

Measuring the Frictional Costs of Online Transactions: The Case of a Name-Your-Own-Price Channel

Il-Horn Hann • Christian Terwiesch

Marshall School of Business, University of Southern California, Los Angeles, California 90025

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

hann@marshall.usc.edu • terwiesch@wharton.upenn.edu

We study the offers submitted by consumers to a large Name-Your-Own-Price (NYOP) online retailer. A distinctive feature of this retailer is that it allows consumers to repeatedly submit offers on one and the same product. While consumers could identify the threshold price (the minimum price for which the retailer is willing to sell) by incrementing their offer in small steps in each consecutive round, such a strategy would require them to go through many additional online transactions. We define frictional cost as the disutility that the consumer experiences when conducting an online transaction, such as submitting an offer. Thus, in our setting, consumers trade off a direct financial value (lower price) for frictional costs. Based on a consumer choice model capturing this trade-off, we use the observed consumer behavior to reconstruct the frictional cost parameters for the consumers in our sample. We show that, perhaps contrary to the general wisdom, frictional costs in electronic markets are substantial, with median values ranging from EUR 3.54 for a portable digital music player (MP3) to EUR 6.08 for a personal digital assistant (PDA). We find that consumers who have gathered experience with the NYOP channel in previous transactions exhibit lower frictional costs than consumers who use the channel for the first time. Surprisingly, sociodemographic variables do not help to explain the variation in frictional costs. (*Electronic Markets; Frictional Costs; Name-Your-Own-Price Channel; Online Haggling; Price Dispersion*)

1. Introduction

The role of search costs in the online transactions of consumers has attracted much attention in recent e-commerce research. Most of this interest stems from a big promise of the Internet: When shopping on the World Wide Web, the costs associated with comparing products and prices should substantially decrease relative to shopping in brick-and-mortar stores. Facing better-informed consumers, online retailers selling homogeneous goods would have to engage in head-to-head price competition, leading to substantially lower prices for consumers.

Empirical research has not confirmed this prediction, however. Researchers have consistently observed price dispersion, both for homogenous goods such as books and music CDs and for differentiated services such as airline tickets (Bailey 1998, Brynjolfsson and Smith 2000, Clay et al. 2001, Clemons et al. 2002, Png et al. 2000). The potential reasons for the observed price dispersion fall into three categories: (1) product differentiation, (2) brand effects, and (3) frictional costs. Product differentiation allows retailers to mitigate price competition by segmenting consumers according to their willingness to pay for customer

service. Brand effects capture perceived or real quality differences, allowing retailers with a valuable brand to extract additional rents. Finally, while search costs on the Internet may be lower than in brick-and-mortar environments, online shopping still requires time and effort from the consumer. We refer to this part of consumer search cost as frictional cost. This includes the disutility of investing time and effort for interacting with a website (such as keying in payment information) and the disutility of interacting with various user interfaces.

Amazon's attempt to patent its one-click shopping technology illustrates the importance of frictional costs and their potential role in creating a competitive advantage in e-markets. Consider a customer at a shopbot who has identified a product of interest to her, which is available on Amazon's website (amazon.com) and Barnes & Noble's online unit barnesandnoble.com (bn.com). The rational, utility-maximizing consumer considers the total cost of purchase and compares the Amazon price, combined with the frictional cost of purchasing the product from Amazon (which corresponds to a single mouse click), with the price at barnesandnoble.com and the frictional cost she would incur there. One-click shopping allows Amazon to reduce the frictional cost of its website, making the overall cost of obtaining the product from Amazon (defined as the purchase price plus frictional cost) lower than the cost of obtaining the product from its competitors. Therefore, the price premium that Amazon is able to charge in a heterogeneous consumer market increases with its ability to decrease the frictional cost of shopping. This competitive advantage will erode if competitors are also able to reduce their respective frictional costs. It has been speculated that this was Amazon's motive behind enforcing this technology patent.¹

¹ Amazon was awarded U.S. patent (5,960,411) for its one-click shopping technology on September 29, 1999. On October 21, 1999, Amazon.com sued barnesandnoble.com for allegedly infringing on the technology with its "Express Lane" feature, which refers to single-action ordering of items in a client/server environment such as the Internet. In its case, Amazon documented the development of the "shopping cart model" purchase system for e-commerce purchasing events. In December 1999, a U.S. District Court judge granted injunctive relief barring barnesandnoble.com from using

The objective of our research is to demonstrate the significance of frictional costs by quantifying them in the case of a German NYOP online retailer. Most implementations of such reverse buying mechanisms allow consumers to submit only one offer for a given product. This policy was first established by industry leader Priceline.com, and was initially also implemented by the intermediary examined in this study. Subsequently, our research site changed its policy and allowed consumers to increment their offer if their earlier offer has failed.

The NYOP retailer we study is an intermediary between a wholesaler and consumers. The wholesaler sets a wholesale price and the intermediary sets a threshold price. The consumer submits an offer that is accepted if it exceeds the threshold price. From the perspective of the intermediary, there are two sources that contribute to profit. First, every successful transaction provides the intermediary with a basic profit (threshold price minus wholesale price). Second, the intermediary also obtains an information rent, the spread between the submitted offer and the threshold price. The retailer underlying this study chose a threshold price at the wholesale price and, thus, relied on the information rent as a source of profit.

If online transactions were frictionless, consumers could identify the threshold price following an "epsilon strategy," i.e., by incrementing their offer in small steps in each round, leaving little or no information rent to the seller. However, would consumers really be willing to incur the effort of repeatedly keying in an offer and waiting for feedback from the retailer if their offer was successful, just to save a penny? Or, would the presence of frictional costs induce them to increment in larger steps, thereby leaving positive information rents to the retailer? Our research setting, in which consumers need to trade off direct financial value (lower price) with the frictional cost of an additional offer, provides an interesting empirical setting to estimate the magnitude of frictional costs in online transactions. This is described in §3.

what Amazon.com termed its "copycat version." A federal court overturned the preliminary injunction in February 2001. In March 2002, the parties settled the infringement lawsuit without disclosing the detailed terms of the settlement.

We present an economic model of consumer behavior that captures the trade-off between the total frictional costs a consumer incurs, which is minimized if the consumer submits only one offer, and the desire for paying a price as close to the threshold price as possible, which can be realized if the consumer frequently increments her offers in small steps (§5). We use the transaction data from our research site to reconstruct the frictional cost parameters for the consumers in our sample. This allows us to make the following three contributions. First, we measure frictional cost using the approach previously described. We show that they are significant in absolute terms and, thus, sufficient for avoiding a complete erosion of the information rent collected by the retailer even if consumers are allowed to repeatedly increment offers (§6). Second, we identify consumer experience with the NYOP channel as the main driver of frictional cost. The frictional cost of submitting an offer is lower for consumers who have previously placed offers on other NYOP products, thereby exhibiting a learning curve pattern. Surprisingly, sociodemographic variables, including income and education, have no significant impact on frictional cost. Third, the fact that our research site initially implemented the “one-shot” model and then later changed to the “iterative/haggling” model currently in place, provides us with a unique quasi-natural experiment.² We contrast the information rent before and after the policy change, and find that the new model leads to lower information rents for a successful offer and does not significantly increase the number of successful offers.

2. Previous Literature

Based on the economic theory of search, early work has characterized the nature of competition in e-commerce as head-to-head competition among retailers (Bakos 1991, 1997).³ However, empirical studies have found ample evidence of price dispersion

in e-markets. Price dispersion has been studied for homogeneous goods such as books, CDs, and software (Bailey 1998, Brynjolfsson and Smith 2000, Clay et al. 2001) and for well-specified services such as air travel (Clemons et al. 2002). These studies consistently find dispersion in prices. Prices vary from 19% across online travel agents for qualitatively identical tickets, 25% for identical CDs, and 33% for identical books.

Various reasons have been offered to explain the observed price dispersion. One explanation is vertical product differentiation among retailers. The economics literature has shown that competing vendors will choose different price and quality levels to soften price competition (Gabszewicz and Thisse 1980, Shaked and Sutton 1982). For example, physical stores can charge a premium for ambience and knowledgeable sales people. Similarly, online retailers can focus on service-oriented customers and excel with targeted offerings, speedy delivery, and great customer service (Clemons et al. 2002). Further, brand effects often mirror such (perceived or real) quality differences. Smith and Brynjolfsson (2001) analyze the last click-through of book shoppers on a shopbot. Even for these price-sensitive and well-informed customers, more than 50% choose not to visit the lowest price retailer. Similarly, Clay et al. (2001) find that Amazon.com can charge higher prices than other established rivals such as Borders Online (borders.com) or barnesandnoble.com.

Another possible reason for the observed price dispersion is that the true cost of search to the consumer includes a certain amount of frictional cost. We define *frictional cost* as the disutility related to learning to navigate through websites, the disutility of keying in order and payment information, the cognitive costs of comparing different offerings, and the opportunity cost of time for the online transaction. A more operational definition of frictional cost for our research setting is provided below.

The Amazon example in the introduction shows how frictional costs impact switching costs and thereby competition. Consider a consumer who is on the site of retailer 1. While frictional costs are specific to each website, *switching costs* is the difference between the sum of frictional costs for a complete transaction incurred at retailer 1 and the

² To have this as a “clean” natural experiment, we have to assume that the different periods did not attract different types of consumers.

³ A more detailed discussion on the theory of search and its effect on pricing can be found in an online appendix at mansci.pubs.informs.org/ecompanion.html.

sum of frictional costs for a complete transaction at retailer 2, plus any additional search costs for identifying retailer 2. Websites such as Amazon.com, which create lower frictional costs through one-click shopping, thus make it more attractive to the consumer to purchase at retailer 1, even when search engines or shopbots make the search cost of identifying retailer 2 negligibly small.

Research by Johnson et al. (2002b) provides indirect support for the existence of frictional costs. Based on a 12-month-long panel data set covering 10,000 households from Media Metrix, they find that users, on average, visit only 1.1 different book sites, 1.2 different CD sites, and 1.8 different travel sites, suggesting that the average consumer engages in little search activity. One potential explanation for this observation is that consumers prefer to save the frictional cost associated with additional online transactions over realizing lower prices for the items purchased. Johnson et al. (2000b) also find that search propensity does not change over time for book and CD sites. For travel sites, the search propensity actually decreases. This is consistent with the fact that travel sites require more input from the user and that their offerings require more cognitive effort to make a purchasing decision when compared to book or CD sites.⁴

While past research has used frictional costs as an explanation for the observed prices on the retailer side and the lack of search behavior on the consumer side, no prior research has been undertaken to directly measure the frictional costs and analyze the sources of these frictional costs. This is the objective of the present study.

3. Research Setting

Our research site is an online intermediary in Germany that has adopted the reverse buying model and has enjoyed a dominant position among NYOP online retailers. Since its inception, the firm has

⁴ Books and CDs are usually well defined and easily specified. In contrast, a flight ticket can be much more complex, as it is typically offered by several airlines at different times and with a different number of stops. This requires a buyer to specify departure and arrival city, date, time, airline, and class and seating preferences on all segments.

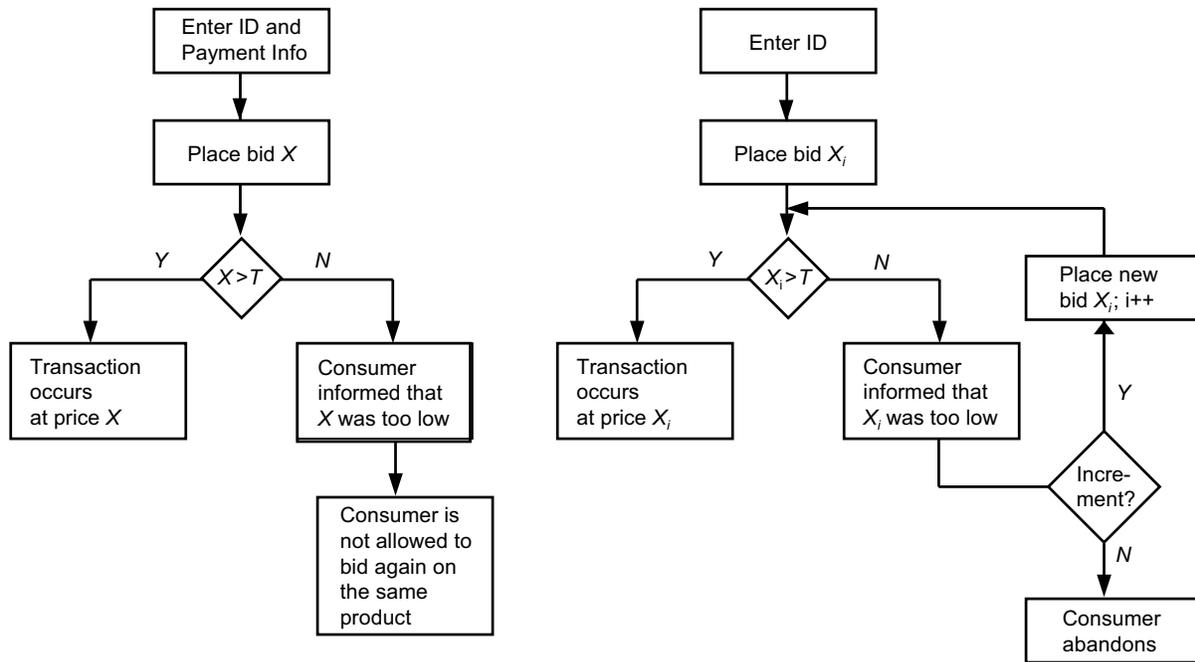
expanded its offerings from airline tickets to computers, consumer electronics, videos, DVDs, and music CDs. While the reverse buying mechanism is used in business-to-consumer (B2C) settings and in some business-to-business (B2B) settings, the firm is exclusively serving end consumers. A detailed discussion of the differences of B2B versus B2C markets can be found in the online appendix.

Business Rules: Old and New Policy

When our research site launched its service, it had implemented a business policy similar to most U.S.-based reverse buying sites, popularized by industry leader Priceline.com. We will refer to this policy as the "old policy." The distinctive element of this policy is that it allows a consumer to submit only one offer for a given product. Figure 1 describes in detail how a consumer interacts with the retailer under the old policy. The interaction begins with the consumer registering at the site. Registration involves entering personal information such as name, address, e-mail account, and credit card information. Subsequently, a consumer offers a price for a desired product. The retailer compares this price with a threshold price that is unknown to the consumer. If the price offered by the consumer is above the threshold price, the transaction occurs at the value of the consumer's offer. For example, if the consumer offers a price of EUR 120 for a VCR and the site holds a threshold price of EUR 100, the consumer would pay EUR 120. We refer to the difference between the offered price (EUR 120) and threshold price (EUR 100) as the retailer's information rent. If the offer is below the threshold price, the consumer is informed that her offer is not accepted and the consumer is prohibited from submitting further offers for the same product.

While many industry observers have praised this business model as the ultimate consumer-driven commerce, others have noticed that these settings are the only market where consumers could pay more (EUR 20 in the above case) than what is asked for by the seller. After being notified that the submitted offer was successful, the consumer knows that she is likely to "have left money on the table" and, ideally, would like to renegotiate the price. As the consumer anticipates this effect, many, especially price-sensitive

Figure 1 Old vs. New Policy



consumers, are likely to either suggest extremely low prices, so that a transaction does often not occur, or abstain from submitting offers completely.

To at least partially overcome this problem, our research site decided to revise the old policy. The change was implemented after about six months of conducting business under the old policy. Under the new policy, customers are allowed to increment rejected offers. This is illustrated in Figure 1. The interaction begins with the consumer registering on the site (entering user ID and password, potentially registering as a new user). She can then suggest a price for one of the featured products. Similar to the old policy, the retailer then compares the offered price with the threshold price. After about five minutes, the consumer is notified whether or not her offer was successful (i.e., sufficiently high). The distinctive feature of the new policy is that a consumer who is not successful is allowed to increment her offer by submitting a second (third, fourth, and so on) offer. This “haggling” between the consumer and the retailer ends if the consumer decides to no longer increment her offer, thereby terminating the haggling process, or if the (incremented) offer is successful.

Hypotheses Development

The history of successful and unsuccessful offers at our research site provides us with a unique opportunity to quantify the frictional costs previously discussed. Given the specific nature of the online transactions faced by consumers in the present context, frictional costs result from (1) the opportunity cost of time required for the website activities, (2) cognitive efforts to navigate through the website, (3) cognitive efforts to key in offers, and (4) the disutility of waiting for the feedback from the retailer. We do not consider the time and effort required for registration as a part of the frictional cost.⁵

Consumers differ with respect to the frictional costs outlined above along two dimensions. First, consumers differ in the amount of time and effort they incur when making an offer at the NYOP retailer. Second, consumers differ in their valuation of time and effort relative to potential price savings. The process of submitting an offer will not be perceived as equally

⁵ The need to register as well as the frictional cost itself results in a self-selected sample. Hence, our findings present a lower bound for the frictional cost of the overall population.

difficult by all consumers. There exists an extensive literature on learning on the individual, group, and organizational level. These studies indicate a strong relationship between learning resulting from experience and increases in efficiency (Argote 1993). This effect has been observed in various settings, ranging from manufacturing to human-computer interaction (Card et al. 1983). More recently, Johnson et al. (2002a) discuss how performance in website-specific browsing tasks improves with practice. After a consumer has repeatedly interacted with our research site, thereby having submitted offers for numerous products before, we expect the consumer to become more efficient in this process. Hence, in the context of the NYOP retailer, we expect her frictional cost to decline.

HYPOTHESIS 1. Consumers with a higher level of experience are expected to exhibit lower frictional costs in their offering sequences.

Prior research suggests that learning can be disrupted if the task to be learned about is subject to change (Argote 1993). Thus, changes in site design should theoretically disturb the learning process. As the site design underlying this research remained constant during the period of our data collection for the "new policy," we cannot empirically test this prediction in this study. Thus, future research is needed to explore the effect of disruption on the consumer's learning process. Another variable that would be interesting to study is the degree of consumer online literacy and the technical details of their Internet access. We could not obtain a sufficient amount of data for these constructs to warrant including them in our estimation, however. Thus, we also have to leave their effect on frictional cost as a subject for future research.

More generally, the time and effort related to making an offer will depend on the consumer's cognitive abilities. Economic theory has motivated equilibria in price dispersion in terms of differences in information processing capabilities across people (e.g., Salop and Stiglitz 1976). Consumers with high cognitive abilities will learn to navigate through a website with more ease and, hence, at lower frictional costs than consumers with low cognitive abilities. In general, cognitive abilities are attained through education.

HYPOTHESIS 2. Consumers with a higher level of education are expected to exhibit lower frictional costs in their offering sequences.

Consider how consumers differ with respect to their valuation of time and effort relative to potential price savings. Holding the amount of effort related to making an offer constant, affluent consumers are less likely to invest effort in the form of frictional cost to achieve a price saving. This reflects their higher opportunity cost of time. In contrast, a less affluent consumer is more likely to engage in multiple rounds within an offer sequence of the haggling process outlined above, as she would be less willing to leave an information rent to the retailer.

HYPOTHESIS 3A. Consumers with a higher level of income are expected to exhibit higher frictional costs in their offering sequences.

A similar argument can be made with respect to those consumers who face a significant opportunity cost of time. Specifically, we expect self-employed consumers, who have the opportunity of increasing their income if they dedicate a larger amount of time to their work, to have a higher frictional cost compared to consumers employed by someone else or currently unemployed.

HYPOTHESIS 3B. Consumers who are self-employed are expected to exhibit higher frictional costs in their offering sequences.

While Hypotheses 1, 2, 3a, and 3b focus on the drivers of frictional costs, thereby looking at variables specific to the consumer, Hypothesis 4 focuses on the information rent, a variable of great interest to the NYOP retailer. Starting at midnight of July 20, 2000, our research site introduced the new policy described above. This transition from old policy to new policy provides us with a unique quasi-natural experiment. Specifically, we are interested in how the introduction of the new policy has impacted the information rents collected by the retailer. The new policy provides the consumer with the option to increment a rejected offer. Rejected offers in this setting can be seen as valuable information. More information acquisition in the form of rejected offers reduces the retailer's information rent.

HYPOTHESIS 4. Information rents will decrease when consumers are allowed to increment unsuccessful offers.

This type of information acquisition requires the consumer to invest effort. In our setting, we labeled this information acquisition effort as frictional costs. Hence, we expect that, for the new policy, frictional costs will prevent a complete erosion of information rents.

4. Data and Methodology

For the purpose of this study, our research site provided us with a complete history of submitted offers, the threshold price, and the cost of goods for specific products.⁶ In addition, we obtained the corresponding customer identification, names, and addresses. The unique customer identification allows us to link the offers a consumer makes on a given product into a sequence of offers for this consumer and this product. Moreover, by being able to recognize customers across different products, we can measure how frequently the customer has previously interacted with our research site, which provides our measure of consumer experience. We also obtained education, employment, and income data from Claritas (Germany), a provider of marketing information. Based on the consumer addresses, we obtained aggregate sociodemographic data for the corresponding local community of every consumer. Local communities are defined at the subzip code level, with the average size of a community being 7,850 individuals living in 3,733 households.

For this study, we chose data for three specific product categories to measure frictional cost. The product categories we chose had to fulfill two criteria. First, they had to be a consumer durable product as opposed to a hotel reservation or an airline ticket. The consumer's valuation for the latter is likely to change, even within a relatively short period of time, which would make it difficult to distinguish between a consumer's frictional cost and some external events influencing the valuation of the good. In contrast, the consumer's valuation of a durable good is not likely

to change within the duration of a sequence of submitted offers, which typically occurred in fewer than two days. Second, the product categories had to have attracted a number of consumers sufficiently large for econometric analysis. Based on these two criteria, we decided to focus on PDAs, optical storage devices (CD-RW and DVD drives), and MP3 players. In addition to these three product categories, we obtained data for products that were offered under the old and new policy. This will be the basis for the quasi-natural experiment and the test of Hypothesis 4.

One of the products we studied as part of the PDA category was the Palm Pilot M100. For this specific product, our sample includes 2,507 offers made by 1,210 different consumers. As indicated by Table 1, from the 1,210 offers consumers placed in their first round, 52 offers were successful, i.e., exceeded the threshold price. From the remaining, unsuccessful offers, 629 consumers immediately abandoned, thus, abstained from incrementing their offer. The remaining 529 suggested a second, higher price, of which in turn 15 were successful. As indicated by Table 1, some consumers incremented their offer many more times, although few actually purchased after round 4. Table 2 provides the corresponding average increments.

Figure 2a (left) summarizes the distribution of the initially suggested prices. An overproportionally large number of offers are made in "round" numbers, resulting in a "lumpy" graph. For example, we observe many offers overproportionally of exactly EUR 100.⁷ The 529 offers in the second round were distributed according to Figure 2b (right). Figure 2b looks remarkably similar to Figure 2a, with two important differences. First, the graph is "thinner" as it includes only consumers with their second offer, thus, not including those who were successful in the first round or have decided to exit. Second, it is shifted upwards, as the average offer in the second round (EUR 141.52) is significantly higher compared to the average offer in the first round (EUR 117.68). As discussed above, the NYOP retailer collects an information rent from each consumer who has submitted an offer exceeding the threshold price. In the

⁶ Given that our analysis is only based on submitted offers, consumers potentially self-select into our sample.

⁷ At the time of our study, EUR 1 corresponded roughly to US \$0.9.

Table 1 Offers Over Rounds Within an Offer Sequence for the Palm Pilot M100

Round	Number of offers	Number of successful offers
1	1,210	52
2	529	15
3	280	10
4	160	2
5	113	1
6	64	1
7	44	1
8	34	1
>8	73	2

Table 2 Offering Increments for the Palm Pilot M100

Length of sequence	Median increase from 1 to 2	Median increase from 2 to 3	Median increase from 3 to 4	Median increase from 4 to 5
3	25.56	15.34		
4	20.46	11.25	7.92	
5	25.56	6.65	5.37	5.12

case of Figures 2a and 2b, the information rent corresponds to the shaded area between the graph of submitted offers and the threshold price. As we can see already from the descriptive data in Figures 2a and 2b, some consumers provide the retailer with a substantial information rent.

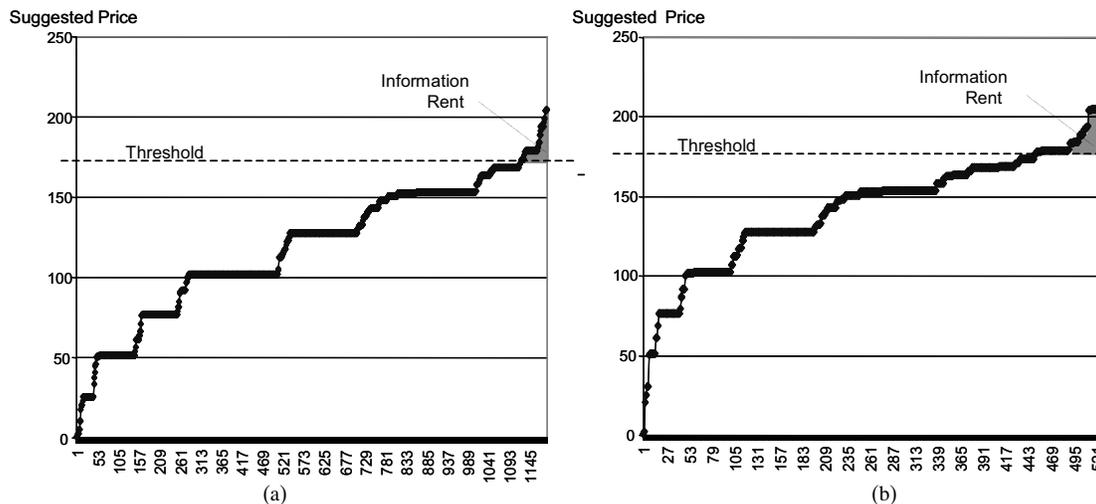
5. Decision Model and Analysis

Under the new policy, a consumer considering to submit an offer faces the following decision situation. If the offer she submits in the current round is *lower* than the unknown threshold price, she incurs frictional costs, but does not receive any direct reward. However, given that the new policy allows for resubmission of an incremented offer, she can learn from

her unsuccessful offer and enter a new round of the same “haggling game,” just with better information. She might also decide to abstain from any further haggling. If the offer she submits is *higher* than the threshold price, the consumer incurs frictional cost and then receives 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 information rent to the retailer.

Our model of consumer haggling is built using a rational consumer framework. We also assume that the consumer’s utility is linear in transaction price. In this setting, a consumer should value a 5-Euro saving when purchasing a 20-Euro music CD equally high as a 5-Euro saving when buying a 1,000-Euro personal computer (PC). From a behavioral perspective, consumers might view a 5-Euro saving on a 20-Euro product as more significant than a 5-Euro saving on a 1,000-Euro product. Thus, in our setting, consumers

Figure 2 Offers for the Palm Pilot M100 in the First and Second Round



interested in purchasing an expensive product might be willing to leave the retailer with a—in absolute terms—higher information rent. On the other hand, one could also argue that consumers seek a “better” price for high-priced items than for low-priced items, i.e., displaying lower frictional costs for more expensive goods. Such a behavioral element, while interesting in its own right, is difficult to combine with a model of rational consumer choice.

Dynamic Choice Model

Define c as the frictional cost for submitting an offer. We expect the consumer’s frictional cost to be constant over the course of a haggling process and independent of the submitted offers. Thus, while we consider learning effects across haggling processes (products), we do not account for learning effects within the sequence of offers for a given product. For example, we do capture the experience of a consumer who has submitted offers for a DVD player and a PC (or any other product offered by the NYOP retailer) and now submits an offer for a Palm Pilot. In contrast, we do not capture the learning effects that the same consumer would experience by moving from the third to the fourth offer for the Palm Pilot. More general formulations of our model are possible, but would require the introduction of additional parameters.

Without any knowledge about the supplier’s decision-making process, a consumer initially assumes that the threshold price is a random variable distributed over an interval from \underline{R} to \bar{R} . This interval describes the consumer’s subjective beliefs about the threshold, assuming that consumers know their \underline{R} and \bar{R} with certainty. The true threshold, which is only known to the retailer, might or might not be in the interval $[\underline{R}, \bar{R}]$. We assume that the consumer perceives the threshold price held by the retailer as constant over the course of the haggling process. The information the retailer provided to the consumer via its own website and via other media outlets, clearly articulated a policy of constant threshold prices. We also obtained data on threshold prices for all products considered in our analysis, which confirmed that threshold prices were, indeed, constant over the course of any given haggling sequence. We also

assume that the consumer’s prior distribution is uniform over the interval between \underline{R} and \bar{R} . Similarly, we assume that the posterior distribution after a rejected offer is uniform between the last offer and \bar{R} . Note that given the nature of the Bayesian updating process formally defined below (the consumer can assign probability 0 to the event that the threshold price lies below her last offer), the distribution needs to be truncated. Moreover, consumers face a high degree of uncertainty with respect to the location of the threshold price in the interval. Analytically, a high variance distribution, which is truncated on both ends, does resemble a uniform distribution, while not requiring any additional shape parameter.

Let R denote the consumer’s reservation price. R might be within the interval ($\underline{R} \leq R \leq \bar{R}$) or it might be above the upper bound of the interval ($\bar{R} < R$). If R were to be below the lower bound ($R < \underline{R}$), the consumer would not submit any offer at all. In our model, the consumer is fully characterized by the parameter quadruplet $(c, \underline{R}, \bar{R}, R)$.

Initially, the consumer submits an offer $x_1 \geq \underline{R}$ and incurs frictional cost c . When deciding about x_1 , the consumer needs to trade off the cost of offering too little with the cost of offering too much. With a subjective probability $(\bar{R} - x_1)/(\bar{R} - \underline{R})$, the offer is lower than the threshold price. In this case, there is no direct reward to the consumer, except the option value of being allowed to submit an incremented offer x_2 . For the second offer, the consumer can rule out the event that the threshold price is below her first offer, x_1 . She can, thus, submit her second offer with better information, which takes the form of a narrower interval $[x_1, \bar{R}]$. With a subjective probability $(x_1 - \underline{R})/(\bar{R} - \underline{R})$, the offer is higher than the threshold price. In this case, the consumer collects a reward of $R - x_1$. Together with the frictional cost of submitting an offer, the consumer’s net utility of submitting an offer would be $R - x_1 - c$. This expression is always positive, as otherwise, the consumer would not submit an offer of x_1 in the first place.

To formulate the dynamic program modeling, the consumer’s behavior, we first consider $V(x)$, the optimal expected incremental surplus earned by a consumer whose last offer x was rejected. This function

can be written based on a stochastic dynamic program with the following Bellmann equation:

$$V(x_i) = \max \left\{ 0, \max_{x_i \leq x_{i+1} \leq R} \left(-c + \left(\frac{x_{i+1} - x_i}{\bar{R} - x_i} \right) (R - x_{i+1}) + \left(\frac{\bar{R} - x_{i+1}}{\bar{R} - x_i} \right) V(x_{i+1}) \right) \right\}, \quad (1)$$

where $((x_{i+1} - x_i)/(\bar{R} - x_i))$ is the subjective probability that a new offer x_{i+1} is successful (exceeds the threshold price), given that the last offer x_i was rejected (did not exceed the threshold price). The consumer incurs frictional costs for every offer she submits, independent of whether or not the offer is accepted by the retailer.

Equation (1) completely specifies the optimal haggling behavior. We can, therefore, write the optimal i th offer $x_i^*(c, \underline{R}, \bar{R}, R)$ as a function of the four parameters frictional cost c , willingness to pay R , and prior distribution \underline{R} and \bar{R} . For example, the first offer the consumer submits, $x_1^*(c, \underline{R}, \bar{R}, R)$, is determined as

$$x_1^*(c, \underline{R}, \bar{R}, R) = \arg \max_{\underline{R} \leq x_1 \leq R} \left(-c + \frac{x_1 - \underline{R}}{\bar{R} - \underline{R}} (R - x_1) + \left(\frac{\bar{R} - x_1}{\bar{R} - \underline{R}} \right) V(x_1) \right), \quad (2)$$

and we can define the other $x_i^*(c, \underline{R}, \bar{R}, R)$ accordingly. Moreover, there exists a stopping round $n^*(c, \underline{R}, \bar{R}, R)$ at which the consumer terminates the haggling process, if none of the previous offers were successful. Note that the rational consumer will calculate all offers x_1^* to $x_{n^*}^*$ prior to submitting the first offer. Note further that the true location of the threshold price, which is known to the retailer but not to the consumer, does not influence the calculation of the offers x_1^* to $x_{n^*}^*$. This calculation is purely based on the consumer's subjective beliefs as captured by the interval $[\underline{R}, \bar{R}]$. For this reason, our model characterizes also those consumers who hold beliefs that are inconsistent with the true threshold price of the retailer. For example, a consumer, who expects the threshold price to be in the interval $[80, 100]$, will behave according to our model, even if the true threshold price lies at 120. An illustration of the consumer trade-off is given in the online appendix.

Imputing Frictional Costs

While simple economic intuition suggests that consumers with low frictional costs are more likely to submit offers using small increments, and consumers with high frictional costs are more likely to submit offers in large increments, the notion of bounded rationality (see, for example, Simon 1955, Cyert and March 1963) suggests that a consumer, fully intending to be rational, will end up with a haggling strategy that resembles the optimal haggling strategy with some noise. Hence, an exact quantification of frictional cost directly from the sequence of offers is not possible. Instead, we have to impute frictional cost by using our decision model (1) and calibrate it based on the observed sequence of submitted offers.

Label the consecutive offers of a consumer on a specific product as x_1 to x_L , where x_1 is the first offer and x_L the last. Assume that the consumer is rational and is fully characterized by the four parameters frictional cost c , willingness to pay R , and prior distribution \underline{R} and \bar{R} . Then, following (1), rational behavior is characterized by $x_1 = x_1^*(c, \underline{R}, \bar{R}, R), \dots, x_L = x_L^*(c, \underline{R}, \bar{R}, R)$, and $L = n^*(c, \underline{R}, \bar{R}, R)$.

When imputing consumer characteristics from an observed sequence of offers, we will, therefore, search for the parameter quadruplet $(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ that achieves the best fit with the observed data. We define the concept of best fit based on two criteria. First, we only consider those potential consumer characteristics $(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ that lead to the observed stopping behavior of the haggling process. For example, when observing a sequence of submitted offers $\{x_1 = 125, x_2 = 150, x_3 = 170, x_4 = 180\}$, we would not consider the consumer to be of type $(20, 100, 200, 200)$, as—following (1)—the optimal number of offers submitted by consumer $(20, 100, 200, 200)$ would be a single offer of 150 instead of a sequence consisting of four offers. Note that the stopping behavior is only observable for those consumers who were unsuccessful. For these cases, we can, thus, restrict ourselves to those parameters where the predicted stopping behavior, $n^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$, is consistent with the observed stopping behavior L :

$$n^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}) = L. \quad (3)$$

If, in contrast, the last offer was successful, we do not observe the consumer's $n^*(c, \underline{R}, \bar{R}, R)$. However, we can impose a constraint on the estimated consumer characteristics, which captures that the consumer's $n^*(c, \underline{R}, \bar{R}, R)$ has to be at least as big as the observed number of offers, L :

$$n^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}) \geq L. \tag{4}$$

An additional constraint reflects the outcome of the haggling process. If the last offer was successful, $x_L^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ had to be above the threshold price, while for an unsuccessful offer, $x_L^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ had to be below the threshold price.

Second, the observed offer should be reasonably close—as measured by the sum of their squared distances—from the optimal offers placed by a consumer of type $(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$. For example, for the above sequence of $\{x_1 = 125, x_2 = 150, x_3 = 170, x_4 = 180\}$, both consumers (5, 100, 200, 200) and (2, 60, 100, 100) fit in terms of the observed stopping behavior ($n^*(2, 60, 100, 100) = n^*(5, 100, 200, 200) = 4$). However, the optimal sequence of submitted offers for the former based on (1) is $\{x_1^*(5, 100, 200, 200) = 129, x_2^*(5, 100, 200, 200) = 153, x_3^*(5, 100, 200, 200) = 172, x_4^*(5, 100, 200, 200) = 185.25\}$ and seems to “fit the observed sequence of submitted offers better” compared to the latter, which would result in offers of $\{x_1^*(2, 60, 100, 100) = 71.6, x_2^*(2, 60, 100, 100) = 81.2, x_3^*(2, 60, 100, 100) = 88.8, x_4^*(2, 60, 100, 100) = 94.1\}$.

More formally, we expect the estimator $(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ to minimize the sum of squared differences between actual and predicted offers. Given that the consumer type is characterized by four parameters, providing our model four degrees of freedom when searching for the best possible fit, we initially restricted our sample to sequences of length four and above ($L \geq 4$). For these sequences, we search for the parameter quadruplet $(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ that achieves the best fit with the observed data. For the unsuccessful offer sequences, we solve the following optimization problem:

$$\min_{\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}} \sum_{i=1}^L (x_i - x_i^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}))^2 \tag{5}$$

subject to $n^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}) = L$.

For the successful offer sequences, we solve

$$\min_{\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}} \sum_{i=1}^L (x_i - x_i^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}))^2 \tag{6}$$

subject to $n^*(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R}) \geq L$.

The imputed consumer characteristics $(\hat{c}, \hat{\underline{R}}, \hat{\bar{R}}, \hat{R})$ we obtained from (5) and (6) exhibit an interesting pattern. The consumer's expectation of the upper bound \bar{R} is found to be similar to the consumer's valuation of the product R . This pattern is intuitive, as it suggests that the consumer has seen a posted price for the product in another channel, which determines both her expectation of the upper bound \bar{R} and her valuation of the product R .

Given that the two parameters \bar{R} and R coincided in our estimation, we can reduce the number of parameters describing a consumer from four to three, thereby moving from a quadruplet $(c, \underline{R}, \bar{R}, R)$ to a triplet (c, \underline{R}, R) . The triplet is obtained by adding the constraint $\bar{R} = R$ to the optimization problems (5) and (6). This reduction in parameters allows us to include sequences of length three into our samples ($L \geq 3$), which substantially increases our sample sizes. All of the following results were analyzed with both the three and the four parameter formulation. Except for the difference in sample size, no systematic differences could be detected. In the remainder of the article, we therefore focus on the three parameter formulation.

To investigate the validity of the specific assumptions outlined above and the overall formulation of our consumer decision model, we compare the actual offers submitted by the consumers with the predicted offers based on our models. In other words, we compare the actual offers x_i with the optimal offers, $x_i^*(\hat{c}, \hat{\underline{R}}, \hat{R})$, a rational consumer with the imputed parameters $(\hat{c}, \hat{\underline{R}}, \hat{R})$ should place. Based on these comparisons, we did not detect any systematic deviation between predicted and actual values. A simple regression analysis between actual and predicted values shows that more than 95% of the variance in consumer behavior is captured by our model. Moreover, the slope of the corresponding regression line is estimated as 0.97, thereby close to identity. The average deviation between predicted and actual offers was

EUR 2.50. A comparison of residuals across rounds of offer sequences (e.g., a comparison of all third offers across all offer sequences) also did not indicate any systematic differences. This analysis provides strong support for the validity of our consumer haggling model.

6. Econometric Analysis

Based on the estimation defined by (5) and (6), we now turn to the measurement of frictional costs for those consumers submitting offers on PDAs, CD-RW and DVD drives, and MP3 players. Given that we can measure frictional costs only for sequences of length $L = 3$ and longer, we obtain a sample size of 327 for PDAs, 690 for CD-RW and DVD drives, and 202 for MP3 players. The resulting descriptive statistics for the frictional cost estimator \hat{c} are summarized by Table 3. First, and most importantly, we observe that frictional costs are, indeed, substantial. The median frictional cost in the sample is EUR 6.08 for the PDA (2.70% of the threshold price), EUR 4.29 for the CD-RW and DVD drives (3.15% of the threshold price), and EUR 3.54 for the MP3 player (2.81% of the threshold price). Second, we also observe that the frictional cost substantially varies across consumers. The standard deviation of frictional cost is EUR 7.57 for the PDA, EUR 4.42 for the CD-RW and DVD drives, and EUR 3.81 for the MP3 player. Thus, consumers are heterogeneous in their frictional costs.

Sources of Frictional Costs: Model Specification

When modeling the impact of experience on frictional cost, we rely on the conventional model, which postulates a constant elasticity between frictional cost and experience. In other words, the relationship between

an increase of experience at a specific experience level and the corresponding decrease in frictional cost at the specific cost level is constant. In its simplest form, we estimate the following model:

$$\hat{c}_i = \alpha * xp_i^{-\beta} * e^{\varepsilon_i}, \quad (7)$$

$$\log(\hat{c}_i) = \log(\alpha) - \beta * \log(xp_i) + \varepsilon_i. \quad (8)$$

In this model, the frictional cost of consumer i is denoted as \hat{c}_i . The variable xp represents the cumulative number of offers the consumer has submitted to our site *before* submitting the first offer for the haggling sequence under analysis. There are situations in which experience could also be accumulated by using similar websites from competing retailers. The NYOP retailer we studied pioneered the business model in Germany and had, at the time of our study, a dominant market position. Thus, it is unlikely that our measure is confounded by “spillovers” reflecting consumer experience with other retailers.

We extend this base model by adding several sociodemographic variables that allow us to test the hypotheses on education, income, and self-employment. To control for the differences in price sensitivity across the product variants in our sample, we also introduce a fixed effect variable $prod$ into our regression model.⁸

Thus, the regression model extends to

$$\log(\hat{c}_i) = \log(\alpha) + \beta_1 * \log(xp_i) + \beta_2 * inc_i + \beta_3 * edu_i + \beta_4 * selfemp_i + \sum \gamma_j * prod_{i,j} + \varepsilon_i. \quad (9)$$

We measure income (inc) as the average purchasing power of a person living in the community. Education (edu) is measured as the percentage of people with a degree preparing for a university education. In Germany, only 30% of adults have such an “Abitur,” thus, this measure is different from measuring education in the United States based on a high-school degree. The

⁸ Producers of consumer electronics have long engaged in the practice of price discrimination over time. A producer can charge a higher price at the beginning of a product’s life cycle than at later stages. Buyers of recently introduced models such as the Palm Pilot M100 are typically less price sensitive compared to buyers who intend to purchase a model that is already in the more mature stage of its life cycle.

Table 3 Descriptive Statistics of Frictional Costs (in Euros)

	Mean	Std dev	Min	Q1	Median	Q3	Max
CD-RW and DVD drives	5.51	4.42	0.26	2.34	4.29	8.37	32.06
MP3 player	4.84	3.81	0.44	1.70	3.54	8.88	22.33
PDA	7.95	7.57	0.44	2.93	6.08	10.69	65.59

percentage of self-employed people is measured by *selfemp*. We have compared our data with the most recent sociodemographic census data from the Statistische Bundesamt Deutschland and found them to be consistent. We tested for interaction effects between independent variables and did not find any significant results. Hence, the model above assumes homogeneous effects across the population.

Sources of Frictional Costs: Test of Hypotheses 1, 2, 3a, and 3b

We begin our analysis with a model relating the frictional cost to the haggling experience while controlling for the product. The first hypothesis states that greater haggling experience should lead to lower frictional costs, hence, we expect $\beta_1 < 0$. The corresponding regression results are summarized in Table 4.

First, consider the CD-RW. Our first econometric model includes only an intercept and the fixed effects to capture differences across product variants. The intercept of 1.13 is significant at the 1% level, so is the overall model ($F = 5.14$). Yet, the overall explanatory power is rather low ($\text{Adj. } R^2 = 7.8\%$). This changes as we add experience to the model. Based on Model CDR-2, we find strong support for the hypothesis that an increase in haggling experience lowers frictional costs. The coefficient for learning is negative and highly significant ($\beta_1 = -0.35, p < 0.0001$). Moreover, the explanatory power of the model substantially increases ($\text{Adj. } R^2 = 17.0\%$).

Next, we extend Model CDR-2 by adding the socio-demographic variables. If frictional cost is lower with higher levels of education (H2), higher with higher levels of income (H3a), and higher for self-employed consumers (H3b), we expect $\beta_2 < 0, \beta_3 > 0$, and $\beta_4 > 0$. As before, we find support for the hypothesis that haggling experience lowers frictional cost. However, we do not find support for the hypotheses that any of the sociodemographic factors influence frictional cost.

The results for the PDAs and for the MP3 player are similar to the ones for the CD-RW and DVD drive. Experience is significantly decreasing frictional costs and increasing the explanatory power of the model. Despite these similarities, a couple of differences deserve further discussion. First, the intercept in the regression varies across product categories. The MP3 player has the lowest intercept, 1.07 in Model MP3-3, while the PDA has the highest intercept (2.33 in Model PDA-3). This reflects the differences in absolute frictional costs already observed in Table 3. It is interesting to observe that consumers exhibit the highest frictional costs for the most expensive product we studied. The average threshold prices for the three product categories were EUR 203.92 for the CD-RW and DVD drives, EUR 153.07 for the MP3 players, and EUR 281.55 for the PDAs. This observation supports our estimation approach of keeping the three product categories underlying this study separate, as opposed to pooling them into one single model. It also opens

Table 4 Regression on Frictional Cost and Experience for CD-RW and DVD Drives, MP3 Player, and PDA (with Product Fixed Effect)

Variables	CDR-1	CDR-2	CDR-3	MP3-1	MP3-2	MP3-3	PDA-1	PDA-2	PDA-3
Intercept	1.13*** (7.56)	1.73*** (10.97)	1.75*** (7.57)	0.43* (1.89)	0.76*** (2.97)	1.07** (2.60)	1.42*** (21.87)	1.70*** (12.47)	2.33*** (6.96)
Experience	—	-0.35*** (-8.71)	-0.35*** (-8.74)	—	-0.20*** (-2.74)	-0.21*** (-2.75)	—	-0.17** (-2.30)	-0.21*** (-2.86)
Income	—	—	-0.00 (-0.32)	—	—	-0.00 (-1.05)	—	—	0.00* (-1.7)
Education	—	—	-0.29 (-0.60)	—	—	0.24 (0.27)	—	—	0.54 (0.75)
Self-employed	—	—	0.84 (0.70)	—	—	0.23 (0.11)	—	—	-1.25 (-0.71)
<i>N</i>	690	690	690	202	202	202	327	327	327
<i>F</i>	5.14	10.40	8.69	6.48	6.79	5.15	13.90	11.89	8.68
Adj. <i>R</i> ²	7.77%	16.98%	16.74%	17.91%	20.59%	19.86%	10.61%	11.78%	14.43%

up interesting opportunities for future research that we describe in the last section.

Second, we observe slight differences across product categories concerning the explanatory power of our model specification. The adjusted R^2 of the regression with experience and demographic variables ranges between 14.4% (PDA) and 19.9% (MP3 player). Moreover, the increase in the adjusted R^2 that we obtain by adding experience to the base model is highest for the CD-RW and DVD drives (increase of 9.2%) and lowest for the PDA (increase of 1.2%).

Third, we observe a slightly negative effect of income on frictional costs in the case of PDA. However, this effect is only significant at the 10% level and is not observed in the other two product categories. Otherwise, parameter estimates are relatively stable across the three product categories, except that the coefficient for experience slightly differs from the ones obtained for the MP3 player and the PDA (-0.35 compared to -0.21).

Our regression results clearly indicate that experience is a main driver of frictional costs. These findings are consistent across three product categories we analyzed. We also find that sociodemographic variables are, in general, inadequate predictors of frictional cost. These results, while contradicting popular beliefs about the characteristics of the Internet population, reaffirm Montgomery's (1999) findings. Based on Media Metrix panel data, he analyzed the Web usage of 5,000 household over 16 months. He concluded that sociodemographic factors such as age, education, and income are poor predictors of Web usage, accounting only for less than 6% of the variation across users.

A potential limitation of our effort to establish experience as the primary source of frictional costs relates to the direction of the causality between these two variables. Thus, despite a 1% significant effect of the experience coefficient in a regression with frictional costs as the dependent variable and an increase in the explanatory power of the model, we cannot rule out the possibility that the true causality is reverse, i.e., frictional costs drive experience. One could argue that customers with lower frictional costs are likely to obtain a higher experience from submitting offers on other products, which leads to a higher experience

score. To fully establish causality, one would have to design a controlled experiment. Such an experiment could randomly assign subjects to two groups, force one of the groups to submit offers on some products (leading to higher experience in this group), and then compare both groups with respect to their haggling pattern for other products.

Quasi-Natural Experiment: Old vs. New Policy

In our analysis of the quasi-natural experiment, we find that the information rent significantly decreases once consumers are given the option to increment rejected offers. Before the policy change, the average margin for the retailer was 32.7%. After the policy change, the margin was reduced to 20.9%. This strongly supports Hypothesis 4. A more detailed summary of our analysis of the quasi-natural experiment can be found in the online appendix.

7. Key Decisions in Managing the Iterative Policy

Based on our experience with the NYOP retailer, we identified five important decisions related to the implementation of the iterative policy. We also contrast how the decisions were made at our research site with how the management of our site should have made these decisions in retrospect, based on the experience described in this paper.

First, when moving to an iterative policy, one needs to decide if one changes the threshold price. Our research site has historically focused on the information rent as its source of profits. This is in line with economic intuition for the old policy: if the consumer is only allowed to submit one single offer, it is optimal for the retailer to accept every offer above cost. The intuition changes with the introduction of the iterative policy. Allowing the consumer to iterate actually *strengthens* the position of the retailer. Without iteration, a consumer effectively makes a "take-it-or-leave-it" ultimatum to the NYOP retailer. If iteration is allowed, the consumer is not committed to her offer. This can make it optimal for the retailer to reject even potentially profitable offers in the hope of achieving a higher price from the same consumer in the following round. Our research site, however, did not

increase threshold prices when introducing the new policy. It thereby failed to substitute the decrease in information rent with a potential increase in the baseline profit margin.

Second, the extent to which a consumer incurs frictional cost can be partly controlled through the design of the user interface. Specifically, the retailer can increase frictional costs by requiring the consumer to enter more information every time she submits an offer. Moreover, the response time with which the consumer is informed if her offer has been accepted is also under the control of the retailer. The sooner the feedback is provided, the lower the consumer's frictional cost becomes. When calibrating frictional costs, the retailer needs to balance two competing effects. On the one hand, a decrease in frictional costs will make it more likely that consumers participate and submit offers in the first place. On the other hand, lower frictional costs increase the number of iterations, n^* , and thereby decrease the information rent. Finding the optimal balance between these two effects is difficult, as it requires data about how the information rent decreases with frictional costs, and how many additional offers the site receives. When designing the new policy, none of this information was available to the management of our research site, so that decisions such as waiting times were made in a rather ad hoc fashion. Our study presented in this paper has explicitly measured the effect of information rent erosion. Additional market research is needed to discover if and to what extent the new policy will be able to attract new consumers beyond the time frame analyzed above.

A third question relates to if and how the threshold price in an iterative policy should be adjusted over the course of the haggling sequence. While, for the products described in this study, the threshold price was constant across offers, our findings triggered a discussion at our research site as to what extent iteration should be made less attractive by increasing the threshold price once a consumer has submitted a certain number of unsuccessful offers. In this case, iterative haggling as described by (1) would become more costly as the consumer would not only incur frictional cost c when submitting an offer, but would also partially forego the opportunity of future price

savings. We further discuss this question in the context of future research opportunities.

Fourth, the iterative policy provides new opportunities of price discrimination. For example, if data from prior interactions with a consumer indicates that the consumer is likely to increment her offer, the intermediary can increase profits by rejecting the first offer even if her first offer was above the threshold price.⁹ At the time of our study, such an approach would have been a violation of German law. According to German law concerning rebates (*Rabattgesetze*), providing different prices and especially different processes for obtaining the price might be illegal. During the last two years, the laws have been differently interpreted in different cases, leading to a somewhat ambiguous legal basis. One major price intermediary was found guilty of illegal price discrimination, which resulted in 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.

Finally, NYOP retailers need to prepare their sites against an increasing number of consumers who do not submit offers directly but, instead, employ electronic agents for this purpose. The electronic agent can automatically increment rejected offers, which would lead to an almost complete erosion of the information rent. Delaying the response time with which the consumer (or the agent) is informed about the outcome of her offer provides one defense against this scenario. Alternatively, a retailer could increase the threshold price (see above) if a consumer extensively iterates. We did not find any evidence of agent-based automated price incrementing in our data set.

⁹ Nevertheless, applying different threshold prices or response times based on consumer recognition 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 up publicly apologizing and refunding all customers who had paid higher prices.

In summary, despite our findings of a decrease in profits in this implementation of an iterative policy, the iterative policy has substantial economic potential. More practical experience is needed to calibrate the decision making on the five dimensions outlined above. The fact that, following the example of our research site, U.S. price intermediary Priceline.com is now partly embracing the iterative policy for certain product categories suggests that other NYOP retailers also believe in the potential of the iterative approach.

8. Conclusion and Future Research

In contrast to early visions of frictionless markets on the Internet, we find that frictional costs are substantial in absolute terms. A consumer considering a Palm Pilot purchase worth about EUR 200 assigns a median disutility of EUR 6.08 to a simple online transaction of keying in a number. Results for other product categories, including MP3 players and CD-RW and DVD drives, were consistent with this finding. Our sample was limited to consumers who had submitted at least three offers and, moreover, had to go through an initial registration process. This obviously limits the generalizability of our findings to a larger population. Yet, given that especially consumers with low frictional costs are more likely to self-select into our sample (that is, they are more likely to provide multiple offers and tolerate a registration process), frictional costs of the entire Internet population would be even *higher* than the estimates reported in this study.

In addition to determining the average frictional cost, we find that frictional costs substantially vary across consumers. We are able to explain a significant portion of this variance through the consumer's experience with the NYOP channel. We did not find income or education to explain variation in frictional cost. The absence of significant effects of income on online shopping behavior is consistent with a previous study by Montgomery (1999). Potentially, there are other variables such as computer literacy, speed of Internet access, and online experience, that negatively affect frictional cost, thereby negating the positive effect of income.

Our model of iterative haggling between a consumer and a NYOP channel together with our measurement of frictional costs opens up several avenues

for future research. First, our study restricts itself to measuring frictional costs in a specific domain. Such a restriction is necessary, as frictional costs directly depend on the task the consumer confronts. Nevertheless, replicating the measurement of frictional costs in other settings and comparing the results with our findings from a NYOP channel is an important research extension. Second, fully rational consumers should display similar frictional costs, independent of the value of the product under consideration. Interestingly, we observed that consumers exhibited the highest frictional costs for the most expensive product. The difference in absolute frictional costs across models resembles an experimental result reported in Kahneman and Tversky (1984) and Thaler (1999). It would be interesting to study if frictional costs of one and the same consumer are, indeed, higher for more expensive products. Third, in this paper, we have focused our modeling effort on the consumer haggling process. An interesting extension of our work would be to "switch sides" and take the perspective of the retailer.

In conclusion, our research shows that frictional costs in e-markets are nonnegligible in B2C settings. This may contribute to the existence of price dispersion on the Internet, a phenomenon that has puzzled researchers and managers alike.

Acknowledgments

The authors to thank the management team of the German price intermediary for providing their internal data, which are the empirical foundation of this study. The authors also thank Claritas (Germany) for providing sociodemographic data from German consumers. Sergei Savin provided substantial help in programming the consumer choice model. They are also grateful to Pete Fader, Rebecca Hann, Paul Kleindorfer, and Ivan Png for their helpful comments. The second author gratefully acknowledges financial support from the Wharton e-Business Initiative. Finally, the authors are grateful to the two editors, the associate editor, and three referees for a constructive review process.

An online appendix to this paper is available at <http://mansci.pubs.informs.org/ecompanion.html>.

References

- Argote, L. 1993. Group and organizational learning curves: Individual, system, and environmental components. *British J. Soc. Psych.* 32(1) 31-51.

- Bailey, J. 1998. Electronic commerce: Prices and consumer issues for three products: Books, compact discs, and software. *Organ. Econom. Co-Operation Development* 98(4).
- Bakos, J. Y. 1991. A strategic analysis of electronic marketplaces. *MIS Quart.* 15(3) 295–310.
- . 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Sci.* 43(12) 1676–1692.
- Brynjolfsson, E., M. D. Smith. 2000. Frictionless commerce: A comparison of internet and conventional retailers. *Management Sci.* 46(4) 563–585.
- Card, S. K., T. P. Moran, A. Newel. 1983. *The Psychology of Human Computer Interaction*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Clay, K., R. Krishnan, E. Wolff. 2001. Prices and price dispersion on the Web: Evidence from the online book industry. *J. Indust. Econom.* 49(4) 521–539.
- Clemons, E. K., I. H. Hann, L. Hitt. 2002. Price dispersion and differentiation in online travel: An empirical investigation. *Management Sci.* 48(4) 534–549.
- Cyert, R., J. G. March. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.
- Gabszewicz, J., J.-F. Thisse. 1980. Entry (and exit) in a differentiated industry. *J. Econom. Theory* 22(2) 327–338.
- Johnson, E. J., S. Bellman, J. Lohse. 2002a. Cognitive lock in and the power law of practice. Working paper, Columbia Business School, Columbia University, New York.
- , W. Moe, P. Fader, S. Bellman, J. Lohse. 2002b. On the depth and dynamics of online search behavior. Working paper, The Wharton School, University of Pennsylvania, Philadelphia, PA.
- Kahneman, D., A. Tversky. 1984. Choices, values, and frames. *Amer. Psych.* 39(4) 341–350.
- Montgomery, A. 1999. Using clickstream data to predict WWW usage. Working paper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA.
- Png, I., T. S. Y. Lee. 2001. Buyer switching costs.
- , ———, C. Yan. 2000. The competitiveness of on-line vis-à-vis conventional retailing. Working Paper, School of Computing, National University of Singapore, Singapore.
- Pratt, J. W., D. A. Wise, R. Zeckhauser. 1979. Price differences in almost competitive markets. *Quart. J. of Econom.* 93(2) 189–211.
- Salop, S., J. Stiglitz. 1976. Bargains and ripoffs: A model of monopolistically competitive price dispersion. *Rev. Econom. Stud.* 44(3) 493–510.
- Shaked, A., J. Sutton. 1982. Relaxing price competition through product differentiation. *Rev. Econom. Stud.* 49(1) 3–13.
- Simon, H. A. 1955. A behavioral model of rational choice. *Quart. J. Econom.* 69(1) 99–118.
- Smith, M., E. Brynjolfsson. 2001. Customer decision-making at an Internet shopbot: Brand matters. *J. Indust. Econom.* 49(4) 541–558.
- Thaler, R. H. 1999. Mental accounting matters. *J. Behavioral Decision Making* 12(3) 183–206.

Accepted by Arthur Geoffrion and Ramayya Krishnan; received February 2002. This paper was with the authors 16 months for 6 revisions.