Price negotiation in differentiated product markets: An analysis of the market for insured mortgages in Canada

Jason Allen — Bank of Canada
Robert Clark — HEC Montréal
Jean-François Houde — UW-Madison

February 15, 2010

Abstract

In many differentiated product markets prices are determined through a negotiation process between buyers and sellers. Sellers post a price, but consumers may be able to negotiate a discount. The extent to which sellers negotiate may depend on consumer characteristics; ignoring the actual pricing mechanism can lead to an incomplete and biased analysis. Moreover, discounting is an important form of price discrimination, and therefore an interesting phenomenon worth measuring and understanding. Despite its prevalence, this form of pricing has been largely ignored by researchers studying market power in differentiated products markets.

We analyze detailed transaction-level data on a large set of approved mortgages in Canada between 1992 and 2004 and administered by either the Canadian Mortgage and Housing Corporation or Genworth Financial. These data provide information on features of the mortgage, household characteristics (including place of residence), and market-level characteristics. The richness of these consumer data in combination with lender-level location data (MicroMedia ProQuest) allow us to empirically examine the functioning of this important market.

We propose a model of bank choice and price negotiation that incorporates three key features of the market: differentiated services, heterogeneous bank valuations of consumers, and search costs. More generally, our goal is to build and estimate a tractable empirical model to analyze differentiated markets in which prices are negotiated rather than posted.

1 Introduction

Introduction

In many differentiated product markets prices are determined through a negotiation process between buyers and sellers. Sellers post a price, but consumers may be able to negotiate a discount. The extent to which sellers negotiate may depend on consumer characteristics; ignoring the actual pricing mechanism can lead to an incomplete and biased analysis. Moreover, discounting is an important form of price discrimination, and therefore an interesting phenomenon worth measuring and understanding. Despite its prevalence, this form of pricing has been largely ignored by researchers studying market power in differentiated products markets.

Our case study is the market for mortgages. In this market most consumers negotiate a discount from the price posted by a financial institution located in their neighborhood. Therefore, the size of the discount a consumer can negotiate should depend on the number and characteristics of lenders present in his/her local market. In practice, however, predicting discount sizes is complicated by the fact that financial institutions are selling services that are differentiated. For instance, the location of retail branches determines the cost of shopping for mortgages (horizontal differentiation), while the quality of complementary services (e.g. convenience of teller/ATM network) affects the value of contracting with each institution (vertical differentiation). Consumers also differ in terms of their preferences for these amenities, as well as their ability to negotiate the best deal.

To shed light on these issues, we analyze detailed transaction-level data on a large set of approved mortgages in Canada between 1992 and 2004 and administered by either the Canadian Mortgage and Housing Corporation or Genworth Financial. These data provide information on features of the mortgage, household characteristics (including place of residence), and market-level characteristics. The richness of these consumer data in combination with lender-level location data (MicroMedia ProQuest) allow us to empirically examine the functioning of this important market.

Ultimately our goal is to understand how the interplay of concentration, differentiation and information frictions determines the ability of banks to exert market power and potentially distort housing and credit markets. In light of the recent U.S. mortgage crisis, this represents a unique opportunity to use and extend state-of-the-art empirical IO methods to answer policy-relevant questions.

In this paper we propose a model of bank choice and price negotiation that incorporates key features of the mortgage market:

1. **Differentiated services.** Although mortgage contracts can be reasonably viewed as homogeneous, not all banks provide the same local branch-service network. Similarly, banks offer an array of financial services, which consumers use since they typically conduct their day-to-day banking at the branch that holds their mortgage.

2. **Heterogeneous valuations.** Not all consumers are equally valuable to banks. The risk of default is not an important factor in our context since we focus on fully insured contracts. Instead profitability differences originate from the complementary of services offered by banks. Richer households, for instance, are more attractive than poorer households since they are more likely to 'purchase' other services such as savings accounts, credit cards, investment accounts, property and life insurance, etc.

3. **Search costs.** An important source of dispersion in transaction prices is the heterogeneity in the ability of consumers to negotiate. We capture this heterogeneity by assuming that consumers differ in the cost of negotiating with multiple banks simultaneously, and model the search process sequentially: consumers first receive a quote from their "home" bank, and then decide whether or not pursue a more extensive search.

More generally, our goal is to build a tractable empirical model to analyze differentiated markets in which prices are negotiated rather than posted. Several important markets fit this description but the negotiation aspect has been largely ignored by the previous literature. Two approaches have typically been taken by earlier authors: (i) imputation or (ii) a monopoly assumption. For instance, in order to predict the price of cars that were not chosen by consumers, Berry, Levinsohn, and Pakes (2004) replace transaction prices by the mode price of each model, thereby ignoring the correct price determination mechanism. Adams, Einav, and Levin (2009) on the other hand ignore competition and product choice by modeling each transaction as a monopoly pricing problem. Similar examples can be found in insurance, housing, and consumer loans markets. We believe that the price-setting mechanism choice is crucial, since it leads to different counter-factual conclusions.

The paper also contributes to the search literature by explicitly introducing observed product differentiation into a search problem. Important related papers include Sorensen (2001), Hortacsu and Syverson (2004), and Hong and Shum (2006). Similarly, a related literature in Marketing and IO is concerned with heterogeneous consideration-sets (e.g. Goeree (2008)). Our approach differs from these literatures, however, since we do not assume that prices are posted. In that respect our approach is closer to the labor search literature, and in particular to Postel-Vinay and Robin (2002).

2 The Market for insured mortgages in Canada

The mortgage environment in Canada is currently dominated by the "Big 6" Canadian banks, that hold 70% of mortgages outstanding. Other players in the mortgage market include two large provincial credit unions, and hundreds of small foreign banks (that entered relatively recently), credit unions, and trust companies. We have mortgage-contract data for all lenders from 1992 and 2004, and observe the exact identity of the twelve largest lenders (i.e. the remaining lenders are classified as other trust, other credit-union, and other banks). The mortgage-market share of the Big 6 has risen sharply since 1992 when it was 38%, following deregulation of the banking sector. Prior to deregulation trust companies were the major lenders in the residential mortgage market. Deregulation removed the reserve requirements on banks that trusts did not face, making them more competitive. It also allowed banks to acquire trusts, leading to a wave of mergers between 1992 and 2000.

Here we provide details on the Canadian mortgage market, and contrast it to the American experience. The big Canadian banks operate nationally and post national prices. Posted prices of the same institution are therefore identical across markets, although the degree of discounting varies by market. This practice is different than in the United States, where even national lenders post regional mortgage rates. Local branch managers in Canada, however, have substantial authority to discount the posted price, under general guidelines from headquarters.

The biggest difference between the Canadian and American mortgage markets are interest-rate deductibility and product types. Mortgage interest payments in the United States are tax-deductible, which is not the case in Canada. Importantly, unlike their American counterparts, Canadian lenders do not use a point system. Discounts are negotiated between consumers and

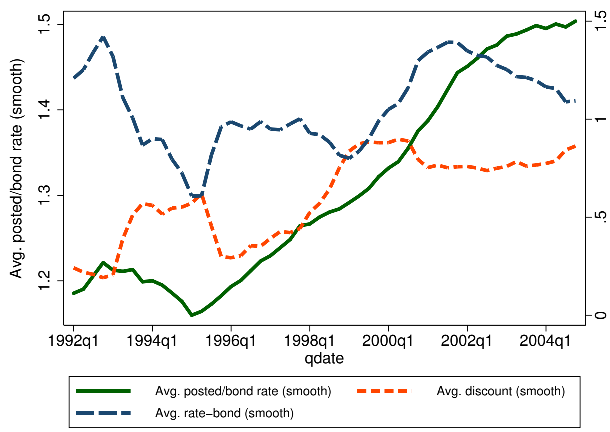

Figure 1: Discounting on 5 years mortgages between 1992 and 2004

lenders, and borrowers cannot purchase discount points to lower the mortgage rate.

In terms of the product type, the majority of Canadians sign 5-year fixed rate mortgages. At the end of the 5 years, the mortgage is rolled over with a new 5-year fixed rate contract (there are substantial pre-payment penalties in Canada as most lenders charge a fee equal to three times the monthly mortgage payment to renegotiate). The typical mortgage has a 25 year amortization period. Every 5 years, therefore, the rate is renegotiated, which in effect acts like an adjustable rate mortgage with a fixed time-frame to renegotiate. Note that the longer term mortgages that are the norm in the United States were phased out in the late 1960's in Canada after lenders experienced difficulties with volatile interest rates and maturity mismatching.

As is the case in the United States, the mortgage products offered by Canadian lenders fall into two broad categories: conventional mortgages (low loan-to-value), which are typically uninsured but can be privately insured, and high loan-to-value, which require insurance. In Canada over 50% of mortgages outstanding are insured, and over 80% of new home-owners purchase mortgage insurance. Our data are restricted to insured mortgages and we focus our attention on new-home buyers.

Mortgages are insured by either the government insurer, the Canadian Mortgage and Housing Corporation (CMHC), or a private insurer, Genworth Financial. Financial institutions require potential borrowers to buy insurance for the life of the mortgage if the amount of the loan is greater than 25% of the value of the home. The mortgage insurer charges a premium based on the household's loan-to-value ratio. The insurance premium is typically rolled into the mortgage and financed over the full amortization. In addition to the 25% rule, there are also a number of other eligibility guidelines that should be met for mortgage loan insurance. A household should have a debt service to gross income ratio of less than 32 and a total debt service ratio of less than 40.[1]

[1] Gross debt service is defined as principal and interest payments on the home, property taxes, heating costs, annual site lease in case of leasehold, and 50% of condominium fees. Total debt service is defined as all payments for housing and other debt.

To manage their insurance portfolios, CMHC and Genworth Financial collect contract and household data for each mortgage that they insure. These institutions have, subject to a confidentiality agreement, provided us with a sample of over 800,000 such contracts. This includes information on the contract itself, the borrower, and the lender. Depending on the year and the insurer, we have up to 22 variables describing the mortgage transaction, including the complete contract (price, loan, term, amortization), as well as household characteristics at the time of the mortgage application – such as income, credit score, borrower-lender relationship duration, residential status of the borrower, whether or not a broker was involved in the negotiation process, etc. The median household in our data, for example, has an annual income of 62,100 Canadian dollars,[2] a FICO score over 650 (considered above average), and a total debt service ratio of 32 (well below the limit of 40). The median household borrows 133,800 dollars to buy a 147,700 dollar home, implying a loan-to-value ratio of 90. We discuss prices in the following section.

2.1 Descriptive analysis of price distribution

An interesting feature of today's Canadian mortgage market is that transaction prices on most mortgages represent a substantial discount from the posted price. This has not always been the case. Figure 1 presents three time-series describing the 5-year fixed rate mortgage market: the average difference between the posted mortgage rate and bond rate (solid), the average difference between the bond and transaction rate (long-dash), and the average difference between the posted and transaction rate (dash). We label the last variable the *interest rate discount* negotiated between banks and consumers.

Two trends are particularly important. First, the average discount is upward trending, starting near 0 in 1992 and reaching close to 100 basis points (1%) in 2004. The second trend is the sharp increase in the gap between the average posted rate and the 5 year bond rate. Both trends suggest that banks increased the amount of discounting while also increasing their posted rate. The increase in interest rate discounts is therefore misleading, since the average margin on mortgage contracts remained fairly constant over the period (long-dash blue line). This suggests, therefore, that the increase in mortgage discounting was not a result of more intense competition between banks, but rather an attempt to better discriminate between consumers.

Figures 2(a) and 2(b) further illustrate this point by displaying the cross-sectional distribution of discounts offered by the eight largest institutions at two points in time. Both figures show a substantial amount of dispersion in transaction prices. However, in the first three periods of our sample 25% of consumers did not receive a discount, and very few received a discount exceeding 1%. In contrast, in 2000 the mode discount is very close to 1%, and only 10% of consumers paid the posted rate.[3] Despite these changes, the shape of the overall price distribution suggests that a large number of consumers do not search intensively when shopping for their mortgage.

[2] All dollar amounts are in 2002 Canadian dollars.

[3] Both graphs suggest that a portion of consumers paid an interest rate above the posted rate. This is misleading because of "90-day" rate guarantees which are common, but not always available to Canadians, but which need to be taken into consideration when calculating discounts.

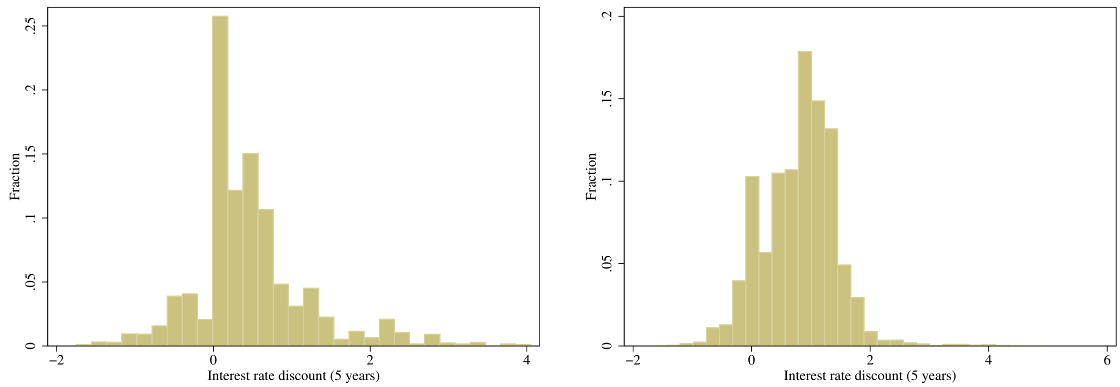

(a) Discounts between 1992 − 1995

(b) Discounts between 2000 − 2003

Figure 2: Distribution of interest rate discounts for five years mortgages at two points in time

Next we describe the way in which consumer characteristics drive discounting, controlling for trends in posted prices and cost of funding. We focus our attention on the period between February 2001 and October 2002, which corresponds to the post-merger period. Our estimating equation regresses discounts on household characteristics (such as income, credit score, whether a broker was used or not used), consumers' lender branch network density (relative to the average), lender fixed-effects, term-specific month fixed-effects, and location fixed-effects.

Parameter estimates for the key explanatory variables are presented in Table 1. The first two rows present the mean and coefficient of variation of each variable. The third row presents the parameter estimates for the full specification, including market, month, and bank fixed-effects, as well as consumer characteristics. The market fixed-effect accounts for 10% of the variation in discounts. Adding month effects helps us explain 23% of the variation in discounts. Adding bank fixed-effects does not explain very much variation in the discounting (2.3%). Finally, consumer characteristics help explain close to 5% of the variation in discounts.

We find a negative correlation between relative branch size and discounting. We interpret this result as a network quality effect. Institutions with greater branch network densities on average can offer smaller discounts because consumers care about the quality of complementary branch-services in addition to prices.

With respect to household-level characteristics, we find that discounting is increasing in income as well as FICO score. That is, consumers with higher income and better credit receive larger discounts. This might be because these clients have higher future value than those with low income and low credit scores and lenders compete more for these customers. It could also be that high income and high FICO score consumers search and bargain more than low-income and low-score consumers. In addition we find that new consumers receive larger discounts than existing clients. This is similar to what Goldberg (1996) finds in the automobile market. Younger clients (proxied as people living with their parents at the time of their mortgage application) also receive larger discounts than older clients - suggesting either greater sophistication in searching for a good mortgage rate (e.g. via the internet) or that institutions want to lock-in

Table 1: Evidence of price discrimination

Dependent variable: Discount

	Branch network	Income (log)	FICO score	New customer	Parents	Broker	House /income	Down	R^2
Mean	0.571	11.04	4.5	0.42	0.08	0.333	2.29	-2.53	
Coeff. Var.	0.52	0.04	0.31	1.2	3.3	1.45	0.33	-0.18	
Estimates	-0.0388**	0.124***	0.0455***	0.0660***	0.0636***	0.210***	0.0970***	0.0832***	0.301
	(0.0188)	(0.0157)	(0.00332)	(0.0108)	(0.0169)	(0.0115)	(0.00947)	(0.0106)	

R^2 can be broken down into: 0.102 with nbh fixed-effects; 0.129 from month-effects; 0.023 from bank fixed-effects and 0.047 from household characteristics. *** p<0.01, ** p<0.05, * p<0.1

clients at an early stage, or both.

We also find that people who use brokers on average receive a 21 basis point discount beyond what people who negotiate directly with their financial institution receive. Brokers act as a matching agent between consumers and lenders (unlike in the U.S. where brokers can originate loans), receiving volume discounts from lenders. Given that brokers on average lead borrowers to receive larger discounts, we are interested in understanding the demand for brokers across consumer-types.

In addition, we find that consumers purchasing more expensive homes, relative to their income, receive larger discount. This is true holding the down payment constant. One possibility is that consumers buying expensive homes search and bargain more than those who purchase cheaper homes. Another possibility is that financial institutions offer larger discounts to consumers buying more expensive homes because they face less pre-payment risk from these individuals. Pre-payment is the risk lenders face that consumers will contribute over their monthly payment to pay down the mortgage at a quicker rate than the lender expected. Consumers that are more leveraged are less likely to pre-pay their mortgage, and given a fixed debt level, these consumers are less risky than consumers who are less leveraged.

Finally we provide evidence that most consumers negotiate a discount from the price posted by a financial institution located in their neighborhood. We define the location of each house as the center of its Forward Sortation Area (or FSA), which corresponds to the first three digit of the postal-code. There are over $1,300$ FSAs in Canada. Figure 3 illustrates the distribution of distances from the FSA centroid of each house to the closest branch of the chosen bank, and the average distance distance to the closest branch of competing banks. The average consumer transacts with a bank located within 2KM of their house, and 95% of transactions are done with a bank located within 5KM. On the other hand, the average distance to closest branch of competing banks is over 3KM.

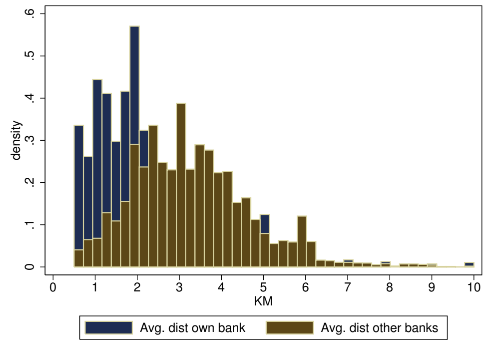

Figure 3: Distribution shortest distances between homes and banks

3 A model of price negotiation in differentiated product markets

3.1 Bank choice and price negotiation

The setup of the model is as follows. Each consumer i has the option to buy from n_i banks differentiated along two dimensions: quality and profitability. The profitability of consumer i for bank j is composed of a deterministic cost component c_{ij}, an independent private value term ϵ_{ij}, and a transaction price p_{ij}. The net value of serving consumer i for bank j is thus given by:

$$\Pi_{ij} = p_{ij} - c_{ij} + \epsilon_{ij}. \tag{1}$$

Profitability represents the match value for bank j of consumer i. It is meant to approximate the expected discounted value of transacting with consumer i. We model the deterministic component c_{ij} as a reduced form function of observed bank and consumer characteristics, and normalize the units to match the annual interest rate. For instance, in the empirical application: $c_{ij} = \bar{c}_j + W_j\alpha_0 + Z_i\alpha_1$, where $\bar{c}_j$ is a bank j fixed-effect, W_j are observed bank characteristics, including the cost of loan financing, and Z_i are demographic characteristics affecting the expected profit from complementary benefits sold to consumer i.

From the consumer's point of view, banks are differentiated in terms of the quality and the convenience of the portfolio of services that they offer. Building on the differentiated-products literature, we assume that consumers take the contractual characteristics as given (i.e. house price, term, and down-payment). Prices are determined by a multilateral bargaining process in which consumers simultaneously obtain n_i quotes. Once a consumer receives quotes $\{p_{i1}, ..., p_{in_i}\}$, he chooses the bank that offers the highest utility:

$$\max_{j \leq n_i} \quad U_{ij} = X_{ij}\beta_i + \delta_j - p_{ij} \equiv \mu_{ij} - p_{ij}, \tag{2}$$

where X_{ij} is a vector of bank j characteristics, δ_j is an unobserved vertical characteristic associated with bank j, and β_i is a vector of preference parameter specific to consumer i. As in Berry, Levinsohn, and Pakes (1995) or Nevo (2001), we model β_i as a function of observed demographics D_i and normally distributed random coefficients η_i: $\beta_i = D_i\pi + \Sigma\eta_i$.

Our setup departs from the traditional discrete choice literature since prices are negotiated rather than posted, and observed to the econometrician only for options that are not rejected by consumers. We therefore need to augment the model with a mechanism that determines the equilibrium prices offered to each consumer. In particular, we model this process as an ascending auction with independent private value (IPV), in which banks bid for a mortgage contract after observing consumer's valuations $\{\mu_{i1}, ..., \mu_{in_i}\}$ and the deterministic component of profits $\{c_{i1}, ..., c_{in_i}\}$.[4]

This is not the only mechanism to allocate products and split surplus, and we are open to considering alternative options in the future. We are not aware however of other tractable multilateral bargaining models with heterogeneous players and asymmetric information. Moreover, this modeling choice approximates the idea that consumers who choose to shop for more than one quote have the option to go back and forth between banks until an agreement is reached, thereby allowing banks to respond directly to their opponents' quotes (unlike a sealed-bid auction). The IPV assumption also guarantees the existence of a simple Nash equilibrium price.

The major difference between the standard ascending auction and our model is the fact that consumers value banks differently. In the standard model (i.e. $\mu_{ij} = \mu_{ik}$) the equilibrium price is determined by the profitability of the second-highest bidder. With differentiated sellers, the most productive bank does not automatically win the auction, since consumers might prefer to pay more for higher quality.[5] To characterize the equilibrium it is useful to think of firms as bidding "utility" levels rather than prices. For instance if the highest utility bid is $\tilde{u}$, bank j can win the contract by offering a price that leaves the consumer indifferent between buying from bank j or the next best alternative: $p_{ij} = \mu_{ij} - \tilde{u}$. As long as $p_{ij} - c_{ij} + \epsilon_{ij} = \mu_{ij} - \tilde{u} - c_{ij} + \epsilon_{ij} > 0$, it will be in bank j's interest to offer this quote to consumer i. Using a similar argument it is clear that the winning bank will be the one offering the largest *potential* surplus to the consumer:

$$f(i) = \arg\max_{j \leq n_i} \mu_{ij} - c_{ij} + \epsilon_{ij}. \tag{3}$$

Similarly $s(i)$ indexes the second-highest surplus bank. Note that by offering a quote equal to its own valuation, bank $s(i)$ is indifferent between serving and not serving this customer. The highest surplus bank can therefore win the auction by offering the following price:

$$p_i^* \quad = \quad \mu_{i,f(i)} - \big(\mu_{i,s(i)} - c_{i,s(i)} + \epsilon_{i,s(i)}\big) = \mu_{i,f(i)} - \nu_i, \tag{4}$$

where $\nu_i = \max_{j \neq f(i)} \mu_{ij} - c_{ij} + \epsilon_{ij}$ denotes the consumer surplus from the second highest bidder, and with a slight abuse of notation n_i denotes the choice set of consumer i. The statistical distribution of the winning price is therefore obtained from the distribution of the second-order statistics of surplus $\nu_i(n_i)$. Similarly the probability that consumer i contracts with bank

[4] A similar modeling choice is made in Postel-Vinay and Robin (2002) to model wage determination with on-the-job search.

[5] This setup is similar to the score auction common in the procurement literature. In our setup however, the quality of a product is fixed.

j, conditional on n_i, is given by:

$$\Pr(i \to j|n_i) = \Pr\Bigg(\mu_{ij} - c_{ij} + \epsilon_{ij} > \mu_{ik} - c_{ik} + \epsilon_{ik}, \forall j \neq k\Bigg). \quad (5)$$

Notice that the model so far does not have an outside option. We ignore the possibility for consumers of not buying a house, and focus only on a set of consumers who qualify for loan insurance. On the other hand, consumers might find it profitable to reject the highest utility offer if it yields a lower utility than: (i) the posted price of the winning bank, or (ii) the initial quote described in the next section. To save space, we are not formally describing these possibilities here, but it is straightforward to introduce them to the model. In short, the posted price and the initial quote constrain the strategy of banks very much like a reserve price in a standard auction.

3.2 Adding search frictions

We model the search process sequentially. First, each consumer visits their "home" bank, denoted by h_i, to receive an initial quote. Then, conditional on this quote and the cost of inviting more banks to the negotiation table, he decides to shop or to accept the offer. If he chooses to shop, he invites n_i banks to the auction and picks the option that yields the highest utility, as described above. For simplicity we fix n_i to the number of banks located in a constant distance radius around the house of each consumer. This number, and the identity of banks competing for consumer i is therefore common knowledge and deterministic. A similar assumption is used in the non-sequential search literature (e.g. Burdett and Judd (1983)).[6]

The cost of inviting n_i banks is heterogeneous across consumers and depends on a parameter λ_i distributed according to CDF $F(\lambda)$ in the population. This parameter is privately observed by consumers, and banks use the first quote to screen consumers with high search cost. For a consumer of type λ_i, the optimal search decision is binary: accept the value if the initial offer p^0 exceeds the expected value of searching. Formally, the type $\bar{\lambda}_i(p^0)$ is indifferent between shopping and accepting p^0:

$$\mu_{i,h_i} - p^0 = \sum_{j=1}^{n_i} \Pr(i \to j|n_i)\left[\mu_{ij} - E(p_{ij}|j, n_i)\right] - \bar{\lambda}_i(p^0) \quad \Leftrightarrow \quad \bar{\lambda}_i(p^0) = E(U_i|n_i) - \mu_{i,h_i} + p^0, \quad (6)$$

where $E(U_i|n_i)$ is the expected utility of inviting n_i banks. Notice that the threshold $\bar{\lambda}_i$ is indexed by i since consumer heterogeneity in μ and c affects the expected utility of searching. The profit maximization problem of the home bank can be written as follows:

$$\max_{p^0 \leq \bar{p}_0} \left(p^0 - c_{i,h_i} + \epsilon_{i,h_i}\right)\left[1 - F(\bar{\lambda}_i(p^0))\right],$$

where $\bar{p}_{h_i}$ is the posted price offered by bank h_i.

[6] It is straightforward to allow consumers to invite a different number of banks based on their search cost. For now, we assume this process because it simplifies the pricing decision of the home bank.

3.3 Adding search frictions (alternative model – very preliminary)

The model described previously provides a framework to understand the relationship between differentiation, market structure, and the transaction price of mortgage contracts. Another important feature of this market is the fact that consumers differ in their ability to search and negotiate the best deal.

We model the search process sequentially. Initially, each consumer visit its "home" bank to get an initial quote. Then, conditional on this quote and the cost of inviting more banks to the negotiation table, he decides to shop or accept the offer. If he chooses to shop, this consumer invite the n_i banks to the auction and pick the option that yields the highest utility as described before.

The cost of inviting n_i banks is heterogenous across consumers and depends on a parameter λ_i distributed according to CDF $F(\lambda)$ in the population. We assume that consumers rank banks according the observed surplus $\mu_{ij} + \pi_{ij}$, and use this ordering to decide which bank to invite. This ordering is deterministic and observed by banks and consumers. However, the parameter λ_i is privately observed by consumers, and banks use the first quote to screen consumers with high search cost.

For a consumer of type λ_i, the optimal search decision consists of choosing the number of banks to invite by trading off the shopping cost and the expected utility gains of requesting more bids. Formally, letting j denotes the rank of each bank with respect to $\mu_{ij} + \pi_{ij}$, the indirect utility of shopping is:

$$\begin{aligned} V_i(\lambda_i) &= \max_{n \leq \bar{n}} \sum_{j=1}^{n} \Pr(i \to j|n) \left[\mu_{ij} - E(p_{ij}|j,n)\right] - \lambda_i n \\ &= \max_{n \leq \bar{n}} E\left[U_i|n\right] - \lambda_i n, \end{aligned} \tag{7}$$

where $E[U_i|n]$ is the utility of inviting n banks, and $\lambda_i n$ is the shopping cost. It is easy to show that the both terms are monotonically increasing functions of n. Therefore for each λ_i there exists a unique number of banks that maximizes this function. The following inequalities define $n_i(\lambda_i)$:

$$E\left[U_i|n_i(\lambda_i)\right] - \lambda_i n_i(\lambda_i) \geq 0,$$
$$E\left[U_i|n_i(\lambda_i)+1\right] - \lambda_i (n_i(\lambda_i)+1) < 0.$$

These inequalities generate $\bar{n}-1$ cutoffs $\{\lambda_i^1, ..., \lambda_i^{\bar{n}-1}\}$ that define the range of types inviting n banks to the negotiation table. These cutoffs are indexed by i since consumer heterogeneity in μ and π affects the expected utility of searching.

A consumer will decide to search if the expected value of searching exceeds the value of accepting the initial offer p^0. This gives rise to a reserve price rule $p(\lambda_i)$ that is invertible: for any p^0 there exists a unique λ such that all consumers with $\lambda_i < \lambda(p^0)$ decide to search. The profit maximization problem of the home bank can therefore be written as followed:

$$\max_{p^0} \quad (p^0 + \pi_{i0} + \epsilon_{i0}) \left[1 - F(\lambda(p^0))\right] \tag{8}$$

The difficulty however is that since the number of banks is discrete, the value function defined in equation 7 is not a smooth function of λ. As a result, to solve the profit maximization problem of the home bank we cannot differentiate the inverse reserve price function.

4 Functional form assumptions and empirical implementation

Our model involves a number functions and distributions that need to be parameterized before bringing it to the data. We are particularly interested in choosing distributional assumptions for the private value shock ϵ_{ij} and the search cost λ_i such that the outcomes can be written as analytical functions of the primitives: $\beta_i, \delta_j, \bar{c}_j, \alpha$. This will greatly simplify the estimation procedure and allow for the inclusion of rich sources of consumer heterogeneity.

Following Brannan and Froeb (2000), we assume that $\{\epsilon_{i1}, ..., \epsilon_{in_i}\}$ are independently distributed according to an extreme-value distribution with location and spread parameters $(0, \sigma)$. Given this assumption, the bank-specific cost component, $c_{ij} - \epsilon_{ij}$, has an extreme-value distribution, as does the "consumer surplus offer" $\mu_{ij} - c_{ij} + \epsilon_{ij} = \pi_{ij} + \epsilon_{ij}$; the latter with parameters (π_{ij}, σ). Similarly, the maximum surplus offer, $\max\{\pi_{ij} + \epsilon_{ij} : 1 \leq j \leq n_i\}$, is also distributed according to an extreme-value variate with parameters $(\pi_{i,\max}, \sigma)$ where: $\pi_{i,\max} = \frac{1}{\sigma}\log\left(\sum_{k=1}^{n}\exp\left(\sigma\pi_{ik}\right)\right)$. With this we can determine the winning probability of bank j:

$$\Pr(i \to j|n_i) = \frac{\exp\left(\sigma\pi_{ij}\right)}{\exp\left(\sigma\pi_{i,\max}\right)} = \frac{\exp\left(\sigma\pi_{ij}\right)}{\sum_{k=1}^{n_i}\exp\left(\sigma\pi_{ik}\right)}. \tag{9}$$

Next, we can determine the expected winning bid of the auction. We are interested in determining the expected value of the second highest surplus offer conditional on bank j winning. This is the discount or premium that j will have to pay relative to consumer i's willingness to pay (i.e. ν_{ij} in equation 4). The conditional mean of the jth bank's winning bid is given by:

$$E(p_{ij}|j, n_i) = \mu_{ij} - E\left(\nu_{ij}|j, n_i\right) = \mu_{ij} - E(\pi_{i\max}|j, n_i) - \frac{\log\left((1 - \Pr(i \to j|n_i)\right)}{\sigma\Pr\left(i \to j|n_i\right)}. \tag{10}$$

Finally, assuming that the search cost parameter λ_i is distributed according to an exponential distribution with mean $\frac{1}{\kappa}$ yields an analytical expression for the initial offer p^0.

In particular, in the interior the optimal first quote from the home bank of consumer i solves the following first-order condition:

$$\begin{aligned} &\left(1 - F\left(\bar{\lambda}_i(p^0)\right)\right) - f\left(\bar{\lambda}_i\right)\left(p^0 - c_{i,h_i} + \epsilon_{i,h_i}\right) = 0 \\ &\Rightarrow p^0_{i,h_i} = \frac{1}{\kappa} + c_{i,h_i} - \epsilon_{i,h_i} \end{aligned} \tag{11}$$

The other unobservable from the econometrician's point of view is the distribution of consumer heterogeneity. We have already discussed the distribution of β_i. Another form of consumer heterogeneity is the identity of the home bank, denoted h_i. As we discussed in the data section, this information is only partly observed, and only for households who chose to remain with their main financial institutions. Rather than parameterize the distribution of this unobservable we will use another source of information that provides the joint distribution of financial institution affiliation and consumer characteristics. In particular, we have access to an extensive survey of consumer finances run by Ipsos-Ried. This survey tracks, among other things, the identity of the main financial institution of a random sample of Canadian households between 1999 and 2007.[7] This data-set will allow us to calculate the empirical distribution of bank affiliation conditional on covariates observed in the mortgage data. Let $\hat{g}(j|D)$ denote this density function for bank j and consumer characteristics D, most notably income and location.

We are now ready to describe the estimation procedure. In principle we could estimate the model using aggregate data, by employing a GMM procedure similar to Berry, Levinsohn, and Pakes (1995). For instance, defining the market as a Province and time period, the market share (equation 9) and average winning price of each bank (equation 10) can be inverted to recover the average quality and cost terms: δ_j and $\bar{c}_j$. One could then form empirical moment conditions to identify the remaining non-linear parameters. However since we have rich micro-data matching transaction prices with consumer characteristics, we believe that a likelihood approach is more appropriate. We therefore envision estimating the model using a similar estimation procedure as the one used by Goolsbee and Petrin (2004).

Here we briefly describe the likelihood function of the key outcome variables. Recall that we take as given the demographic characteristics of households and the contract terms other than the interest rate. We focus on homogeneous mortgage contracts with 5 year fixed-rate terms and 25 years amortization periods, and ignore contracts that are negotiated through a broker. This last selection criteria will be relaxed in the actual empirical implementation (see discussion below).

The model predicts three outcomes: the transaction price P_i, the identity of bank B_i, and whether or not consumers transact with their home bank H_i. Let $Y_i = \{P_i, B_i, H_i\}$ denote the vector of outcome variables. Conditional on consumer heterogeneity $\{\beta_i, h_i\}$, observable characteristics $O_{ij} = \{X_{ij}, W_j, Z_i, D_i\}$, and parameter vector θ, the density of transaction price and bank choice is given by the following expression:

$$l(P_i, B_i|H_i, O_{ij}, \theta, \beta_i, h_i) = \begin{cases} \phi_\nu\left(\mu_{i,B_i} - P_i|n_i, \theta, \beta_i\right) \Pr(i \to B_i|n_i, \theta, \beta_i) & \text{If } H_i = 0, \\ \psi_\epsilon\left(\frac{1}{\kappa} + c_{i,h_i} - P_i|\theta\right) & \text{If } H_i = 1. \end{cases}$$

where $\phi_\nu(\cdot)$ is the density of the second-order statistics defined in equation 4, and $\psi_\epsilon(\cdot)$ is the density ϵ.[8] The probability of accepting the initial quote is derived directly from the distribution of search cost. The joint likelihood the observed outcome Y_i is thus given by:

$$\begin{aligned} l(Y_i|O_i, \theta, \beta_i, h_i) \quad &= \quad H_i \times \left(1 - F\left(\bar{\lambda}_i(p^0_{i,h_i})\right)\right) l(P_i, B_i|1, O_{ij}, \theta, \beta_i, h_i) \\ &\quad + (1 - H_i) \times F\left(\bar{\lambda}_i(p^0_{i,h_i})\right) l(P_i, B_i|0, O_{ij}, \theta, \beta_i, h_i). \end{aligned} \tag{12}$$

We will then use simulation methods to evaluate the likelihood contribution of each observed individual. In particular, if S denotes the number of simulated draws for β_i, the simu-

[7] See Allen, Clark, and Houde (2008) for more details on the data.

[8] Notice that the previous expression is incomplete since it ignores the possibility of a corner solution in prices. We ignore this possibility to save on space.

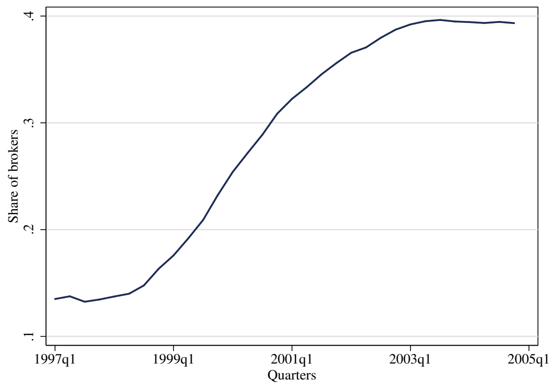

(a) Market share of broker transactions

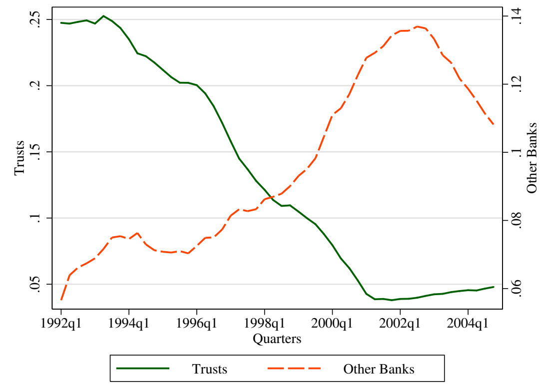

(b) Market share of Trusts and Other banks

Figure 4: Trends in the market share of broker transactions and other financial institutions (smoothed)

lated likelihood of individual i's outcome is given by:

$$\mathcal{L}_i(Y_i|O_i,\theta) \approx \frac{1}{S}\sum_s \left[\sum_{j=1}^{J} \hat{g}(j|D_i) l(Y_i|O_i,\theta,\beta_s,j)\right], \tag{13}$$

where J is the maximum number of available options.

5 Extensions and policy related questions

The model described above provides a novel framework in which to empirically analyze markets with differentiated products, search frictions, and negotiated prices. In developing the model one objective was tractability, and as result several features of the market have been omitted. In this section we present four features of the market that deserve to be incorporated in future iterations. All four lead to important policy questions that we wish to explore.

Mortgage brokers. During our sample period, the penetration of mortgage brokers has become very important. Since 1997 the share of transactions negotiated through a broker has more than doubled, from less than 15% to nearly 40% at the end of 2004. Importantly, this diffusion process varies across region and coincides with the rapid increase in discounting documented in Section 2.1. Our model allows us to study the diffusion of brokers, evaluate the extent to which this diffusion is optimal and to evaluate its benefits for consumers and banks. We can therefore contribute to the literature on the benefits of intermediation in markets (see for instance Rubinstein and Wolinsky (1987), Rust and Hall (2003), or Hendel, Nevo, and Ortalo-Magne (2010)).

In Canada brokers do not issue or sell mortgages, but rather serve as an intermediary to shop for the lowest interest rate available. The financial cost is borne by the lender, therefore it is not surprising to see a rapid diffusion of brokers in this market. There are two inefficiencies, however, that may limit this diffusion. First, brokers may not optimally search and therefore only receive a small number of quotes. Second, since brokers are negotiating on the behalf of consumers, the preference of the consumer might not be fully reflected in the quotes. Taking into account these inefficiencies, it is straightforward to incorporate brokers into our model by allowing consumers to choose brokerage services if they do not like their initial quote.

After observing the initial quote, consumers might have an additional option of using a broker rather than searching the market. Consumers with intermediate search costs will prefer this option if the expected utility of using a broker is larger than accepting the initial quote and smaller than the gain of searching. The value of using a broker can be made smaller by reducing the number of searches (i.e. $n_{broker} \leq n_i$). In addition, if consumers cannot communicate perfectly their preferences (i.e. the μ_{ij}'s), the broker will implement the negotiation as a pure ascending auction. As a result the winning bank will be the one offering the lowest interest rate, rather than the highest utility. This will automatically reduce the value of using a broker, since consumers could end up with a lower value match relative to the search option.

Internet banking. Another important trend in the mortgage market is the arrival of virtual banks at the end of the 1990's, and especially the entry of ING in 1997. ING, for example, holds approximately 5% of mortgages outstanding. While this number might appear small, virtual banks bring a new business model to this market: interest rates are not negotiated and banks post a common rate across the country that is significantly below the posted rate of major banks. As a result, the entry of ING can distort the market by attracting richer households with high search costs and low valuation for complementary services. Incorporating virtual banks into the model allows us to contribute to the ongoing debate about the effect of the diffusion of the internet on the functioning of retail markets.[9] In terms of the model, virtual banks constrain the bidding strategy of brick-and-mortar banks by providing a credible outside option to consumers who are familiar with the Internet and who do not value the density of a branch network.

Alternative pricing mechanisms. A related question is the profitability and welfare consequences of alternative pricing strategies. It is clear that by directly negotiating discounts with consumers, banks can discriminate between consumers based on their valuations and search costs. On the other hand, consumers willing to search can extract sizable discounts. The overall welfare consequence of having negotiated rather than posted prices is therefore an empirical question.

There exists an extensive literature, starting with Riley and Zeckhauser (1983), that compares the profitability of posted-price and haggling pricing strategies. When it is costly for sellers to haggle with consumers, Riley and Zeckhauser (1983) argue that it is always more profitable to commit to a no-haggle strategy in order to reduce the probability of high-price offer rejection. Subsequent authors have enriched the analysis by adding search frictions on the

[9] See for instance Scott-Morton, Zettelmeyer, and Silva-Risso (2001), Ellison and Ellison (2009), Goldmanis, Hortacsu, Syverson, and Emre (2010), Allen, Clark, and Houde (2008).

part of consumers, product differentiation, and alternative bargaining mechanisms.[10] Interestingly, the inability to commitment on the part of the seller typically prevents firms from using strictly posted prices in equilibrium (see for e.g. Corts (1998)). A large fraction of this literature is therefore interested in explaining the fact that some industries are able to co-ordinate on posted-price transactions (e.g. large department stores), while others, like the mortgage industry, have settled on negotiated prices. The car industry has received special attention since a group of retailers have recently succeeded in using no-haggle pricing policies while competing with haggling retailers.[11]

Our setup is ideal for studying these questions. Posted prices naturally impose an upper bound on the negotiation process, similar to a reserve price in an auction. We can therefore use our model to compute various counter-factual experiments in which one or more banks commit to a no-haggle policy. Fortunately, one Canadian bank did experiment with a no-haggle pricing policy towards the end of our sample period. The strategy was unprofitable however, and the company reverted to negotiated prices after 2 years.

Another advantage of studying these questions is that between 2002 and 2004, the leading bank in the market, TD-Canada Trust, committed to a no-haggle policy based a national posted interest rate. During this period, TD-Canada Trust posted a rate significantly below its competitors' posted rates and reduced markedly the dispersion of transaction prices. The strategy was not profitable however, and the company reverted to fully negotiated prices in 2004. It was reported at the time that the strategy would have been profitable for TD-Canada Trust if other major banks had followed with similar announcements. This is consistent with the theory literature, since in an asymmetric situation the no-haggling price can be used by consumers to negotiate a better deal with other competitors, thereby reducing the market share of the non-haggling firm and increasing the average discount to consumers.

Endogenous loan size with non-linear budget constraints. The theoretical model assumes that consumers negotiate interest rates by taking as given the other characteristics of the contract, including the term, the loan size and the house price. This assumption is justified because the objective of the project is to measure market power in this market and therefore estimate consistently substitution patterns between banks. However, this assumption ignores an important channel through which firms can exert market power, namely the reduction of loan sizes and distortion in housing choices.

We plan on modeling the additional margin of loan size and house price in a future paper. This is not a simple task because the insurance premium is set as a piece-wise linear function of the down-payment, and therefore the optimal loan size can change discontinuously when consumers receive lower rate offers. Figure 5(a) shows that the distribution of down-payments is discontinuous around the points where the insurance premium changes. This is strong evidence that the insurance premium matters in determining the optimal down-payment. Figure

[10] A non-exhaustive list of contributions include: Wang (1993), Bester (1993), Wang (1995), Arnold and Lippman (1998), Corts (1998), and Desai and Purhit (2004).

[11] CarMax in the used-car market and Saturn are two notable example. Desai and Purhit (2004) provide a theoretical discussion, Huang (2009) analyze empirically the case of CarMax, and Zeng, Dasgupta, and Weinber (2008) study the introduction a no-haggle policy by Toyota in Canada.

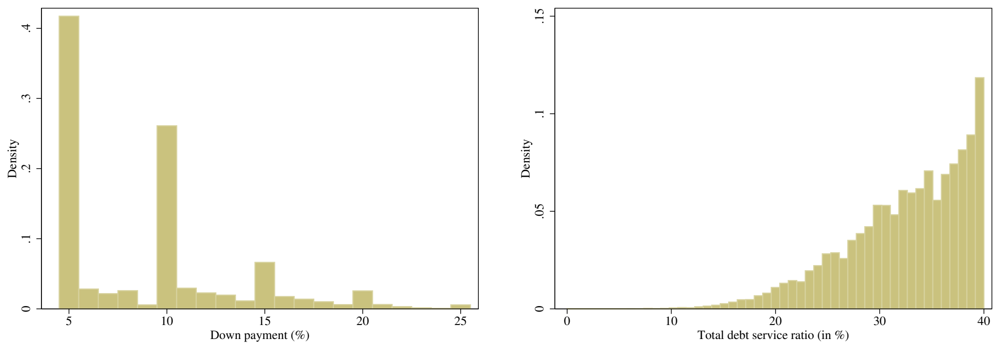

(a) Down-payment

(b) Debt service to income ratio

Figure 5: Distribution of down-payment choice and debt service ratio between 2000 and 2004

5(b) on the other hand plots the density of the debt service to income ratio, which is constraint to be below 40% by the insurance company. That limit and the minimum down-payment of 5% are clearly binding for a large number of households.

Moreover, the regulation of these constraints and the price schedule of mortgage insurance was changed between 1992 and 2005. These changes and the fact that we observe variation in consumer financial characteristics will allow us to identify important preference parameters related to loan and housing demand elasticities.

References

Adams, W., L. Einav, and J. Levin (2009). Liquidity constraints and imperfect information in subprime lending. *American Economic Review 99*, 49–84.

Allen, J., R. Clark, and J.-F. Houde (2008). Market structure and the diffusion of e-commerce: Evidence from the retail banking industry. Bank of Canada Working Paper No. 2008-32.

Arnold, M. and S. A. Lippman (1998, July). Posted prices versus barganing in markets with asymmetric information. *Economic Inquiry*, 450–457.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society 63*(4), 841–890.

Berry, S., J. Levinsohn, and A. Pakes (2004, Feb). Differentiated products demand systems from a combination of micro and macro data: The new car market. *The Journal of Political Economy 112*(1), 68.

Bester, H. (1993). Bargaining versus price competition in markets with quality uncertainty. *The American Economic Review 83*(1), 278–288.

Brannan, L. and L. M. Froeb (2000, May). Mergers, cartels, set-asides, and bidding preferences in asymmetric oral auctions. *The Review of Economic and Statistics 82*(2), 283–290.

Burdett, K. and K. Judd (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society 51*, 955–970.

Corts, K. S. (1998, Summer). Third-degree price discrimination in oligopoly: All-out competition and strategic commitment. *The Rand Journal of Economics 29*(2), 306–323.

Desai, P. S. and D. Purhit (2004). "let me talk to my manager": Haggling in a competitive environment. *Marketing Science 23*(2), 219–233.

Ellison, G. and S. Ellison (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica: Journal of the Econometric Society 77*(2), 427–452.

Goeree, M. S. (2008). Limited information and advertising in the us personal computer industry. *Econometrica: Journal of the Econometric Society 76*(5).

Goldberg, P. K. (1996, Jun.). Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy 104*(3), 622–654.

Goldmanis, M., A. Hortacsu, C. Syverson, and O. Emre (2010). E-commerce and the market structure of retail industries. *Economic Journal*.

Goolsbee, A. and A. Petrin (2004, March). The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica 72*(2), 351.

Hendel, I., A. Nevo, and F. Ortalo-Magne (2010). The relative performance of real estate marketing platforms: Mls versus fsbomadison.com. *The American Economic Review*.

Hong, H. and M. Shum (2006). Using price distributions to estimate search costs. *RAND Journal of Economics 37*(2), 257–275.

Hortacsu, A. and C. Syverson (2004, May). Product differentiation, search costs, and competition in the mututal fund industry: A case study of s&p 500 index funds. *Quarterly Journal of Economics*, 403–456.

Huang, G. (2009). Posted price and haggling in the used car market. Working paper, John Hopkins University.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307.

Postel-Vinay, F. and J.-M. Robin (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica 70*(2), 2295–2350.

Riley, J. and R. Zeckhauser (1983, May). Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics 98*(2), 267–289.

Rubinstein, A. and A. Wolinsky (1987). Middlemen. *The Quarterly Journal of Economics 102*(3), 581–593.

Rust, J. and G. Hall (2003). Middlemen versus market makers: A theory of competitive exchange. *Jounal of Political Economy 111*.

Scott-Morton, F., F. Zettelmeyer, and J. Silva-Risso (2001). Internet car retailing. *Journal of Industrial Economics 49*(4), 501–519.

Sorensen, A. (2001, October). An empirical model of heterogeneous consumer search for retail prescription drugs. Working Paper 8548, NBER.

Wang, R. (1993). Auctions versus posted-price selling. *American Economic Review 83*(4), 838–851.

Wang, R. (1995). Bargaining versus posted-price selling. *European Economic Review 39*, 1747–1764.

Zeng, X., S. Dasgupta, and C. B. Weinber (2008, September). The competitive implications of a "no-haggle" pricing policy: The access toyota case. Working Paper, Univerity of British Columbia.