cases}.$$

Example of the second case: When F is the decreasing-triangle distribution $F(x)=x(2-x)$ and H is the increasing-triangle distribution $H(c) = \left(\frac{c}{p}\right)^2$, then $x_m = \frac{3m}{2}$ and $\beta_0(x) = \min\left(\frac{2x}{3}, p\right)$. Therefore, $m^* = \frac{4}{7}$ when $p > \frac{6}{7}$ and $m^* = \frac{3}{5} > \frac{1}{3} = \psi^{-1}(0)$ otherwise (detailed derivation of the latter omitted). The optimal minimum bid and the corresponding behavior are thus as follows:

$$\begin{cases} p \le \dfrac{3}{5} : \text{ optimal minimum bid } = p, \text{ everyone with } x \ge p \text{ bids } p \\ \dfrac{3}{5} < p < \dfrac{6}{7} : \text{ optimal minimum bid } = \dfrac{3}{5}, \text{ everyone with } x > \dfrac{3}{5} \text{ bids } \dfrac{3}{5} \\ p > \dfrac{6}{7} : \text{ optimal minimum bid } = \dfrac{4}{7}, \text{ people with } x \in \left[\dfrac{4}{7}, \dfrac{6}{7}\right] \text{ bid } \dfrac{4}{7}, x > \dfrac{6}{7} \text{ bid } \beta_0(x) \end{cases}$$

The $\psi(m^*)=E(c|c<m^*)$ condition in Proposition 3 is intuitively related to monopoly posted pricing, in which the optimal price to post when facing a production cost of k satisfies $\psi^{-1}(k)$. The following corollary clarifies the relationship:

Corollary to proposition 3 *When the optimal minimum bid is high enough that all bidders bid it, the optimal minimum bid is set at the level of the monopoly posted price that would be optimal for a posted-price retailer facing a procurement cost that depends on price charged, such that the procurement cost is the draw from distribution H truncated above at the posted price.*

In other words, when all bidders bid the minimum bid amount, the optimal minimum bid feels like a posted price. But the minimum-bid strategy is clearly not equivalent to charging a posted price, because the minimum-bid retailer is not obliged to sell to every buyer willing to pay m. Instead, the sale is made only when m happens to exceed the actual procurement cost c. The corollary casts the minimum-bid problem into a posted-price problem, in which the procurement cost is endogenous to price in a very particular manner.

To illustrate the difference between the minimum bid and the posted price, note that a retailer forced to charge a fixed posted price irrespective of cost would maximize the expected profit $[1-F(r)][r-E(c)]$ subject to $r \le p$, and so charge $r^*=\min[\psi^{-1}(E(c)),p]$. A comparison with the $\psi(m^*)=E(c|c<m^*)$ condition of Proposition 3 shows the optimal fixed posted price r^* exceeds the optimal minimum bid m^* as long as they are both below p. Intuitively, the fixed posted price needs to be higher than the minimum bid to compensate for higher average costs that the fixed-price seller incurs. The lack of flexibility also makes it obvious that the fixed-price strategy is less profitable then. When $r^* = m^* = p$, the two strategies and their profits obviously coincide, because all bids are accepted by construction (we assumed $c \le p$ in defining H).

Uniform-uniform example: When F is uniform on $[0,1]$ and H is uniform on $[0,p]$, the optimal posted price would thus be $\min\left(\frac{1}{2}+\frac{p}{4},p\right)$. The $r^*>m^*$ inequality noted above manifests itself in the example because $\frac{1}{2}+\frac{p}{4}<p \Leftrightarrow p>\frac{2}{3}$.

6.3 Participation fee: a two-part-tariff strategy

Spann et al. (2010, 2015) assume F = Uniform[0,1] and H = Uniform[0,p], and show the retailer profits more from charging a participation fee than from charging a minimum markup. They also show the participation fee alone is more profitable than any combination of a participation fee and a minimum markup. We now repeat their main result in our notation, focusing on $p>\psi^{-1}(0)=\frac{1}{2}$. The optimal fee e^* is $\sqrt{e^*(p)}=\frac{2}{7\sqrt{p}}$, and the resulting retailer profit is $\Pi_e(p)=\frac{4}{147p}+\Pi_0(p)$ (Proposition 2 of Spann et al. 2015). The participation fee screens low-valuation buyers out of the market, and the implied entry threshold is $x_e=\frac{4}{7}$. In other words, buyers with $x \ge \frac{4}{7}$ pay the fee and participate. The retailer then accepts all bids above cost, so his NYOP selling strategy is $A(b)=\Pr(b>c)$, and the uniform assumption on H implies buyers who enter bid $\beta_0(x)=\frac{x}{2}$.

Given a general F and H, the retailer's profit is most conveniently expressed in terms of the entry threshold, keeping in mind that the fee paid by all participants equals the expected surplus of the marginal entrant:

$$\Pi_e(x_e,p)=[1-F(x_e)]\underbrace{H(\beta_0(x_e))(x_e-\beta_0(x_e))}_{e=S(x_e)\equiv \text{ expected surplus of } x_e}+\int_{x_e}^{1}\pi_0(\min(x,p))dF(x) \tag{11}$$

Our next proposition chararacterizes the optimal entry threshold:

Proposition 4 *The candidate optimal entry threshold in a participation-fee strategy satisfies* $\psi(x_e)=E(c|c<\beta_0(x_e))$, *and the optimal entry threshold is* $\min(x_e,p)$.

Note that although elegant, the condition in Proposition 4 involves the bidding strategy, and so it does not characterize x_e purely in terms of the model parameters F and H (in contrast to previous propositions in this paper). Comparing Propositions 3 and 4, we can conclude the following immediate corollary from $\beta_0(x)\leq x$:

Corollary to proposition 4 *When the optimal minimum bid is high enough that all bidders bid it, it exceeds the optimal entry threshold of the participation-fee strategy.*

We now turn to profit comparisons among all the selling strategies considered in this paper.

6.4 Comparing the four alternative strategies: analytical results

Closed-form analytical profit comparisons are available for the F = Uniform[0,1] and H = Uniform[0,p] assumption used throughout the preceding examples. Before looking at profits, considering the allocation rules implied by the alternative strategies is useful because the revenue-equivalence theorem links allocation rules tightly to profits. To summarize the theoretical results so far, when cost is c, a customer x is allocated the when:

Optimal mechanism: $\psi(x)\geq c$ or $x\geq p$
Minimum bid: $\psi(x)\geq E(c|c<x)$ *and* $m\geq c$
Fixed price: $\psi(x)\geq E(c)$
Participation fee: $\psi(x)\geq E(c|c<\beta_0(x))$ and $\beta_0(x)\geq c$

where we assume for simplicity that the first case of Proposition 3 applies, namely, that all bidders submit the minimum bid and that p is high enough that it does not "take over" as the minimum bid, fixed price, or the entry threshold. Figure 3 illustrates these allocation rules given the uniform-uniform assumption with the outside price fixed at p= 4/5.

Figure 3 shows that the participation fee and the passive selling strategies can never (under any p) lead to the same allocation as the optimal policy of Proposition 1 because they never allocate the good to the buyer when $x < 2c$, but the optimal policy allocates it to all $x > p$ irrespective of c. In addition, Fig. 3 illustrates that when p= 4/5, all the allocation rules are different from each other. The fact that all buyer types pay the same amount under the optimal minimum-bid and fixed-price strategies makes the allocation rules of those strategies look like rectangles. Because the rule of the optimal policy is not a rectangle, it must be strictly more profitable than the minimum-bid and fixed-price strategies. However, recall from section 6.2 that $\psi^{-1}(0)\leq m^*\leq r^*\leq p$. This inequality implies that as p drops down toward $\frac{1}{2}$, the allocation rules of these two simple strategies approach the optimal allocation rule. Therefore, and unlike the passive and participation-fee strategies, the minimum-bid and fixed-price strategies can be as profitable as the optimal strategy when p is low. The reason is obvious: the first best strategy reduces to a fixed-price strategy when p=$\psi^{-1}(0)$.

Figure 4 plots profits of all four alternative selling strategies and the first best strategy as a function of the price of the outside option (under the F = Uniform[0,1] and H = Uniform[0,p] assumption used throughout this section), and illustrates our last proposition:

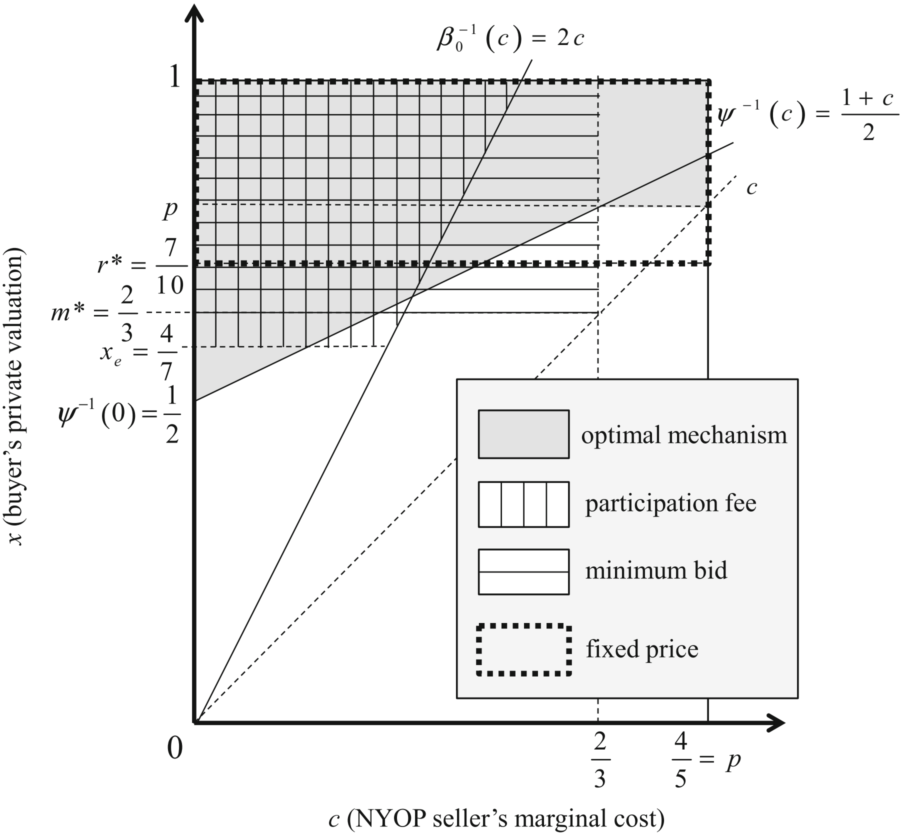

Fig. 3 Allocation rules: the optimal mechanism vs. simpler strategies. Note to Figure: The posted price p is set to 4/5 throughout the figure. The support of (x,c) is in the px1 rectangle, and $F(x)$ is uniform. The shaded area indicates the allocation rule of the optimal mechanism. The area filled by vertical lines indicates the allocation rule of the participation-fee strategy with the fee set optimally at 5/49 and the resulting entry-valuation threshold of 4/7. The area filled by horizontal lines indicates the allocation rule of the minimum-bid strategy with the minimum bid set optimally at 2/3. The area outlined with a thick dotted line indicates the allocation rule of the fixed-price strategy with the price level set optimally at 7/10. The passive-strategy allocation rule corresponds to the triangle above the $x = 2c$ line

Proposition 5 *When F is uniform on* $[0,1]$ *and H is uniform on* $[0,p]$, *as p approaches* $\psi^{-1}(0) = \frac{1}{2}$ *from above, the relative profit advantage of the first-best strategy over the minimum-bid strategy and the fixed-price strategy vanishes. The relative ranking of the simple strategies is as follows:*

a) *the minimum-bid strategy strictly dominates the participation-fee strategy for all p, and the profit advantage is decreasing in p.*
b) *the minimum-bid strategy is equivalent to the fixed-price strategy for* $p \leq \frac{2}{3}$ *and strictly dominates the fixed-price strategy otherwise.*
c) *a p* exists such that the fixed-price strategy strictly dominates the participation-fee strategy for p<p*, and* vice versa.
d) *the passive strategy is strictly dominated by all other strategies.*

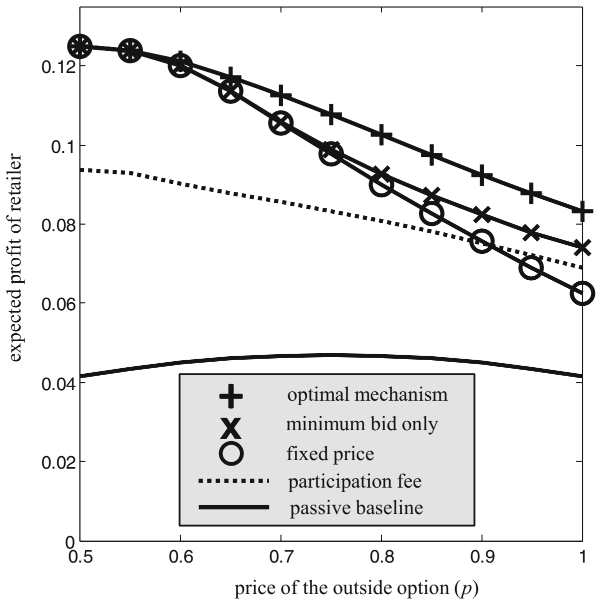

Fig. 4 Expected profit of NYOP retailers: four alternative strategies. Note to Figure: Illustration of Proposition 5. F is uniform on [0,1] and H is uniform on $[0,p]$

To gain intuition for the minimum bid approaching full optimality as p approaches half, note that for $p=\frac{1}{2}$, both the first best contingent bid-acceptance rule and the optimal minimum-bid strategy reduce to posted pricing. Both the first best NYOP retailer and his minimum-bid counterpart abandon low-valuation buyers completely, focusing on stealing the high-valuation buyers from the outside market and charging them p.

The full optimality of the minimum-bid strategy at $p=\frac{1}{2}$ also immediately implies its dominance over participation fees for low p. The difference between these two simple NYOP strategies declines as p rises, because the retailer who uses a participation fee is better able to extract profit as p increases. Both retailers face rising costs as p increases, but the minimum-bid retailer does not benefit from a compensating increase in bids beyond $p=2/3$. To see why, note that for $p>2/3$, the outside price does not affect entry and bidding when a retailer commits to a minimum bid. Therefore, the profit $\Pi_m(p)$ is decreasing in p purely because of rising retailer costs. The participation-fee retailer faces the same rising costs, but p also increases his buyers' bids because the bidding strategy is $\beta_0(x)=\frac{\min(x,p)}{2}$. In other words, the bidding subsequent to paying a participation fee is better at price discriminating than bidding at a minimum-bid retailer. Therefore, $\Pi_e(p)$ decreases more slowly than $\Pi_m(p)$. Proposition 5 shows they never cross over: even at $p=1$, the minimum-bid strategy is more profitable.

The intuition behind claim b) is clear from the fact that the minimum-bid and fixed-price strategies coincide exactly for low p, because the retailer sets the minimum bid at p, effectively guaranteeing a sale to all buyers with $v>p$ because it is common knowledge that $c<p$. In other words, the minimum-bid strategy for $p\leq\frac{2}{3}$ is effectively a fixed-price strategy. For higher p, the fixed-price seller charges a higher price than the minimum bid to hedge against higher average costs as shown in section 6.2. The

minimum-bid seller reaps part of the resulting efficiency benefit, because he can still turn bids down when cost happens to be too high.

Claim c) regarding the relationship between the fixed-price strategy and the participation-fee strategy is good news for participation fees: if they were dominated by the much simpler fixed price, recommending them to NYOP sellers would be dubious advice. Instead, the claim shows participation fees can be relatively profitable, but only when the retailer is a near-monopolist and cannot commit to a minimum-bid strategy.

6.5 Profit comparison of the four alternative strategies: simulation results

This section explores how the results of Proposition 5 generalize to other distributional assumptions beyond uniform F and H, with a special focus on the minimum-bid strategy that won the horserace among simpler strategies in Proposition 5. To explore a wide variety of distributional shapes while maintaining as much tractability as possible, we fix $p=4/5$ throughout and we consider the Kumaraswamy distribution for both F and H, parametrized as

$$F(x) = 1-(1-x^{\alpha})^{\gamma} \text{ on}[0,1] \text{ and } H(c) = 1-\left(1-\left(\frac{c}{p}\right)^{\rho}\right)^{\tau} \text{ on the}[0,p] \qquad (12)$$

The Kumaraswamy distribution is closely related to the Beta family, can capture the same variety of density shapes (flat, increasing, decreasing, concave, convex, U, inverted-U), and, unlike the Beta, has a closed-form *cdf,* simplifying many of the computations. In the simulation study, we varied all four parameters $\{\alpha,\gamma,\rho,\tau\}$ independently of each over the range $\{0.75,1,2,4\}$ in a full factorial design. Therefore, we considered $4^4=256$ pairs of (F,H) distributions. Please see Fig. 5 for an illustration of distribution shapes of F that result from varying both α and γ (the corresponding H shapes are analogous, with the support reduced to $[0,p]$). The relatively high value of $p=4/5$ was chosen because several of the NYOP selling strategies (including the first best one) boil down to uninteresting posted-price selling when p is low relative to the distribution of valuations. To assess when $p=4/5$ is "low," Fig. 5 shows the proportion of buyer valuations above $p=4/5$ for every α and γ.

The main goal of the simulation study is to explore the boundary conditions of the main takeaways from Fig. 4 and Proposition 5 that the minimum-bid strategy

1) achieves much of the first best strategy's profit lift relative to passive selling and
2) dominates all other simple strategies under consideration.

We answer these questions in turn by comparing the profitability of the minimum-bid strategy first to that of the optimal strategy, and second to that of the participation-fee strategy. In each study, we analyze the following ratio, and we call it relative profit lift:

$$\textit{Relative Profit Lift}(\text{minimum bid, alternative}; p) = \frac{\Pi_m(p)-\Pi_0(p)}{\Pi_{alternative}(p)-\Pi_0(p)} \qquad (13)$$

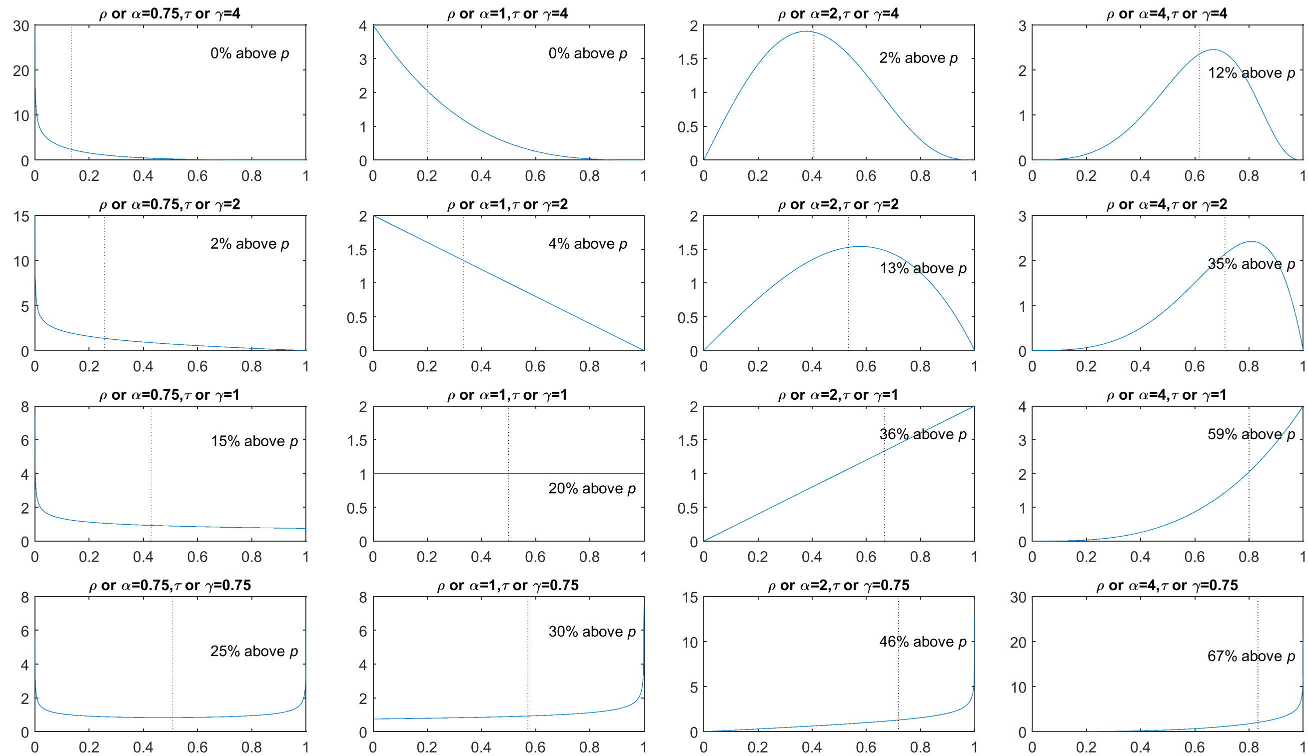

Note to Figure: Each plot shows one of the considered distributions of valuations (indicated by parameters α,γ in the plot title). The dashed vertical line indicates the mean. The considered distributions of cost (indicated by ρ,τ) have the same shape, but their support is only $[0,p]$. The text inside each plot indicates the percentage of valuation probability mass above $p=4/5$, namely $1\text{-}F(4/5;\,\alpha,\gamma)$ rounded to integer level.

Fig. 5 Kumaraswamy distribution shapes considered for both F and H in the simulation

Minimum bids vs. *optimal NYOP selling*: Figure 6 shows the relative profit lift of the minimum-bid strategy versus the optimal mechanism for all combinations of F and H under consideration. The optimal mechanism weakly dominates the minimum-bid strategy, so the relative profit lift is between 0 and 1 and can be interpreted as a percentage of the maximum theoretical profit improvement that the minimum-bid strategy achieves. Each subplot in Fig. 6 fixes a particular distribution of valuations (indicated by parameters α,γ in the plot title) and considers all distributions of cost (all combinations of cost-distribution parameters ρ,τ). The contours are isoquants of relative profit lift spaced 0.1 apart, i.e. deciles of the percentage of maximum lift achieved. To help interpret the results, Fig. 6 also delineates regions of the parameter space when the optimal minimum bid is equal to the outside price p and regions when the optimal minimum bid is so low that some bidders bid more than the minimum bid (this situation is called the "second case" in Proposition 3).

The number inside each plot shows the magnitude and approximate location of the lowest relative lift in the plot. For example, the lowest relative profit lift under uniformly distributed valuations (α=1,γ=1) is 0.76 when ρ= 2 and τ=0.75. Given that low point, the two contour lines shown in the (α=1,γ=1) plot correspond to the 0.8 and 0.9 levels of the relative profit lift. We can thus conclude at a glance that when valuations are distributed uniformly, the minimum-bid strategy achieves at least 76 % of the profit lift achieved by the optimal strategy, and the relative profit lift rises as ρ falls and τ rises, that is (see Fig. 5), as costs become concentrated at the low end and the mass just below p approaches zero. Having explained the plots, we now turn to the substantive results.

The previous section found that, at least in the uniform-uniform case, much of the profit-increasing potential of the first best policy is achievable via just rejecting low bids in a simpler minimum-bid strategy. The relative profit lift of the minimum-bid strategy in the uniform-uniform case can be expressed analytically as 496/603≈0.82 (the location of this value in the $\{\alpha,\gamma,\rho,\tau\}$ space is shown by the star symbol in the (α=1,γ=1) plot of Fig. 6). The first result of our simulation study is that the finding generalizes to situations when $\gamma\leq 1$, that is, when the valuation distribution has sufficient mass at the top. Because we are considering p= 4/5 throughout, the exact shape of F above 4/5 does not matter, and the percentage of "high-valuation" bidders who can afford the outside option is a relevant metric for the amount of mass at the top of the distribution. We find the percentage of bidders who can afford the outside option has a high correlation of 0.83 with the average (over different cost parameters) relative lift of the minimum-bid strategy. Intuitively, the minimum-bid strategy always increases bids of people with valuations above the optimal minimum bid, but this increase only boosts retailer profits when enough bidders are above the minimum. Because the minimum bid is weakly below p, the percentage of high-valuation bidders is a proxy for the amount of bidders above the optimal minimum (which depends on both distributions in complicated ways).

Although the percentage of high-valuation bidders goes a long way toward explaining the effectiveness of the minimum-bid strategy, the shape of the valuation distribution below p also matters. Consider the (α=2,γ=2) and (α=0.75,γ=1) situations: they have almost the same number of high-valuation bidders, but the relative effectiveness of minimum bids is both lower on average (75 % vs. 85 % relative lift) and also less robust to different cost distributions (much lower when costs are concentrated at the top) when (α=2,γ=2). Consider the (ρ= 4, τ=0.75) cost distribution for

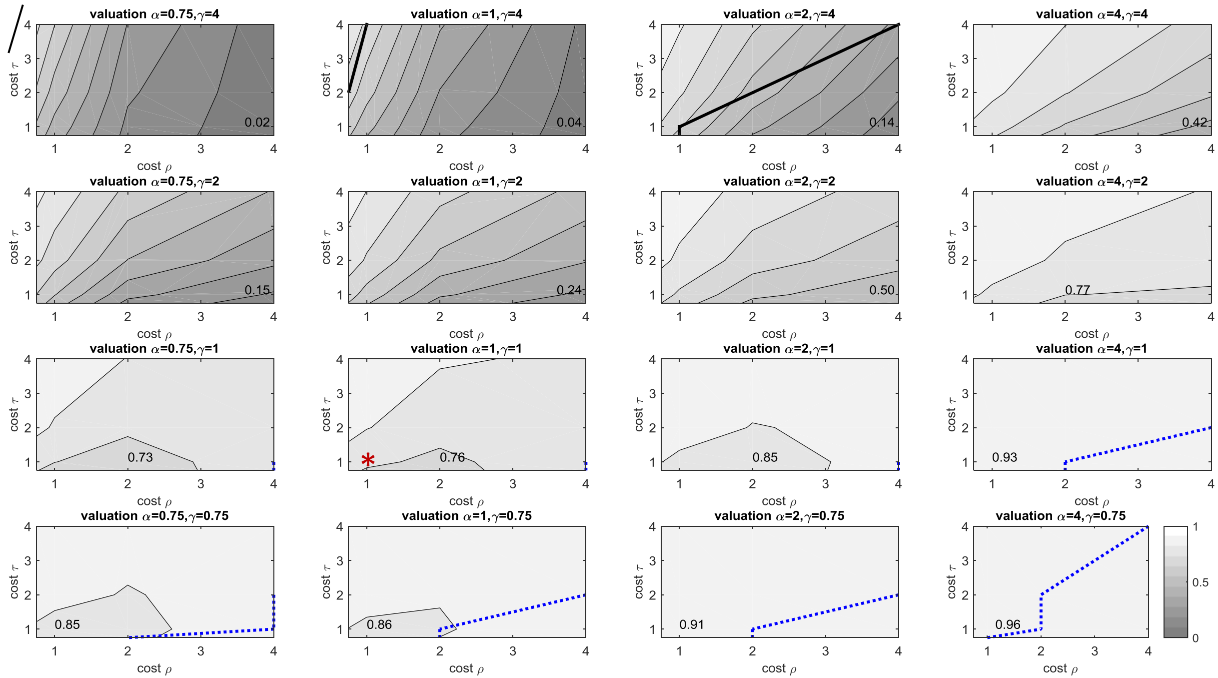

Note to Figure: The posted price in the outside market is set at p=4/5 throughout. Each contour plot fixes a particular distribution of valuations (indicated by parameters α,γ in the plot title) and all considered distributions of cost (all combinations of cost-distribution parameters ρ,τ). The contours are isoquants of relative lift. The number in each plot shows the magnitude and location of the lowest relative lift in the plot. Cost-parameter combinations below the **thicker dashed** (blue) line involve a minimum bid equal to p, the lack of a dashed line in a plot indicates that all optimal minimum bids are strictly below the outside market price p. Cost-parameter combinations below the **thicker solid** (black) line involve optimal minimum bid low enough that some bidders bid more than the minimum bid. When α=0.75 and γ=4, this is true for all (ρ,τ).

Fig. 6 Relative profit lift of the minimum-bid strategy: *(minimum bid – passive) / (optimal – passive)* which the relative lift increases from 50 % for (α=2,γ=2) to 88 % for (α=0.75,γ=1): the optimal minimum bid equals p for (α=0.75,γ=1), so trade is guaranteed by virtue of $c \leq p$. On the other hand, the optimal minimum bid is 0.76 for the (α=2,γ=2) case, which leaves 19 % of the costs too high, preventing trade. The downward-sloping density of valuations just below p is clearly putting a downward pressure on the optimal minimum bid, and in turn reducing efficiency of the market. In other words, the minimum-bid strategy does not perform well (relative to the optimal strategy) when lowering the minimum results in a lot more additional bidders, that is, when the demand is steep near the level of the optimal minimum bid, which in turn depends on both the demand and the cost distributions.

The intuition that locally steep demand hurts the minimum-bid strategy explains much of the pattern of results in Fig. 6 beyond just the pair of valuation distributions studied in the previous paragraph. Consider the steep hill-shaped F(α=4,γ=4) case: when costs are concentrated at the top and the optimal minimum bid is thus on the "downhill" side of f, the minimum-bid strategy achieves only 42 % relative lift. However, when the costs are concentrated at the bottom and the optimal minimum bid is thus on the "uphill" side of f, the minimum-bid strategy achieves 98 % of the relative lift. We summarize the above findings in the following result:

Simulation result 1 Regardless of the cost distribution, the minimum-bid strategy captures much of the maximum theoretical profit available from the first best NYOP selling strategy when a lot of buyers can afford the outside option. Across all cost distributions considered, the minimum-bid strategy increases profits at least *73 %* as much as the first best strategy when $\gamma \leq 1$. In addition, the minimum-bid strategy captures more of the maximum theoretical profit when the density of valuations around the minimum bid level is not downward sloping.

Minimum bids vs. participation fees To study question 2), we first note that the minimum-bid strategy will always at least weakly dominate the fixed-price and passive strategies, so question 2) boils down to a comparison between the minimum-bid and the participation-fee strategies. Figure 7 shows the profit lift of the minimum-bid strategy relative to the participation-fee alternative (see Equation 13 for the definition of relative lift). The equal-lift line is highlighted in thickness, and the lighted shading highlights situations in which the minimum-bid strategy is more profitable. All other nomenclature is analogous with that in Fig. 6.

It is immediate from Fig. 7 that the finding from the uniform-uniform case does not generalize to valuation distributions concentrated at the bottom of the support. Instead, we find regions of the parameter space in which participation fees dominate minimum bids. Not surprisingly, those regions correspond to the regions analyzed above in which the minimum-bid strategy is weak relative to the optimal mechanism. Another observation we make based on Fig. 7 is that participation fees dominate minimum bids only when $\alpha < \rho$ and $\gamma < \tau$. In words, participation fees outperform minimum bids only when valuations are concentrated at the bottom of the support and tend to be substantially lower than costs. We summarize these findings as:

Simulation result 2 The minimum-bid strategy is more profitable than the participation-fee strategy whenever valuations tend to be higher than the costs, that is, when the gains

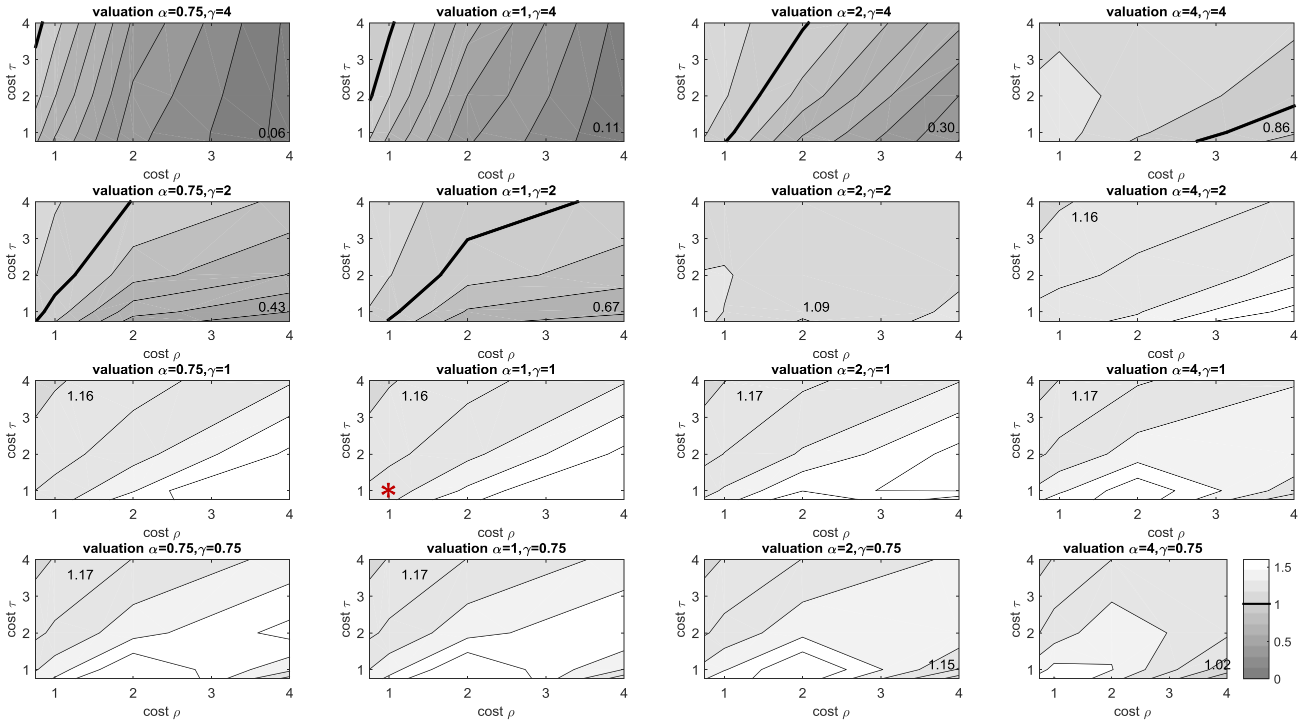

Note to Figure: The posted price in the outside market is set at p=4/5 throughout. Each contour plot fixes a particular distribution of valuations (indicated by parameters α,γ in the plot title) and all considered distributions of cost (all combinations of cost-distribution parameters ρ,τ). The contours are isoquants of relative lift. The 1 isoquant above which minimum bid outperforms the participation fee is highlighted by **thickness**. The number in each plot shows the magnitude and location of the lowest relative lift in the plot.

Fig. 7 Minimum bid vs. participation fee: Profit lift *(minimum bid – passive)/(participation fee – passive)* from trade are large. Across all cost distributions considered, the minimum-bid strategy increases profits at least 15 % more than the participation-fee strategy when $\alpha \leq 2$ and $\gamma \leq 1$. Conversely, the participation-fee strategy dominates when valuations are concentrated at the bottom of the support ($\gamma > 2$) and the gains from trade are small.

To gain intuition for why participation fees can dominate the minimum bids, recall that the participation-fee strategy is good at realizing gains from trade because it captures the marginal entrant's surplus (see proof of Proposition 4 for the role of gains from trade in the retailer's optimization). One way to see how the participation-fee strategy is better at realizing gains from trade, note that although both strategies exclude buyers with very low valuations from the market, the minimum-bid strategy usually excludes more buyers (always true when all bidders bid the minimum bid; see Corollary to Proposition 4). When not many high-valuation buyers are present (and the minimum-bid strategy is thus weak) and the gains from trade are low (most valuations are below costs), the participation-fee strategy can thus be more profitable overall. For example, when valuations are concentrated at the bottom of the support ($\gamma = 4$), the correlation between the minimum bid's relative lift and gains from trade is 0.71.

For a concrete example of participation fees outperforming minimum bids, consider the $\{\alpha=1, \gamma=2, \rho=2, \tau=1\}$ situation in the second closed-form example of the minimum-bid strategy. Proposition 4 and the fact that $\beta_0(x) = \min\left(\frac{2x}{3}, p\right)$ imply the entry threshold is $x_e = \frac{9}{19} \approx 0.47 < m^* = \frac{3}{5}$. The optimal participation fee is the expected surplus of the marginal entrant, which is only about 0.024 here. Given $p = 4/5$, all bidders facing the participation fee bid less than their counterparts facing the minimum-bid strategy, so the minimum-bid strategy raises more bidding revenue. However, the seemingly small fee paid by the roughly 27 % of buyers with $x > x_e$ is enough to tip the profit balance in favor of the two-part-tariff strategy.

7 Discussion

Name-your-own-price (NYOP) selling accommodates buyer activism whereby buyers submit bids for products the retailer may or may not be able to subsequently procure at a low-enough cost. We define a selling strategy of an NYOP retailer and optimize it. A selling strategy is a schedule of bid-acceptance probabilities that depend only on the bid level and on the commonly known outside spot-market price. The retailer needs to set a single such strategy for a range of possible costs upfront, but he learns the current cost realization before making each bid-acceptance decision. We use mechanism-design techniques to craft the optimal bid-acceptance probability schedule. Our solution is a pair of the optimal bid-acceptance probability function and the buyer bidding function that best responds to it, and the only assumption we need is that the distribution of buyer valuations be regular in the sense of Myerson (1981).

Neither the optimal selling strategy nor the bidding strategy that best responds to it are trivial: the bid-acceptance probability is increasing for small bids, equal to unity for very high bids, and only partially determined for intermediate bids: optimality dictates only an upper bound of the policy. The reason is that the optimal allocation calls for "high" bidders who could potentially afford the outside option to receive the object with certainty, but bidders who barely cannot afford it receive the object with a probability strictly less than unity (otherwise, the high bidders would mimic lower-valuation bidders). The equilibrium bidding function thus needs to involve a jump discontinuity whereby the bidders who could potentially afford the outside option all pool at a particular bid level (which we prove is unique). Pooling occurs among high-valuation bidders because the outside posted-price market is a real option for them, and so they all mimic the buyer with valuation equal to the outside price.

We find the optimal selling strategy allows the NYOP retailer to make as much profit as he would if he could learn his cost first and use the optimal mechanism contingent on it (it is well known that the contingent optimal mechanism is to make a take-it-or-leave-it offer at a fixed price). In other words, we show the optimal strategy can achieve first-best ex-post profits despite the retailer not knowing his cost realization at the time of announcing his strategy. The intuition for this result is that at the time of making the bid-acceptance decision, the retailer knows both his cost and the buyer's valuation, and so he can base his decision on both pieces of information available to the first-best seller. The retailer learns the buyer's valuation by inverting the bidding function at the observed bid. Seemingly, the retailer could use any increasing bid-acceptance probability to achieve his goal of inverting the bidding function, but only the bid-acceptance probability implied by the ex-post optimal allocation is sustainable in equilibrium (buyers are not deceived).

In the proposed static model, NYOP is the only chance to buy the good for the buyer who cannot afford the outside market price. One may conclude the bids of those buyers are thus artificially inflated via a restriction on supply. This assumption is not critical to the revenue-equivalence results: in a more complicated dynamic model, the buyer can also wait to see whether posted prices drop, or even try NYOP bidding again at a later stage. Either of these real options to obtain a substitute product in the future would reduce buyer bids. However, the same options would also reduce the buyer's willingness to pay a posted price today as in the durable goods monopoly literature. Therefore, the revenue equivalence between dynamic NYOP and the corresponding dynamic posted pricing would likely continue to hold.

The NYOP retailer needs to credibly communicate the strategy to prospective buyers. Credibility requires commitment to a particular acceptance probability for every possible bid level, but is does *not* require the retailer to credibly communicate his cost realization or commit to any action contingent on a cost realization. The competition with the outside posted-price market requires the optimal retailer to sometimes subsidize some unprofitable high bids (e.g., bids by high-value bidders who all bid as if their valuation were equal to the outside posted price). This type of commitment is a new requirement brought about by the NYOP format in that it is not required for standard monopoly posted pricing.

The complexity of the first best strategy may be too impractical for some real-world markets. We consider three types of simpler second-best NYOP strategies: first, retailers may be able to commit to considering bids only above some minimum level, but not to a probability of bid acceptance above that level. Second, the retailers may be able to charge a participation fee, that is, ask to be paid for considering bids, as proposed by Spann et al. (2010, 2015). Finally, we also consider the non-NYOP strategy of just posting a fixed selling price before knowing the cost realization. We compare the profitability all these strategies to both the profitability of the first-best strategy and the profitability of the passive strategy of accepting all bids above cost. We perform this comparison numerically for a wide range of possible cost and valuation distributions given a particular level of the outside price, and we perform it analytically for all levels of the outside price when both distributions are uniform.

Not surprisingly, we document that all "active" strategies strictly outperform the passive strategy of accepting all bids above cost, and the minimum-bid strategy weakly dominates the fixed posted-price benchmark. More surprisingly, we find the minimum-bid strategy achieves much of the maximum theoretical profit (i.e., the profit of first best strategy) in most of our simulation scenarios. Specifically, the minimum-bid strategy does particularly well when the distribution of valuations involves a lot of consumers who can afford the outside option, and when the density of valuations around the optimal minimum bid level is not downward sloping. The intuition for these drivers of minimum-bid's profitability is that the minimum bid strategy is good at increasing high-valuation buyers' bids. When enough such buyers are present (probability of affording the outside option is a proxy for this scenario) and when the optimal minimum bid is relatively high (as is the case when the density of valuations is non-decreasing), the increase in high-value buyers' bids results in a lot of additional profit relative to the passive strategy.

In terms of bidding behavior, the minimum-bid strategy results in many, if not all, buyers pooling at the minimum bid level. This behavior is different from the effect of a reserve price in an auction, where only the buyer with valuation equal to the minimum bid bids exactly the minimum bid. The reason for the difference is that an NYOP bidder who can afford the minimum bid does not have to compete with other bidders who can also afford it, and so he faces a positive probability of winning the item.

Regarding the comparison of the proposed simpler strategies to each other, we find that both the minimum-bid strategy and the participation-fee strategy can dominate. The participation-fee strategy is good at realizing gains from trade; that is, it is more efficient than the minimum-bid strategy. The intuition is the same as for the standard two-part tariff in the posted-price setting: a seller using a two-part tariff wants to make trades happen because he captures the full gain from trade of the marginal entrant. When the potential gains from trade are small, that is, when most costs exceed most valuations, the minimum-bid strategy is not very profitable (see previous paragraph), and the participation-fee strategy thus dominates it because it is good at capturing those gains. In summary, we conclude that much of the "heavy lifting" of the first-best strategy is often accomplished by the much simpler minimum-bid strategy, but NYOP retailers in thin markets can do better by charging participation fees.

Regarding practical implementation, the mechanism proposed here is more in line with "select your price" than with true "name your own price" in Chernev (2003). Moreover, we suggest the retailer should present not only a menu of prices, but also the acceptance probabilities. Doing so facilitates commitment (either via a third-party auditor or via reputation, as we describe in detail) and simplifies bidding.

Outside the scope of this paper, credibly communicating the acceptance probabilities also enables the retailer to better learn consumer preferences by controlling consumer beliefs about chances of bid acceptance. Specifically, the retailer should be able to easily use the bid-inversion approach designed for analysis of first-price sealed-bid auctions (e.g. Guerre et al. 2000). Instead of using the empirical probability of winning at any given bid level (which needs to be estimated in the case of auctions), the retailer can simply plug in his selling strategy. We propose that this feature of optimal NYOP could generate a very clean source of demand information as a by-product of selling. At the minimum, this information would be useful to the retailer by allowing him to fine-tune his strategy over time.

Also beyond the scope of this paper, the modeling approach proposed here could be use to derive optimal procurement strategies under ex-ante valuation uncertainty in the "reverse" RFQ scenario, in which bids are offers to sell at a given price, and the procurement strategy needs to be set before the buyer learns about his valuation of the good or service being bought.

Appendix: notation table and proofs of propositions

Proof of proposition 1 Let $m(x)$ be the expected payment by a buyer with valuation x. From risk neutrality, the utility of a buyer x who reports type z is $U(x)=xq(z)-m(z)$, and standard incentive-compatibility arguments (see Myerson 1981 for details) imply the expected payment is the following function of $q(x)$:

$$m(x) = m(0) + xq(x) - \int_0^x q(t)dt \qquad \text{(IC)}$$

Table 1 Notation

x : buyer's valuation of the object
b : buyer's bid
$f(x),F(x)$: density and the cumulative distribution function of the distribution of x
$\psi(x)\equiv x-\frac{1-F(x)}{f(x)}$: virtual value of type x
c : retailer's procurement of the object
$h(c),H(c)$: density and the cumulative distribution function the distribution of c
p : lowest price posted in an outside spot market for the object
$\beta(x)$: buyer's bidding function
$A(b)$: the *NYOP selling strategy,* i.e., probability that retailer accepts bid b
$U(b,x)$: expected surplus of a buyer x who bids b
$q(x)$: ex-ante probability that retailer allocates the object to a buyer with valuation x. Subscripts on q indicate various second-best strategies in section 6.
$\pi(x,c)$: probability that retailer with cost c allocates the object to a buyer with valuation x
e : participation fee
m : minimum bid
r : posted price fixed upfront, i.e., before knowing cost c
$\Pi_s(p) = \max_s \Pi_s(s,p)$, where $\Pi_s(s,p)$ is the expected profit of a retailer who uses a selling strategy $s\in\{0,A,m,e,r\}$

When (IC) does not hold, the buyers do not have the incentive to report their x truthfully.

Consider the direct-revelation retailer who can set an arbitrary bid-acceptance rule $\pi(x,c)$. Plugging the implied bid-acceptance rule $q(x) = \int_0^p \pi(x,c)dH(c)$ into (IC) implies that *on average over all c*, such a retailer receives a payment of

$$m(x) = m(0) + x\int_0^p \pi(x,c)dH(c) - \int_{\underline{x}}^{x}\int_0^p \pi(t,c)dH(c)dt \tag{A1}$$

Note the rule π can use c as an input, so the retailer can set the rule for all possible c levels upfront. However, the buyers do not know c at the time of submitting their bids, so incentive compatibility only restricts the average payment of a given buyer type. Because all buyers with $x \geq p$ pay $m(p)$, the expected profit of the retailer is

$$\begin{aligned} \Pi(\pi) = \Pr(x < p)\left(E_{x|x<p}[m(x)] - E_{c,x|x\leq p}[c\pi(x,c)]\right) \\ + [1 - \Pr(x < p)][m(p) - E_c[c\pi(p,c)]] \end{aligned} \tag{A2}$$

Plugging the m function from (A1) into the profit expression (A2) yields

$$\begin{aligned} \Pi(\pi) = m(0) + \int_{\underline{x}}^{p} x\int_0^p \pi(x,c)dH(c)dF(x) - \int_{\underline{x}}^{p}\int_0^x\int_0^p \pi(t,c)dH(c)dtdF(x) - \int_{\underline{x}}^{p}\int_0^p c\pi(x,c)dH(c)dF(x) \\ + [1-F(p)]\left[p\int_0^p \pi(p,c)dH(c) - \int_0^p\int_0^p \pi(t,c)dH(c)dt - \int_0^p c\pi(p,c)dH(c)\right] \end{aligned} \tag{A3}$$

where the second row corresponds to the profit from high buyers ($x \geq p$), and the last term in each row is the expected cost of goods sold.

As in other mechanism-design settings, the term $\int_{\underline{x}}^{p}\int_0^x\int_0^p \pi(t,c)dH(c)dtdF(x)$ in the first row can be simplified by first changing the order of integration from c,t,x to x,t,c, and noting that $\pi(t,c)$ does not depend on x:

$$\begin{aligned} \int_{\underline{x}}^{p}\int_0^x \pi(t,c)dtdF(x) = \int_0^p\left(\pi(t,c)\int_t^p dF(x)\right)dt == \int_0^p \pi(t,c)[F(p)-F(t)]dt \\ = \int_{\underline{x}}^{p} \pi(x,c)\left(\frac{F(p)-F(x)}{f(x)}\right)dF(x) \end{aligned} \tag{A4}$$

where the last equality simply renames the t variable as x and changes variables. Finally, change the order of integration to be first over x and then over c throughout, and collect terms:

$$\begin{aligned}\Pi(\pi) = m(0) + \int_0^p \int_{\underline{x}}^p \left(x - \frac{F(p) - F(x)}{f(x)} - c \right) \pi(x,c) dF(x) dH(c) \; + \\ + [1-F(p)] \left[\int_0^p (p-c)\pi(p,c) dH(c) \right] - [1-F(p)] \int_0^p \int_{\underline{x}}^p \pi(x,c) dx dH(c)\end{aligned} \tag{A5}$$

The last term in (A5) is the expected surplus of the high buyers ($x \geq p$). It obviously depends on the allocation rule for all $x \leq p$, and after rewriting it as $\int_0^p \int_{\underline{x}}^p \pi(x,c) \left(\frac{1-F(p)}{f(x)} \right) dF(x) dH(c)$, we can incorporate it into the first row of (A5) to result in (Standard individual rationality arguments also imply $m(0)=0$):

$$\Pi(\pi) = E_c \left[\int_{\underline{x}}^p \left(x - \frac{1-F(x)}{f(x)} - c \right) \pi(x,c) dF(x) + [1-F(p)](p-c)\pi(p,c) \right] \tag{A6}$$

Equation (A6) implies the retailer profit is *as if* all high customers paid p and all low customers delivered the same profit they would in the absence of the posted-price competitor. In other words, the surplus of high-value buyers implied by the IC constraint affects the payments of low-value buyers exactly as it would in the absence of the posted-price competitor.

The optimal allocation rule is obvious, and it maximizes the expected profit pointwise:

$$\pi(x,c) = 1 \Leftrightarrow \left(x < p \text{ and } c < x - \frac{1-F(x)}{f(x)} \right) \text{ or } x = p \tag{A7}$$

To see the optimality of always selling to high-value buyers with $x=p$, note that although the term $\pi(p,c)$appears in (A6) twice, its impact on profits inside the integral is measure zero, whereas its impact on profits in the $[1-F(p)](p-c)\pi(p,c)$ term has positive measure.

To derive the c-contingent profit shown in the proposition, use integration by parts to show that $[1-F(p)](p-c) = \int_p^{\bar{x}} (\psi(x) - c) dF(x)$, and verify that the right-hand expression results when we plug A7 into the expression in the large square brackets in A6. *QED Prop* 1

Proof of Lemma 1 Because $q(x) = U'(x), \beta'(x) = \frac{U(x)U''(x)}{\left(U'(x)\right)^2} > 0 \Leftrightarrow U''(x) > 0 \Leftrightarrow q'(x) > 0$. From Proposition 1, $q'(x) = \psi'(x) h(\psi(x)) > 0 \Leftrightarrow \psi'(x) > 0$. To derive the optimal bidding function, plug q and U from Proposition 1$q(x) = \int_0^p \pi(x,c) dH(c) = H(\psi(x))$ and $U(x)$

$= \int_{\underline{x}}^{x} q(t)dt = \int_{\psi^{-1}(0)}^{x} H(\psi(t))dt$ into Eq. 2 $:\beta(x) = \int_{\psi^{-1}(0)}^{x} t\frac{dH(\psi(t))}{H(\psi(x))} = \int_{0}^{\psi(x)} \psi^{-1}(c)\frac{dH(c)}{H(\psi(x))}$ $= E\left[\psi^{-1}(\text{cost})\middle|\psi^{-1}(\text{cost}) < x\right]$, where the second equality follows from a change in variables $c=\psi(z)$. *QED Lemma* 1

Proof of Lemma 2 To be incentive compatible, $\beta(p)$must satisfy, for every $x<p$:

$$\begin{array}{ll} p-\beta(p)\geq A(\beta(x))(p-\beta(x)) & (IC1) \\ U(x)\geq x-\beta(p) & (IC2) \end{array}$$

where the first inequality($IC1$) ensures type p bids $\beta(p)$ and the second inequality ($IC2$)ensures types $x<p$ do not deviate to $\beta(p)$. The deviation surplus for type p is

$$A(\beta(x))(p-\beta(x)) = H(\psi(x))\left(p-\int_0^{\psi(x)} \psi^{-1}(w)\frac{h(w)}{H(\psi(x))}dw\right) = \int_0^{\psi(x)} \left(p-\psi^{-1}(w)\right)dH(w)$$

which is obviously increasing in x, so the best deviation from bidding p is to bid $\bar{\beta_p}$. Therefore, ($IC1$) reduces to $\beta(p)\leq p-\int_0^{\psi(p)} \left(p-\psi^{-1}(w)\right)dH(w) = p-\int_{\psi^{-1}(0)}^{p} H(\psi(z))dz$. The LHS of ($IC2$)is the expected equilibrium surplus of type x, which Proposition 1 pins down as $U(x) = \int_{\psi^{-1}(0)}^{x} H(\psi(t))dt$. Therefore, ($IC2$)is $\beta(p)\geq p-\int_{\psi^{-1}(0)}^{p} H(\psi(z))dz$ because $x-\int_{\psi^{-1}(0)}^{x} H(\psi(z))dz$ is increasing in x. Therefore, the two incentive-compatibility constraints together uniquely determine the bid of type p as$\beta(p) = p-\int_{\psi^{-1}(0)}^{p} H(\psi(z))dz$. It is easy to show that the invertibility constraint $\beta(p)>\bar{\beta_p}$is always satisfied. *QED Lemma* 2

Proof of Lemma 3 The upper bound on probability acceptance must simply satisfy, for every x,

$A(b)(x-b)\leq H(\psi(x))(x-\beta(x))$

The RHS is maximized by $x=p$, so using the same arguments as in the proof of Lemma 2, the constraint is thus $A(b)\leq\left(\frac{1}{p-b}\right)\int_{\psi^{-1}(0)}^{p} H(\psi(z))dz$ *QED Lemma*3

Proof of Proposition 4 The expected profit of a retailer who uses a participation fee such that buyers with $x\geq x_e$ enter solves the following problem:

$$x_e = \underset{z}{\text{argmax}}\,[1-F(z)]\underbrace{H(\beta_0(z))(z-\beta_0(z))}_{S(z)\equiv \text{ expected surplus of} z} + \int_z^1 \pi_0(\min(x,p))dF(x)$$

where $\pi_0(x) = \int_0^{\beta_0(x)} (\beta_0(x)-c)dH(c) = H(\beta_0(x)) \, [\beta_0(x)-E(c|c < \beta_0(x))]$is the expected profit from a buyer with valuation x. The envelope theorem implies $S'(z)=H(\beta_0(z))$, so the FOC of the retailer's problem is $[1-F(x_e)]H(\beta_0(x_e))=f(x_e)[S(x_e)+\pi_0(x_e)]$.

In words, raising the fee increases the payment of everyone above the entry threshold (LHS) while decreasing the number of entrants, which results in a marginal loss of the fee and bidding profit (RHS). In other words, the RHS is the marginal loss of the gains from trade, because the $H(\beta_0(x_e))\beta_0(x_e)$ cancels out in $S(x_e)+\pi_0(x_e)$:

$[1-F(x_e)]H(\beta_0(x_e))=f(x_e)H(\beta_0(x_e))[x_e-E(c|c<\beta_0(x_e))]$

Both sides of the FOC are weighted by $H(\beta_0(x_e))$because no trade occurs when $c>\beta_0(x_e)$. Canceling out the weights results in the final FOC equation: $\psi(x_e)=E(c|c <\beta_0(x_e))$. *QED*

Proof of Proposition 5 For $p<2/3$, the relative profit advantage of the first-best strategy over the minimum-bid strategy is $\Pi_A(p)-\Pi_m(p) = \frac{(2p-1)^3}{12p} \to 0$ as p approaches $\frac{1}{2}$.

Claim a): Let $\Delta(p)=\Pi^*_m(p)-\Pi^*_e(p)$. There are three regions of p: First, for $p < \frac{4}{7}$, $\Delta(p) = \frac{p(1-p)}{8} > 0$ and obviously decreasing in p. Second, for $\frac{4}{7} < p < \frac{2}{3}$, $\Delta(p)$ is decreasing on the interval because $\Delta'\left(\frac{4}{7}\right) = -\frac{1}{56} < 0$ and $\Delta''(p) = -\frac{5}{6}-\frac{8}{147p^3} < 0$, and $\Delta(p)>0$ because and $\Delta\left(\frac{2}{3}\right) = \frac{127}{5292} > 0$ and $\Delta(p)$ is decreasing. Third, for $p > \frac{2}{3}$, $\Delta(p)$ is decreasing on the interval because $\Delta'(1) = -\frac{55}{10584} < 0$ and $\Delta''(p) = \frac{1}{6}+\frac{124}{1323p^3} > 0$, and $\Delta(p)>0$because $\Delta(1) = \frac{55}{10584} > 0$ and $\Delta(p)$ is decreasing.

Claim b): The equivalence for $p\leq\frac{2}{3}$ is obvious from both strategies charging p. The dominance for $p > \frac{2}{3}$ follows from $\Pi^*_m(p)-\Pi^*_{fix}(p) = \frac{2}{27p}-\left(\frac{1-E(c)}{2}\right)^2 = (8-3p)(3p-2)$ $\frac{2}{432p>0}$.

Claim c): Let $\Delta(p)=\Pi^*_{fix}(p)-\Pi^*_e(p)$. Consider $p > \frac{2}{3}$, where $\Delta(p) = \frac{1}{4}-\frac{4}{147p}-\frac{3p}{8}+\frac{7p^2}{48}$. We see that $\Delta(1) = -\frac{5}{784} < 0$ and $\Delta\left(\frac{2}{3}\right) = \frac{127}{5292} > 0$, so there must be at least one p^* s.t. $\Delta(p^*)=0$. To prove uniqueness, note $\Delta(p)$ is decreasing on the interval because $\Delta'(1) = -\frac{11}{196} < 0$ and $\Delta''(p) = \frac{7}{24}-\frac{8}{147p^3} > 0$.

Claim d): Given the ranking established in a)-c), it is enough to show the passive strategy is strictly less profitable than the participation-fee strategy for all p and the fixed-price strategy for$p > \frac{2}{3}$. The former follows from $\Pi^*_e(p)-\Pi_0(p) = \left\{ p < \frac{4}{7} : \frac{p(6-7p)}{24} > 0 p\geq\frac{4}{7} : \frac{4}{147p} > 0 \right.$. The latter follows from $\Pi^*_{fix}(p)-\Pi_0(p) = \frac{12-p(18-7p)}{48}$, which is decreasing for $p > \frac{2}{3}$ and positive at $p=1$. *QED Proposition 5.*

References

Amaldoss, & Jain. (2008). Joint bidding in the name-your-own-price channel: a strategic analysis. *Management Science, 54*(10), 1685–1699.

Anderson, C. (2009). Setting prices on priceline. *Interfaces, 39*(4), 307–315.

Anderson, C., & Wilson, J. G. (2011). Name-your-own price auction mechanisms – modeling and future implications. *Journal of Revenue and Pricing Management, 10*, 32–39.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science, 9*(3), 226–232.

Chen, C.-H. (2012). Name Your Own Price at Priceline.com: Strategic Bidding and Lockout Periods. *Review of Economic Studies*. published online 2012.

Chernev, A. (2003). Reverse pricing and online price elicitation strategies in consumer choice. *Journal of Consumer Psychology, 13*, 51–62.

Ding, M., Eliashberg, J., Huber, J., & Saini, R. (2005). Emotional bidders: an analytical and experimental examination of consumers' behavior in a priceline-like reverse auction. *Management Science, 51*(3), 352–364.

Fay, S. (2004). Partial-repeat-bidding in the name-your-own-price channel. *Marketing Science, 23*(3), 407–418.

Fay, S. (2008). Selling an opaque product through an intermediary: the case of disguising one's product. *Journal of Retailing, 84*(1), 59–75.

Fay, S. (2009). Competitive reasons for the name-your-own-price channel. *Marketing Letters, 20*, 277–293.

Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica, 68*(3), 525–574.

Hann, I.-H., & Terwiesch, C. (2003). Measuring the frictional costs of online transactions: the case of a name-your-own-price channel. *Management Science, 49*(11), 1563–1579.

Krishna, V. (2002). *Auction theory*. London: Academic Press of Elsevier Science.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research, 6*(1), 58–73.

Riley, J., & Zeckhauser, R. (1983). Optimal selling strategies: when to haggle, when to hold firm. *Quarterly Journal of Economics, 98*(2), 267–289.

Shapiro, D. (2011). Profitability of name your own price mechanisms in the case of risk-averse buyers. *Marketing Science, 30*(2), 290–304.

Shapiro, D., & Shi, X. (2008). Market segmentation: the role of opaque travel agencies. *Journal of Economics & Management Strategy, 17*(4), 803–837.

Shapiro, D., & Zillante, A. (2009). Naming your own price mechanisms: revenue gain or drain? *Journal of Economic Behavior and Organization, 72*, 725–737.

Spann, M., & Tellis, G. J. (2006). Does the internet promote better consumer decisions? Case of name-your-own-price auctions. *Journal of Marketing, 70*(1), 65–78.

Spann, M., Zeithammer, R., & Haübl, G. (2010). Optimal reverse-pricing mechanisms. *Marketing Science, 29*(6), 1058–1070.

Spann, M., Haübl, G., Skiera, B., & Bernhardt, M. (2012). Bid-Elicitation Interfaces and bidding behavior in retail interactive pricing. *Journal of Retailing, 88*(1), 131–144.

Spann, M., Zeithammer, R., & Haübl, G. (2015). Erratum to "optimal reverse-pricing mechanisms". *Marketing Science, 34*(2), 297–299.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance, 16*(1), 8–37.

Wang, T., Gal-Or, E., & Chatterjee, R. (2009). The name-your-own-price channel in the travel industry. *Management Science, 55*(6), 968–979.