THE EFFECT OF MERGERS IN SEARCH MARKET:
EVIDENCE FROM THE CANADIAN MORTGAGE INDUSTRY

Jason Allen
Robert Clark
Jean-François Houde

The Effect of Mergers in Search Market: Evidence from the Canadian Mortgage Industry
Jason Allen, Robert Clark, and Jean-François Houde
NBER Working Paper No. 19126
June 2013
JEL No. L11,L41,L85

ABSTRACT

We examine the relationship between concentration and price dispersion using variation induced by a merger in the Canadian mortgage market. Since interest rates are determined through a search and negotiation process, consolidation eliminates a potential negotiation part- ner, weakening consumers bargaining positions. We combine reduced-form techniques to es- timate the mergers distributional impact, with a structural model to measure market power across consumers with different search costs. Our results show that competition benefits only consumers at the bottom and middle of the transaction price distribution. Estimates from a search and negotiation model attribute these differences to the presence of large search frictions.

Jason Allen
Bank of Canada
234 Wellington Street
Ottawa, ON K1A 0G9
jallen@bankofcanada.ca

Robert Clark
Institute of Applied Economics
HEC Montreal
3000, chemin de la Cote-Sainte-Catherine
Montreal (Quebec) H3T 2A7
Canada
Robert.Clark@hec.ca

Jean-François Houde
Wharton School
University of Pennsylvania
3620 Locust Walk
3008 Steinberg Hall - Dietrich Hall
Philadelphia, PA 19104
and NBER
houde@wharton.upenn.edu

Introduction

Like many consumer-finance products, mortgage contracts are negotiable: lenders post a common sticker price, and contract terms are determined through a search and negotiation process between local branch managers and individual borrowers. This allows informed consumers to gather multiple quotes and obtain an interest rate that reflects the expected lending cost, even with a small number of competing lenders. In practice, however, consumers differ in their ability to understand the subtleties of financial contracts and their willingness to negotiate and search for multiple quotes. Indeed, recent surveys in North America show that while some buyers get multiple quotes when shopping for their mortgage contract, nearly half only get one.[1]

These features lead to price dispersion. In Allen et al. (2013) we document that the Canadian mortgage market exhibits substantial dispersion, even for contracts with homogeneous terms and for which lenders are fully protected against the risk of default by a government-backed insurance program. The inter-quartile range of the net transaction interest spread approaches 100 basis points and the coefficient of variation is about 59%. Importantly, more than 60% of the dispersion in margins is unexplained by standard borrower and contract characteristics.

Similar phenomena have been documented by a growing literature analyzing price dispersion in retail markets: new cars (e.g. Goldberg (1996), Scott Morton et al. (2003), Busse et al. (2006), and Langer (2012)), mortgage broker fees (Hall and Woodward (2012)), real-estate (Hendel et al. (2009)), and health care services (e.g. Sorensen (2001), Grennan (2013)). Building on Stigler (1961), many of these papers have provided evidence that a significant fraction of the observed dispersion in prices is caused by the inability of (some) consumers to gather information and negotiate discounts, suggesting that search and information frictions are important factors in these markets.

In contrast, we know relatively little about the impact of market structure and competition on the distribution of negotiated prices.[2] It remains an empirical question whether the benefit of competition is spread equally across consumers, and whether competition raises or lowers the dispersion of transaction prices. Indeed, search-theoretic models of price dispersion provide ambiguous predictions with respect to the impact of concentration on price dispersion, as discussed in Janssen and Moraga-González (2004) and Baye et al. (2006).[3]

Our objective is to measure the role of competition in determining the distribution of prices in markets with price negotiation. Unlike in markets without search frictions, in our context the

---

[1] In Canada, the Canadian Association of Accredited Mortgage Professionals conducts an annual survey on the state of the mortgage market and on average have reported that about 45% of new home buyers visit one lender. For similar U.S. evidence see Lee and Hogarth (2000).

[2] An exception is the recent literature on negotiation between buyer and seller networks (e.g. Town and Vistnes (2001), Crawford and Yurukoglu (2012), Lewis and Pflum (2013), Gowrisankaran et al. (2013)). While it analyzes the impact of market structure on the relative bargaining leverage of parties, it does not provide a convincing interpretation of residual price dispersion, and abstracts from the role of information and search frictions.

[3] The same is true for discrimination-based theories: the "textbook" price discrimination model suggests a negative relationship between dispersion and competition, while Borenstein (1985) and Holmes (1989) show that an increase in competition in markets with third-degree price discrimination lead to more dispersion.

effect of competition on prices depends not only on the relative market power of firms, but also on the willingness and ability of consumers to haggle with sellers. If there is heterogeneity in these abilities across consumers and markets, then variation in transaction prices induced by changes in competition may not reveal the importance of market power. Because of this ambiguity, we combine reduced-form econometrics techniques to estimate the impact of losing a competitor on negotiated prices, with a structural econometric model to measure the extent of market power across consumers with different search and negotiation costs.

To perform our analysis, we use administrative data on insured mortgage contracts, and take advantage of the quasi-experimental variation created by a horizontal merger between two important mortgage lenders in Canada. Our empirical strategy relies on the idea that the merger of two banks' networks creates discrete changes in the choice-set of consumers located near the branches of both merging parties, while the number of options offered to consumers living close to only one or neither remains unchanged.

This variation naturally leads to a difference-in-difference estimator with which we infer the average price changes among consumers directly affected by the merger by looking at the evolution of prices in comparable local markets in which the number of lenders remained constant.[4] Our first set of results shows that the loss of a competitor led to an increase in the average interest rate in treated markets of approximately 6 basis points (bps), which corresponds to about 15% of residual price dispersion in our sample.

We then estimate the distributional effect of the merger, an aspect of merger evaluation that has been ignored up to now by the literature. To do so, we use the *change-in-change* estimator proposed by Athey and Imbens (2006) to recover the counter-factual distribution of negotiated rates that would have been observed if the merger had not been approved by antitrust authorities. This allows us to document substantial heterogeneity in the impact of the merger along the distribution of negotiated rates. In particular, our second set of results show that the loss of a competitor increased interest rates between 7 and 9 bps for consumers in the lower and middle percentiles of the distribution, and had no effect on consumers in the top 30%. This result implies that borrowers who are the most adversely affected by a negotiation-based pricing policy do not benefit from an increase in the degree of competition between lenders.

Having estimated the pre- and post-merger price distribution, we measure the impact of the merger on residual price dispersion. Our third set of results shows that the merger led to a 16% decrease in the inter-quartile range, while the coefficient of variation decreased by 15%. This establishes a positive relationship between the number of firms and residual price dispersion in markets with negotiated prices.

We then develop and estimate a model of search and price negotiation that is consistent with

---

[4]This strategy has recently been used to study the impact of mergers in gasoline markets (e.g. Hastings (2004), and Houde (2012)), cement industry (e.g. Hortacsu and Syverson (2007)), health-care industry (e.g. Dafny et al. (2011)), and durable goods (e.g. Weinberg et al. (2013)).

our three reduced-form results. The parameters of the structural model are estimated by inverting the observed post-merger price distribution and matching the predicted merger effect obtained from the *change-in-change* estimator. This approach allows us to estimate a non-parametric search-cost distribution, as done in recent papers estimating search models in markets with price dispersion (e.g. Hortaçsu and Syverson (2004), Hong and Shum (2006), and Wildenbeest (2011)).

Finally, we use the estimated model to identify the contribution of consumers' search-costs in generating market power among mortgage lenders. In our setting, because consumers endogenously respond to changes in market structure by reducing their search effort, the price effect of losing a competitor is the sum of a reduction in search costs and an increase in the price that consumers expect to pay if negotiation fails. For instance, at the estimated parameters, the merger reduced the bargaining leverage of the median consumer by about 6 bps, which corresponds to a 9 bps increase in the price of the next best alternative, and a 3 bps search-cost reduction.

Moreover, we show that the net effect of mergers is very different in markets with lower search costs. In particular, reducing consumer search costs by half would lead to a 30% *increase* in the average effect of the merger, and 46% *more* homogeneous effects across consumers. These results imply that the reduced-form impact of mergers in markets with large search-costs is not an accurate representation of firms' market power.

Our paper is related to an extensive literature on bank mergers (see Berger et al. (1999) for a discussion). A lack of consumer-level data has made it difficult to analyze the effect on transaction prices, and therefore most studies have focused on the impact of mergers on average transaction prices or posted deposit services (fees and rates). For instance Prager and Hannan (1998) find that bank mergers in the U.S. led to a decrease in deposit rates. Using Italian data, Focarelli and Panetta (2003) find a decrease in the short-run, but an increase in the long run due to efficiency gains resulting from the merger. Finally, our results are related to those of Sapienza (2002) who, in the context of business lending, finds that borrowers with few outside banking relationships are significantly less affected by mergers, while those with an intermediate number of banking relationships are affected the most. Although the methods for identification are different, her interpretation of the economic channel through which the results are derived is similar to ours.

The rest of the paper is structured as follows. Section 1 describes the Canadian mortgage markets, focusing on market structure and pricing. Section 2 presents our sample and defines the variables used in the analysis. In Section 3 we provide we describe our empirical approach and explain our identification strategy. Section 4 presents results of the reduced-form analysis. In section 5 we develop and estimate a model of search and price negotiation. Section 6 concludes and presents implications of our results for mortgage-market policies.

3

1 The Canadian mortgage market

1.1 Market structure and mergers

The Canadian mortgage market is currently dominated by six national banks (Bank of Montreal, Bank of Nova Scotia, Banque Nationale, Canadian Imperial Bank of Commerce, Royal Bank Financial Group, and TD Bank Financial Group), a regional cooperative network (Desjardins in Québec), and a provincially owned deposit-taking institution (Alberta's ATB Financial). Collectively, they control 90 per cent of assets in the banking industry and we conveniently call them the "Big 8."

The market was not always this concentrated. Until the early 1990s the Canadian residential-mortgage market also featured a large number of trust companies. Trusts are like Savings & Loans in the U.S. At the time the main difference between trusts and banks was that trusts were more lightly regulated with regards to reserve requirements. In particular, trusts did not have to hold reserves against mortgages, while chartered banks did. This provided trusts with a competitive advantage in the mortgage market due to a lower cost of funding. Cross-ownership between the two types of institutions was not permitted until the 1992 revisions to the Bank Act. Following these revisions banks and trusts were granted almost identical powers, making them undifferentiated products from the point of view of consumers.[5]

As a result of the Bank Act revisions and a series of bad residential and commercial loans that created solvency and liquidity issues for the trusts in the 1980s, Canadian chartered banks acquired the majority of trust companies over the course of the following decade. The merger wave led to the six largest banks controlling approximately 80 per cent of the mortgage market – almost double their 1980s market share. These mergers all resulted in significant expansion of the merged entity's branch network since in each case the Canadian Competition Bureau required little or no forced divestiture of branches.

Figure 1 presents the evolution of the mortgage-market share of the main lending groups – The Big 8, Trusts, Credit Unions and other banks as well as the major mergers, including the number of branches acquired. The figure also lists the major trust acquisitions along with the number of branches acquired in each case. The major acquisitions were: Canada Trust & Toronto-Dominion (2000), National Trust & Scotia Bank (1997), Montreal Trust & Scotia Bank (1994), Royal Trust & Royal Bank (1993), and Central Guaranty Trust & Toronto-Dominion (1993). A more detailed discussion of the major Canadian bank mergers is presented in the Appendix.

Our empirical analysis focuses on the short-run impact of one of the major mergers between a bank and a trust. As a result we study contracts signed within a year of the merger. This is

---

[5]There were still differences in ownership structure (trust companies could be closely held - and commercial ownership of trusts became common, while banks had to be widely held to prevent ownership concentration) as well as in supervisory authority (banks are federally regulated whereas trust companies can be federally or provincially regulated), but these differences are unlikely to affect consumer demand. In 1992 trusts were given full consumer lending powers, and banks were permitted to offer in-house wealth management advice (fiduciary services).
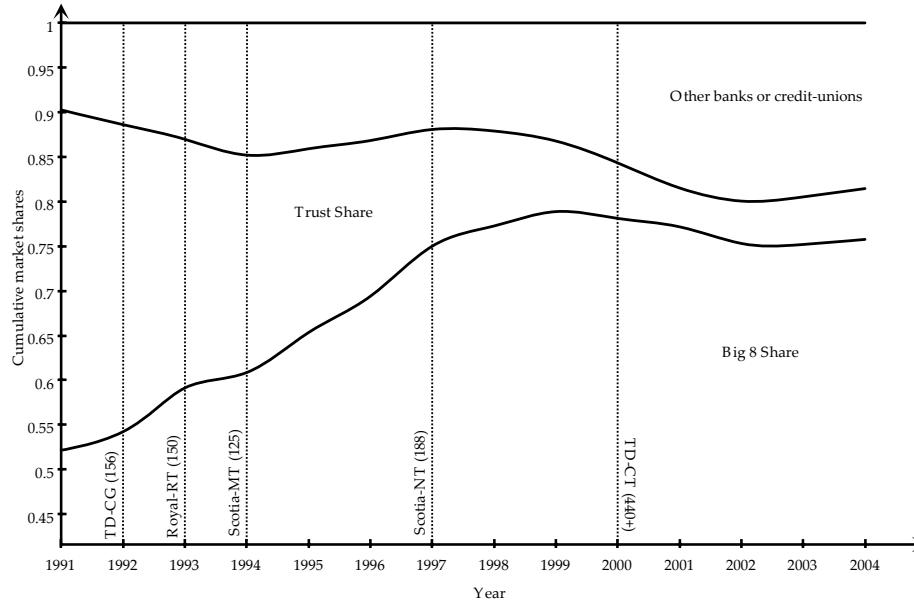
Figure 1: Evolution of financial institution market shares for newly insured mortgages (smoothed)

appropriate, since we know that the merged entity did not start closing duplicate branches until approximately a year after the official merger date. For confidentiality reasons cannot reveal the parties involved in the merger. We therefore label the two institutions A and B, and hide the exact timing of the merger.

1.2 Pricing and negotiation

The large Canadian banks operate nationally and post prices that are common across the country on a weekly basis in both national and local newspapers, as well as online. There is little dispersion in posted prices, especially among the largest banks. In contrast there is a significant amount of dispersion in transaction rates. This comes about because borrowers can search for and negotiate individual discounts. One option for borrowers is to visit local branches and negotiate directly with branch managers who have the authority to offer borrowers discounts below the posted price under general guidelines from headquarters.

Negotiating larger discounts is costly for the local bank manager, reducing the commissions earned by branch employees (see KPMG (2008)), but worthwhile if a consumer is likely to switch to another financial institution without a discount. Local branch managers compete against rival banks, but not against other branches of the same bank.[6] Survey evidence from the Canadian

---

[6]Borrowers must present credible quotes to competing branches, and branch managers are explicitly told not to compete against each other. We know from private conversations with loan officers that only in extremely rare occasions will a branch manager deviate from this directive. There is also important record keeping that prevents branch managers from the same institution from making competing offers. First, a pre-approved mortgage is typically good for

Association of Accredited Mortgage Professionals reports that about 52% of new home buyers visit more than one lenders when shopping for a mortgage (Dunning (2010)).

Alternatively borrowers can hire brokers to search for the best rates on their behalf. Unlike in the United States (except California), brokers in Canada have fiduciary duties. Brokers are "hired" by borrowers to gather the best quotes from multiple lenders, but are compensated by lenders an amount equal to 1-1.3 per cent the value of the mortgage. According to detailed survey evidence collected by Taddingstone in 2005 brokers on average contact 5.9 lenders for their clients, suggesting they do, in fact, assist in gathering multiple quotes. See Allen et al. (2013) for a more detailed description of the pricing strategies used in the market.

2 Sample selection and variable definition

Our analysis focuses on insured mortgages, and we use administrative data obtained directly from the two insurers operating in Canada: Canada Mortgage and Housing Corporation (CMHC), and Genworth Financial. CMHC is a crown corporation with an explicit backstop from the federal government. Genworth also receives an explicit government of Canada guarantee, but for 90 per cent of the amount owing. During our sample period, both insurers used the same strict federal approval guidelines: (i) borrowers with less than 25% equity must purchase insurance, and (ii) borrowers with monthly gross debt payments that are more than 32% of gross income or a total debt service ratio greater than 40% are rejected.[7] The government also sets an insurance premium that is solely a function of the loan-to-value ratio, and ranges from 1.75 to 3.75 per cent of the loan. The qualifying rules and premiums are common across lenders and based on the posted rate. Borrowers qualifying at one bank, therefore, should assume that they can qualify at other institutions, given that the lender is protected in case of default.

We construct a 10% random sample of all contracts issued between 1992 and 2004. We have access to 20 household/mortgage characteristics, including all of the financial characteristics of the contract (i.e. rate, loan size, house price, debt-ratio, risk-type), the lender identity (for the 12 largest lenders), some demographic characteristics (e.g. income, prior relationship with the bank, residential status, dwelling type), as well as the house location up to the forward-sortation area (FSA).[8] Table 9 in the Appendix lists all of the variables included in our data-set.

---

90 days, therefore a bank is committed to a consumer for those 90 days and this information is locked into a banks (and therefore all branches) database. Second, when mortgage insurance is provided the insurers' automated underwriting program flags whether an individual has qualified for insurance. Third, when individuals apply for a mortgage there is a credit check and all lenders can see that this has been done.

[7]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 per cent of condominium fees. Total debt service is defined as all payments for housing and other debt.

[8]This unit of aggregation is defined as the first three letters of the postal-code. It corresponds to about 4 to 6 census-tracts in urban areas (or between 10,000 and 40,000 households), or one small town in more rural areas. The median population size per FSA is about 16,000. There are over 1,300 FSA's in Canada, and over 850,000 postal codes.

We follow a quasi-experimental approach to study the impact of the merger on the distribution of negotiated interest rates. Over the next three subsections we describe the construction of our sample: (i) definition of the treatment and control group, (ii) sample selection criteria, and (iii) variable definitions and summary statistics.

2.1 Treatment and control groups

Our empirical strategy relies on the idea that when two bank networks merge, the pre-merger location of branches creates discrete changes in the structure of local markets. In particular, when two neighboring branches merge local competition is immediately reduced, since loan managers stop competing for the same borrowers. Importantly, since retail mergers are negotiated nationally, these changes can be viewed as exogenous relative to local market conditions, at least in the short run.

A retail merger can generate two types of changes. First, it can reduce the number of available options for consumers who had both brands in their neighborhood. Second, if the two merging firms were ex-ante different in terms of their product characteristics, the merger can change the characteristics of the available options for consumers who had only one of the two merging firms in their neighborhood. There is no effect for consumers who had neither brand in their neighborhood pre-merger.

Bank A is a national bank, and therefore present in nearly all local markets pre-merger. Trust B, on the other hand, is smaller, and isolated from A in only 2% of its markets. Given the extent of branch co-location we focus only on the first effect of the merger, namely the reduction in the number of available options. The *treatment group* is therefore defined as the set of consumers who had both lenders in their choice-set prior the merger, while the *control group* is the set of consumers who had only one or none of of the merging firm.[9]

To operationalize this definition, we need to formally define the choice-set of consumers. To do so, we exploit the fact that the pricing decision is decentralized, and therefore that consumers negotiate directly local branch managers. This allows us to define a consumer's choice-set as the set of lenders present in a neighborhood around the house's FSA, denoted by $\mathcal{N}_i$. In the empirical analysis we use a fixed radius of 5KM to define neighbrohood boundaries, and measure distances as the Euclidian distance from FSA centroids to the closest branch of each bank. We make use of a yearly panel obtained from the Financial Services Canada directory (Micromedia ProQuest). This data-set contains the location of active branches for all financial institutions.

Figure 2 motivates this definition, by illustrating the distribution of the minimum distance to chosen and competing lenders. The average distance to chosen lenders is much smaller than

---

[9]Ideally we would use as control group local markets in which none of the merging banks' branches were present. However, the fact that Bank A is present in nearly all neighborhoods limits the amount of variation to identify the model. In Section 4, we study the robustness of our results to this alternative definition of the control group.

(a) Distance to chosen bank (km) (b) Distance to competing banks (km)
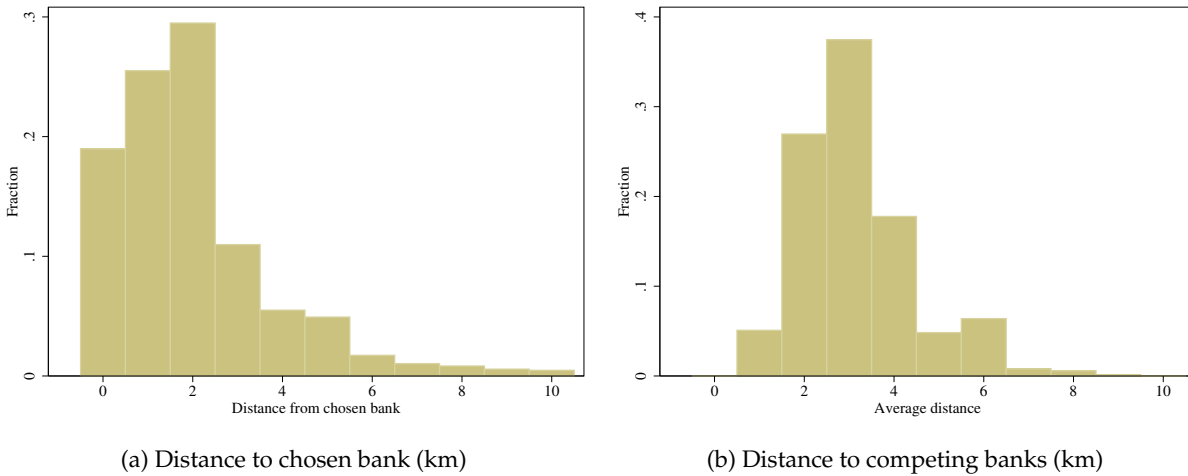
Figure 2: Distribution shortest distances between homes and banks

(a) Markets with A & B (b) Markets with A or B alone (or none)
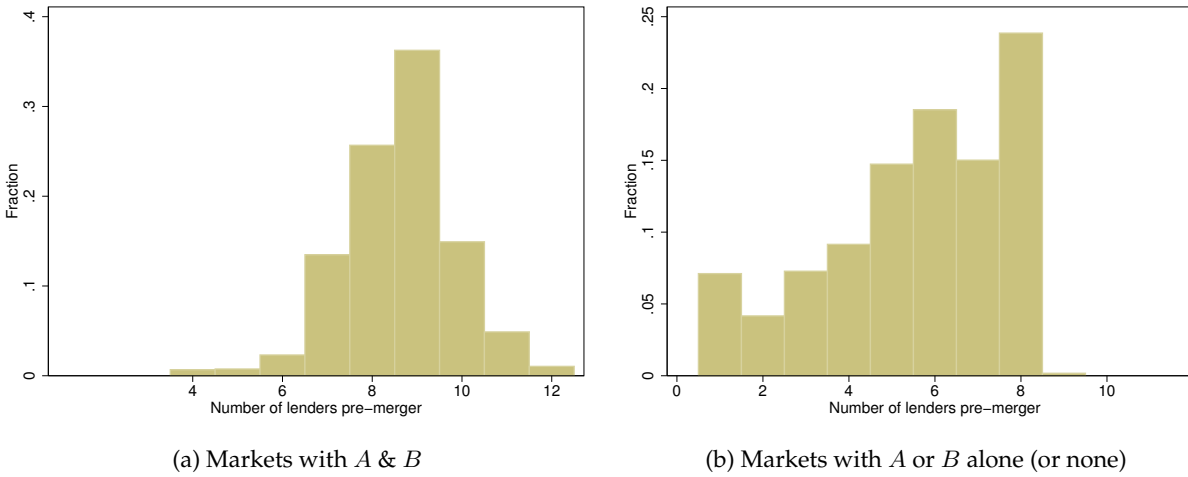
Figure 3: Distribution of the number lenders per local markets before the merger

the average distance to other financial institutions, and nearly 80% of transactions occur with a bank located within 2KM of consumers. This provides strong evidence that consumers most often deal with a bank that has a large presence in their region. Indeed, the 5KM threshold includes more than 90% of transactions. Notice that we measure distance as the Euclidian distance, which under-estimates the actual driving distance from the center of each FSA to a branch. A 5KM radius is therefore larger than the radii of most FSAs, and broadly corresponds to each consumer's municipality.[10]

---

[10]In Section 4.1 we evaluate the robustness of our results to alternative neighborhood sizes.

2.2 Sample selection

We construct our sample based on three criteria: (i) timing of the merger, (ii) homogeneity of contractual characteristics, and (iii) comparability of market structure.

Time period We first select contracts signed within a year of the merger date. We further eliminate contracts for which the closing date is less than 90 days after the merger, to avoid including in the post-merger period rates that were negotiated pre-merger. We use 90 days because in Canada lenders tend to guarantee price quotes for 90 days. This leads to a slightly uneven split of observations before and after the merger: 42% of the transactions take place post-merger.

Contract characteristics We select our sample by limiting heterogeneity in contractual characteristics (other than prices), and across consumer attributes. To do so, we select newly issued mortgages, excluding home-owners that are either refinancing or renewing their mortgage contract, and contracts with a 25 years amortization period and 5 year fixed-rate term. During our sample period, nearly all mortgage contracts were fixed rate, and over 85 per cent had a 5 year term. A 5 year fixed-rate mortgage contract must be renegotiated every five years, and banks impose substantial penalties to refinance before the end of the term.[11] This has been the standard contract offered by Canadian banks since the late 1960's. Similarly, almost all contracts have 25 year amortization periods.

Market structure Our choice-set definition creates a split between contracts signed in areas where both A and B were present pre-merger (i.e. treatment group), and areas in which only A or B, or neither were present (i.e. control group). Not all of these neighborhoods are directly comparable, and local markets with both A and B tend to be larger and have more lenders.

Figures 3a and 3b illustrate that the treatment and control groups significantly overlap only in medium-sized markets with five to eight lenders. Less than 1% of control markets have more than eight lenders, while less than 1% of treatment markets have fewer than five. Ideally, we would estimate the treatment effect separately for consumers facing similar choice-sets, but the distribution of lenders only allows us to do so for local markets with an intermediate level of concentration.[12]

---

[11]Unlike in the United States, refinancing is uncommon in Canada. This is largely because of the relatively short term of the mortgage contract (5 years versus 30), which makes the benefits from refinancing, that might come from lower interest rates relative to the large penalties imposed, less attractive compared to simply waiting to renewal. In addition refinancing in the U.S. happens when borrowers move. In Canada borrowers can port their mortgage, i.e. their mortgage can be transferred to the new home.

[12]This difference in the structure of local markets between treatment and control groups can bias our results, since merger effects are unlikely to be constant across markets of different sizes. Similarly, markets with a larger number of lenders tend to be more urban, and therefore possibly subjected to correlated aggregate trends. The fact that the compositions of two groups differ implies that these unobserved factors can be confounded with the causal effect of the merger.

Table 1: Summary statistics on consumer choice-sets prior to the merger

	Full sample		$5 \leq N \leq 8$	
	A or *B* or none	*A* & *B*	*A* or *B* or none	*A* & *B*
Number of Lenders	5.855	9.627	6.465	7.421
	(2.43)	(1.96)	(1.07)	(0.74)
Number of Branches	18.037	64.452	14.272	22.173
	(20.98)	(71.68)	(8.56)	(12.37)
Branch HHI	1.567	1.728	1.620	1.627
	(0.34)	(0.35)	(0.26)	(0.23)
Share of bank A	0.059	0.139	0.063	0.156
	(0.09)	(0.06)	(0.07)	(0.06)
Share of trust B	0.008	0.095	0.012	0.121
	(0.06)	(0.06)	(0.05)	(0.05)
Δ Branch HHI		0.179		0.230
		(0.22)		(0.25)

Each entry corresponds the sample average and standard-deviations (in parenthesis), calculated using the observation weights from the mortgage contract data-set. Local markets are defined as 5KM euclidian distance around each FSA centroid. Markets "*A* or *B*, or none" do not have *A* and *B* together, and markets "*A* & *B*" have both merging parties.

Given this additional constraint, we restrict our sample to borrowers facing choice-sets with five to eight lenders. In the estimation sample, therefore, households with only one or with neither of the two firms are under-represented (i.e. 37% versus 67%), but less so than in the full sample (i.e. 10% versus 90%). The final sample includes slightly more than eighteen thousand observations over approximately 400 different locations.

Table 1 describes the structure of choice-sets. The first two columns illustrate the distribution of the number of lenders and branches in the full sample, while the second two consider only overlapping markets. Excluding non-overlapping markets leads to comparable neighborhoods: the average number of branches and Herfindahl-Hirschman Index (HHI) are similar across the two groups, as is the number of lenders. The last three rows show that, absent other concurrent changes, concentration would increase significantly in markets with both *A* and *B*. Both institutions had a large presence in *A* & *B* markets, with a cumulative average market share of 28%. The merger alone therefore corresponds to an average increase of 0.23 points in the branch HHI.

2.3   Variable definitions and descriptive statistics

Our analysis focuses mostly on two outcome variables: discounts and margins. These two variables provide direct measures of the cross-sectional dispersion of rates, by controlling for intertemporal variation in aggregate interest rates. Transaction margins are measured as the interest rate paid by consumers less the swap-adjusted 5 year government of Canada bond rate measured

Table 2: Summary statistics on mortgage contracts and household characteristics

	Control/Before				Control/After			
	Mean	S.D.	$P(25)$	$P(75)$	Mean	S.D.	$P(25)$	$P(75)$
Margin	1.07	0.46	0.73	1.42	1.43	0.56	1.02	1.82
Residual margin	1.07	0.43	0.78	1.35	1.67	0.46	1.36	2.02
Zero discount - %	36.38	48.12			26.92	44.36		
Income (/1000)	61.95	25.00	43.70	74.77	62.84	24.49	45.37	75.29
House price (/1000)	121.11	55.49	82.18	145.30	118.25	52.16	81.18	143.74
Loan (/1000)	113.4	49.7	78.7	137.2	110.0	46.9	76.9	133.5
LTV - %	91.58	4.26	90.00	95.00	91.25	4.34	90.00	95.00
FICO $\geq$ 600 - %	67.23	46.95			64.27	47.93		
Status: Renter - %	68.30	46.55			70.05	45.82		
Status: Parents - %	5.73	23.24			6.64	24.90		
Switch - %	30.21	45.93			36.84	48.25		
Broker - %	21.91	41.37			30.05	45.86		
	Treatment/Before				Treatment/After			
	Mean	S.D.	$P(25)$	$P(75)$	Mean	S.D.	$P(25)$	$P(75)$
Margin	0.93	0.49	0.65	1.27	1.44	0.56	1.12	1.82
Residual margin	1.06	0.43	0.78	1.35	1.72	0.47	1.44	2.00
Zero discount - %	23.67	42.51			22.23	41.58		
Income (/1000)	69.33	26.48	50.71	82.23	70.99	26.49	52.46	84.11
House price (/1000)	162.93	63.38	116.77	201.83	161.07	64.58	114.84	200.72
Loan (/1000)	152.3	57.6	110.5	188.2	149.9	57.6	108.2	187.0
LTV - %	91.35	4.25	90.00	95.00	90.99	4.48	89.60	95.00
FICO $\geq$ 600 - %	62.40	48.44			62.56	48.40		
Status: Renter - %	68.28	46.54			71.15	45.31		
Status: Parents - %	8.31	27.61			9.32	29.08		
Switch -%	26.68	44.23			38.43	48.65		
Broker - %	15.66	36.34			27.73	44.77		

The sample size is 18,121 divided between the control and treatment group, pre- and post-merger with 62.8% of contracts in the treatment and 42.2% of the contracts observed post-merger. It includes a random sample of homogenous term and amortization contracts insured by CMHC or Genworth within one year of the merger. Margins and residual margins are defined in the text. $FICO \geq 600$ is an indicator variable equal to one if a consumers credit score is greater than 600. Renters and parents correspond to new home buyers exiting from renting and living with parents, respectively. Switcher is an indicator variable equal to one if consumers have no prior experience with the chosen financial institution. The sample is restricted to households with 5 to 8 lenders located within 5KM of their FSA centroid.

at the week of negotiation, and discounts are measured relative to the bank-specific posted rate at the week of negotiation. We use the discount variable to calculate the fraction of households receiving zero discounts, defined as an indicator variable equal to one if the observed discount is less than 10 basis points.

Note that we only observe the closing date on the house sale. For each contract, we identify the negotiation week by calculating the absolute difference between the transaction rate and the

posted rate for the weeks within 90 days of the closing date, and take the smallest value, and/or the closest in time in case of a tie. From this, we estimate that 44% of contracts are negotiated within 1 week of the closing week, and the remaining are approximately uniformly distributed between 1 week and 90 days. This definition also increases the fraction of consumers paying a rate equal or close to the posted-rate. Less than 10% of consumers pay a rate equal to the posted-rate valid at the closing date, while we estimate that roughly 25% pay the posted rate valid at the "negotiation" week. Many of the consumers that we categorize as receiving zero discounts actually benefit from a discount relative to the posted rate valid at the closing date, while others end up paying more than this posted rate.

We use the negotiated margins as our main outcome variable. To measure residual price dispersion, we proceed by decomposing the observed margins into a deterministic function of borrowers' characteristics, and an idiosyncratic component m_i:

$$\text{Margin}_i = \beta'\mathbf{X}_i + \mu_i^{\text{week}} + \underbrace{\theta_{G_i,T_i} + e_i}_{=m_i}, \tag{1}$$

where $\text{Margin}_i = \text{Rate}_i - \text{Bond}_i$ is the transaction margin, and θ_{G_i,T_i} is a group/period fixed-effect, G_i and T_i index the group (treatment or control) and time period (before or after the merger) for borrower i, μ_i^{week} is a closing-week fixed-effect, and $\mathbf{X}_i$ is a vector of control variables.[13] In the empirical analysis we mostly focus on the impact of the merger on the level and dispersion of *residual margins, m_i*.

Notice that we index each observation by i, with the understanding that i captures: (i) the individual borrower (or household), (ii) the time period of the contract, and (iii) the location of the purchased house. Although we observe each household only once, we observe most FSAs pre- and post-merger. This allows us in some specifications to control for location fixed-effects in (1).

Table 2 describes the sample, split across treatment/control groups and pre/post-merger. Overall the average margin is 117 basis points, and it has increased over time. In contrast, the fraction of consumers paying the posted rate (i.e. zero-discounts) remained fairly stable at around 26%. Since the remaining 74% of consumers negotiate discounts, most of the dispersion in margins is unexplained by variation in the level of interest rates. This is illustrated by the residual margin dispersion on the second line of each panel of Table 2. Slightly more than half of the margin variance in our data comes from cross-sectional residual dispersion alone. The magnitude of dispersion is large compared to other financial markets, considering the homogeneity of the contract terms and the presence of insurance. For instance, Hortaçsu and Syverson (2004) show that the cross-

---

[13]The exact set of control variables is: income, loan size, loan to income ratio, other debts, debts to income ratio, loan-to-value categorical variables, credit score categories (4), residential status category (4), switcher, and FSA-level census characteristics (i.e. house value, income, education, average age, and migration). To control for the non-random nature of missing household characteristics we interact a missing value dummy with: province indicators, treatment group, after merger, and bank indicator variables.

sectional standard-deviation among mutual fund transaction fees was equal to 60 bps in 2001. Mutual funds exhibit substantially more heterogeneity in observable characteristics and returns than the 5-year fixed rates mortgage contracts studied here.

Comparing the treatment and control groups pre-merger, we can see that approximately 23% of consumers in the treatment group and 36% of those in the control group received zero discounts. Margins were also higher in the control groups. These differences mostly reflect the fact that markets in the treatment group are more urban and feature more competition.

These differences can be seen by comparing income and loan size in the two groups. The average home-owner in the treatment group earns $69,330 and contracts a loan of $152,300, while in the control group the mean income and loan size are $61,950 and $113,400 respectively. However, loan-to-value ratios are similar across the two groups. The other characteristics are all quite similar across the two groups prior to the merger.

The effect of the merger on rates can already be seen in Table 2. Rates rise in both the treatment and control markets, but the increase is about 7 bps greater in the treatment. Similarly, the fraction of consumers paying the posted rate falls everywhere, but by less in the treatment.

3 Estimation and identification strategy

Our objective is to study the effect of the merger on the distribution of negotiated rates. Equation 1 characterizes a reduced-form pricing function. Using the language of the treatment effects literature, we think of the merger as an intervention (I_i) that changes residual margins negotiated between borrowers and lenders:

$$m_i = I_i m^I(u_i, T_i) + (1 - I_i) m^N(u_i, T_i), \tag{2}$$

where T_i is a before/after period indicator, $I_i = T_i \times G_i$ is a merger indicator variable, and $m_i^I(u_i, T_i)$ and $m_i^N(u_i, T_i)$ correspond to the post- and pre-merger residual margin functions, respectively. Equation 2 allows for the possibility that mergers have heterogeneous effects by expressing transaction prices as a non-separable function of the consumers' unobserved heterogeneity component u_i. This variable determines the size of negotiated discounts, and represents, for example, unobserved search effort or negotiation skills.

Since the *post-merger* prices are observed from the data, the goal of our empirical analysis is to identify the counter-factual *pre-merger* prices given by the function $m^N(u_i, T_i = 1)$ for "treated" consumers ($G_i = 1$). This allow us to measure the effect of the merger along the distribution of residual margins: $m^I(u_i, T_i) - m^N(u_i, T_i)$. In contrast, the merger evaluation literature has focused on estimating changes in average prices: $E\left(m^I(u_i, T_i) - m^N(u_i, T_i)\right)$. More formally, our empirical analysis focuses on estimating the effect of the merger along three dimensions:

1. Average prices: $\bar{\alpha} = E\left[m^I(u_i, T_i = 1) - m^N(u_i, T_i = 1) | G_i = 1\right]$,

2. Price distribution: $\alpha_q = m^I(u_i, T_i = 1) - m^N(u_i, T_i = 1)$, such that $\Pr(u < u_i | G_i = 1) = q$,

3. Price dispersion: $\Delta\text{Disp}(m) = \text{Disp}\left[m^I(u_i, T_i = 1) | G_i = 1\right] - \text{Disp}\left[m^N(u_i, T_i = 1) | G_i = 1\right]$,

where $\text{Disp}(x)$ is one of three measures of price dispersion: standard-deviation, coefficient of variation, or inter-quartile range.

To perform the estimation, we impose three assumptions on the residual margin function and the distribution of u_i. First, in order to compare treated and control markets, we assume that the function $m^N(u, T)$ is common across treated and control consumers (hence the importance of using comparable markets with intermediate levels of competition). Second, residual margins are monotonically increasing functions of u_i. Ceteris paribus, a high u_i consumer pays a high rate, and receives a small discount. Finally, the distribution of unobserved heterogeneity is assumed to be constant over time, but can differ across control and treated markets. This assumption implies that the effect of the merger could differ across the two groups, even though the residual margin function is common. We let $H_1(u)$ and $H_0(u)$ denote the distributions of u within treated and control areas, respectively.

All three assumptions are fairly standard in the difference-in-difference literature (e.g. Abadie (2005)), and are adapted to our context, following Athey and Imbens (2006), to reflect the fact that the errors are non-separable. Below we describe the three estimators used to characterize the average and distributional impact of the merger.

Linear difference-in-difference estimator We first use a linear difference-in-difference (DiD) approach to measure the average treatment effect (ATE) of the merger. In our context, the DiD model linearly approximates the residual margin function above, and estimates equations (1) and (2) jointly. Moreover, the linearity assumption allows us to control for location (FSA) fixed-effects. This leads to the following estimating equation:

$$\text{Margin}_i = \bar{\alpha}I_i + \beta'\mathbf{X}_i + \mu_i^{\text{week}} + \mu_i^{\text{fsa}} + e_i, \tag{3}$$

where $\mathbf{X}_i$ is a vector of control variables, and μ_i^{week} and μ_i^{fsa} are week and FSA fixed-effects.

An important concern is that the merger predominantly affected a small number of provinces, and therefore might be correlated with location-specific housing or financial variables affecting negotiated interest rates. For instance, house prices and the timing of purchases follow regional cycles and trends that could be confounded with the merger. Furthermore, local housing market trends may impact mortgage rates differently in treated and control markets. We consider two sets of control variables to address this question.

Our baseline specification includes controls that describe the financial and demographic characteristics of the contract, and the identity of lenders.[14] We use the vector $\mathbf{X}_i^{\text{base}}$ to denote these

---

[14]The exact set of control variables is listed in footnote (13).

controls. In an effort to further control for confounding factors we also estimate $\bar{\alpha}$ using a richer set of variables, labelled $\mathbf{X}_i^{\text{trend}}$, that control for heterogenous observable trends across our control and treatment groups. We do so by interacting every consumer and contract characteristics in $\mathbf{X}_i$ with T_i, our "post-merger" dummy variable. We also include in $\mathbf{X}_i^{\text{trend}}$ province-level linear and quadratic trends. We present the empirical results using both $\mathbf{X}_i^{\text{base}}$ and $\mathbf{X}_i^{\text{trend}}$ specifications.

In addition to estimating the ATE for margins, we study the impact on the probability of paying the posted rate. Using the same two specifications, we replace p_i by an indicator variable equal to one if the consumer pays a rate greater than or equal to the posted-rate plus 10 bps.

Matching estimator An alternative approach to control for differences between control and treatment groups is to estimate $\bar{\alpha}$ using a matching DiD estimator. Since we want to compare local markets along multiple dimensions, we implement a matching DiD estimator based on a propensity score function, and measure the degree of similarity between local markets using a Kernel smoothing function. In addition to the three assumptions imposed on u_i already mentioned, this estimator also imposes the restriction that treated and control neighborhoods have a significant overlap in terms of the probability of being affected by the merger. We therefore exclude neighborhoods in the top and bottom 2% of the propensity score distribution.

In order to control for borrower-specific observed covariates, we calculate the matching estimator using residual margins. This leads to a sequential estimator: (i) estimate m_i by OLS from equation 1, (ii) estimate the propensity score function via a Logit model (i.e. probability of having both A and B in a neighborhood), and (iii) estimate the ATE of the merger using the estimator of Heckman et al. (1997). The propensity score controls for time-invariant demographic characteristics and branch concentration of each neighborhood.[15]

This estimator relaxes the assumption that the residual margin function $m^N(u_i, T_i)$ is common across control and treated local markets. Instead, we are implicitly assuming that neighborhoods with similar propensity scores have homogenous pricing functions.

Change-in-change estimator We apply the change-in-change (CiC) estimator proposed by Athey and Imbens (2006) to study the distributional impact of the merger on negotiated margins. In addition to the three assumptions above, we also assume that the support of u_i in the treatment group overlaps with the support of u_i in the control group. We impose this restriction by dropping the non-overlaping residual margins across the two groups for each time period.

In practice, we also assume that the observed characteristics of the contracts $\mathbf{X}_i$ are independent of u_i. Like with the matching estimator, this allows us to perform the estimation using the

---

[15]The list of controls are demographic characteristics measured using 2001 census tables at the FSA level and branch location data and include: average house value and income to house value ratio, fraction of households that rent, average age, inter-province migration ratio, fraction of people with a university degree, number of dwellings, pre-merger branch-HHI, and number of lenders pre-merger.

residual margins distribution defined in equation 1, and to control for observable characteristics of the contracts in the parametric form suggested by equation 1.[16]

Under these assumptions, the change in the q^{th} percentile of the margin distribution in the control group identifies the change in the markup function that is strictly due to time. Using this logic, and the fact that the distributions are invertible (under the monotonicity assumption), we can recover the counter-factual distribution of prices in the treatment group for any point m_i in the common support:

$$F_{1,1}^c(m_i) = F_{1,0}\left(F_{0,0}^{-1}\left(F_{0,1}(m_i)\right)\right), \tag{4}$$

where the superscript c identifies the counter-factual distribution, and $F_{G,T}(m)$ is the CDF of margins in the sub-population (G, T). Intuitively, we obtain the counter-factual distribution by transforming the observed residual margin distribution in the treatment group at time 0 (i.e. $F_{1,0}(m)$) to mimic the change in the distribution observed in the control group. We then use the empirical counterparts of $F_{1,1}^c(m)$ and $F_{1,1}$ to calculate the impact of the merger on various moments of the margin distribution, such as the average, the standard-deviation and the inter-quartile range. We also recover an estimate of the effect of the merger for different percentiles q:

$$\hat{\alpha}(q) = \hat{F}_{1,1}^{-1}(q) - \hat{F}_{0,1}^{-1}\left(\hat{F}_{0,0}\left(\hat{F}_{1,0}^{-1}(q)\right)\right), \tag{5}$$

where $\hat{F}(\cdot)$ refers to the empirical distribution function (EDF), and $\hat{F}^{-1}(\cdot)$ is the inverse of the EDF. Since we can normalize the scale of u_i to be between zero and one, $\alpha(q)$ measures the estimated effect of the merger on a consumer of type $u_i = q$.

4   Results

4.1   Average impact of the merger

In Table 3 we present the average effect of the merger on margins and discounting, estimated by OLS, matching, and using the CiC estimator. We present results for both the "baseline" (including all the covariates in $\mathbf{X}_i^{\text{base}}$) and "trends" (including all the covariates in $\mathbf{X}_i^{\text{trend}}$) specifications.[17]

Columns (1) and (2) present the average treatment effect of the merger on margins in the baseline and trends cases, respectively. Adding the trend variables increases the point estimate from about 6 bps to 7 bps in the OLS specification. This increase is not uniform across the three methods of estimation. The matching estimator yields slightly higher estimates of the ATE, but the baseline

---

[16]This is the approach suggested by Athey and Imbens (2006) (pages 465-466). See Linton et al. (2005) for a related two-step procedure.

[17]In the appendix we report OLS estimates of the merger effect, along with the marginal effect of other covariates to provide the interested reader with some information on how these are priced. For a more detailed discussion of the pricing of these contracts see Allen et al. (2013).

Table 3: Average effect of the merger on margins and discounts

	Margin		Zero discount	
	Baseline	With trend	Baseline	With trend
Linear DiD (OLS)				
Merger ATE	0.0607^a	0.0719^a	0.0646^a	0.0477^a
	(0.0183)	(0.0242)	(0.0154)	(0.0188)
N	18,121	18,121	18,121	18,121
R^2	0.408	0.420	0.181	0.189
Matching DiD				
Merger ATE	0.0739^a	0.0692^b	0.0670^a	0.0538^b
	(0.0249)	(0.0273)	(0.0206)	(0.0218)
N	17,220	17,220	17,220	17,220
Change-in-change				
Merger ATE	0.057^a	0.0661^a		
	(0.015)	(0.021)		
N	18,103	18,103		

Standard-errors clustered at the FSA level are in parenthesis. Significance levels: ^{a} p<0.01, ^{b} p< 0.05, ^{c} p< 0.1. The dependent variable in columns (1) and (2) is the transaction rate minus bond rate, in columns (3) and (4) is an indicator variable for the transaction rate within 10 basis points of the posted rate, which we take to imply zero discount. All specifications include borrower characteristics and bank characteristics as well as FSA and week fixed effects. The trend specifications include province trends as well as the borrower covariates interacted with the "after-merger" dummy. The matching estimator is calculated using the propensity score with four (4) nearest neighbors. Standard errors and hypothesis tests are calculated by bootstrapping the original sample 1000 times. See Huynh et al. (2011) for analysis of the bootstrap performance in the context of the CIC estimator.

estimate is higher than the one with trends (i.e. 7.39 versus 6.92 bps). The average effect estimated with the CiC estimator is the lowest, between 5.7 and 6.6 bps.

In columns (3) and (4) we estimate the impact of the merger on the probability of not receiving a discount. On average, consumers are 5-6 percentage points less likely to receive a discount post-merger. This effect corresponds roughly to a 20% increase in the probability of not receiving a discount. Once again, the matching results mostly confirm the OLS estimates of the merger, suggesting that our choice of sample and control variables accurately correct for systematic differences between the treated and control neighborhoods. Notice that we do not report an estimate for the change-in-change estimator since it does not provide a point estimate for the discrete outcome case (only bounds).

Overall, all three methods produce remarkably similar estimates, ranging from 5.7 to 7.39 for margins, and from 4.77 to 6.7 percentage points for the probability of not receiving a discount. The estimates are also precisely estimated in most specifications; only the matching ATEs with trends have a p-value lower than 1%. We can therefore reject the null hypothesis that the merger did not change market power, at least for the subsample of neighborhoods with five to eight lenders pre-merger. This implies that residual dispersion is not solely driven by risk- or cost-based pricing.

Recall that this is the standard model assumed in the finance literature to explain the observed dispersion of lending rates (e.g. Edelberg (2006), and Einav et al. (2012)).

In terms of magnitude, the point estimates for the ATE on margins correspond to between 10%-15% of the observed standard-deviation in our sample. For an average loan size of $152,000, we estimate that the merger led to a $5.73 increase in monthly payments (evaluated at the baseline OLS estimate).

This relatively small price increase suggests on the one hand that the merger did not cause substantial harm to the average borrower. This is consistent with the notion that, featuring five to eight lenders, the markets were fairly competitive to begin with. In addition, several institutional features support this interpretation: contracts are homogeneous, rates are negotiable, and, due to loan securitization, for a given consumer costs are mostly common across lenders. Moreover, in the market that we study, lenders are fully protected against the risk of default by a government insurance program, which standardizes the lending conditions across financial institutions. These features allow informed consumers to gather multiple quotes, and obtain an interest rate that reflects the expected lending cost, even with a small number of competing lenders.

On the other hand, our regression methodology likely provides a lower bound evaluation of $\bar{\alpha}$. This is because the treatment effect approach to merger analysis assumes fixed market boundaries, and in our context that consumers living in the same neighborhood have a known and common choice set. The presence of heterogeneity in the sets of lenders considered by consumers translates into measurement error in the treatment variable, which would in general attenuate our estimate of the merger effect. Assuming that this measurement error is independent of the timing of the merger itself, our estimates represent a lower bound of the effect of the merger on consumers who did have both lenders in their choice-set.

Robustness analysis We analyse next the robustness of the average merger effect to alternative estimation methods, and sample selection. We present robustness results with respect to the choice of controls, the size of local neighborhoods, and different event windows. Due to space constraints, we focus on the average effect on transaction margins, but the robustness analysis for discounting yields similar results, and is available upon request.

In Table 4 we present the average effect of the merger estimated with alternative sets of control variables. Our baseline specification is reproduced in column (2) and the trend results are reproduced in column (3). Column (1) shows that failure to control for characteristics of the transaction does not lead to a substantial change in the ATE. In column (4) we introduce provincial fixed effects interacted with the post-merger dummy variable. Since a large fraction of observations are clustered in a few provinces this specification adds substantial strain to the data. However, we still estimate an ATE of 5 bps that is statistically significantly different from zero at the 5% confidence level. The point estimates remain fairly constant across all four specifications, which suggests that our conclusions are not sensitive to the choice of controls.

Table 4: Alternative linear difference-in-difference specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Merger	0.0623^a	0.0607^a	0.0719^a	0.0500^b	0.0564^a	0.0602
	(0.0188)	(0.0183)	(0.0242)	(0.0251)	(0.0193)	(0.0374)
Observations	18,121	18,121	18,121	18,121	15,770	13,216
R-squared	0.344	0.408	0.416	0.417	0.415	0.420
Location (FSA) FE	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓
Control group	Full	Full	Full	Full	A	None
X Controls		✓	✓	✓	✓	✓
D × X & Province trends			✓	✓		
D × Province FE				✓		
% Treated	0.628	0.628	0.628	0.628	0.722	0.862

The dependent variable is margins. Significance levels: ^{a} p<0.01, ^{b} p< 0.05, ^{c} p< 0.1. Standard errors are clustered at the location level (FSA).

In columns (5) and (6) we modify the definition of the control group. First we restrict the control group to markets in which only bank A is present, and then to markets where neither A nor B are present. The "only A" specification raises the fraction of treated observations to 72%, but leaves the point estimate largely unchanged (i.e. 5.64 versus 6.04 bps). In contrast, in the "no A or B" specification the point estimate is equal to 6.02 bps, but the standard-error nearly doubles. This decrease in precision is due to the fact that the fraction of observations belonging to the control group goes down from 43% to less than 14%. Despite this, the results from these two specifications suggest that the magnitude of the estimated average merger effect is not driven by changes occurring in the control markets, such as the re-branding of bank A post-merger.

Our second set of robustness tests considers the impact of different neighborhood sizes. Our main specification assumes that consumers shop within a 5KM radius around the centroid of their FSA's. Increasing this threshold tends to raise the number of FSA's directly affected by the merger, at the cost of including lenders that are not considered by the average consumer (falsely treated neighborhoods). Using a smaller threshold reduces the number of treated neighborhoods, and includes in the control group areas that are affected by the merger. Therefore, over-estimating or under-estimating the size of consumers' choice-sets exacerbates the measurement error problem in our treatment variable, and biases our results towards zero. The estimation results, reproduced in Table 5a, confirm this intuition. The effect of the merger is statistically different from zero for most neighborhoods, but is smaller and less precisely estimated for smaller and larger radii. This is particular true when we control for province-level trends in the larger neighborhood specifications, since the number of provinces present in both treated and control groups shrinks (i.e. it becomes harder to distinguish between the merger effect and other province-level trends).

Table 5: Robustness analysis results on the average merger effect

(a) Alternative neighborhood sizes

	(1) 3KM	(2) 4KM	(3) 5KM	(4) 6KM	(5) 7KM
Baseline	0.0405^b	0.0529^a	0.0607^a	0.0478^b	0.0304
	(0.0184)	(0.0198)	(0.0183)	(0.0198)	(0.0204)
With trends	0.0242	0.0496^c	0.0719^a	0.0414	0.0173
	(0.0213)	(0.0262)	(0.0242)	(0.0258)	(0.0266)
Observations	18,121	18,121	18,121	18,121	18,121
% Treated	0.479	0.589	0.628	0.706	0.728

(b) Alternative event windows and falsification tests

	(1) 6-Months	(2) 1-Year	(3) 18-Months	(4) + 6 Months	(5) - 6 Months
Baseline	0.0969^a	0.0607^a	0.0172	-0.0135	0.0166
	(0.0279)	(0.0183)	(0.0152)	(0.0254)	(0.0177)
With trends	0.0677^c	0.0719^a	0.0444^b	-0.0250	0.0100
	(0.0391)	(0.0242)	(0.0200)	(0.0343)	(0.0217)
Observations	8,040	18,121	27,978	17,296	17,148

Dependent variable is margins. Significance levels: ^{a} p<0.01, ^{b} p< 0.05, ^{c} p< 0.1. All specifications include week and location (FSA) fixed effects, these are included in X_i^{base}. Trend specifications also include the variables in X_i^{trend} Standard errors are clustered at the location level (FSA).

Finally, we consider different event windows and present the results of a falsification exercise in Table 5b. Recall that throughout the paper we use contracts signed at most one year before or after the official merger date. In columns (1) and (3) we compare our main specification with a shorter window of six months and a longer window of eighteen months. The results suggest that the effect of the merger was strongest within six months of the announcement, but diminished over time, possibly due to the closing of branches.

In columns (4) and (5) we consider two alternative merger dates: six months before and six months after the actual date. In both cases, we fail to find any evidence of rate increase, suggesting that our estimates are not confounded with the presence of unobserved factors affecting the treated neighborhoods within a year of the merger.

4.2 Distributional impact of the merger

In this section, we estimate the impact of the merger on the distribution of transaction rates. We use the change-in-change estimator to evaluate the counter-factual distribution of residual margins absent the merger in the treated neighborhoods, denoted by $\hat{F}_{1,1}^c(m)$.

Figure 4 plots the counter-factual (red) and observed (blue) empirical distribution functions
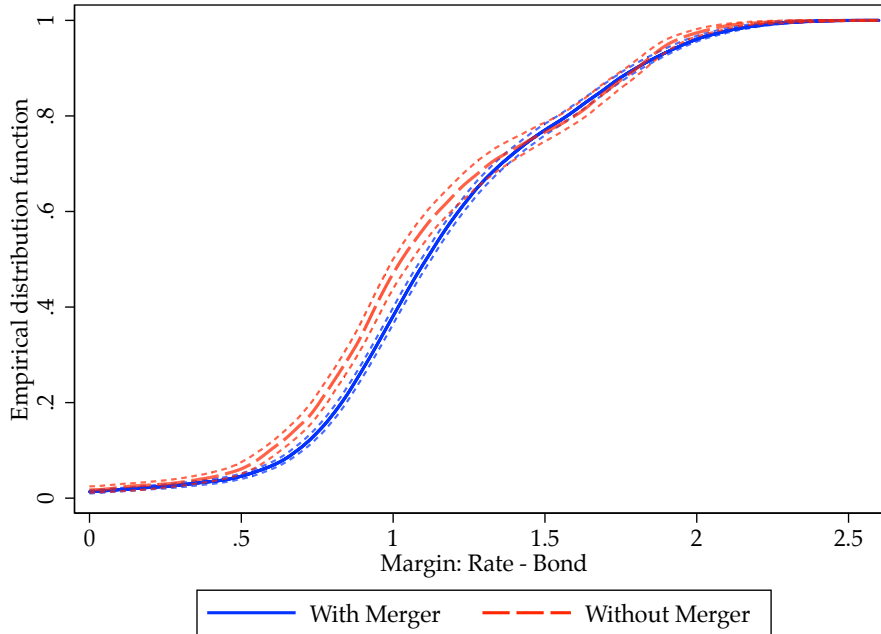
Figure 4: Empirical distribution of residual margins with and without the merger

(EDFs) for consumers with $G_i = 1$ and $T_i = 1$. The dashed lines represent bootstrapped 95% confidence intervals. For about three-quarters of the range, the counter-factual distribution is shifted to the left, confirming that most consumers pay higher interest rates as a result of the merger. The figure also reveals that the two distributions are nearly identical for consumers located in the top percentiles, suggesting that consumers paying higher rates were unaffected by the merger. Therefore, the increase in average rates is entirely due to the fact that consumers who, without the merger, were paying relatively low rates, experienced a significant decline in their ability to negotiate discounts. The same pattern is observed if we use a specification with trend, or if we vary the dimension of the choice-sets.[18]

Table 6 summarizes the effect of the merger on various moments of the margin distributions. At the bottom of the table, the Kolmogorov-Smirnov (KS) statistic is a measure of the difference between two EDFs, and tests the null hypothesis of equality of the rate distributions with and without the merger. As suggested by Figure 4, the test statistic easily rejects this null hypothesis. The counter-factual rate distribution without the merger therefore first-order stochastically dominates the post-merger rate distribution observed in the data.

The first seven rows of Table 6 calculate the difference between the two EDFs, and list the quantile treatment effect (QTE) of the merger. Consistent with the figure, we find that the effect of the merger ranges from around 9.1 bps for borrowers at the 10th quantile to zero at the 90th

---

[18]These results are available upon request.

quantile. We cannot reject the null hypothesis that the merger had no effect on negotiated rates for consumers beyond the 30th percentile. The QTE is similar once we control for trends (column 4).

This nonlinearity has two consequences. First the median borrower experienced a 2.7 bps larger rate increase than the average (i.e. 8.42 versus 5.7 bps). Second, the merger led to a compression of the overall distribution of rates, since they increased for borrowers in the lower part of the distribution, but were unchanged for those at the top.

Therefore, an important consequence of the merger is to reduce the amount of residual price dispersion. We estimate that the standard-deviation declined by 2.8 bps, while the interquartile range fell by 8.7 bps. Moreover, since average margins were higher post-merger, the change in the coefficient of variation shows a more pronounced decline in price dispersion: the merger led to a 4.6 percentage point decrease in the standard-deviation relative to the average margin. Similarly, the inter-quartile and inter-decile changes correspond roughly to 15% of the observed dispersion measured in the post-merger data.

This finding is in line with the predictions of Borenstein (1985) and Holmes (1989) regarding the relationship between competition and price dispersion. According to their setup, an increase in competition lowers the prices paid by the price-sensitive segment of the market, while leaving more or less unchanged the prices paid by consumers at the top of the distribution that tend to be loyal to one company. In contrast, it is at odds with the "textbook" price discrimination model, which suggests a negative relationship between dispersion and competition (see for example Stole (2007)). Note also search theoritical models of price dispersion have ambiguous predictions with respect to the impact of concentration on price dispersion, as discussed in Janssen and Moraga-González (2004) and Baye et al. (2006).

Our approach differs substantially from most of the existing empirical literature testing the relationship between market structure and price dispersion. The typical approach has been to compare, either with or without the use of instrumental variables, measures of residual dispersion across markets with different numbers of competitors or degrees of concentration. This has led to a wide array of empirical results, some of which have supported the traditional interpretation that market power enhances the ability of firms to price discriminate between consumers (see for example Gerardi and Shapiro (2009), Busse and Rysman (2005), Marvel (1976) and Barron et al. (2005)) and others of which found a negative or zero correlation between dispersion and concentration (see for instance Borenstein and Rose (1994) and Stavins (2001)). Still others find support for both interpretations, depending on market definition (Lewis (2008)).

In contrast, we measure the impact of losing one competitor on residual price dispersion in a comparable group of local markets. By exploiting the variation induced by the merger, our approach does not require the use of an instrumental variable, nor does it compare firms operating under different market structures.[19] We therefore think that our method is more credible and

---

[19]See Gerardi and Shapiro (2009) for a discussion of the potential pitfalls of relying on market-structure comparisons to estimate the relationship between concentration and price dispersion.

Table 6: Distributional effect of the merger on residual margins

	Baseline			With trend		
	Estimates	95% Conf. interval		Estimates	95% Conf. interval	
Quantile merger effects						
q_5	0.0766	0.0123	0.157	0.0782	0.0080	0.18
q_{10}	0.0914	0.0505	0.13	0.0844	0.0379	0.138
q_{25}	0.0762	0.0454	0.0994	0.0767	0.0397	0.119
q_{50}	0.0842	0.0433	0.114	0.0874	0.0446	0.128
q_{75}	-0.0008	-0.0621	0.0686	0.0185	-0.0473	0.0885
q_{90}	-0.0042	-0.0485	0.0383	0.0069	-0.0393	0.0591
q_{95}	0.056	0.0053	0.0949	0.0513	-0.0125	0.0973
Price dispersion						
Δ Std.Dev.	-0.0287	-0.0587	-0.0066	-0.0256	-0.056	-0.0034
Δ Coef. variation	-0.0462	-0.0766	-0.0211	-0.0455	-0.0792	-0.0184
$\Delta\ q_{75} - q_{25}$	-0.077	-0.131	-0.0099	-0.0582	-0.114	-0.0044
$\Delta\ q_{90} - q_{10}$	-0.0957	-0.153	-0.0457	-0.0776	-0.134	-0.0261
$H_0 : F_{1,1}^c(x) = F_{1,1}(x)$						
KS statistics	4.7^a			4.84^a		

Significance levels: [a] p<0.01, [b] p< 0.05, [c] p< 0.1. The dependent variable is residual margins. Confidence intervals were calculated by bootstrapping the sample 1,000 times. The KS statistic is a test of the equality of the observed and counter-factual empirical distribution functions.

could be used in other contexts as well to study the link between competition and dispersion.

5 Mergers when prices are negotiated

In this section, we interpret the empirical results from the previous section through the lense of a search and price negotiation model. We assume that the transaction rate is determined by deterministic functions of the bond-rate and borrower characteristics, and by a bilateral negotiation process generating residual margins:

$$\text{Rate}_i = \text{Bond}_i + \beta' \mathbf{X}_i + \mu_i^{\text{week}} + m_i, \tag{6}$$

where m_i is, as above, the residual margin. Our interest is in modelling the determination of m_i.

Our objective is to provide a theoretical framework that can replicate both qualitatively and quantitively the effects of the merger. In order to do so, we develop a model of search and price negotiation that is fully consistent with the identification assumptions of the CiC estimator described in Section 3. In particular, we use a model that yields a reduced-form pricing equation that is a monotonic function of a single-dimension unobserved heterogeneity term, u_i, measuring the marginal search cost of consumers. This allows us to invert the observed pre- and post-merger

margin distribution in order to recover a semi-parametric estimate of the search-cost distribution in the population of consumers directly affected by the merger. Finally, we use our estimates to perform counterfactual experiments in which we decompose the merger effect into a search-cost effect and an expected price effect, and in which we simulate the effect of lowering the search-cost of consumers.

5.1 Model

We consider a model with $n + 1$ lenders, in which negotiation takes place over three stages. In the first stage, consumers receive a take-it-or-leave-it offer m_i from one lender. In the second stage, if the initial offer is rejected, consumers put forth a search effort e_i at cost κ_i to gather additional quotes. The number of quotes that consumers obtain is stochastically determined by their search effort: with probability $s(e_i)$ they are randomly matched with n banks, and with probability $1 - s(e_i)$ they receive 2 quotes.[20] In the final stage, competition takes place, and consumers choose the lowest price offer. We characterize each stage sequentially, starting with the final competition game.

Competition stage As in Hall and Woodward (2012) and Allen et al. (2012), we model the competition stage as an English auction between at most n lenders. Alternatively, one can think of this stage as a Bertrand game between a random number of lenders, where the randomness is created by consumers' search effort. This modeling strategy differs from the standard search model often used in I.O. (e.g. Varian (1980)), which is based on a price-posting assumption. However, it is a common way of introducing negotiation in on-the-job-search environments (e.g. Postel-Vinay and Robin (2002)).

Lenders are ex-ante identical with a common marginal cost c, but face a mean-zero additive idiosyncratic cost shock ϵ_j that is privately observed after consumers are matched with lenders. The common component, c, should be interpreted as a lending cost on top of the bond rate and observable characteristics (for example processing fees). Note that, without loss of generality, we can also interpret ϵ_j as the idiosyncratic willingness to pay for each lender, or a combination of cost and value differences. The data do not allow us to differentiate between the two interpretations.

The game has a unique Bayesian-Nash equilibrium: banks are willing to offer up to their privately observed cost $c + \epsilon_j$, and the most efficient bank wins the contract by offering a rate equal to the second lowest cost, $c_{(2)} = c + \epsilon_{(2)}$, where $\epsilon_{(2)}$ is the second order statistic of $\{\epsilon_1, .., \epsilon_k\}$. Let $E(m^*|\tilde{n}) = c + E\left(\epsilon_{(2)}|\tilde{n}\right)$ denote the expected second-stage transaction price, where $\tilde{n}$ is

---

[20]We assume that the initial lender does not participate in the competition stage. We also assume that consumers get a minimum of two quotes in order to generate a finite price even for consumers with high search-costs. Alternatively we could have explained the reduced-form results by adding heterogeneous reservation values, for instance due to the posted rate or the value of renting a house. We chose not to include these features since they would require more than one source of unobserved heterogeneity, and therefore do not satisfy the identification assumptions of the reduced form model (i.e. the CiC estimator only accommodates a scaler unobserved heterogeneity).

the number of quotes generated (either two or n). Given our matching assumption, the gain from searching is summarized by the expected cost difference between obtaining two or n quotes: $\Delta(n) = E\left(\epsilon_{(2)}|2\right) - E\left(\epsilon_{(2)}|n\right) > 0$. This function in increasing in n.

Search effort At this stage, consumers incur a search cost κ_i to gather additional quotes. We use a simple search/matching technology to describe this process. Consumers invest in costly effort e to increase the probability $s(e)$ of receiving the maximum number of quotes n. Consumers putting forth zero effort are automatically matched with just two lenders, and we use a Pareto distribution to characterize the matching probability function: $s(e) = 1 - (1 + e)^{\frac{\gamma}{\gamma - 1}}$.

Consumers choose an optimal level of search effort to minimize the sum of the expected transaction price and the search cost. We assume that the search cost is an increasing function of effort e, and heterogeneous across consumers: $\kappa_i(e) = u_i \cdot (1 + e)$, where u_i is the marginal search cost of consumer i. That is, even consumers who choose not to exert any search effort must incur a cost for getting two quotes. This cost measures both the cost of acquiring information, and the haggling or time-cost of setting up the auction.

A consumer of type u_i facing n possible options chooses an optimal effort level that minimizes the net borrowing cost:

$$r(u_i, n) = \min_{e \geq 0} u_i \cdot (1 + e) + c + E\left(\epsilon_{(2)}|2\right) - s(e)\Delta(n). \tag{7}$$

In Appendix C we show that, in the interior, the solution to this problem gives rise to effort and matching probability functions that are decreasing in u_i and increasing in n, and that the optimal effort level exhibits a threshold property: consumers with marginal costs larger than a threshold $\bar{u}(n)$ invest zero effort.

Initial offer In the first stage, information is symmetric and the initial lender offers a rate, m_i, such that the consumer is indifferent between accepting this rate and searching for an extra quote: $m_i = r(u_i, n)$. In equilibrium, consumers do not search, but pay a price that reflects the value of their reserve price including the search cost, given by equation (7). This leads to the following reduced-form pricing equation:

$$m_i = c + \pi(u_i, n) = \begin{cases} c + \kappa(u_i, n) + E\left(\epsilon_{(2)}|2\right) - s(u_i, n)\Delta(n) & \text{if } u_i < \bar{u}(n), \\ c + u_i + E\left(\epsilon_{(2)}|2\right) & \text{if } u_i \geq \bar{u}(n). \end{cases} \tag{8}$$

where $\kappa(u_i, n)$ and $s(u_i, n)$ are the equilibrium search cost and matching probability functions, respectively, and $\pi(u_i, n)$ is the equilibrium first-stage margin.

In Appendix C, we show that the pricing function has the following properties: (i) it is monotonically increasing in u_i, (ii) it is weakly decreasing in n, and (iii) the marginal effects are increasing, $\partial^2 m_i / \partial u_i \partial n \geq 0$. As we discuss below, these properties have implications for the predicted

impact of mergers.

The first property implies that the pricing function satisfies the monotonicity assumption used to estimate the counter-factual price distribution (see Section 3). Therefore, as in Hortaçsu and Syverson (2004) and Hong and Shum (2006), the distribution of search costs can be recovered from the observed post-merger residual margin distribution:

$$F_{1,1}(m) = \Pr(c + \pi(u_i, n) < m) = H_1\left(\pi^{-1}(m - c, n)\right). \tag{9}$$

See Appendix C for derivation of the inverse function.

The second and third properties of the pricing function lead to the following proposition with respect to the impact of mergers in negotiated price markets.

Proposition 1. *The effect of a merger is heterogeneous in the consumer's type, u:*

1. *Positive price differential only among searchers:*

$$\pi(u_i, n - 1) - \pi(u_i, n) = \begin{cases} > 0 & \text{if } u_i < \bar{u}(n), \\ = 0 & \text{if } u_i \geq \bar{u}(n). \end{cases} \tag{10}$$

2. *Price differential is decreasing in u:*

$$\pi(u, n - 1) - \pi(u, n) \geq \pi(u', n - 1) - \pi(u', n), \text{ for all } u' \geq u. \tag{11}$$

The first element of the proposition follows directly from the fact that prices are weakly decreasing in n, and implies an increase in the average transaction prices. Moreover, if the median consumer exerts effort, it also provides a rationale for the result that the average merger effect is smaller than the effect for the median consumer. This is because the average is calculated over both searchers and non-searchers – two groups that are differentially affected by the merger.

The second element follows from the fact that the pricing function exhibits an increasing marginal effect of u_i with respect to n: $\partial^2 m_i / \partial u_i \partial n \geq 0$. Intuitively, relative to high u_i consumers, low search cost consumers are more likely to receive n quotes, and are therefore more affected by a reduction in $\Delta(n)$. This prediction is consistent with the observed non-linear effect of the merger: consumers at the top of the distribution were not impacted by the merger, while borrowers below the 70th percentile experienced rate increases between 7 and 9 bps.

The increasing difference property can also be used to show that mergers decrease residual price dispersion. In particular, we can re-write the second element of the proposition to establish

the following prediction with respect to the change in the inter-quartile range of residual margins:

$$
\begin{aligned}
\pi(u_{25}, n-1) - \pi(u_{25}, n) &\geq \pi(u_{75}, n-1) - \pi(u_{75}, n) \\
0 &\geq \text{IQR}(n-1) - \text{IQR}(n),
\end{aligned}
$$

where u_q denotes the q^{th} percentile of the marginal search-cost distribution, and IQR denotes the inter-quartile range. In other words, dispersion of prices in equilibrium is an *increasing* function of n. This prediction is therefore consistent with our empirical results, that the merger led to a contraction of the empirical margin distribution which translates into a significant reduction in residual price dispersions (i.e. 8% or 16% depending on the measure).

5.2 Model estimation

In this section we estimate the parameters of the bargaining model using the observed and counter-factual distributions estimated using the CiC estimator. The model parameters include: (i) the non-parametric distribution of consumer types $H_1(u)$ in the treated areas, and (ii) the preference parameter vector $\{c, \gamma, \sigma_\epsilon\}$, where σ_ϵ is the parameter determining the importance of firm lending cost heterogeneity. We assume the idiosyncratic component of lenders' costs is uniformly distributed between $-\sigma_\epsilon$ and σ_ϵ, which leads to expected second-stage prices of $c + E\left(\epsilon_{(2)}|\tilde{n}\right) = c + \sigma_\epsilon(3 - \tilde{n})/(\tilde{n} + 1)$ (where $\tilde{n}$ is 2 or n), and gains from searching of $\Delta(n) = 4\sigma_\epsilon(n - 2)/3(n + 1)$.

Notice also that the gain from searching is a function of the maximum number of quotes pre-merger (i.e. n). We set this number to 5, which corresponds to the lower-bound in the actual number of lenders in our estimated sample. While this number is somewhat arbitrary, it roughly corresponds to the maximum number of quotes consumers report in surveys of shopping behavior (Clayton Research (2005)).

To understand the scale of the model parameters, recall that we use the residual margin, defined in equation 1, as the dependent variable to estimate $\hat{F}_{1,1}^c$. Let m_i^I and m_i^N denote the observed (with merger) and counter-factual (without merger) transaction margins for consumer i. The lending cost parameter c therefore measures the average lending cost above the bond-rate for consumers with representative characteristics.

We consider a nested fixed-point estimator that minimizes the distance between the estimated counterfactual margins and the model predicted margins, conditional on the restriction that the search-cost distribution is consistent with the observed margin distribution. In practice we impose these restrictions at S random percentiles. This leads to the following non-linear least-square (NLS) estimator:

$$
\min_{\theta = \{c, \gamma, \sigma_\epsilon\}} \quad \sum_{i=1}^{S} \left(c + \pi(u_{q_i}, n|\theta) - m_i^N\right)^2 \quad \text{s.t.} \quad u_{q_i} = \pi^{-1}(m_i^I - c, n-1|\theta), \tag{12}
$$

27

Table 7: Model estimation results

Panel A: Estimation results					
	c	γ	σ_ϵ	n	NLS
Model parameters	0.726	0.585	0.858	5	0.0339
	(0.037)	(0.11)	(0.14)	(0)	(0.021)

Panel B: Simulation results					
	Marginal	Profit margins		Merger effect $\alpha(u_i)$	
	search-cost u_i	Without merger	With merger	Predicted	Estimated
Average	0.259	0.364	0.422	0.0577	0.0585
	(0.065)	(0.042)	(0.034)	(0.014)	(0.015)
Std-deviation	0.221	0.336	0.308	0.028	0.0319
	(0.018)	(0.007)	(0.004)	(0.0071)	(0.0073)
25^{th} Percentile	0.0911	0.099	0.179	0.0803	0.0664
	(0.05)	(0.045)	(0.033)	(0.014)	(0.015)
Median	0.199	0.322	0.382	0.0606	0.0841
	(0.073)	(0.042)	(0.034)	(0.014)	(0.019)
75^{th} Percentile	0.409	0.641	0.673	0.0324	0.0235
	(0.09)	(0.041)	(0.037)	(0.018)	(0.035)

Bootstrap standard-errors are in parenthesis. Number of bootstrap replications = 1000. Number of simulated prices used in the estimation = 100. The top and bottom 5% are eliminated from the estimated price distributions. Parameter definitions: c = lending cost, γ = matching probability parameter, σ_ϵ = cost heterogeneity, n = maximum number of quotes pre-merger (fixed).

where $q_i \sim U[0,1]$ denotes a simulated percentile, and $m_i^I = \hat{F}_{1,1}^{-1}(q_i)$ and $m_i^N = \hat{F}_{1,1}^{c-1}(q_i)$ are the inverses of the post-(intervention) and pre-merger margin distributions, respectively.

Conditional on the distribution $H_1(u)$, the model parameters are identified from the residual margins at different percentiles of u_i. This is because we invert point-by-point the observed post-merger margin distribution. The objective function is therefore equivalent to minimizing the distance between the predicted and observed quantile treatment effect of the merger: $c + \pi(u_{q_i}, n|\theta) - m_i^N \equiv \alpha(u_{q_i}|\theta) - \hat{\alpha}_{q_i}$.

The distribution of search-costs is non-parametrically identified up to an additional scale restriction. This is because the minimum search cost is not separately identified from the intercept c. We therefore impose a lower bound on the search cost distribution, $u_0 = 0$. This is equivalent to assuming that consumers receiving the largest discounts in our sample would have obtained the maximum number of quotes with probability one.[21]

Estimation results and goodness of fit Table 7 presents the model estimation results. Panel A

---

[21]We implement this restriction by adding a penalty to the NLS objective function such that the lowest simulated percentile $u_0 = \min_i\{u_i\}$ is close to zero.

shows the estimated parameters and the value of the NLS objective function. The standard errors are calculated using the bootstrapped price distributions estimated in Section 4.2, holding fixed the S simulated percentiles. Panel B summarizes the model predictions at the estimated parameters (columns 1-3), and compares the predicted and estimated quantile treatment effects of the merger (columns 4 and 5).

The average transaction margin can be calculated as the sum of the estimated average profit, 42 bps, and the estimated lending cost $\hat{c} = 0.726$, which can be interpreted as the average lending cost over the 5-year bond-rate. Scaled up by the bond-rate observed over our sample period, as in equation 6, this leads to predicted markups ranging between 5% and 10%.

The second and third columns of panel B show that profit margins vary widely across consumers. Because banks are able to perfectly discriminate, profit margins are less than 18 bps below the 25^{th} percentile of the post-merger distribution, and more than 67 bps above the 75^{th}. Undoing the merger would have led to a marked reduction in profit margins, but mostly among consumers with low search-costs (e.g. 8 bps at the 25^{th} percentile and only 3.2 bps at the 75^{th}). Similarly, consistent with our empirical results, the merger significantly reduced the dispersion of profit margins across borrowers. Our simulation results show that the coefficient of variation went from 92% before the merger, to 72.8% after (i.e. 0.336/0.364 versus 0.308/0.422).[22]

The estimated parameter σ_ϵ determines the expected gain from searching. At the estimated value, the gain from obtaining five quotes instead of two corresponds to $\Delta(n) = 57.2$ bps, which leads to an upper bound on the effect of mergers equal to $\Delta(n) - \Delta(n-1) = 11.44$ bps.[23] Taken literally, this parameter also implies important idiosyncratic cost differences across lenders: the range of ϵ_i corresponds to $2 \times 0.858 = 1.716$ bps. However, as we discussed above, heterogeneity across lenders in the competition stage can arise from differentiation and cost heterogeneity without changing the model formulation. Also, this parameter is not used measure the importance of residual dispersion. In other words, because information is symmetric in the first stage, consumers do not reject the initial quote, and the estimated parameter σ_ϵ is not restricted to match the observed distribution of prices. Therefore, one should not conclude that the model predicts substantial cost differences between lenders.

The matching probability parameter γ determines the marginal effect of search effort on the probability of receiving the maximum number of quotes. Since the estimated marginal cost of effort u_i varies greatly across consumers (standard-deviation is 0.221), this marginal product is highly dispersed. Evaluated at the average cost 0.259, a 1 percent increase in effort leads to a 1.8 percent increase in the probability of receiving 5 quotes. In comparison, the search effort elasticity is close to one at the 90^{th} percentile of the search cost distribution, and above 6 at the 10^{th}.

---

[22]It should be noted that at margins at the very bottom of the price distribution can actually be negative. We interpret this to mean that financial institutions make lending decisions based on average and not individual profits and that high-search cost borrowers subsidize low-search cost borrowers. According to our estimates less than 10% of borrowers have negative margins pre merger, and almost none do post merger.

[23]This is calculated by evaluating the change in $\Delta(n)$: $\Delta(n) - \Delta(n-1) = \sigma_\epsilon \times \left[4(4-2)/(3(4+1)) - 4(5-2)/(3(5+1)) \right]$.

(a) Marginal search cost distribution   (b) Post-merger search cost distribution (levels)
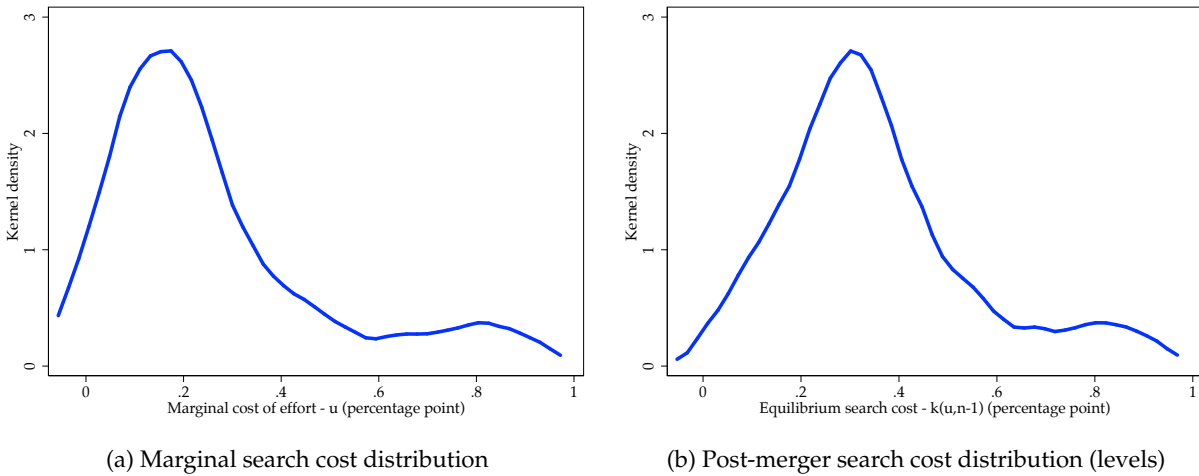
Figure 5: Estimated search cost distributions

The goodness of fit of the model can be studied by comparing the predicted and estimated effect of the merger (columns 4 and 5). As we discussed in Section 4.2, the data suggest a fairly flat merger effect among consumers in the middle and low percentiles of the price distribution. In contrast, the model predicts a monotonically decreasing merger effect, which tends to over estimate its effect at the bottom of the distribution, while under-estimating it at the median.

These differences arise because the estimated effect is fairly noisy at the top of the distribution potentially due to outliers, as illustrated by the quantile treatment effect beyond the 90% percentile (see Table 6). In contrast the model predict a strictly monotonically decreasing effect: If the merger has a zero effect at percentile q, the effect must also be zero for all percentiles $q' \geq q$. As a result, the NLS estimator best fits the counter-factual price distribution by predicting relatively few consumers investing zero search effort (i.e. $u_i > \bar{u}$), which leads to a smooth decline in the merger effect as a function of u_i.

Finally, Figures 5a and 5b plot the density of the marginal search cost and search cost level post-merger, respectively. As the table suggested, the distribution of consumers' marginal cost of effort u_i is highly dispersed, and skewed to the right. For the modal consumer the marginal cost corresponds to about 17 bps, while the average is equal to 25.9 bps.

The cost that consumers incur should they search is obtained by calculating the optimal level of effort for each type before and after the merger. Overall, that search cost levels are fairly large for most consumers, and highly dispersed. Post-merger, the average cost is equal to 35 bps, with a standard deviation of nearly 20 bps. For the average loan size in our data set, a 35 bps mortgage rate increase leads to an increase in monthly payments of $60, or $4,789.61 over five years.

To put this number into context, in Allen et al. (2013) we estimate that the marginal effect of brokers on the average transaction rate, after correcting for endogenous selection, was 40 bps.

Since brokers also provide slightly different services than multi-product lenders (i.e. lower value given they are a monoline business), it is reasonable to conclude that 40 bps is an upper bound on the average search cost in the population.

Merger effect decomposition Recently, researchers and policy makers have sought empirical evidence on the direction of price changes following *approved* mergers so as to inform antitrust authorities of the potential impact of *prospective* mergers (see Ashenfelter et al. (2009)). In posted-price markets without search frictions, reduced-form estimates provide a direct measure of the change in firms' market power caused by a decrease in competition, and can be used to predict the impact of counter-factual mergers in markets with similar structure. This is not the case in markets where consumers and firms individually bargain over prices.

In price-negotiation environments the reduced-form effect of mergers on prices depends not only on the change in market power, but also on the willingness and ability of consumers to search and haggle. If, as we have argued, consumers' abilities to negotiate discounts vary and are heterogenous in unobserved dimensions across markets, the true effect of counter-factual mergers will differ, sometimes substantially, from the effects predicted using only retrospective analysis.

We illustrate this point in two ways. First, we show that the net effect of mergers on the distribution of prices has two components: an endogenous adjustment of consumers' search effort, and a pure expected price change. Second, we simulate a counter-factual merger in an environment with lower search costs to highlight the impact of the search cost distribution on the predicted merger effects and price dispersion.

Within our search and negotiation framework, the distribution of search costs interacts with the degree of market power by determining the search effort and matching probability of consumers, as well as the price they expect to pay if negotiation fails. In particular, following a reduction in the number of competitors, the expected competition-stage price increases, while the optimal search effort decreases, leading to a reduction in the search cost and matching probability (i.e. $e(u_i, n)$ is increasing in n). The effect of a merger on the bargaining leverage of consumers is therefore the result of two opposite forces: (i) a reduction in search cost post-merger, and (ii) an increase in the expected competition-stage price. Formally, we can decompose the reduced-form merger effect that we estimated in Section 4 as follows:

$$
\begin{aligned}
\alpha(u_i) \;=\;& m(u_i|n-1) - m(u_i|n) = [\kappa(u_i, n-1) + E(m^*|u_i, n-1)] - [\kappa(u_i, n) + E(m^*|u_i, n)] \\
=\;& \underbrace{u_i \times [e(u_i, n-1) - e(u_i, n)]}_{\text{Search cost effect}<0} + \underbrace{\underbrace{[s(u_i, n) - s(u_i, n-1)]\,\Delta(n)}_{\text{Sorting effect}>0} + \underbrace{[\Delta(n) - \Delta(n-1)]\,s(n-1)}_{\text{Market power effect}>0}}_{\text{Expected competition-stage price effect}>0}.
\end{aligned}
$$

The last two terms are positive and measure an increase in the average price that consumers expect to pay if negotiation fails.

31

Table 8: Counterfactual and decomposition analysis

	Estimated search cost				Counterfactual (low) search cost		$\Delta\alpha$ (5)-(3)	Δm (6)-(4)
	$\Delta\kappa$	$\Delta E(m^*)$	$\alpha(u)$	$m(u\|n)$	$\alpha(u')$	$m(u'\|n)$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Distribution								
$u_{5\%}$	-0.71	10.93	10.21	57.81	10.62	53.20	0.41	-4.60
$u_{25\%}$	-1.99	10.02	8.03	82.48	9.17	69.65	1.13	-12.83
$u_{50\%}$	-3.14	9.21	6.06	104.74	7.85	84.49	1.79	-20.25
$u_{75\%}$	-4.79	8.03	3.24	136.66	5.97	105.76	2.73	-30.89
$u_{95\%}$	-0.15	0.15	0.00	181.42	3.33	135.60	3.33	-45.82
Average	-2.82	8.59	5.77	108.97	7.60	87.31	1.84	-21.67
Coefficient of variation	-50.82	29.46	48.59	30.81	26.01	25.63	-22.58	-5.19

The first four columns present the simulation results using the estimated search cost distribution. $\Delta\kappa$ is the change in search costs in going from 5 to 4 lenders. $\Delta E(m^*)$ is the change in the expected competitive transaction margin. Columns (5) and (6) present the counter-factual net merger effect and pre-merger margins for a low search-cost distribution (i.e. $u' = \hat{u}/2$). Columns (7) and (8) show the change in the merger effect and price levels from a reduction in the marginal search-cost. All results are presented in basis points.

As discussed above, given our estimate of σ_ϵ, the pure market power effect corresponds to an 11.44 bps increase in the competition-stage price for consumers expecting to receiving the maximum number of quotes. However, the adjustment in the matching probability and search cost vary greatly across consumers with different values of u_i. In Table 8 we illustrate the result of this decomposition for consumers located at five percentiles of the marginal search cost distribution (i.e. 5^{th}, 25^{th}, 50^{th}, 75^{th} and 95^{th}), as well as the overall average effects in the population.

In the first two columns, we estimate that the average effect of going from five to four lenders of 5.77 bps is the sum of a change in the expected competition-stage price of 8.59, and a decrease in search cost of 2.82 bps. Therefore, on average, the effect of the merger on the expected competition-stage price is roughly fifty percent higher than the reduced-form estimate (i.e. $E\left(\Delta\kappa(u_i)/\alpha(u_i)\right) = 2.82/5.77 = 48\%$), since the merger also decreases the search costs of consumers.

This decomposition is highly heterogeneous, since consumers with high marginal search cost shrink their levels of effort by a larger margin than consumers at the bottom of the distribution. The same price change ratio is near 100% for consumers at the 75^{th} percentile of the search cost distribution, and less than 18% for consumers below the 5^{th} percentile. Notice also the level changes in these two components are not monotonically increasing in u_i, since consumers at the very top of the price distribution react to the merger by reducing their search effort completely, which leads to a near zero net merger effect as illustrated by the consumers located at and above the 95^{th} percentile of the distribution.

In columns (5) and (6) we simulate the equilibrium price distribution and merger effects for an alternative lower search cost distribution. Such a decrease might come about if, for example, the

development of online mortgage shopping platforms lowers the cost of search for all consumers. We simulate the effect of a decrease by half for every u. Columns (7) and (8) calculate the difference between the estimated pre-merger price levels and merger effects, and the counter-factual ones.

The simulated change in the pre-merger price distribution (column (8)) shows that a 50% reduction in search cost would lead to large reductions in the level and dispersion of mortgage rates. We estimate that on average negotiated margins would fall nearly 20%, while the dispersion of prices would fall by 15%. These changes mostly come from the fact that consumers at the top of the price distribution are now able to negotiate substantially lower rates; more than 45 bps at the very top.

While this suggests that consumers benefit greatly from a reduction in search cost, we also estimate that the price effect of mergers in low search-cost environments would be much more important. On average, the effect of the merger is 32% larger in a low-search cost market. The effect of the merger is also much more homogenous across consumers, as illustrated by the 50% reduction in the coefficient of variation of the merger effect (last row). Once again, this reduction reflects the fact that consumers are the top of the distribution are now much more affected by the merger, since they are expecting to exert more search effort.

Overall, our results suggest that regulators evaluating mergers in markets with negotiated prices should worry, not only, about the pure market-power effect of mergers, measured here by the number of lenders and importance of heterogeneity (i.e. σ_ϵ), but also about the distribution of search frictions. In an environment with high search costs, firms enjoy large profits from being able to discriminate across consumers, despite the fact that mergers have little impact on prices. The small effects of mergers in this context hide the importance of market power. In contrast, the effect of losing a lending option can be important in markets featuring sophisticated consumers, which therefore increases the importance of the antitrust authority's evaluation of the merger.

6 Conclusion

In contrast to most of the literature studying the effects of horizontal mergers that focuses on posted prices or average transaction prices, we take advantage of individual transactions to document important heterogeneity in the reactions of firms and consumers to a merger. Our empirical analysis exploits observed differences in the choice sets of consumers and their financial characteristics to estimate heterogenous treatment effects across unobserved consumer types.

We find that the average effect of the merger yields a statistically significant increase in interest rates. This finding, however, masks important heterogeneity. Some borrowers pay significantly higher rates following the merger, while others are barely affected. The evidence we present suggests that much of the heterogeneity in rates can be explained by differences in search costs and negotiation ability. Borrowers at the top of the price distribution, those with high search

costs/bargaining abilities, are not affected by the merger, while those lower in the price distribution are affected. Together these results imply that price dispersion falls as a result of the merger.

Our findings on the heterogeneous effect of the merger imply that competition appears to have no impact on rates for consumers who are the most adversely affected by price discrimination; that is borrowers at the top part of residual price distribution. These results have important implications for the design of mortgage-market policies. If the objective is to support vulnerable borrowers (those paying the highest rates), policies designed to increase competition, or to prevent increases in concentration, such as restrictions on merger activity or prevention of bank failure may be ineffectual. Instead, what would be required would be policies designed to help borrowers search for and negotiate better terms. For example, policies that improve the financial literacy of borrowers may help them in their negotiations. Bertrand and Morse (2011), for instance, found using a randomized experiment at payday lenders that increasing information available to consumers led to a reduction in the amount people borrowed. In the housing market, Gerardi et al. (2010) argue that financial illiteracy played an important role in the rate of foreclosures in the U.S. housing crisis.

In Canada, these borrowers could be provided with more information on the distribution of discounts, or be informed as to the benefits of using brokers. Since brokers have fiduciary duties towards borrowers, their use helps borrowers search for and negotiate better terms. This is in contrast to the U.S. where, except in California, brokers do not have fiduciary duties. As a result, and as documented by Hall and Woodward (2012) there is considerable confusion surrounding the mortgage process generally, and the role of brokers in particular. Although there has been some discussion about the possibility of assigning fiduciary duties to brokers in the U.S., language initially appearing in the Dodd-Frank Act that would have done so was ultimately removed.

Another potential avenue for lowering search costs is through the use of the internet. The development of the internet and other technological improvements may lower the costs of gathering information and of getting approval for particular rates. This could result in a shift in the distribution, similar to what is described in Hortaçsu and Syverson (2004). We would expect that this would increase bargaining power for all borrowers, but especially those with greater search costs, resulting in a higher degree of competition.

References

Abadie, A. (2005). Semiparametric difference-in-differences estimators. *Review of Economic Studies 72*, 1–19.

Allen, J., R. Clark, and J. F. Houde (2012). Price negotiation in differentiated product markets: The case of insured mortgages. Mimeo, University of Pennsylvania.

Allen, J., R. Clark, and J. F. Houde (2013). Price dispersion in mortgage markets. Forthcoming *Journal of Industrial Economics*.

Ashenfelter, O., D. Hosken, and M. Weinberg (2009). Generating evidence to guide merger enforcement. CEPS working paper No. 183.

Athey, S. and G. Imbens (2006). Identification and inference in nonlinear difference-in-difference models. *Econometrica 74*, 431–497.

Barron, J., B. Taylor, and J. Umbeck (2005). Number of sellers, average prices, and price dispersion. *International Journal of Industrial Organization 22*, 1041–1066.

Baye, M., J. Morgan, and P. Scholten (2006). *Information, Search, and Price Dispersion*. Handbook on Economics and Information Systems.

Berger, A., R. Demsetz, and P. Strahan (1999). The consolidation of the financial services industry: Causes, consequences, and implications for the future. *Journal of Banking & Finance 23*, 135–194.

Bertrand, A. and A. Morse (2011). Information disclosure, cognitive biases, and payday borrowing. *Journal of Finance 66*, 1865–1894.

Borenstein, S. (1985, Autumn). Price discrimination in free-entry markets. *The Rand Journal of Economics 16*(3), 380–397.

Borenstein, S. and N. Rose (1994). Competition and price dispersion in the u.s. airline industry. *Journal of Political Economy 102*(4), 653–683.

Busse, M. and M. Rysman (2005). Competition and price discrimination in yellow pages advertising. *Rand Journal of Economics 36*, 378–390.

Busse, M., J. Silva-Risso, and F. Zettelmeyer (2006). $1,000 cash back: The pass-through of auto manufacturer promotions. *The American Economic Review 96*(4).

Clayton Research (2002-2005). Residential mortgage survey. Survey.

Competition Bureau (2003). The merger enforcement guidelines as applied to a bank merger. Technical report.

Crawford, G. and A. Yurukoglu (2012). The welfare effects of bundling in multichannel television market. *American Economic Review 102*, 643–685.

Dafny, L., M. Duggan, and S. Ramanarayanan (2011). Paying a premium on your premium? consolidation in the health insurance industry. forthcoming *American Economic Review*.

Dunning, W. (2010). Annual state of the residential mortgage market in Canada.

Edelberg, W. (2006). Risk-based pricing of interest rates for consumer loans. *Journal of Monetary Economics 53*, 2283–2298.

Einav, L., M. Jenkins, and J. Levin (2012). Contract pricing in consumer credit markets. *Econometrica 80*(4), 1387–1432.

Focarelli, D. and F. Panetta (2003). Are mergers beneficial to consumers? Evidence from the market for bank deposits. *American Economic Review 93*, 1152–1172.

Gerardi, K., L. Goette, and S. Meier (2010). Financial literacy and subprime mortgage delinquency: Evidence from a survey matched to administrative data. mimeo.

Gerardi, K. and A. Shapiro (2009). Does competition reduce price dispersion? *Journal of Political Economy 117*(1), 1–37.

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

Gowrisankaran, G., A. Nevo, and R. Town (2013). Mergers when prices are negotiated: Evidence from the hospital industry. working paper, Northwestern University.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review 103*, 145–177.

Hall, R. E. and S. Woodward (2012). Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. *American Economic Review 102*, 3249–3276.

Hastings, J. (2004). Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in southern california. *American Economic Review 94*, 317–328.

Heckman, J., H. Ichimura, and P. Todd (1997). Matching as an econometric evaluations estimator: Evidence from evaluating a job training programme. *Review of Economic Studies 64*, 605–654.

Hendel, I., A. Nevo, and F. Ortalo-Magne (2009, December). The relative performance of real estate marketing platforms: Mls versus fsbomadison.com. *American Economic Review 99*(5), 1878–1898.

Holmes, T. J. (1989). The effects of third-degree price discrimination in oligopoly. *American Economic Review 79*, 244–250.

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

Hortaçsu, A. and C. Syverson (2004). 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.

Hortacsu, A. and C. Syverson (2007). Cementing relationships: Vertical integration, foreclosure, productivity, and prices. *Journal of Political Economy 115*, 250–301.

Houde, J.-F. (2012, August). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review 102*(5), 2147–2182.

Huynh, K. P., D. T. Jacho-Chávez, and M. C. Voia (2011). Nonlinear difference-in-difference treatment effect estimation: A distributional analysis. *In Chapter 9, pages 251-272, Advances in Econometrics: Missing-Data Methods 27A*.

Janssen, M. and J. Moraga-González (2004). Strategic pricing, consumer search and the number of firms. *Review of Economic Studies 71*, 1089–1118.

KPMG (2008). Canada mortgage bonds program evaluation. Audit and Evaluation Services - Final Report.

Langer, A. (2012, August). Demographic preferences and price discrimination in new vehicle sales. working paper, University of Arizona.

Lee, J. and J. Hogarth (2000). Consumer information search for home mortgages: Who, what, how much, and what else? *Financial Services Review 9*, 277–293.

Lewis, M. (2008). Price dispersion and competition with differentiated sellers. *Journal of Industrial Economics 56*(3), 654–678.

Lewis, M. and K. Pflum (2013). Diagnosing hospital system bargaining power in managed care networks. working paper, Ohio State University.

Linton, O., E. Massoumi, and Y.-J. Whang (2005). Consistent testing for stochastic dominance under general sampling schemes. *Review of Economic Studies 72*, 735–765.

Marvel, H. (1976). The economics of information and retail gasoline price behavior: An empirical analysis. *Joural of Political Economy 84*, 1033–1080.

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

Prager, R. and T. Hannan (1998). Do substantial horizontal mergers generate significant price effects? Evidence from the banking industry. *Journal of Industrial Economics 46*, 433–452.

Sapienza, P. (2002). The effects of banking mergers on loan contracts. *Journal of Finance 57*, 329–367.

Scott Morton, F., F. Zettelmeyer, and J. Silva-Risso (2003). Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quantitative Marketing and Economics 1*, 65–92.

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

Stavins, J. (2001). Price discrimination in the airline market: The effect of market concentration. *Review of Economics and Statistics 83*, 200–202.

Stigler, G. (1961). The economics of information. *Joual of Political Economy 69*(3), 213–225.

Stole, L. A. (2007). *Chapter 34: Price Discrimination and Competition*, Volume 3. Amsterdam: Elsevier. Handbook of Industrial Organization.

Town, R. and G. Vistnes (2001). Hospital competition in hmo networks. *Journal of Health Ecnomics 20*, 733–753.

Varian, H. (1980). A model of sales. *The American Economic Review 70*(4), 651–659.

Weinberg, M., O. Ashenfelter, and D. Hosken (2013). The price effects of a large merger of manufacturers: A case study of maytag/whirlpool. Forthcoming, American Economic Journal: Economic Policy.

Wildenbeest, M. R. (2011). An empirical model of search with vertically differentiated products. *RAND Journal of Economics 42*, 729–57.

A   Large Canadian Mergers

1993

On January 1, 1993 TD Bank acquired, under duress, Central Guaranty Trust. Even though Central Guaranty Trust had a poor balance sheet, there was substantial interest by several financial institutions in acquiring its assets from Central Capital Corporation, which owned 87 per cent of trust. At the official auction TD won over the joint bid from National Bank, Canada Trust, and Montreal Trust. Given the conditions of the Central Guaranty Trust balance sheet (they had very high risk mortgages and commercial lending activities), the Canadian Deposit Insurance Corporation provided financial support to TD in the takeover. In terms of branches, TD acquired 156 locations. TD also inherited nearly 11 billion dollars in deposits.

Soon after TD Bank's acquisition of Central Guaranty Trust, the second largest acquisition (in terms of assets) in Canada occurred, with Royal Bank acquiring Royal Trust. In 1992 Royal Trust's parent company, Royal Trustco, had experienced liquidity issues, and in early 1993 announced it was looking for a buyer of the trust company. RBC's takeover of Royal Trust was announced on March 18, 1993 and consummated on September 1, 1993. The Royal Trust brand was well-known and well-respected in the financial industry. Furthermore, most people believed that RBC was a perfect match for Royal Trust. They shared the same name, colors, and both had distinguished histories. Royal Trust had 150 branches at the time of the acquisition, largely in Ontario and Quebec, but also with a significant presence in Alberta and British Columbia. According to Competition Bureau (2003), the RBC-Royal Trust merger was analyzed by the Bureau when it was first proposed. The Competition Bureau did not place any restrictions on the merger.

1994

The Competition Bureau also did not place restrictions on Bank of Nova Scotia's (BNS) acquisition of Montreal Trust on April 12, 1994. In this instance BNS acquired 9 billion dollars in deposits and 125 branches located mostly in Quebec and Ontario. Montreal Trust had experienced some losses in the early 1990s because of market value deficiencies in the investments and assets, but was considered a sound financial institution and a good purchase by BNS.

1997

On August 14, 1997 Bank of Nova Scotia acquired National Trust and Victoria and Grey Mortgage Corporation, without restrictions by the Competition Bureau. In this case BNS acquired nearly 12 billion dollars in deposits and 199 branches located in Quebec, Ontario, and Manitoba. Most of what is known about National Trust is confidential. We do know, however, that the transaction was valued at approximately 1.25 billion dollars even though National Trust was considered a poorly run institution that had acquired an excessive number of small, failing trusts throughout the 1990s. It was largely the mis-management of the infrastructure that led National Trust to look for a buyer.

2000

The last bank merger to be approved in Canada was Toronto-Dominion Banks' (TD) acquisition of Canada Trust in 2000. The price tag was roughly 8 billion dollars (TD financed the purchase by issuing 700 million dollars in equity) and it resulted in 440 branches located across the country being acquired as well as a strong share of the mortgage market. The merger was analyzed by the Competition Bureau and allowed to be completed under minor conditions. For example, TD had to divest in some of its branches in three of the seventy-four markets defined by the Bureau (Kitchener-Waterloo-Cambridge-Elmira, Port Hope, and Brantford-Paris). TD also had to sell CT's MasterCard credit card business (they sold the consumer credit card business to Citibank in November 2000). TD was issuing Visa credit cards at the time of the acquisition, and it was not until 2009 that Canadian banks could sell both brands simultaneously.

B   Additional tables

Table 9: Definition of Household / Mortgage Characteristics

Some variables were available only from one of the mortgage insurers.

Name	Description
FI	Type of lender
Source	Identifies how lender generated the loan (branch, online, broker, etc)
Income	Total amount of the borrower(s) salary, wages, and income from other sources
TDS	Ratio of total debt service to income
Duration	Length of the relationship between the borrower and FI
R-status	Borrowers residential status upon insurance application
FSA	Forward sortation area of the mortgaged property
Market value	Selling price or estimated market price if refinancing
Dwelling type	10 options that define the physical structure
Close	Closing date of purchase or date of refinance
Loan amount	Dollar amount of the loan excluding the loan insurance premium
Premium	Loan insurance premium
Purpose	Purpose of the loan (purchase, port, refinance, etc.)
LTV	Loan amount divided by lending value
Price	Interest rate of the mortgage
Term	Represents the term over which the interest rate applies to the loan
Amortization	Represents the period the loan will be paid off
Interest type	Fixed or adjustable rate
$CREDIT$	Summarized application credit score (minimum borrower credit score).

Table 10: Average Treatment Effect of the Merger with Covariates

VARIABLES	Margin		Zero discount	
	Baseline	With trend	Baseline	With trend
Merger	0.0607^a	0.0719^a	0.0646^a	0.0477^a
	(0.0183)	(0.0242)	(0.0154)	(0.0180)
log of Monthly income (X 100K)	0.150^a	0.123^a	0.109^a	0.0909^a
	(0.0181)	(0.0237)	(0.0165)	(0.0220)
Log of house price (X 100K)	-0.310^a	-0.320^a	-0.220^a	-0.231^a
	(0.0194)	(0.0237)	(0.0171)	(0.0219)
Log of other debt (X 1K)	-0.0511^a	-0.0493^a	-0.0310^a	-0.0364^a
	(0.00759)	(0.00985)	(0.00665)	(0.00883)
FICO$\geq$600	-0.225^a	-0.235^a	-0.197^a	-0.198^a
	(0.0354)	(0.0435)	(0.0305)	(0.0402)
Status: Renter	-0.0101	-0.0112	-0.00935	-0.0221
	(0.0113)	(0.0159)	(0.0101)	(0.0140)
Status: Parents	-0.0547^a	-0.0541^b	-0.0517^a	-0.0630^a
	(0.0190)	(0.0274)	(0.0155)	(0.0225)
$0.85 \leq$ LTV ≤ 90	0.0509^a	0.0404^a	0.0257^b	0.0212
	(0.0113)	(0.0150)	(0.0103)	(0.0133)
$0.90 \leq$ LTV ≤ 0.95	0.0650^a	0.0358^c	0.0410^a	0.0279
	(0.0147)	(0.0204)	(0.0130)	(0.0176)
I(Min.down)	0.118^a	0.107^a	0.0874^a	0.0824^a
	(0.0114)	(0.0155)	(0.0112)	(0.0142)
Switcher	-0.0899^a	-0.0488^b	-0.0564^a	-0.0119
	(0.0129)	(0.0189)	(0.0123)	(0.0185)
Broker	-0.0615^a	-0.0378^b	-0.0945^a	-0.0904^a
	(0.0141)	(0.0185)	(0.0122)	(0.0176)
R-squared	0.408	0.416	0.181	0.189

Significance levels: ^{a} p<0.01, ^{b} p< 0.05, ^{c} p< 0.1. Number of observations ins 18,121.

C Theory Appendix

In the interior, the solution to the effort-choice problem in equation (7) gives rise to search effort and matching probability functions that are decreasing in u_i and increasing in n:

$$e(u_i, n) = \left[(u_i/\lambda \Delta(n))^{-1/(1+\lambda)} - 1 \right] \tag{13}$$

$$s(u_i, n) = \begin{cases} 1 - (u_i/\lambda \Delta(n))^{\lambda/(1+\lambda)} & \text{if } \lambda \Delta(n) > u_i \\ 0 & \text{otherwise,} \end{cases} \tag{14}$$

where $\lambda = \gamma/(1-\gamma)$.

To study the effect of a merger we consider the effect of u_i and n on the equilibrium margin:

$$\pi_1(u_i, n) = 1 + e(u_i, n)\left(1 - \frac{1}{1+\lambda} + \frac{1}{1+\lambda}\right) - \frac{1}{1+\lambda} + \frac{1}{1+\lambda} = 1 + e(u_i, n) > 0$$

$$\pi_2(u_i, n) = \frac{u_i(e(u_i, n)+1)}{\Delta(n)(1+\lambda)}\Delta'(n) - \frac{u_i(e(u_i, n)+1)}{\Delta(n)(1+\lambda)}\Delta'(n) - s(u_i, n)\Delta'(n) = -s(u_i, n)\Delta'(n) < 0.$$

Therefore price is increasing in u_i, and decreasing in n. We can also use equation (14) to determine the threshold level $\bar{u}(n)$, associated with positive search effort: $e(u_i) = 0$ if $u_i > \lambda \Delta(n)$. If the search function $s(u_i, n)$ is such that $s(u_i, n) = 0$ for all $u_i > \bar{u}$, the marginal effect of changing n is zero for consumers at the top of the price distribution.

Using these two results leads to our first prediction: A merger raises prices for all $u_i < \bar{u}$: $\pi(u_i, n-1) - \pi(u_i, n) > 0$ for all $u_i < \bar{u}(n)$.

We are also interested in the effect of the merger along the distribution. From the above results we can calculate the following cross-partial derivates:

$$\frac{\partial^2 \pi(u_i, n)}{\partial u \partial n} = \frac{\partial^2 \pi(u_i, n)}{\partial n \partial u} = \frac{e(u, n)}{\Delta(n)(1+\lambda)}\Delta'(n) > 0. \tag{15}$$

From these we see that because the marginal effect of n is increasing in u_i, the effect of a merger is larger (more negative) at the bottom of the price distribution:

$$0 \le \pi(u_{q_1}, n-1) - \pi(u_{q_1}, n) \le \pi(u_{q_2}, n-1) - \pi(u_{q_2}, n), \quad \text{if } q_2 < q_1. \tag{16}$$

Since the function is continuous and the cross-partial derivatives are strictly positive (in the interior), we can establish the following effect of a merger on price dispersion as measured by the interquartile range (IQR):

$$\text{sign}\left(\text{IQR}(n-1) - \text{IQR}(n)\right) = \text{sign}\left(\frac{1}{2}\int_{u_{25}}^{u_{75}} \frac{\partial \pi(u_i, n-1)}{\partial u}du_i - \frac{1}{2}\int_{u_{25}}^{u_{75}} \frac{\partial \pi(u_i, n)}{\partial u}du_i\right) \le 0$$

The second equality comes from the fact that:

$$\pi(u_{75}, n) = \pi(u_{25}, n) + \int_{u_{25}}^{u_{75}} \frac{\partial \pi(u_i, n)}{\partial u}du_i, \tag{17}$$

and its sign is negative since from equation (15):

$$\frac{\partial \pi(u_i, n)}{\partial u} \ge \frac{\partial \pi(u_i, n-1)}{\partial u}. \tag{18}$$