

Online Haggling at a Name-Your-Own-Price Retailer: Theory and Application

Christian Terwiesch

Il-Horn Hann

Sergei Savin

The Wharton School

GSIA

Columbia Business School

March 17, 2003

Abstract

We present a formal model of haggling between a name-your-own-price retailer and a set of individual buyers. Rather than posting a price, the retailer waits for potential buyers to submit offers for a given product and then chooses to either accept or reject them. Consumers whose offers have been rejected can invest in additional haggling effort and increment their offers. The main advantage of this pricing model is that it allows the name-your-own-price retailer to engage in price discrimination: as haggling is costly for the potential buyer, customers with a high willingness to haggle will achieve lower transaction prices. Thus, haggling effort can be used as a self-selection mechanism to implement price discrimination. Our study is motivated by several name-your-own-price retailers that have recently emerged on the Internet. Based on detailed transaction data of a large German name-your-own-price retailer, we present a model of consumer haggling. We then show how this model can be used to improve the decision making of the retailer, who needs to choose a threshold price above which all offers are accepted. Another decision variable for the retailer lies in the user interface design, which allows the retailer to either facilitate or to hinder the haggling of the consumer.

1 Introduction

The emergence of the Internet and its extensive usage for electronic commerce has given companies the opportunity to experiment with a number of innovative pricing models. A well-known example for this is the name-your-own-price (NYOP) model and, more generally, the concept of online haggling. In the NYOP setting, instead of posting a price, the seller waits for an offer by the potential buyer that she can then either accept or reject.

Historically, haggling over prices was the common mode of doing business in the western world, and it is still common practice in pre-industrial societies today. When in 1653 George Fox, the founder of Quakerism, proposed the radical notion that everyone should pay the same price for the same good, his proposal was by and large ignored. Two hundred years later, Aristide Boucicaut introduced fixed prices at his Paris dry goods shop, Bon Marché. The innovation quickly diffused and was adopted by the first new mass retailers in the United States, managed by R. H. Macy and F. W. Woolworth, and J. Wanamaker, who were rapidly replacing the small owner-operated stores. Given the size and organizational structure of these new retailers, store owners had to rely on clerks to interact with customers. This created three important advantages of fixed prices relative to the existing practice of haggling. First, in absence of posted prices, the owner of the retail store had to provide detailed instructions to his clerks on how to conduct the haggling process, requiring extensive training. Second, the capacity of the clerk in checking out customers was constrained by the lengthy haggling for every transaction, requiring the owner to hire additional clerks. Third, principal-agent problems between owner and clerk required close supervision and monitoring.

These disadvantages have lead to the almost complete extinction of haggling in retail settings in most industrialized societies. However, all three of these effects can be dramatically reduced in electronic markets. Instead of having clerks haggle with customers, the firm can create an electronic agent that receives offers from the customer and then considers autonomously which offers to accept and which to reject. The main advantage of this approach is that it allows the seller to engage in price discrimination: as haggling

is costly for the potential buyer, customers with a high willingness to haggle will achieve lower transaction prices. Thus, haggling effort can be used as a price discrimination mechanism, which - under broad circumstances - leads to higher profits compared to a uniform, posted price.

Unlike some early claims of the business press, predicting that the NYOP-model will soon replace fixed retail prices¹, we doubt that the pendulum of economic history will swing back completely towards an economy of collateral bargaining. Nevertheless, the relative ease of transacting in electronic markets will make haggling attractive for market environments currently relying purely on posted prices. This has been most visible in the emergence of several new price intermediaries, including Priceline.com, which has implemented a NYOP model for selling airline tickets, new vehicles, long distance calls, and home loans. In this article, we present a formal model of the haggling process between consumers and a German NYOP retailer. The NYOP retailer we study is an intermediary between a wholesaler and consumers. In our setting, the wholesaler set a wholesale price at which products are sold to the intermediary, and the intermediary sets a threshold price for consumer offers. The consumer submits an offer which is accepted if it exceeds the threshold price.

We begin our analysis by proposing a consumer haggling model. Consumers incur haggling effort for every offer they submit to the NYOP retailer. When making an offer, a consumer balances the cost of offering too much, leading to an extensive information rent for the seller, with the cost of offering too little, leading to additional or wasted haggling cost. After establishing an appropriate consumer model, we turn to the problem faced by the NYOP retailer. The retailer's problem is to determine a threshold price above which it accepts the consumer offer. In addition to these tactical decisions, the retailer also needs to decide to what extent he should hinder or facilitate the haggling of the consumer.

The proprietary data set we built based on our collaboration with the German NYOP retailer together with our analytical results allows us to make the following contributions:

¹ "Haggling goes High-Tech", April 10, 2000, Wall Street Journal

- First, we model the consumer haggling process using a stochastic dynamic programming formulation (**Proposition 1**) and illustrate its predictive power based on the transaction data we obtained from the German NYOP retailer.
- Second, we derive the profit maximizing threshold price (**Proposition 2**) and show how it compares with the threshold prices that were used by the retailer we studied. Based on the consumer data we collected, we also show that our optimal threshold price would have increased retailer profits substantially.
- Third, we discuss under what conditions haggling can be advantageous compared to posted prices. We find that if consumers are very heterogeneous concerning valuation and haggling ability, haggling can lead to higher profits compared to posted prices (**Proposition 3**). However, this was only the case for 1 out of 3 products we studied. At the higher level, the main advantage of haggling is that it allows the wholesaler to complement sales from traditional retail channels with sales from the NYOP channel that has on average lower transaction prices. However, since the NYOP retailer never posts a price, these additional sales do not necessarily come at the expense of lower sales at the higher posted price.
- Finally, we show how changes in the required effort to increment an offer influences consumer haggling, which leads to interesting observations concerning the NYOP interface design (**Proposition 4**).

2 An NYOP Application

All NYOP applications attempt to discriminate consumers according to their willingness to pay, yet there exist different ways this price discrimination is implemented. A closer look at Priceline’s business model across product categories reveals two methods of price discrimination. In the first method, potential buyers place offers on a product, facing uncertainty about some of the product’s attributes. For example, customers placing offers for air travel face uncertainty about the detailed travel schedule and do not know which

carrier will fulfill their demand. This allows Priceline to screen consumers according to their type, while allowing airlines to serve customers that they were previously not able to distinguish from less price sensitive customers. In general, this price discrimination method works well for multi-attribute products, which are fairly close substitutes (air travel, hotel accommodation).

Priceline uses a different method of price discrimination for the sales of undifferentiated goods. Here, Priceline uses haggling effort - representing consumer effort and time loss from the online haggling process - as a way to discriminate between consumers. For example, a consumer placing an offer for calling capacity can start with a low offer and then - upon being rejected after a 60 second wait period - increment the offer. Priceline allows customers to submit 3 consecutive offers for the same phone number and capacity before barring customers to submit additional offers for 24 hours. While this approach can lead to an attractive price for the consumer, the consumer needs to invest in haggling effort (extra offers) to realize a lower price. As price sensitive consumers are more likely to tolerate such disutility, price discrimination is achieved.

Research Setting

The research site underlying this study uses this second method of price discrimination. Figure 1 describes in detail how a consumer interacts with this NYOP retailer. After providing an identification (or registering as a new user), the consumer makes an offer for a product. The retailer then compares the offer with an internal threshold price. If the offer exceeds the threshold, the transaction occurs at the price offered by the consumer. If the offer is below the threshold, the consumer is informed that her offer was too low and is given the opportunity to submit an incremented offer after a certain delay period.

The NYOP set-up we study brings together three perspectives: those of the consumer, the retailer, and the wholesaler. We illustrate these different perspectives based on an example of a hypothetical consumer X, who is interested in buying a personal digital assistant (PDA). Consumer X has seen a posted price of 222 Euro for the product at a large computer discounter. However, she is not willing to spend this amount. For the

sake of argument, assume her willingness to pay is 200 Euro. Consumer X is aware of the NYOP channel and expects to find the exact product there at a lower price. She is uncertain though how much lower this price would be. For this reason, she first submits an offer of 155 Euro, and - upon being notified that her offer has been rejected - increments the offer further to 173, 186, and 196 Euro. At this point, she is informed by the NYOP retailer that her offer has been accepted and receives the product for 196 Euro.

Second, consider the perspective of the wholesaler. The wholesaler has traditionally relied on posted price retailers, who have purchased the product from the wholesaler at a wholesale price of 193 and marked the product up to the 222 Euro mentioned above. Instead of lowering the price in the hope of attracting more customers, the wholesaler can use the NYOP channel to reach people like consumer X who currently abstain from purchasing. At the same time, the effort required for haggling with the NYOP retailer limits the cannibalization between channels.

Finally, consider the position of the NYOP retailer, who acts as an intermediary between consumers and the wholesaler. The intermediary receives offers from consumers and needs to determine a threshold price above which he is willing to accept an offer from the consumer. In the case of the Palm IIIc, this threshold was 193 Euro. A successful offer from the consumer, e.g. the offer of 196 of consumer X, will lead to two sources of profit. First, the intermediary obtains an information rent, the spread between the submitted offer and the threshold price ($196 - 193 = 3$ Euro). Second, if the NYOP retailer chooses a threshold price above the wholesale price, he also obtains an additional profit, consisting of the threshold price minus wholesale price. The NYOP retailer we studied decided to set the threshold price equal to the wholesale price and thus relied on the information rent as the source of profits.

In addition to setting the threshold price, the NYOP retailer also has control over the user interface design. Specifically, he can influence the haggling effort of the consumer via the amount of information the consumer has to key in for every offer. The NYOP retailer also chooses the time delay between receiving an offer and informing the consumer about the outcome of the offer, which directly impacts the consumer's haggling effort. In our

research setting, consumers were informed after five minutes about the outcome of their offer.

Research Questions

Motivated by our interactions with this German NYOP retailer, we seek to answer the following two questions:

(1) How should the NYOP retailer set the threshold price, above which he accepts offers submitted by consumers?

(2) To what extent should the haggling effort from the consumer be increased or decreased via the interface design and feedback mechanisms of the haggling process?

In addition to these two questions that are geared directly to improve decision making at the NYOP retailer, we are also interested in the modeling the consumer haggling process as well as in comparing retailer profits of an NYOP retailer with profits obtained for a traditional posted price setting.

Data Collection and Research Methodology

For the purpose of our research, the German NYOP retailer provided us with a complete history of submitted offers as well as with information about the corresponding threshold prices. In addition we obtained the corresponding customer identifications, which allows us to link the offers a consumer makes for a given product into a sequence of offers for this consumer and this product. Such sequences of offers are the unit of analysis in our research.

Our analysis is based on several consumer electronic products that were offered at our research site in May 2001. We collected data for a personal digital assistant ($B_{PDA} = 246$ offers, $N_{PDA} = 46$ haggling sequences), a CD-rewriter ($B_{CDR} = 365$ offers, $N_{CDR} = 63$ haggling sequences), and a DVD-Player ($B_{DVD} = 363$ offers, $N_{DVD} = 45$ haggling sequences).

We use these data in the two ways. First, we used the haggling sequences for each product to validate our consumer model. This will be discussed in Section 4. Second, we divided the set of haggling sequences for each product into a calibration sample and a hold-

out sample. The hold-out sample was then used to evaluate the performance of our optimal threshold derived in Proposition 3 and to compare it with the actual profits obtained by the NYOP retailer. This will be discussed in Section 5. A potential shortcoming of our study is that our sample is limited to consumer’s who have incurred the effort of registering at the NYOP retailer, which could lead to a sample selection bias. This limits our ability to generalize our findings to the entire consumer population.

3 Related Literature

The NYOP setting we study relates to existing literature on auctions as well as bargaining. This research typically assumes that decision makers are currently making optimal decisions at equilibrium and know that others are doing the same and that best decision responses will be made (Rubinstein 1982).

Consider the literature on *auctions* (Klemperer 1999) and their application to electronic commerce (Pinker *et al.* 2002) first. Auction theory is concerned with the efficient allocation of a scarce good. Sellers are typically interested in an auction mechanism that yields the highest price for this scarce good and research has focused on optimal auction design and comparison of different types of auctions. A famous result, known as the revenue equivalence theorem, states that for the four independent private value auctions (first price, second price, English, and Dutch auction) the expected revenue for the seller is identical. This result is driven by the bidders “competing” for the purchase of the scarce good. However, in the context we study, the seller does not face a supply constraint. Over the one year interaction we had with the retailer, there was not a single case in which a consumer who had submitted an offer above the threshold price was not rewarded with the product. Another difference that our research context has with most of the auction research is the standard assumption made in the analysis of auctions that there are no costs associated with submitting a bid. Typically, it is assumed that in open bid auctions, each bidder in turn submits a bid equal to the previous bid plus the minimum bid increment unless the resulting bid would be higher than her valuation, in which case she exits. Such

an approach is not optimal if the bidder incurs a cost for every bid she places (Daniel and Hirshleifer 1998). Taken together these two observations, we conclude that our setting does not correspond to an auction. For this reason, we refer to the prices submitted by consumers as “offers” as opposed to “bids”.

Next, consider the relationship between our research setting and models of *bargaining*. Bargaining refers to situations where (i) individuals have the possibility of concluding a mutually beneficial agreement, (ii) there is a conflict of interests about which agreement to conclude, and (iii) no agreement may be imposed on any individual without his approval (Osborne and Rubinstein 1990). Bargaining was originally formalized using an axiomatic approach, which has become known as the Nash Bargaining Problem (Nash 1950). Using the notion of extensive games Rubinstein (1982) was the first to formulate bargaining through a procedural approach in which players make decisions sequentially in a pre-specified order. One key determinant of the outcome in many bargaining models is the cost of impatience or delay. In our model, this corresponds to the consumer’s haggling cost.

In the setting we study, the seller uses a common threshold price for all consumers and does not change this threshold price over the course of the haggling sequence. This reflects legal constraints from German trade laws at the time of our study (“Rabattgesetze”)² and the fear of the NYOP retailer that treating consumers differently might be perceived as unfair³.

²Over the last two years, the laws have been interpreted differently in different cases, leading to a somewhat ambiguous legal basis. One major price intermediary was found guilty of illegal price discrimination which resulted significant legal costs and, more importantly, severe damage to the firm’s brand name. The Rabattgesetze were eliminated in the summer of 2001, making online haggling possible from a legal perspective. However, within the time-span of our research cooperation, the site had not yet implemented such additional price discrimination.

³Applying different threshold prices or response times could, if detected, create to perceptions of unfairness from the public. Such a case was experienced by Amazon.com when the firm charged higher prices for DVDs to more loyal customers (in an attempt to leverage their frictional cost advantage). However, customers discovered that they had paid more for the same product than other customers at the same point in time, leading to major negative publicity for the company. Amazon.com ultimately ended

Given a constant threshold price, the bargaining problem corresponds to a game, in which the seller chooses the threshold price and members of the consumer population choose their offering sequences. However, even with this simplification, a characterization of the resulting equilibrium is almost impossible to achieve in practice, as it involves multiple parties with private information. Already, solutions to relatively simple bargaining problems often offer multiple equilibria. The number of possible solutions can be further multiplied when mixed strategies are applied (Kennan and Wilson 1993). For example, Rubinstein’s full information bargaining problem of dividing a pie of 1 between 2 players can be shown to have one perfect equilibrium. However, multiple equilibria arise, if the pie is an amount of money denominated in discrete units (van Damme et al. 1990). Sutton (1986) illustrates that Rubinstein’s game does also not provide a unique equilibrium when a third player is introduced.

The mathematical complexity of bargaining models further increases if one considers multiple agents. Since, in our setting, the NYOP retailer has to choose a common threshold price for all consumers, strategic consumers would take each others action into account when submitting offers. Bargaining settings with many agents require that each agent has some knowledge about each other’s private information. The typical assumption is that each agent’s valuation follows a certain distribution function. In other words, the actual realization is unknown to an agent, however, the distribution of the random variable is common knowledge (see for example Mailath and Postlewaite 1990a; 1990b). Thus, in order to formulate an equilibrium model, we would have to obtain data not only on the private information about each consumer, but also about the consumers’ beliefs about the private information of others.

In addition to the mathematical complexity associated with a bargaining equilibrium model in our setting, the assumption of the auction and bargaining literature that every decision maker already acts optimally is difficult to combine with the objective of advising decision makers on how to improve their actions. In a recent editorial of *Marketing Science*,

up publicly apologizing and refunding all customers who had paid higher prices.

Steven Shugan, editor-in-chief of *Marketing Science*, provides a thoughtful discussion about the different perspectives of equilibrium models on the one hand and models geared towards decision support and improvement on the other hand. As Shugan observes: “It is difficult to advise players who already act optimally or to explore improvements at equilibrium. [...] Given this argument, we might ask when is it appropriate to assume optimal behavior. The answer is the same as with any other assumption. An assumption is appropriate when it provides a good approximation within the context of the research objective (Shugan 2002).”

The objective of this research, as stated in our research questions above, is to improve the decisions from the NYOP retailer. Given this objective, we need to make two simplifying assumptions. First, we restrict the decision space of the NYOP retailer to constant threshold prices. This reflects legal constraints and the need for a “fair process” discussed above. Second, we assume that the consumer “lumps together” the various forms of uncertainty she faces (product valuations of other consumers, beliefs of other consumers, wholesale price) into a single distribution over the NYOP retailer’s threshold price. This reduces the parameters required to describe a consumer substantially enabling us to perform econometric analysis.

Based on these two assumptions, consumer haggling becomes a problem of search. The consumer balances search effort, which is minimized if she only submits one offer, with obtaining a price that lies only slightly above the threshold price, which is achieved if she submits many offers with small increments.

Consumer models of search have been analyzed extensively in the economic literature. As pointed out by Stigler (1961), consumers trade-off their disutility of search with potential price savings resulting from gathering price information. While Stigler considered a model of consumer search across *multiple* sellers, the qualitative effect of search on pricing is similar: firms are able to charge above marginal costs as long as the cost of obtaining price information is non-zero. Following Stigler, several papers (Salop and Stiglitz 1976; Varian 1980; Bakos 1997) examine the effect of consumer search costs on pricing behavior of producers. Models of consumer search have been used repeatedly in empirical stud-

ies of consumer behavior in electronic commerce settings. In a related paper (Hann and Terwiesch 2001), we use a consumer search model to estimate the consumer’s disutility of making a single offer in the context described above. For the average consumer, the disutility of making or incrementing a single offer (haggling cost) was found to be equivalent to about 5.5 Euro for a 200 Euro product (1 Euro=0.9 US\$ at the time of the study).

4 The Consumer Model

A consumer making an offer on a product featured at the NYOP retailer faces the following decision. If the offer she makes in the current round is *lower* than the unknown threshold price held by the retailer, she incurs haggling costs, but does not receive any direct reward. However, an unsuccessful offer does provide additional information, which is valuable for the consumer if she decides to continue the haggling process. If the offer she makes is *higher* than the threshold price, the consumer realizes a reward consisting of the difference between her valuation of the product and her offer. Despite this reward, the consumer now knows that she is likely to have paid too much and thereby left an extra profit margin to the retailer.

Dynamic Programming Formulation

In the following model formulation, we assume that the consumer’s utility function is additive in haggling effort and potential price savings. We also assume that the consumer is risk neutral and that her haggling costs do not change within the course of a haggling sequence. Define c as the cost for the consumer of making (incrementing) one offer. Moreover, let T be the - to the consumer unknown - threshold price. A consumer initially assumes that T is a random variable distributed over $[R_{\min}, R_{\max}]$.

A consumer with a reservation price r submits an offer $x \in [R_{\min}, R^*]$, where $R^* = \min(r, R_{\max})$. If $x \geq T$, the offer is successful and the consumer receives the product, realizing a surplus of $r - x - c$. However, if $x < T$, the offer is rejected. In this case, the consumer can either terminate the haggling process or continue it by incrementing her offer and incurring additional haggling cost c .

If the consumer decides to make a new offer, her information about the threshold price T has improved to an interval of $[x, R_{\max}]$. Given the nature of the updating process resulting in a two-side truncated prior distribution for T and an overall high level of uncertainty about T , we assume that T is uniformly distributed over the corresponding interval. Below, we will show that a consumer model with this assumption provides a good fit with the data we collected from the German name-your-own-price retailer.

In order to model the behavior of a consumer described by the set of characteristics $(c, R_{\min}, R_{\max}, r)$, we consider the optimal expected incremental surplus $V(x)$ earned by a consumer whose last offer x was rejected. The dynamic programming optimality equation for $V(x)$ can be expressed as:

$$V(x) = \max \left(0, \max_{x \leq y \leq R^*} \left(-c + \left(\frac{y-x}{R_{\max}-x} \right) (r-y) + \left(\frac{R_{\max}-y}{R_{\max}-x} \right) V(y) \right) \right), \quad (1)$$

where $\left(\frac{y-x}{R_{\max}-x} \right)$ is the probability that a new offer y exceeds the retailer's threshold price T , given that the last offer x did not exceed T . The outer maximization operator reflects the choice that the consumer has after each unsuccessful offer: to terminate the haggling process or to submit an incremented offer. The recursion (1) is complemented by the boundary condition

$$V(R^*) = \max(0, r - R^* - c) \quad (2)$$

Expressions (1) and (2) define the consumer haggling problem. The recursion (1) can be re-written in a more convenient form by introducing the new variable $v = (R^* - x)/c$ and new value function $L(v) = (R_{\max} - R^* + vc)V(R^* - vc)/c^2$. Then, (1) becomes

$$L(v) = \max(0, -(A+v) + M(v)), M(v) = \max_{0 \leq u \leq v} ((v-u)(B+u) + L(u)), \quad (3)$$

with

$$A = \frac{R_{\max} - R^*}{c}, B = \frac{r - R^*}{c}. \quad (4)$$

Note that $A, B \geq 0$, and $AB = 0$. The boundary condition (2) can be re-expressed as

$$L(0) = M(0) = 0. \quad (5)$$

Properties of the Optimal Solution

The dynamic program defined above does not have a well-defined horizon. This makes it impossible to apply standard backward induction technique as in the case of the finite-horizon problems. For this reason, obtaining closed-form expression for the value function is typically not feasible

However, the specific structure of (1)-(2) allows us to construct a closed form solution in this particular setting. The building blocks for the optimal value function are presented by the following family of recursive value functions.

Definition 1: For any $n \in N$, let $L_n(v) = \alpha_n v^2 + \beta_n v + \gamma_n$, where $\alpha_0 = 0$, $\beta_0 = \max(0, B - 1)$, $\gamma_0 = 0$, and

$$\begin{aligned}\alpha_n &= \frac{n}{2(n+1)}, \beta_n = B + \frac{1}{n+1} \left(\beta_0 - B - \frac{n(n+3)}{2} \right), \\ \gamma_{n+1} &= \gamma_n - A + \frac{\left(\beta_0 - B - \frac{n(n+3)}{2} \right)^2}{2(n+2)(n+1)}, n \in N.\end{aligned}\tag{6}$$

Also, define

$$v_{n+1}^* = \min(v | L_{n+1}(v) \geq L_n(v)), n \in N, v_0^* = 0.\tag{7}$$

Using (6), we can re-express (7) as:

$$v_{n+1}^* = \left(\sqrt{1 + A + \beta_0 - B} + \sqrt{1 + \frac{n(n+3)}{2}} \right)^2 - A, n \in N.\tag{8}$$

The recursive family $\{L_n(v)\}$ is directly related to the optimal consumer haggling strategy as expressed by equations (3)-(5):

Proposition 1 (consumer haggling model):

a) Let $v_{\max} = \frac{R^* - R_{\min}}{c}$, $z = \sqrt{v_{\max} + A} - \sqrt{1 + A + \beta_0 - B}$, and

$$\bar{n} = \lfloor \frac{1}{2} \left(\sqrt{1 + 8z^2} - 1 \right) \rfloor = \max \left(n \in N | n \leq \frac{1}{2} \left(\sqrt{1 + 8z^2} - 1 \right) \right).\tag{9}$$

Then, the solution to the consumer haggling problem is given by

$$L^*(v) = \begin{cases} L_{\bar{n}}(v), & v_{\bar{n}}^* \leq z \leq v_{\max}, \\ L_n(v), & v_n^* \leq v \leq v_{n+1}^*, n = 0, \dots, \bar{n} - 1. \end{cases}\tag{10}$$

b) The optimal number of offers made by a consumer is:

$$n^* = \begin{cases} \bar{n} + 1, & r > R_{\max} + c, \\ \bar{n}, & r \leq R_{\max} + c. \end{cases} \quad (11)$$

The values of optimal offers can be computed as

$$\begin{aligned} x_k^* &= \begin{cases} R_{\min} + \frac{ck(\bar{n}-k+1)}{2} + (R_{\max} - R_{\min}) \left(\frac{k}{\bar{n}+1} \right), & r > R_{\max} + c, \\ R_{\min} + \frac{ck(\bar{n}-k+1)}{2} + (r - c - R_{\min}) \left(\frac{k}{\bar{n}+1} \right), & r \leq R_{\max} + c, \end{cases} \\ k &= 1, \dots, \bar{n}, x_{\bar{n}+1}^* = R_{\max}. \end{aligned} \quad (12)$$

The closed-form expressions for the total number and the values of consumer offers (11)-(12) provide an important tool for the analysis of the consumer haggling behavior.

Numerical Example

To illustrate our notation as well as the consumer trade-off, consider a consumer of type (5,100,200,200), corresponding to haggling cost $c = 5$, valuation $r = 200$, and initial information $[R_{\min}, R_{\max}] = [100, 200]$. In other words, the consumer incurs a disutility equivalent to 5 Euro for every offer she places, has seen a quoted price from another channel for the same product at 200 Euro (or has a no-purchase utility of 200 Euro) and expects the threshold price to be between 100 and 200 Euro.

Using our decision model (1)-(2), she would initially offer $x_1^* = 129$. Assuming this first offer, as well as the following offers were rejected by the retailer, the consumer would subsequently make the following offers: $x_2^* = 153$, $x_3^* = 172$ up to $x_4^* = 185.25$. If the last offer of 185.25 was not successful, it is optimal for the consumer to terminate the haggling process, as the potential benefits (savings relative to 200) would not justify the additional haggling effort ($n^* = 4$). This pattern is illustrated by Figure 2.

Now, compare this consumer with a consumer of equal characteristics, except a lower haggling cost, $c = 2$. As can also be seen in Figure 2, the decrease in haggling cost motivates the consumer to increase her haggling effort, leading to smaller increments between consecutive offers. The overall haggling effort, capturing the maximum number of offers the consumer would make prior to terminating the haggling process, increases to

$n^* = 8$. Finally, consider the impact of initial information on the haggling process. If the consumer has a “sharper” prior about the threshold price ($R_{\min} = 150$ instead of 100), the first offer submitted by the consumer will be higher and n^* decreases to 2.

Estimation of Consumer Characteristics

Instead of computing x_i^* and n^* for a given quadruplet of $(c, R_{\min}, R_{\max}, r)$, we can also use the mapping procedure in the opposite direction. In other words, upon observing empirically a sequence of consumer offers x_1 to x_L , where x_1 is the first offer and x_L the last, we can attempt to estimate the parameters $(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r})$ that best fit the observed sequence. Consider for the moment only haggling processes that included four or more offers by the consumer. For illustrative purposes, Figure 3 provides two actual haggling processes we observed at our research site. A comparison of the two sequences suggests that the former comes from a consumer with lower R_{\min} and a higher c .

In general, we can search for the parameter quadruplet $(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r})$ that achieves the best fit with the observed data by solving the following optimization problem:

$$\min_{\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r}} \sum_{i=1}^L \left(x_i - x_i^*(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r}) \right)^2 \quad (13)$$

The optimization (13), which resembles the method of least squares in econometrics, must be carried out under the constraint reflecting our knowledge about the number of offers the consumer placed. In particular, if the last offer was rejected, and, thus, the consumer opted out of the haggling process, the number of offers we observe, L , is the maximum number of offers the consumer was willing to place, n^* . Thus, we add the following constraint on potential estimators $(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r})$:

$$n^*(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r}) = L. \quad (14)$$

If the last consumer offer was accepted, L serves as the lower bound for the number of offers a consumer was willing to place, n^* . In this case, the constraint becomes:

$$n^*(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r}) \geq L. \quad (15)$$

Our optimization approach allows us to reduce a sequence of offers to four parameters.

Validation of Consumer Model

To demonstrate that reducing the haggling sequences to a parameter quadruplet does not lose too much of the underlying information, we can recompute the optimal haggling sequence from the four parameters via the dynamic programming formulation and then compare the resulting “predicted” offers with the actual offers. Figure 4a compares the predicted values (vertical axis) with the actual offers (horizontal axis) from our sample. A perfect fit would correspond to a straight line through the origin and unit slope. A simple regression analysis between actual and predicted values shows that 97.5% of the variance in consumer behavior is captured by our model. Moreover, the slope of the corresponding regression line is estimated as 1.01, thereby close to identity.

The length of the sequences ranges between 4 and 12. Given the ratio between parameters and observation, special attention is given to haggling sequences of length 7 and longer. Limiting the regression analysis between actual and predicted offers to haggling sequences of length 7 or longer does not significantly change the corresponding fit (adj. R^2 at 97.0%).

Each of the haggling sequences we use for the validation of the consumer model provides only a small number of data points. When assessing the fit of our consumer model, it is important that we not only look at the *absolute* fit, but also compare the fit *relative* to other models. To allow for such a relative comparison, we define two competing models of consumer haggling. Unlike our consumer model outlined in Proposition 1, which is based on rational consumer behavior and thereby exhibits a strong face validity, the two competing models are “greedy” heuristics of consumer behavior.

- In the *constant increment model*, the consumer is defined by her first offer, x_1 , the number of offers she placed, N , and her last offer, x_N . Based on the triplet (x_1, N, x_N) we predict the i -th offer of the consumer as $\hat{x}_i = x_1 + i \frac{x_N - x_1}{N-1}$. In other words, the increments are evenly spaced between the first and the last offer.
- In the *population based increment model*, the consumer is defined by her first offer, x_1 , and the number of offers she placed, N . The model also uses the average increment of

the consumer population Δ_i for $i = 1..N$. We predict the i -th offer of the consumer as $\hat{x}_1 = x_1$ and $\hat{x}_i = \hat{x}_{i-1} + \Delta_i$. Thus, we calibrate the haggling sequence based on the first offer and the number of submitted offers, but otherwise, assume that the consumer behaves similar to the rest of the population.

Table 1 compares the explanatory power of these two models with our consumer model defined by Proposition 1. The comparison between actual and predicted data is done for the offers (350, 390, 420, 440 for consumer X's haggling for the PDA) as well as for the increments (40, 30, 20). We look at all haggling sequences as well as the sub-sample consisting of sequences with $N \geq 7$. Note that offers are easier to predict than increments, as is reflected in the better fit of the validating regression analysis. Based on Table 1, we can make the following observations. First, the fit of our model clearly dominates the fit of the other two models. In all four cases, our model explains significantly more variance. Second, we observe that our model is out-performing the relative comparisons especially for the most difficult validation setting: when predicting increments (opposed to absolute offers) for long haggling sequences (seven offers and more), our model explains 52.4% of the variance, while the comparison models only predict 21% and 16.5% respectively.

Taken together the good fit of our model with its face validity based on the consumer decision problem described in Proposition 1, we will use this representation of consumer behavior in our effort to support the decision making of the NYOP retailer.

5 The Optimal Threshold Price

A retailer operating a NYOP site needs to set the threshold price T to maximize the cumulative profit from all successful offers. We assume that the wholesaler of the product charges the retailer a wholesale price w . Moreover, we assume that the physical inventory is owned entirely by the wholesaler and there exists no binding supply constraint. Both assumptions were clearly fulfilled in the case we studied.

We assume that the consumer market consists of a number of consumer types each de-

scribed by a unique quadruple of parameters $(c, R_{\min}, R_{\max}, r)$. In particular, assume that haggling offers coming from a consumer with parameters $(c, R_{\min}, R_{\max}, r)$ are described by Proposition 1. We begin our analysis by looking at a homogeneous consumer market such that the parameters $(c, R_{\min}, R_{\max}, r)$ are the same for all consumers. Following the notation of the previous Section, we use n^* and $\hat{x} = x_{n^*}^*$ to denote the optimal number of offers and the value of the highest offer, respectively. We observe that \hat{x} would be the last offer produced by a consumer if the threshold price T is set higher than \hat{x} . For a fixed threshold price level $T \leq \hat{x}$, define the index of the first offer exceeding the threshold price:

$$\bar{k}(T) = \min(k | x_k^* \geq T) \quad (16)$$

If the threshold price is set at T , the revenue generated from a consumer is equal to the smallest offer, if such exists, exceeding T , $x_{\bar{k}(T)}^*$. Consequently, it is optimal to equate the threshold to the value of the highest potential offer \hat{x} . We summarize these observations in the form of a Lemma:

Lemma (optimal threshold price in a homogeneous market): *Let w be the wholesale price for the offered product. The expected profit per consumer for a threshold price $T > w$ in a homogeneous market is given by*

$$\Pi_h(T, w) = \begin{cases} 0, & \text{for } T > \hat{x}, \\ x_{\bar{k}(T)}^* - w, & \text{for } T \leq \hat{x}. \end{cases} \quad (17)$$

Consequently, the profit maximizing threshold price T^ corresponds to the value of the highest consumer offer \hat{x} .*

The Lemma states that the highest offer $\hat{x} = x_{n^*}^*$ determines the optimal threshold price. It is clear, however, that the problem of maximizing (17) is degenerate, since any threshold value in the interval between the second highest offer and the highest one $(x_{n^*-1}^*, x_{n^*}^*]$ would generate the same profit. The Lemma assumes that $w < \hat{x}$, i.e. the wholesale price paid by the retailer is consistent with the consumer preference for the product.

In a heterogeneous market, the optimal pricing problem is more complex. We consider a

consumer market consisting of I groups, so that the i -th group, $i = 1, \dots, I$, is characterized by the set of parameters $(c^i, R_{\min}^i, R_{\max}^i, r^i)$ common for all members of the group. We also assume that a consumer belongs to group i with a probability p_i , $\sum_{i=1}^I p_i = 1$.

The consumer market is thus defined by a set of I vectors $(p_i, c^i, R_{\min}^i, R_{\max}^i, r^i)$. Define n_i^* as the optimal number of offers and $\hat{x}_i = x_{n_i^*}^*$ as the highest offer from a consumer belonging to group i (without loss of generality, we assume that $\hat{x}_1 \leq \hat{x}_2 \leq \dots \leq \hat{x}_I$). In addition, define x_{ik}^* , $i = 1, \dots, I$, $k = 1, \dots, n_i^*$ as the k -th offer from a customer in the i -th group. Similarly to (16), we introduce

$$\bar{k}_i(T) = \min(k | x_{ik}^* \geq T) \quad (18)$$

For a wholesale price w , we introduce the smallest index of a group whose highest offer exceeds w :

$$i(w) = \min(j | \hat{x}_j \geq w). \quad (19)$$

Finally, for pairs of consumer groups $i = 1, \dots, I$ and $j = i + 1, \dots, I$ let \hat{x}_{ji} be the lowest offer of group j exceeding the highest offer \hat{x}_i of group i :

$$\hat{x}_{ji} = x_{j\bar{k}_j(\hat{x}_i)}^* \quad (20)$$

As suggested by the Lemma above, the threshold price in the case of a heterogeneous market should be set to one of the highest offers \hat{x}_i .

Proposition 2 (optimal threshold price in a heterogeneous market): *The expected profit per consumer for a wholesale price w and a threshold price $T > w$ in a heterogeneous market is given by*

$$\Pi_n(T, w) = \begin{cases} 0, & \text{for } T > \hat{x}_I, \\ \sum_{j=l+1}^I p_j (x_{j\bar{k}_j(T)}^* - w), & \text{for } \hat{x}_l < T \leq \hat{x}_{l+1}, l = i(w), \dots, I - 1 \\ \sum_{j=i(w)}^I p_j (x_{j\bar{k}_j(T)}^* - w), & \text{for } w < T \leq \hat{x}_{i(w)}. \end{cases} \quad (21)$$

Let

$$j^*(w) = \arg \max_{j=i(w), \dots, I} \left(p_j(\hat{x}_j - w) + \sum_{l=j+1}^I p_l(\hat{x}_{lj} - w) \right). \quad (22)$$

Then, the profit-maximizing threshold price T^* is equal to $\hat{x}_{j^*(w)}$.

Proposition 2 states that selecting the optimal threshold price from the set of best offers from each consumer group requires a trade-off between potential profits in the event that the consumer offer exceeds the threshold and the probability of this event. Thus, a lower threshold price will lead to a lower profit from a given consumer while increasing the number of incidences in which offers are accepted.

In Figure 5 we illustrate this trade-off for the case of $I = 2$ consumer groups such that $w < \hat{x}_1 < \hat{x}_2$. Assuming that a consumer belongs to group 1 with a probability p_1 and to group 2 with a probability $p_2 = 1 - p_1$, we observe that the expected profit under thresholds $T = \hat{x}_1$ and $T = \hat{x}_2$ is given by $\Pi_n(T = \hat{x}_1) = p_1\hat{x}_1 + p_2\hat{x}_{21} - w$ and $\Pi_n(T = \hat{x}_2) = p_2(\hat{x}_2 - w)$, respectively. Consequently, $\Pi_n(T = \hat{x}_1) \geq \Pi_n(T = \hat{x}_2) \Leftrightarrow p_1 \geq p_1^c = \frac{\hat{x}_2 - \hat{x}_{21}}{\hat{x}_1 - w + \hat{x}_2 - \hat{x}_{21}}$. Thus, it is optimal to set the threshold at the level of the “lower” offer \hat{x}_1 , provided that the probability for a consumer to belong to group 1 is high enough. As this probability drops below the critical level p_1^c , it becomes optimal to switch the threshold to the “higher” level \hat{x}_2 , thus lowering the probability of the “buy” event in anticipation of much higher potential profits. We note that the “switching” probability level p_1^c is an increasing function of the wholesale price w , so that higher wholesale prices increase the importance of the “higher offer” class, inducing higher optimal threshold values.

Validation Procedure for the Optimal Threshold Price

We can use the transaction data we collected from the German NYOP retailer to test the performance of the optimal threshold price derived above. In our analysis, we assumed that each haggling sequence was generated by a consumer belonging to a distinct consumer group.

Unlike in the validation of our consumer model, where we used the entire sample to assess the goodness of fit, an evaluation of the optimal threshold price requires a different approach. Choosing a threshold price which performs better than the one chosen by our research site ex-post would be of little value in improving the actual decision making. For this reason, we divided the haggling sequences for each product into two equal or

nearly equal sub-samples. The first sub-sample (calibration sample) was used to generate consumer types and to calibrate the decision rules. The calibration sample consisted of the 50% of the haggling sequences the NYOP retailer received first and was used to estimate the parameters of the consumer population. The second sub-sample (hold-out sample) was used as a testing ground for the decision rules. Dividing the sample based on arrival time opposed to a random split allows us to retrospectively recreate the managerial setting as faced by the management of the NYOP retailer.

In the calibration sample, we used each completed haggling sequence to compute the parameter triplet $(\hat{c}, \hat{R}_{\min}, \hat{R}_{\max}, \hat{r})$ that achieved the best fit as defined by (13). The resulting parameter quadruplets were used to characterize consumer groups and to support the optimization for the threshold price as outlined in Proposition 2. For example, in the case of the PDA, $N_{PDA}^1 = 23$ consumer groups were selected to reflect this consumer market.

Haggling sequences in the hold-out sample were used to create hold-out consumers following the ex-post classification approach described in the validation section. Assuming equal probabilities across consumer groups, we then randomly drew “virtual” consumers from the hold-out population and created haggling sequences using Proposition 1. For each sequence, we can establish its contribution to retailer profits for a given threshold price. The advantage of working with virtual consumers, opposed to the actual offers we collected, is as follows. Consider a consumer, who placed offers (in Euro) 180 and 195 unsuccessfully, yet achieved a successful offer in the third round for 201. Assume that the threshold price was at $T = 200$. This sequence of observed offers is *endogenous* with respect to the threshold price: if the threshold price had been $T = 205$, the third offer would have been unsuccessful and we can only speculate if the consumer might have incremented her third offer further. Thus, when evaluating the performance of a decision rule different from the one that was actually used at our research site, it is important to analyze consumers at the level of their consumer characteristics $(c, R_{\min}, R_{\max}, r)$ opposed to working with their actual offers, x_i . Given the good ex-post fit of our model (see Figure 4), little information is distorted this way.

Validation Results for the Optimal Threshold Price

Table 2 illustrates the performance of the our optimal threshold pricing rule. The second and third column summarize the threshold price used by our research site and the corresponding profits. The third and fourth column present threshold and profits computed based on Proposition 2. The relative profit improvements from the policies recommended by our static thresholds are sizable: 10% for DVD-Player, 11% CD-Rewriter, and a notable 72% for the PDA. We observe that for all three products, Proposition 2 recommended a higher threshold price than what was used in practice.

To validate this pattern, we computed the *ex-post* optimal threshold and the corresponding profits (last two columns in Table 2). While such a number does not reflect the decision situation as faced by our research site (it uses information that was only available in hind-sight), it does provide an upper bound on profits. A comparison between the actual threshold, the threshold suggested by Proposition 2, and the *ex-post* optimal threshold reveals that our results indeed come much closer to this *ex-post* optimal solution.

Comparison to Posted Prices

In addition to validating our optimal threshold prices of Proposition 2, we can also use the estimates of consumer characteristics to compute the optimal posted price and the corresponding profits. Towards this end, we use the valuations we estimated for the consumers in the calibration sample to create a demand curve and then compute the optimal posted price for a given wholesale price w .

For the DVD player, we obtain an optimal posted price of 246 Euro and a resulting profit of 13.32. For the PDA, the optimal posted price is 207 (profits of 8.61) and for the CD-RW the optimal posted price is 166 (profits of 5.97).

Based on these results, we observe that for one out of the three products, the profits obtained based on haggling exceeded the profits the retailer would have obtained from posting prices. The intuition for this result is as follows: if the market is very heterogeneous and there exists a market segment with high haggling costs and high product valuations, not posting prices allows the NYOP retailer to obtain a substantial profit from a few

customers. This compensates for the overall inefficiencies that haggling creates for the entire consumer population. The comparison between profits under haggling vs. profits under posted prices is formalized in Proposition 3. Unfortunately, providing analytical results for any number of customer types is not possible, and we have to limit our analysis to the special case of two consumer groups.

Define two groups of consumers, $(p_1, c_1, R_{\min}^1, R_{\max}^1, r_1)$ and $(p_2, c_2, R_{\min}^2, R_{\max}^2, r_2)$. Here $p_i, i = 1, 2$ expresses the probability that a consumer belongs to group i with characteristics $(c_i, R_{\min}^i, R_{\max}^i, r_i)$, $p_1 + p_2 = 1$. To simplify analysis, we assume $R_{\min}^1 = R_{\min}^2 = 0$, $R_{\max}^1 = r_1$, $R_{\max}^2 = r_2$. Without loss of generality, we assume that $r_2 > r_1$. Then, the following general result can be obtained with respect to relative profit values under fixed pricing vs. haggling:

Proposition 3 (Posted prices vs. haggling)

Let $\frac{r_2}{4+2\sqrt{3}} < c_2 \leq \frac{r_2}{4}$ and $w + \frac{p_2 r_2}{2p_1} < r_1 < \frac{r_2}{2}$. Then there exists a threshold value of the frictional cost $c_1^ > 0$ such that for any $0 < c_1 < c_1^*$ optimal haggling profits Π_h are higher than the optimal fixed-price profits Π_f .*

To illustrate the intuition of Proposition 3, consider a retailer facing the following two groups: $(p_1, c_1, R_{\min}^1, R_{\max}^1, r_1) = (0.995, 0.01, 0, 101, 101)$ and $(p_2, c_2, R_{\min}^2, R_{\max}^2, r_2) = (0.005, 50, 0, 300, 300)$. Let the wholesale price be $w = 100$. Given the reservation prices of both groups, the best posted price would be at $R = 300$ which would lead to expected profits of $\Pi_f = (R - w) p_2 = 1$. We can compute the maximum offers that each consumer group would be willing to submit in a haggling situation as $\hat{x}_1 = 100.974$ and $\hat{x}_2 = 150$ (because of high haggling costs, consumers of the second group actually place only one offer at 150). Setting a static threshold at $T = \hat{x}_1$, the retailer earns an expected profit of $\Pi_h = (\hat{x}_1 - w) p_1 + (\hat{x}_2 - w) p_2 = 0.974 \times 0.995 + 50 \times 0.005 = 1.219 > \Pi_f$. Thus, we observe that in a market of strong heterogeneity concerning reservation prices as well as haggling costs, haggling can lead to higher profits than posted prices.

While profits for the CD-Rewriter exceeded the optimal posted price profits, the corresponding difference was relatively small. Moreover, for the other two products we observed

that posted price profits clearly dominated profits under haggling. Thus, at first sight it seems that posted prices are more profitable than haggling. However, this ignores one very important element of our research setting discussed in Section 2. The role of the NYOP retailer is not to provide the primary sales channel for the wholesaler, but to allow a set of customers, who currently abstain from purchasing, to obtain the product at a lower price. The haggling format allows the wholesaler to minimize the cannibalization of the existing retail channel as no price is ever posted. Consequently, haggling enables the retailer to price discriminate in two different forms:

- It lies in the nature of haggling that customers who are willing to invest more haggling effort are able to achieve a better transaction price. As long as haggling costs are mildly correlated with the consumer’s willingness to pay, this creates price discrimination *within* the population of customers interacting with the NYOP retailer (Proposition 3).
- Using an NYOP retailer allows the wholesaler to segment the overall consumer population into a conventional retail channel with posted prices and a haggling channel for customers with a lower willingness to pay. This creates price discrimination *across* channels. For the products we analyzed, we found that the threshold prices used by the NYOP retailer were substantially lower than typical posted prices in traditional retail channels.

6 Design of Haggling Mechanism

In the context we study, the NYOP retailer is able to influence the haggling cost of a consumer by choosing the time delay with which a consumer is notified about an unsuccessful offer. If a consumer received instantaneous feed-back on her offer, she would incur less effort following an incremental search strategy compared to a one day delay. Thus, the retailer is able to scale the consumer’s frictional costs upwards or downwards.

Proposition 4 (Design of Haggling Mechanism): *a) The optimal number of offers, n^* , is a non-increasing function of c . In particular, the consumer engages in haggling if and only if her haggling cost c does not exceed the critical value $c_0 = \frac{(r-R_{\min})^2}{2(R_{\max}-R_{\min})}$.*

b) For valuation r , and prior information $[R_{\max}, R_{\min}]$ such that $r \leq R_{\max} + c$, define

$$c_i = \min(c | n^*(c) = i), i \in N \quad (23)$$

and let $\hat{x}(c) = x_{n^(c)}^*$ denote the highest consumer offer. Then, $\hat{x}(c) < \hat{x}(c = 0) = \min(r, R_{\max})$ for any $c > 0$, and*

$$\begin{aligned} \frac{\partial \hat{x}(c)}{\partial c} &\geq 0, c_{i+1} < c < c_i, \\ \lim_{c \rightarrow c_i - 0} \hat{x}(c) &> \hat{x}(c_i), i \in N. \end{aligned} \quad (24)$$

As indicated by Proposition 4a, a consumer with a higher product valuation will be willing to accept a higher per-round haggling cost c to engage in haggling with the NYOP retailer. Similarly, a consumer with a higher perceived upper bound for the threshold price, R_{\max} , will be less willing to accept a high cost of haggling. Note that the influence of the R_{\min} on c_0 is different from that of r and R_{\max} . In particular, as R_{\min} grows for fixed R_{\max} and r , two trends are at work. On the one hand, growing R_{\min} indicates that the expected bargain value associated with buying from the NYOP retailer is decreasing. On the other hand, as $[R_{\min}, R_{\max}]$ “shrinks”, the uncertainty of that value also diminishes, and the perceived effort of identifying the threshold is reduced. For small values of R_{\min} ($R_{\min} < 2R_{\max} - r$) the first trend is more pronounced, and the growth of R_{\min} results in the decline in the participation barrier ($\frac{\partial c_0}{\partial R_{\min}} < 0$). For larger values of R_{\min} , the second trend dominates, resulting in $\frac{\partial c_0}{\partial R_{\min}} > 0$.

We observe from Proposition 4b that the value of the highest potential offer made by a consumer exhibits an interesting non-monotone behavior with respect to her haggling cost c . Small changes in c typically do not influence the optimal number of offers. Thus, for a small increase in haggling cost the consumer will make a more “aggressive” (higher) terminal offer. However, a larger increase in c can lead to the “loss” of an offer (a decrease in n^*). In this case, the value of the highest bid can actually decrease, reflecting the

willingness of a consumer to terminate the haggling process “earlier”. Note that, despite such non-monotonicity, the largest value of the last offer is observed when the cost of haggling decreases to zero. In this case, the highest offer submitted by a consumer reflects either her estimate of R_{\max} or her product valuation r , whichever is smaller.

The result that the NYOP retailer can “squeeze” a maximum surplus out of the consumer if haggling costs are *decreased* is certainly counter-intuitive. At first glance, we would expect that the additional flexibility for the consumer to increment her offer with little effort should move surplus from the retailer to the consumer. To better understand why a low haggling cost puts the consumer at a disadvantage, consider the case where haggling costs are sufficiently high so that the consumer only submits one offer. Upon receiving the offer, the NYOP retailer is in a position in which it is optimal to accept every offer above the wholesale price w . Because of the high cost of haggling, the buyer has a credible commitment that her offer will not be further incremented. If, however, haggling is made easier for the consumer, the consumer loses this opportunity of credibly committing herself to not increment her offer, which is in the advantage of the NYOP retailer. Thus, lower haggling cost allows for a higher granularity in implementing the price discrimination method.

As haggling costs approach zero, in the limit the haggling channel resembles posted prices. However, interpreting this limiting case deserves some further discussion. First, the limiting case is a purely hypothetical case, as haggling would always require some minimum effort of the consumer, most importantly reflecting the consumer’s disutility of keying in information. Second, as indicated by Proposition 3, under certain conditions, haggling may actually generate higher profits than posted prices. Third, our analysis only focuses on the effect that changes in haggling effort have on the NYOP channel in isolation. If haggling becomes too easy, it is likely that other consumers, who are currently using the posted price channel switch to the NYOP channel. Thus, while from the perspective of the NYOP retailer small haggling costs are desirable, the wholesaler is likely to be concerned with the cannibalization of the existing channels.

The Role of Electronic Agents

Many industry experts predict that electronic agents will soon dramatically decrease the consumer's haggling costs. Proposition 4 allows us to predict some of the consequences of electronic agents, both from the perspective of the consumer, as well as from the perspective of the NYOP retailer. Everything else constant, electronic agents will be advantageous to the consumer. One could imagine that the electronic agent knows its principal's three parameters R_{\min} , R_{\max} , and r as well as her payment information. This would leave the consumer without any interaction with the retailer.

For any static threshold price, early adopters of electronic agents would gain the most as they are able to negotiate the lowest price for them without having to incur the actual haggling effort. As electronic agents become more widely used, the retailer is likely to respond. Specifically, if electronic agents are widely adopted, the retailer might abandon any price discrimination strategy and move to fixed pricing.

Another possible response to electronic consumer agents is for the retailer to delay the response to submitted offers. If every round of haggling would take a day, haggling would be costly to the consumer, even if it would not require actual effort. In this scenario, the electronic agent would have to trade-off response time to the principal with the principal's desire for an attractive price. Thus, price-discrimination based on haggling is sustainable if the NYOP retailer chooses an appropriate response-time.

7 Conclusion

In this paper, we presented a model of online haggling between a NYOP retailer and a set of consumers. While investing effort in haggling is wasteful from a welfare perspective, it does allow both retailer and wholesaler to engage in a finer market segmentation. The wholesaler uses a NYOP retailer as an additional channel to serve parts of the consumer population who is not willing to purchase the product at the posted price. Similarly, the NYOP retailer is able to engage in price discrimination within the set of consumers who visit his web-site: customers who are willing to haggle extensively (low c in our model) are - on average - obtaining the product at a lower price, as they are able to erode the

NYOP retailer’s information rent further than consumers who only submit a few offers (higher c in our model). Thus, an NYOP channel allows the wholesaler to increase profits, compared to selling only at a posted price.

We began our analysis by looking at the consumer’s decision problem, trading-off haggling effort with price savings. We model this trade-off as a search problem. The analytical solutions we derive (Proposition 1) explain more than 97% of the variance in the offers submitted by consumers and 66% of the variance in offer increments. The fact that consumers follow such a predictable pattern in their haggling policy benefits the retailer. After collecting data about consumer characteristics, the retailer can use the results of Proposition 2 to optimally set her threshold price. Based on the empirical transaction data we collected from our research site, we found that the threshold price we derive this way improves profits of the NYOP retailer substantially.

In addition to the tactical problem of choosing an appropriate threshold price, we also analyzed the more strategic question to what extent the NYOP retailer should support the haggling effort of the consumer. Such changes in haggling effort can be achieved by adjusting the time the retailer takes to inform a consumer about the outcome of her offer or by the design of the interface itself. Counter to initial intuition, we find that lowering haggling costs may actually hurt the consumer, as it eliminates her opportunity to credibly commit towards not incrementing her offer. We also compare the optimal threshold price with the optimal posted price. We find that under some conditions, haggling can lead to higher profits compared to posted prices. However, the main advantage of the haggling model is not that it out-performs posted pricing within a channel, but that it provides the wholesaler an opportunity to reach customers who currently abstain from purchasing. Developing an integrated, multi-channel strategy for the wholesaler, including a joint optimization across channels would be an interesting extension of our research.

A second opportunity for future research lies in extending our analysis of the NYOP retailer’s decision rule of which offers to accept. The focus of our model is one of decision support for the NYOP retailer, not of developing an equilibrium model of bargaining. Future research is needed to derive equilibrium models for the bargaining process between

the NYOP retailer and a set of heterogeneous consumers. This is an especially challenging task, if the NYOP retailer would be allowed to change the threshold price over the course of a haggling sequence.

Finally, future research is needed to better understand the impact of electronic agents on haggling in a NYOP context. Electronic agents would reduce haggling effort for the consumer, yet, due to the response delay of the retailer, not entirely. Moreover, as discussed in conjunction with Proposition 4, a reduction in haggling effort is not necessarily in the interest of consumers. However, as long as only few consumers have access to electronic agents, their usage would allow a consumer with high haggling costs to behave as if she were a consumer with low haggling costs, i.e., by following small increments in her offers. Thus, it would be interesting to study how haggling would change in a population where some consumers use electronic agents, while others still haggle manually.

In summary, we believe that online haggling - while not replacing posted prices - will become a common element of online business transactions. Understanding the underlying theory and its potential application thus becomes of fundamental importance for online retailers and management scholars alike.⁴

8 References

Bakos, J. Y. "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, Vol. 43, No. 12, 1997, pp. 1676-1692.

Bertsimas, D., J. Hawkins, G. Perakis, "Optimal Bidding in On-Line Auctions over the Internet", 1st Annual INFORMS Revenue Management Section Conference

Chiang, J., "Competing Coupon Promotions and Category Sales," *Marketing Science*,

⁴The authors would like to thank the management team of the German name-your-own price retailer for providing their internal data, which is the empirical foundation of this study. We are also grateful to the associate editor and the two referees of Hamm and Terwiesch 2002, who helped us to better understand the consumer's problem. The first author gratefully acknowledges financial support from WeBI, the Wharton e-Business Initiative.

14 (Winter), 1995, pp. 105-122

Daniel, K. D., D. Hirshleifer, "A Theory of Costly Sequential Bidding", Working Paper at the Kellogg Graduate School of Management, July 1998

Hann, I.H., C. Terwiesch, "Measuring Frictional Cost of Online Transaction: The Case of a NYOP-Retailer," Working Paper, GSIA and Wharton, 2001.

Kennan, J. and Wilson, R., "Bargaining with Private Information," *Journal of Economic Literature*, Vol. 31, No. 1. (March), pp. 45-104, 1993.

Klemperer, P., "Auction Theory: A Guide to the Literature", *Journal of Economic Surveys*, 1999, Vol. 13 (3), pp. 227-286

Mailath, G.J. and Postlewaite, A., "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, 57, 1990a, pp. 351-367.

Mailath, G.J. and Postlewaite, A., "Workers versus Firms: Bargaining over a Firm's Value," *Review of Economic Studies*, 57, 1990b, pp. 369-380.

Narasimhan, C., "A Price Discrimination Theory of Coupons", *Marketing Science*, Vol. 3, No. 2 (Spring), pp. 128-147, 1984

Nash, J.F., "The Bargaining Problem," *Econometrica*, Vol. 18, No. 2. (April), pp. 155-162, 1950.

Osborne, M.J. and Rubinstein, A., "Bargaining and markets," San Diego: Academic Press, 1990.

Pinker, E. J., A. Seidmann, Y. Vakrat, "Managing Online Auctions: Current Business and Research Issues", Working Paper, June 2002, University of Rochester

Rubinstein, A., "Perfect Equilibrium in a Bargaining Model", *Econometrica*, Vol. 50, 1982, pp. 97-109.

Salop, S., Stiglitz, J. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, Vol. 44, 1976, pp. 493-510.

Shugan, S. M., "Editorial: Marketing Science, Models, Monopoly Models, and Why We Need Them", *Marketing Science*, Vol. 21, No. 3 (Summer), pp. 223-228, 2002

Stigler, G. J. “The Economics of Information,” *The Journal of Political Economy*, Vol. 69, No. 3, June 1961, pp. 213-225.

Sutton, J., “Non-Cooperative Bargaining Theory: An Introduction,” *The Review of Economic Studies*, Vol. 53, No. 5. (October), pp. 709-724, 1986.

Varian, H. R. “A Model of Sales,” *The American Economic Review*, Vol. 70, No. 4, September 1980, pp.651-659.

Van Damme, E., Selten R., and Winter E., “Alternating Bargaining with Smallest Money Unit,” *Games and Economic Behavior*, 4, pp. 132-144, 1990.

9 Appendix

Proof of Proposition 1: For the proof of the statements of Proposition 1 we need two additional results:

Lemma A1: For $n \in N$, $(n+1)v_{n-1}^* - nv_n^* \leq n(\beta_{n-1} - B) \leq nv_{n+1}^* - (n+1)v_n^*$.

Proof: We start with the inequality $(n+1)v_{n-1}^* - nv_n^* \leq n(\beta_{n-1} - B)$. Using (6) and (8), this is equivalent to $\sqrt{1+A+\beta_0-B} \times \left((n+1) \sqrt{1+\frac{(n-2)(n+1)}{2}} - n \sqrt{1+\frac{(n-1)(n+2)}{2}} \right) \leq 0$, which is satisfied for any $n \in N$, since the expression in brackets is non-positive and $1+\beta_0-B = \max(0, 1-B) \geq 0$.

Similarly, $n(\beta_{n-1} - B) \leq nv_{n+1}^* - (n+1)v_n^*$ is equivalent to $\beta_0 - B - (n^2 + n - 2) \leq \sqrt{1+A+\beta_0-B} \left(n \sqrt{1+\frac{n(n+3)}{2}} - (n+1) \sqrt{1+\frac{(n-1)(n+2)}{2}} \right)$ which is true for any $n \in N$, since the left-hand side here is non-positive and the right-hand side is non-negative. \square

Lemma A2: For $n \in N$, $v_n^* \leq v \leq v_{n+1}^*$,

$$L_n(v) = -(A+v) + \max_{v_{n-1}^* \leq u \leq v_n^*} ((v-u)(B+u) + L_{n-1}(u)). \quad (25)$$

$$v_{n+1}^* = \min \left(v \mid \max_{v_n^* \leq u \leq v} ((v-u)(B+u) + L_n(u)) \geq \max_{v_{n-1}^* \leq u \leq v_n^*} ((v-u)(B+u) + L_{n-1}(u)) \right). \quad (26)$$

Proof: Using (6), we establish after some algebra that

$$L_n(v) = -(A+v) + \max_{u \in R} ((v-u)(B+u) + L_{n-1}(u)), \quad n \in N \quad (27)$$

Introducing

$$\begin{aligned}
u_{n-1}^g(v) &= \arg \max_{u \in R} ((v-u)(B+u) + L_{n-1}(u)) = \frac{v + \beta_{n-1} - B}{2(1 - \alpha_{n-1})}, \\
u_{n-1}^l(v) &= \arg \max_{v_{n-1}^* \leq u \leq v_n^*} ((v-u)(B+u) + L_{n-1}(u)) = \begin{cases} v_{n-1}^*, & u_{n-1}^g \leq v_{n-1}^*, \\ u_{n-1}^g, & v_{n-1}^* \leq u_{n-1}^g \leq v_n^*, \\ v_n^*, & v_n^* \leq u_{n-1}^g, \end{cases} \quad (28)
\end{aligned}$$

for $n \in N$, we observe that in order for (25) to be valid, we need to show that $v_n^* \leq v \leq v_{n+1}^*$ implies $v_{n-1}^* \leq u_{n-1}^g(v) \leq v_n^*$. Using the result of Lemma A1 and (28), we obtain $v_{n-1}^* \leq u_{n-1}^g(v_n^*)$ and $u_{n-1}^g(v_{n+1}^*) \leq v_n^*$. Combining this with the observation that $\frac{\partial u_{n-1}^g(v)}{\partial v} = \frac{1}{2(1-\alpha_{n-1})} > 0$, we get (25). Finally, (26) follows from (7) and (25). \square

We start the proof of Proposition 1 from the statement in part a). We note that (8) implies that for \bar{n} defined in (9), $v_{\bar{n}}^* \leq v_{\max} < v_{\bar{n}+1}^*$. Below we demonstrate that the function (10) is a solution to (3) for $0 \leq v \leq v_{\max}$. Indeed, consider first the interval $0 \leq v \leq v_1^*$. For any v from this interval, we have $L^*(v) = L_0(v)$. Let us assume that $A > 0$, $B = 0$. Then, according to (6), $L_0(v) = 0$, and $\max_{0 \leq u \leq v} ((v-u)(B+u) + L_0(u)) = \frac{v^2}{4} \leq A + v$ for $v \leq v_1^*$, so that $\max(0, -(A+v) + \max_{0 \leq u \leq v} ((v-u)(B+u) + L_0(u))) = 0 = L_0(v)$. The same result is obtained for $A = 0$, $B < 1$. For $B \geq 1$, we get, according to (6) and (8), $\max_{0 \leq u \leq v} ((v-u)(B+u) + L_0(u)) = vB$, and, as before,

$\max(0, -(A+v) + \max_{0 \leq u \leq v} ((v-u)(B+u) + L_0(u))) = (B-1)v = L_0(v)$. Further, for $v_n^* \leq v \leq v_{n+1}^*$, $n = 1, \dots, \bar{n} - 1$,

$$\begin{aligned}
& \max_{0 \leq u \leq v} ((v-u)(B+u) + L^*(u)) \\
&= \max \left(\max_{0 \leq u \leq v_1^*} ((v-u)(B+u) + L_0(u)), \dots, \max_{v_n^* \leq u \leq v} ((v-u)(B+u) + L_n(u)) \right) \\
&= \max_{v_{n-1}^* \leq u \leq v_n^*} ((v-u)(B+u) + L_{n-1}(u)) \quad (29)
\end{aligned}$$

where we have used (26). Thus,

$$\begin{aligned}
& -(A+v) + \max_{0 \leq u \leq v} ((v-u)(B+u) + L^*(u)) \\
&= -(A+v) + \max_{v_{n-1}^* \leq u \leq v_n^*} ((v-u)(B+u) + L_{n-1}(u)) = L_n(v) = L^*(v) \quad (30)
\end{aligned}$$

for $v_n^* \leq v \leq v_{n+1}^*$. Similarly, for $v_n^* \leq v \leq v_{\max} < v_{n+1}^*$,

$$-(A+v) + \max_{0 \leq u \leq v} ((v-u)(B+u) + L^*(u)) = L_{\bar{n}}(v) = L^*(v). \quad (31)$$

Turning to part b), we first would like to express the optimal bidding sequence in the notation compatible with the definition of the recursive family $\{L_n(v)\}$. In particular, we would like to show that the optimal number of offers made by a consumer whose decision model is expressed by (3)-(5) is given by

$$n^* = \begin{cases} \bar{n}, & B < 1, \\ \bar{n} + 1, & B \geq 1, \end{cases} \quad (32)$$

and the optimal sequence of offers is given by

$$\mathbf{b}^* = \begin{cases} \{b_1^*, \dots, b_{\bar{n}}^*\}, & B < 1, \\ \{b_0^*, b_1^*, \dots, b_{\bar{n}}^*\}, & B \geq 1, \end{cases} \quad (33)$$

where

$$b_k^* = \frac{kv_{\max}}{\bar{n}+1} - \frac{k(\bar{n}-k+1)}{2} + (\beta_0 - B + 1) \left(1 - \frac{k}{\bar{n}+1}\right), k = 1, \dots, \bar{n}, b_0^* = 0. \quad (34)$$

In order to establish (32)-(34), we note that in the beginning of the haggling process, a consumer evaluates $L(v_{\max})$ and the optimal first offer is given by

$$b_{\bar{n}}^* = \arg \max_{v_{\bar{n}-1}^* \leq u \leq v_{\bar{n}}^*} ((v_{\max} - u)(B+u) + L_{\bar{n}-1}(u)) = \frac{v_{\max} + \beta_{\bar{n}-1} - B}{2(1 - \alpha_{\bar{n}-1})}, \quad (35)$$

where we have used (25) and the fact that $v_{\bar{n}-1}^* \leq \frac{v + \beta_{\bar{n}-1} - B}{2(1 - \alpha_{\bar{n}-1})} \leq v_{\bar{n}}^*$ for $v_{\bar{n}}^* \leq v \leq v_{\bar{n}+1}^*$.

Utilizing (6), we get

$$b_{\bar{n}}^* = \frac{v_{\max} + \beta_{\bar{n}-1} - B}{2(1 - \alpha_{\bar{n}-1})} = \frac{\bar{n}v_{\max} + \beta_0 - B - \frac{(\bar{n}-1)(\bar{n}+2)}{2}}{\bar{n}+1}. \quad (36)$$

If the first offer $b_{\bar{n}}^*$ is rejected, a consumer establishes the value of the next offer $b_{\bar{n}-1}^*$ by evaluating $L(b_{\bar{n}}^*)$. Since $v_{\bar{n}-1}^* \leq b_{\bar{n}}^* \leq v_{\bar{n}}^*$, $b_{\bar{n}-1}^* = \arg \max_{v_{\bar{n}-2}^* \leq u \leq v_{\bar{n}-1}^*} ((b_{\bar{n}}^* - u)(B+u) + L_{\bar{n}-2}(u)) = \frac{b_{\bar{n}}^* + \beta_{\bar{n}-2} - B}{2(1 - \alpha_{\bar{n}-2})} = \frac{(\bar{n}-1)b_{\bar{n}}^* + \beta_0 - B - \frac{(\bar{n}-2)(\bar{n}+1)}{2}}{\bar{n}}$. In the same fashion, $b_k^* = \frac{kb_{k+1}^* + \beta_0 - B - \frac{(k-1)(k+2)}{2}}{k+1}$, $k = 1, \dots, \bar{n} - 1$. Using this result recursively, we get

$$b_k^* = \frac{kv_{\max}}{\bar{n}+1} + k \left(\frac{1}{k(k+1)} + \dots + \frac{1}{\bar{n}(\bar{n}+1)} \right) (\beta_0 - B) - \left(\frac{\varepsilon_k}{k(k+1)} + \dots + \frac{\varepsilon_{\bar{n}}}{\bar{n}(\bar{n}+1)} \right), \quad (37)$$

where $\varepsilon_k = \frac{(k-1)(k+2)}{2}$. Since $\frac{1}{k(k+1)} = \frac{1}{k} - \frac{1}{k+1}$ and $\frac{\varepsilon_k}{k(k+1)} = \frac{1}{2} + \frac{1}{k+1} - \frac{1}{k}$, (37) becomes

$$b_k^* = \frac{kv_{\max}}{\bar{n} + 1} - \frac{k(\bar{n} - k + 1)}{2} + (\beta_0 - B + 1) \left(1 - \frac{k}{\bar{n} + 1}\right). \quad (38)$$

We note that $v_0^* = 0 \leq b_1^* \leq v_1^*$. Thus, if $B < 1$, $L(b_1^*) = 0$, which implies that if offer b_1^* is rejected, a consumer terminates the haggling process. On the other hand, if $B \geq 1$, $L(b_1^*) = (B - 1)b_1^* > 0$, and an extra offer $b_0^* = 0$ is placed.

Finally, the above results are easily converted into the original notation used in (1) and (2). In particular, the parameter z which defines the optimal number of offers made by a consumer characterized by the quadruplet $(c, R_{\max}, R_{\min}, r)$ can be re-expressed as

$$z = \begin{cases} \sqrt{\frac{R_{\max} - R_{\min}}{c}}, & r > R_{\max} + c, \\ \sqrt{\frac{R_{\max} - R_{\min}}{c}} - \sqrt{\frac{c + R_{\max} - r}{c}}, & r \leq R_{\max} + c, \end{cases} \quad (39)$$

To avoid degenerate cases, we limit ourselves to those cases where the consumer engages in haggling, which requires $r > R_{\min} + c$. This, in turn, implies non-negativity of z in (39). The values of optimal offers are then given by $x_k^* = R^* - cb_{\bar{n}-k+1}^*$, $k = 1, \dots, \bar{n} + 1$, which reduces to (12).

Proof of Proposition 2: We start by observing that (21) is a generalization of (17) for the case of a non-homogeneous consumer market: the generated profit is a weighted sum of profit values generated by each consumer group, and, for a given threshold value, the profit contribution of a particular group is equal to the smallest offer exceeding the threshold minus the wholesale price w . According to (18), $x_{j\bar{k}_j(T)}^*$ is a non-decreasing function of T , and, therefore, $T = \hat{x}_l$ dominates in terms of generated profits any threshold from the interval $(\hat{x}_{l-1}, \hat{x}_l]$ for all $l = i(w), \dots, I$. Thus, the optimal threshold price should be equal to one of the highest consumer offers \hat{x}_i . Comparison between $\Pi_n(T = \hat{x}_j)$ for $j = i(w), \dots, I$ leads to (22).

Proof of Proposition 3: The optimal value of the fixed-price profits Π_f can be computed as

$$\Pi_f = \max(r_1 - w, p_2(r_2 - w)). \quad (40)$$

On the other hand, if haggling is used, the value of the frictional cost for consumers belonging to the second group is such that they only make a single bid at $\frac{r_2}{2}$. Under

this conditions, the best haggling profit is obtained by setting the haggling threshold at $T = \hat{x}_1$, where \hat{x}_1 is the highest bid made by consumer of the first group. The optimal value of the haggling profit is

$$\Pi_h = p_1 (\hat{x}_1 - w) + p_2 \left(\frac{r_2}{2} - w \right). \quad (41)$$

For the analysis below, it is convenient to define $z(c_1) = \sqrt{\frac{r_1}{c_1}} - 1$ and

$$x(c_1) = \frac{c_1}{2} \left(\frac{1}{2} \left(\sqrt{1 + 8z^2(c_1)} - 1 \right) - 1 \right) + (r_1 - c_1) \left(1 - \frac{2}{\sqrt{1 + 8z^2(c_1)} - 1} \right) \quad (42)$$

Note that $x(c_1)$ is a continuous function of c_1 , and, according to the results of Proposition 1, for $c_1 > 0$,

$$x(c_1) < \hat{x}_1 < r_1, \quad (43)$$

while

$$\lim_{c_1 \rightarrow 0} x(c_1) = r_1. \quad (44)$$

Given the continuity of $x(c_1)$ and the limit value (42), there exist $\bar{c} > 0$ such that $p_1(x(c_1) - w) + p_2(\frac{r_2}{2} - w) > r_1 - w$ for all $0 < c_1 < \bar{c}$. Now, since our assumption $r_1 - w > \frac{p_2 r_2}{2p_1}$ is equivalent to $p_1(r_1 - w) + p_2(\frac{r_2}{2} - w) > p_2(r_2 - w)$, by the same argument as above, there exists $\hat{c} > 0$ such that $p_1(x(c_1) - w) + p_2(\frac{r_2}{2} - w) > p_2(r_2 - w)$ for all $0 < c_1 < \hat{c}$. Selecting $c_1^* = \min(\bar{c}, \hat{c})$, and using (43), we get the statement of the Proposition.

Proof of Proposition 4: a) Due to monotone relation between \bar{n} and z , as indicated by (9), it is sufficient to prove that the statement of the Proposition applies to z . The monotonicity of z with respect to r and R_{\min} directly follows from (39). Further, considering the only non-trivial case of $r \leq R_{\max} + c$, we get $\frac{\partial z}{\partial c} = \sqrt{\frac{R_{\max} - R_{\min}}{c + R_{\max} - r}} \frac{1}{2c^{\frac{3}{2}}} \times \left(\frac{R_{\max} - r}{\sqrt{R_{\max} - R_{\min}}} - \sqrt{c + R_{\max} - r} \right) < 0$, since the expression in the brackets is a decreasing function of c and it is negative for smallest possible values of c no matter whether $r \leq R_{\max}$ or $r \geq R_{\max}$. Clearly, as (11) indicates, for $r > R_{\max} + c$ a consumer places at least 1 offer (the highest one equal to R_{\max}). If, however, $r \leq R_{\max} + c$, a consumer participates in the

auction if and only if $\bar{n} \geq 1$. This last condition is equivalent, according to (9), to $z \geq 1$. Using (39), we have $\sqrt{\frac{R_{\max}-R_{\min}}{c}} - \sqrt{\frac{c+R_{\max}-r}{c}} \geq 1 \Leftrightarrow c \leq \frac{(r-R_{\min})^2}{2(R_{\max}-R_{\min})}$.

b) Using (9) and (39), we observe that $\lim_{c \rightarrow 0} \bar{n}(c) = +\infty$, and $\lim_{c \rightarrow 0} c\bar{n}(c) = 0$. Thus, $\hat{x}(0) = \min(r, R_{\max})$. Considering the only non-trivial case of $r \leq R_{\max} + c$, we get

$$\begin{aligned} \hat{x}(0) - \hat{x}(c) &\geq \frac{1}{\bar{n} + 1} \left(r - R_{\min} - c - \frac{c\bar{n}(\bar{n} + 1)}{2} \right) \geq \frac{r - R_{\min} - c - cz^2}{\bar{n} + 1} \\ &= \frac{2\sqrt{c + R_{\max} - r} (\sqrt{R_{\max} - R_{\min}} - \sqrt{c + R_{\max} - r})}{\bar{n} + 1} > 0, \end{aligned} \quad (45)$$

since, by assumption, $r > R_{\min} + c$.

For $c_{i+1} < c < c_i$, $i \in N$, the optimal number of offers $n^* = \bar{n}$ remains unchanged, and according to (12), $\frac{\partial \hat{x}(c)}{\partial c} = \frac{\partial x_{\bar{n}}^*}{\partial c} = \frac{\bar{n}}{2} - \frac{\bar{n}}{\bar{n}+1} = \frac{\bar{n}(\bar{n}-1)}{2(\bar{n}+1)} \geq 0$. In the neighborhood of c_i we get $\lim_{c \rightarrow c_i-0} \hat{x}(c) = R_{\min} + \frac{c_i(i+1)}{2} + (r - c_i - R_{\min}) \left(\frac{i+1}{i+2} \right)$, and $\hat{x}(c_i) = R_{\min} + \frac{ic_i}{2} + (r - c_i - R_{\min}) \left(\frac{i}{i+1} \right)$, so that $\lim_{c \rightarrow c_i-0} \hat{x}(c) - \hat{x}(c_i) = \frac{c_i}{2} + \frac{(r-c_i-R_{\min})}{(i+2)(i+1)} > 0$.

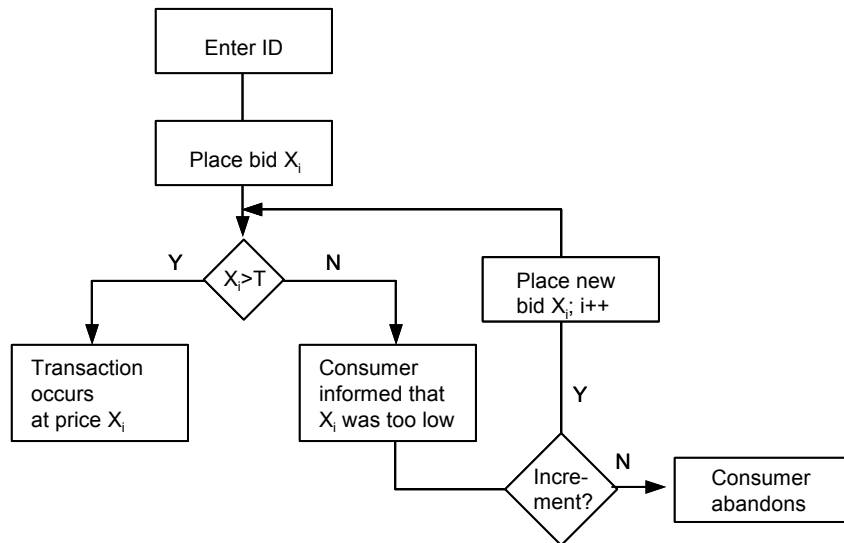


Figure 1: Description of consumer model

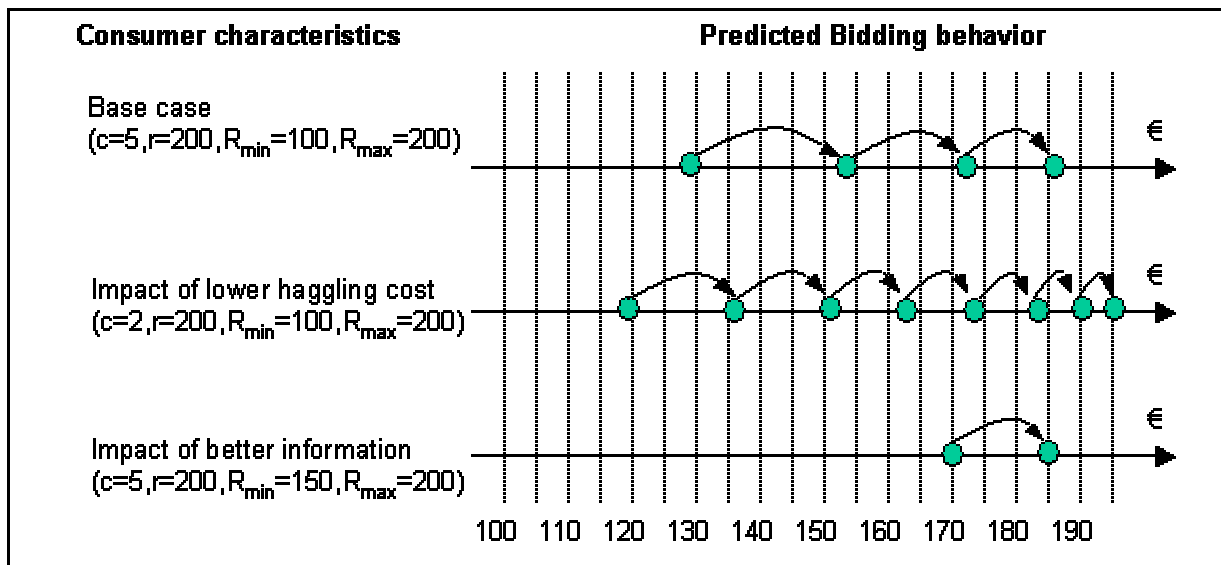


Figure 2: Illustration of consumer haggling model




Figure 3: Estimating consumer attributes based on observed offers

Dependent variable in validation regression	Length of haggling sequences		Constant increment model	Population typical increments	Consumer search model (Proposition 1)
Increments	N*≥4	$Adj.R^2$	47.5%	17.5%	66.4%
		β_l	1.00	1.03	1.01
Increments	N*≥7	$Adj.R^2$	21%	16.5%	52.4%
		β_l	1.00	0.89	1.01
Offers	N*≥4	$Adj.R^2$	95.2%	74.9%	97.5%
		β_l	0.95	0.62	1.02
Offers	N*≥7	$Adj.R^2$	93.1%	77.2%	97.0%
		β_l	0.88	0.69	0.94

Table 1: Validation of the consumer model based on a regression model
 $Actual_i = \beta_0 + \beta_l Predicted_i + \varepsilon_i$




Figure 4: Model validation by comparing actual bids with predicted

Product	Actual Threshold	Actual Profit	Static Threshold	Profit under static Threshold	Ex-post optimal Threshold	Ex-post optimal Profit
DVD-Player	208	9.40	247	10.23	231	13.07
PDA	193	4.96	220	8.54	237	9.58
CD-Rewriter	144	5.41	160	6.02	153	6.55

Table 2. Profit per consumer (in €) for different static pricing rules.

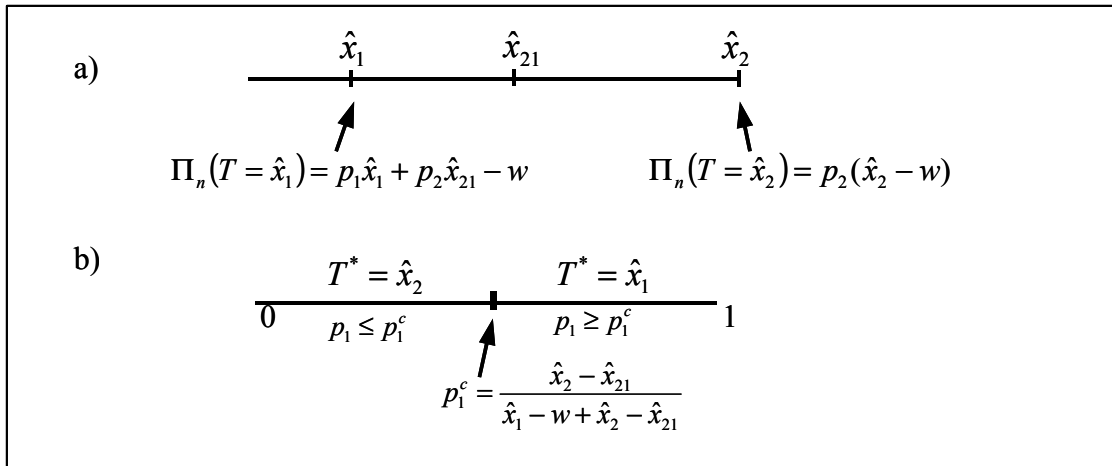


Figure 5: Non-homogeneous consumer market consisting of 2 groups: (a) expected profits (per consumer) for two competing threshold values; (b) optimal threshold as a function of the probability that a consumer belongs to the 1st group.