WORKING PAPERS

When the Price You See Is Not the Price You Get: A Bargaining Study

Sandro Shelegia

Joshua Sherman

November 2014

Working Paper No: 1410

DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

When the Price You See Is Not the Price You Get: A Bargaining Study*

Sandro Shelegia† Joshua Sherman‡

November 6, 2014

Abstract

Although (or because) it is uncommon to observe consumers bargaining at retail stores in the Western world, the circumstances under which retail firms are actually willing to bargain is largely unknown. We construct a theoretical model in order to better understand how price and firm characteristics influence a firm's incentives to bargain and test the model's predictions by conducting a field experiment at nearly 300 stores throughout Vienna, Austria. In particular, we analyze the extent to which retail firms throughout Vienna consent to granting a discount when asked. A discount was granted approximately 40% of the time, and the average positive discount was approximately 10% off of a product's posted price. We relate firms' willingness to bargain to price and firm characteristics, in line with our theory.

JEL Classification: L81; D12; C78; C93.
Keywords: Bargaining; Pricing, Margins, Field Experiments.

1 Introduction

Given the prevalence of posted prices in the Western world, it is easy to forget that posted prices are a relatively recent phenomenon in the history of commerce. The first documented instance in the world occurred in Tokyo in 1673 and the first documented instance in the West occurred in New York in 1823.[1] Since then, posted prices have been

*We owe many thanks to 19 RAs who worked in various capacities on this project. We would like to thank Daniel Garcia, Maarten Janssen, Wieland Mueller, Karl Schlag, James Tremewan, and Jean-Robert Tyran for helpful suggestions. Assistance provided by Sylvie Hansbauer and Andrea Neidhart was also very much appreciated.

†Department of Economics, University of Vienna. Email: sandro.shelegia@univie.ac.at
‡Department of Economics, University of Vienna. Email: joshua.sherman@univie.ac.at.

[1] See Mahoney and Sloane (1966).

the norm for retail goods sold in developed countries. Furthermore, it is relatively rare to observe consumers seeking to bargain for products with a posted price in the West.[2]

We can conceive of two plausible explanations for consumers' general reluctance to bargain for retail goods. One explanation is that consumers believe that firms are unwilling to bargain, and therefore consumers view a posted price as a final offer. A second possibility is that consumers find the effort or embarrassment associated with bargaining to be too costly. While this study will not be able to address the latter explanation, it will seek to assess retail firms' willingness to bargain and relate their propensity to bargain to price and firm characteristics.

The rarity of bargaining has important implications for welfare. Generally speaking, when the cost of bargaining for firms is low, transactions that do not occur due to the fact that a consumer assumes that bargaining will be fruitless result in deadweight loss whenever a consumer's valuation for a good is greater than the firm's marginal cost and less than the posted price.

Given that we know nearly nothing about bargaining in the West, this study aims to make a first cut at understanding the issue. To this end, we sent hypothetical consumers to nearly 300 stores throughout Vienna to ask for discounts. We then use the data collected from these interactions in order to test our theoretical predictions regarding the likelihood and the extent to which discounts are granted by a given firm selling a particular product.

Our study is very different from previous empirical studies of bargaining from several standpoints. That is, whereas previous studies have analyzed firm bargaining behavior in individual markets, we study a large collection of diverse stores throughout one city. In addition, while previous studies have studied environments in which bargaining is the norm, little is known regarding the extent to which firms are willing to bargain in Western retail stores, the firms of focus in our study. The products sold by such firms comprise a non-trivial component of total consumer expenditures. According to Statistik Austria, products sold by retail firms of the type that we examine account for approximately 20 percent of total consumer expenditures.[3] And finally, while previous studies have focused on race, gender, or statistical discrimination as explanations for firm discounting behavior, our study focuses on how price and firm characteristics influence the extent and the degree to which discounts are offered (while controlling for consumer heterogeneity).

Bargaining with lengthy delays is unnatural for most retail products. Therefore, for purposes of understanding bargaining in the retail sphere, we view bargaining as a means for a retailer that sells many units of a particular good to price discriminate amongst

[2]The NY Times (December 15, 2013) and other news outlets have recently addressed the phenomenon of bargaining at retail stores in a casual, anecdotal manner and found that retail stores such as Best Buy, Home Depot, Nordstrom, and Bloomingdales were willing to give a discount when asked. Nevertheless, the phenomenon still appears to be relatively rare outside of the real estate and automobile industries.

[3]According to 2013 data provided by Statistik Austria regarding the main aggregates of National Accounts in Austria.

consumers with heterogeneous valuations. In such an environment, bargaining is mainly an attempt to extract additional surplus based on the consumer's willingness to pay; surplus above and beyond that which is extracted by the firm with a uniform price. Another well-studied instrument in economics that may be employed to achieve this goal is price discrimination. For example, third degree price discrimination exploits the firm's a priori knowledge regarding the willingness to pay of subsets of consumers in order to extract surplus. In contrast, a consumer's initiation of a request for a discount may allow the firm to obtain information about her willingness to pay and extract surplus a posteriori. Therefore, a key common driver of price discrimination and bargaining is the firm's ability to learn the consumer's valuation for its product in order to extract surplus.

We model a firm selling a good that posts a price and decides whether to bargain with willing consumers. If a consumer initiates a request for a discount, the firm receives additional information about the consumer's willingness to pay and can decide whether or not to grant a discount based on this information. In this simple environment, the firm's decision to offer a discount will depend on a number of factors, including the size of its margins, its ability to monitor and incentivize its clerks, and its cost of bargaining. Regarding the firm's margins, clearly if a firm chooses to bargain, its price and therefore its margin will depend on the proportion of consumers that are willing to request a discount. However, if the proportion is small, as we believe to be the case for most products that we examine, then the effect of this proportion on the firm's pricing decision will be negligible. Therefore, for most products we will examine, rather than margins and the bargaining decision being determined simultaneously, the bargaining decision will be determined according to the firm's margins.

On the surface, one might ask why we would study the firms propensity to bargain in an environment in which the proportion of consumers which ask for a discount is small. It is important to note that this study does not seek to understand firm bargaining behavior in environments in which bargaining is common. The focus of our study, as discussed earlier, is on understanding how firms which typically are *not* approached by consumers with a request for a discount would respond to such a request.

While we only consider a monopolist, it shall become evident from the model that competition affects bargaining decisions insofar as it reduces margins, and therefore, controlling for competition, the qualitative predictions of the monopoly model can be extended to oligopolistic environments.[4]

Our model predicts that a firm is more willing to bargain when its margins are large, its bargaining cost is small and its ability to monitor and incentivize clerks is high. Furthermore, the effect of a firm's margins on its propensity to bargain will increase in its

[4]Implicit in this is our assumption that bargaining is infrequent enough such that the strategic effects of firms' bargaining decisions on their competitors are small, and therefore a firm's bargaining decisions can be analyzed in isolation.

ability to monitor and incentivize its clerks. We test these predictions by examining a large variety of products at various price levels in order to exploit variation in presumed margins. In particular, we sent 12 research assistants to pose as consumers at nearly 300 stores in four demographically diverse areas of Vienna. The research assistants were assigned to record characteristics of the stores that they were assigned as well as information about their interactions with salespeople in an attempt to receive a discount. Each store was assigned to be visited by three RAs and each store was assigned to be visited at least once before Christmas (December 9-14, 2013) and at least once after Christmas (December 27, 2013-January 4, 2014). Products surveyed ranged from a posted price of 30 EUR to 999.99 EUR, with a median posted price in the data of 135 EUR. To our knowledge, nearly all products that were observed were new products.

While the comparative statics of our theoretical model will allow us to form empirical predictions regarding the direction in which we expect firm and product characteristics to influence a firm's likelihood of offering a discount, we were quite uncertain a priori with regards to the extent to which firms would grant a discount overall. Nevertheless, we were largely surprised to learn that a discount was offered approximately 40% of the time overall (303 of 751 products). In order to get a sense for the types of products for which a discount was granted, here we provide a list of 10 products taken from 303 such instances: a backpack, a blanket, a cordless screwdriver, a kite, a manicure set, a bottle of perfume, a scarf, a stuffed animal, a surveillance camera for babies, and a sweater.[5] Including instances in which no discount was granted, the average discount was approximately 10 EUR (7,715 EUR off of 159,516 EUR, or 4.8%). Conditional on a discount being granted, the average discount size was approximately 25 EUR (7,715 EUR off of 80,725 EUR, or 9.5%) and the simple average discount percentage was approximately 10%. For perspective, consider that individual interactions were relatively brief and transpired in a matter of minutes, if not seconds. Along these lines, it is interesting to compare the average discount of 10 EUR to estimates of savings in the search literature. Pinna and Seiler (2014) track consumers paths in a Northern California supermarket and find that extending search by 56 seconds lowers their total expenditure by 7%, or $1.90 off of an average expenditure of $27. While we did not ask our RAs to record the duration of their interactions, an average discount of 10 EUR for interactions that anecdotally did not last for more than a few minutes is comparable with this estimate of returns to consumer search behavior.

We may now summarize our main empirical findings. On average, we find that the probability of a discount increases with the price of the good, is lower for sale items, and is lower at large-scale firms (firms with many branches or firms with multinational reach). The first finding is consistent with our theory under the plausible assumption that in a cross-section of product markets, a higher price is associated with a higher

[5]A complete list of every product observed is available from the authors upon request.

(absolute) margin. The second finding suggests that products on sale have lower margins than identically priced goods that are not on sale. Such an inference would be consistent with an assumption that sales reflect demand rather than cost shocks. The third finding is also in line with our theory, given a presumption that large-scale firms such as Zara tend to organize themselves in such a way as to make it harder to monitor and incentivize giving out individualized discounts. In addition, we find that the difference in probability of earning a discount at a small-scale firm vs. a large-scale firm tends to be larger at higher price levels, as consistent with our theory that predicts a larger effect of margins (as proxied by price) on the propensity to bargain for firms with a higher ability to monitor and incentivize their clerks.

We also find that it is less likely to receive a discount for non-sale items after Christmas, whereas we do not find a difference in discount probabilities for sale items before and after Christmas, suggesting that sales immediately after Christmas do not affect margins differently than sales before Christmas. When a discount is granted, the amount of a discount increases with the price of the good; we estimate that the elasticity of discount size with respect to price is approximately unitary. Interestingly, stores which were observed with no customers for a three minute period give substantially larger discounts. Generally speaking, firms did not decide to grant a discount on the basis of RA identity or RA gender.

The paper is structured as follows. Section 2 provides a description of a general theoretical framework that will allow us to predict the effect of price and firm characteristics on the probability that a firm will grant a discount. Section 3 describes the design of the empirical study, and in Section 4 we discuss the empirical analysis and results. Section 5 concludes.

1.1 Literature

Following the pioneering bargaining work of Nash Jr (1950), the literature studying the decision of whether to commit to a uniform price or bargain begins in earnest approximately thirty years later with Rubinstein (1982). Shortly thereafter, several additional studies address the bargaining problem, including Riley and Zeckhauser (1983). They examine the uniform price versus bargaining decision made by the firm and find that committing to one price to each risk-neutral buyer is more profitable than any alternative pricing strategy. Other contemporaneous studies addressing the bargaining problem include Fudenberg and Tirole (1983), Myerson and Satterthwaite (1983), Sobel and Takahashi (1983), and later Perry (1986), Fudenberg, Levine and Tirole (1987), and Hart and Tirole (1988). Later studies include Bester (1994), Wang (1995), and Arnold and Lippman (1998). While our study addresses the firm's cost of bargaining as a per-customer cost of bargaining, Wang (1995) models it as a fixed cost. We find the per-consumer cost to be more realistic in a retail environment because Wang (1995) models bargaining as an all-or-nothing decision, while we model bargaining as a mechanism the firm chooses to employ with some fraction of consumers who seek to bargain, and so accounting for a (marginal) per-customer cost is natural. It also allows us to relate firm-specific characteristics that capture bargaining cost to the probability a firm agrees to grant a discount. When engaging in behavior-based discrimination, the firm determines the second price offer to an individual consumer based on whether he or she purchased at the first price offer. Fudenberg and Villas-Boas (2006) provide a survey of papers that study behavior-based discrimination. For purposes of our study, however, such theories will not be directly applicable because a firms pricing strategy will be observed only once for a given consumer. This shall rule out applications of behavior-based discrimination in which a second price is charged to a consumer that has already purchased. Whereas the majority of the theoretical literature on bargaining and price discrimination is devoted to a monopoly, most empirical studies (to be reviewed shortly) take place in oligopoly settings. Along these lines, a small collection of such theoretical papers lends insight into the question of whether uniform pricing or price discrimination is more profitable in oligopoly settings, including Holmes (1989), Corts (1998), and Armstrong and Vickers (2001). The three studies provide evidence that price discrimination is more profitable than uniform pricing in homogenous good oligopoly markets, whereas conclusions regarding which setting yields higher profits in differentiated product models are more sensitive to model parameters.

Empirical evidence of discriminatory price offers has been investigated thoroughly in the labor economics literature. For example, the suggestion that discriminative salary offers can be explained by race or gender differences, and whether such discrimination is animus based or statistically based, has long been a topic of empirical study. Existing empirical studies of price-related bargaining focus on the consumer side of bargaining, and each study examines one particular product or service. Along these lines, studies that have found evidence of statistical-based discrimination in price offers to consumers include Ayres and Siegelman (1995) (cars), Goldberg (1996) (cars), List (2004), and Castillo et al. (2013) (taxi service in Peru).

2 Theory

The aim of this section is to present a conceptual framework that we believe captures the main features of retail bargaining decisions. Along these lines, we begin by providing a concise description of our empirical setting in order to contextualize the theoretical model presented below.

A consumer enters a retail store with an interest to purchase a particular product. Upon noting the posted price of the product, the consumer approaches the nearest salesperson and engages in a brief interaction in order to request a discount off of the posted price. The salesperson then either denies the consumer's request or approves it with a lower price offer. Empirically, we observe a wide variety of retail stores selling a diverse range of products at various price levels. Therefore, we need to formulate a flexible theoretical model in which a firm chooses a posted price and decides how to respond to discount requests.

It is important to note outright that in our setting, we consider bargaining as a firm's attempt to price discriminate among heterogeneous consumers. This view is shared in the theoretical literature on endogenous choice of pricing mechanisms,[6] where bargaining is studied as a costly discrimination tool because it allows a firm to obtain additional information (via interaction with the consumer) about the consumer's willingness to pay. Based on this tradition, and given that it closely corresponds to our empirical setting, we will assume away any issues related to delays and discounting with which one must exercise caution when modeling bargaining over a single item with a costly delay (e.g. firms bargaining over an investment project).

While most, if not all, retailers in our sample have multiple competitors, we shall focus on a monopoly model that can be extended to competition, but which itself provides enough richness to highlight the main issues involved. Focusing on monopoly is largely without loss in terms of how prices are determined (we may always vary firm's demand as means of capturing its competitive environment). More importantly, a firm's choice regarding whether to bargain with consumers at all, and if so under what conditions, will be assumed to be independent of its competitors' decisions for reasons discussed in the Introduction. Furthermore, we shall study a monopolist that sells a single good rather than multiple goods because we believe (and observe empirically) that pricing and bargaining decisions are product-specific.[7]

2.1 Model

This leads us to a monopoly model in which a firm selling a good chooses a price and whether to bargain with consumers who request discounts. The firm is required (by law) to post a price p per unit at which it will be obliged to sell upon the consumer's request. If prompted, the firm may also choose to bargain with a consumer (we assume that the firm never initiates bargaining). In order to give the firm a reason to bargain, we assume that each consumer demands one unit of the good, and each consumer i has a valuation v_i that is drawn from a continuous distribution $F(v)$ on $[\underline{v}, \bar{v}]$. We shall assume that bargaining and non-bargaining consumers have identical preferences. This assumption is made for simplicity, and indeed one could easily allow for $F(v)$ to differ depending on

[6] E.g. see Bester (1994), Wang (1995), and Arnold and Lippman (1998).

[7] Of course, in our empirical analysis we will account for within-store correlation across observations, which we will address in further detail in Section 4.

the consumer's type.[8] Without interaction (through bargaining), the firm does not know anything about a particular realization v_i and thus cannot condition its posted price on it. That is, third degree price discrimination is not feasible. We assume that if consumer i decides to bargain, the firm obtains a noisy signal τ_i about v_i, but in what follows we shall concentrate on the simple case in which the signal is perfectly informative so that $\tau_i = v_i$.[9]

Furthermore, we assume that a fraction β of consumers may bargain and experience no cost of doing so,[10] the remaining consumers never bargain.

We consider two parameters which prevent a firm that agrees to grant a discount from obtaining the entire amount $v_i - c_i$ from a given consumer. The first parameter, $\lambda \in [0, 1]$, will represent the firm's ability to maximize the proportion of potential profit $v_i - c_i$ when granting a discount. In our model, for simplicity we consider this parameter to be exogenously given and resulting from a firm's decision regarding its scale. That is, given a firm's decision regarding its scale (i.e. number of firms, multinational reach), we presume that small-scale firms will feature salespeople that are better motivated and whom have higher incentives to maximize firm profits than salespeople at large-scale firms. For example, an owner who serves as a salesperson at her own store with only one location is likely more motivated and more highly skilled at extracting surplus from customers then the average salesperson at a multinational chain. At a large chain, salespeople either may be forbidden from granting discounts due to the chain's difficulty in monitoring the circumstances under which discounts are granted, or salespeople may be offered commissions and therefore the salesperson's interest in maximizing revenues will not be perfectly aligned with the firm's interest in maximizing profits. In either case, we presume that a salesperson at a large scale firm will be less effective in maximizing $v_i - c_i$ than a salesperson at a small scale firm. In our empirical analysis, we will approximate λ with various store characteristics, such as how many stores are operated by the same owner, whether the store is part of a multinational chain, and the physical size of the store. Note that endogenizing the firm's decision regarding its scale (i.e. endogenizing λ) is theoretically possible but is beyond the scope of this paper.

The second parameter that prevents a firm that agrees to grant a discount from obtaining the entire amount $v_i - c_i$ is a fixed cost of bargaining with a consumer, which we denote as b. In our model, b is only incurred if the salesperson agrees to grant a discount. It may be interpreted as the opportunity cost associated with paperwork or electronic manipulation required to sell a product for a price that is below the posted price as well as the cost of granting a discount given that another customer in the store may witness

[8] The firm would use hagglers' distribution to compute gains from bargaining, and use both consumer types' distributions to set the price.

[9] This assumption can be relaxed at the expense of a substantial loss in tractability.

[10] We can allow for such a cost and endogenize the bargaining decision, but for the most part we shall assume that β is small, thus endogenizing it is of no particular importance.

the interaction and seek a discount. In addition, the firm's owner cannot guarantee that the clerk will only negotiate with consumers from whom the firm earns positive profits. In fact, we assume that when the clerk grants a discount, he does so for any consumer for whom $v \geq c$, while the owner would like him to grant a discount only to consumers with valuation above $c + b$. That is, the firm will expend the fixed bargaining cost b on some consumers for whom such an expenditure is ultimately wasteful (consumers for whom $v < b + c$);

The fixed bargaining cost b is different than λ in that λ measures the firm's ability to extract surplus. While λ is essentially related the organization of a firm and not specific to the interaction, b is related to the specific circumstances surrounding the interaction. The distortion due to λ has a particular form - a firm extracts a fraction λ of the profit from each consumer who bargains. That is, due to the lack of managerial control of the salesperson's effort or skill, not all possible profit is extracted due to errors in giving appropriate discounts.

Now, assume a firm charges price p. By law, a firm cannot charge a price higher than p to any consumer with whom it bargains and who has a valuation above p. Thus, regardless of whether a firm bargains or not, all consumers with valuations above p will not bargain and will pay p. The firm will earn $(1 - F(p))(p - c)$ from such consumers. A fraction β of consumer with valuations below p will attempt to bargain. These will include consumers with a valuation below c, and for simplicity we shall assume that such consumers do not bargain knowing that there is no room for bargaining. If a firm bargains, then for every consumer whose valuation v is between c and p it will make a profit of $\lambda(v - c) - b$. Given all the above, the firm's profit can be written as

$$\pi = (1 - F(p))(p - c) + I_B \beta \int_c^p (\lambda(v - c) - b) f(v) dv, \tag{1}$$

where I_B is an indicator function that is equal to one if the firm bargains, and equal to zero if it does not.

2.2 Equilibrium

The firm must decide which price to charge and whether to allow a salesperson to bargain with consumers. In order to simplify analysis, we shall assume that π is concave in p for $I_B = 0, 1$.[11] Denote by p^M the monopoly price without bargaining. As is well known, p^M solves

$$p^M - c = \frac{1 - F(p^M)}{f(p^M)}. \tag{2}$$

If a firm bargains, then it sets $I_B = 1$ and charges p^B that solves

[11] This assumption is satisfied for various demand functions, including linear demand.

$$p^B - c = \frac{1 - F(p^B)}{f(p^B)(1 - \beta\lambda)} - \frac{\beta b}{(1 - \beta\lambda)}. \tag{3}$$

In order to better understand the firm's maximization problem, and subsequently its bargaining decision, denote by $\bar{p}$ a price such that $\int_c^{\bar{p}} (\lambda(v - c) - b) f(v) dv = 0$. Then the firm will bargain if $p^m \geq \bar{p}$. It may also be the case that a firm bargains even if $p^m < \bar{p}$. That is, since $p^B > p^m$, p^B may exceed $\bar{p}$ and result in a higher profit than not bargaining and charging p^M to all consumers. This possibility, however, is of little empirical relevance in environments where β is small.

Now let us define the firm's propensity to bargain $\omega(p^*)$ as

$$\omega(p^*) = \lambda \int_c^{p^*} (v - c) \frac{f(v)}{F(p^*) - F(c)} dv - b \tag{4}$$

Noting that $\frac{f(v)}{F(p^*)-F(c)}$ is the conditional density of v on $[c, p^*]$, $\omega(p^*)$ has a simple interpretation: $\omega(p)$ is the average profit the firm earns per bargaining consumer after it extracts surplus and pays b.

Clearly, the firm's propensity to bargain is closely related to its margin in that it is equal to the average valuation within the margin minus the bargaining cost. In fact, for a class of demand functions $(A - p)^\eta$, the integral in (4) is exactly equal to a fixed fraction of the firm's absolute margin.

We can now state the following sufficient condition for bargaining (necessity requires further assumptions on $F(v)$)

Proposition 1 *The firm sets $p^* = p^B > p^M$ when $\pi^B \geq \pi^M$ and $p^* = p^M$ otherwise. It bargains with a consumer iff $\omega(p^*) > 0$.*

Proof. We already assumed that the profit function has an interior maximizing price regardless of whether the firm bargains. Now assume that $\omega(p^M) > 0$. Then even when charging p^M, the firm wants to bargain with consumers. Then it should charge p^B because it is strictly larger than p^M. To see this, note that if $\omega(p) > 0$, then its derivative with respect to p is positive, and therefore $p^B > p^M$. ■

If the firm bargains, then the larger β, the higher is $p^* = p^B$. The intuition is the following - a firm that chooses to bargain does so because it makes extra profit from bargaining. When the proportion of consumers who bargain increases, the firm will increase its price in order to extract surplus from an even larger proportion of bargainers. The optimal bargaining price p^B also increases in λ for a similar reason - the more the firm can extract from bargaining, the more it will want to increase the price in order to bargain with an even larger proportion of bargainers whose valuations fall below the price. Finally, the larger is b, the bargaining cost, the lower is p^B.

We note p^* is endogenous to the bargaining decision, but given that we expect β to be small, and therefore for p^* to be close to p^M, this issue is of little practical importance. Regardless of this, we can state:

Proposition 2 *A firm is less likely to bargain when b large, λ small, and β small.*

Proof. The profit associated with bargaining for any given p is profit without bargaining plus $\beta \int_c^p (\lambda(v-c)-b)f(v)dv$. The latter is reduced when b increases or when λ decreases. Therefore when b increases or λ decreases, a firm that does not bargain will never begin to bargain. When β increases, the firm's profit increases for $p > \bar{p}$ and decreases for $p < \bar{p}$. Thus in the relevant range of p (where bargaining is more profitable than not bargaining), an increase in β increases profit from bargaining. ■

We now turn to the issue of an average observed discount, conditional on firm bargaining. Therefore now let us assume that $\omega(p^*) > 0$ is satisfied.

The average discount amount will depend on the salesperson's assessment of a particular RA's willingness to pay. Let τ_i denote the average signal that a clerk receives about the perceived willingness to pay of RA i. Then the average discount received by i, denoted by z_i^*, is given by

$$z_i^* = p^* - \tau_i.$$

We shall note that τ_i must be between p^* and c, and thus the average discount will be related to the firm's margins. The exact relation will depend on the stance we take regarding τ_i.

Let $\bar{v}$ denote the average valuation across all consumers who ask for a discount, thus $\bar{v} = \int_c^{p^*} \frac{vf(v)}{F(p^*)-F(c)}dv$. Denoting ϵ_i as the difference between $\bar{v}$ and the salesperson's perceived valuation of RA i, we assume that $\tau_i = \bar{v} + \epsilon_i$, where ϵ_i is independent of $\bar{v}$ and across RAs. We can then compute z_i^* as

$$z_i^* = p^* - \int_c^{p^*} \frac{vf(v)}{F(p^*)-F(c)}dv + \bar{\epsilon}_i, \tag{5}$$

where $\bar{\epsilon}_i$ is the mean of ϵ_i. Therefore $\bar{\epsilon}_i$ captures the difference between the average valuation across all consumers who bargain and the average (perceived) valuation of RA i. The first two terms in (5) are closely related to the firm's absolute margin $p^* - c$, which henceforth we shall denote by m^*. As an example, for a linear demand function, $z_i^* = m^*/2 + \bar{\epsilon}_i$.

We may therefore conclude that firms with higher margins will give larger discounts. If a case can be made that higher prices are associated with higher margins, then higher prices will also be associated with deeper discounts.[12]

[12] If higher prices imply higher margins, higher prices may lower $\bar{\epsilon}_i$. In such a case, a positive relationship between price and discount size may be attributed to a reduction in $\bar{\epsilon}_i$. We have taken several measures in the study design to prevent such an issue from occurring (e.g. we instructed RAs to never bargain for

2.3 Theory in the context of a cross-section of markets

Our empirical analysis is closely tied to our theoretical model and analyzes two primary outcomes. First, we seek to understand the circumstances under which a discount will be granted. Second, when a discount is granted, we seek to understand the determinants of the size of the discount. According to our theory, a firm will grant a discount when $\omega^* > 0$, and conditional on $\omega > 0$, the size of the discount is z^*.

The propensity to bargain depends on the firm's ability to extract surplus, λ, the per-transaction bargaining cost, b, the proportion of consumers who bargain, β, and demand and cost conditions, F and c. In our empirical analysis we observe a number of firm and product specific variables that help us capture b, λ and β. However, we do not observe c or F. Thus one empirical challenge is to find proxies for values contained in the integral in (4), which, as we show momentarily, is closely related to the firm's absolute margin.

Recall that we observe a large variety of products sold by various firms, and so it is useful to consider a cross-section of markets. We are thus interested in the relationship between p^* and ω^* in a cross-section of product markets with various realizations of F and c. If it is the case that there is a positive relationship between c and m (and hence p) in a *cross-sectional* sense, then there will also be a positive relationship between p and ω. In such a case, a higher price would be associated with a higher propensity to bargain.[13]

In order to better understand the relationship between ω and p, we shall consider β to be close to zero, which we suspect is the case for the majority of the products that we observe empirically. Let demand be of the form $D(p) = (A - p)^\eta$, so that $F(v) = 1 - \left(1 - \frac{v}{A}\right)^\eta$, where total demand is then $D(p) = A^\eta(1 - F(p))$.[14] Given that A^η simply scales demand up, all our previous results apply. So

$$p^* = p^M = \frac{A + \eta c}{1 + \eta}. \tag{6}$$

It shall be useful to define $m^* \equiv p^* - c$ as the equilibrium absolute margin, which is given by

$$m^* = \frac{A - \eta c}{1 + \eta}. \tag{7}$$

a product for which the posted price exceeded 1000 EUR), and also attempt to control for it empirically. See a more detailed discussion of this issue in Section 4.1.

[13]If, for example, one generates markets with random linear demand and cost, the relationship between price and margins will be the same as the relationship between price and the propensity to bargain. While it is possible to generate a cross-section of markets with F and c pairs such that the relationship between price and margins will be the opposite of the relationship between price and the propensity to bargain, such a construction would rely on particular shifts of distributions with cost that are unlikely to be empirically relevant.

[14]Recall that in our model total demand is at most 1, thus we have to normalize $D(p)$ by dividing it by A^η.

For $\beta = 0$, the propensity to bargain takes the following simple form:

$$\omega^* = \lambda\gamma(\eta)m^* - b > 0 \tag{8}$$

where $\gamma(\eta) = \frac{1+\eta\left(2-\frac{1}{1-\left(\frac{\eta}{1+\eta}\right)^{\eta}}\right)}{1+\eta}$. Note that $\gamma(\eta)$ does not depend on any parameter other than η, is positive, decreasing in η, and converges to 1 when $\eta \to 0$. Thus, the higher is η, the less willing is the firm to bargain. As per Proposition 2, condition (8) is satisfied when λ is sufficiently high, or when b is sufficiently low.

In a cross-section of product markets, the realizations of demand intercepts and marginal costs may be correlated, in particular because markets where $A < c$ do not exist. Assume that each product market has a draw from some distribution of marginal costs and a correlated draw from a distribution of intercepts. If we write $E(A|c)$ for the conditional mean of A given c, we can then take the derivative of $E(\omega^*)$ with respect to c to get

$$\frac{\partial E(\omega^*)}{\partial c} = \lambda\gamma(\eta)\left[\frac{\partial E(A|c)}{\partial c} - \eta\right]. \tag{9}$$

We can now link, in a cross-sectional sense, the firm's propensity to bargain to the firm's marginal cost in a simple way. If the mean of the intercept increases with c at a faster rate than η (for linear demand η is unity), than the propensity to bargain will increase with higher realizations of c, with the opposite holding true when $\frac{\partial E(A|c)}{\partial c} < \eta$. This follows from the fact that a higher realization of c leads to an even higher average realization of A.

The same exercise for p^* yields

$$\frac{\partial E(p^*)}{\partial c} = \frac{1}{1+\eta}\left[\frac{\partial E(A|c)}{\partial c} + \eta\right]. \tag{10}$$

Whenever $\frac{\partial E(A|c)}{\partial c} > \eta$, on average one will observe higher prices and higher margins for goods with higher cost realizations. Under these circumstances, there will be a positive relationship between price and the propensity to bargain. On the other hand, when $-\eta < \frac{\partial E(A|c)}{\partial c} < \eta$, a higher cost realization will lead to higher prices but lower margins and thus one will observe a negative relationship between price and the propensity to bargain.[15]

Now consider the average discount. Using demand of the form assumed in this section,

[15] The case where $\frac{\partial E(A|c)}{\partial c} < -\eta$ is of little economic interest.

we may write (5) as

$$z_i^* = \frac{\eta}{(1+\eta)^2\left(\left(\frac{1+\eta}{\eta}\right)^\eta - 1\right)}(A-c) + \bar{\epsilon}_i \equiv \phi(\eta)(A-c) + \bar{\epsilon}_i. \tag{11}$$

Assuming that $\partial\bar{\epsilon}_i/\partial c = 0$,[16] the relationship between c and the average discount is

$$\frac{\partial E(z_i^*)}{\partial c} = \phi(\eta)\left[\frac{\partial E(A|c)}{\partial c} - 1\right]. \tag{12}$$

Thus we conclude that whenever $\frac{\partial E(A|c)}{\partial c} > 1$, there is a positive cross-sectional relationship between p^* and z_i^*.

The statistical relationship between the realization of intercept A and cost c, as well as demand curvature η, is unknown to us. We do however, believe that $\frac{\partial E(A|c)}{\partial c} > \eta$, which is equivalent to stating that goods with higher prices have higher absolute (not percentage) margins. Assuming that this is the case, then our previous analysis leads to a conclusion that firms are more likely to give discounts on more expensive items. Furthermore, to the extent that the intercept grows larger than marginal cost ($\frac{\partial E(A|c)}{\partial c} > 1$) across products, conditional on a firm giving a discount, it will give larger discounts on more expensive goods.

To conclude our discussion on the link between the cross-sectional variation of p^* and our two primary variables of interest, ω^* and z^*, we note that when β is small, variation in b and λ only affects ω^* but does not affect p^* or z^*. Thus, when β is small, variables that proxy for the firm's bargaining cost and the firm's ability to extract surplus provide predictions regarding a firm's willingness to bargain but are not meaningful for purposes of predicting the size of a discount.

2.3.1 Interpretation of observed sales

We observe whether a product is on sale, which intuitively is closely related to a firm's margin. The inference about margins from observing that the product is on sale is not straightforward, and indeed may imply higher or lower margins depending on the circumstances.

A sale in our data is recorded whenever a list price of a product is higher than the posted price. In our simple model, this can only arise if there is a change in the primitives of the model, such as a demand or cost shock.[17]

In what follows, we provide very brief intuition with respect to how sales due to demand or cost shocks will influence a firm's propensity to bargain. In the Appendix we provide a formal two-period extension of our model in which sales may result from either

[16] See a discussion of this assumption in Section 4.1.

[17] When β is close to zero, changes to λ and b do not result in price changes.

i.i.d. demand shocks or *i.i.d.* cost shocks.

Consider two goods with the same posted price, where one item is a sale item and the other is not. We are interested in understanding the effect of an item's sale status on the firm's propensity to bargain. First suppose that sales reflect demand shocks. Intuitively, due to the fact that the items share the same posted price, the sale item should have a higher marginal cost because it has suffered a downward demand shock whereas the non-sale item has not. In such a circumstance, there will be a lower propensity to bargain for the sale item than the non-sale item.

Now consider the same example as above where prices fluctuate due to cost shocks rather than demand shocks. Then one would predict that there will be a higher propensity to bargain for the sale item because given two items at the same posted price, the sale item must have a higher demand because it has suffered a downward cost check yet shares its posted price with the non-sale item.

IO literature has also extensively studied temporary price reductions that are conducted in order to attract consumers who either have more information about prices than others (e.g. Varian (1980)), or are more patient (Sobel (1984)). While these explanations for sales are beyond the scope of our simple model, it is still illustrative to consider their impact. In these models, price reductions occur while consumer valuations and costs remain constant. Therefore, in the context of our model, one may refer to Equation 4 in order to see that any such reduction in price reduces the firm's willingness to bargain because all variables are held constant except for p^*, which is reduced.

It is worth emphasizing that a sale affects a firm's propensity to bargain through the probabilistic inference we derive about margins. Because the overall propensity to bargain is λ times the margin, any effect sales may have on propensity to bargain will be stronger for firms with higher λ.

We close our discussion of this section by relating sales to the predicted amount of a discount. Whenever sales reduce the propensity to bargain due to smaller margins, their effect on the discount amount will be in the same direction. That is, on average, smaller discounts will be granted on sale items relative to non-sale items. The opposite holds if sales increase the firm's propensity to bargain.

3 Study design

First and foremost, Vienna served as a reasonable location for purposes of conducting this study because it was the authors' place of residence prior to and during the data collection and because the authors are well-acquainted with the retail areas of the city. In addition, the retail shopping environment in modern-day Vienna is comparable to many large Western European and North American cities in that it consists of thousands of retail stores, from small independently owned stores to large multinational chains. The retail shopping culture in Vienna is also comparable to other Western cities in that it is uncommon to observe consumers seeking discounts from retailers. Therefore Vienna served as an appropriate venue for purposes of studying how price and firm and characteristics influence a Western firm's propensity to bargain.

The study design proceeded in several stages. First, we selected the geographic areas of Vienna to be studied. Vienna is comprised of 23 districts. We chose four distinct geographic areas that vary in average annual net income per person: the 1st district (34,333 EUR, the wealthiest district), the 2nd (18,838 EUR, the 17th wealthiest district) and 20th (17,334 EUR, the 22nd wealthiest district) districts, the 18th (23,771 EUR, the 4th wealthiest district) and 19th (25,372 EUR, the 3rd wealthiest district) districts, and the 6th (21,989 EUR, the 12th wealthiest district) and 7th (22,659 EUR, the 8th wealthiest district) districts. The city's main shopping thoroughfare, Mariahilferstrasse, is located on the border between the 6th and 7th districts. Not only does the area around Mariahilferstrasse feature more stores than the other three areas, it is arguably the heart of commercial Vienna. We shall refer to the 1st district as Area I, the 2nd and 20th districts as Area II, the 18th and 19th districts as Area III, and the 6th and 7th districts as Area IV.

In order to construct a sample of stores to observe in each district, RAs were instructed to record the name of every retail store on the main thoroughfares of these areas that met certain criteria. That is, stores that were service-focused (e.g. restaurants, salons, etc.), stores that primarily sell food or beverages, pharmacies, and highly specialized stores (e.g. hearing aids, orthopedic shoes) were not considered. Furthermore, the second highest price of a store was required to be at least 120 EUR; this would rule out stores such as "Tabak" shops, for example. Approximately 750 stores were recorded in total; approximately 150 stores were recorded in each of the Areas I-III and approximately 300 stores were recorded in Area IV due to its importance as the primary commercial area in Vienna. The RAs assigned to this task were not subsequently assigned to ask for discounts at stores. Then, from each geographic area, a sample of 40% of the stores was selected at random for purposes of observation. Therefore the final sample consisted of 300 stores, although several stores needed to be discarded during the data collection phase for reasons to be noted later in the study.

We searched for RAs who would be asked to seek discounts at stores by posting an advertisement in the building of the Faculty of Economics and Business at the University of Vienna and by sending an email with the same advertisement to the co-authors' former students. We hired the first six male and first six female German-fluent RAs for which we were able to schedule an interview. In order to prevent intra-RA communication during the project, RAs were not told the names of the other students who were hired nor were their names ever displayed on group e-mails. These RAs' first task was to visit approximately 20 stores each in order to record several pieces of information about each store; we refer to this as the "Information Gathering" stage and an RA who recorded information from a particular store was referred to as an "Information Gatherer" (IG). These observations are summarized in the next section.

Then, three separate RAs were randomly assigned to each store in the sample using a stratified approach. More specifically, a random assignment was made with the following restrictions: each store was assigned a visit by at least one RA of each gender, each store was assigned to be visited at least once both before and after Christmas, an IG associated with a particular store was forbidden from visiting the same store as a bargainer, the observations of a given RA were divided roughly evenly across the four geographic areas, and each RA was assigned roughly the same number of observations before and after Christmas. Roughly half of the stores were assigned to be visited twice before Christmas and once following Christmas, and roughly half of the stores were assigned to be visited once before Christmas and twice following Christmas. If a store was assigned to be visited twice during the same shopping period, then assignments were made in order to ensure that no store would be visited twice on the same day. The size of the price range observed, the number of stores visited, the number of visits per store, the areas of the city observed, and the number of RAs employed were dictated by budgetary constraints.

Each assignment required the RA to find a product for which he or she could express credible interest within a given price range. That is, for each store there were three distinct price ranges assigned. The ranges were calculated using the lowest and second highest prices recorded during the IG stage. While one would certainly not expect prices to be distributed uniformly within a store, in order to approximate a representative price range in any given retail store we imagined that prices were distributed uniformly and calculated a threshold at the 20th percentile and 50th percentile of such a distribution. Our thinking was that the mass of prices is likely to be dense at the lower end of the distribution and less dense as price increases. In this way, we sought to capture categories that were roughly similar to terciles of a given store's price distribution.

For each assignment, RAs were given a store name, address, the nature of the products sold at the store (e.g. clothing) in order to prevent RA errors, the posted price range within which they should find a product, and the date range on which they should visit the store. If an RA could not easily find a product in the price range assigned, they were asked to find a price that was as close as possible to the range, but not below 30 EUR and not above 1000 EUR. In practice, it was uncommon for this to occur. Unfortunately, resource limitations prevented us from purchasing products if a discount was granted. For this project, we would have needed to spend 73,010 EUR, recalling that discounts in the amount of 7,715 EUR were granted on products that totaled 80,725 EUR in posted price. In addition to the large expense associated with purchasing products, it also would have been impossible to know which products would be discounted *a priori*, and therefore it would have been challenging to estimate an appropriate budget in such a case.

The fact that we allowed RAs to choose the product with which to bargain carried with it both advantages and disadvantages. First and foremost, it would have been very time-consuming and costly for us to pre-select products from each store. Furthermore, asking an RA to find a particular product with which to bargain carries several risks. As a practical matter, it may take a long time for an RA to find a particular product, and in some cases a pre-selected product may be sold out by the time of observation. Also, it is conceivable that a given RA might not be able to easily express credible interest for a pre-assigned product. For these reasons, it was not realistic to pre-select products for this project. However, with the appropriate time and resources this might be feasible and interesting for a future project.

Of course, a disadvantage of allowing RAs to select the product for which to ask for a discount is that the RA may somehow bias the sample according to products for which he/she thinks a discount would be granted. Along these lines, RAs were told that the only consideration that should be made when choosing a product is whether one could express credible interest in it. That is, neither should an item's "sale" status nor should an item's price level within the assigned range affect its potential for selection. RAs were instructed to always approach the nearest salesperson to the product which they had chosen regardless of the salesperson's level of responsibility rather than intentionally seeking or asking for a manager. In any case, when a salesperson referred an RA to a manager, they were instructed to note such an incident, although in practice this occurred rarely.

Practically speaking, bias could be potentially introduced due to RA selection of products if RAs systemically chose products based on an unobserved characteristic that affects the likelihood of receiving a discount. Along these lines, RAs were asked to record characteristics of the product and the salesperson that were observable to them that could have potentially influenced the probability of a discount (list price, sale amount, salesperson gender and age). If RAs believed that one of these characteristics made a discount more likely, then we would be faced with a less representative sample than desired but such RA preferences should not bias our results. We provide a detailed analysis of RA-specific outcomes in the Appendix in order to understand whether RAs were deliberately targeting particular types of products.

Along these lines, given that RAs performed their observations unsupervised, several measures were in place in order to minimize the possibility that RAs fabricated their results. First, as noted earlier, RAs were not told the identity of the other RAs hired in order to prevent fabricated results. Secondly, it was emphasized that the objective of the project was not to obtain a discount; rather, the objective was to express credible interest in a product in the manner of an interested consumer, and that as researchers we are interested in the outcome of such an interaction and not in maximizing the success of such an interaction. Note that providing an incentive for successful bargaining would also have created a perverse incentive for RAs to fabricate their results. Thirdly, RAs were informed that two additional RAs would visit the same store in order to implicitly remind them that the store's bargaining practices would also be documented by other individuals. And finally, RAs were told that at some point following the completion of the project that there was a possibility that the bargaining practices of the stores would be surveyed. In Section 5 we shall revisit the performance of RAs both individually and jointly at the store level.

Prior to the beginning of their work, each RA was individually trained by one of the co-authors. The training entailed discussing the guidelines for their work, answering their questions, and bringing them to five stores that would not be observed during the study in order for them to practice requesting discounts. After each practice observation, the RA discussed the interaction and any ambiguities surrounding the work was clarified. The training had several objectives. First, training was instituted in order to minimize the amount of learning that would occur once the actual work began. Second, since aspects of any interaction are unobserved in the data, the training sought to homogenize these interactions such that unobserved, non-random impressions or language used by the RAs would be minimized. Even so, we attempt to account for such unobserved aspects of the interactions by incorporating fixed effects at the RA level in our empirical analysis. As noted before, while deliberate variation in RA approach would be interesting to study in its own right, we felt that heterogenizing RA behavior would have compromised our ability to perform statistical inference on the issues of primary focus for this study given our sample size and the number of firm and product characteristics in which we are interested.

4 Empirical analysis

4.1 Bridging theory and empirics

The data collection yields two primary outcomes of interest - whether or not a firm granted a discount for a particular product, and if a discount was granted, the size of the discount. In our theoretical model, the decision regarding whether to grant a discount is governed by ω, the propensity to bargain. The variable ω depends on four primary variables: the average valuation between the posted price and marginal cost, the proportion of consumers who bargain (β), the ability of the firm to extract surplus (λ), and the firm's bargaining cost (b). Note that the average valuation between posted price and marginal cost is in general given by the integral in (4), but for the class of demand functions $(A-p)^{\eta}$ which includes linear, concave and convex demands, it is proportional to the absolute margin m. In what follows, we shall refer to the average valuation between the posted price and marginal cost as the margin, recognizing that in general that these two values may be different.

4.1.1 Margins

We begin by discussing the way in which we empirically proxy for margins. The proportion of consumers who bargain, β, is likely to be close to zero at most of the stores that we examine.[18] According to our theory, when β is small, the price and hence margin is almost entirely determined by the distribution of consumer valuations (F) and the marginal cost (c). While we do not observe these variables empirically, we do observe the list price, which corresponds to p_0 in the Appendix, and the posted price, which corresponds to p_1 in the Appendix. Therefore, in order to predict the relationship between the list price and the posted price on the one hand, and the propensity to bargain on the other, we understand variation in the list price and the posted price to be the outcome of underlying variation in F and c, which result in variation in the firm's propensity to bargain.

Both the list price and the posted price may serve as proxies for margins. First, for items that are not on sale, if it is the case that price and margins are correlated in the cross-section of product markets that we observe, then we may use price to proxy for margins. Second, in addition to list price, for items that are on sale, the amount of sale s, defined as $s = p_0 - p_1$, may also be informative about the margin. As discussed in Section 2 and in the Appendix, the relationship between sale amount and propensity to bargain may be positive or negative, but we do expect a negative relationship. This is because most plausible explanations for sales in our dataset are price reductions due to (i) demand shocks, (ii) price randomization or (iii) inter-temporal price discrimination. All three predict that observing an item on sale reduces the posterior estimate of ω, holding list price fixed. Therefore we shall employ the list price and sale amount as our proxies for retail margins. However, one could also perform the same inference using the posted price rather than the list price.

4.1.2 The proportion of consumers who bargain

When examining the determinants of a firm's propensity to bargain, in order to exclude the effect of β via price on the propensity to bargain from the effects of F and c via price, it is useful if we observe a proxy for variation in β in the data. In our case, we believe that the only product category for which β is likely to be non-negligible is jewelry, and therefore indicator variables for store type will be included in our empirical analysis. Our theory predicts that stores facing a non-negligible proportion of consumers who seek to bargain will set a higher price in order to profit from the larger proportion of bargainers. Therefore, if we view product categories as proxies for variations in β, it is important to account for product categories separately from price in the regression in

[18] See a discussion later in the section.

order to distinguish the effects of underlying demand and cost variation (via price) from the effects of β. Since β is positively correlated with margins (and by extension, the price level), a failure to include product category dummies would result in an inflated estimate of the effect of price on bargaining propensity.

In our analysis we classify observations into one of five categories: clothing, shoes/leather goods, jewelry, household, and other (summary statistics are provided in Table 2). We further divide each product category into two subcategories: low-priced stores within a category and high-priced stores within a category. We classify the price level of a store within a category according to the second highest price observed at that store. Our reasoning for this split is explained below.

Beyond capturing variation in the proportion of consumers that seek to bargain (β), an additional advantage of including variables that classify products into categories is that such variables may capture crude differences in competition, which are otherwise very difficult to measure given the heterogeneity of products in the dataset. Competition reduces margins in most theoretical models of firm behavior,[19] and therefore to the extent that two products share the same price but a product for which more competition exists has a higher marginal cost, product category indicator variables will seek to control for differences in margins across industries. What is important to remember is that we are somewhat less interested in the coefficient estimates of the product indicator variables per se than we are in controlling for factors that may result in mismeasurement of the price variable coefficient estimates.

Finally, the likelihood that a particular RA receives a discount may depend on a clerk's judgment regarding that RA's willingness to pay for a product, namely, whether the clerk judges the RA's valuation to be below the posted price. Thus, it may be that the probability that an RA receives a discount is increasing in the posted price simply because for higher priced items, an RA's willingness to pay for the product (as estimated by a clerk) is relatively low. If this were the case, then RAs would be more likely to receive a discount at high-priced stores within a category. Therefore, dividing each product category into low-priced stores and high-priced stores may capture such a phenomenon.

4.1.3 Ability of the firm to extract surplus

We utilize data on the number of stores owned by the same firm as our primary empirical proxy for λ, the salesperson's ability to capture surplus via bargaining. Our reasoning for using the number of stores as a proxy for λ is that owners of firms with a larger scale generally have more difficulty monitoring and incentivizing their salespeople. Even though some large stores may partially overcome this problem by incentivizing clerks with sales commissions, such a tool does not eliminate the basic conflict between the owner's

[19]In our monopoly model competition is captured by a leftward shift in the distribution of valuations.

desire to maximize profits and the clerk's desire to maximize revenue. For example, clerks who sell on commission may grant discounts to too many consumers, thus extracting less than the profit-maximizing amount of surplus. Recall that in our theoretical model, λ is determined exogenously, and therefore it essentially follows from the firm's decision regarding the scale of the store. As two alternative measures for λ, we will also distinguish small scale firms from large scale firms according to a) whether the firm in question is a multinational firm b) the physical size of the store.

We also examine observation-specific salesperson gender and salesperson age, as these two variables may be loosely associated to the firm's ability to extract surplus. As a practical matter, young salespeople and saleswomen are observed at much higher proportions at large-scale firms than at small-scale firms.

4.1.4 The firm's bargaining cost

In our context, we consider the cost of bargaining to be any time or expense that is incurred upon agreeing to grant a discount. These costs may entail the time and effort related to the administrative procedures necessary to sell the product at a discount as well as lost revenues due to the possibility of other customers witnessing the interaction and asking for a comparable discount. We include two variables that we believe capture the cost of bargaining b in our regression: the number of employees observed and the number of customers observed entering the store during a three minute period. These observations were recorded by an information-gathering RA at the beginning of December 2013, approximately one week before the bargaining observations commenced. The advantage of recording these observations separately from the bargaining observations lied in the fact that a given bargaining observation was not burdened by an excessive amount of data collection, and therefore an RA was allowed to focus on selecting a product of interest, asking for a discount from a salesperson, and recording information from the interaction. The downside of this approach is that the number of employees observed and the number of customers entering the store may have differed in the information-gathering period from the bargaining periods, although the largest potential differences would likely be associated with the number of customers that were observed.

The reasoning behind observing the number of employees is that in environments in which some customer service is provided (supermarkets are an example of an environment in which there is little to no customer service offered), the number of employees is often endogenously chosen according to the number of customers in the store. As the average number of customers increases, the employee to customer ratio typically decreases. That is, at least one employee is required in any store even if no consumers enter the store, whereas stores with large demand are likely to have a employee to customer ratio that is less than one. Therefore the opportunity cost of an employee granting a discount due to bargaining is likely to be lower at a store with a small number of employees. However, one difficulty associated with this variable is that the number of employees may also be a signal of the firm's ability to extract surplus, which is captured by the parameter λ in our theoretical model.

We should note that computing employee-customer ratios using our measure of customers observed is problematic given the potential variation in customers observed in the IG stage relative to the bargaining stage. This is particularly true for observations in which there are several employees and a small number of customers observed. An additional possibility would be to construct a variable that indicates the difference between the number of employees and number of customers observed; however such a variable suffers from the drawback that a given difference at a store with a small number of employees and customers would be treated in the same manner as the same difference at a store with a larger number of employees and customers. In practice, there is likely to be a difference in bargaining costs on behalf of the salesperson across two such observations.

Our primary interest in observing the number of customers that enter in a three minute period was to distinguish between stores that were observed to be empty with stores that were observed with customers. If at least one other customer is present in a store, a firm that would grant a discount when asked may give no discount smaller discount (or no discount) when other customers are watching.

4.1.5 Other controls

We also include controls which relate to our theoretical model but for which predictions are less clear. We collected data in two periods, pre-xmas and post-xmas, because we wanted to test for the possibility that stores view their inventory post-xmas as sunk; if this were the case, marginal costs for stores would effectively decrease following Christmas, and by extension sale items would not be less likely to be granted a discount post-xmas.

In our study design we assigned RAs to collect observations in different areas of the city in order to ensure that the bargaining behavior we observe was not limited to stores that typically face one type of clientele, as measured in terms of willingness to pay or in terms of the curvature of the demand curve. Therefore we include geographic area effects in our analysis.

Although we sought to minimize the extent to which RAs gave heterogeneous impressions to the salesperson, it may be that the salesperson's judgment of RA's willingness to pay varied across RAs. Therefore we will account for fixed effects for RAs. This issue is discussed in the context of the theoretical model in Section 2.

4.1.6 Discount amounts

Our theory predicts that the size of the discount granted will only depend on margins, the RAs perceived willingness to pay, and the nature of consumer demand. More specifically, a sufficient condition for discount size to rise faster than marginal cost is if the conditional mean of A given c rises faster than c, per (12). Since this would also imply that price rises faster than costs (per (10)), then it would follow that discount size increases with price, as discussed in Section 2. And as can be seen in (11), the amount of a discount may be individual specific.

When β is not small, variation in b and λ also influence the size of the discount. However, just as we do not interact proxies for λ and b with β in the participation equation due to sample size considerations, we refrain from doing so in the amount equation as well. Unlike in the participation equation, when β is close to zero, the effect of b and λ on the discount size is negligible.

Therefore, in our amounts regression we use all proxies for the margin as well as RA fixed effects. The margin-related variables include the list price, the sale amount, and industry controls which partially capture crude differences in competition across industries, as discussed earlier.

Finally, we also take into account the fact that a smaller discount may be granted if another customer witnesses the interaction. While such a prediction is outside of our model, we nevertheless account for the number of customers observed in the store for purposes of controlling for such an effect.

4.2 Summary statistics

In Table 1 we report summary statistics for the quantitative variables of primary interest in our analysis: list price, posted price, discount amount, and discount percentage. The posted price is the price that the consumer would pay without explicitly asking for a discount. If the product is on sale, the posted price will be below the list price. Otherwise, the posted price is identical to the list price. Note that the maximum list price of 5900 EUR was recorded for a carpet, which was associated with the maximum sale amount and sale percentage of 5240 EUR and 89%, respectively. The next highest list price was 1299 EUR. Removing this observation has very little effect on our forthcoming results. The third and fourth row of Table 1 report summary statistics of discounts granted due to bargaining; these cases are conditional on a discount being granted and are calculated off of the current price. One may infer from the table that 303 of 751 observations entailed a discount being granted due to bargaining and that 169 of 751 items in the data were on sale.

Figure 1 illustrates the distribution of discounts in absolute and relative size. There were no discounts granted in approximately 60 percent of the observations, as can be seen in Figures 1A and 1C. Furthermore, note that Figure 1B only uses strictly positive discounts as the horizontal axis measures the log of discount amounts. In Figure 1C, horizontal lines were added to show the relatively higher frequency of 3 percent, 5 percent, and 10 percent discounts. Figure 2 illustrates the distribution of list prices observed in absolute and relative sizes, and Figure 3 does the same for posted prices (after any sale reductions).

We can generally divide our variables into two categories - store-specific variables and observation-specific variables. Store-specific data that relate to characteristics observed at the store itself were recorded during the weeks prior to the bargaining interactions. Store-specific data that relate to institutional characteristics of the store, such as multinational presence and the existence of additional branches of the same store, were recorded during the months following the bargaining interactions. Of course, observation-specific variables pertain to the bargaining interaction itself, and these data were recorded immediately following an interaction. In Table 2 we report summary statistics for our store-specific variables. The authors of this study made a subjective store category designation for each store in the absence of any preferable, more objective available designation.

Table 3 lists all of the multinational firms in our dataset and the incidence with which a discount was granted at each firm. We define a firm as multinational either if it owns stores both within Austria and outside of Austria or if it franchises its stores internationally. Whereas the names of domestic firms may not carry much meaning for readers outside of Austria, the recognizable names of many of the multinational firms should assist the reader to understand the types of large-scale firms that were observed in our study.

4.3 The empirical model

Recall once again that we recorded two primary outcomes - whether or not a firm granted a discount for a particular product, and if a discount was granted, the size of the discount. The first outcome is essentially a participation decision - whether or not to grant a discount, and the second outcome is an amount decision. Generally speaking, when data is censored, a threshold value is observed for each observation at or beyond the threshold - this may occur, for example, when a maximum income is reported for each observation in which income exceeds a particular value. In our data, the preponderance of observations in which no discount is granted is not strictly speaking an issue of censoring. Therefore we adopt Wooldridge (2010) slightly more appropriate "corner solution response" terminology for purposes of understanding our data, because we do observe the entire possible range of the response variable.

Most commonly, the type I Tobit model has been employed in order to address corner solution responses. However, one important restriction of the type I Tobit model is that the partial effects of an explanatory variable on the participation decision and the amount decision must have the same signs. In principle, our theoretical framework allows for the possibility of signs to differ across the two equations for a particular variable, and therefore it would appear to be more appropriate to utilize a more flexible model, whereby separate mechanisms are allowed to determine the participation decision and the amount decision.

Let d be a binary variable that indicates whether the observed discount size z is zero or strictly positive, and let z^* be a nonnegative, continuously distributed latent variable that represents the size of the discount that the salesperson will offer. Then we may write:

$$z = dz^* \tag{13}$$

When a discount is strictly positive, $d = 1$ and $z = z^*$; otherwise $d = z = 0$. Note that in principle that the salesperson's discount offer may not be their final offer, but uncovering this information would entail a much longer bargaining interaction. Of course, as noted earlier, a drawback of a more protracted bargaining is that the firm's perception of the consumer's discount factor would introduce additional complication into the analysis. So while it is possible that z is underestimated in our framework, it would seem that such underestimation is minimal because in most cases it is reasonable to assume that in a typical retail environment, persisting to bargain would not yield a greater discount than the one initially received. Obviously, if the environment we were examining was one in which bargaining is the rule rather than the exception, as is the case at car dealerships and flea markets, such underestimation of z due to our bargaining instructions would be more severe.

The appropriate model for analysis depends in part on our assumptions regarding the relationship between d and z^*. If they are independent conditional on a set of explanatory variables, then we may analyze the firm's decision using what is commonly referred to as a two-part model. However, if some common unobserved factors affect both d and z^*, then one should consider analyzing the firm's decision using what Wooldridge (2010) refers to as an exponential type II (ET2T) model. Here, we shall investigate both possibilities and compare our results.

First, let us suppose that d and z^* are independent conditional on a set of explanatory variables, and the participation decision is modeled in terms of the probit model, where we denote r as a vector of covariates affecting whether or not a particular firm will choose to bargain for a particular product if $\omega = r\alpha + v > 0$:

$$P(d = 1|r) = \Phi(r\alpha) \tag{14}$$

Note that the model cannot predict negative outcomes for z because the support of z^* is $(0, \infty)$. One possibility is to define the amount equation as:

$$z^* = x\beta + u \tag{15}$$

where u given x follows a truncated normal distribution with lower truncation point $-x\beta$. This equation, together with (14), is commonly referred to as the truncated normal hurdle

(TNH) model and was first proposed by Cragg (1971). Another possibility is to define

$$z^* = e^{(x\beta + u)} \tag{16}$$

where u given x follows a normal distribution with mean zero and variance σ^2. Also proposed by Cragg (1971), this equation together with (14) is commonly referred to as the lognormal hurdle (LNH) model. Here, we may express z as:

$$z = dz^* = 1[r\alpha + v > 0]e^{(x\beta + u)} \tag{17}$$

where v is unobservable with a standard normal distribution, u and v are independent, and (u, v) is independent of x with a bivariate normal distribution. In this case, due to our previous assumption on the distribution of u, we may say that $z^* = e^{(x\beta+u)}$ has a lognormal distribution and z conditional on $(x, z > 0)$ has a lognormal distribution as well because we assume that the errors of the participation and amount equations are independent of each other. Whether the amount equation is modeled according to (14) or (16), the participation equation and the amount equation may be modeled independently from one another.

If, however, we relax the assumption that $Cov(u, v) = 0$, then we may modify the lognormal hurdle model in order to obtain the ET2T model. In this case, note that we may not use (15) as the amount equation because the ET2T allows for negative outcomes on s; this would be a particular concern if $Corr(u, v) = \rho$ is estimated to be negative, as $E[log(z)|x, z > 0] = x\beta + \rho\sigma\lambda(r\alpha)$, where $\lambda(\cdot)$ is the inverse Mills ratio, $Corr(u, v) = \rho$, and $Var(u) = \sigma^2$. Therefore we may only apply the type II model to $ln(z)$. Of course, this is only a potential problem if the amount equation is expressed as (15) rather than (16). It is rather obvious but nevertheless useful to note that the LNH model is equivalent to an ET2T model in which ρ is constrained to equal zero.

However, the literature has noted that the more general ET2T model also carries a potential risk of poor identification of the amount equation. In particular, in cases when $r \equiv x$, $Cov(u, v) = \rho\sigma$ is not separately identified by $E(z|x)$.

4.4 Empirical results

We specify the participation equation using (14). The right-hand-side variables of the participation equation, introduced earlier in this section, are listed in Table 4. We also include three interactions that are not shown in Table 4 but which we will report separately: list price with firm scale (our proxy for λ), firm scale with sale size (s), and sale size with observation period (pre-xmas or post-xmas). In our forthcoming discussion of the results we shall explain our interest in these interactions.

We employ two approaches for specifying the amount equation. One approach entails including all of the variables contained in the participation equation in the amount equation. We refer to this as the full amount equation. No interactions are included in the amount equation, however, due to sample size considerations. Furthermore, the theoretical predictions with regards to our interactions in the participation equation are ambiguous with respect to discount size when β is close to zero. A second approach entails including variables in the amount equation for which we have predictions regarding discount amounts according to our theoretical model. We refer to this as the parsimonious amount equation.

We analyze the participation and amount equations using the TNH model, the LNH model, and ET2T model. Because these models are non-nested, we may apply Vuong's (1989) test to check whether the difference in log-likelihoods between the TNH model and the LNH model is statistically significant. The test finds that the average difference in the log likelihood between the LNH model and the TNH model is 0.163 and statistically significant at the .01 level for the full specification. This is also the case when applying the Vuong test to the parsimonious specification. Therefore, because we have strong evidence that the TNH model is inappropriate for our empirical application, we only report the results of the LNH model and ET2T model in Table 4. It is also interesting to note that a Type I tobit using discount amount as the dependent variables and the right-hand-side variables from the full specification yields a log-likelihood value that is significantly lower than the TNH model, further evidence that a flexible two-part specification is most appropriate in our setting.[20]

The results of the four specifications in Table 4 are qualitatively quite similar. Therefore, before discussing the effects of individual variables, it is useful to compare the specifications in an effort to understand which specification is most appropriate for our analysis.

When we include the same variables in the participation and amount equations, we obtain a negative and significant estimate for ρ. While such an estimate is not implausible, it is somewhat suspicious because upon removal of variables in the amount equation for which we have no theoretical predictions in terms of discount size (when β is close to zero), we do not reject the null hypothesis that $\rho = 0$. Furthermore, when these variables are included in the amount equation, none of them are estimated to be statistically significant. Estimates of $\rho = 0$ are provided at the bottom of Table 4.

Given that the LNH model fits better than the TNH model, and given that we cannot make a strong statistical claim that $\rho \neq 0$ nor can we easily address the difficulties of properly identifying ρ when the amount equation contains many of the same variables as the participation equation, in what follows we shall refer to the results of the LNH model. Since we are not faced with the identification concerns associated with the ET2T

[20] Applying a χ^2 test with a number of restrictions equal to the number of variables in the Tobit model, the LR statistic is 2(LL of Tobit - LL of TNH)=2(-1649+1523)=252, which has a p-value of essentially zero. Therefore the TNH model is a superior fit than the Type I Tobit model.

model when using a model in which we constrain $Cov(u, v)$ to equal zero, we report the results of the full specification when conducting robustness checks later in the section. In practice, estimates of the LNH model are very similar to estimates of the ET2T model in all specifications to follow, and estimates obtained from the full amount equation are very similar to the estimates associated with the same variables in the parsimonious amount equation.

We now proceed by interpreting our empirical results in the context of the theoretical framework in Section 2 by examining Table 4. First, we notice that discounts are more likely to be granted for higher priced products. In Table 4 we divide prices into quartiles; this allows for a certain degree of flexibility in estimating the relationship between price and bargaining propensity. This finding is consistent with the theoretical notion that margins are increasing with costs. Indeed, while it is certainly plausible that margins decrease as cost increases for a particular good, it would be surprising to observe such a phenomenon in a diverse cross-section of products, as this would imply that 1000 EUR items typically have smaller absolute margins than 30 EUR items (whose margins may be no greater than 30 EUR). Furthermore, the results from the amount equation indicate that the size of discounts granted are roughly proportional to a given product's list price. When using $ln(listprice)$ instead of list price quartiles in an alternative specification in Table 5, the elasticity of discount size with respect to list price is estimated to be nearly unitary (the 95 percent confidence interval is estimated as [0.811,1.009]). In other words, the percentage discount given by a firm is predicted to be nearly constant at all price levels in the data. This finding in the amount equation is consistent with the participation equation result that discount probability increases with price, as only an increasing relationship between margin and price would allow for the percentage discount to remain constant as price increases.

We may also draw conclusions regarding percentage margins. Recall from our theory that under the assumption that RAs occupy a fixed percentile position in the distribution of consumers between price and marginal cost,[21] the amount of a discount is proportional to the absolute margin. In this case our estimate of the elasticity of discount size with respect to price suggests that percentage margin is nearly constant across the range of list prices.

We categorize sale items according to the absolute size of the sale. That is, we distinguish between 85 observations for which the sale size is up to 45 EUR from 84 observations for which the sale size is greater than 45 EUR. We find that both categories of sale items are significantly less likely to earn a discount than non-sale items by a substantial difference. Such a finding would be consistent with the hypothesis that for products with the same list price, absolute margins are lower for sale products. However, we do not find a

[21]This may be the case if RAs were judged to be average consumers with enough willingness to pay to afford to pay marginal cost but not the posted price.

significant difference in probability of earning a discount between the two categories of sale items.

Given the period for which we collected data, one question that arises is whether sales immediately after Christmas reflect firms' desire to clear inventory. If one assumes that sales that occur 2-3 weeks prior to Christmas do not reflect a desire to clear inventory, we may get a sense of whether sales during the period after Christmas reflect firms' desire to clear inventory by examining whether firms were more likely to bargain on sale items after Christmas. Figure 4C shows that the predicted probability of receiving a discount after Christmas is not higher than before Christmas for non-sale items or sale items. This is not very surprising, as it is more realistic to expect that season-ending inventory clearance sales begin in earnest later in January, continuing into February.

Unfortunately, we only observed 26 sale items for which a discount is granted, and therefore our ability to draw conclusions regarding the relationship between sale size and discount size is limited. However, there is some evidence from the amount equation that the very small subset of sales for which bargaining was successful may have included items that were on clearance-related sales. We observe that amongst items which are granted discounts, items that we categorize as sold at a "small" sale size are granted discounts that are approximately 55% smaller in magnitude than items that are not on sale, which is consistent with our finding from the participation equation. However, large sale size items are estimated as generating discounts that are approximately 80% larger than small sale items. One possible explanation for this finding is that items designated for large sales and for which discounts are granted are different in nature than other sale items in that they represent inventory that the store seeks to clear.

As expected, we find that discounts are significantly less likely at large firms overall, and this finding is robust to various thresholds established for a firm's number of stores. This finding is also robust to our specification in which we consider domestic firms to be small-scale firms and multinational firms to be large-scale firms. This follows directly from our theory, which predicts that the propensity to discount increases with λ, the salesperson's ability to capture surplus via bargaining.

Our theory also predicts that the effect of λ on the probability of obtaining a discount will depend on the price of the product. In particular, the theory predicts that as the size of the margin increases, the difference in the probability of obtaining a discount from a firm with a high value of λ and obtaining a discount from a firm with a low value of λ will increase. This implies that it would be appropriate to interact the list price with our measure of λ because we claim that the list price level serves as a proxy for the size of the margin, as discussed earlier. We therefore interact the price quartile variable with the indicator for how many of the same stores are owned by the same firm. We find that as the price increases, the difference in probabilities of obtaining a discount between stores who are effective at extracting surplus versus stores that are less effective at extracting surplus widens. Note that for the highest price quartile, these differences are statistically different for most λ. This can be seen from our main specification in Figure 4 as well as when using alternate ways of specifying the list price (Figure 5) and alternate ways of specifying λ (Figure 7). In Figure 7C we display results related to a specification in which we proxy for λ according to the number of seconds required by an RA to walk through the store premises, a measure for the physical size of the store. At stores that we examine (in which there are no car dealerships), it is possible that organizational structures in which the salesperson's ability to extract surplus is poor are likely to characterized by large physical premises. Therefore an alternative possible measure for λ is a store's physical size.

An additional prediction from our theory is a second-order effect relating to the interaction between λ and the size of a sale. That is, while any sale reduces the propensity to bargain, an increase in the amount of a sale reduces the propensity to bargain of small firms (for whom λ is presumably closer to one) by a larger amount than for large firms (for whom λ is presumably further from one). Therefore we would expect that small firms' decreases in propensity to bargain will be more substantial than large firms' decreases in propensity to bargain as the size of a sale increases. This may be observed in Figure 4B, in which the slope of the line representing large firms that stretches from 0 to 1 on the x-axis is .08 larger than the slope of the line representing small firms over the same horizontal domain; from 1 to 2, the slope of the line representing large firms is .23 larger than the slope of the line representing small firms. We can also reject at the .01 level a joint hypothesis test that these differences of the slopes (.08 and .23) are both equal to zero.

We predict a nearly 70% probability of earning a discount on what we classify as expensive jewelry - the jewelry items with a list price in the top half of the distribution of list prices for jewelry products. No other product category exceeds a predicted probability of 45%. In addition, jewelry is the only product category for which RAs are predicted to earn a discount with a significantly higher probability for more expensive stores within the category. Whereas it is anecdotally unusual for consumers to ask for discounts for most of the products that we observe in our dataset, we suspect that consumers do occasionally seek discounts on jewelry. In the context of our theory, this would imply a higher β for jewelry than for other product categories. Therefore, using our theory, we may interpret this empirical finding as evidence that a jewelry item at the same list price as a product in a different product category is more likely to earn a discount due to the fact that a higher β for expensive jewelry implies a higher margin for the same observed list price.

It is important to note that one potential explanation for RAs receiving discounts with a higher probability for products at expensive jewelry stores relative to inexpensive ones is that the perceived percentile at which the RA is located within the distribution of customer valuations may be lower at high-priced jewelry stores. While in principle we are unable to disentangle the effect of β from the effect of RAs falling in the perceived distribution of customer valuations as prices increase, this explanation seems unlikely because this is not observed for any other product category, and we see no obvious reason why RAs would fall in percentiles for jewelry but not for any other product category.

Recall that our theory predicts that a higher bargaining cost on behalf of the firm will decrease a firm's propensity to bargain. Empirically we find that the number of customers present, which we view as a proxy for the firm's cost of bargaining, does influence the firm's bargaining behavior. As noted previously, while these customer observations were recorded separately from the bargaining interactions, nevertheless there should be a high degree of correlation between the number of customers observed two weeks prior to the first bargaining period and the number of customers present at the time of bargaining. When classifying the number of customers in categories, we predict that positive discounts are significantly smaller in size at stores at which no customers were observed. When classifying the number of customers as a quantitative variable, we find that an increase in the number of customers reduces the likelihood that a discount will be granted. Both findings are consistent with our theoretical prediction that the presence of other customers increases the salesperson's cost of bargaining; and therefore at stores at which there are typically a very small number of other customers, on average the salesperson's cost of bargaining will be lower. With respect to our other proxy for a firm's bargaining cost, the number of employees observed at a store, stores in which one employee is observed employees are predicted to grant discounts with a higher probability than stores that employ four or more employees.

Indicator variables for the district of Vienna in which the store is located and the age range of the salesperson were found to be insignificant. Examining the 12 fixed effects for RA identity, we find that of the 66 comparisons across all possible RA pairings, only 5 of the 66 fixed-effect pairs were significantly different than one another. In the amount equation, 10 of the 66 fixed-effect pairs were significantly different than one another. In other words, RAs appear to have been perceived quite homogeneously by the salespeople who were approached for this study.

It is worthwhile to note that our theory also predicts that β will affect the propensity to bargain via the price, λ, and b. Given that we only believe that β is meaningfully large enough for expensive jewelry stores, unfortunately we have too few observations to interact jewelry with our other variables.

5 Research Assistant Behavior

In a study such as this one, not only are RAs assigned to collect observations in the field, the outcome of their observations is partially dependent upon their own behavior. Therefore the "human element" is a non-trivial issue in our study. While the previous section attempted to capture the salesperson's perception of an RA via RA fixed effects and RA gender effects, in this section we turn to examine the behavior of the RAs. While this is not the focus of our study, it is nevertheless important to understand whether RA behavior jeopardized the integrity of our analysis.

5.1 RA performance over time

5.1.1 Hot hand

We perform several tests in order to examine the extent to which discount outcomes for a particular RA vary over time. First, for each individual RA, we run a simple vector autoregression (VAR) in which the dummy variable indicating whether a discount was granted is the dependent variable. On the right hand side we analyze five lags of the dependent variable as well as all of the explanatory variables in Table 4 (but without any interactions). In addition to the fact that all lags of the dependent variable are found to be insignificant for the majority of RAs, lag-order selection statistics using the Akaike Information Criterion (AIC) support the inclusion of no lags for 9 of the 12 RAs, one lag for 2 of the 12 RAs, and two lags for 1 of the 12 RAs. When rerunning the VAR for these three RAs, only the second lag for one of the RAs is significant (and negatively so). Therefore we have strong evidence that previously observed discount outcomes do not affect an RA's subsequent discount outcomes.

Next, we run a Prais-Winsten regression using the same variables (again without lags of the dependent variable and without interactions) in order to check for the presence of unobserved serial correlation; in one instance we combine all observations and include fixed effects for each RA and we also run 12 separate regressions for each RA. In no cases do we find evidence of any type of serial correlation.

5.1.2 Learning and fatigue

In order to address the question of whether there is a deterministic trend with regards to the rate with which RAs obtain a discount, we run a simple linear probability model with the same regressors as in the previous specification and also include a variable that tests for the existence of a deterministic trend, the order in which the RA visited the store. We run this regression separately for each of the 12 RAs. The trend variable is negative and significant for two of the RAs at the .05 level (magnitudes of -.014 and -.021), one of the RAs at the .1 level (magnitude of .013), and insignificant for the remainder of the RAs. A positive and significant estimate might have suggested evidence of learning. We cannot rule out the possibility that a negative coefficient might suggest that a particular RA exerted less effort in obtaining a discount over time.

5.2 Effect of first RA visit on subsequent visits to the same store

It is also worthwhile to check whether a particular RA, by visiting a store first, somehow contaminated the observations that followed at the same store. This might occur if, for example, an RA inadvertently angered a salesperson during his or her interaction by asking for a discount (for one of a variety of possible reasons). This could result in a very small number of discounts granted at stores at which that particular RA visits prior to other RAs. Although our sample size is relatively small for any given "leading" RA, restricting our analysis to stores which appear three times in our dataset in Table 9 we do not find evidence that a particular RA visiting a store first led to a substantially lower incidence of discounts afterward at the same store.

In Table 10 we restrict follower RAs to those that observed a store within the same week as the leading RA. When analyzing the data in this fashion we are dealing with a particularly small sample size, but we report these summary statistics nevertheless. In Columns 2 and 3 we report the outcomes related to observations at the same store between Dec 9 - Dec 14 and in Columns 4 and 5 we report the outcomes related to observations at the same store between Dec 27 - Jan 4. RAs who follow RA #10 do record a particularly low number of discounts in both periods (2 of 10 observations and 0 of 7 observations, respectively), however this does not appear to occur for any of the other RAs.

5.3 RA product choices

It is also important to check whether RAs sought to choose products that they thought would generate a discount. The RAs were given the freedom to choose any product within a particular price range. This means that there were two variables which the RA could in effect influence in their search for a product - the product's price and whether the product was a sale item. Our concern would be if there was a large degree of heterogeneity in the percentile of the price range chosen across RAs; this might imply that individual RA effects could contaminate the price coefficients, and vice-versa. Likewise, if individual RAs chose a significantly different percentage of sale items relative to other RAs, the same concern would apply. While the nature of the product chosen is also up to the discretion to the RA, the degree of heterogeneity of products in our data set is far too large to investigate systematic biases in this direction. In fact, any combination of variables which we do not interact but for which RAs favored (e.g. high prices in the second period) might lead to spurious estimates.

In Table 8 we display summary statistics of the average price percentile at which each RA chose a product within the assigned price range. While RAs #1, #4, and #6 chose relatively low prices within the assigned range on average, none of the remaining RAs' price percentiles chosen are significantly different from one another. Here we use the term "percentile" loosely as we do not know the distribution of prices at a given store within a particular price range; however we surmise that in any given price range that there will be typically be a larger number of products at the lower end of the range than at the higher end.

In Table 8 we also display summary statistics of the average number of instances in which an RA chose a product that was on sale. Eight of these 66 pairs' price percentile choices are statistically significantly different from one another, and seven of these eight pairs include either RA #1 or RA #2, who observed sale items the least frequently.

Of secondary concern is whether RAs favored products which they (correctly or incorrectly) thought would generate a discount via price preferences or sale/non-sale preferences. While this would not influence our coefficient estimates due to the fact that we control for these variables, this would influence the overall incidence of discounts. This is not the primary interest of our study, however we did seek to construct a representative sample of stores and products of a certain nature that retail consumers face in the West; therefore we are interested in how representative our prediction of bargaining is evaluated at the average values of the variables.

5.4 Pairwise behavior of RAs / honesty

Although we attempted to conceal the identities of the RAs from one another during the project, it is possible that some of our RAs knew the identities of other RAs working on the project during the data collection period. If this were the case, one concern that might arise is that an RA who was assigned the same store as another RA would use data collected by that RA in order to fabricate what he or she considered to be realistic results from that store without actually carrying out the observation assignments. This would only be a concern in the second period because first period data were submitted prior to the beginning of the second period; during the first period the median number of stores for which a given RA pair was mutually assigned was only one, the mode was zero, and the maximum number of mutual store assignments was seven.

Over all observations, a given RA pair was allocated a median of 10 mutually assigned stores, with a maximum of 17 mutually assigned stores. For each of the 66 RA pairs, we calculated the percentage of observations for which they both recorded a discount for the same store, the percentage of observations for which they both recorded not receiving a discount for the same store, and the percentage of observations for which one RA received a discount and the other RA did not receive a discount. On average, both RAs received a discount approximately 26 percent of the time whereas neither RA received a discount approximately 46 percent of the time. There were no RA pairs which always recorded a discount at mutually visited stores and one out of 66 RA pairs always recorded not receiving a discount at mutually visited stores (8 stores, or 16 total observations). Therefore we feel comfortable claiming that there does not appear to be evidence that

RAs fabricated observations based on data collected by a fellow RAs.

It should be noted that a priori it is not entirely obvious that cheating is more likely to manifest itself as the same observation at the same store. For example, if RAs were concerned that the co-authors would analyze the data for cheating, they might try to avoid recording the exact same outcome for all mutually visited stores.

5.5 Missed observations and errors

We now address the observation assignments that were not performed and the errors that RAs in the observations that they recorded. Our analysis uses 751 out of 861 total observation assignments. Of the 110 observations that we do not analyze, 86 observations were never performed. Of these 86 observations, in approximately 75 percent of these cases the RA visited the store when it was closed. In certain cases this was avoidable, as a visit to the store's website or a phone call would have indicated the shop's hours. This was usually the case when the RA visited on a Saturday when the shop was closed. In other cases, the shop closed during the week following Christmas, and advance information in this regard was not always obtainable. In yet other cases, stores were closed during times for which it should have been open according to its own publicized business hours. In some cases reasons were given on the storefront (e.g. illness), but in most cases no explanation was given. Nearly all of the remaining 25 percent of observations were not performed either because the RA could not find a product with a posted price for which he or she could credibly bargain or because the RA did not properly locate the store for one reason or another. Furthermore, there were five observations that were mistakenly unassigned. One characteristic of nearly all of these 86 observations is that they pertain to small-scale stores, and given that small stores are more likely to give a discount, 40% may underestimate the probability of obtaining a discount.

The 24 observations which were recorded but which were not analyzed contained errors, ambiguities, inconsistencies, or applied to stores which in retrospect should not have been included in the sample due to the nature of the product sold. For example, one store's focus is bathroom and pool installation (3 obs), and another store is a non-profit organization that sells used clothes (3 obs). Eight observations utilized products outside of the instructed posted price range of 30-1000 EUR, four observations were recorded on products for which there was no posted price, and eight observations were recorded with ambiguities or inconsistencies (e.g., a discount was granted but the size of the discount was not stated explicitly or recorded).

6 Conclusion

Unlike previous studies, our study theoretically and empirically analyzes how price and firm characteristics will influence a firm's willingness to bargain. In order to examine this issue, we chose an empirical setting in which one rarely observes consumers asking for discounts due to our interest in understanding whether it is correct for a consumer to presume that an attempt to bargain at a "typical" retail store in the Western world will be unsuccessful. To the contrary, not only do we find that retail firms in Vienna grant discounts in 40 percent of instances in which a discount is requested, we find that price and firm characteristics are strong predictors of a firm's likelihood of granting a discount.

In particular, a firm's willingness to bargain will depend on its margins, its ability to extract surplus from consumers, and its cost of bargaining. We proxy for these variables empirically by utilizing data on observed prices and firm characteristics. In fact, while the relationship between price and a firm's propensity to bargain is interesting in its own right, we also demonstrate that asking firms for discounts is a vehicle for inferring the relationship between price levels and margins in the absence of cost data.

Outside of automobiles, this is clearly only a first step towards understanding the circumstances under which a retail firm will agree to bargain with a consumer. In order to further address this question, it would be useful to obtain more detailed product and firm information, seasonal variation in sales activity, usage of cash discounts, and cross-country differences. It would also be useful to collect more detailed information directly from firms regarding what they would claim to be their bargaining practices.

Furthermore, we cannot completely answer the larger question of why bargaining is generally not observed in the West, nor can we make strong statements regarding whether bargaining at retail stores would be welfare-enhancing relative to a uniform price policy. In order to explore this larger question, we believe that retail consumers should be approached in the field.

Figure 1: Quantile plots

(a) Discount amount

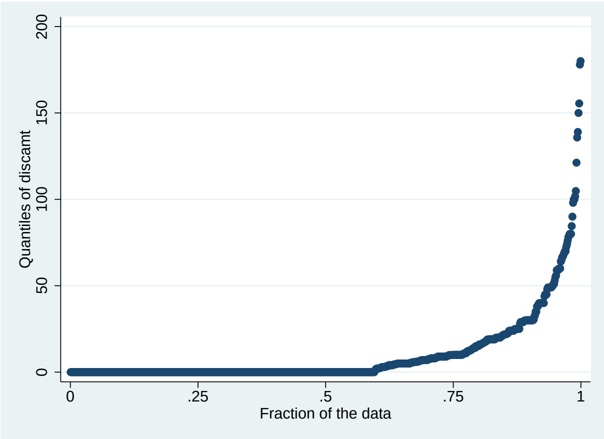

(b) ln(discount amount)

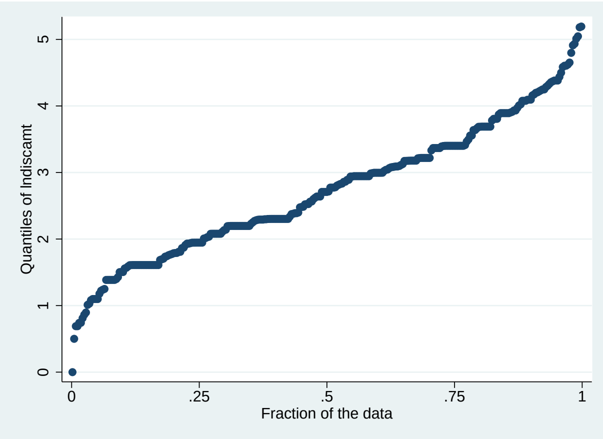

(c) Discount percentage

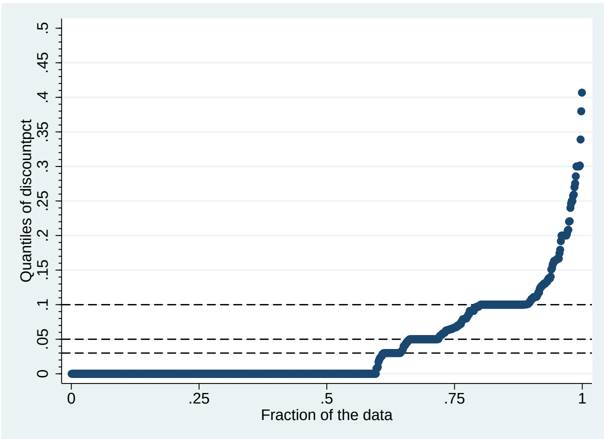

Figure 2: Quantile plots

(a) List price

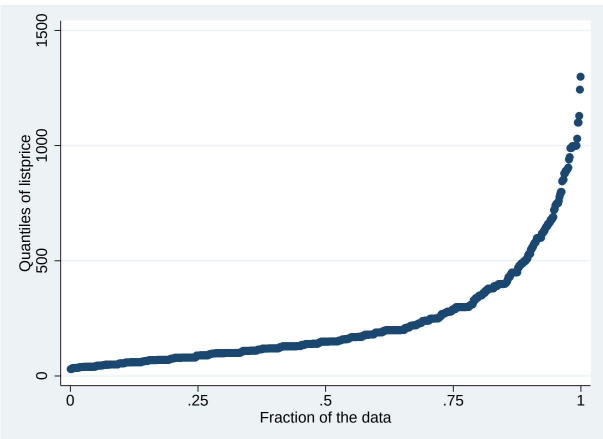

(b) ln(list price)

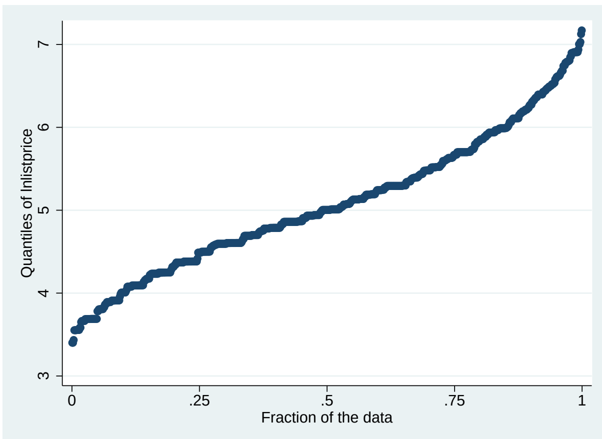

Figure 3: Quantile plots

(a) Posted price

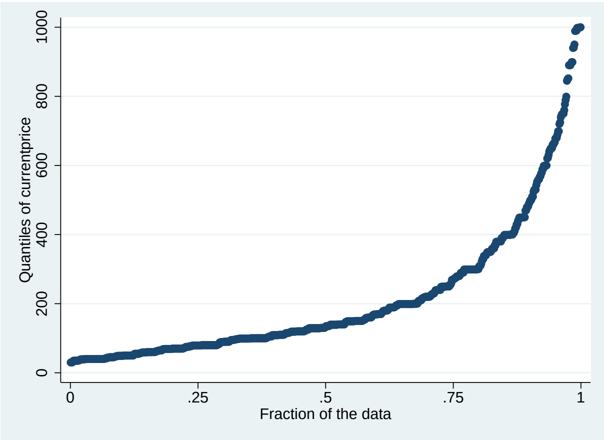

(b) ln(posted price)

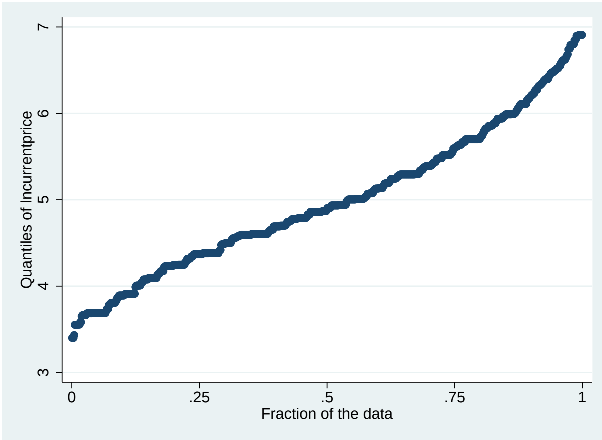

Figure 4: Predicted probability of a discount: Interactions not shown in LNH (full) specification in Table 4

(a)

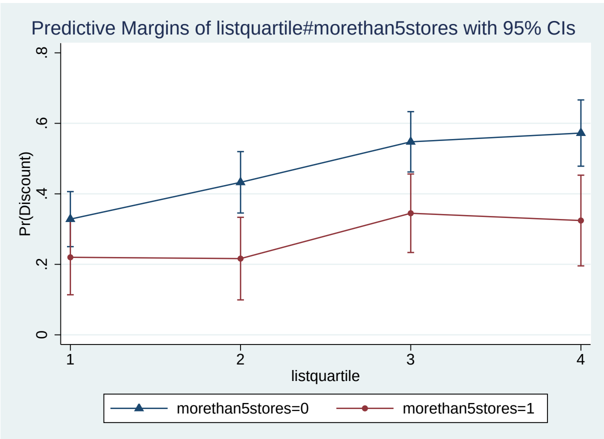

(b)

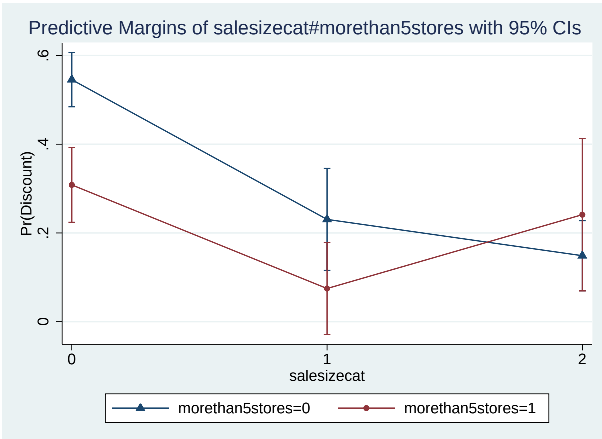

(c)

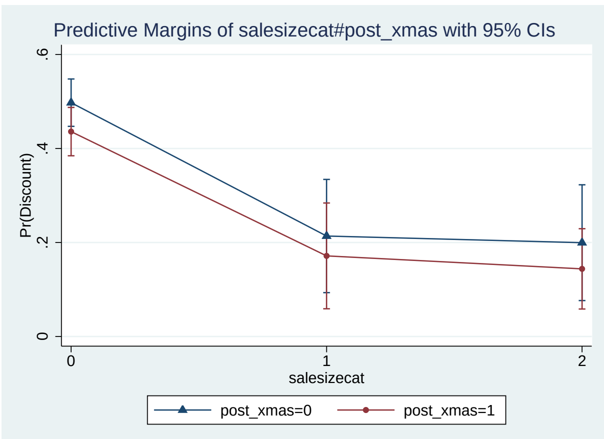

Figure 5: Predicted probability of a discount: Interactions of alternate specifications of price with firm scale

(a)

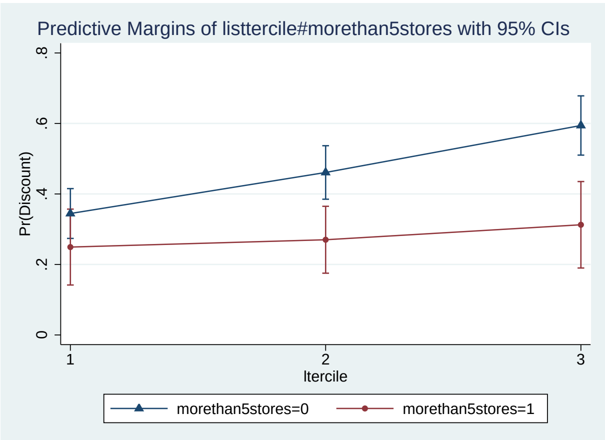

(b)

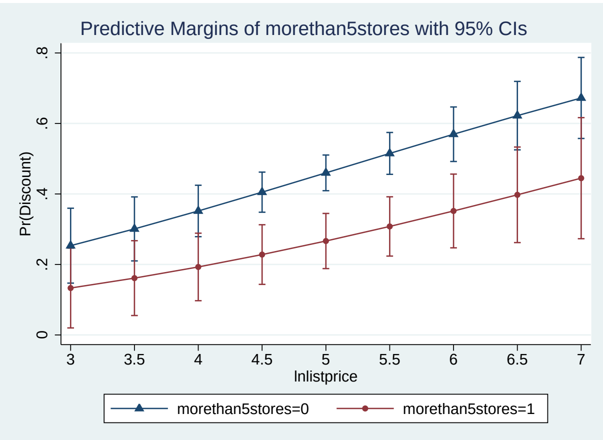

(c)

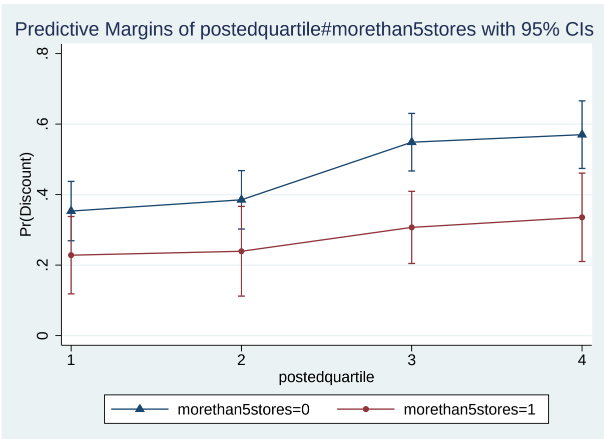

Figure 6: Predicted probability of a discount: Interactions of price with alternate specifications of firm scale

(a)

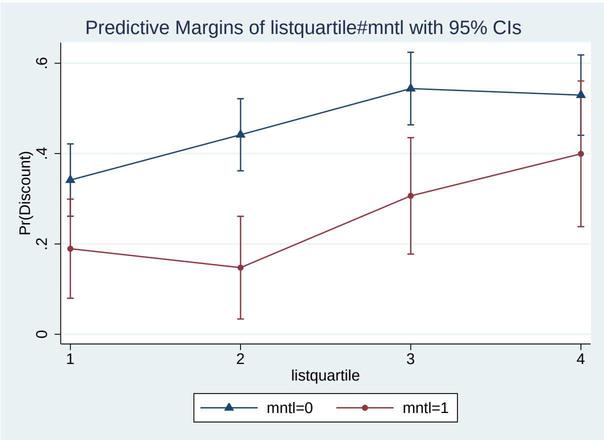

(b)

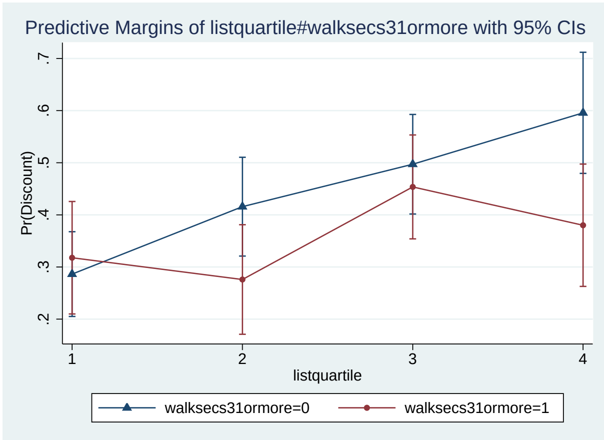

(c)

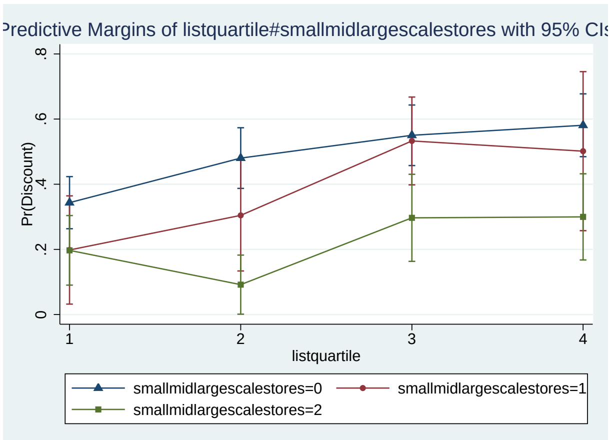

Figure 7: Predicted probability of a discount: Interactions of sale size category with alternate specifications of firm scale

(a)

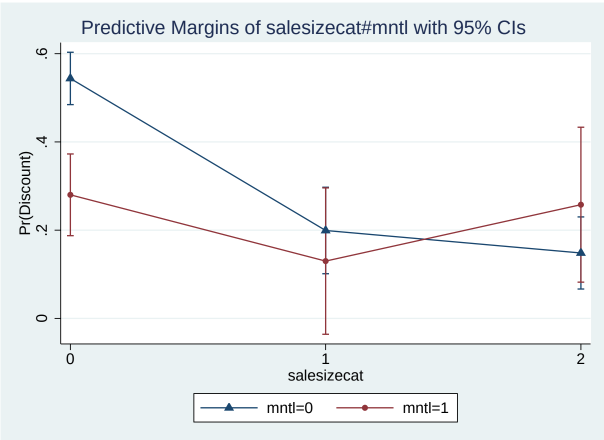

(b)

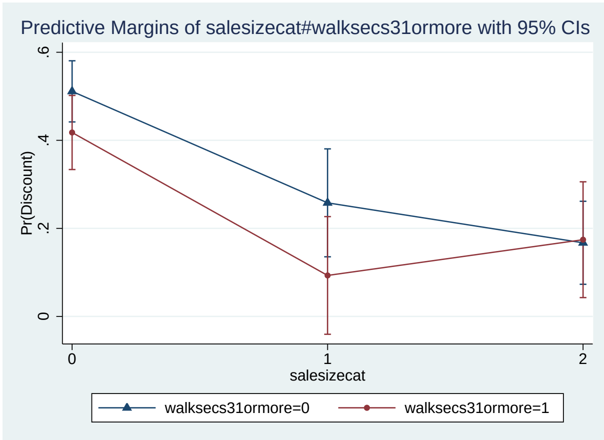

(c)

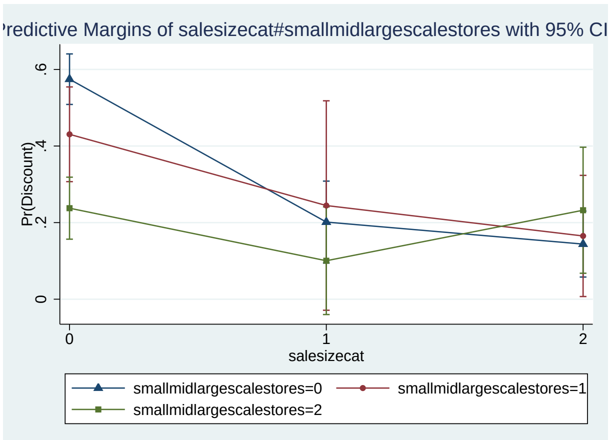

Figure 8: Predicted probability of a discount: Interactions when jewelry observations are excluded from the analysis

(a)

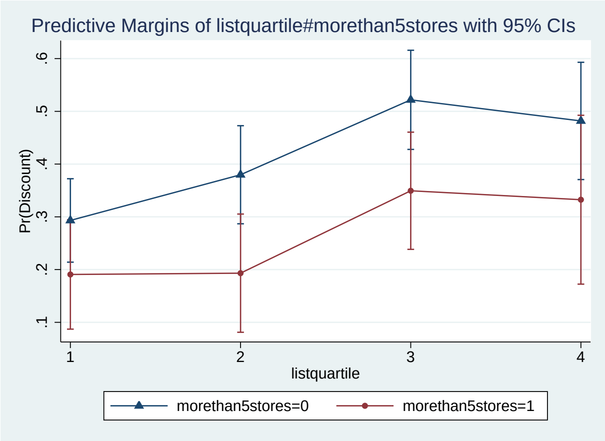

(b)

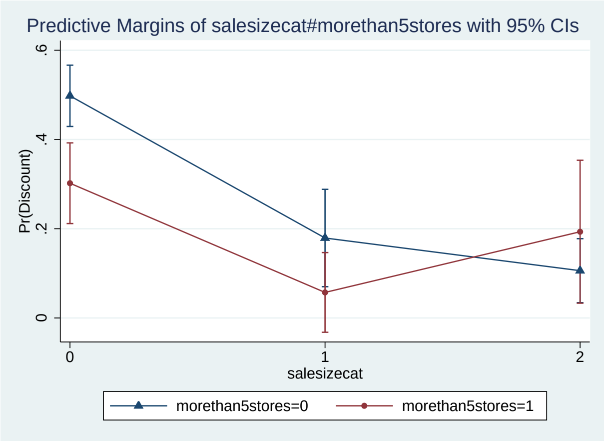

(c)

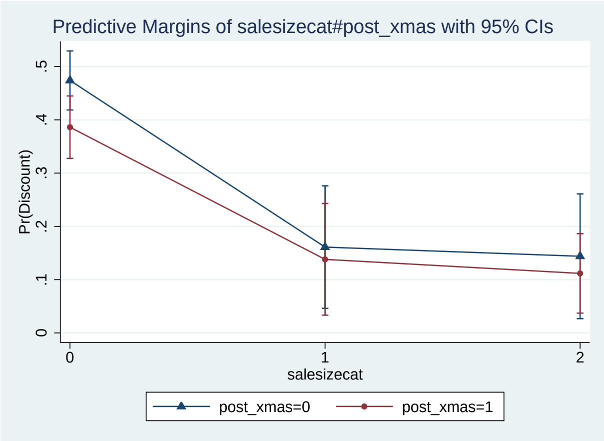

Figure 9: Average percentage of RAs who obtained a discount by order in which store was visited

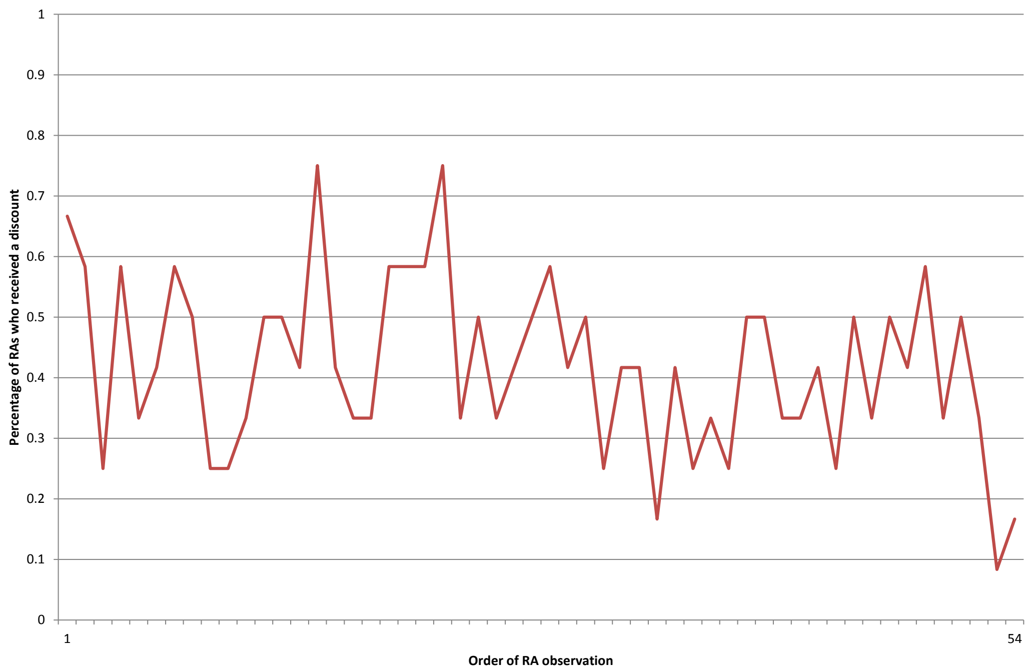

Table 1: Summary price, sale, and discount statistics

Variable	Mean	St. Dev.	Median	Mode	Min	Max	Obs
List price	237.133	302.48	149	199	30	5900	751
Posted price (after any sale)	212.405	204.52	135	199	30	999.99	751
Positive sale amount (off of list price)	109.884	411.653	45	20	5	5240	169
Positive sale percentage (off of list price)	0.306	0.147	0.3	0.3	0.034	0.888	169
Positive discount amount (off of posted price)	25.464	29.286	15	5	1	180	303
Positive discount percentage (off of posted price)	0.098	0.066	0.098	0.1	0.008	0.407	303

Table 2: Summary statistics of categorical variables

Variable	Unique stores in the dataset	Observations
Number of stores in the world, per firm		
1	130	323
2	30	85
3	13	36
4	10	30
5	5	13
6	5	15
7	5	13
8	2	6
10-15	6	17
>15	74	213
Scale		
Domestic	204	532
Multinational	76	219
Geographic area		
1st district	55	157
2nd and 20th districts	57	143
18th and 19th districts	54	137
6th and 7th districts	114	314
Employees		
1	117	293
2	67	187
3	30	87
4	22	61
5	9	24
> 5	35	99
Number of customers observed		
0	138	349
1	42	120
2	26	72
3	20	56
4	4	12
5	6	17
6-10	26	73
> 10	18	52
Store category		
Clothing	96	270
Shoes	33	94
Jewelry	36	96
Household	26	66
Other	89	225
Walking seconds		
Up to 30 seconds	158	411
31-60 seconds	47	130
61-120 seconds	28	76
More than 2 minutes	47	134

Table 3: Discounting behavior of multinational firms

Granted a discount at least once	Incidence	Never granted a discount	Obs
Bucherer	2 of 2 obs	Aldo (franchise store)	3
Cadenza	1 of 3 obs	Bonita	3
Colli	3 of 3 obs	Boss	3
Dorotheum	2 of 3 obs	Butlers	2
Douglas	2 of 3 obs	Camper	3
EMI Music	1 of 2 obs	Casa (franchise store)	3
Fogal	1 of 3 obs	Comma	3
Freytag&Berndt	2 of 3 obs	Cos	3
Högl	2 of 3 obs	Desigual	3
J&L Lobmeyr	2 of 2 obs	Diesel	3
Jacques Lemans (franchise store)	3 of 3 obs	Energie	3
Le Clou	1 of 3 obs	Escada	2
Levis (franchise store)	1 of 3 obs	Esprit	3
Marionnaud	1 of 3 obs	Footlocker	3
Matratzen Concord	3 of 3 obs	Fossil	3
Pandora (franchise store)	1 of 3 obs	Gabor	3
Samsonite	2 of 3 obs	Georges Rech	3
Saturn	1 of 3 obs	Geox	3
Sidestep	2 of 3 obs	Gerry Weber	3
Sport 2000 (franchise store)	3 of 3 obs	Gloriette	3
Stiefelkönig	1 of 3 obs	Grüne Erde	3
Tchibo	1 of 3 obs	G-Star (franchise store)	3
Vans	1 of 3 obs	Gucci	3
Wmf	2 of 3 obs	H&M	3
		Hallhuber	3
		Humanic	3
		Jack Wolfskin	2
		Jack Jones	2
		Joseph Ribkoff (franchise store)	3
		Kare (franchise store)	3
		Lacoste	3
		Mango	3
		Meblik Kids	2
		Moncler	3
		Mont blanc	3
		Müller	3
		Nespresso	3
		Peak Performance	3
		Pearle	2
		Puma	3
		S. Oliver	3
		Salamander	3
		Sisley (franchise store)	3
		Sports Experts	3
		Stefanel	3
		Swarovski	3
		Swatch	3
		Tom Tailor	3
		Triumph	3
		Weltbild	3
		Wolford	3
		Zara	3

Note: The designation as a franchise store was determined based on a telephone conversation with at least one salesperson at each of the above stores. As noted in Table 2, observations at multinational stores account for 219 of 751 observations overall.

Table 4: Average partial effects in the participation and amount equations: Full and Parsimonious Specifications

	LNH (Full)		ET2T (Full)		LNH (Pars)		ET2T (Pars)	
	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq
List price: 2nd quartile	0.071*	0.525***	0.073*	0.529***	0.071*	0.493***	0.072*	0.491***
	(0.039)	(0.101)	(0.038)	(0.102)	(0.039)	(0.101)	(0.039)	(0.1)
List price: 3rd quartile	0.182***	1.139***	0.182***	1.155***	0.182***	1.102***	0.181***	1.113***
	(0.044)	(0.117)	(0.043)	(0.121)	(0.044)	(0.115)	(0.044)	(0.118)
List price: 4th quartile	0.192***	1.86***	0.191***	1.87***	0.192***	1.831***	0.19***	1.837***
	(0.05)	(0.123)	(0.05)	(0.126)	(0.05)	(0.122)	(0.051)	(0.124)
Small sale item	-0.274***	-0.446**	-0.281***	-0.502**	-0.274***	-0.471**	-0.277***	-0.496**
	(0.05)	(0.227)	(0.047)	(0.222)	(0.05)	(0.234)	(0.049)	(0.233)
Large sale item	-0.295***	0.15	-0.296***	0.145	-0.295***	0.191	-0.295***	0.201
	(0.046)	(0.201)	(0.044)	(0.175)	(0.046)	(0.201)	(0.045)	(0.202)
More than five stores	-0.185***	-0.033	-0.18***	0.005	-0.185***		-0.181***	
	(0.052)	(0.128)	(0.051)	(0.134)	(0.052)		(0.053)	
Dec 27 - Jan 4	-0.059**	0.098	-0.057**	0.099	-0.059**		-0.051*	
	(0.027)	(0.072)	(0.026)	(0.071)	(0.027)		(0.031)	
Female salesperson	-0.066*	-0.132	-0.068*	-0.136	-0.066*		-0.073*	
	(0.039)	(0.09)	(0.039)	(0.091)	(0.039)		(0.04)	
One customer	-0.057	-0.27**	-0.064	-0.305**	-0.057	-0.227*	-0.061	-0.219*
	(0.054)	(0.126)	(0.051)	(0.125)	(0.054)	(0.12)	(0.053)	(0.12)
Two customers	-0.086	-0.276*	-0.087	-0.301*	-0.086	-0.28*	-0.088	-0.283*
	(0.061)	(0.15)	(0.061)	(0.151)	(0.061)	(0.149)	(0.061)	(0.15)
Three or more customers	-0.067	-0.361**	-0.066	-0.37**	-0.067	-0.328**	-0.072	-0.279*
	(0.071)	(0.162)	(0.07)	(0.156)	(0.071)	(0.133)	(0.072)	(0.166)
Two employees	-0.027	0.003	-0.022	0.003	-0.027		-0.026	
	(0.051)	(0.101)	(0.05)	(0.099)	(0.051)		(0.051)	
Three employees	-0.108	-0.164	-0.107	-0.193	-0.108		-0.115*	
	(0.068)	(0.155)	(0.065)	(0.155)	(0.068)		(0.067)	
Four or more employees	-0.178**	0.124	-0.172**	0.109	-0.178**		-0.168**	
	(0.076)	(0.185)	(0.074)	(0.185)	(0.076)		(0.078)	
2nd and 20th districts	0.099	0.091	0.103	0.135	0.099		0.107	
	(0.067)	(0.156)	(0.066)	(0.158)	(0.067)		(0.07)	
18th and 19th districts	0.101	-0.035	0.106*	0.005	0.101		0.102	
	(0.066)	(0.15)	(0.064)	(0.149)	(0.066)		(0.064)	
6th and 7th districts	0.022	-0.068	0.024	-0.037	0.022		0.022	
	(0.055)	(0.129)	(0.053)	(0.13)	(0.055)		(0.054)	
Salesperson age: 35-50	0.005	-0.069	0.005	-0.077	0.005		0.003	
	(0.038)	(0.095)	(0.037)	(0.092)	(0.038)		(0.038)	
Salesperson age: > 50	-0.027	-0.012	-0.021	-0.019	-0.027		-0.025	
	(0.05)	(0.132)	(0.049)	(0.131)	(0.05)		(0.049)	

Continued on next page

Table 4 : Continued from previous page

	LNH (Full)		ET2T (Full)		LNH (Pars)		ET2T (Pars)	
	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq
Expensive clothing	-0.126**	0.072	-0.126**	0.055	-0.126**	0.105	-0.124**	0.097
	(0.062)	(0.162)	(0.06)	(0.159)	(0.062)	(0.158)	(0.061)	(0.156)
Inexpensive shoes	-0.118	-0.004	-0.11	0.008	-0.118	-0.007	-0.117	-0.005
	(0.09)	(0.267)	(0.089)	(0.26)	(0.09)	(0.266)	(0.09)	(0.267)
Expensive shoes	-0.08	-0.21	-0.063	-0.178	-0.08	-0.156	-0.075	-0.15
	(0.087)	(0.281)	(0.091)	(0.285)	(0.087)	(0.273)	(0.088)	(0.275)
Inexpensive jewelry	-0.025	0.31*	-0.024	0.303*	-0.025	0.363**	-0.024	0.358**
	(0.071)	(0.181)	(0.07)	(0.178)	(0.071)	(0.171)	(0.071)	(0.171)
Expensive jewelry	0.283***	0.188	0.268***	0.257	0.283***	0.218	0.282***	0.265
	(0.093)	(0.187)	(0.094)	(0.191)	(0.093)	(0.171)	(0.093)	(0.173)
Inexpensive household	-0.007	0.286	-0.01	0.247	-0.007	0.317	-0.008	0.305
	(0.112)	(0.217)	(0.11)	(0.198)	(0.112)	(0.21)	(0.112)	(0.208)
Expensive household	-0.09	0.379	-0.075	0.366	-0.09	0.452*	-0.085	0.441*
	(0.094)	(0.245)	(0.093)	(0.244)	(0.094)	(0.249)	(0.095)	(0.254)
Inexpensive “other”	-0.105	-0.105	-0.105*	-0.115	-0.105	-0.013	-0.107*	-0.014
	(0.064)	(0.165)	(0.063)	(0.168)	(0.064)	(0.146)	(0.064)	(0.153)
Expensive “other”	-0.014	-0.156	-0.016	-0.149	-0.014	-0.04	-0.016	-0.039
	(0.063)	(0.181)	(0.062)	(0.181)	(0.063)	(0.163)	(0.062)	(0.164)
Observations	751	303	751	303	751	303	751	303
Log likelihood	-1473.399		-1472.042		-1478.83		-1478.29	
ρ	Zero-constrained		-0.664**		Zero-constrained		-0.283	
			(0.207)				(0.393)	

Note: Default categories are 1st quartile, non-sale items, firms with up to five stores, Dec 9 - Dec 14, male salesperson, no customers observed, one employee observed, 1st district, salesperson age < 35, and inexpensive clothing. Interactions and RA fixed effects are not shown for space reasons. Interaction effects are reported in Figure 4.

Table 5: Average partial effects in the participation and amount equations (LNH model): Alternative price variables

	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq
List price: 2nd tercile	0.085**	0.684***				
	(0.035)	(0.095)				
List price: 3rd tercile	0.183***	1.603***				
	(0.045)	(0.104)				
ln (list price)			0.093***	0.91***		
			(0.022)	(0.051)		
Posted price: 2nd quartile					0.025	0.491***
					(0.04)	(0.112)
Posted price: 3rd quartile					0.152***	1.05***
					(0.045)	(0.116)
Posted price: 4th quartile					0.174***	1.821***
					(0.053)	(0.128)
Small sale item	-0.267***	-0.469**	-0.273***	-0.532**	-0.258***	-0.288
	(0.049)	(0.245)	(0.049)	(0.222)	(0.05)	(0.238)
Large sale item	-0.285***	0.021	-0.304***	-0.028	-0.256***	0.348*
	(0.047)	(0.205)	(0.045)	(0.169)	(0.049)	(0.179)
More than five stores	-0.181***	0.056	-0.184***	-0.079	-0.182***	-0.019
	(0.052)	(0.127)	(0.051)	(0.125)	(0.052)	(0.128)
Dec 27 - Jan 4	-0.063**	0.026	-0.058**	0.085	-0.057**	0.081
	(0.027)	(0.075)	(0.027)	(0.066)	(0.027)	(0.073)
Female salesperson	-0.066*	-0.094	-0.062	-0.108	-0.064	-0.106
	(0.039)	(0.09)	(0.039)	(0.087)	(0.039)	(0.091)
One customer	-0.049	-0.158	-0.054	-0.217*	-0.059	-0.278**
	(0.054)	(0.13)	(0.053)	(0.117)	(0.053)	(0.124)
Two customers	-0.092	-0.322**	-0.054	-0.217**	-0.091	-0.304*
	(0.061)	(0.151)	(0.053)	(0.117)	(0.061)	(0.155)
Three or more customers	-0.066	-0.362**	-0.062	-0.362**	-0.071	-0.428***
	(0.07)	(0.165)	(0.071)	(0.161)	(0.07)	(0.158)
Two employees	-0.026	-0.011	-0.024	0.008	-0.027	0.042
	(0.052)	(0.101)	(0.052)	(0.098)	(0.052)	(0.105)
Three employees	-0.109	-0.136	-0.106	-0.17	-0.111	-0.171
	(0.069)	(0.155)	(0.068)	(0.142)	(0.068)	(0.158)
Four or more employees	-0.176	0.184	-0.18**	0.101	-0.178**	0.132
	(0.075)	(0.187)	(0.076)	(0.181)	(0.075)	(0.177)
Observations	751	303	751	303	751	303
Log likelihood	-1486.753		-1460.289		-1478.4	

Note: Standard errors in parentheses. Default categories are 1st tercile (list), 1st quartile (posted), non-sale items, firms with up to five stores, Dec 9 - Dec 14, male salesperson, no customers observed, and one employee observed. Interactions, area, age, product type, and RA effects are not shown.

Table 6: Average partial effects in the participation and amount equations (LNH model): Alternative firm-scale variables

	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq
List price: 2nd quartile	0.063 (0.039)	0.525*** (0.101)	0.062 (0.039)	0.517*** (0.1)	0.075** (0.037)	0.528*** (0.101)
List price: 3rd quartile	0.174*** (0.045)	1.135*** (0.117)	0.18*** (0.044)	1.134*** (0.117)	0.188*** (0.044)	1.147*** (0.116)
List price: 4th quartile	0.188*** (0.053)	1.858*** (0.123)	0.202*** (0.052)	1.847*** (0.122)	0.203*** (0.05)	1.863*** (0.122)
Small sale item	-0.287*** (0.048)	-0.451** (0.225)	-0.266*** (0.049)	-0.456** (0.231)	-0.282*** (0.048)	-0.447** (0.225)
Large sale item	-0.296*** (0.046)	0.132 (0.191)	-0.299*** (0.046)	0.139 (0.201)	-0.302*** (0.043)	0.152 (0.201)
Multinational firm	-0.196*** (0.055)	0.053 (0.154)				
Walking seconds > 30			-0.083 (0.055)	-0.105 (0.109)		
Between 4-15 stores					-0.105* (0.062)	-0.083 (0.124)
More than 15 stores					-0.26*** (0.05)	-0.02 (0.16)
Dec 27 - Jan 4	-0.065** (0.027)	0.101 (0.071)	-0.054** (0.027)	0.094 (0.07)	-0.058** (0.027)	0.096 (0.072)
Female salesperson	-0.098** (0.038)	-0.133 (0.091)	-0.09** (0.039)	-0.143 (0.091)	-0.055 (0.038)	-0.13 (0.09)
One customer	-0.101** (0.05)	-0.269** (0.129)	-0.078 (0.055)	-0.254** (0.129)	-0.056 (0.051)	-0.265** (0.127)
Two customers	-0.091 (0.05)	-0.271* (0.129)	-0.055 (0.055)	-0.236 (0.129)	-0.089 (0.057)	-0.281* (0.15)
Three or more customers	-0.038 (0.057)	-0.367*** (0.149)	-0.08 (0.068)	-0.339** (0.158)	-0.052 (0.066)	-0.371** (0.166)
Two employees	-0.037 (0.05)	-0.002 (0.1)	-0.024 (0.052)	0.01 (0.099)	-0.038 (0.048)	0.001 (0.1)
Three employees	-0.106 (0.067)	-0.17 (0.159)	-0.106 (0.074)	-0.132 (0.16)	-0.104 (0.069)	-0.166 (0.156)
Four or more employees	-0.199*** (0.072)	0.109 (0.181)	-0.193** (0.079)	0.166 (0.186)	-0.167** (0.073)	0.135 (0.182)
Observations	751	303	751	303	751	303
Log likelihood	-1471.108		-1481.014		-1463.161	

Note: Standard errors are in parentheses. Default categories are 1st quartile, non-sale items, walking seconds ≤ 30, firms with ≤ 3 stores, Dec 9 - Dec 14, male salesperson, no customers observed, and one employee observed. Interactions, area, age, product type, and RA effects are not shown.

Table 7: Average partial effects in the participation and amount equations (LNH model): Alternative product category specifications

	Part Eq	Amt Eq	Part Eq	Amt Eq	Part Eq	Amt Eq
List price: 2nd quartile	0.075*	0.554***	0.059	0.578***	0.059	0.581***
	(0.04)	(0.1)	(0.04)	(0.106)	(0.04)	(0.106)
List price: 3rd quartile	0.179***	1.154***	0.197***	1.127***	0.183***	1.132***
	(0.044)	(0.106)	(0.045)	(0.122)	(0.045)	(0.11)
List price: 4th quartile	0.252***	1.978***	0.167***	1.752***	0.172***	1.846***
	(0.047)	(0.11)	(0.056)	(0.143)	(0.051)	(0.133)
Small sale item	-0.277***	-0.512**	-0.281***	-0.469*	-0.266***	-0.549**
	(0.05)	(0.214)	(0.048)	(0.251)	(0.049)	(0.229)
Large sale item	-0.329***	0.149	-0.303***	0.169	-0.309***	0.174
	(0.043)	(0.231)	(0.044)	(0.225)	(0.042)	(0.25)
More than five stores	-0.208***	-0.032	-0.147***	0.067	-0.157***	0.115
	(0.051)	(0.137)	(0.056)	(0.142)	(0.055)	(0.147)
Dec 27 - Jan 4	-0.059**	0.094	-0.073**	0.129	-0.072**	0.12
	(0.028)	(0.075)	(0.03)	(0.088)	(0.03)	(0.09)
Female salesperson	-0.057	-0.067	-0.053	-0.093	-0.05	-0.051
	(0.039)	(0.087)	(0.042)	(0.108)	(0.041)	(0.101)
One customer	-0.047	-0.234*	-0.1*	-0.347**	-0.101*	-0.387**
	(0.055)	(0.132)	(0.06)	(0.166)	(0.058)	(0.169)
Two customers	-0.087	-0.317**	-0.11*	-0.174	-0.089	-0.235*
	(0.065)	(0.134)	(0.064)	(0.15)	(0.068)	(0.14)
Three or more customers	-0.069	-0.416***	-0.089	-0.417**	-0.065	-0.449**
	(0.07)	(0.159)	(0.075)	(0.171)	(0.075)	(0.164)
Two employees	-0.013	-0.017	-0.015	0.061	-0.007	0.068
	(0.053)	(0.102)	(0.058)	(0.111)	(0.059)	(0.122)
Three employees	-0.093	-0.171	-0.131*	-0.256	-0.134*	-0.313
	(0.069)	(0.174)	(0.079)	(0.224)	(0.077)	(0.225)
Four or more employees	-0.169**	0.082	-0.163**	0.165	-0.152*	0.157
	(0.077)	(0.178)	(0.082)	(0.197)	(0.084)	(0.192)
Observations	751	303	655	235	655	235
Log likelihood	-1499.656		-1137.992		-1148.181	

Note: Standard errors are in parentheses. Default categories are 1st quartile, non-sale items, firms with $\leq$ 5 stores, Dec 9 - Dec 14, male salesperson, no customers observed, and one employee observed. Interactions, area, age, product type (pertains only to (9)), and RA effects not shown.

Table 8: Average price percentile and percentage of sale items chosen by RA identity

RA ID	Pctile of assigned range (avg)	St. Dev.	Sale items observed (pct)	St. Dev.
1	0.365	0.271	0.149	0.359
2	0.425	0.342	0.133	0.343
3	0.513	0.341	0.279	0.452
4	0.372	0.355	0.323	0.471
5	0.525	0.304	0.182	0.389
6	0.376	0.308	0.305	0.464
7	0.515	0.32	0.233	0.427
8	0.425	0.339	0.185	0.392
9	0.512	0.312	0.231	0.425
10	0.509	0.305	0.292	0.458
11	0.441	0.312	0.172	0.38
12	0.469	0.343	0.213	0.413

Note: RAs #1-#6 are females and RAs #7-#12 are males.

Table 9: Discounts granted to RAs that followed a particular RA at the same store (three observations per store required)

	First follower		Second follower		Both followers	
Leader ID	Discounts	Observations	Discounts	Observations	Discounts	Observations
1	4	10	4	10	8	20
2	8	18	6	18	14	36
3	9	16	5	16	14	32
4	12	23	6	23	18	46
5	5	19	6	19	11	38
6	7	15	5	15	12	30
7	6	19	5	19	11	38
8	7	16	6	16	13	32
9	13	25	9	25	22	50
10	8	19	5	19	13	38
11	7	16	6	16	13	32
12	5	15	3	15	8	30

Note: RAs #1-#6 are females and RAs #7-#12 are males.

Table 10: Discounts granted to RAs that followed a particular RA at the same store within the same week

	Leader & follower visit pre-xmas		Leader & follower visit post-xmas	
Leader ID	Follower Discounts	Follower Obs.	Follower Discounts	Follower Obs.
1	2	6	7	20
2	5	8	7	10
3	7	10	1	11
4	5	10	1	9
5	4	11	3	9
6	4	8	4	10
7	7	13	1	8
8	5	10	2	5
9	9	16	2	7
10	2	10	0	7
11	4	10	2	5
12	3	9	1	7

Note: RAs #1-#6 are females and RAs #7-#12 are males.

References

Armstrong, Mark, and John Vickers. 2001. "Competitive Price Discrimination." *RAND Journal of Economics*, 32(4): 579–605.

Arnold, Michael A, and Steven A Lippman. 1998. "Posted Prices Versus Bargaining in Markets with Asymmetric Information." *Economic Inquiry*, 36(3): 450–57.

Ayres, Ian, and Peter Siegelman. 1995. "Race and Gender Discrimination in Bargaining for a New Car." *American Economic Review*, 85(3): 304–21.

Bester, Helmut. 1994. "Random Advertising and Monopolistic Price Dispersion." *Journal of Economics & Management Strategy*, 3(3): 545–59.

Castillo, Marco, Ragan Petrie, Maximo Torero, and Lise Vesterlund. 2013. "Gender differences in bargaining outcomes: A field experiment on discrimination." *Journal of Public Economics*, 99: 35–48.

Corts, Kenneth S. 1998. "Third-degree price discrimination in oligopoly: all-out competition and strategic commitment." *The RAND Journal of Economics*, 306–323.

Cragg, John G. 1971. "Some statistical models for limited dependent variables with application to the demand for durable goods." *Econometrica: Journal of the Econometric Society*, 829–844.

Fudenberg, Drew, and Jean Tirole. 1983. "Sequential bargaining with incomplete information." *The Review of Economic Studies*, 50(2): 221–247.

Fudenberg, Drew, and J Miguel Villas-Boas. 2006. "Behavior-based price discrimination and customer recognition." *Handbook on economics and information systems*, 1: 377–436.

Fudenberg, Drew, David K Levine, and Jean Tirole. 1987. "Incomplete information bargaining with outside opportunities." *The Quarterly Journal of Economics*, 37–50.

Goldberg, Pinelopi Koujianou. 1996. "Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey." *Journal of Political Economy*, 622–654.

Hart, Oliver D, and Jean Tirole. 1988. "Contract renegotiation and Coasian dynamics." *The Review of Economic Studies*, 55(4): 509–540.

Holmes, Thomas J. 1989. "The effects of third-degree price discrimination in oligopoly." *The American Economic Review*, 244–250.

List, John A. 2004. "Neoclassical theory versus prospect theory: Evidence from the marketplace." *Econometrica*, 72(2): 615–625.

Mahoney, Tom, and Leonard Sloane. 1966. *The great merchants: America's foremost retail institutions and the people who made them great.* Harper & Row.

Myerson, Roger B, and Mark A Satterthwaite. 1983. "Efficient mechanisms for bilateral trading." *Journal of economic theory*, 29(2): 265–281.

Nash Jr, John F. 1950. "The bargaining problem." *Econometrica: Journal of the Econometric Society*, 155–162.

Perry, Motty. 1986. "An example of price formation in bilateral situations: A bargaining model with incomplete information." *Econometrica: Journal of the Econometric Society*, 313–321.

Riley, John, and Richard Zeckhauser. 1983. "Optimal selling strategies: When to haggle, when to hold firm." *The Quarterly Journal of Economics*, 267–289.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica*, 50(1): 97–109.

Sobel, Joel, and Ichiro Takahashi. 1983. "A multistage model of bargaining." *The Review of Economic Studies*, 50(3): 411–426.

Vuong, Quang H. 1989. "Likelihood ratio tests for model selection and non-nested hypotheses." *Econometrica: Journal of the Econometric Society*, 307–333.

Wang, Ruqu. 1995. "Bargaining Versus Posted-Price Selling." *European Economic Review*, 39(9): 1747–1764.

Wooldridge, Jeffrey M. 2010. *Econometric analysis of cross section and panel data.* MIT press.

Appendix

In order to understand sales we shall look at a two-period version of our model where the monopolist determines its price p_0 in period zero and its price p_1 in period one. We will abstract from issues of inter-temporal discrimination (see discussion at the end of this section), and will assume that p_0 does not affect demand in period one. Thus in every period the monopolist chooses the profit-maximizing price for that period. We then refer to a sale as a situation in which $p_0 > p_1$. Conversely, when $p_0 \leq p_1$, the associated item will not be considered to be on sale, and thus we will consider such observations as non-sale observations.

More formally, consider a two-period version of our model with β small. Let A be a parameter that defines a family of distribution functions F_A that are ordered by A in such a way that p^M is increasing in A, thus higher A implies higher price (this can be due to shift and/or rotation of demand). To further simplify analysis, we shall assume that holding c constant, an increase in A, increases $\omega(p^M)$, thus A increases the average valuation between marginal cost and the monopoly price. For example, if demand is given by $D(p) = (A - p)^\eta$, an increase in A increases p^M but also ω.

Let there be two time periods, $t = 0, 1$, and let A_t and c_t denote realizations of the parameters for time period t. Assume c_0 and A_0 are jointly distributed according to some distribution function. Conditional on (c_0, A_0), next period's realizations are then jointly distributed according to another distribution function. This specification allows for arbitrary time dependence between the two parameters of the model.

Now we consider inference after observing p_0 and p_1 about A_1 and c_1, which is necessary for assessment of propensity to bargain today, denoted by ω_1. Formally, we are interested in deriving the posterior distribution of ω_1 conditional on p_0 and p_1 and we are interested how this distribution changes with p_0 holding p_1 fixed.

Given our assumptions about the relationship between F and A, for a given price p_t, there is a downward-sloping schedule of A_t and c_t that are compatible with such a price. Let $\tilde{A}(c, p)$ be the demand coefficient such that for marginal cost c the monopoly price is p. It is trivial to show that $\tilde{A}_c(c, p) < 0$ and $\tilde{A}_p(c, p) > 0$. Thus in the (A, c) space, $\tilde{A}(c, p)$ is downward-sloping and shifts up with p. Also, holding a price p fixed, any inference that makes it more likely that c_1 is high and thus A_1 is low, also makes is more likely that ω_1 is low.

One observes a sale when $p_0 > p_1$ and no sale when $p_0 \leq p_1$. The main question now is whether increasing p_0 over p_1 changes our inference about ω in such a way that higher p_0 implies lower expected propensity. Recall that for $\beta \to 0$ the propensity is given by

$$\omega = \lambda \int_c^{p^M} (v - c) \frac{f_A(v)}{F_A(p^M) - F_A(c)} dv - b. \tag{18}$$

Thus for an increase in p_0 over p_1 to cause reduction in propensity to bargain, it has to be the case that an increase in p_0 increases our inference about c and/or reduces our inference about A. Because we are primarily interested in $p_0 \geq p_1$, we shall consider this case without loss of generality.

It should be immediately clear that if shocks are uncorrelated over time, then p_0 cannot affect the inference about ω_1. However, shocks may well be correlated over time and the this inference may go either way depending on the distribution of parameters and the demand function. To illustrate this point, we provide two examples where in one the fact that a good is on a sale implies that the propensity to bargain is lower than when identically priced good is not on sale (holding p_1 constant, increasing p_0 above p_1 results in a reduction of the propensity), and another where exactly the opposite is true.

Consider two possibilities. First, c is random but the same across periods, and A is drawn independently in both periods.[22] We shall call this demand shocks example. Second, A is random but the same across periods and c is random and drawn independently in each period. This is the cost shocks example.

Intuitively, in the demand shocks example observing a relatively high p_0 makes it more likely that it originates from a relatively high c, thus today $A_1 - c$ has to be low. Then we have to conclude that observing a sale reduces inferred margin. The opposite is true in the cost shocks example because there high p_0 makes it more likely that it resulted from a high A, and thus $A - c_1$ is likely to be high.

Take the demand shocks example first. Let c be the same over both periods and distributed according to some $H(c)$ on $[\underline{c}, \bar{c}]$ and A_t is independent over time and distributed according to some $G(A)$ on $[\underline{A}, \bar{A}]$ (lowercase letters will denote densities).

In order to derive the conditional density of ω_1, one needs to find the conditional distribution over (A_1, c). Because for a fixed p_1 and c there is only one A_1 that is consistent with p_1, it is sufficient to derive the conditional distribution of c given p_0 and p_1.

Note first that some realizations of c may be inconsistent with p_0 and p_1 for any A_0 and A_1 becuase $A_t \in [\underline{A}, \bar{A}]$. Given that $p_0 \geq p_1$, c may not be lower than c_1, defined as the marginal cost at which price is p_0 for the highest possible realization of A_0, i.e. c_1 is the solution to $\tilde{A}(p_0, c_1) = \bar{A}$. Similarly, c may not exceed c_2, defined as the solution to $\tilde{A}(p_1, c_2) = \underline{A}$. Thus c has to belong to the interval $[\max\{\underline{c}, c_1\}, \min\{\bar{c}, c_2\}]$. Given a $c \in [\max\{\underline{c}, c_1\}, \min\{\bar{c}, c_2\}]$, using the Bayes rule, the conditional density of c conditional on p_0 and p_1 is

$$h(c|p_0, p_1) = \frac{g(\tilde{A}(c, p_0))g(\tilde{A}(c, p_1))h(c)}{\int_{\max\{\underline{c}, c_1\}}^{\min\{\bar{c}, c_2\}} g(\tilde{A}(x, p_0))g(\tilde{A}(x, p_1))h(x)dx}. \tag{19}$$

The formula above indicates two channels through which p_0 will matter for ω_1. First, for sufficiently high p_0, increasing p_0 will increase the lower bound of the interval of the posterior distribution, $\max\{\underline{c}, c_1\}$. This is because when p_0 is high, $c_1 > \underline{c}$, and so $\max\{\underline{c}, c_1\} = c_1$ and c_1 is increasing in p_0. Intuitively, if p_0 is high, it could not have come from very low marginal cost because, even the highest demand shock would not generate such a high price, thus high p_0 excludes very low cost realizations as a possibility. Second, the inference changes with p_0 due to the relative probabilities of various cost and demand parameters, as summarized by the Bayes rule. The second effect will in general depend on the shape of H and G and is a priori impossible to sign. It is possible to show,

[22] Introducing serial correlation can only obscure the simple point we want to make here

however, that if both distributions are uniform, then the second channel will not exist and therefore increasing p_0 will either not change inference about ω_1 (when $\underline{c} > c_1$) or will shift the distribution to the left.

To illustrate the above, consider a special case for the linear example above where $F(v) = \frac{v}{A}$ and uniform shocks such that $H = U(0,2)$ and $G = U(2,4)$. We know that $p_t = \frac{A_t+c}{2}$ and $\omega_1 = \frac{\lambda(A_1-c)}{4}$. The first can be rewritten as $\tilde{A}(c,p) = 2p - c$, so $\omega_1 = \frac{\lambda(p_1-c)}{2}$ and so holding p_1 fixed, and increase in c leads to lower ω. Because the distributions are uniform, densities will play no role, i.e $h(c|p_0,p_1) = \frac{1}{2}$ for all permissible c. Thus the main issue is conditional on prices, which cost levels are possible. $p_t = \frac{A_t+c}{2}$, and so $2p_t - 4 \leq c \leq 2p_t - 2$. Given that $p_0 \geq p_1$, we have $2p_0 - 4 \leq c \leq 2p_1 - 2$, which immediately implies that the biggest price difference consistent with equal marginal cost is 1. For example, if $p_1 = 2$ and $p_0 = 2.5$, then the marginal cost is constrained to be in the interval $[1,2]$.

There are two cases. If $2p_0 - 4$ is below 0 for $p_0 < 2$, and for such p_0 a marginal increase in the period 0 price will not lead to any change in the posterior distribution of c or ω_1. On the other hand, when $p_0 \geq 2$, an increase in p_0 leads to a rightward shift of the posterior distribution of c, and thus a leftward shift of the distribution of ω_1.

This simple example with inter-temporal demand shocks and a persistent marginal cost shows that observing an item on a sale may imply that the propensity to bargain on this item is lower than for another item with the same current price that is not on sale.

Next we provide an opposite example using the same uniform distributions.[23] Thus assume that everything is as before, but now c_t is drawn in each period and A is fixed. If we redo everything in terms of A, we get that, conditional on prices, A is uniform on $[\max\{2p_1, 0\}, \min\{2p_0 - 2, 2\}]$. Given that ω_1 can be rewritten as $\frac{\lambda(A-p_1)}{2}$, we arrive at the opposite conclusion regarding sales - an increase in p_0, provided that $p_0 \geq 2$, leads to an increase in the posterior distribution of A, thus to a higher expected bargaining propensity ω_1.

These two examples illustrate that depending on the nature of shocks, observing a sale may increase or decrease the expected propensity to bargain. While we do not pursue a general result here, it seems intuitive that environments with demand shocks and persistent costs will tend to be such that firms will be less willing to bargain for items that are on sale, while the opposite will hold in environments with cost shocks and persistent demand. Given that in our empirical environment sales are more likely to arise due to demand shocks,[24] we would expect that sales tend to reduce propensity to bargain.

[23] A general analysis can be conducted as above.

[24] As discussed in detail in the empirical section, sales due to cost shocks such as clearance sales (where effective marginal cost is zero) should be more prevalent post-Christmas, but we find no such evidence.