

PRICE COMPETITION AND CONCENTRATION IN DECENTRALIZED MARKETS: EVIDENCE FROM A MERGER BETWEEN CANADIAN MORTGAGE LENDERS

JASON ALLEN
BANK OF CANADA

ROBERT CLARK
HEC MONTRÉAL

JEAN-FRANÇOIS HOUDE
UNIVERSITY OF WISCONSIN-MADISON

ABSTRACT. This paper examines the impact of bank consolidation on price competition in mortgage markets. Mortgage markets are decentralized and so rates are determined through a search and negotiation process. The primary effect of consolidation is therefore to reduce the number of partners available with whom to negotiate, although it can also change the characteristics of the product, and impact the search effort of consumers. Using a Canadian merger as a case study, we find that consolidation had a small, but statistically significant effect on mortgage interest rates. A decomposition of this aggregate effect reveals important heterogeneity in the impact of the merger. We find that consumers able to negotiate sizeable discounts were adversely affected by the merger, while consumers at the top of the conditional price distribution were unaffected. These results suggest that, on average, the mortgage market is fairly competitive, despite the presence of a small number of banks in the Canadian market. However, not all consumers benefit from this competition, and our results suggest that banks enjoy substantial market power over consumers with large search costs.

Date: This version: March 7, 2012. First version: January 14, 2010.

Correspondence: Jason Allen: Bank of Canada, Ottawa, Ontario, K1A 0G9; Phone: (613) 782-8712; Email: jallen@bankofcanada.ca, Robert Clark: HEC Montréal, CIRANO and CIRPÉE, Montreal, Quebec; Phone: (514) 340-7034; Email: robert.clark@hec.ca, Jean-François Houde: University of Wisconsin-Madison, Madison, Wisconsin and CIRANO; Phone: (608) 262-3805; Email: houdejf@ssc.wisc.edu. Jason Allen thanks UW-Madison and HEC Montreal for their hospitality while writing this paper. Robert Clark thanks the FQRSC and HEC Montréal for funding. Jean-François Houde thanks the UW Graduate School research grant program and NSF (SES-1024840) for funding. We thank the Canada Mortgage and Housing Corporation for providing us with the data. We also thank Allan Crawford, Walter Engert, and Ron Morrow for their consistent effort on this project. We thank seminar participants at the Bank of Canada, FTC, DePaul University, Boston University, Duke University, New York University, the IIOC (Vancouver), Invitational Choice Symposium (Miami), SEDs (Montréal), AEA (Denver and Chicago) for their comments. We thank Peter Newberry for research assistance. We greatly benefited from discussions with Dan Akerberg, Jean Boivin, Michèle Bourque, Liran Einav, Gautam Gowrisankaran, Kim Huynh, Ken Hendricks, Jack Porter, Matt Shum, Alan Sorensen, and Chris Taber. The views in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada. All errors are our own.

1. INTRODUCTION

North American and European mortgage markets have become increasingly concentrated over the past two decades. Understanding the effect of concentration on interest rates is central to the design and evaluation of policies meant to regulate mortgage markets. These include policies regarding the approval of bank mergers, constraints on the scope of bank activities, price discrimination, and banks' costs of funding (e.g. capital or securitization). In this paper, we take advantage of the quasi-experimental variation created by a merger between two important lenders in Canada to evaluate the effect of concentration on the distribution of mortgage interest rates paid by consumers.¹

Like many consumer finance product markets, the Canadian mortgage market is fully decentralized: each week financial institutions post a common interest rate that does not vary across regions, and contract terms are determined through a search and negotiation process between local branch managers and individual borrowers. This feature 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 Canada and in the United States show that, while some buyers get multiple quotes when shopping for their mortgage contract, nearly half only get one.²

The price setting mechanism and heterogeneity across consumers induce a substantial amount of price dispersion, even for contracts with homogeneous maturities and terms, and for which lenders are fully protected against the risk of default by a government insurance program. Indeed, approximately ninety percent of consumers receive a discount off the posted interest rate when negotiating their first mortgage, and on a typical week discounts range from zero to 200 basis points (bps), with an average of 100 bps. This dispersion is substantial, both relative to the average transaction interest rate (7.15%), and the average transaction margin relative to federal bonds with the same maturity (108 bps).

In this context, a lender's capacity to exercise market power vis-a-vis a borrower depends not just on differentiation as in standard posted-price markets, but also on the borrower's ability to generate competitive outside offers. A merger can enhance market power by reducing the number of negotiating partners accessible to consumers. We are interested in quantifying the impact of this loss on consumer-level transaction rates. If the market is sufficiently competitive to start with, the loss of a bargaining partner should not raise prices significantly for the average borrower. This is the essence of our first econometric test. Moreover, if consumers differ in their ability to negotiate discounts, the merger should have heterogeneous effects on consumers facing the same set of lending options. To investigate this possibility, our econometric analysis focuses on estimating

¹For confidentiality reasons we cannot reveal the exact details of the merger.

²In Canada, the Canadian Association of Accredited Mortgage Professionals conducts an annual survey on the state of the mortgage market and reports that about 45% of new home buyers visit one lender. In the U.S., LendingTree reported that 39% of new home buyers gather only one quote. Gerald et al. (2010) study the role of financial illiteracy in explaining the rate of foreclosures during the recent U.S. housing crisis.

the distribution of the treatment effect of the merger on transaction rates across observationally equivalent groups of consumers.

To conduct our empirical analysis we exploit a detailed data-set containing a random sample of new mortgages administered by the Canada Mortgage and Housing Corporation (CMHC); a federal crown corporation responsible for insuring the loans of home buyers who require insurance under the National Housing Act. These data provide administrative information on contract terms, household characteristics, and market-level characteristics. The richness of this consumer information, in combination with lender-level location data, allow us to identify the causal effect of the loss of a lending option on transaction interest rates.

We use a difference-in-difference estimation strategy to quantify this effect. This relies on the fact that most consumers shop for their mortgage contract locally, and that the merger differentially impacted the structure of local-markets across the country. In particular, roughly sixty per cent of consumers in our data-set had both merged entities in their neighborhoods prior to the merger, while the remaining had only one or neither. The latter set of borrowers represent a natural control group. This strategy has recently been used to study the impact of mergers in gasoline markets (e.g. Hastings (2004), Hastings and Gilbert (2005), and Houde (2011)), in the cement industry (e.g. Hortacsu and Syverson (2007)) and in the health-care industry (e.g. Dafny et al. (2011)), in deposit banking markets (e.g. Focarelli and Panetta (2003)), and, closely related to us, in the market firms loans (e.g. Sapienza (2002), and Panetta et al. (2009)). See Ashenfelter et al. (2009) for a survey of the retrospective merger evaluation literature.

Our paper contributes to this literature in two ways. First, the scope of the merger induced important changes in the structure of local markets, and as a result provides a credible source of variation to test for the presence of a positive relationship between concentration and prices. Second, our study provides important insights on the sources of market power in markets with negotiated prices. While several papers have documented the importance of price dispersion and search frictions in consumer markets with negotiated prices, the literature studying the relationship between concentration and prices has mostly focused on posted price markets.³ Our paper helps fill this gap, and provides the first empirical analysis of the impact of a merger on the distribution of prices in a decentralized market.

The results can be summarized as follows. We estimate that the merger led to an increase in the average interest rate in treated markets of approximately 6 basis points (bps), which corresponds to 12% of the standard-deviation of rates in our cross-section. For an average loan size of \$131,000, we estimate that the merger led to a \$6.35 increase in monthly payments (or \$381 over the five years of the contract). This effect is small, and can be thought of as a lower bound because of the likely presence of measurement error in the choice-set of consumers. Moreover, decomposing the effect of the merger by observed transaction characteristics reveals important asymmetries. Consumers transacting with the merged entity experienced the largest price increase: up to 14 bps

³Here is a non-exhaustive list of markets with negotiated prices that have been studied in the literature: insurance (Dahlby and West (1986)), new cars (Goldberg (1996), Busse et al. (2006), Scott Morton et al. (2003)), car loans (Adams, Einav, and Levin (2009)), mutual funds (Hortaçsu and Syverson (2004)), housing (Hendel et al. (2010)), and mortgage brokers (Hall and Woodward (2010)).

(\$890), compared to 4 bps (\$254) for consumers selecting a competing institution. The effect is also significantly larger for consumers who are not facing borrowing constraints due to the insurance program; suggesting that consumers facing more and better lending options were more likely to be adversely affected. We also find that the effect of the merger is increasing in the overall presence of the merged entity; up to 30 bps in markets in which the merging parties control fifty percent of retail branches.

Using the estimator proposed by Athey and Imbens (2006) we estimate the effect of the merger across the different quantiles of the residual price distribution. Conditional on a rich set of characteristics, consumers in the low percentiles of the rate distribution receive a relative large discount while borrowers in the high percentiles of the rate distribution pay close to the posted price. We find that losing a lender option raised interest rates between 7 and 8 bps for consumers in the lower and middle percentiles of the conditional transaction rate distribution, but is statistically indistinguishable from zero in the top 35%. In other words, borrowers receiving little to no discount are not affected by the merger.

These results are consistent with the presence of sizable and heterogenous search costs among borrowers. When negotiating with a lender the rate a borrower pays is a function of his/her outside option, which in turn depends on his/her search cost and negotiation ability, as well as the structure of the local market. Our results suggest that, on average, losing a lending option has little effect. Since margins for these borrowers are quite low, our findings imply that the market is fairly competitive and that the average borrower is able to extract a large share of the transaction surplus through search and negotiation.

However, our quantile results inform us that there is important heterogeneity in the effect of the merger. Since we condition on a rich set of characteristics, the different percentiles of the distribution can be interpreted as the unobserved search and negotiation ability of consumers – borrowers receiving big discounts are those with low search cost and/or high negotiation ability. Our results confirm that only consumers gathering multiple quotes are adversely affected by the merger. In contrast, consumers who are unable or unwilling to negotiate pay a price that is strictly a function of their willingness to pay, and not directly affected by the loss of a lending option.

These results have important policy implications. On the one hand, our analysis suggests that, from the point of view of consumers who search, the market is fairly competitive. Although rates increase, the extra interest cost for the average searcher caused by the merger is relatively small, especially considering the fact that our analysis focuses exclusively on local markets with a small number of lenders (i.e. between five and eight). On the other hand, we also find that borrowers at the top of the conditional price distribution, and therefore less likely to search, are not affected by the degree of competition between lenders: access to fewer lender options does not lead to higher rates for these borrowers. Therefore, policies designed to increase competition, or at least to prevent increases in concentration, may not be effective for those consumers unwilling to search or unable to negotiate. Instead, policies aimed at improving financial literacy, or reducing the ability of lenders to price discriminate between consumers based on negotiation ability are more likely to be effective at reducing interest rates for these consumers.

The paper is structured as follows. Section 2 describes the Canadian mortgage markets, focusing on market structure, mortgage insurance and contract type, pricing, negotiation, and shopping habits. Section 3 presents the data. Section 4 discusses our identification strategy, and Section 5 presents the results. Section 6 discusses the results through the lens of a simple bargaining model. Section 7 concludes. We relegate some tables to the Appendix.

2. THE CANADIAN MORTGAGE MARKET

2.1. 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 Quebec), and a provincially owned deposit-taking institution (Alberta's ATB Financial). Collectively, they control 90 per cent of assets in the banking industry and are called 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 make mortgage loans, funding them by issuing guaranteed investment certificates and accepting deposits. 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.⁴

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. Today, there are still many trusts operating in Canada, but they are small and their influence on the mortgage market is much less than it was prior to 2000. 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),

⁴There 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 the market share of financial institutions among new insured mortgage contracts (smoothed)

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 one of these major acquisitions.

2.2. Mortgage contracts and mortgage insurance. There are two types of mortgage contracts in Canada – conventional mortgages which are uninsured since they have a low loan-to-value ratio, and high loan-to-value mortgages, which require insurance (for the lifetime of the mortgage). Today, approximately 80% of newly issued mortgages fall into the latter category. The primary insurer is the Canada Mortgage and Housing Corporation (CMHC), a crown corporation with an explicit backstop from the federal government.⁵ Our analysis focuses on insured mortgages.

All insurers use the same strict guidelines for insuring mortgages.⁶ First, borrowers with less than 25% equity must purchase insurance.⁷ Second, borrowers with monthly gross debt payments that are more than 32% of gross income or a total debt service ratio of 40% will almost certainly be

⁵There are a number of private insurers as well, the only one in existence during our sample was Genworth Financial, which also has an explicit government of Canada guarantee, albeit for 90 per cent. CMHC's market share during our sample averages around 80 per cent.

⁶The reference parameters are set by the Federal Government. See Traclet (2005) for more details.

⁷This is true during our sample. Today borrowers with less than 20% equity must purchase insurance.

rejected.⁸ CMHC charges the lenders an insurance premium, ranging from 1.75 to 3.75 per cent of the value of the loan – lenders pass this premium on to borrowers. Insurance qualifications (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. According to a recent report by the Royal Bank of Canada (Hardy and Mun (2009)), the historical average cost of mortgage default in Canada is 2 basis points per year. Historical delinquency rates in Canada are less than 0.4%, about ten times smaller than in the U.S.

During our sample period, nearly all mortgage contracts were fixed rate, among which over 85 per cent had a 5 year term (the second most common term is 36 months). A 5 year fixed-rate mortgage contract must be renegotiated every five years, which in effect acts like an adjustable rate mortgage with a fixed time-frame to renegotiate. This has been the standard contract offered by Canadian banks since the late 1960's. Most contracts also have 25 year amortization periods.

2.3. 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 at the big banks: the coefficient of variation on posted rates for the Big six is close to zero.

In contrast there is a significant amount of dispersion in transaction rates. In Allen et al. (2011a) we document that the coefficient of variation in margins between 1999 and 2004, for example was 60%. Approximately 10% of borrowers pay the posted rate. The remainder receive a discount below the posted price. This comes about because borrowers can search for and negotiate better rates. 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.⁹ Given advances in technology, software in a central location can now process large volumes of applications and return a standard price for a particular set of consumer characteristics to a local manager. The local manager, however, has authority to negotiate. 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. Survey evidence from the Canadian Association of Accredited Mortgage Professionals reports that about 55% of new home buyers (new home buyers are defined as individuals who get a new mortgage as opposed to homeowners who refinance or renew an existing mortgage) visit multiple lenders when shopping for a mortgage, but the remainder visit only one.¹⁰

⁸Gross 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.

⁹A household survey on residential shopping patterns by Altus Clayton Group and Ipsos-Reid found that between 55% and 70% of respondents claimed to walk into a local branch to search out rate options while 35% to 40% called a local branch.

¹⁰In the U.S., LendingTree reports that 39% of new home buyers gather only one quote. See also Lee and Hogarth (2000).

Alternatively borrowers can hire brokers to search for the best rates on their behalf. Unlike in the United States, brokers in Canada have fiduciary duties. Brokers are compensated by lenders, but “hired” by borrowers to gather the best quotes from multiple lenders. Detailed survey evidence by Taddingstone in 2005 (MortgageBrokerReport@taddingstone.com) found that brokers on average contact 5.9 lenders for their clients, suggesting they do, in fact, assist in gathering multiple quotes.

3. DATA

3.1. Mortgage-contract data. Our data set is a sample of insured mortgage contracts obtained directly from CMHC.¹¹ We obtained a 10% random sample of new mortgage contracts issued between 1992 and 2004, sampled by Census Metropolitan Area (CMA).¹² We further restrict the sample to contracts signed within a year of the merger. This is in part for convenience, since our branch data are annual. In addition, we know that the merged entity did not start closing duplicate branches until approximately a year after the official merger date. Our analysis therefore focuses on the short-run impact of the merger, holding fixed the distribution of branches.

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), and some demographic characteristics (e.g. income, prior relationship with the bank, residential status, dwelling type).¹³ In addition, we observe the location of the purchased house up to the forward sortation area (FSA). This unit of aggregation 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 16,000.¹⁴

During our sample period, some of these characteristics were missing for a fraction of contracts (mostly risk, residential status, and financial intermediary). We include these observations in the regression analysis, but add a series of interaction terms with group dummy variables to control for the fact that the number of missing values was decreasing over time and unequal across lenders (see footnote 20).

We restrict our sample to newly issued mortgages, excluding home-owners that are either re-financing or renewing their mortgage contract. We also focus on contracts with homogenous lengths and terms. In particular, our analysis focuses on contracts with a 25 years amortization period, and 5 year fixed-rate term.¹⁵

Table 1 describes the key variables that we use in our analysis. The key outcome variable is the transaction interest rate paid by consumers. The first row describes the distribution of the interest

¹¹Although we also have access to Genworth contracts, which we use in Allen et al. (2011a) to highlight the importance of discounting in the Canadian mortgage market, the Genworth data does not include contracts from Trust companies or cover the entire sample period, and therefore is of limited use for our merger analysis.

¹²Breslaw et al. (1996) have previously used this data to study mortgage term and amortization choice in the 1980s.

¹³Table 8 in the Appendix lists all of the variables included in data-set.

¹⁴The FSA is the first half of a postal code. There are over 1,300 FSA's in Canada, and over 850,000 postal codes.

¹⁵This sample selection choice, however, could bias our estimates if banks reacted to the merger by offering more aggressive discounts on short terms contracts, or variable interest rates. In results not reported here we formally test the null hypothesis of exogenous contract choice with respect the merger, and fail to find evidence of endogenous selection.

TABLE 1. Summary statistics on mortgage contracts and household characteristics

VARIABLES	Mean	Std-dev.	P_{25}	P_{50}	P_{75}
Transaction rate deviation	7.17	.519	6.83	7.07	7.56
Monthly Payment	948	401	635	882	1195
Income (x 1000\$)	66.6	26	47.9	62.8	80.1
Loan (x 1000\$)	131	56.3	86.9	121	166
Loan-to-value	.913	.0433	.9	.927	.95
FICO score ≥ 600	.636	.481	0	1	1
Renter	.699	.459			
Parents	.0804	.272			
Broker	.25	.433			
Switcher	.339	.473			

The sample size is equal to 17,139. It includes a random sample of homogenous term and amortization contracts insured by CMHC within one year of the merger. Transaction rate is expressed in difference from weekly average (scaled up by the unconditional average). Renters and parents correspond to new home buyers (the omitted category is home owners). 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 five to eight lenders located within five kilometers of their FSA centroid (see section 2.1 for more details).

rate paid by a consumer, expressed in deviation relative to the weekly average (and then scaled up the unconditional average rate between 1991 and 2004). This variable highlights the large amount of cross-sectional dispersion in rates across consumers.¹⁶ Interestingly, the week-to-week variation corresponds to 48% of the total variance in interest rates: the standard-deviation goes from 0.74 to 0.52 when we express rates in deviation from their weekly average. Therefore, slightly more than half of the total interest rate variation in our data is coming from cross-sectional dispersion. In comparison the average margin, measured relative to the 5-year government bond rate, is equal to 1.08 during our sample period.

This magnitude of dispersion is large compared to other financial markets, considering the homogeneity of the contract terms. For instance, Hortaçsu and Syverson (2004) show that the cross-sectional standard-deviation among mutual fund transaction fees was equal to 60 bps in 2001. Mutual funds exhibit substantially more heterogeneity in returns than the 5-year fixed rates mortgage contracts that we study.

The average home owner in our data earns \$67,000, and contracts a loan of \$131,000. A large fraction (40%) of households are constrained by the minimum down-payment requirement of 5% of the house price. Also, most consumers represent relatively low default risks to the CMHC, since 65% of borrowers have credit scores greater than 600. This is partly due to the fact that households are constrained to have a total debt service ratio below 40%.

In our sample 25% of contracts were negotiated through a broker, and 34% of consumers switch to a financial institution with whom they did not have any prior experience. Notice that the vast majority of broker transactions are labeled “switchers” (i.e. 74%), since brokers most commonly

¹⁶We omit the summary statistics on the level to avoid revealing the sample period, and preserve the confidentiality of the transaction. We use rates in levels in the econometric analysis, and control for week fixed-effects.

deal with smaller institutions (i.e. trust or insurance companies). This is because some of the major banks refuse, sometimes explicitly, to deal with brokers.¹⁷

3.2. Lender location data. We have compiled branch-location information for all financial institutions in Canada from the Financial Services Canada directory produced by Micromedia ProQuest. This data-set provides a yearly panel of branch addresses for every lender. We use this information together with the house locations to count the number of branches and number of lenders present in a neighborhood of each contract. We also construct measures of market concentration using the branch share of lenders located around each house. These measures are used to characterize the choice-set of consumers, as we discuss in greater details below. For more details on the evolution of branch networks in Canada see Allen, Clark, and Houde (2008).

4. ESTIMATION AND IDENTIFICATION STRATEGY

4.1. Measuring the impact of the merger. Our analysis focuses on one of the major acquisitions that took place in Canada between the early 1990s and early 2000s. For confidentiality reasons we label the two merging institutions A and B . Firm A is a national bank, and therefore has branches in nearly all local markets in Canada prior to the merger. Firm B , a Trust, is smaller, and prior to the merger less than 2 per cent of its observed contracts were in markets where A was not one of its competitors.

Our objective is to measure the treatment effect of the merger on transaction interest rates, denoted by p_i . We distinguish between three effects: (i) the average treatment effect: $E(\Delta p_i) = \alpha$, (ii) the treatment effect conditional on observed characteristics: $E(\Delta p_i|X) = \alpha(X)$, and (iii) the distribution of the treatment effect: $F_\alpha(x) = \Pr(\Delta p_i < x)$.

We follow a difference-in-difference estimation approach to identify these three objects. In particular, we use observed contracts before and after the merger in “treated” neighborhoods that were directly impacted by the merger, and in “control” neighborhoods that did not experience the same decrease in the number of lenders. The control transactions are used to estimate the counter-factual distribution of rates in the treated areas absent of the merger.

In order to define a set of consumers affected by the merger, we assume that consumers shop for mortgages in a neighborhood around the house they purchase. This is akin to defining market boundaries in a traditional merger retrospective analysis, and in our context defines a consumerss choice-set. Treated consumers are defined as purchasing a house in a neighborhood in which both A and B were present prior to the merger, and consumers located in neighborhoods with only A or B , or none are part of the control group. We let G_i denote the treatment group indicator. Similarly, T_i is a period indicator variable equal to one for contracts signed after the merger, and $M_i = T_i \cdot G_i$ is the treatment (merger) indicator variable.

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

¹⁷Brokers in Canada are compensated by lenders an amount equal to 1-1.3 per cent the value of the mortgage. Unlike in the United States, mortgage brokers are not compensated as a function of the interest rate. Furthermore, mortgage brokers have fiduciary duties with respect to the borrower.




FIGURE 2. Distribution shortest distances between homes and banks

purchased house. Although we observe each household only once, we observe most FSAs before and after the merger.

In the empirical analysis, we use a fixed radius of 5 KM around the center of each FSA to define choice-sets.¹⁸ This assumption appears to be reasonable in the case of mortgage contracts since consumers transact with an institution that is on average located within 2 kilometers of the center of their FSA.¹⁹ The two histograms in Figure 2 show that the average distance to consumers' chosen lenders is much smaller than the average distance to other financial institutions.

4.2. Average treatment effects. Our baseline econometric specification is a linear pricing equation relating the observed transaction interest rates with the group indicator variables defined above, and observed characteristics of the contracts:

$$(1) \quad p_i = \alpha M_i + \gamma G_i + \beta' \mathbf{Z}_i + \lambda' \mathbf{Week}_i + u_i,$$

where α measures the average effect of the merger on transaction interest rates, \mathbf{Z}_i is a vector of control variables describing the financial and demographic characteristics of the contract (i.e. bank and province indicator variables), and \mathbf{Week}_i is a vector of week dummy variables (i.e. λ absorbs the effect of T_i). We control for variables that we believe are predetermined at the negotiation stage (i.e. financial, market, and demographic characteristics), and include period fixed-effects to isolate the effect of the merger on the cross-sectional distribution of rates (i.e. β absorbs potential aggregate effects of the merger). We also include in \mathbf{Z}_i time-varying variables measuring the

¹⁸In section 5.4 we relax this assumption and evaluate the robustness of our results to the size of competing neighborhoods.

¹⁹We define distance as the Euclidian distance between the center of each forward sortation area of the household and the closest branch, by postal code, associated with the chosen institution.

distribution and number of branches among competing networks to account for changes in the number retail branches over the merger period.²⁰

To estimate heterogenous treatment effects we allow α to be a linear function of X_i :

$$(2) \quad p_i = \alpha M_i + \alpha_X M_i \cdot X_i + \gamma G_i + \lambda_X T_i \cdot X_i + \gamma_X G_i \cdot X_i + \beta' \mathbf{Z}_i + \lambda \text{Week}_i + u_i,$$

and $\alpha + \alpha_X X_i = \alpha(X_i)$ measures the average treatment effect conditional on X_i . In the empirical analysis we are interested in distinguishing between the effect of the merger on consumers dealing with the merged entity (i.e. **direct**), and on clients of competing institutions (i.e. **indirect**). In order to test for presence of a dominance effect, we also allow $\alpha(X)$ to vary with the branch share of the merging firms in each local market, as well as with the number of lenders present. We will also estimate the effect of the merger across two groups of consumers: borrowers putting the minimum downpayment (i.e. $LTV = 95\%$), and unconstrained consumers (i.e. $75\% \leq LTV < 95\%$).

4.2.1. Identification assumptions. Unlike other forms of program evaluation, retail mergers are not endogenously chosen by local market participants. Since bank mergers are negotiated nationally, these changes can be viewed as exogenous relative to local market conditions, at least in the short run.

This does not mean that the treatment is necessarily independent of unobserved transaction attributes u_i , since the timing and location of local mergers can be correlated with other aggregate variables affecting interest rates. The key identifying assumption that allows us to interpret α (and α_X) as the causal effect of the merger, is that conditional on \mathbf{Z}_i , differences between the treatment and control groups are the same on average in different time periods (pre- and post-merger). That is, the merger is not *confounded* with aggregate trends specific to the control or treatment groups.

A related problem is heterogeneity in choice-sets of borrowers, which would introduce measurement error in the treatment variable. This will lead to an attenuation bias if the error in the measurement of the actual choice-set of consumers is independent of G_i . If that is the case, the OLS coefficient can be thought as a lower bound estimate of the true average treatment effect.

4.2.2. Comparability of treatment and control groups. An often overlooked issue in the evaluation of mergers is the comparability of treatment and control groups. In our context, the most important difference is that local markets with both A and B tend to be larger and have more lenders. Figures 3a and 3b illustrate that the two groups significantly overlap only in medium-size markets with five to eight lenders. Less than 1% of control markets have more than 8 lenders, while less than 1% of treatment markets have fewer than 5.

These differences in the structure of local markets between treatment and control groups can bias our results for at least two reasons. First, if the merger increases market power by reducing the number of competing quotes that consumers can obtain, the average treatment effect should

²⁰The exact set of control variables are: 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), FSA demographic characteristics (i.e. income, house value, fraction of renters, fraction with university degree, fraction of inter-provincial migration, age), number of lenders (other than A or B), and average number of branch per lender (other than A or B). To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.




FIGURE 3. Distribution of the number lenders per local markets before the merger

be different for markets with a small versus large number of lenders. Ideally, we would estimate the treatment effect separately for consumers facing similar choice-sets, but the distribution of lenders only allow us to do so for local markets with an intermediate level of concentration.

A second validity threat comes from the fact that markets with a larger number of lenders tend to be more urban, and therefore possibly subjected to correlated aggregate trends. These unobserved factors would be confounded with the causal effect of the merger, since nearly fifty percent of treated observations are located in areas with more than eight lenders. This is directly related to the concern exposed in Heckman et al. (1997): “the failure to locate participants and comparison group members in the same labor market is a major source of evaluation bias”, even more so than selection bias caused by endogenous participation.

Although controlling for observable differences by including Z_i in the regressions solves part of this problem, the linear specification imposes very strong restrictions when the two groups are not comparable. When treated and control observations are too different, the linear specification will fail to accurately predict the evolution of rates in the treated markets absent of the merger.

Our main approach to deal with this issue is to restrict the sample to local markets with five to eight lenders, and to estimate equation 1 by OLS. This restriction comes at the cost of weakening the external validity of our estimates, since we can only measure the effect of a merger in markets with a moderate level of competition. With this additional restriction, households with only one or none 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 17,139 observations over 427 different locations (including 265 with both institutions).

Table 2 describes the structure of choice-sets, defined as the 5 KM buffer around each observed contract. The first two columns illustrate the distribution of the number of lenders and branches in the full sample; including markets with more than 8 and less than 5 lenders. Excluding the non-overlapping local markets leads to comparable neighborhoods: the average number of branches and Herfindahl-Hirschman Index (HHI) are similar across the two groups, as the number of

TABLE 2. Summary statistics on consumer choice-sets prior to the merger

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

Each entry corresponds the sample average and standard-deviation (in parenthesis), calculated using the observation weights from the mortgage contract data-set. Local markets are defined as 5 KM 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.

lenders. The last three rows show that, absent of any other simultaneous change, 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 27%. The merger alone therefore corresponds to an average increase of 0.213 points in the branch HHI.

Our second approach to deal with the comparability of treated and control neighborhoods is to use the difference-in-difference matching estimator proposed by Heckman et al. (1997, 1998). This estimator computes the counter-factual average interest rate changes in the treatment group, using a weighted average of changes observed in the control group. We follow Heckman et al., and use a Kernel estimator based on the propensity score of being “treated” by the merger.

In our context, although we observe each household only once, we observe most neighborhoods before and after the merger. We compare contracts according to the probability that their respective neighborhood had branches from bank *A* and trust *B* prior to the merger. To calculate the propensity score, we estimate a Logit model at the contract-level that controls for time-invariant demographic characteristics and branch concentration of each neighborhood. Moreover, we compare neighborhoods with the same number of lenders pre-merger. By construction this restricts our analysis to markets with five to eight lenders, and therefore uses the same sample as the OLS estimator.

This leads to the following estimator of the average treatment effect on treated consumers:

$$(3) \quad \hat{\alpha}_M = \frac{1}{|\mathcal{N}_{1,1}|} \sum_{i \in \mathcal{N}_{1,1}} \left\{ p_i - \sum_{j \in \mathcal{N}_{0,1}} \omega_1(i, j) p_j \right\} - \frac{1}{|\mathcal{N}_{1,0}|} \sum_{i \in \mathcal{N}_{1,0}} \left\{ p_i - \sum_{j \in \mathcal{N}_{0,0}} \omega_0(i, j) p_j \right\},$$

where $\mathcal{N}_{G,T}$ denotes the set of observations in group $G = \{0, 1\}$ (i.e. $G = 1$ for treatment group) and period $T = \{0, 1\}$ (i.e. $T = 1$ after the merger), $|\mathcal{N}_{g,t}|$ measures the number of observations in (G, T) , and $\omega_T(i, j) = K_b(i, j) / \sum_{j \in \mathcal{N}_{0,T}} K_b(i, j)$ is a weight measured with a kernel with bandwidth b over the propensity score of observations i and j . Note that $K_b(i, j) = 0$ if i and j do not face the same number of lenders per-merger.

To control for time-varying observed characteristics Z_i , we implement the regression-adjusted estimator proposed by Heckman, Ichimura, and Todd (1997). We first project rates on the contract characteristics, week and province fixed-effects using the control group observations. We estimate the matching DID estimator in a second stage using the residual dispersion. Standard-errors are calculated by bootstrapping the two-step procedure.

4.3. Distribution of treatment effect. To estimate the distributional effects of the merger we use the change-in-change estimator proposed by Athey and Imbens (2006). Specifically, they develop an alternative non-parametric estimator for discrete treatments that allows for heterogenous and non-linear effects.

In our context, we rewrite the reduced-form pricing function in equation 1 as the sum of two components: (i) a linear function of observed consumer characteristics Z_i , and (ii) a non-linear function of consumers unobserved attributes u_i and aggregate factors summarized by T_i and M_i . This leads to a semi-parametric pricing equation:

$$(4) \quad p_i = \beta' \mathbf{Z}_i + \lambda' \mathbf{Week}_i + \mu(u_i; T_i, M_i).$$

This specification allows the non-linear function to change over time (i.e. T_i), for instance, with the diffusion of mortgage brokers into the market, or aggregate business cycles. The effect of the merger is possibly non-linear, and changes the shape of the $\mu(\cdot)$ function. As we discuss in Section 6 below, $\mu(u_i)$ can be thought of as the markup that a consumer of type u_i is able to negotiate over the common lending cost $c_i = \beta' \mathbf{Z}_i + \lambda' \mathbf{Week}_i$. Under this interpretation, it is reasonable to expect that the effect of the merger will differ across consumers with different unobserved bargaining skills. Our objective is to measure the distribution of the merger effect across consumer types:

$$(5) \quad \alpha(u_i) = \mu(u_i; T_i = 1, M = 1) - \mu(u_i; T_i = 1, M = 0), \quad \forall i \text{ such that } G_i = 1.$$

In order to estimate the distribution of $\alpha(u_i)$, we impose three assumptions discussed in Athey and Imbens (2006).

Assumption 1. *The function $\mu(u; T_i, M_i)$ is strictly monotonic in u_i , and constant across groups G .*

Assumption 2. *The distribution of types $u_i \sim H_G(u)$ differs across groups G , but is constant over time.*

Assumption 3. *The support of u_i in the treatment group overlaps with the support of u_i in the control group: $\mathcal{U}_1 \subseteq \mathcal{U}_0$.*

The first assumption provides a strict ranking of markups with respect to consumer types, while the other two assumptions allow us to compare the change in markups over time in the treatment and control groups. The objective is to calculate the change in the function $\mu(\cdot)$ that would have occurred over time without the merger by constructing the counterfactual margin

$\mu(u_i; T_i = 1, M = 0)$ for all i such that $G_i = 1$. This transformation is implemented using the fact that the change in the q^{th} percentile of the control price distribution identifies the change in the markup function that is strictly due to time. Using this logic, and the fact that the empirical price distributions are invertible (under assumption 1), we can recover the counter-factual distribution of prices, denoted $\tilde{F}_{1,1}$, in the treatment group for any price p in the common support:

$$(6) \quad \tilde{F}_{1,1}(p) = F_{1,0} \left(F_{0,0}^{-1} (F_{0,1}(p)) \right) = F_{1,0} (k_{1,1}(p)),$$

where $F_{G,T}(p)$ is the CDF of prices in the sub-population (G, T) , and $k_{1,1}(p)$ is the counter-factual price transformation. Intuitively, we obtain the counter-factual distribution of prices by transforming the observed price distribution in period zero (i.e. $F_{1,0}(p)$) in a such a way that mimics the change in the price distribution observed in the control group. From this transformation we can recover the average treatment effect on the treated:

$$(7) \quad \begin{aligned} \alpha_C &= E [P_i | i \in \mathcal{N}_{1,1}] - E [k_{1,1}(P_i) | i \in \mathcal{N}_{1,0}] \\ &\approx \frac{1}{|\mathcal{N}_{1,1}|} \sum_{i \in \mathcal{N}_{1,1}} p_i - \frac{1}{|\mathcal{N}_{1,0}|} \sum_{i \in \mathcal{N}_{1,0}} \hat{F}_{0,1}^{-1} \left(\hat{F}_{0,0}(p_i) \right) = \hat{\alpha}_C, \end{aligned}$$

where $\hat{F}_{G,T}$ corresponds to the empirical CDF of prices in sample (G, T) . We use a subscript C to label this the “change-in-change” estimator (CIC).

Similarly, we can recover an estimate of the effect of the merger for each percentile of the distribution:

$$(8) \quad \hat{\alpha}_C(q) = \hat{F}_{1,1}^{-1}(q) - \hat{F}_{0,1}^{-1} \left(\hat{F}_{0,0} \left(\hat{F}_{1,0}^{-1}(q) \right) \right).$$

Since we can normalize the scale of u_i to be between zero and one, the previous expression measures the estimated effect of the merger on a consumer with outside option $u_i = q$.

In practice, we construct both estimators while controlling for covariates in the parametric form suggested by equation 4, assuming that u_i is independent of Z_i . We follow the suggestion of Athey and Imbens (2006) (pages 465-466), and construct the empirical price distributions using the residuals of a regression of transaction interest rates on Z_i . The standard-errors are obtained by bootstrapping.²¹

5. RESULTS

5.1. Average treatment effects. In Table 3 we present OLS regression results of the average treatment effect of the merger. The dependent variable is measured as the transaction rate. We describe the sample and control variables in the footnote of each table.

In column (1) we estimate that the merger led to an average interest rate increase in treated markets of 5.5 basis points (bps). This corresponds to 12% of the cross-sectional standard-deviation of interest rates, and 5.5% increase in retail margins relative to the five-year bond rate. Evaluated at the average of \$131,000, we estimate that the merger led to a \$6.35 increase in monthly payments which translates into \$381 over the five years of the contract.

²¹See Huynh et al. (2011) for analysis of the bootstrap performance in this context.

TABLE 3. Effect of the merger on transaction interest rates

VARIABLES	(1)	(2)	(3)	(4)	(5)
Merger	0.0556 ^a (0.0174)	0.142 ^a (0.0444)	-0.0389 (0.0444)	0.0779 ^a (0.0199)	0.0874 ^a (0.0216)
Merger X Indirect		-0.107 ^b (0.0470)			
Merger X A+B branch share			0.651 ^a (0.239)		
Merger X 1(Lenders=8)				-0.0365 ^b (0.0158)	
Merger X 1(LTV=95)					-0.0745 ^b (0.0310)
Constant	7.933 ^a (0.139)	8.222 ^a (0.143)	7.948 ^a (0.143)	7.897 ^a (0.140)	7.947 ^a (0.139)
Observations	17,139	17,139	17,139	17,139	17,139
R-squared	0.602	0.602	0.602	0.602	0.602

Heteroscedasticity-robust standard-errors are in parenthesis. Regressions include a constant. Significance levels: ^a p-value < 0.01, ^b p-value < 0.05, ^c p-value < 0.1. The dependent variable is the transaction rate. The sample includes all contracts one year before and one year after the merger. Neighborhoods with fewer lenders than the 1st percentile in the treatment group and more than 99th percentile in the control group are excluded (i.e. less than 5 or more than 8). The control variables in columns (2)-(5) include: income, loan-size, loan/income interaction, loan-to-value categories (3), credit score categories (5), residential status (i.e. parents or renter), census characteristics (income, age, population, house-value), and bank and month dummies. Specification (1) uses the same controls as specification (2), but is estimated on the full sample, including local markets with more than eight or less than five lenders. All specifications include bank, week and province fixed effects. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

Columns (2) to (5) show that the effect of the merger is heterogeneous across consumers and markets. Column (2) decomposes the aggregate merger effect into a direct effect and an indirect effect. Consumers transacting with the merged entity experienced the largest price increase: over 14 bps, compared to about 4 bps for consumers selecting a competing institution. Again supposing an average loan size of \$131,000, this works out to an increase of \$890 at the merged entity and just \$254 at competing institutions.

In column (3) and (4) we study how the effect depends the structure of each consumer's choice-set. We find that the average treatment effect is an increasing function of the combined number of branches of the two institutions. Specifically, rates increased by 30 basis points in markets where the merged entity controlled 45% of the branches (95th percentile). Over five years this works out to almost \$1908 in increased payments. Although we restrict our sample to local markets with intermediate levels of concentration, the results in column (4) suggest that the effect of the merger is significantly smaller with eight lenders. The average effect in neighborhoods with eight lenders

TABLE 4. Average treatment effect of the merger on transaction interest rates for different estimators and sub-populations

ESTIMATORS	(1) Aggregate	(2) Direct	(3) Indirect	(4) LTV=95%	(5) LTV<95%	(6) N<8
OLS DID	0.056 ^a (0.017)	0.142 ^a (0.044)	0.035 ^c (0.019)	0.013 (0.025)	0.087 ^a (0.022)	0.078 ^a (0.020)
Matching DID	0.055 ^c (0.033)	0.121 (0.081)	0.035 (0.037)	0.008 (0.051)	0.091 ^b (0.042)	0.083 ^b (0.036)
Change-in-change	0.064 ^a (0.018)	0.116 ^b (0.056)	0.047 ^b (0.019)	0.021 (0.027)	0.099 ^a (0.025)	0.098 ^a (0.026)
Nb. Observations	17,139	3,141	13,998	7,218	9,921	8,352
% Treated	0.62	0.73	0.60	0.61	0.62	0.48

Point estimates correspond to the average treatment effect on treated calculated using OLS (row 1), propensity score matching (row 2), and the Change-in-Change estimator (row 3). Matching and CIC standard-errors are obtained by bootstrapping with 1,000 replications. Significance levels: ^a p-value<0.01, ^b p-value < 0.05, ^c p-value < 0.1. The dependent variable is the transaction rate. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1st percentile in the treatment group and more than the 99th percentile in the control group are excluded (i.e. less than 5 or more than 8). The control variables include: income, loan-size, loan/income interaction, loan-to-value categories (3), credit score categories (5), residential status (i.e. parents or renter), and bank and month dummies. All specifications include bank, week and province fixed effects. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

is equal to 4.14 basis point (statistically different from zero), compare to 7.79 in neighborhoods with less than eight lenders.

The last specification in column (5) suggests that less financially constrained buyers experience larger rate increases as a result of the merger. The point estimate for consumers putting less than 75% downpayment is 8.74 basis points, and is not statistically different from zero for borrowers investing the smallest possible amount (i.e. 5%).

In Table 4 we study the robustness of our main results to alternative estimation methods. The first row corresponds to the OLS estimates calculated from Table 3, and the second and third rows correspond to the average treatment effect estimated with the matching and change-in-change (CIC) estimators using alternative subsamples (i.e. across columns). Overall all three estimation methods yield very similar conclusions. The change-in-change estimator produces larger aggregate effects than the other two (i.e. 6.4 bps compared to 5.6 or 5.5), but all three yield overlapping confidence intervals. The estimates are also very similar across different groups of consumers (i.e. columns). The matching estimator is less precise in general, and some of the estimates are not significantly different from zero at standard levels as the sample shrinks. The results suggest that matching the sample based on a common range of number of lenders, as in Table 3, corrects for most, if not all, of the bias generated by systematic differences across control and treatment




FIGURE 4. Distribution of the average treatment effect of the merger on transaction interest rates

groups. Moreover, the CIC estimator relaxes the linear functional form, without significantly affecting our estimate of the average treatment effect. Our conclusions therefore remain unchanged: $A + B$ consumers were more affected by the merger than competing institutions' consumers, financially un-constrained consumers experienced significant price increases (between 8.3 and 9.8 bps), and the rate increase is more important in markets with fewer lenders.

5.2. Distribution of treatment effect. In this section, we analyze the distribution of treatment effects across consumers' unobserved types. Low quantiles are borrowers who received a relatively big discount, while high quantiles are borrowers paying high rates. Recall that we control for a rich set of financial characteristics, and therefore what we estimate is the treatment at different quantiles of the conditional rate distribution.

Figure 4 presents the estimated treatment effect at all percentiles between 10% and 90% of the conditional rate distribution.²² The aggregate effect becomes statistically indistinguishable from zero starting at the 70th percentile, and peaks around the 62th percentile. Over the reported range, we estimate that the merger raised interest rates between 6 bp and 8.5 bp for consumers below

²²We chose to drop the very top and bottom percentiles from the figure since the estimates are not very stable beyond the 90% percentile. This is likely due to the fact that we need to trim the sample in order to meet the overlapping support assumption of the estimator, and as a result are left with fewer observations outside of the 10-90 range. It is also likely that the CIC estimator is more sensitive at the top, since the empirical price distribution is not invertible for consumers paying the posted price (i.e. violates the monotonicity assumption).

TABLE 5. Distribution of treatment effects of the merger on transaction interest rates

GROUPS	Nb. Obs.	% Treated	Distribution of treatment effects				
			20%	40%	50%	60%	80%
Aggregate	17,139	0.62	0.065 ^b	0.067 ^b	0.072 ^b	0.084 ^b	0.033
Direct	3,150	0.74	0.108	0.147 ^b	0.167 ^b	0.188 ^b	0.124
Indirect	13,989	0.60	0.060 ^b	0.054 ^b	0.059 ^b	0.062 ^b	-0.029
LTV= 95%	7,209	0.62	0.030	0.030	0.041	0.056	0.024
LTV< 0.95%	9,930	0.63	0.096 ^b	0.091 ^b	0.088 ^b	0.094 ^b	0.052
N<8	8,301	0.49	0.062 ^b	0.077 ^b	0.085 ^b	0.103 ^b	0.079

The treatment effects at different percentiles of the distribution are calculated using the Change-in-Change estimator (see text). The dependent variable is the transaction rate. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). Confidence intervals were calculated by bootstrapping the sample 1,000 times. A *b* superscript is added to indicate estimates that are significantly different from zero at the 5% level. All specifications include bank, week and province fixed effects.

the 70th percentile, and by an average of 3.8 over the next group of consumers (not statistically significantly from zero).

In Table 5 we present estimates of the treatment effects at different percentiles of the distribution for several specifications in addition to the aggregate effect. Results are calculated using the change-in-change estimator, and we report the treatment effect at six percentiles between 20 and 80. As in Figure 4 we find that the effect of the merger ranges from around 6.5 to 8.4 bps for borrowers in the 20th to 60th percentiles of the conditional transaction rate distribution, but is statistically indistinguishable from zero in the 80th percentile. In rows (2) through (5) we decompose the aggregate result, and again study the effect of the merger on different sub-populations. The results show that the same non-linear pattern that we observed in the aggregate sample emerges for the direct and indirect effects of the merger. These are measured to be around 15 and 6 bps respectively, in the lower and middle percentiles, and not significant in the top 30%. The decrease is more important in the competing banks subsample. Similarly, we measure the effect of the merger to be about 9 bps for less financially constrained borrowers in the lower and middle percentiles, and zero in the top part of the distribution.

Importantly, the shape of the quantile treatment effect is the same across the five subsamples reported in the table: the effect is linear or slightly increasing for consumers below the 70th percentile, and then decreases sharply by about half and becomes statistically indistinguishable from zero at the top of the conditional price distribution.

It should be noted that for these results to be meaningful, the mix of consumer types must be constant across sub-groups. Specifically, the distribution of unobservables across sub-populations must be the same over time, which is only the case if the merger does not lead borrowers with different unobservables to sort themselves endogenously into different sub-populations. Otherwise, Assumption 2 would be violated.

TABLE 6. Sorting effects of the merger

VARIABLES	(1) Pr(A)	(2) Pr(A or B)	(3) Broker	(4) Broker
Merger	0.0117 (0.0120)	-0.0243 ^c (0.0129)	0.0404 ^a (0.0148)	
Merger X Direct				-0.145 ^a (0.0286)
Merger X Indirect				0.0763 ^a (0.0171)
Constant	0.315 ^a (0.0962)	0.362 ^a (0.0992)	0.0369 (0.109)	0.0939 (0.108)
Observations	17,139	17,139	12,233	12,233
R-squared	0.058	0.079	0.347	0.354

Heteroscedasticity-robust standard-errors are in parenthesis. Significance levels: ^a p-value < 0.01, ^b p-value < 0.05, ^c p-value < 0.1. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables. The control variables include: bank, month and province fixed effects. Regressions in columns (3) and (4) exclude observations with missing broker information.

5.3. Sorting effects of the merger. Part of the heterogeneity in price effects of the merger is associated with the endogenous sorting of consumers. To investigate this we study the effect of the merger on outcomes other than rates. Specifically, we consider the impact of the merger on the probability of choosing one of the merging lenders, or of using a broker.²³ The results are presented in Table 6. In column (1) we restrict attention to the probability of choosing Bank *A*. We find no significant effect on *A*'s market share despite the fact that it has absorbed *B* following the merger. Looking at their combined share we see that it decreased by 2.43% (see column (2)). This decrease is only statistically different from zero at the 10% level. Part of this can be explained by the fact that transaction rates increased more among *A* + *B* consumers, than among competing lenders.

We also find that borrowers transacting with the merged entity are fifteen percent less likely to do so through a broker. In comparison, competing institutions increased their share of broker transactions by 7.6%. In the aggregate sample (i.e. first row), the probability of using a broker increased by about 4%. The asymmetry between the direct and indirect effects is not surprising, since Trust companies are more likely to deal with brokers than national banks. Moreover, since consumers dealing with brokers tend to receive larger discounts, this sorting effect tends to bias upward our estimate of the effect on rates for *A* or *B* consumers, and bias downwards our estimate

²³We have also tested for the presence of merger effects on other financial characteristics such as loan size, risk type and down-payment, and failed to find any statistically significant effects. These additional results are available upon request.

for other lenders' clients. Overall, since the net effect on broker is positive, controlling for a broker indicator variable in our main regression specification would increase the aggregate point estimate in Table 3 by 0.5 basis points (unreported regression).

5.4. Robustness. In this subsection we analyze the robustness of our results to different sample assumptions and treatment-group definitions. We present robustness results with respect to the inclusion of provincial trends, the size of local neighborhoods, and different event windows.

First, we reproduce all specifications controlling for provincial linear trend variables. These trends are aimed at capturing unobserved factors that are common to a given market, and serially correlated. If present, these factors would bias our results, since most of the borrowers affected by the merger are concentrated in a small number of provinces. For some specifications, adding these variables puts tremendous strain on the data. This is especially true for specifications estimating the effect of the merger across aggregate groups of consumers (such as our measure of $A + B$ concentration, or the presence of brokers). Despite this, nearly all the results are unchanged, both qualitatively and statistically.

Tables 9 to 12 in Appendix B reproduce the results of Tables 3 to 6 with the addition of provincial trends. In general, the average treatment effects with provincial trends are larger, and slightly less precisely estimated. The aggregate impact of the merger is estimated to be 6.3 bps compared to 5.6 bps in the main linear regression and 7.2 bps using the change-in-change estimator compared to 6.4 bps.

One area where the provincial trends do have an impact is on the matching DID estimates. The point estimates of the average treatment in the different subsamples are systematically larger than their non-trends counterparts. It is not clear why we obtain this result. One possible interpretation is that the first-stage of the matching estimator is sensitive to the inclusion of aggregate controls such as the province-level trends since a set of provinces are over-represented in the treatment group and the observables are "conditioned out" using only the control subsample. This is not the case in the OLS and CIC estimators. Also, including a provincial trend reduces the importance of the merger on broker selection.

Second, we consider different neighborhoods sizes. Recall that we assumed that consumers shop within a 5 KM 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 areas that are too large (falsely treated). Using a smaller threshold reduces the number of treated neighborhoods, and defines as control neighborhoods those that are affected by the merger. Therefore, over-estimating or under-estimating the size of shopping areas should bias our results towards zero.

The results varying the role of neighborhood size are presented in Panel A of Table 13 of the Appendix. The estimation results confirm our intuition about local markets. The effect of the merger is statistically different from zero for distance radii between 3 and 6 KMs in Panel A. Below and above these levels the point estimates decrease towards to zero, and we cannot reject

the null hypothesis of no merger effect. Panel B shows that controlling for provincial trends reduces the precision of the estimates, but that the treatment effect is still statistically significant for neighborhoods between 3 and 5 KMs.

It should be noted that varying the size of consumer neighborhoods does not address the possibility of an attenuation bias in the estimation of α , caused by unobserved heterogeneity in consumers' choice-sets. As we discussed in Section 4.2.1, this source of heterogeneity would introduce measurement error in the treatment variable. Without an instrumental variable that would be independent of our merger indicator variable, we must interpret our estimate as a possible lower bound of the true causal effect of the merger.

Finally, in Table 7 we consider different event windows, and present the results of a falsification exercise. Recall that throughout the paper we used contracts signed at most one year before or after the official merger date. This is a natural choice since the merged entity started closing duplicate branches about a year after the merger. Our estimates therefore capture the effect of the merger holding fixed (roughly) the distribution of branches. In the first three columns we compare our main specification (column 2), with a shorter window of six months (column 1) and a longer window of one year and a half (column 3). In the last two columns we also estimate the model assuming two alternative merger dates: 6 months before and 6 months after the actual date. For both exercises, we present the results with and without controlling for province-level trends.

Without controlling for the trend variables, the results in Panel A suggest that the effect of the merger using a six month window is about the same as the one we get using the one year window (i.e. 7 versus 6 bps), but expanding our sample to cover an eighteen-month period reduces the point estimate to two basis points. Similarly, the two falsification specifications in columns (4) and (5) suggest that the effect of the merger is positive and statistically significant at a 1% level only using the actual merger date. The point estimate is lower and less precise using the "6 months earlier" specification, and negative and significantly different from zero using the "6 months after" specification.

If nothing else happened during the eighteen-month extended period, these two sets of results would suggest that that the merger led to a short-run price increase and a longer-run price decrease due to efficiency gains; similar to the results of Focarelli and Panetta (2003) for the Italian banking market. A more likely interpretation, however, is that the treated markets were affected by a common omitted variable. Both the larger time window and the falsification exercise suggests that something happened towards the end of the one year period and during the following year, such that interest rates went down in the treated markets relative to the control markets. If this event occurred during our main sample period, it would bias our results against finding any price increase. Likely candidates for such variables include shocks affecting the evolution of housing markets, which are correlated across space and over time.

To verify this, we also present the same robustness results controlling for province-level trends, aimed at capturing other events occurring predominantly in the treated markets. The results in Panel B confirm the second interpretation, and clearly show that the merger created a discontinuity in the evolution of prices around the merger date. Conditioning on the trends, we find

TABLE 7. Effect of the merger on transaction prices for alternative period windows and merger dates

VARIABLES	(1) +/- 6 Months	(2) +/- 1 Year	(3) +/- 1.5 Year	(4) - 6 Months	(5) + 6 Months
Panel A: Without province-level trends					
Merger	0.0735 ^a (0.0266)	0.0556 ^a (0.0174)	0.0219 (0.0147)	0.0430 ^b (0.0168)	-0.0499 ^a (0.0173)
Observations	8,352	17,139	25,665	17,968	16,986
R-squared	0.445	0.602	0.582	0.617	0.494
Panel A: With province-level trends					
Merger	0.0302 (0.0365)	0.0627 ^a (0.0230)	0.0595 ^a (0.0182)	0.00661 (0.0207)	-0.0122 (0.0214)
Observations	8,352	17,139	25,665	17,968	16,986
R-squared	0.447	0.602	0.583	0.618	0.497

Heteroscedasticity-robust standard-errors are in parenthesis. Significance levels: ^a p-value<0.01, ^b p-value < 0.05, ^c p-value < 0.1. The dependent variable is the transaction rate. All specifications in Panel A and B include bank, week and province fixed effects. In addition, Panel B includes province-level trends. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). The control variables include: income, loan-size, loan/income interaction, loan-to-value categories (3), credit score category (5), residential status (i.e. parents or renter), bank, month, province dummies, and provincial trends. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

that changing the period window does not significantly affect the point estimate (i.e. first three columns of panel B).²⁴ In addition, the two falsification specifications strongly reject the presence of a discontinuity six months prior or six months after the realized merger. Indeed, the positive and negative coefficients found in columns (4) and (5) of Panel A are almost entirely explained by the inclusion of the regional linear trend variables. We view these results as a confirmation that the merger raised prices significantly around the merger date in the treated markets.

6. DISCUSSION

6.1. Theory model. The effect of the merger can be understood through the lens of a simple two-stage negotiation model. In the initial stage, each borrower i is matched with a lender j and negotiates a mortgage interest rate p_{ij} . If negotiation fails with j , consumer i explores other lenders present in his/her neighborhood. Assuming that lenders have symmetric costs c_i and earn zero

²⁴Notice that the merger effect in column (1) of Panel B is not statistically significantly different from zero. This is due to a lack of variation in the data when we use the shorter window. In other words, since most of the treated observations are concentrated in a few provinces, it is hard to distinguish between province-specific unobserved linear trends, and our dummy variable capturing the event of the merger. The larger windows allows the two variables to be more easily distinguished.

profits if the negotiation fails, the outcome of the negotiation in the first stage depends on the consumer's outside option generated in the second stage. We assume that information is symmetric across both parties, and use a Nash bargaining solution concept to approximate the outcome of the initial stage negotiation.

Let $\theta_j - p_{ij}$ denote the net consumer surplus of transacting with bank j if the interest rate offer is p_{ij} , and $\bar{\theta} - p^0(c_i, u_i)$ the outside option of consumer i . The value of the outside option is determined by the level of competition in the market, and the ability of consumers to gather competitive quotes. Rather than explicitly modelling this process, we assume that the outcome of the competition stage can be expressed as a reduced form function of the lending cost c_i and an index of consumers' negotiation skills u_i .²⁵ In particular, competition between lenders creates a constant margin over c_i which varies across consumers according to u_i :

$$p^0(c_i, u'_i) = c_i + \mu(u'_i) > c_i + \mu(u_i) = p^0(c_i, u_i) \Leftrightarrow u'_i < u_i.$$

The value of u_i summarizes the negotiation skills or search cost of consumers, along with the value of rental opportunities for consumer i . Under this interpretation, consumers with high u_i expect to receive offers that approach the marginal cost of lending in the second stage. Similarly, consumers with poor negotiation skills (i.e. low u_i) expect to receive offers that reflect the minimum of the common posted interest rate (maximum quote), and their willingness to pay for renting a house. Formally, we assume that the markup function is weakly decreasing in u_i , converges to zero for large u_i , and to a positive constant $\bar{\mu}$ for consumers with low bargaining abilities (those with $u_i < \bar{u}$). Importantly, those consumers obtain a discount that is independent of market structure.

The value of the consumer's outside option, net of the relative willingness to pay term ($\Delta_j = \theta_j - \bar{\theta}$), determines the level of the transaction surplus that is split between lender j and borrower i . Assuming equal bargaining weights, the solution to the Nash bargaining problem takes the following linear form:²⁶

$$(9) \quad p_{ij} = \begin{cases} c_i + \frac{1}{2} (\Delta_j + \mu(u_i)) & \text{if } u_i > \bar{u} \\ c_i + \frac{1}{2} (\Delta_j + \bar{\mu}) & \text{if } u_i \leq \bar{u}. \end{cases}$$

In this framework, if the market is not perfectly competitive, the effect of a merger can be thought of as increasing $\mu(u_i)$ for consumers whose outside option consists of gathering additional quotes. This is because a merger between national financial institutions leads to a reduction in the number of potential bargaining partners since local branch administrators who are now part of the new extended network will stop responding to each other's offers. This market power effect of the merger can be thought of as changing the shape of the markup function from $\mu(u_i; M_i = 0)$ to $\mu(u_i; M_i = 1)$, where M_i is an indicator variable equal to one in the post-merger period. A merger that leads to an increase in market power implies an increase in the average transaction

²⁵In Allen et al. (2011b) we provide a more complete model of negotiation in which the second stage competition is explicitly modelled, and where there is asymmetric information between lenders and borrowers.

²⁶The Nash-Bargaining problem is given by: $\max_{p_{ij}} (\Delta\theta_j + p^0(c_i, u_i) - p_{ij})^{1/2} (p_{ij} - c_i)^{1/2}$.

interest rate:

$$(10) \quad \alpha = \frac{1}{2} E [\mu(u_i; M_i = 1) - \mu(u_i; M_i = 0)] > 0$$

Importantly, this market power effect of the merger should vary across consumers with different u_i . While it is conceivable that consumers with good negotiation skills are less affected (i.e. $p_i^0 \approx c_i$), the Nash bargaining model predicts that consumers with poor negotiation skills should be unaffected since they are offered a rate that does not reflect the competitiveness of their local market (i.e. $\mu(u_i) = \bar{\mu}$). Formally, the distribution of the treatment effect across u_i takes the following form:

$$(11) \quad \alpha(u_i) = \begin{cases} \frac{1}{2} E (\mu(u_i; M_i = 1) - \mu(u_i; M_i = 0)) & \text{If } u_i > \bar{u}, \\ 0 & \text{Otherwise.} \end{cases}$$

Another, more indirect, effect of the merger is to change the retail network configuration of the merging parties. In Canada, nearly 70% of consumers combine their mortgage and day-to-day banking services with the same financial institution. To the extent that consumers value the complementary services offered by their mortgage lenders, a merger can increase the willingness to pay of consumers for the new entity, while decreasing the relative value of banking with a competing institution. Within the model, this quality increase changes the relative willingness to pay term, Δ_j , in favor of the merging firms. If this is the case, the average treatment effect of the merger will be different for AB consumers (direct effect) and competing banks' consumers (indirect effect):

$$(12) \quad \begin{aligned} \alpha_{AB} &= \frac{1}{2} E [\Delta_j^1 - \Delta_j^0 + \mu(u_i; M = 1) - \mu(u_i; M = 0) | j_i = A, B] \\ &> \frac{1}{2} E [\Delta_j^1 - \Delta_j^0 + \mu(u_i; M = 1) - \mu(u_i; M = 0) | j_i \neq A, B] = \alpha_{\text{Other}}. \end{aligned}$$

Note that from the perspective of borrowers transacting with a competing lender, the effect of the merger is to improve the quality of their outside option and therefore put downward pressure on the rate quote from their lenders. A similar asymmetry can be interpreted as efficiency gains resulting from the merger. In the current setting, quality improvements and firm-specific cost reductions are isomorphic.

It should also be pointed out that a similar asymmetry between lenders can be caused by a sorting effect: the distribution of consumer types within A or B pre-merger is different than within AB post-merger. For instance, it is possible that consumers dealing with trust companies are more price sensitive (i.e. higher u_i) than consumers dealing with national banks. Unless we observe u_i , we cannot distinguish between a quality improvement and a change in the mix of consumer types across lenders. In the empirical analysis, we use proxy variables for u_i to empirically distinguish between these two sources for heterogeneous treatment effects. For instance, we estimate the effect of the merger on the probability of using a broker.

6.2. Interpretation of the results. The average treatment effect of the merger provides us with some sense of the degree of market power in this industry. Our finding that losing a lender option

has a positive effect on rates is consistent with the idea that the market is not perfectly competitive since if it were, the loss of a lender would not raise prices significantly for the average borrower. We also find that although the average effect is not that large especially compared with the average monthly payment of consumers, the effect of losing a lending option is quite pronounced in local markets in which the merging parties were dominant pre-merger (i.e. up to 30 basis points).

These results, coupled with the fact that for the average borrower margins over the five-year bond rate are quite low (about 108 basis points), lead us to conclude that many local markets are fairly competitive. Several institutional features are consistent with 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.

Although our findings are consistent with the interpretation that firms operate under tight margins, we must also point out that our regression analysis methodology provides a lower bound evaluation of $\bar{\alpha}$, and therefore on the degree of market power. This is because, the difference-in-difference methodology imposes strong assumptions regarding the consideration sets of consumers. The likely presence of heterogeneity in the sets of lenders considered by consumers translates into measurement error in the treatment variable, which attenuates our estimate of the merger effect.

Our results also suggest that there are important differences in the effect of the merger for different percentiles of the rate distribution. In our analysis we condition on a rich set of characteristics, and so the percentiles of the conditional rate distribution can be interpreted as identifying different consumers' unobserved negotiation skills indexed by u_i in the Nash bargaining model above. This interpretation is exact if we assume that u_i 's are independent of consumer attributes, and that the common lending cost is accurately measured by the observed characteristics of each transactions. Under this assumption, the equilibrium pricing function corresponds to our estimating equation defined in equation 4:

$$(13) \quad p_i = c_i + \frac{1}{2}\mu(u_i; M_i, T_i) = Z_i\gamma + \lambda\mathbf{Week}_i + \frac{1}{2}\mu(u_i, M_i, T_i).$$

This formulation allows us to interpret the distribution of the treatment effect estimated with the Change-in-Change estimator as tracing-out the distribution of the change in the markup function defined in equation 11.

Our findings show that only consumers gathering multiple quotes are adversely affected by the merger. In contrast, consumers who are unable or unwilling to negotiate pay a price that is strictly a function of their willingness to pay and/or the common posted interest rate, which is not directly affected by the loss of a lending option. We conclude that lenders price discriminate between borrowers, exercising market power vis-a-vis those with poor negotiation ability or high search costs.

These findings also give a measure of the fraction of consumers unable to haggle (i.e. \bar{u}). Recent surveys show that about 45% of consumers gather only one price quote when shopping for their first mortgage. Assuming that these consumers are at the top of the price distribution, our results suggest that two-thirds of consumers who accept the initial quote do so because they are unwilling or unable to gather competing offers.²⁷ However, this does not mean that these consumers pay the posted interest rate, since we estimate that only 10% of consumers pay a rate within 10 bps of the posted price.

Interestingly, these results are consistent with those of Sapienza (2002) who, in the context of business lending, finds that those borrowers with either many or very few outside options for loans (as measured by the number of other banking relationships they have) are unaffected by the merger. Although the methods for identification in her paper and ours are different, the interpretation of the economic channel through which the results are derived is the same. Those borrowers that are either in a strong or weak bargaining position are not affected when there is one less lender with whom to negotiate.

A related interpretation is that u_i proxies for the number of banks at which consumer i would qualify for a loan, or at least his/her perception of this number. If banks use different lending rules, the presence of unobserved risk factors across borrowers could generate similar heterogeneous treatment effects. Some of our results are consistent with this interpretation. In particular, we find that consumers putting the minimum payment down are unaffected by the merger, while the median financially unconstrained borrower experienced a 9 bp increase. However, it is unlikely that unobserved risk characteristics fully explain the results, since the insurance program tends to minimize the importance of risk considerations in the lending decision (i.e. the government uses common criteria across all lenders). In addition, we document the same non-linearity in the treatment effects across quantiles in the subsample of non-financially constrained households (see Table 5 row 5).

Finally, we also find that the merger had an asymmetric effect on consumers transacting with the merged entity and on consumers transacting with competing financial institutions. We find that borrowers contracting with the merged entity are significantly more affected by the merger than borrowers contracting with rival lenders. We associate this result with the fact that the merger increased the relative willingness to pay of consumers for the merged entity. As we discussed above, this is consistent with the fact that the merger led to an expansion in the network of bank AB , which improved the relative quality of its financial services.

Again, it is important to point out that a complementary interpretation is that the merged entity began serving a different mix of consumers after the merger. Indeed, part of the direct price increase that we document is due to the fact that prior to the merger Trust B consumers were more likely to obtain large discounts (conditional on their Z s, i.e., observable characteristics).²⁸ These

²⁷The fraction of non-searchers with positive price response to the merger is:

$$33\% = [70\% (\text{positive price effect}) - 55\% (\text{searchers})] / 45\% (\text{non-searchers}).$$

²⁸Excluding B consumers from the sample reduces the direct effect in column (3) of Table 3 to 9 bps.

price sensitive types are more likely to transact with competing institutions post-merger, as evidenced by the fact that those transacting with the merged entity are less likely to have transacted with brokers. This could either be because the merged entity is not offering competitive rates to brokers, or because its more dominant position in the market allows it to refuse to deal with brokers altogether.²⁹ Borrowers who transact with the merged entity post merger are those with a high valuation for the range of services or the network size of their mortgage lender.

7. CONCLUSION

In contrast to most of the literature studying the effects of horizontal mergers that focuses on posted prices, we take advantage of transaction level data 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 heterogeneous treatment effects, and estimates the distribution of treatment effects across unobserved consumer types.

We find that the average effect of the merger yields a small but statistically significant increase in mortgage interest rates, suggesting that the average consumer is able to extract a large share of the transaction surplus through search and negotiation even when there is one fewer lending option. However, this finding masks important heterogeneity. Some borrowers pay significantly higher rates following the merger, while others are barely affected. Although in some cases the higher rates may be associated with higher market power resulting from the decrease in the number of lender options, many of those paying higher rates are contracting with the merged entity and may be benefiting from the improved level of services offered. Even more importantly, 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 very high search costs/bargaining abilities, are not affected by the merger, while those lower in the price distribution are. Therefore our results imply that lender competition appears to have no impact on rates in the top part of the 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 instance, 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, Geraldi et al. (2010) argue that financial illiteracy played an important role in the rate of foreclosures in the U.S. housing crisis.

²⁹At least one national bank has at times explicitly refused to deal with brokers.

In Canada, these borrowers could 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 (2010) there is considerable confusion surrounding the mortgage process generally, and the role of brokers in particular.³⁰ 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

- Adams, W., L. Einav, and J. Levin (2009). Liquidity constraints and imperfect information in subprime lending. *American Economic Review* 99, 49–84.
- Allen, J., R. Clark, and J. F. Houde (2008). Market structure and the diffusion of e-commerce: Evidence from the retail banking industry. Bank of Canada Working Paper No. 2008-32.
- Allen, J., R. Clark, and J. F. Houde (2011a). Discounting in mortgage markets. Bank of Canada Working Paper No. 2011-3.
- Allen, J., R. Clark, and J. F. Houde (2011b). Price negotiation in differentiated product markets: The case of insured mortgages. mimeo.
- 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.
- Bertrand, A. and A. Morse (2011). Information disclosure, cognitive biases, and payday borrowing. *Journal of Finance* 66, 1865–1894.
- Breslaw, J., I. Irvine, and A. Rahman (1996). Instrument choice: The demand for mortgages in Canada. *Journal of Urban Economics* 39, 282–302.
- 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).
- Competition Bureau (2003). The merger enforcement guidelines as applied to a bank merger. Technical report.
- Dafny, L., M. Duggan, and S. Ramanarayanan (2011). Paying a premium on your premium? consolidation in the health insurance industry. forthcoming *American Economic Review*.

³⁰The Wall Street Journal (Hagerty (2007)), for example, highlights the debate in the U.S. over mortgage brokers – most consumers believe that brokers represent them, however, as the president of the Colorado Mortgage Lenders Association states “The mortgage broker does not represent the borrower.”

- Dahlby, B. and D. West (1986). Price dispersion in an automobile insurance market. *Journal of Political Economy* 94, 418–438.
- Focarelli, D. and F. Panetta (2003). Are mergers beneficial to consumers? Evidence from the market for bank deposits. *American Economic Review* 93, 1152–1172.
- Geraldi, K., L. Goette, and S. Meier (2010, April). Financial literacy and subprime mortgage delinquency: Evidence from a survey matched to administrative data. Technical Report 2010-10, Federal Reserve of Atlanta.
- 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.
- Hagerty, J. (2007). Mortgage brokers: Friends or foes. *Wall Street Journal*.
- Hall, R. E. and S. Woodward (2010). Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. NBER working paper 16007, forthcoming *American Economic Review*.
- 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.
- Hastings, J. and R. Gilbert (2005). Market power, vertical integration and wholesale price of gasoline. *Journal of Industrial Economics* 53, 469–492.
- Heckman, J., H. Ichimura, and P. Todd (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a training program. *Review of Economic Studies* 64(4).
- Heckman, J., H. Ichimura, and P. Todd (1998, September). Matching as an econometric evaluation estimator. *Econometrica* 66.
- Hendel, I., A. Nevo, and F. Ortalo-Magne (2010). The relative performance of real estate marketing platforms: Mls versus fsbomadison.com. *The American Economic Review*.
- Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual 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. (2011). Spatial differentiation and vertical mergers in retail markets for gasoline. Forthcoming, *American Economic Review*.
- 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*.
- KPMG (2008). Canada mortgage bonds program evaluation. Audit and Evaluation Services - Final Report.
- 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.
- Panetta, F., F. Schivardi, and M. Shum (2009). Do mergers improve information? evidence from the loan market. *Journal of Money, Credit and Banking* 41(4), 673–709.

- 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.
- Traclet, V. (2005). Structure of the canadian housing markets and finance system. Bank of Canada, at [http : //www.bis.org/publ/cgfs26cbpapers.htm](http://www.bis.org/publ/cgfs26cbpapers.htm).

APPENDIX 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. 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. 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 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.

APPENDIX B. ADDITIONAL TABLES

TABLE 8. Definition of Household / Mortgage Characteristics

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
Applicant type	Quartile of the borrowers risk of default
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
<i>CREDIT</i>	Summarized application credit score (minimum borrower credit score).

Some variables were only included by one of the mortgage insurers.

TABLE 9. Effect of the merger on transaction interest rates with provincial trends

VARIABLES	(1)	(2)	(3)	(4)	(5)
Merger	0.0627 ^a (0.0230)	0.138 ^a (0.0464)	-0.0140 (0.0518)	0.0827 ^a (0.0251)	0.0943 ^a (0.0264)
Merger X Indirect		-0.0956 ^b (0.0469)			
Merger X A+B branch share			0.502 ^c (0.277)		
Merger X 1(Lenders=8)				-0.0332 ^b (0.0161)	
Merger X 1(LTV=95)					-0.0755 ^b (0.0310)
Constant	7.766 ^a (0.209)	8.054 ^a (0.211)	7.803 ^a (0.214)	7.725 ^a (0.210)	7.778 ^a (0.209)
Observations	17,139	17,139	17,139	17,139	17,139
R-squared	0.602	0.603	0.602	0.602	0.603

Replication of Table 3 with province-level trends. Heteroscedasticity-robust standard-errors are in parenthesis. Significance levels: ^a p-value<0.01, ^b p-value < 0.05, ^c p-value < 0.1. The dependent variable is the transaction rate. Regressions include a constant. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). The control variables include: income, loan-size, loan/income interaction, loan-to-value categories (3), credit score category (5), residential status (i.e. parents or renter), bank, month, province dummies, and provincial trends. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

TABLE 10. Effect of the merger on transaction interest rates with provincial trends estimated with different estimation methods

ESTIMATORS	(1) Aggregate	(2) Direct	(3) Indirect	(4) LTV=95%	(5) LTV<95%	(6) N<8
OLS DID	0.063 ^a (0.023)	0.138 ^a (0.046)	0.042 ^c (0.024)	0.019 (0.029)	0.094 ^a (0.026)	0.083 ^a (0.025)
Matching DID	0.169 ^a (0.034)	0.267 ^a (0.085)	0.140 ^a (0.037)	0.118 ^b (0.054)	0.209 ^a (0.044)	0.181 ^a (0.037)
Change-in-change	0.072 ^a (0.023)	0.071 (0.067)	0.076 ^a (0.024)	0.006 (0.036)	0.121 ^a (0.034)	0.092 ^a (0.030)
Nb. Observations	17,139	3,141	13,998	7,218	9,921	8,352
% Treated	0.62	0.73	0.60	0.61	0.62	0.48

Replication of Table 4 with province-level trends. Heteroscedasticity-robust standard-errors are in parenthesis. Significance levels: ^a p-value<0.01, ^b p-value < 0.05, ^c p-value < 0.1. The dependent variable is the transaction interest rate. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). The control variables include: income, loan-size, loan/income interaction, loan-to-value categories (3), credit score category (5), residential status (i.e. parents or renter), bank, month, province dummies, and provincial trends. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

TABLE 11. Distribution of treatment effects of the merger on transaction prices with provincial trends

GROUPS	Nb. Obs.	% Treated	Distribution of treatment effects				
			20%	40%	50%	60%	80%
Aggregate	17,139	0.62	0.073 ^b	0.073 ^b	0.076 ^b	0.088 ^b	0.047
Direct	3,141	0.73	0.073	0.105	0.120	0.135	0.089
Indirect	13,998	0.60	0.087 ^b	0.079 ^b	0.081 ^b	0.083 ^b	0.008
LTV= 95%	7,218	0.61	0.015	0.015	0.026	0.037	0.012
LTV< 0.95%	9,921	0.62	0.120 ^b	0.109 ^b	0.104 ^b	0.109 ^b	0.081
N<8	8,352	0.48	0.055	0.068 ^b	0.076 ^b	0.094 ^b	0.079

Replication of Table 5 with province-level trends. The treatment effects at different percentiles of the distribution are calculated using the Change-in-Change estimator (see text). The dependent variable is the transaction rate. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). Confidence intervals were calculated by bootstrapping the sample 1,000 times. A ^b superscript is added to indicate estimates that are significantly different from zero at the 5% level. All specifications include bank, week and province fixed effects as well as provincial trends. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

TABLE 12. Sorting effects of the merger

VARIABLES	(1) Pr(A)	(2) Pr(A or B)	(3) Broker	(4) Broker
Merger	0.00797 (0.0169)	-0.0166 (0.0181)	-0.00404 (0.0180)	
Merger X Direct				-0.170 ^a (0.0296)
Merger X Indirect				0.0339 (0.0211)
Constant	0.0338 (0.163)	0.167 (0.163)	0.0131 (0.143)	0.0158 (0.143)
Observations	17,139	17,139	12,233	12,233

Replication of Table 6 with province-level trends. Heteroscedasticity-robust standard-errors are in parenthesis. Significance levels: ^a p-value < 0.01, ^b p-value < 0.05, ^c p-value < 0.1. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). The control variables include: bank, month, province dummies, and provincial trends. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.

TABLE 13. Effect of the merger on transaction prices for alternative competitive neighborhood assumptions

VARIABLES	(1) R=1	(2) R=2	(3) R=3	(4) R=4	(5) R=5	(6) R=6	(7) R=7	(8) R=8	(9) R=9	(10) R=10
Panel A: Without province-level trends										
Merger	0.0143 (0.0180)	0.0152 (0.0136)	0.0309 ^b (0.0145)	0.0440 ^a (0.0157)	0.0556 ^a (0.0174)	0.0398 ^b (0.0187)	0.0321 (0.0204)	0.0192 (0.0181)	0.0184 (0.0202)	0.0127 (0.0204)
Observations	21,062	24,346	22,625	20,050	17,139	15,376	13,044	22,946	20,613	20,122
R-squared	0.602	0.603	0.602	0.604	0.602	0.603	0.608	0.599	0.600	0.599
Panel B: With province-level trends										
Merger	0.0152 (0.0187)	0.00421 (0.0150)	0.0295 ^c (0.0170)	0.0350 ^c (0.0197)	0.0627 ^a (0.0230)	0.0267 (0.0271)	0.0140 (0.0293)	0.0251 (0.0204)	0.0330 (0.0235)	0.0321 (0.0239)
Observations	21,062	24,346	22,625	20,050	17,139	15,376	13,044	22,946	20,613	20,122
R-squared	0.603	0.604	0.602	0.605	0.602	0.604	0.610	0.600	0.601	0.599
% Treated	0.207	0.425	0.575	0.603	0.620	0.632	0.626	0.805	0.832	0.830

Heteroscedasticity-robust standard-errors are in parenthesis. Significance levels: ^a p-value < 0.01, ^b p-value < 0.05, ^c p-value < 0.1. The dependent variable is the transaction interest rate. The sample includes all contracts one year before and after the merger. Neighborhoods with fewer lenders than the 1% percentile in the treatment group and more than 99% percentile in the control group are excluded (i.e. less than 5 or more than 8). Each column corresponds to a different competitive neighborhood assumption measured in Euclidian distance. The control variables include: income, loan-size, loan/income interaction, loan-to-value categories (3), credit score category (5), residential status (i.e. parents or renter), bank, month, province dummies, and provincial trends. To control for the non-random nature of missing household characteristics we interact a missing value dummy with: treatment group, after merger, and bank indicator variables.