Name-Your-Own Price Auction Mechanisms – Modeling and Future Implications

Chris K. Anderson, Cornell University
John G. Wilson, University of Western Ontario

A popular method for selling excess inventory over the Internet is via a Name-Your-Own Price auction, where the bidder bids on an item and the seller immediately decides on whether or not to accept the bid. The analytical modeling of such auctions is still in its infancy. A number of papers have appeared over the last few years making various assumptions about buyers and sellers. The intent of this article is to carefully delineate the various assumptions and modeling approaches and, consequently, suggest avenues for further research.

Introduction

In an Name-Your-Own Price (NYOP) auction, a potential buyer places at bid for an item such as a hotel room, rental car and so on. If the seller judges in the bid to be acceptable, the buyer's credit card is immediately charged. An NYOP auction is for what are termed 'opaque' goods. Here the buyer is unaware of the brand of good before the auction is completed. For instance, a bid on Priceline could be for a three star hotel in a certain area at a certain price. The name of the hotel will not be revealed until the room has been purchased.

In analyzing such an auction, there are a number of factors that can be considered. Is the auction single or multiple bid? What are the assumptions regarding the amount that buyers are willing to pay? What strategy does the seller employ in deciding whether or not to accept a bid? What does the buyer know about the seller's strategy? What inventory does the seller have? Is the NYOP channel being considered in isolation or in conjunction with other channels?

Although the models and situations will differ from paper to paper, the assumptions are made regarding such issues as a customer's reserve prices may often be evaluated independent of the model. Some of the key facts and results from various approaches to NYOP auctions will be summarized. Much of the research on NYOP is motivated by Priceline's NYOP model as they are the first and largest NYOP seller. We will briefly outline how Priceline's NYOP mechanism functions as motivation for our later discussions. The literature does not directly address the Priceline mechanism. However, the various models provide insight into various aspects of NYOP auctions and modeling issues. We start by summarizing how customers' reservation prices are modeled. The retailers strategies (for example, using a threshold prices) and consumers knowledge of them is addressed and various strategies that a customer might follow are listed next. The section on models lists some of the results from the various approaches. Some preliminary work on comparing different models is then described. We finish with some general comments on future implications.

Priceline's NYOP

Priceline was first formed in 1998 during the glory days of the dot.com boom. Initially focused on selling airline seats they quickly expanded into other segments of the travel space as well as other products (mortgages and gas and others). Following the crash of the dot.com bubble, Priceline needed rationalize their product lines and focused on the travel. Today, hotels remain the core product of Priceline's NYOP model. Customers looking to acquire services through Priceline's NYOP mechanism submit bids or offer prices for service without knowing the actual service provider. Using hotels as an example, a customer submits a price for an overnight stay arriving 1st November for a three-star hotel in Midtown Manhattan. Priceline then determines if that offer can be met. If the offer is declined, the customer can alter their bid by changing an attribute, changing the three-star to a four-star or changing Midtown to Times Square or changing the stay date. Alternatively, the customer could rebid on the three-star in Midtown after 24 hours has lapsed. Priceline's method of determining if a bid is accepted is unique and greatly favors the service provider. Following a customer's bid, pc, Priceline creates a list of all n qualifying properties, for example all three-star or better hotels in Midtown Manhattan. Priceline then randomly selects one of these qualifying properties, each with an equally likely probability of being selected. Priceline then checks the prices that the hotel has loaded to Priceline (all prices are loaded into the Worldspan global distribution system). Properties are allowed to load multiple rates into Worldspan. If the selected property has a price, p, lower than the consumer's bid then a transaction occurs. Generally speaking the hotel receives p the customer pays pc and Priceline keeps p_c-p as its share (Priceline has set minimum margins, that is, p_c-p>minimum margin). Under most circumstances if the property has more than one room rate less than the customer's bid the property receives the highest price less than pc. For example a customer bids $100 and the selected property has loaded rates of $100, $90 and $70, the customer would pay $100, the property would receive $90 and Priceline would keep $10.

If the property selected in this 'first round' does not have a price $p<p_c$ then Priceline goes back to the n-1 remaining properties and again randomly selects a property and checks prices. In this 'second' round each property does not have an equal chance of being selected. Instead, their probability of being selected depends upon their success rate in the first round, that is, the percentage of time when randomly selected that they have a rate less than the customers bid. A property that when selected in the first round has a rate $p<p_c$ 50 per cent of the time would have twice as likely a chance of being randomly selected in the second round as one which has a rate low enough 25 per cent of the time. This success rate or batting average is only calculated based on the first round, not on subsequent rounds. This second round is repeated until either a property is found which meets the customer's bid or until all properties are exhausted and no sale occurs.

The Priceline property selection and bid matching mechanism greatly favors the property as it selects the highest price that yields Priceline a profit. The random nature of property selection does not require the properties to compete on price with each other only to 'compete' with the customer as one firm's price relative to another does not impact their probability of being selected (at least in the first round).

Customer Reserve Prices

Each buyer is assumed to have a reservation price, the maximum price at which the buyer will consider buying an item from the opaque channel. Fay (2004) models this reserve price as the value of a uniform random variable over an interval [a, b]. Cai et al (2009) suppose that a bid of x costs a buyer $x+\theta$, where θ is uniformly distributed over an interval [c, d] and represents the disutility to the customer of bidding on the opaque channel. Wang et al (2009) assume that the reservation price for a customer is uniformly distributed over the normalized interval [0, 1]. Almadoss and Jain (2008) consider the case where a basket of items may be bid on either singly or individually and assume that valuations are uniform over (0, 1) and additive over items. There are some papers that do not require a uniform reservation price. Terwiesch et al (2005) allow for multiple bids with a transaction cost, which may be different for different customers, associated with making a bid. They assume a discrete model where there are a fixed number of possible reserve price-transaction cost combinations, each with a known fixed probability. Wilson and Zhang (2008) allow a general density over a range (0, H), where H represents the lowest price of the item on a non-opaque channel. Fay and Laran (2009) look at modeling from an individual customer's point of view, assume that multiple bids are allowed but that the value to a customer is discounted by a factor d^t, where $d<1$ and t is the number of the successful bid. Anderson and Xie (2010) also model consumers with valuations as uniform over (0, 1) in a setting where firms optimally set prices at a full information retailer, a posted price opaque retailer as well as an NYOP retailer. Based on their valuations, posted prices and the probability of bids being accepted (also uniform) consumers optimally choose their purchase channel or sequence of channels (in the case of failed bids).

Retailer's Policy and Customer's Beliefs

The most common policy in the literature is for the retailer to have a threshold above which bids will be accepted. Terwiesch et al (2005), Hann and Terwiesch (2003) and Wang et al (2009) assume that a customer's prior belief for the threshold is uniformly distributed over an interval. Fay (2004) assumes that the retailer announces two thresholds with the lower one applying to the case where there will be a limited supply of items and the higher threshold applies to the case where there are more items available. The customer is assumed to know the probabilities of these two eventualities. Fay and Laran (2009) allow multiple bids, assume that the threshold is uniformly distributed over an interval and that there is a fixed know probability that, in any given future period, the threshold will be redrawn from the same interval. In Almadoss and Jain (2008) the threshold depends on the (random) costs to the retailer, the equilibrium value of the game to the retailer and the probability that a new mid will materialize after the current one. Here, the customer also knows the probabilities for the threshold values. Terwiesch et al (2005) assume that there is a fixed number of customer classes each with known probability of occurring. Customers from each class assume that the threshold is drawn from a uniform distribution over an interval that can depend on the customer class. Anderson and Xie (2010) assume that consumers believe the threshold is uniform drawn from a distribution that is a function of posted prices and a regular and opaque posted price retailer. Cai et al (2009) and Wilson and Zhang (2008) assume that the probability that a bid of x will be accepted is known to the customer but need not be from a specific distribution such as the uniform.

Customer's Strategy

In Cai et al (2009), a customer who has a disutility of θ (the 'penalty' involved in bidding on the NYOP channel) will choose a value to minimize the expected cost

$$(x+\theta)\,F(x) + B(1-F(x))$$

where p(x) is the uniform distribution will be accepted and B is the price of the item if the bid fails and the customer buys the item at the regular price.

Wilson and Zhang (2008) assume the customer will choose x to minimize

$$xp(x) + H(1 - p(x))$$

where H is the regular price of the item and p(x) is the probability that a bid of x will be accepted. A difference here is that the retailer knows that the customers will use the above approach and thus chooses a strategy that will produce a p(x) that works in the retailer's favor.

In Hann and Terwiesch (2003), customers may bid multiple times but have transaction costs associated with each bid. The goal is to choose bids to maximize expected (bit not discounted) value over the sequence of bids using a dynamic programming approach.

In Fay (2004), the customer knows the thresholds used by the retailer (which correspond to different inventory amounts) and the probability that any threshold will be in play will be at one of these price thresholds, in particular the one that maximizes the expected surplus.

Fay and Laran (2009) allow multiple bids. They formulate a model where over a fixed sequence of bids that customers goal is to maximize expected discounted value. The complexity of the analysis is such that they focus on a number of particular betting patterns.

Anderson and Xie (2010) allow consumers to optimally choose between a full information channel with posted prices, an opaque channel with posted prices and an opaque NYOP channel. As a function of opacity, consumer valuations are discounted on the opaque channels with consumers choosing the channel sequence which maximizes their surplus. Using a model of consumers choosing to maximize their surplus Anderson and Xie illustrate how opaque selling maximizes firm revenues while naturally segmenting consumers.

Wang et al (2009) suppose that customer i chooses a bid B_i to maximize $(dv\text{-}B_i)\Pr[B_i \geqslant \textit{Threshold}]$, where d represents a discount to the customer's value v that reflects the inconvenience due to the opacity of the site. Note in this formulation that unlike Cai et al (2009),Wilson and Zhang (2008) and Anderson and Xie (2010), for instance, the customer is not assumed to buy on a posted price site if a bid on the opaque site is not successful.

Models

Almadoss and Jain (2008) look at the case where there is a package of items – for example, a hotel and a flight. The retailer's costs are random with a known distribution and customers reserve prices are uniform over (0,1). When there is a non-zero probability that another bid will be received, they show that requiring a bid on both items together is better than allowing a customer to decide on whether to bid on them jointly or separately.

Cai et al (2009) consider a multiple channel environment with both constrained and unconstrained capacity. They also consider the situation where a second bid is allowed. The customer is assumed to increase the first bid x by Δ. The probability of this second bid is assumed to be given by Δ/(B-x).

In Wilson and Zhang (2008), the retailer assumes that the customers are strategic (see Anderson and Wilson, 2003 and Shen and Su, 2007) and plan accordingly. Specifically, the retailer can provide a p(x), the probability a bid of x will be accepted, that will induce customers to bid close to their reserve prices, thus maximizing the retailers profits.

Fay (2004) considers the issue of multiple bidding. Under the assumption of uniform reservation prices and that the retailer will announce threshold values that correspond to inventory levels, he shows that the expected retailer's profit is the same when all customers are restricted to one bid as it is when they are allowed to rebid. However, the retailers profits can suffer if a single bid is assumed but customers get around the system by bidding a second time.

Fay and Laran (2009) formulate a model where a customer may bid multiple times, where the value of winning a bid is discounted by factor d^t, where $d<1$ and t is the number of the successful bid. The customer's knowledge of the bid is that it is uniformly distributed over an interval and may be reset by drawing from the same distribution with a fixed probability for each bid.

In Terwiesch et al (2005), customers have transaction costs associated with each bid and uniform distributions over the value of the threshold that the retailer is using. Under these conditions, they derive the optimal number of offers and associated values for a customer with given characteristics. For the case of homogenous customers (same values of transaction costs and same distribution for threshold value), they determine the optimal threshold price for the retailer. As the retailer can, to some extent, influence the transaction cost for the consumer this can lead to some interesting design issues. Hann and Terwiesch (2003) also consider this model. Fay (2004) also addresses the issue of transaction cost and shoes that in the absence of consumers who break the one bid rule, requiring a single bid auction can be optimal when transaction costs are not too large.

Wang et al (2009) assume that there are two channels — a NYOP price channel and posted price channel. Early demand from consumers on the posted price channel is used to inform the retailer on decisions for the NYOP channel. Under this conditions and assuming that customer valuations are uniform, optimal threshold prices are derived and Bayesian Nash Equilibrium strategies investigated.

Comparison of Models

Little work has been done on comparing models. Ogonowska et al (2009) take first steps in comparing a system such as Hotwire's with one such as Priceline's. They use an illustrative example to try and compare. Hann and Terwiesch (2003) compare a fixed posted price with a fixed threshold price auction, where in the latter a customer is allowed multiple bids but with a transaction cost associated with each bid. Wilson (2010) looks at the case where the reserve price distribution had a piecewise linear density function. For this case, he compares the optimal threshold policy with the policy of providing customers with a function p(x), the probability a bid of x will be accepted, that induces them to bid close to their reserve prices. Anderson and Xie (2010) is the first paper that looks at firm trying to set optimal prices (both full information and opaque) simultaneously with optimal threshold policies at the NYOP retailer. They illustrate that it is always optimal to post both full information and opaque prices with the opaque prices converging to full information prices as a function of the opacity of the opaque channel. Whereas NYOP selling is not always optimal with the firm converging to a two channel strategy under certain circumstances.

Implications for Revenue Management Further Research

Revenue management (RM) fundamentally hinges upon a firm's ability to segment consumers; early forms of RM simply used purchase timing to segment consumers with decision rules focused on when to curtail discounted demand. Fences or restrictions (Saturday night stay over being the most effective at separating leisure from business demand) were added over time as a means to simultaneously sell at different prices. As the internet reduced costs for consumer shopping, and reduced information asymmetry, many active users of RM became dynamic pricers increasing the commoditization of many services and encouraging consumers to be more strategic in purchase (timing) decisions. Travel demand post 9/11 and the recent recession dramatically illustrates some of the shortfalls of traditional RM as often demand is considerably less than capacity with firms reducing prices but not stimulating demand. Many airlines reduced route frequencies, moving aircraft capacity out of markets, allowing traditional RM approaches to function adequately. Conversely hotel capacity remained in markets (often operated by banks post foreclosure) with firms reducing prices with little occupancy impact as competitors responded.

As much of the research shows, NYOP and opaque selling naturally segment price sensitive (brand agnostic) consumers from brand loyal (price inelastic) consumers providing a efficient mechanism for sellers to simultaneously sell at multiple prices to segment consumers – the heart of RM. Much of the research to this point has focused on stylized approaches to modeling consumer behavior without general modeling frameworks.

As noted by Anderson and Wilson (2003), retailers need to be increasingly aware that their profits can be adversely affected by customers who are aware of the RM models being used and act strategically. Wilson and Zhang (2008) found a way to maximize the retailer's profit even when customers are provided with full information regarding the model, thus negating any negative effect from strategic behavior on the part of bidders. Almadoss and Jain (2008) 'examined the retailer's equilibrium bid acceptance strategy in the presence of strategic bidders'. This effect has been investigated in the general RM literature. Elmaghraby et al (2008) look at markdown mechanisms in the presence of strategic buyers. Su (2007) looks at the effect of customers who are strategic and will wait for discounts. Wilson et al (2006) look at the effect on booking limits of strategic customer behavior (See Shen and Su (2007), for a review of the relevant literature). The Internet has empowered customers. Those who wish to can find strategic information on various sites. For instance, biddingfortravel.com provides information on successful bids on Priceline. A recent paper by Anderson (2009) elicited a lively discussion biddingfortravel.com[1] on how to make bidding decisions strategically. Given that, in the future, even naïve users of the Internet can access websites of varying degrees of sophistication that will provide them with advice, it would seem important that future research should explicitly incorporate this strategic customer behavior.

Strategic consumer behavior arises from repeat bidding in NYOP selling. Multiple bidding has received little attention with conflicting results. Cai et al (2009) assume prob of first bid x being accepted is F(x), the uniform distribution and probability of second bid of $x+\Delta$ equals $\Delta/(B-x)$. Under this scenario: a double bid scenario can be attractive to customers different from Fay's (2004) earlier results under more restrictive conditions. More recently, Fay and Laran (2009) formulate an N period problem where bids are made at each period. However, given the computational complexity they focus on a restricted set of bidding patterns.

With very few exceptions (for example, Terwiesch et al, 2005 and Wilson and Zhang, 2008) most of the analytical modeling for NYOP options makes heavy use of uniformity assumptions. In particular, customers' reserve prices and opinions of the retailer's threshold are generally assumed to be uniform random variables. As observed by Cai et al (2009) 'although the assumptions of uniform distribution functions is common in the literature, it is desirable to relax this assumption to obtain more insights'. In many ways, one might consider these assumptions to be the weakest part of the literature and, possibly, a fruitful area for future research. In general, for instance, it seems to be a very strong assumption to assume that reservation prices are uniform. For instance, consider a population of made up of students, professors and businesspeople bidding for hotel rooms on an NYOP channel. In this case, a tri- modal distribution would seem most appropriate. Should this be the case, the literature on NYOP channels is not particularly informative. As pointed out by Wilson and Sorochuk (2009) assumption of a uniform distribution function can impose considerable limitations on the type of population being modeled.

The issue of dynamically changing NYOP auctions presents some particularly challenging technical problems. Fay and Laran (2009) take some first steps in allowing the threshold to change over time.

Generally, the analytical models assume that customers are rational in an economics sense. For instance, the assumption is often made that customers seek to maximize expected profit or maximize their surplus. A number of studies indicate that customers do not necessarily behave in the manner assumed by these models. For instance, Spann and Tellis (2006) make the comment 'Our results provide doubts about the assumption of strict rationality even in the context of the Internet'. This indicates that, in developing models, it can make sense to assume a wide variety of behavior among consumers.

Much of the NYOP literature has failed to incorporate competition and the resulting opacity into the optimal use of opaque selling. NYOP selling allows a firm producing a branded product to simultaneously sell it as a commodity to price sensitive consumers. Different opaque mechanisms offer differing levels of opacity with Hotwire.com informing consumers of specific details about their pending opaque purchase (list of amenities as well as user feedback), whereas Priceline chooses to offer neither. Lastminute.com's Top Secret opaque posted selling of hotels is non-binding as it allows consumers to cancel (as required by European law) their hotel room once they find out the 'opaque' service provider.

NYOP, opaque selling and auctions in general will continue to evolve within a RM context. The recent success of all pay auctions like Swoopo.com indicates how online selling and auction mechanism are a natural match. This is further evidenced by newer offerings, such as PriceWhispers.com, where consumers are NYOP for branded products with only uncertainty around fulfillment and not around good/service provider.

References

Anderson, C.K. (2009) Setting prices on priceline. *Interfaces* 39(4): 307–215.

Anderson, C.K. and Wilson, J.G. (2003) Wait or buy? The strategic consumer: Pricing and profit implications. *Journal of the Operational Research Society* 54(3): 299–306.

Anderson, C.K. and Xie, X. (2010) A Choice Based Dynamic Programming Approach for Setting Opaque Prices. Cornell University. Working Paper.

Almadoss, W. and Jain, S. (2008) Joint bidding in the name-your- own-price channel: A strategic analysis. *Management Science* 54(10): 1685–1699.

Cai, G., Chao, X. and Li, J. (2009) Optimal reserve prices in name-your-own-price auctions with bidding and channel options. *Production and Operations Management* 18(6): 653–671.

Elmaghraby, W., Gulcu, A. and Keskinocak, P. (2008) Optimal markdown mechanisms in the presence of rational customers with multi-unit demands. *Manufacturing and Service Operations Management* 10: 126–148.

Fay, S. (2004) Partial-repeat-bidding in the name-your-own-price channel. *Marketing Science* 23(3): 207–418.

Fay, S. and Laran, J. (2009) Implications of expected changes in the seller's threshold price in name-your-own-price auctions. *Management Science* 55(11): 1783–1796.

Hann, I. and 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.

Ogonowska, M., Longhi, C. and Torre, D. (2009) *Is There a Best Opaque Booking System for an Online Travel Agency?* University Nice Sophia-Antipolis. Working paper.

Shen, Z.J.M. and Su, X. (2007) Customer behavior modeling in revenue management and auctions: A review and new research directions. *Production and Operations Management* 16(6): 713–728.

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

Su, X. (2007) Intertemporal pricing with strategic customer behavior. *Management Science* 53(5): 726–741.

Terwiesch, C., Savin, S. and Hann, I. (2005) Online haggling at a name-your-own-price retailer: Theory and application. *Management Science* 51(3): 339–351.

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

Wilson, J.G. (2010) Optimal Threshold Versus Probabilistic Auctions for a Name-Your-Own-Price Channel. Ivey School of Business. Working paper.

Wilson, J.G., Anderson, C.K. and Kim, S.W. (2006) Optimal booking limits in the presence of strategic consumer behavior. *International Transactions in Operational Research* 13(2): 99–110.

Wilson, J.G. and Sorochuk, C. (2009) Optimal pricing and inventory decisions when individual customer preferences are explicitly modeled. International *Journal of Revenue Management* 3(2): 119–132.

Wilson, J.G. and Zhang, G. (2008) Optimal design of a name-your-own price channel. *Journal of Revenue Pricing Management* 7(3): 281–290.

---

[1] http://biddingfortravel.yuku.com/topic/98782/t/The-Curtain-is-Parted-More-or-Less.html.