POMS PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT POMS

Vol. 0, No. 0, xxxx–xxxx 2015, pp. 1–18
ISSN 1059-1478 | EISSN 1937-5956 | 15 | 00 | 0001
DOI 10.1111/poms.12412

Strategic Consumer Cooperation in a Name-Your-Own-Price Channel

Tatsiana Levina, Yuri Levin, Jeff McGill, Mikhail Nediak
School of Business, Queen's University, Kingston, Ontario K7L3N6, Canada, tlevin@business.queensu.ca, ylevin@business.queensu.ca, jmcgill@business.queensu.ca, mnediak@business.queensu.ca

Supplier reluctance to openly advertise highly discounted products on the Internet has stimulated development of "opaque" name-your-own-price sales channels. Unfortunately (for suppliers), there is significant potential for online consumers to exploit these channels through collaboration in social networks. In this paper, we study three possible forms of consumer collaboration: exchange of bid result information, coordinated bidding, and coordinated bidding with risk pooling. We propose an egalitarian total utility maximizing mechanism for coordination and risk pooling in a bidding club and describe characteristics of consumers for whom participation in the club makes sense. We show that, in the absence of risk pooling, a plausible bidding club strategy using just information exchange gives almost the same benefits to consumers as coordinated bidding. In contrast, coordinated bidding with risk pooling can lead to significantly increased benefits for consumers. The benefits of risk pooling are highest for consumers with a low tolerance to risk. We also demonstrate that suppliers that actively adjust for such strategic consumer behavior can reduce the impact on their businesses and, under some circumstances, even increase revenues.

Key words: strategic consumers; consumer cooperation; social networks; value of information; coordinated bidding
History: Received: April 2013; Accepted: April 2015 by William Cooper, after 2 revisions.

---

1. Introduction

The growth of the Internet has created unprecedented opportunities for firms to use revenue management and dynamic pricing techniques to obtain revenue from inventories that might otherwise go unsold. Unfortunately, the features of the Internet that enable rapid price changes also encourage consumers to behave strategically in their search for bargains, and this creates significant risks for suppliers. One key concern is cannibalization of demand—loyal customers who are willing to pay regular rates may divert to discount prices if they are visible on a supplier's web site or in the online tracking services. For example, an airline may wish to dramatically lower fares on an undersold flight as the departure date approaches, but this is the same period during which they anticipate receiving full fares from late-booking business travelers. Along with direct revenue dilution effects, there is the danger of brand erosion; for example, perceived "cheapening" of a brand image—a particular concern for hotel chains.

One answer to this predicament is *opaque* sales channels such as *Priceline.com* and *Hotwire.com* through which firms can offer discounted goods or services in generic categories without revealing the supplier identity until after a purchase is completed. For example, an opaque channel may offer hotel bookings by city, a range of availability dates, and "star-rating" but not the identities of the hotels; i.e. a "bundle" of similar products. Such markets can help firms dispose of excess inventory while minimizing the dangers. Some services, like Hotwire, function with posted prices while others, like Priceline, also allow consumers to place bids for their desired product or service—a so-called *name-your-own-price* (NYOP) channel.

It is crucial for the success of opaque NYOP product markets that the market operator does not reveal the inventory status of suppliers and the prices they are willing to accept, or set up too simple a pricing scheme. For example, if the operator maintains a single threshold price for any extended period, a consumer or group of consumers can quickly determine the threshold by sequential sampling. In practice, market operators hamper bid sampling by making accepted bids non-revocable and by requiring a delay after each bid during which an individual cannot submit a bid for an identical product bundle (24 hours in the case of Priceline). In addition, through one mechanism or another, market operators effectively randomize prices. (For readers interested in more details, the method used by Priceline to randomly sample from available hotel rooms and determine successful bids is described in Anderson (2009)). We believe that any viable NYOP market operator needs to

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

"scramble" the information reaching consumers, whether by a mechanism similar to Priceline's (which has been patented) or through some other form of randomized strategy.

Whatever pricing and acceptance mechanism is used, the key aspect of a NYOP market from the consumer viewpoint is the probabilistic nature of success or failure in bidding. As different suppliers add or remove availabilities and suppliers are randomly sampled, the prices at which bids are successful can be extremely volatile. Bids at any particular price can be successful for some customers and unsuccessful for others within a short time span. Higher bids are more likely to succeed, but there may still be successes and failures at any discount price level.

Time constraints and limited resources restrict the number of opportunities for consumers to gain information if they work alone, but social networking offers the possibility of exploiting NYOP channels. Such networks can range from informal information-sharing among users of "chat" sites, through webpages offering commentary and advice on bidding strategies, to fully developed web services like *Bidlesstravel.com* that orchestrate information exchanges among consumers about successful and unsuccessful bids. (Other examples are BiddingForTravel.com and BetterBidding.com.) There is also real potential for computer-guided bidding; for example, *Biddingtraveler.com* experimented for a time with automated Priceline bidding.

This raises several questions. Can collaboration among consumers genuinely improve biding outcomes in spite of the barriers against strategic bidding? How should collaboration be implemented? If collaboration is worthwhile, is simple exchange of bid history information sufficient or should more coordinated approaches be taken? Is it beneficial to "pool" risks by redistribution of benefits among participants. Providing answers to these questions is the objective of this study.

We take the approach of designing a robust mechanism that satisfies four key requirements: first, the underlying model must avoid complex assumptions about pricing distributions or paths; second, there should be some way to incorporate historical data or prior information; third, it must be simple for consumers to submit the information required; and fourth, it must be feasible for the operator of a cooperative bidding club to compute a bidding strategy that is easy to convey to consumers and prove optimal within the limitations of the information.

The mechanisms proposed in this study are based on a variant of Bayesian updating of success probabilities at a limited number of bid prices, guided by a dynamic program. We require only that the probabilities of success at the given prices are relatively stable over the period of time that the club is bidding for a particular bundle. Beyond this, no specific distributional assumptions are required, and there is no stipulation of any detailed stochastic process for prices. This simple approach has the merit that improved outcomes for consumers can serve as a lower bound on the benefits that might be achieved by more elaborate methods. Our subsequent conclusions about the relative merits of different levels of collaboration do not necessarily extend to other methods but are, nonetheless, suggestive of outcomes that might be observed.

The study focuses on cooperation among strategic consumers in NYOP product markets. Such consumers seek bidding strategies that balance the benefit of a low purchase price against the risk of failing to acquire the product (supply risk). A consumer's chance of getting a low price is enhanced if they have the opportunity to learn bid success probabilities from other bidders. We call the informal information exchange forums considered in this study *strategic bidder clubs* to distinguish them from organized *bidding rings* widely studied in the economics literature on auctions. (See, for example, the discussion on bidding rings in Talluri and van Ryzin (2004) and the references therein.) In our view, bidding clubs may be much less formal and must cope with much more ill-structured auction settings than conventional bidding rings.

We assess the potential benefits of three types of cooperation in bidder clubs, defined as follows:

Information Exchange: A group of consumers agree upon a range of discrete prices from which they will choose their bids but otherwise bid independently. They share the size of their bid and their subsequent success or failure in obtaining the product. This process essentially bypasses the restriction on multiple simultaneous samples set by the opaque channel operator and should improve estimates of probabilities of success at different prices.

Bid Coordination: Consumers exchange information, as above, but each consumer also commits to make their bid at a price assigned to them by the club. This allows optimizing the price sampling strategy but introduces the risk for individual consumers of missing an acquisition or paying more than others in the club because their assigned price was high.

Risk Pooling: Consumers coordinate their bids, as above, but also agree to sharing of benefits or costs after the full process has been completed.

Simple information exchange is already widespread in the online services mentioned above. We are not aware of current services that permit organized bidding by a club, but this can occur through informal contacts in social networks. An even more likely possibility is an automated bidding service similar to those in existing online auctions. Such automated coordinated bidding has the potential to increase member utilities through more efficient learning of successful bid probabilities. Risk pooling could improve this form of cooperation by introducing a payment system that permits consumers to redistribute some of the surplus in a pre-agreed fashion to increase their average utilities. Although the main advantage of risk pooling is unrelated to improvements in information or bidding strategies, as it simply transfers surplus between consumers to spread it more equally, the three aspects can be used synergistically.

We propose an egalitarian cooperation mechanism for bid coordination and risk pooling that is feasible for a group of consumers who are sufficiently homogeneous with regard to individual rationality (IR). The proposed mechanism can be extended to more heterogeneous populations but with a loss of its straightforward egalitarian properties.

We show that a plausible bidding club strategy using just information exchange results in almost the same benefits to consumers as coordinated bidding if risk pooling is not available. In contrast, when risk pooling is available in the form of bidding outcome-contingent transfer payments, consumers can significantly increase their benefits from cooperation. The increase in benefits of risk pooling for consumers with a low tolerance to risk can be quite dramatic. We also demonstrate how, at the cost of certain restrictions on bidding policies, the proposed mechanisms can also be used in large-scale settings.

The proposed benchmark setting of individual bidding provides a technical contribution to the literature on strategic experimentation (see, e.g. Bolton and Harris 1999) by addressing a problem of learning with ordered success probabilities until the first successful action or the number of attempts is exhausted.

2. Literature Review

Our work is related to a general area of research on strategic consumer behavior, which can be described as rational planning of purchases over time. A prime example of such behavior is strategic delay of a purchase in anticipation of lower prices. Behavior of this type was analyzed in the seminal work of Coase (1972), and surveys of subsequent work on strategic consumers can be found in Shen and Su (2007), Aviv et al. (2009) and Aviv and Vulcano (2010). The research in this area typically does not consider cooperation between consumers, but instead focuses on the implications of strategic consumer behavior *for the firms* and the various ways to mitigate it.

In our view, strategic behavior is particularly relevant when firms employ opaque products and NYOP sales channels, and there is a growing literature dealing with these topics because they are an important recent phenomenon in many industries. For general background, we refer the reader to the following papers and the references therein: Jerath et al. (2009, 2010), Jiang (2007), Fay and Xie (2008), Jiang (2007), Shapiro and Shi (2008), Fay (2004, 2008), Amaldoss and Jain (2008), Wang et al. (2009), Huang et al. (2014).

NYOP channels for opaque products of the type considered in our study are discussed by Anderson (2009) and Fay (2009). The first article analyzes the mechanism used by PriceLine to determine if the bid is successful and the probabilistic inventory release policy of a specific product provider. Fay (2009) examines whether a firm can use the NYOP format to soften price competition with its posted-price rival in a market where consumers are differentiated in their costs of using the NYOP channel and finds that this is indeed the case. The author emphasizes that it is beneficial for the firm to maintain the consumers' impression of randomness in the NYOP bid outcome. Our study shows that, even in such a randomized environment, consumers can improve their chances of successful bids through independent or cooperative learning.

Fay and Laran (2009) study (theoretically and empirically) a general model of consumer behavior which considers variations in the threshold price in the course of consumer bidding. Given a large number of competitors selling their products through firms like PriceLine, the assumption of a variable threshold is arguably more realistic than the fixed one, and Fay and Laran (2009) justify the idea that consumers expect to see a new threshold each day. When the threshold price changes vary rarely or very frequently, bids are increasing over time, and the same pattern also emerges for individual bidders in our study. This result is intuitive, since our model corresponds to that of very frequent threshold price changes.

A number of papers primarily focus on consumer behavior and learning in NYOP/opaque product markets. Hann and Terwiesch (2003) describe an empirical study to measure consumer costs of repeated bidding (frictional costs) in a setting where consumers are learning a fixed unknown threshold price set by the NYOP retailer. The measurement is based on a normative dynamic bid planning model of individual consumers. While the authors find that frictional costs in the NYOP electronic market may be substantial, one of their findings is that consumer experience reduces the costs. This suggests that the availability of consumer experience and bid results

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

from data sources like BidLessTravel.com can reduce frictional costs. Our study extends this work to the more common case of random threshold prices, allows for sharing of information, and generalizes reduced friction costs to improved expected utility for consumers.

Spann et al. (2004) present a normative model of consumer bidding behavior and derive a closed form solution under the assumption of a fixed threshold price. The model is used to simultaneously estimate willingness-to-pay and frictional costs of consumers. Related studies of Hann and Terwiesch (2003) and Terwiesch et al. (2005) also calculate the optimal bidding strategy, given a fixed set of expectations about the acceptance probabilities. The current article extends these models to account for how new information about the acceptance probabilities alters one's bidding strategy.

In the papers mentioned at this point, there is no cooperation between the consumers. Such cooperation in secret reserve price auctions (such as eBay's Best Offer and other NYOP auctions) is considered by Hinz and Spann (2008, 2010). These works consider a seller setting a *fixed* secret reservation price before bidding begins, and consumers who are not forward-looking but can use information about past bid results obtained from a social network. The model of the market in Hinz and Spann (2010) is more general than that of Hinz and Spann (2008) and incorporates a possibility that information about some past bids may be false. Moreover, Hinz and Spann (2010) present a decision support system that can aid the seller in setting the optimal reserve price and analyzing the impact of different social network structures. Both papers present empirical studies of bidder behavior in the laboratory and real-world conditions and evaluate the effects of the network structure on bid values. Unlike Hinz and Spann (2008, 2010), we consider forward-looking bidders who rationally anticipate the consequences of their own and others' bids. Moreover, the setting is more challenging for the consumers since the effective reservation prices in the market are sufficiently volatile that they appear as random up to a discrete probability distribution over several ranges. This randomness in bid success is driven by the presence of multiple capacity providers and a randomization of bid acceptance by the market operator. Indeed, even eBay's Best Offer allows for potential variation in offer acceptance over time by an auctioneer. Consumers seek to rationally maximize their utility through various levels of cooperation in bidding. Bid coordination and risk pooling constitute the next logical step in cooperation after information exchange, and we evaluate their theoretical feasibility by explicitly addressing the issue of rationality of consumer participation. We also discover the most beneficial settings for consumers in terms of their attitude to risk, time horizon and prior information. Stylized mechanisms described in this study can be used as a foundation for scalable bidding cooperation tools.

3. Model

We focus on the interplay of consumer learning and cooperation by considering a particular opaque bundle or a non-opaque product with an uncertain threshold price and a group of consumers who are interested in acquiring the product within a common search time. The group under consideration is small compared to the full population of potential consumers and does not exert a significant market power. This aspect permits us to avoid in explicitly modeling the market operator since the impact of the group is limited and it is unlikely that the operator would be able to distinguish the actions of the group from the rest of the market. The model captures three of the important aspects of market operation and consumer behavior discussed above: first, a bidding process with uncertain success probabilities at different prices; second, strategic consumers who have a limited number of opportunities to bid; and third, cooperative behavior of the consumers.

The bidding process is described in section 3.1, and strategic consumers, whose individual bidding strategies in the absence of information exchange serve as a benchmark, are described in section 3.2. Several levels of cooperation are considered. In the first, the consumers have access to the bidding history of others but do not coordinate their bidding strategies. In the second, bidding strategies are coordinated with the objective of maximizing the expected total utility of club participants. Finally, we discuss a deeper level of coordinated bidding in which consumers redistribute benefits at the end of the bidding process. For all three levels, we show that, under stated assumptions, it is rational for consumers to participate in cooperation. All technical proofs are in Online Appendix.

3.1. Bidding Process

We assume that each consumer can submit a limited number T of bids to the market operator. Bids are submitted sequentially, and each bid is processed independently from other bids (including previous bids of the same consumer). These assumptions are reasonable considering the total size of the market, the negligible impact of individual consumers, and the potential for consumer backlash if there was evidence of price manipulation against individuals. Based on PriceLine's 24 hour waiting period, we can think about bid submission times as days. For convenience, we index bids in the reverse order by $t = T, T-1, \ldots, 1$.

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

Bid independence also applies to bids submitted within one day by consumers from our group of interest—a strategic bidders club whose members may exchange and combine bidding information. Since, in existing practice, bidding and information exchange are done in different forums, we make the minimal assumption that information exchanges occur at the end of the day after the results of bids are known to all consumers in the club. Thus, we do not model the dynamics of bids during the day. In general, our assumptions are directed toward establishing lower bounds for the benefit of collaboration; thus, for example, information exchange during the day is an increase in information and, in an optimal bidding strategy, cannot lead to lower expected utility than simple end of day exchange. Deeper levels of cooperation can be implemented via automated bidding (such as "bidding robots" discussed later), but, using a lower bounding perspective, we still assume that results of the bids placed on a particular day are known only at the beginning of the next day.

As discussed in the introduction, in the unlikely event that the market operator maintains a constant "threshold" price over a consumer's search time, approximating that price would be a simple matter of starting a stepwise search from a low starting bid and increasing by increments on successive days until a bid is successful. The increments in the search process and the accuracy of the approximation will be determined by the number of bidding periods the consumer is willing to undertake. If several consumers agree to collaborate, the increments can be much smaller, and a fixed threshold price can be determined quite quickly and accurately.

Unfortunately, consumers are unable to know whether the price is constant or slowly moving during their search time. It is far more likely that the market operator (as illustrated by Priceline) will be accepting or rejecting bids dynamically based on a complex of factors like remaining booking time, fluctuating inventories of rooms from multiple sources across many accommodation classes, and demand forecasts. From the point of view of the consumer, the minimum acceptable price for a particular room class is an unobservable, discontinuous, and volatile stochastic process—a moving target. There is really no reason for the market operator to maintain a single threshold price. In fact, the operator can control inventories by controlling the "rates" or probabilities of sales at two or more prices. A fixed-price policy is a special case of a randomized policy and, hence, optimal randomized policies cannot do worse than fixed policies, and they are an effective way for market operators to 'scramble' the information reaching consumers.

Much of the uncertainty confronting consumers is captured by a model with several given price levels in which the *probabilities* of success at these prices are constant. Such a situation would correspond to a secret threshold or reserve price that is sufficiently volatile and can easily transit across multiple levels in either direction in any given day, or if the operator is employing randomized pricing with relatively fixed randomization probabilities. In the latter case, it is not necessary that the operator be randomizing over the same candidate prices agreed upon by consumers, only that the probabilities of acceptance or rejection of bids at the consumer's prices be relatively constant.

The design of an effective search process in such a setting is an interesting research question, but the focus is on the benefits that might or might not accrue to consumers in a bidding club rather than the search process itself. Accordingly, we construct a plausible search method that respects the volatility in prices while remaining simple to implement for consumers via an automated system.

We assume that consumers can agree on a discrete set of discount prices $\mathcal{P} = \{p_1, \ldots, p_q\}$, where $p_1 > \ldots > p_q$, that may be obtained in the opaque market and a full price, $p_0 \geq p_1$, obtainable by direct purchase. The full price could be a representative or average posted price usually obtainable from price-shopper web sites and other public sources. If the full price is not an option, we set $p_0 = \infty$ with an understanding that consumers can forgo the purchase altogether if they cannot secure the product at a discount price. Individual consumers bid prices in $\mathcal{P}$ and, given the 'black-box' nature of the market, it is possible that there will be both successful and unsuccessful bids at any of these prices on the same day. We approximate this condition by assuming that the outcomes of different bids are statistically independent. We thus assume that, for each p_j, there is a fixed probability θ_j the bid at p_j is successful, and that the probabilities are non-decreasing with p_j, i.e. $\boldsymbol{\theta} \in \Theta = \{\boldsymbol{\theta} \in \mathbb{R}^q : 0 \leq \theta_q \leq \cdots \leq \theta_1 \leq 1\}$. These probabilities are unknown—the only observations available to consumers will be the frequencies of successful bids at the different prices over time, which will allow progressively more precise estimates of θ_j if they are constant during the search time. This assumption of constant θ_j is justified in typical revenue management settings where marginal values of products are driven by the opportunity costs. The NYOP channels are often used to dispose of "distressed inventory" that effectively has zero opportunity cost. If the main driver of product prices and bidding success rates is the information "scramble" by the opaque channel operator, the prices and success rates are likely to be stationary over time.

Generally, there is a finite number N_T of bidders who rationally plan their bidding strategies. Each consumer's strategy is a sequence of bids at prices in the price set $\mathcal{P}$ that continues until a bid is successful or the search time finishes. Any consumers who are successful drop from the process. At time $t \in \{T, \ldots, 1\}$, there are N_t consumers still bidding. Consumers' bidding strategies vary with their willingness-to-pay (WTP), attitude to risk, and amount of knowledge about the market. We begin by analyzing a reference case of a single independent strategic consumer.

3.2. Strategic Consumers as Independent Bidders

Consumer WTP and attitudes to risk can be captured by *utility*w_j of purchase at p_j, $j = 1, \ldots, q$. Without loss of generality, the highest price p_1 in $\mathcal{P}$ is below consumer WTP (specifically, $p_1 < WTP$), and it is natural to assume that w_j, $j = 1, \ldots, q$ are finite and $0 \le w_0 < w_1 < \cdots < w_q$ (where $0 \le w_0$ because consumers can simply forgo the purchase if $p_0 > WTP$). Since there are no other restrictions on w_j's, this model can capture arbitrary consumer preferences for lower prices in the context of his attitude to risk. It is possible that the bundle of opaque products contain some with higher valuations than others. Without loss of generality, the w_j's can be interpreted as expected utilities across possible valuations for products included in successful bids. The objective of the consumer is to maximize the expected utility of purchase. This problem can be viewed as general learning where rewards are discrete random variables with supports $\{0, w_j\}$. The learning proceeds until the first non-zero outcome or all attempts are exhausted. This structure is reminiscent of the *multi-arm bandit problem* where the task is to select one of the available actions with uncertain outcomes (typically, in the absence of distributional information). However, the fundamental differences are as follows:

- It is *impossible to exploit* previously selected actions once bidding success is achieved.
- The exploration is restricted to observing failures for past action in the presence of the *ordered* structure of success probability set Θ.

The asset selling problem with posted prices and unknown success probabilities has the same form.

The expected utility and the chosen course of action depend on beliefs of consumers about the market. With the prices in the market simplified to discrete price levels, such beliefs can be represented as a prior distribution $\pi(\boldsymbol{\theta})$ on the bid success probabilities $\boldsymbol{\theta} \in \Theta$. This set-up is reminiscent of sequential estimation of the parameters of several different Bernoulli distributions, which could be performed with independent Bayesian updating processes using beta/binomial conjugate distributions. However, since the bid success events are *nested* and, consequently, the success probabilities are *ordered*, the estimation must be done jointly. The multinomial generalization of the binomial cannot be used either because bid successes at different prices are not mutually exclusive. Nonetheless, Bayesian updating can be accomplished with a generalization of the Beta-binomial estimation procedure.

Consider the following generalization of a Beta distribution to multiple dimensions:

$$\pi(\boldsymbol{\theta}|\mathbf{a}, \mathbf{b}) = C \prod_{j=1}^{q} \theta_j^{a_j - 1} (1 - \theta_j)^{b_j - 1}, \quad \boldsymbol{\theta} \in \Theta, \qquad (1)$$

where C is a normalizing constant and $\mathbf{a} = (a_1, \ldots, a_q)$, $\mathbf{b} = (b_1, \ldots, b_q)$ form a parameter vector $\mathbf{x} = (\mathbf{a}, \mathbf{b})$ with strictly positive elements. Suppose there have been m_j successful and n_j failed bids at price p_j, $j = 1, \ldots, q$ and let $\mathbf{m} = (m_1, \ldots, m_q)$, $\mathbf{n} = (n_1, \ldots, n_q)$. If no bid is made at price p_j, then $m_j + n_j = 0$. In Online Appendix S1, we show that the posterior distribution has the same form (1) but with the updated parameter vector $(\mathbf{a} + \mathbf{m}, \mathbf{b} + \mathbf{n})$. The normalizing constant of Equation (1) plays a major role in our model. In particular, $\frac{1}{C} = I(\mathbf{a}, \mathbf{b})$ where

$$I(\mathbf{a}, \mathbf{b}) = \int_{\boldsymbol{\theta} \in \Theta} \prod_{j=1}^{q} \theta_j^{a_j - 1} (1 - \theta_j)^{b_j - 1} d\theta_j \qquad (2)$$

is a generalization of the Beta function to multiple dimensions. Using Equation (2), the expected value of θ_j and $1 - \theta_j$ with respect to distribution (1) is $\frac{I(\mathbf{a}\mathbf{e}_j, \mathbf{b})}{I(\mathbf{a}, \mathbf{b})}$ and $\frac{I(\mathbf{a}\mathbf{b} + \mathbf{e}_j)}{I(\mathbf{a}, \mathbf{b})}$, respectively (here, $\mathbf{e}_j$ is the jth coordinate vector, $j = 1, \ldots, q$).

For individual consumers who do not participate in information exchange, the problem of finding the optimal bidding policy is formally a finite-horizon Markov decision process (MDP). Suppose at the beginning of day t a consumer with utilities $\mathbf{w} = (w_0, w_1, \ldots, w_q)$ has a belief vector $\mathbf{x} = (\mathbf{a}, \mathbf{b})$ which already incorporates any past bidding history. The vector $\mathbf{x}$ can be thought of as the state in the MDP. Let $U_t(\mathbf{w}, \mathbf{x})$ be the value function of this decision process—the optimal expected utility of the consumer with utilities $\mathbf{w}$ at time t whose current state of beliefs is $\mathbf{x}$. The utilities in $\mathbf{w}$ are fixed in some of the following discussions but can vary across consumers in others. For consistency, we retain $\mathbf{w}$ as an explicit argument throughout. The consumer has q options—bid any $p \in \mathcal{P}$, and two possible outcomes for each bid—acceptance or rejection by the operator of the market. Thus, the value function and the optimal

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

policy are found by solving the following dynamic programming recursion:

$$U_t(\mathbf{w},\mathbf{a},\mathbf{b}) = \max_{j=1,\ldots,q}\left\{\frac{I(\mathbf{a}+\mathbf{e}_j,\mathbf{b})}{I(\mathbf{a},\mathbf{b})}w_j + \frac{I(\mathbf{a},\mathbf{b}+\mathbf{e}_j)}{I(\mathbf{a},\mathbf{b})}U_{t-1}(\mathbf{w},\mathbf{a},\mathbf{b}+\mathbf{e}_j)\right\}, \quad (3)$$

with the terminal condition $U_0(\mathbf{w},\mathbf{a},\mathbf{b}) = w_0$. By letting $u_t(\mathbf{w},\mathbf{a},\mathbf{b}) = I(\mathbf{a},\mathbf{b})U_t(\mathbf{w},\mathbf{a},\mathbf{b})$, we can write recursion (3) more compactly as

$$u_t(\mathbf{w},\mathbf{a},\mathbf{b}) = \max_{j=1,\ldots,q}\{I(\mathbf{a}+\mathbf{e}_j,\mathbf{b})w_j + u_{t-1}(\mathbf{w},\mathbf{a},\mathbf{b}+\mathbf{e}_j)\} \quad (4)$$

with the terminal condition $u_0(\mathbf{w},\mathbf{a},\mathbf{b}) = I(\mathbf{a},\mathbf{b})w_0$. We use $p_t^*(\mathbf{w},\mathbf{a},\mathbf{b})$ to denote the optimal bidding strategy of a consumer with utilities **w** and information vector (**a**,**b**) at time t. (We use the convention that the ties in the maximum are broken in favor of more likely success, that is, a higher price.) A practical implementation of the optimal bidding strategy can be accomplished using a bidding robot. Bidding robots are used widely in online auctions (e.g., http://www.bidderrobot.com) and greatly simplify bid result tracking and new bid submission by the consumers. In our case, automation would permit solution of the dynamic programming recursion (4) for consumers.

An immediate observation from recursion (4) is that the optimal policy is of the open-loop form, that is, deterministic. Indeed, bid failure at price p_j advances the state to $(\mathbf{a},\mathbf{b}+\mathbf{e}_j)$ in a predetermined fashion, and if the consumer still bids at time t, he must have suffered bid failures at times $T, \ldots, t+1$. Moreover, bidding strategies of individual consumers possess the following monotonic property (recall that the time index runs from T down to 0):

PROPOSITION 1. *Consider* $t \geq 2$. *If* $p_t^*(\mathbf{w},\mathbf{a},\mathbf{b}) = p_j$, *then* $p_{t-1}^*(\mathbf{w},\mathbf{a},\mathbf{b}+\mathbf{e}_j) \geq p_j$.

The optimal policy is also deterministic if the actual success probabilities are known. Thus, this proposition implies that the difference between optimal strategies under uncertainty vs. those when success probabilities are known is only in the switching times to progressively higher prices. The optimal bidding strategy under uncertainty will be identical to the optimal strategy with known probabilities whenever they have the same switching times. Moreover, even if they do not have identical switching times, the two policies may differ only slightly. As a result, the strategies obtained under uncertainty may be very robust. It is also true that the value of information/coordination is *zero* if the consumer happens to be using the optimal known probability strategy when bidding individually. Thus, we may expect the information value to be low when the individual bidding strategy under uncertainty remains optimal for moderate shifts in the success probability distribution.

Dynamic programming recursion (3) effectively solves the asset selling problem with posted prices. The formulation and the associated algorithm can be extended to a general dynamic pricing and demand learning problem of perishable items with a discrete set of prices and an unknown Bernoulli demand process, but such generalization is beyond the scope of this study.

Next, we consider consumer cooperation by information exchange (without coordination).

3.3. Strategic Bidders Club with Uncoordinated Strategies

Since the bidders club does not have market power and there is no competition between bidders, it is impossible to decrease performance merely by participation in information exchange. Existence of online forums such as BidLessTravel shows that bidders are generally willing to provide information about past bid results. Thus, we can consider a consumer who has access to the bid results of others but, otherwise, bids independently. The expected value of consumer utility depends on the consumer's belief state and the future evolution of this state driven by the joint bidding behavior of the entire group. In the absence of coordination, bidding behavior of others is uncertain for an individual consumer who has to form expectations about it. The difficulty of forming rational expectations of others' behavior justifies the assumption of limited foresight on the part of individual consumers. Let ξ_j be the probability of other consumers bidding p_j, $j = 1, \ldots, q$. We assume that each consumer has evolving beliefs about ξ_j represented as a probability distribution, starts with a prior distribution on $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_q)$ and performs Bayesian updating based on new information. An appropriate prior distribution for uncertain probability $\boldsymbol{\xi}$ is Dirichlet with parameters $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_q)$.

In the information exchange context, consumers know the total number of consumers N_t who have not acquired the product at time t. From the individual point of view, the number of others k_j bidding p_j, $j = 1, \ldots, q$ form a multinomial random vector on $N_t - 1$ trials with category j probability ξ_j. This observation, along with a conjugate relation between multinomial and Dirichlet distributions, facilitates a dynamic programming solution for the problem of expected utility maximization by bidding under information exchange. The state description is now extended to $(\mathbf{w}, N_t, \mathbf{x}, \boldsymbol{\alpha})$. The value function $U_t^E(\cdot)$ of the dynamic program can be expressed via its weighted equivalent as $u_t^E(\mathbf{w}, N_t, \mathbf{x}, \boldsymbol{\alpha}) = I(\mathbf{x})B(\boldsymbol{\alpha})$ $U_t^E(\mathbf{w}, N_t, \mathbf{x}, \boldsymbol{\alpha})$ where $B(\alpha)$ is the normalizing constant of the Dirichlet distribution. We use the short-hand notation $|\mathbf{z}| = \sum_{j=1}^{q} z_j$ for any vector $\mathbf{z} \geq 0$.

Proposition 2. *Suppose that consumer beliefs about bid levels of other consumers follow a Dirichlet distribution with parameter vector* $\boldsymbol{\alpha}$. *Then the scaled maximum expected utility* $u_t^E(\cdot)$ *can be found through the dynamic program*

$$u_t^E(\mathbf{w}, N_t, \mathbf{a}, \mathbf{b}, \boldsymbol{\alpha}) = \max_{j=1,\ldots,q} \Bigg\{ I(\mathbf{a}+\mathbf{e}_j, \mathbf{b})B(\boldsymbol{\alpha})w_j + \sum_{\substack{|\mathbf{k}|=N_t-1 \\ \mathbf{k}\geq 0}} \sum_{\substack{\mathbf{m}+\mathbf{n}=\mathbf{k} \\ \mathbf{m},\mathbf{n}\geq 0}} \binom{N_t-1}{\mathbf{m},\mathbf{n}} u_{t-1}^E(\mathbf{w}, N_t - |\mathbf{m}|, \mathbf{a}+\mathbf{m}, \mathbf{b}+\mathbf{e}_j+\mathbf{n}, \boldsymbol{\alpha}+\mathbf{k}) \Bigg\} \quad (5)$$

with the boundary condition $u_0^E(\mathbf{w}, N_0, \mathbf{x}, \boldsymbol{\alpha}) = I(\mathbf{x})B(\boldsymbol{\alpha})$ w_0, *where the multinomial coefficient is denoted as* $\binom{N_t-1}{\mathbf{m},\mathbf{n}} = \frac{(N_t-1)!}{m_1!\ldots m_q! n_1!\ldots n_q!}$.

Outline of Proof. The proof follows if it can be shown that the recursion in Equation (5) leads to correct computation of expected utilities under the Dirichlet assumption. □

Consider a situation in which k_j consumers (other than a consumer under consideration) bid p_j and let $\mathbf{k} = (k_1, \ldots, k_q)$. The Dirichlet distribution with parameter vector $\boldsymbol{\alpha} > 0$ is conjugate with multinomial likelihood, thus the posterior distribution is Dirichlet with parameters $\boldsymbol{\alpha} + \mathbf{k} = (\alpha_1 + k_1, \ldots, \alpha_q + k_q)$. This provides a very compact representation of updates to consumer perceptions of $\boldsymbol{\xi}$. Moreover, the distribution of the number of bids at p_j can be obtained by simply taking the expected value of multinomial probabilities with respect to $\boldsymbol{\xi}$.

Consider now an arbitrary belief vector $\mathbf{x} = (\mathbf{a},\mathbf{b})$. A consumer bids p_j and observes either success or failure of his bid. If the bid is successful, which occurs with probability $\frac{I(\mathbf{a}+\mathbf{e}_j,\mathbf{b})}{I(\mathbf{a},\mathbf{b})}$, the consumer obtains the utility corresponding to the price. These elements are the same for consumers in the absence of information exchange.

An unsuccessful bid occurs with probability $\frac{I(\mathbf{a},\mathbf{b}+\mathbf{e}_j)}{I(\mathbf{a},\mathbf{b})}$ and results in updated belief vector $(\mathbf{a}, \mathbf{b} + \mathbf{e}_j)$. In this case, the consumer proceeds to observe the number of bidders k_j at p_j, $j = 1, \ldots, q$ and the successes and failures at these prices. Given $(\mathbf{a}, \mathbf{b} + \mathbf{e}_j)$, and $\mathbf{k}$, the probability of m_j accepted bids (and the corresponding number n_j of rejected bids) at price p_j, $j = 1, \ldots, q$ can be easily computed. This fully specifies the probability model of consumer observations and establishes that the consumer bidding strategy optimization problem can be solved using Equation (5).

The following proposition implies that a consumer manifesting the limited foresight behavior described in this section believes that the expected utility under honest information exchange exceeds that in the absence of information exchange.

Proposition 3. *For a customer with utilities* $\mathbf{w}$, *we have* $U_t^E(\mathbf{w}, N_t, \mathbf{x}, \boldsymbol{\alpha}) \geq U_t(\mathbf{w}, \mathbf{x})$ *for all* $t, N_t, \mathbf{x}, \boldsymbol{\alpha}$.

It remains to be seen, however, if the actual rather than perceived expected utility obtained under this type of behavior is indeed higher. We return to this question in section 4.

Remark 1. Intuition is not a reliable guide with respect to monotonic properties over information states that do not occur consecutively during the optimal bidding process. Consider, a special case of two price levels ($q = 2$) and information states $(\mathbf{a}, \mathbf{b} + \mathbf{e}_1)$ and $(\mathbf{a}, \mathbf{b} + \mathbf{e}_2)$. Compared to $(\mathbf{a},\mathbf{b})$, the first state represents additional information of a failed high-price bid, and the second, a failed low-price bid. The outlook associated with the first state is more "pessimistic" since high-price bids should have higher chance of winning. Therefore, it may seem reasonable that it cannot be optimal to bid a lower price in state $(\mathbf{a}, \mathbf{b} + \mathbf{e}_1)$ than in $(\mathbf{a}, \mathbf{b} + \mathbf{e}_2)$. However, there exists a counterexample to this statement (see Online Appendix S2 for details).

Remark 2. Obtaining additional sample information about the market constitutes a relaxation of the problem and cannot lead to worse average outcomes—consumers benefit in expected utility if they are provided with results of additional bids, *ceteris paribus*. A precise statement to that extent, along with a specific bounding expression, is provided in Online Appendix S2.

3.4. Strategic Bidders Club with Coordinated Strategies

We now describe a situation in which consumers form a "bidders club" by agreeing to cooperate not only on information exchange, but also to coordinate their bidding strategies. In practice, coordination is possible among consumers who know each other personally and develop a certain level of commitment to the process. Alternatively, the coordination

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

may take the form of a bidding robot that places bids for a group of consumers in a concerted fashion. A plausible implementation of the coordinated-bidding robot can ensure consumer commitment by collecting a credit card number at the time of registration for group bidding. Once consumers commit, they relinquish control and the robot submits bids on their behalf. Of course, this approach requires a mechanism to ensure fairness of the outcome for club members. If, under club's strategy, a consumer acquires the product at a high price or fails to purchase, he may feel treated unfairly compared to those who managed to secure the product at a lower price or may be suspicious about "rigged" results. Moreover, the club must provide a measurable gain from participation. The issue of fairness can be addressed if coordinated bidding is *egalitarian*, that is, it treats consumers equally in some sense. The second issue (that of a measurable gain) can be framed as ensuring *individual rationality* of bidders' participation in a coordinated bidding strategy. When the consumer population possesses a degree of homogeneity, it is relatively easy to ensure that the mechanism is egalitarian and satisfies IR. When the consumer population is not homogeneous, it may be more difficult to guarantee these properties by the bidding strategy alone. Therefore, we also consider a deeper level of cooperation where the prices paid by the consumers may be augmented by compensatory payments that depend on the combined outcome of all bidders. On one hand, these transfer payments can address the issue of fairness and, on the other, they represent a form of *risk pooling* that reduces the uncertainty faced by risk-averse consumers. Even if consumers do not acquire any additional information about the acceptance probabilities and do not change their bidding strategies, pooling risk would be beneficial indicating a fundamentally advantageous type of consumer cooperation. We demonstrate that the risk-pooling mechanism can be useful for a relatively homogeneous group of consumers in section 3.4.2 and consider heterogeneous groups in Online Appendix S3.

3.4.1. Coordinated Bidding with a Homogeneous Consumer Population. This section requires a closer examination of the utilities of potential club members. Here, we assume a more specific structure of consumer attitudes to risk: that their utility vectors are obtained from a common utility function of purchase surplus. That is, each consumer's utility of purchase at price p_j is $w(v - p_j)$ where $w(\cdot)$ is the common utility function and v is his WTP, which may vary between consumers. Since WTP is the maximum amount, a consumer is willing to pay for the product, we can treat any purchase option with negative surplus as extremely unattractive and assume that $w(v - p_j) = -\infty$ if $v - p_j < 0$. Thus, such options need not be included into $\mathcal{P}$ for a homogeneous population. We still assume that $w_0 \geq 0$ since the default option need not be taken.

The proposed bidding mechanism is *egalitarian* in the following sense. First, it considers a *representative* WTP level $\bar{v}$ and treats consumers identically by using $w_j = w(\bar{v} - p_j)$ in place of their actual utility vectors. Second, even though the club may bid different prices for different consumers on the same day, it also uses randomization to ensure equally likely assignments of consumers to price levels. This guarantees, a priori, that the randomized bidding strategy of the club is identical for every consumer. While the actual consumer utility vectors may differ from $\mathbf{w}$, we can view it as a *representative* utility vector of the consumer population. We also recognize that consumer WTP for opaque products is inherently uncertain and it is possible that it differs from $\bar{v}$ when the purchase takes place at price p_j. At the end of this subsection, we discuss the rationality of consumer participation in the mechanism and a possible treatment of uncertainty.

The evolution of the club belief vector is similar to that in the case of uncoordinated bidders discussed in section 3.3. However, it is helpful to endow it with a day-specific subscript: $\mathbf{x}_t$. The bids are coordinated, thus the fractions $\boldsymbol{\xi}$ of bidders at different prices and their posterior distribution parameters are not needed. Instead, we need to keep track of the individual outcomes of the bids up to day t which we call the *consumer state*. This consumer state is formed by the number of remaining consumers and by the total numbers of successes/failures at each price by time t: $\mathbf{s}_t = (N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t)$. Given the initial state at time T, the number of remaining consumers at time t can be computed as $N_t = N_T - |\hat{\mathbf{m}}_t|$ and the belief vector as $\mathbf{x}_t = \mathbf{x}_T + (\hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t)$. This redundancy in state representation can be exploited for numerical implementation.

The consumer state uniquely identifies the total utility for each final outcome. Let $U_t^C(\mathbf{w}, N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t)$ be the expected total utility at time t in state $(N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t)$ given utility vector $\mathbf{w}$ and $u_t^C(\mathbf{w}, N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t) = I(\mathbf{x}_t) U_t^C(\mathbf{w}, N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t)$ be its scaled version. A state $(N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t)$, becomes terminal either because time runs out (at time $t = 0$) or because all bidders are successful prior to time 0 ($N_t = 0$). In either case, the terminal value is given by

$$u_t^C(\mathbf{w}, N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t) = I(\mathbf{x}_t)\left(\sum_{j=1}^{q} \hat{m}_{jt} w_j + N_t w_0\right). \quad (6)$$

Given a particular state, the bidding strategy of the club in this state can be described by vector

$\mathbf{k} = (k_1, \ldots, k_q)$ satisfying $|\mathbf{k}| = N_t$ and $\mathbf{k} \geq 0$ that specifies the number of bids at each price level. A number of successes and failures of these bids are described by q-dimensional vectors $\mathbf{m}$ and $\mathbf{n}$, respectively. This is similar to the case of uncoordinated bidding. This strategy is applied by the bidding robot in a randomized fashion, so that the chance that the robot bids p_j for a particular bidder is $\frac{k_j}{N_t}$. One way to generate such random assignments of distinct bidders to bid levels is to sample the permutations of bidders at random and to assign them in the resulting order: k_1 bidders to bid p_1, k_2 bidders to bid p_2, etc.

PROPOSITION 4. *For N_T consumers with the same utility vector $\boldsymbol{w}$, the club's scaled maximum expected total utility $u_t^C(\cdot)$ under coordinated bidding can be found via the dynamic program*

$$u_t^C(\mathbf{w}, N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t) = \max_{\substack{|\mathbf{k}| = N_t \\ \mathbf{k} \geq 0}} \Bigg\{ \sum_{\substack{\mathbf{m}+\mathbf{n}=\mathbf{k} \\ \mathbf{m},\mathbf{n} \geq 0}} \left[\prod_{j=1}^{q} \binom{k_j}{m_j} \right] \times u_{t-1}^C\Big(\mathbf{w}, N_t - |\mathbf{m}|, \hat{\mathbf{m}}_t + \mathbf{m}, \hat{\mathbf{n}}_t + \mathbf{n}, \mathbf{x}_t + (\mathbf{m}, \mathbf{n})\Big) \Bigg\} \tag{7}$$

with the terminal condition (6). Moreover, under the optimal randomized policy, each bidder obtains the same expected value of $\frac{1}{N_T} U_T^C(\mathbf{w}, N_T, \mathbf{0}, \mathbf{0}, \mathbf{x}_T)$.

We now consider populations of consumers for whom it is rational to participate in coordinated bidding.

DEFINITION 1. We call the population homogeneous if either (a) all consumers have WTP equal to p_0, or (b) there is an open, non-empty, interval $(v', v'') \subseteq (p_1, p_0)$ such that all consumers have WTP in (v', v'').

Case (a) describes a group of consumers for whom the default purchase option at p_0 is tolerable. In this case, consumers are not willing to pay more than the price p_0 they can always secure outside the club and we can set $\bar{v} = p_0$. Case (b) describes a group of consumers who cannot afford the default purchase option but for whom all prices in $\mathcal{P}$ are feasible. The narrowness of the interval $v'' - v'$ is a measure of population homogeneity. For each $\bar{v} \in (p_1, p_0)$, we may ask whether it is possible that there is any sufficiently homogeneous population of consumers (described by $(v', v'') \ni \bar{v}$) who would rationally participate in the club using representative WTP level $\bar{v}$ and the resulting utility vector $\mathbf{w}$.

The following proposition shows that it is rational for groups of consumers whose utility vectors are identical to engage in coordinated bidding. This result applies to arbitrary utility vectors.

PROPOSITION 5. *For any T, N_T, $\mathbf{x}_T > 0$ and $\boldsymbol{w} \geq 0$,*

$$\frac{1}{N_T} U_T^C(\mathbf{w}, N_T, \mathbf{0}, \mathbf{0}, \mathbf{x}_T) \geq U_T(\mathbf{w}, \mathbf{x}_T). \tag{8}$$

The following corollary states the implications of Proposition 5 in terms of homogeneous consumer populations and shows that rationality applies to more diverse groups of consumers than those with identical $\mathbf{w}$:

COROLLARY 1. *If consumers are homogeneous of type (a), it is rational for them to commit to coordinated bidding in a club using representative WTP level $\bar{v} = p_0$. Moreover, if inequality (8) is strict for $\boldsymbol{w}$ resulting from representative WTP level $\bar{v} < p_0$ and the utility function $w(\cdot)$ is continuous, then there exists a range $(v', v'') \ni \bar{v}$ as in type (b) homogeneity such that it is rational for any consumer with WTP in (v', v'') to commit to coordinated bidding.*

The first claim of the corollary shows that, in reality, the formation of the club would be attractive for a group of consumers who find the default purchase option acceptable. For such a group, the WTP is likely to reset to the value of the default option resulting in identical utility vectors. The second claim shows that the formation of the club is also plausible for a group of consumers whose WTP levels are sufficiently close to $\bar{v} < p_0$. The condition $\bar{v} < p_0$ implies that all such consumers would not even consider the default option. In the numerical experiments, we see that the sufficient condition of strict inequality (8) is satisfied under a wide range of inputs. This guarantees IR for coordinated bidding in some range of WTP, and shows that the coordinated bidding mechanism is feasible. The commitment to coordinate is a form of self-selection. Further improvement in incentives for collaboration are possible if we consider risk pooling.

REMARK 3. Since the proof of the corollary relies only on the continuity of $w(\cdot)$, the model can be extended by including some degree of randomness in the valuations of consumers. The distribution of valuations may be conditional on the consumer and the price used in the successful bid. The key is that the support of the distribution is sufficiently narrow interval $(v', v'') \ni \bar{v}$. Such an extension covers real-world applications where consumers have varying valuations of the products included in the bundle at each price.

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

3.4.2. Risk Pooling with a Homogeneous Consumer Population. The randomization of prices assigned to different bidders under coordinated strategies ensures that the mechanism is egalitarian but it does so only in a probabilistic sense and does not prevent unfavorable outcomes for individual bidders. Risk pooling is a form of insurance that compensates bidders when unfavorable outcomes do occur. The compensation occurs at the expense of bidders with favorable outcomes so that the total balance of transfer payments is zero. This restriction is known as that of *balanced payments* in the mechanism design theory.

With an egalitarian mechanism and homogeneous consumers, the transfer payment for a particular consumer can depend only on the final outcome of the bidding process and the price paid by this consumer under automated bidding. Formally, the payments are functions of the $\hat{\mathbf{m}}_t$ component of the terminal state (other parts of the state are irrelevant because they do not affect the terminal value) and the price p_j paid in the absence of compensations. We denote the transfer payments by $y_j(\hat{\mathbf{m}}_t)$, $j = 0, 1, \ldots, q$, where $y_0(\hat{\mathbf{m}}_t)$ is the compensation paid to the consumer if all bids placed for him by the robot fail. The balance requirement takes the form

$$\sum_{j=1}^{q} \hat{m}_{jt} y_j(\hat{\mathbf{m}}_t) + \left(N_t - \sum_{j=1}^{q} \hat{m}_{jt}\right) y_0(\hat{\mathbf{m}}_t) = 0 \tag{9}$$

for all terminal $(t, \hat{\mathbf{m}}_t)$.

Recall that the pair $(t, \hat{\mathbf{m}}_t)$ is considered "terminal" if either $t = 0$ or $|\hat{\mathbf{m}}_t| = N_t$. We see that the $y_j(\hat{\mathbf{m}}_t)$ variables are only relevant to those j with $\hat{m}_{jt} > 0$ but this mild redundancy simplifies the notation. To state the optimization problem, instead of the utility vector $\mathbf{w}$, we use an explicit dependence of the club's utility on consumer WTP $\bar{v}$ and on the payments $\mathbf{y}(\cdot)$. Let the expected total utility of the club and its scaled version be $U_t^R(\bar{v}, \mathbf{y}(\cdot), N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t)$ and $u_t^R(\bar{v}, \mathbf{y}(\cdot), N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t) = I(\mathbf{x}_t) U_t^R(\bar{v}, \mathbf{y}(\cdot), N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t)$, respectively. Taking payments in the terminal state $(N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t)$ into account, the club's scaled total utility is

$$\begin{aligned} &u_t^R(\bar{v}, \mathbf{y}(\cdot), N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t) \\ &= I(\mathbf{x}_t)[\sum_{j=1}^{q} \hat{m}_{jt} w(\bar{v} + y_j(\hat{\mathbf{m}}_t) - p_j) \\ &+ \left(N_t - \sum_{j=1}^{q} \hat{m}_{jt}\right) w(y_0(\hat{\mathbf{m}}_t))], \end{aligned} \tag{10}$$

where $w(y_0(\hat{\mathbf{m}}_t))$ is the utility of the default option. (Recall that we have either $\bar{v} = p_0$ or $\bar{v} < p_0$. The latter case means that the default option will not be taken. Thus, in both situations, the resulting purchase surplus is zero.)

For *given* transfer payments, the problem of the club's expected total utility maximization can be solved by a dynamic program of the same form as Equation (7) but using Equation (10) as the terminal condition. However, the full problem is to maximize the club's expected utility over the bidding strategy *along with* the payments:

$$\max_{\mathbf{y}(\cdot)\text{s.t.}(9)} U_T^R(\bar{v}, \mathbf{y}(\cdot), N_T, \mathbf{0}, \mathbf{0}, \mathbf{x}_T). \tag{11}$$

Formulation (11) approaches the club's expected total utility maximization sequentially. Given the payments $\mathbf{y}(\cdot)$, it optimizes the bidding policy, and, assuming optimal bidding for each $\mathbf{y}(\cdot)$, it optimizes over $\mathbf{y}(\cdot)$. Equivalently, one can first optimize over $\mathbf{y}(\cdot)$ for each bidding strategy. With such a sequence, it is obvious that the total expected utility can be improved unless, in each terminal state, every consumer obtains the same utility. We claim the following:

Proposition 6. *Suppose that utility function $w(\cdot)$ is continuously differentiable and strictly concave for nonnegative surpluses. Then the optimal bidding strategy and transfer payments solving Equation (11) are obtained by*

$$y_j^*(\hat{\mathbf{m}}_t) = \frac{1}{N_T} \sum_{j'=1}^{q} \hat{m}_{j't}(\bar{v} - p_{j'}) - (\bar{v} - p_j), \quad j = 1, \ldots, q, \tag{12}$$

$$y_0^*(\hat{\mathbf{m}}_t) = \frac{1}{N_T} \sum_{j=1}^{q} \hat{m}_{jt}(\bar{v} - p_j), \tag{13}$$

and the dynamic program of the same form as Equation (7) using the terminal condition

$$\begin{aligned} &u_t^R(\bar{v}, \mathbf{y}^*(\cdot), N_t, \hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t, \mathbf{x}_t) \\ &= I(\mathbf{x}_t) N_T w\left(\frac{1}{N_T} \sum_{j=1}^{q} \hat{m}_{jt}(\bar{v} - p_j)\right) \text{ for all terminal } (t, \hat{\mathbf{m}}_t). \end{aligned}$$

The payments defined by Equations (12) and (13) achieve the goal of equalizing the net surpluses of all consumers in every terminal state. Exact solution of the utility maximization is possible because we can establish this simple structure. The structure of the optimal payments also implies that their introduction results in a strict improvement in the total expected utility.

The additional improvement in utility provided by risk pooling compared to coordinated bidding can be substantial. This reinforces the result of Corollary 1 and can extend the interval of consumers' WTP in which it is rational for them to commit to cooperation. Moreover, Remark 3 that introduces a possible extension of the model to mildly heterogeneous consumer valuations of the products included in the bundle also applies to the case of risk pooling.

4. Benefits of Information Exchange and Coordination

We now examine several combinations of model settings to address the following questions:

(1) What are the benefits of information exchange, coordination and risk pooling?
(2) How do these benefits depend on consumer attitude to risk, prior information, the number of price levels, the available time and the number of consumers?
(3) Can cooperation approaches be scaled to larger settings?

The expected values of bidding policies depend greatly on the relation between prior information about success probabilities and their actual values. While it makes sense to compare policies using the average over a prior on the success probabilities rather than any fixed values, one needs to appropriately calibrate the consumer and club uncertainty in success probability to the actual uncertainty. To provide fair and consistent comparisons of scenarios, we assume that individual consumers and the club share the same beliefs about success probabilities, and these beliefs match the actual market uncertainty. In technical terms, the actual market uncertainty corresponds to the club's belief vector $\mathbf{x}_T$. The main metric in our study is the *individual benefit* of consumers defined as the relative difference in the certainty equivalent (CE) of a purchase under, respectively, information exchange, coordination or risk pooling; and CE under independent bidding.

The individual benefits are examined under the following prior distributions. First, the uniform distribution corresponding to a belief vector with $\mathbf{a} = \mathbf{b} = \mathbf{e}$. While it may seem counterintuitive, this distribution does not model a complete lack of information. If we examine the interpretation of prior (1), we see that this belief vector matches the observations of exactly one success and failure at each price level. Therefore, we also examine another prior distribution. In Bayesian statistics, a distribution typically called *uninformative* describes a vanishingly small amount of information. In our case, the uninformative distribution would correspond to $\mathbf{a} = \mathbf{b} = \varepsilon\mathbf{e}$ where ε is near zero. This distribution adequately describes a nearly complete lack of prior information, but it has the drawback of permitting arbitrarily high values of the density as success probabilities approach one. This may be unrealistic as we would typically expect near-one success probabilities to be rather unlikely. Therefore, instead of the uninformative prior, we examine a "pessimistic" prior with $\mathbf{a} = \varepsilon\mathbf{e}$ and $\mathbf{b} = \mathbf{e}$ which is a hybrid between the uniform and the uninformative ones. This "pessimistic" density function assumes arbitrarily high values as success probabilities approach zero and is bounded for success probabilities near one. In our experiments, we use $\varepsilon = 0.05$. Other elements of the experimental setup include:

- We treat consumer population as homogeneous with representative $\bar{v} = 1$.
- The number of possible price levels q ranges from 2 to 4 and the prices are equally spaced on the interval from 0 to 1 (e.g., for $q = 2$, the prices are $p_1 = 2/3$ and $p_2 = 1/3$). The default purchase option is priced at $p_0 = 1$.
- Because of the rapid (exponential) growth in the state space size as T, N_T or q increase, we limit the number of periods T to 10, 8 and 6 for q equal to 2, 3 and 4, respectively. (These values of T are realistic given consumer "patience" issues.) The maximum number of consumers N_T is varied depending on the experiment (for coordinated bidding N_T is 10, 8 and 6 for q equal to 2, 3, and 4, respectively).
- Consumers value every monetary amount or purchase surplus z according to an exponential utility function $w(z) = 1 - \exp(-\frac{z}{\rho})$. In this model, the attitude to risk is summarized by single parameter ρ called the *risk tolerance*, which is easy to estimate in practice. We tested levels of ρ at 0.25, 1 and 4 describing, respectively, low, medium, and high tolerance to risk.

All figures including Figures S1–S6 discussed next are provided in the Online Appendix. The detailed graphs of the benefits (in %) are plotted as functions of the length of the time horizon T. Figures S1, S3 and S5 show these graphs for the uniform prior distribution and q = 2, 3, and 4, respectively. Similarly, Figures S2, S4, and S6 show the benefits for the pessimistic distribution. Each figure contains nine subplots for each combination of a cooperation procedure (information exchange, coordinated bidding, and risk pooling) and tolerance to risk.

4.1. Benefits of Information Exchange

In this set of experiments, we evaluate performance of the information exchange procedure discussed in section 3.3. The maximum number of bidders considered was $N_T = 8$, 6 and 4 for q = 2, 3, and 4, respectively. The prior on the bids of others is assumed to be uninformative Dirichlet with parameter vector $\boldsymbol{\alpha} = 0.05\mathbf{e}$. The bidding policy for each combination of experimental settings is computed using dynamic programming and then simulated 10,000,000 times. The average utility obtained in these simulations is used to estimate the CE of the bidding policy and individual benefits of information exchange for the consumers.

Overall, benefits are modest at levels below 4%. The highest benefits are obtained for the uniform prior, and we can see a clear increase in benefits when the number of bidders grow in this case. For the pessimistic prior, information exchange does not produce a reliable improvement over independent bidding except when the time horizon is relatively long, there is sufficient price flexibility (q = 3 or 4) and consumers are not too risk-averse. Generally, consumer tolerance to risk strongly affects the amount of time required to realize noticeable benefits from information exchange, but increased time does not always imply increased benefits. The latter observation can be explained by more effective independent bidding for large T.

4.2. Benefits of Bid Coordination and Risk Pooling

The behavior of the benefits is somewhat similar to those in the case of information exchange, but the observations are more accurate because they are based on the exact calculation rather than a Monte-Carlo simulation. Generally, we see an increase in benefits of coordination as the number of bidders increases. Under the uniform prior, additional benefits of bid coordination compared to information exchange are minimal. However, for the pessimistic prior, a significant boost in benefits appears with increased price flexibility. In particular, consumers with medium and high tolerance to risk seem to benefit from collaboration when the time horizon is relatively long. Compared to information exchange, it takes a shorter time horizon to obtain noticeable benefits. The reason is that success probabilities for low prices are quite small, but risk-tolerant consumers have enough patience to attempt low prices and to learn the corresponding chances of success, resulting in increased benefits for large T.

The biggest beneficiaries of risk pooling are consumers with a low tolerance to risk. The benefits are particularly high for short-time horizons and exceed 30% for the uniform prior and q = 2. Risk pooling also makes all the difference for the pessimistic prior case and short time horizons because, in its absence, the benefits of coordinated bidding are near zero. It is evident that out of the three methods of consumer cooperation, risk pooling provides the most robust incentives for cooperation.

Additional insight into the impact of coordination and risk pooling is obtained from the change in expected consumer surplus associated with cooperation. In general, coordination and risk polling permit consumers to bid more aggressively compared to individual bidding. Risk pooling is usually more effective in that respect but, for many experimental settings, the change in surplus resulting from coordination and risk pooling is the same showing that the actual optimal bidding strategies are identical and risk pooling only redistributes the resulting surplus. In a limited number of settings under the uniform prior, the change in surplus is negative, which implies that the club may prefer to trade some of the total surplus for a reduction in risk. The change in surplus also provides a quick assessment of the club's impact on the market operator and the suppliers because the increase in the surplus of the consumers matches the decrease in the revenues for the operator/suppliers. The maximum change in surplus resulting from coordination and risk pooling (as % of the expected surplus under independent bidding) is provided in Table 1. The impact on the market is limited for the pessimistic prior (below 5.3%) and the uniform prior (below 6%) unless consumers have a low tolerance to risk. If independent consumers have a low tolerance to risk, they bid quite conservatively, resulting in a relatively low expected surplus. As a result, coordination and risk pooling lead to significant changes in their bidding behavior by making it far more aggressive. With risk pooling and very risk-averse consumers, the maximum change for q = 2,3 and 4 is 10.3%, 13.8% and 26%, respectively. These maximum changes occur for T = 2 and large group sizes, and, generally, the changes in surplus are much smaller for $T > 2$.

An increase in the total surplus of consumers in a group reduces the total revenue from this group obtained by the sellers. On the other hand, the sellers offering a default option and those selling through the NYOP channel may be different. Thus, it is impor-

Table 1 The Maximum % Change in Consumer Surplus Due to Coordinated Bidding (C) and Risk Pooling (R)

		q = 2		q = 3		q = 4	
Prior	ρ	C (%)	R (%)	C (%)	R (%)	C (%)	R (%)
Pessimistic	0.25	0.5	0.7	0.8	2.0	0.5	4.0
	1	0.7	1.2	4.1	4.2	5.2	5.3
	4	1.4	1.4	4.2	4.2	4.3	4.3
Uniform	0.25	7.2	10.3	4.9	13.8	9.3	26.0
	1	3.3	3.5	5.4	6.0	2.3	2.6
	4	3.5	3.5	4.4	4.5	2.4	2.5

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412

tant to examine the impact of consumer cooperation on the revenue of the NYOP channel (i.e. expected revenue from sales at $p_1, \ldots, p_q$). Figures S7–S10 show the percentage change in NYOP channel revenue due to collaboration and risk pooling for $q = 3$ and $q = 4$ as compared to the NYOP channel revenue under independent bidding. The change is driven by two components: bid levels used by the consumers and the resulting NYOP sales (the percentage changes in the latter and the former are shown, respectively, in Figures S11–S14 and S15–S18). The combined effect may be positive or negative.

Under the pessimistic prior, which corresponds to a low chance of winning, we find that the change is negative in all experimental settings. The magnitude of the change is somewhat larger under the risk pooling and increases with group size and longer horizon (reaching values around −10%). This phenomenon results from lower bid levels (see Figures S16 and S18) combined with lower sales (see Figures S12 and S14) indicating that, under the pessimistic prior, coordinated bidding and risk pooling both push consumers to place lower and more risky bids. This is an unfavorable outcome both for the market operator and the suppliers who specifically want to dispose of distressed inventory.

Under the uniform prior, the change in revenue is positive in some moderate or high risk tolerance settings: very short time horizons for $q = 3$ (reaching values above 50%), and longer time horizons for $q = 4$ (reaching values above 10%). For low risk tolerance, coordinated bidding (but not risk pooling) also leads to increase in revenue for some group sizes and time horizons. The change was lower or more negative for risk pooling as well as for larger group size. Positive change is effected by both an increase in sales and an increase in the revenue per bid. This indicates consumer shift toward higher bids that are more likely to succeed under the influence of additional information in coordinated bidding. Risk pooling somewhat negates this tendency. The overall conclusion is that the NYOP channel operator and sellers who offer their products through this channel may benefit from consumer cooperation. Moreover, essentially the same bidding policy is optimal for $\bar{v} \to p_0 - 0$ (tending from below) which corresponds to a consumer group that would not buy at the default price. In this case, the observed increase in revenue is pure gain for capacity suppliers.

4.3. Scalable Approaches to Cooperation

Practical implementations of cooperation require scalable approaches that can address deviations from stylized assumptions of the base model. The policy optimality requirements may be relaxed, but the presented model still provides ideas for several heuristic algorithms. We start by constructing a heuristic policy for individual bidding.

According to Proposition 1, the optimal individual bidding policy is monotonic. Any such policy can be described either by a sequence of price level indices $j(t)$ such that $j(t) \geq j(t-1)$, $t = T, \ldots, 2$, or a vector $\mathbf{n} \in \mathbb{Z}_+^q$ where each component n_j specifies the number of times price p_j was used. The latter form stipulates to bid price p_q for n_q times first (i.e. $j(T) = \ldots = j(T - n_q + 1) = q$), then p_{q-1} for n_{q-1} times (i.e. $j(T - n_q) = \ldots = j(T - n_{q-1} + 1) = q - 1$), etc., until p_1 is used n_1 times. If $n_j = 0$ for some j, then p_j is not used.

Proposition 7. *The scaled expected value of a monotonic bidding policy in the form of* $\mathbf{n}$ *vector is*

$$u_T^{\mathbf{n}}(\mathbf{w}, \mathbf{a}, \mathbf{b}) = \sum_{j=1}^{q} w_j \left[I\left(\mathbf{a}, \mathbf{b} + \sum_{k=j+1}^{q} n_k \mathbf{e}_k\right) - I\left(\mathbf{a}, \mathbf{b} + \sum_{k=j}^{q} n_k \mathbf{e}_k\right)\right] + w_0 I\left(\mathbf{a}, \mathbf{b} + \sum_{k=1}^{q} n_k \mathbf{e}_k\right). \tag{15}$$

The closed-form expression (15) provides a recipe for a scalable individual bidding approach that serves as a benchmark for evaluating cooperative heuristic policies and a building block for their construction. In the dynamic programming formulation (4), the major source of complexity is evaluation of every possible future belief state potentially arising under arbitrary monotonically increasing bids. Thus, we propose to use a roll-out heuristic where the future periods (excluding the current one) are split into two subintervals of constant prices. Formally, suppose that the belief vector at time t is ($\mathbf{a}$, $\mathbf{b}$) and that the price value index at $t - 1$ is j_0 (if $t = T$, let $j_0 = q$). The heuristic is to bid $p_{j_1^*}$ where j_1^* is the first part of a triple (j_1, j_2, j_3) delivering the maximum in

$$\begin{aligned} \tilde{u}_t^I(\mathbf{w}, \mathbf{a}, \mathbf{b}, j_0) &= \max_{j_0 \geq j_1 \geq j_2 \geq j_3 \geq 1} \Big\{ w_{j_1} I(\mathbf{a}, \mathbf{b}) + (w_{j_2} - w_{j_1}) I(\mathbf{a}, \mathbf{b} + \mathbf{e}_{j_1}) \\ &+ (w_{j_3} - w_{j_2}) I\left(\mathbf{a}, \mathbf{b} + \mathbf{e}_{j_1} + \left\lceil \frac{t-1}{2} \right\rceil \mathbf{e}_{j_2}\right) \\ &+ (w_0 - w_{j_3}) I\left(\mathbf{a}, \mathbf{b} + \mathbf{e}_{j_1} + \left\lceil \frac{t-1}{2} \right\rceil \mathbf{e}_{j_2} + \left\lfloor \frac{t-1}{2} \right\rfloor \mathbf{e}_{j_3}\right) \Big\}. \end{aligned} \tag{16}$$

The resulting algorithm proceeds as follows:

(1) Initialize $t = T$, $j_0 = q$ and ($\mathbf{a}$,$\mathbf{b}$).

(2) If $t = 0$, take the default option with utility w_0 and stop. Otherwise, bid price with index j_1^* computed according to Equation (16).
(3) If successful, stop. Otherwise, update $\mathbf{b}$ to $\mathbf{b} + \mathbf{e}_{j_1^*}$, j_0 to j_1^*, t to $t-1$, and repeat step 2.

The (scaled) expected utility value $\tilde{u}_t^I(\mathbf{w}, \mathbf{a}, \mathbf{b}, q)$ is lower bound for the optimal expected utility of individual bidding starting from the belief vector ($\mathbf{a}$,$\mathbf{b}$) at time t. In the cooperative settings, this expression can also be used to estimate the future value if we assume that each consumer starts bidding independently from the next period. For information exchange setting, this bounding idea can be combined with the assumption that each consumer thinks that others behave exactly like him/her. As long as consumers have the same utilities, under the common belief state that is characteristic of information exchange, the assumption of identical behavior is self-fulfilling. Thus, tracking the probability $\boldsymbol{\xi}$ of others' bids at each price level is no longer needed. The heuristic information exchange policy is to bid j delivering the maximum in

$$\tilde{u}_t^E(\mathbf{w}, N_t, \mathbf{a}, \mathbf{b}) = \max_{j=1,\ldots,q} \Big\{ I(\mathbf{a} + \mathbf{e}_j, \mathbf{b}) w_j + \sum_{m=0}^{N_t-1} \binom{N_t - 1}{m} \tilde{u}_{t-1}^I(\mathbf{w}, \mathbf{a} + m\mathbf{e}_j, \mathbf{b} + (N_t - m)\mathbf{e}_j, q) \Big\}. \quad (17)$$

The summation is over possible number m of successful bids by others under the assumption that they always bid the same price as the bidder in question (i.e. p_j). If consumer's bid at p_j fails and there are m successes of others, the updated belief state is $(\mathbf{a} + m\mathbf{e}_j, \mathbf{b} + (N_t - m)\mathbf{e}_j)$ and $\tilde{u}_{t-1}^I(\mathbf{w}, \mathbf{a} + m\mathbf{e}_j, \mathbf{b} + (N_t - m)\mathbf{e}_j, q)$ is a (scaled) lower bound on the expected utility of the bidder who stopped participating in information exchange. The resulting approach is tolerant to departures of other bidders and their failures to report results beyond the current period since it uses an independent bidding lower bound for evaluating the future.

In the coordinated bidding setting, different bidders can generally be given different prices to bid. However, this leads to additional complexity (essentially, an additional nested loop for each distinct price level). Therefore, it makes sense to consider a coordinated bidding strategy constrained to a single price level for all bidders. Using the same lower bound, we obtain a coordinated bidding policy that prescribes all remaining bidders to bid p_j where j delivers the maximum in

$$\tilde{u}_t^C(\mathbf{w}, N_t, \mathbf{a}, \mathbf{b}) = \max_{j=1,\ldots,q} \Big\{ \sum_{m=0}^{N_t} \binom{N_t}{m} \Big[m w_j I(\mathbf{a} + m\mathbf{e}_j, \mathbf{b} + (N_t - m)\mathbf{e}_j) + (N_t - m)\tilde{u}_{t-1}^I(\mathbf{w}, \mathbf{a} + m\mathbf{e}_j, \mathbf{b} + (N_t - m)\mathbf{e}_j, q) \Big] \Big\}. \quad (18)$$

The focus on remaining bidders and using the bound $\tilde{u}_{t-1}^I(\mathbf{w}, \mathbf{a} + m\mathbf{e}_j, \mathbf{b} + (N_t - m)\mathbf{e}_j, q)$ allows us to simplify the state description compared to dynamic program (7). In particular, consumer state components $\hat{\mathbf{m}}_t, \hat{\mathbf{n}}_t$ are not needed. It is insightful that bidding policies resulting from Equations (17) and (18) are identical according to the following

Proposition 8. *The index j delivering the maximum in Equations (17) and (18) is the same. Moreover, $\tilde{u}_t^C(\mathbf{w}, N_t, \mathbf{a}, \mathbf{b}) = N_t \tilde{u}_t^E(\mathbf{w}, N_t, \mathbf{a}, \mathbf{b})$ and the result also holds if $\tilde{u}_{t-1}^I(\mathbf{w}, \mathbf{a}, \mathbf{b}, q)$ is replaced by any other function of $\boldsymbol{a}$,$\boldsymbol{b}$.*

This result indicates that information exchange heuristic based on Equation (17) provides a lower bound for coordinated bidding. The bound is reasonable, since information exchange is equivalent to constrained coordinated bidding where all bidders use the same price and assess the future according to the value of independent bidding. Proposition 8 also provides a method for constructing other heuristics with consistent performance under information exchange and bid coordination.

For purposes of these experiments, an analogous rollout heuristic approach to risk pooling is based on a simpler approximation of the future. The expected utility is computed assuming the best constant price for all remaining bidders and time periods. Although, in each period, the price is still flexible and optimized separately, the approximation of the future is potentially less accurate than that of a more flexible lower bound based on Equation (16). The constant price restriction is needed since the bidding outcomes of all consumers jointly affect the utility of risk pooling. Thus, the rollout approach cannot guarantee incentive compatibility. An alternative that guarantees incentive compatibility is to execute independent bidding heuristic for each bidder and to perform risk pooling for the joint final outcomes. The reported risk-pooling result is the maximum of sample averages of these two approaches.

Each of the described heuristic approaches can be made robust with respect to violations of constant success probability assumption. The idea is to implement a limited recall by only including the most recent bid information into the belief state. The amount of recall is a tunable parameter. Old bid results beyond the recall period are much less relevant and can be discarded.

We report results of heuristic coordinated bidding (heuristic information exchange policy is identical) and risk pooling as the percentage improvement in average utility over heuristic independent bidding based on Equation (16). The experiments use the same utility structure and the set of risk tolerances as before, time horizon of $T = 10$ and two combinations of the number of bidders and prices: $N = 15$, $q = 5$ and $N = 8$, $q = 10$. Each of these combinations was examined with two priors: uniform ($\mathbf{a} = \mathbf{e}$) and pessimistic ($\mathbf{a} = 0.1\mathbf{e}$). Moreover, we compare settings where success probability vector $\boldsymbol{\theta}$ is constant and where it undergoes a random walk: the inverse logit transform $\ln \frac{\theta_i}{1-\theta_i}$, $i = 1, \ldots, q$ is subject to Normal(0,1) perturbation, and the resulting new vector is sorted in increasing order. Since too distant history becomes irrelevant in the presence of random perturbations of $\boldsymbol{\theta}$, we consider three levels of recall: up to 2,5 and 10 (full recall) periods in the past. For each resulting combination, we run 1000 Monte-Carlo simulations.

Figures S19–S22 show percentage improvements in average utility as bar charts and 95% confidence intervals as error bars. Each figure is arranged into six plots corresponding to three different levels of risk tolerance and two levels of volatility in $\boldsymbol{\theta}$ (constant or volatile with the SD of 1). In each plot, there are two groups of bars. Bars in a group (coordination or risk pooling) show dependence of benefits on the recall level. The results confirm that coordination leads to persistent benefits in many settings (frequently above 6% and reaching 10%) and even the heuristic risk pooling can provide improvements above 12%. Under the uniform prior, the benefits are lower, but volatility in acceptance probability leads to their increase. The effect of volatility under the pessimistic prior is reversed leading to a decrease in benefits of coordination. In the presence of volatility, it is clearly advantageous to use the recall level of two periods. Also, a higher tolerance to risk and the pessimistic prior generally lead to larger benefits of coordination. Finally, the heuristic risk pooling was quite effective under small risk tolerance.

5. Conclusions

We study the benefits of cooperation between strategic consumers in NYOP product markets. The forms of cooperation include exchange of information as well as bid coordination with and without risk pooling. In the market settings examined experimentally, bid coordination without risk pooling usually does not lead to substantially higher benefits than those from information exchange. An exception is the case of the pessimistic prior (low chances of success) with sufficiently long time horizons. Thus, it is likely not worth the trouble for consumers to implement bid coordination without risk pooling when information exchange forums are already available. On the other hand, introduction of risk pooling by means of outcome-contingent transfer payments leads to a dramatic increase in benefits. Consumers with a low tolerance to risk benefit the most from this form of cooperation. Overall, coordination and risk pooling permit consumers to use more aggressive bidding strategies than the optimal strategies under independent bidding.

We show that bidding strategies of independent consumers are monotonic and, as a result, quite robust even in the presence of bid success rate uncertainty. This is the main reason for limited benefits from cooperation without risk pooling.

From the market operator's point of view, these findings are encouraging. A small increase in consumer's surplus from information exchange and bid coordination is not likely to undermine the revenue of the firms that supply the capacity. Moreover, we have verified that these revenues may increase in some settings. An implementation of a risk-pooling mechanism may make the NYOP market more attractive for risk-averse consumers who want to reduce the uncertainty in the final bidding outcome. This may potentially lead to a surprising outcome that the market operator may not suffer from existence of the bidders' clubs but, rather, benefit from expansion of the market because these clubs attract risk-averse consumers. In our experiments, the only setting where the existence of clubs may seriously affect the revenues is that of a uniform prior (relatively high success probabilities) and consumers with a low risk tolerance.

The coordinated bidding and risk pooling based on egalitarian mechanisms described in this study are feasible for a sufficiently homogeneous group of consumers. Their egalitarian nature ensures that the proposed mechanisms are robust with respect to consumer manipulation of their inputs. We also provide a heuristic extension of the approach for settings where the consumer population is not sufficiently homogeneous. The extended mechanism avoids consumer misrepresentation of WTP by explicitly enforcing incentive compatibility. The practical feature of the mechanisms considered in the study is that consumers are not forced into participation, but rather have an opportunity to self-select, and the related functionality can be made available in social networks along with the standard online forums facilitating information exchange. The mechanisms may be applied not only in opaque settings but also, more generally, in settings with multiple parallel auctions where the reserve price is not certain (such as eBay's Best Offer).

The questions raised in this study are relevant to wider contexts, for example, supply chain analysis. In particular, potential forms of cooperation between small suppliers or retailers who enter into contracts with larger entities by means of B2B auctions or negotiations bear similarity to the scenarios analyzed in this study. Just like strategic consumers, small suppliers or retailers are strategic in their bidding, and can potentially benefit from information exchange, bid coordination, and risk pooling. The main differences between the B2B and consumer settings include the time-scale and the magnitude of resulting contractual relations, the significant market power that a cooperative of retailers or suppliers would wield in the marketplace, as well as ethical and legal implications. A recent example of a B2B auction site with automatic bidding is Groupon Goods Liquidation Auctions (GrouponGoods.bstocksolutions.com) launched by Groupon Goods in Europe.

Acknowledgments

This research was supported by the Natural Sciences and Engineering Research Council of Canada (grant numbers 261512-2009 and 341412-2011). This work was made possible by the facilities of the Shared Hierarchical Academic Research Computing Network (SHARCNET:http://www.sharcnet.ca) and Compute/Calcul Canada. The authors thank the Senior and Associate Editors, and the referees for their constructive suggestions, which helped to better explain the model and improve the results, as well as Andrei Bazhanov for his valuable comments about the manuscript. Guyves Achtari's assistance in implementing computer programs for scalable bidding methods of section 4.3 on SHARCNET facilities is gratefully acknowledged. The fourth author sends prayers of utmost gratitude to God in Whom he sought inspiration during this work.

References

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

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

Aviv, Y., G. Vulcano. 2010. Dynamic list pricing. O. Ozer, R. Phillips, eds. *Oxford Handbook of Pricing Management*. Oxford University Press, 522–584.

Aviv, Y., Y. Levin, M. Nediak. 2009. Counteracting strategic consumer behavior in dynamic pricing systems. S. Netessine, C. Tang, eds. *Consumer-Driven Demand and Operations Management Models*. Springer, 323–352.

Bolton, P., C. Harris. 1999. Strategic experimentation. *Econometrica* **67**(2): 349–374.

Coase, R. 1972. Durability and monopoly. *J. Law Econ.* **15**(1): 143–149.

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

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

Fay, S. 2009. Competitive reasons for the name-your-own-price channel. *Mark. Lett.* **20**(3): 277–293.

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

Fay, S., J. Xie. 2008. Probabilistic goods: A creative way of selling products and services. *Mark. Sci.* **27**(4): 674–690.

Hann, I., C. Terwiesch. 2003. Measuring the frictional costs of online transactions: The case of a name-yourown-price channel. *Man. Sci.* **49**(11): 1563–1579.

Hinz, M., M. Spann. 2008. The impact of information diffusion on bidding behavior in secret reserve price auctions. *Inf. Syst. Res.* **19**(3): 351–368.

Hinz, M., M. Spann. 2010. Managing information diffusion in name-your-own-price auctions. *Decis. Support Syst.* **49**(4): 474–485.

Huang, X., G. Sošić, G. Kersten. 2014. Selling through Priceline? on the impact of name-your-own-price in competitive market. Working paper. Concordia University, Montreal, Canada.

Jerath, K., S. Netessine, S. Veeraraghavan. 2009. Selling to strategic consumers: Opaque selling strategies. S. Netessine, C. Tang, eds. *Consumer-Driven Demand and Operations Management Models*. Springer, 253–300.

Jerath, K., S. Netessine, S. Veeraraghavan. 2010. Revenue management with strategic consumers: Last-minute selling and opaque selling. *Manage. Sci.* **56**(3): 430–448.

Jiang, Y. 2007. Price discrimination with opaque products. *J. Revenue Pricing Manag.* **6**(2): 118–134.

Shapiro, D., X. Shi. 2008. Market segmentation: The role of opaque travel agencies. *J. Econ. Manage. Strategy* **17**(4): 803–837.

Shen, Z., X. Su. 2007. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Prod. Oper. Manag.* **16**(6): 713–728.

Spann, M., B. Skiera, B. Schäfers. 2004. Measuring individual frictional costs and willingness-to-pay via name-your-own-price mechanisms. *J. Interact. Mark.* **18**(4): 22–36.

Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Norwell, MA.

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

Wang, T., E. Gal-Or, R. Chatterjee. 2009. The name-your-own-price channel in the travel industry: An analytical exploration. *Man. Sci.* **55**(6): 968–978.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Figure S1: Benefits (%) of Coordination for $q = 2$ and Uniform Prior as a Function of T.
Figure S2: Benefits (%) of Coordination for $q = 2$ and Pessimistic Prior as a Function of T.
Figure S3: Benefits (%) of Coordination for $q = 3$ and Uniform Prior as a Function of T.
Figure S4: Benefits (%) of Coordination for $q = 3$ and Pessimistic Prior as a Function of T.
Figure S5: Benefits (%) of Coordination for $q = 4$ and Uniform Prior as a Function of T.
Figure S6: Benefits (%) of Coordination for $q = 4$ and Pessimistic Prior as a Function of T.

Figure S7: Change (%) in NYOP Channel Revenue for $q = 3$ and Uniform Prior as a Function of T.

Figure S8: Change (%) in NYOP Channel Revenue for $q = 3$ and Pessimistic Prior as a Function of T.

Figure S9: Change (%) in NYOP Channel Revenue for $q = 4$ and Uniform Prior as a Function of T.

Figure S10: Change (%) in NYOP Channel Revenue for $q = 4$ and Pessimistic Prior as a Function of T.

Figure S11: Change (%) in NYOP Channel Sales for $q = 3$ and Uniform Prior as a Function of T.

Figure S12: Change (%) in NYOP Channel Sales for $q = 3$ and Pessimistic Prior as a Function of T

Figure S13: Change (%) in NYOP Channel Sales for $q = 4$ and Uniform Prior as a Function of T.

Figure S14: Change (%) in NYOP Channel Sales for $q = 4$ and Pessimistic Prior as a Function of T.

Figure S15: Change (%) in NYOP Channel Revenue Per Sale for $q = 3$ and Uniform Prior as a Function of T.

Figure S16: Change (%) in NYOP Channel Revenue Per Sale for $q = 3$ and Pessimistic Prior as a Function of T.

Figure S17: Change (%) in NYOP Channel Revenue Per Sale for $q = 4$ and Uniform Prior as a Function of T.

Figure S18: Change (%) in NYOP Channel Revenue Per Sale for $q = 4$ and Pessimistic Prior as a Function of T.

Figure S19: Change (%) in Average Utility under Heuristic Collaborative Policies for $q = 5$, $N = 15$ and Uniform Prior as a Function of Recall Level.

Figure S20: Change (%) in Average Utility under Heuristic Collaborative Policies for $q = 5$, $N = 15$ and Pessimistic Prior as a Function of Recall Level.

Figure S21: Change (%) in Average Utility under Heuristic Collaborative Policies for $q = 10$, $N = 8$ and Uniform Prior as a Function of Recall Level.

Figure S22: Change (%) in Average Utility under Heuristic Collaborative Policies for $q = 10$, $N = 8$ and Pessimistic Prior as a Function of Recall Level.

Appendix S1: Bayesian Updating for Bid Success Probabilities.

Appendix S2: Proofs and Other Technical Details.

Appendix S3: Coordinated Bidding and Risk Pooling with Nonhomogeneous Consumer Population.

Please Cite this article in press as: Levina, T., et al. Strategic Consumer Cooperation in a Name-Your-Own-Price Channel. *Production and Operations Management* (2015), doi 10.1111/poms.12412