American Economic Review 2013, 104(10): 3365–3396
http://dx.doi.org/10.1257/aer.104.10.3365

The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry†

By Jason Allen, Robert Clark, and Jean-François Houde*

We examine the relationship between concentration and price dispersion using variation induced by a merger in the Canadian mortgage market. Since interest rates are determined through a search and negotiation process, consolidation weakens consumers' bargaining positions. We use reduced-form techniques to estimate the mergers' distributional impact, and show that competition benefits only consumers at the bottom and middle of the transaction price distribution, and that mergers reduce the dispersion of prices. We illustrate that these effects can be explained by the presence of search frictions, and that the average effect of mergers on rates underestimates the increase in market power. (*JEL* G21, G34, K21, L13, L41)

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

These features lead to price dispersion. In Allen, Clark, and Houde (2014c) we document that the Canadian mortgage market exhibits substantial dispersion, even for

* Allen: Bank of Canada, 234 Wellington St, Ottawa, Ontario, K1A0G9 (e-mail: jallen@bankofcanada.ca); Clark: HEC Montréal, 3000, chemin de la Cote-Sainte-Catherine, H3T2A7, CIRANO and CIRPÉE, Montreal, Quebec (e-mail: robert.clark@hec.ca); Houde: Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104 (e-mail: houde@wharton.upenn.edu). This paper previously circulated under the titles: "Price Competition and Concentration in Decentralized Markets: Empirical Evidence a Merger between Canadian Mortgage Lenders" and "The Effect of Mergers in Decentralized Markets: Evidence from the Canadian Mortgage Industry." Jean-François Houde thanks NSF for funding (SES-1024840). We thank the Canada Mortgage and Housing Corporation and Genworth Financial for providing us with the data. We also thank Allan Crawford, Walter Engert, and Ron Morrow for their consistent efforts on this project. We thank seminar participants at numerous institutions. We greatly benefited from discussions with Dan Ackerberg, Jean Boivin, Ulrich Doraszelski, Liran Einav, Matt Grennan, Gautam Gowrisankaran, Joe Harrington, Kim Huynh, Ken Hendricks, Jack Porter, Matt Shum, Katja Seim, Alan Sorensen, Chris Taber, Robert Town, Marcel Voia, and Matt Weinberg. 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. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

† Go to http://dx.doi.org/10.1257/aer.104.10.3365 to visit the article page for additional materials and author disclosure statement(s).

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

contracts with homogeneous terms and for which lenders are fully protected against the risk of default by a government-backed insurance program. The interquartile range of the net transaction interest spread approaches 100 basis points (bps) and the coefficient of variation is about 59 percent. Importantly, more than 60 percent of the dispersion in margins is unexplained by standard borrower and contract characteristics.

Similar phenomena have been documented by a growing literature analyzing price dispersion in retail markets: new cars (Goldberg 1996; Morton, Zettelmeyer and Silva-Risso 2003; Busse, Silva-Risso, and Zettelmeyer 2006; and Langer 2012), mortgage broker fees (Woodward and Hall 2012), real estate (Hendel, Nevo, and Ortalo-Magne 2009), and health care services (Sorensen 2000; Grennan 2013). Building on Stigler (1961), many of these papers have provided evidence that a significant fraction of the observed dispersion in prices is caused by the inability of (some) consumers to gather information and negotiate discounts, suggesting that search and information frictions are important factors in these markets.

In contrast, we know relatively little about the impact of market structure and competition on the distribution of negotiated prices.[2] It remains an empirical question whether the benefit of competition is spread equally across consumers, and how competition affects dispersion of transaction prices. Indeed, search-theoretic models of price dispersion provide ambiguous predictions with respect to the impact of concentration on price dispersion.[3]

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

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

We then estimate the distributional effect of the merger, an aspect of merger evaluation that has been ignored up to now by the literature. To do so, we use the *change-in-change* (CiC) estimator proposed by Athey and Imbens (2006) to recover

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

[3] See Janssen and Moraga-González (2004) for a discussion. The same is true for discrimination-based theories: the "textbook" price discrimination model suggests a negative relationship between dispersion and competition, while Borenstein (1985) and Holmes (1989) show that an increase in competition in markets with third-degree price discrimination leads to more dispersion.

[4] This strategy has recently been used to study the impact of mergers in gasoline markets (e.g., Hastings 2004; and Houde 2012); cement markets (e.g., Hortaçsu and Syverson 2007); health care markets (e.g., Dafny, Duggan, and Ramanarayanan 2012); and durable goods markets (e.g., Ashenfelter, Hosken, and Weinberg 2013).

the counterfactual distribution of negotiated rates that would have been observed if the merger had not been approved by antitrust authorities. This allows us to document substantial heterogeneity in the impact of the merger along the distribution of negotiated rates. In particular, our second set of results shows that the loss of a competitor increased interest rates between 7 and 9 bps for consumers in the lower and middle percentiles of the distribution, and had no effect on consumers in the top 30 percent. This result implies that borrowers who are the most adversely affected by a negotiation-based pricing policy would not benefit from an increase in the degree of competition between lenders.

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

To understand these results we develop a model of mortgage rate determination that is similar to the price negotiation models of Bester (1988) and Wolinsky (1987). Borrowers and lenders negotiate mortgage rates, with the lender taking into account the borrower's search and negotiation ability. This model yields theoretical predictions consistent with our three reduced-form results. Specifically, the presence of search costs reduces the bargaining leverage of consumers by limiting their ability to generate competitive offers, and, therefore, the value of their outside options. As a result consumers with high search costs pay relatively high rates, and face an outside option that is nearly independent of the number lenders in the market. The effect of the merger is felt more strongly among consumers with medium to low search costs, which reduces the importance of price dispersion.

Moreover, the model predicts that the average price effect of a merger can underestimate the underlying increase in market power when search frictions are present. In order to illustrate how large this bias can be, we use the results of the retrospective analysis to estimate the structural parameters of the model. We do so by finding the semi-parametric distribution of search costs that is consistent with the observed post-merger distribution of prices and the counterfactual pre-merger distribution obtained from the reduced-form analysis. This estimation strategy relies on an assumption that the residual distribution of prices is fully explained by the presence of heterogenous search costs across consumers, similar to strategies employed in recent papers estimating equilibrium search models in markets with important price dispersion (e.g., Hortaçsu and Syverson 2004, Hong and Shum 2006, Wildenbeest 2011).

The results show that, in the context of our model, the average price increase of 6 bps corresponds to 50 percent of the actual increase in market power caused by the merger, if all consumers gathered the maximum number of quotes. Moreover, we show that the net effect of mergers is very different in markets with lower search costs. In particular, reducing consumer search costs by half would lead to a 30 percent *increase* in the average effect of the merger, and 46 percent *more* homogeneous effects across consumers. These findings imply that estimates of the average impact of mergers on prices obtained by performing a retrospective merger analysis correspond to a (downward) biased measure of market power when prices are determined through search and negotiation, and therefore highlight the importance of studying the distributional effect of mergers.

This quantitative exercise is based on a stylized model whose objective is to highlight the key factor explaining the reduced-form results; namely heterogeneity in the bargaining leverage of consumers caused by search frictions. In practice, this heterogeneity can also arise from other features of the market, such as product differentiation or asymmetric information. In Allen, Clark, and Houde (2014a) we estimate a model of search and negotiation that uses additional data on consumers' search and switching decisions, and that incorporates multiple sources of residual price dispersion. We use the model to study the welfare cost of search frictions, and measure the ability of firms to price discriminate. Despite the addition of other sources of dispersion, the impact of losing a bargaining partner on the distribution of prices is qualitatively similar to what we document in this paper.

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

The rest of the paper is structured as follows. Section I describes the Canadian mortgage market, focusing on market structure and pricing. Section II presents our sample and identification strategy. Section III presents results of the reduced-form analysis. In Section IV we develop and estimate a model of search and price negotiation. Section V concludes and presents implications of our results for mortgage-market and merger policies. The online Appendix contains additional results.

I. The Canadian Mortgage Market

A. *Market Structure and Mergers*

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

The market was not always this concentrated. Until the early 1990s the Canadian residential mortgage market also featured a large number of trust companies. Trusts are like savings and loan associations in the United States. 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

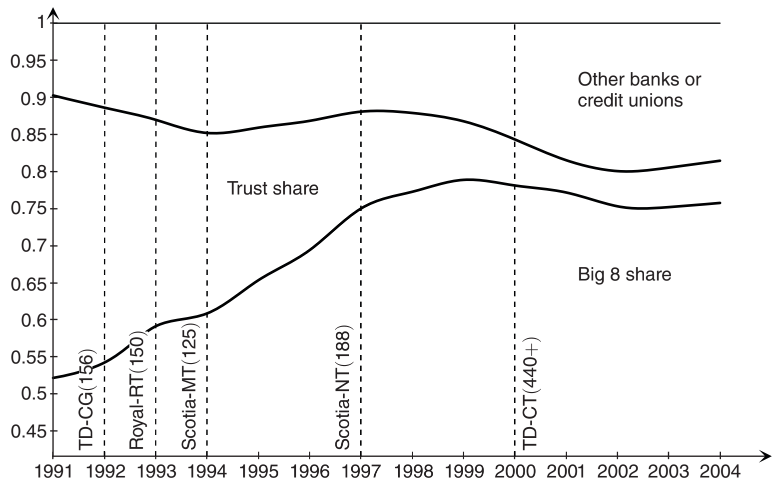

FIGURE 1. EVOLUTION OF FINANCIAL INSTITUTION MARKET SHARES FOR NEWLY INSURED MORTGAGES (*Smoothed*)

competitive advantage in the mortgage market due to a lower cost of funding. Cross-ownership between the two types of institutions was not permitted until the 1992 revisions to the Bank Act. Following these revisions, banks and trusts were granted almost identical powers, making them undifferentiated products from the point of view of consumers.[5]

As a result of the Bank Act revisions and a series of bad residential and commercial loans that created solvency and liquidity issues for the trusts in the 1980s, Canadian chartered banks acquired the majority of trust companies over the following decade. The merger wave led to the six largest banks controlling approximately 80 percent 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 in Canada—The Big 8, Trusts, Credit Unions and other banks—as well as the major mergers, including the number of branches acquired. The major acquisitions were: Canada Trust and Toronto-Dominion, 2000; National Trust and Scotia Bank, 1997; Montreal Trust and Scotia Bank, 1994; Royal Trust and Royal Bank, 1993; and Central Guaranty Trust and Toronto-Dominion, 1993. Online Appendix A describes these mergers.

Our empirical analysis focuses on the short-run impact of one of the major mergers between a bank and a trust. As a result we study contracts signed within a year of the merger. This is appropriate since we know that the merged entity did not start closing duplicate branches until approximately a year after the official merger date. For confidentiality reasons we cannot reveal the parties involved in the merger. We therefore label the two institutions *A* and *B*, and hide the exact timing of the merger.

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

B. *Pricing and Negotiation*

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

Negotiating larger discounts is costly for the local bank manager, reducing the commissions earned by branch employees, but worthwhile if a consumer is likely to switch to another financial institution without a discount.[6] Local branch managers compete against rival banks, but not against other branches of the same bank.[7] Survey evidence from the Canadian Association of Accredited Mortgage Professionals reports that about 52 percent of new home buyers visit more than one lender when shopping for a mortgage.[8]

Alternatively, borrowers can hire brokers to search for the best rates on their behalf. In Canada brokers have fiduciary duties and are compensated by lenders based on volume. According to detailed survey evidence collected by Taddingstone in 2005, brokers on average contact 5.9 lenders for their clients, suggesting they do, in fact, assist in gathering multiple quotes. See Allen, Clark, and Houde (2014c) for a more detailed description of the pricing strategies used in the market.

II. Sample Selection and Identification Strategy

Our objective is to study the effect of the merger on the distribution of negotiated rates. Using the language of the treatment effects literature, we think of the merger as an intervention (I_i) that changes the rates negotiated between borrowers and lenders:

$$Rate_i = I_i f^I(u_i, T_i) + (1 - I_i) f^N(u_i, T_i), \tag{1}$$

where $Rate_i$ denotes the transaction mortgage rate,[9] T_i is a before/after period indicator, and $f_i^I(u_i, T_i)$ and $f_i^N(u_i, T_i)$ correspond to the post- and pre-merger reduced-form pricing functions, respectively. Equation (1) expresses transaction rates as a non-separable function of consumers' unobserved attributes u_i. We chose this formulation to highlight

[6] A recent KPMG consulting report describes the typical compensation of branch manager (KPMG LLP 2008).

[7] Borrowers must present credible quotes to competing branches, and branch managers are explicitly told not to compete against each other. We know from private conversations with loan officers that only in extremely rare occasions will a branch manager deviate from this directive. There is also important record keeping that prevents branch managers from the same institution from making competing offers. First, a pre-approved mortgage is typically good for 90 days, therefore a bank is committed to a consumer for those 90 days and this information is locked into a bank's (and therefore all branches) database. Second, when mortgage insurance is provided the insurers' automated underwriting program flags whether an individual has qualified for insurance. Third, when individuals apply for a mortgage there is a credit check and all lenders can see that this has been done.

[8] The annual reports of the CAAMP are available here: http://www.caamp.org.

[9] We use negotiated rates relative to the five year bond rate to isolate the cross-sectional dispersion in transaction rates. In the empirical analysis we also study the effect of the merger on the probability of receiving a discount off the posted rate. See Section IID for more details.

the possibility that mergers might have heterogeneous effects across consumers with different unobserved ability (or willingness) to negotiate discounts.

We follow a quasi-experimental approach. That is, our objective is to characterize the counterfactual distribution of negotiated interest rates absent the merger among a group of treated consumers ($G_i = 1$), using the observed changes in rates among those unaffected by the merger (i.e., $G_i = 0$). In equation (1), this corresponds to the distribution generated by $f^N(u_i, T_i = 1)$ for consumers directly impacted by the merger. Over the next four subsections we describe the construction of the sample and identification strategy, and provide a set of descriptive statistics.

A. *Mortgage Contract Data*

Our analysis focuses on insured mortgages, and we use administrative data obtained directly from the two insurers operating in Canada: Canada Mortgage and Housing Corporation (CMHC) and Genworth Financial. CMHC is a crown corporation with an explicit backstop from the federal government. Genworth also receives an explicit government of Canada guarantee, but for 90 percent of the amount owing. During our sample period, both insurers used the same strict federal approval guidelines: (i) borrowers with less than 25 percent equity must purchase insurance, and (ii) borrowers with monthly gross debt payments that are more than 32 percent of gross income or a total debt service ratio greater than 40 percent are rejected.[10] The government also sets an insurance premium that is solely a function of the loan-to-value ratio, and ranges from 1.75 to 3.75 percent of the loan. The qualifying rules and premiums are common across lenders and based on the posted rate. Borrowers qualifying at one bank, therefore, should assume that they can qualify at other institutions, given that the lender is protected in case of default.

We construct a 10 percent random sample of all contracts issued between 1992 and 2004. We have access to 20 household/mortgage characteristics, including all financial characteristics of the contract (i.e., rate, loan size, house price, debt ratio, risk type), the lender identity (for the 12 largest lenders), some demographic characteristics (e.g., income, prior relationship with the bank, residential status, dwelling type), as well as the house location up to the forward sortation area (FSA).[11]

B. *Treatment and Control Groups*

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

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

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

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

Bank A is a national bank, and therefore present in nearly all local markets pre-merger. Trust B, on the other hand, is smaller, and isolated from A in only 2 percent of its markets. Given the extent of branch co-location we focus only on the first effect of the merger, namely the reduction in the number of available options. The *treatment group* is therefore defined as the set of consumers who had both lenders in their choice set prior to the merger, while the *control group* is the set of consumers who had only one or none of the merging firms.[12]

To operationalize this definition, we need to formally define consumer choice sets. To do so, we exploit the fact that the pricing decision is decentralized, and therefore that consumers negotiate directly with local branch managers. This allows us to define a consumer's choice set as the set of lenders present in a neighborhood around the house's FSA, denoted by $\mathcal{N}_i$. This neighborhood is defined by measuring the Euclidian distance from each FSA's centroid to the closest branch of each bank. We obtained data on the location of each branch using a yearly panel assembled by Financial Services Canada directory (Micromedia ProQuest). This dataset contains the location of active branches for all financial institutions.

Our main empirical specification uses a fixed radius of 5 kilometers (km) to define neighborhood boundaries. This definition reflects the fact that most consumers choose a lender that has a branch presence near their new house. Indeed, the average distance to chosen lenders is about 2 km, compared to slightly less than 4 km for the average distance to other financial institutions. Moreover, 80 percent of transactions occur with a bank located within 2 km of consumers, and more than 90 percent of transactions are included in a 5 km neighborhood. This provides strong evidence that consumers most often deal with a bank that has a large presence in their region. Notice also that we measure distance as the Euclidian distance, which underestimates the actual driving distance from the center of each FSA to a branch. A 5 km radius is therefore larger than the radii of most FSAs, and broadly corresponds to each consumer's municipality.[13]

C. *Sample Selection*

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

Time Period.—We first select contracts signed within a year of the merger date. We further eliminate contracts for which the closing date is less than 90 days after

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

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

TABLE 1—SUMMARY STATISTICS ON CONSUMER CHOICE SETS PRIOR TO THE MERGER

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

Notes: Each entry corresponds to the sample average and standard deviations (in parentheses), calculated using the observation weights from the mortgage contract dataset. 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* and *B*" have both merging parties.

the merger, to avoid including in the post-merger period rates that were negotiated pre-merger. We use 90 days because in Canada lenders tend to guarantee price quotes for 90 days. This leads to a slightly uneven split of observations before and after the merger: 42 percent of the transactions take place post-merger.

Contract Characteristics.—We select our sample by limiting heterogeneity in contractual characteristics (other than prices), and across consumer attributes. To do so, we select newly issued mortgages, excluding homeowners that are either refinancing or renewing their mortgage contract, and contracts with a 25 year amortization period and five year fixed-rate term. During our sample period, nearly all mortgage contracts were fixed rate, and over 85 percent had a five year term. A five year fixed-rate mortgage contract must be renegotiated every five years, and banks impose substantial penalties to refinance before the end of the term.[14] This has been the standard contract offered by Canadian banks since the late 1960s. Similarly, almost all contracts have 25 year amortization periods.

Market Structure.—Our choice set definition creates a split between contracts signed in areas where both *A* and *B* were present pre-merger (i.e., treatment group), and areas in which only *A* or *B*, or neither were present (i.e., control group). Not all of these neighborhoods are directly comparable, and local markets with both *A* and *B* tend to be larger and have more lenders. Table 1 shows that the average number of lenders pre-merger was slightly less than six in the control neighborhoods, compared to 9.6 in the treated neighborhoods.

The two groups significantly overlap only in medium-sized markets with five to eight lenders. Less than 1 percent of control markets have more than eight lenders,

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

while less than 1 percent of treatment markets have fewer than five. Ideally, we would estimate the treatment effect separately for consumers facing similar choice sets, but the distribution of lenders only allows us to do so for local markets with an intermediate level of concentration.[15]

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

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

D. *Variable Definitions and Descriptive Statistics*

Our analysis focuses mostly on two outcome variables: discounts and margins. These two variables provide direct measures of the cross-sectional dispersion of rates, by controlling for intertemporal variation in aggregate interest rates. Margins are measured as the interest rate paid by consumers less the swap-adjusted five year government of Canada bond rate measured at the week of negotiation, and discounts are measured relative to the bank-specific posted rate at the week of negotiation. The fraction of households receiving zero discounts is defined as an indicator variable equal to one if the observed discount is less than 10 bps.

Note that we only observe the closing date on the house sale. For each contract, we identify the negotiation week by calculating the absolute difference between the transaction rate and the posted rate for the weeks within 90 days of the closing date, and take the smallest value, and/or the closest in time in case of a tie. From this, we estimate that 44 percent of contracts are negotiated within 1 week of the closing week, and the remaining are approximately uniformly distributed between 1 week and 90 days. This definition also increases the fraction of consumers paying a rate equal or close to the posted rate. Less than 10 percent of consumers pay a rate equal to the posted rate valid at the closing date, while we estimate that roughly 25 percent pay the posted rate valid at the "negotiation" week. Many of the consumers that we categorize as receiving zero discounts actually benefit from a discount relative to

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

the posted rate valid at the closing date, while others end up paying more than this posted rate.

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

$$(2) \qquad Margin_i = \boldsymbol{\beta}'\mathbf{X}_i + \mu_i^{\text{week}} + \theta_{G_i,T_i} + e_i = \boldsymbol{\beta}'\mathbf{X}_i + m_i,$$

where $Margin_i = Rate_i - Bond_i$, and θ_{G_i,T_i} is a group/period fixed effect, G_i and T_i index the group (treatment or control) and time period (before or after the merger) for borrower i, μ_i^{week} is a closing-week fixed effect, and $\mathbf{X}_i$ is a vector of control variables.[16] In the empirical analysis we mostly focus on the impact of the merger on the level and dispersion of *negotiated margins*, m_i. Note that we will sometimes refer to this as a residual price, especially when we relate our findings with the literature on residual price dispersion.

In equation (2) and in the rest of the paper, we index each observation by i, with the understanding that i captures: (i) the individual borrower (or household), (ii) the time period of the contract, and (iii) the location of the purchased house. Although we observe each household only once, we observe most FSAs pre- and post-merger. This allows us in some specifications to control for location fixed effects in $\mathbf{X}_i$.

Table 2 describes the sample, split across treatment/control groups and pre/post-merger. Overall the average margin is 117 bps, and it has increased over time. In contrast, the fraction of consumers paying the posted rate (i.e., zero-discounts) remained fairly stable at around 26 percent. Since the remaining 74 percent of consumers negotiate discounts, most of the dispersion in margins is unexplained by variation in the level of interest rates. This is illustrated by the negotiated margin dispersion on the second line of each panel of Table 2. Slightly more than half of the margin variance in our data comes from cross-sectional residual dispersion alone. The magnitude of dispersion is large compared to other financial markets, considering the homogeneity of the contract terms and the presence of insurance. For instance, Hortaçsu and Syverson (2004) show that the cross-sectional standard deviation among mutual fund transaction fees was equal to 60 bps in 2001. Mutual funds exhibit substantially more heterogeneity in observable characteristics and returns than the five-year fixed rate mortgage contracts studied here.

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

This can be seen by comparing income and loan size in the two groups. The average homeowner in the treatment group earns \$69,330 and contracts a loan of \$152,300, while in the control group the mean income and loan size are \$61,950 and \$113,400,

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

Table 2—Summary Statistics on Mortgage Contracts and Household Characteristics

	Control/before				Control/after			
	Mean	SD	$P(25)$	$P(75)$	Mean	SD	$P(25)$	$P(75)$
Margin	1.07	0.46	0.73	1.42	1.43	0.56	1.02	1.82
N-Margin	1.07	0.43	0.78	1.35	1.67	0.46	1.36	2.02
$1(r_i = \overline{r}_i)$	36.38	48.12			26.92	44.36		
Income	61.95	25.00	43.70	74.77	62.84	24.49	45.37	75.29
House	121.11	55.49	82.18	145.30	118.25	52.16	81.18	143.74
Loan	113.4	49.7	78.7	137.2	110.0	46.9	76.9	133.5
LTV	91.58	4.26	90.00	95.00	91.25	4.34	90.00	95.00
FICO	67.23	46.95			64.27	47.93		
Renter	68.30	46.55			70.05	45.82		
Parents	5.73	23.24			6.64	24.90		
Switch	30.21	45.93			36.84	48.25		
Broker	21.91	41.37			30.05	45.86		
	Treatment/before				Treatment/after			
	Mean	SD	$P(25)$	$P(75)$	Mean	SD	$P(25)$	$P(75)$
Margin	0.93	0.49	0.65	1.27	1.44	0.56	1.12	1.82
N-Margin	1.06	0.43	0.78	1.35	1.72	0.47	1.44	2.00
$1(r_i = \overline{r}_i)$	23.67	42.51			22.23	41.58		
Income	69.33	26.48	50.71	82.23	70.99	26.49	52.46	84.11
House	162.93	63.38	116.77	201.83	161.07	64.58	114.84	200.72
Loan	152.3	57.6	110.5	188.2	149.9	57.6	108.2	187.0
LTV	91.35	4.25	90.00	95.00	90.99	4.48	89.60	95.00
FICO	62.40	48.44			62.56	48.40		
Renter	68.28	46.54			71.15	45.31		
Parents	8.31	27.61			9.32	29.08		
Switch	26.68	44.23			38.43	48.65		
Broker	15.66	36.34			27.73	44.77		

Notes: The sample size is 18,121 divided between the control and treatment group, pre- and post-merger with 62.8 percent of contracts in the treatment and 42.2 percent observed post-merger. It includes a random sample of homogenous term and amortization contracts insured by CMHC or Genworth within one year of the merger. Margins and negotiated margins (N-Margin) are defined in the text. $1(r_i = \overline{r}_i)$ corresponds to the percentage of consumers paying the posted rate ($\overline{r}_i$). House measures the house price. FICO is an indicator variable equal to one if a consumer's credit score is greater than 600. Renters and parents correspond to new home buyers exiting from renting and living with parents, respectively. Switcher is an indicator variable equal to one if consumers have no prior experience with the chosen financial institution. All indicator variables are measured in percentage, and dollar expenses are measured in $1,000. The sample is restricted to households with 5 to 8 lenders located within 5 km of their FSA centroid.

respectively. However, loan-to-value ratios are similar across the two groups. The other characteristics are all quite similar across the two groups prior to the merger.

The effect of the merger on rates can already be seen in Table 2. Rates rise in both the treatment and control markets, but the increase is about 6 bps greater in the treatment. Similarly, the fraction of consumers paying the posted rate falls everywhere, but by less in the treatment.

III. Empirical Results

In this section we summarize the main empirical results of the paper. We focus first on measuring the average treatment effect (ATE) of the merger on negotiated rates and on the probability of paying the posted rate. We estimate the ATE using three econometrics techniques: (i) linear difference-in-difference (DiD) estimated by OLS, (ii) the propensity-score matching DiD estimator developed by Heckman, Ichimura and Todd (1997), and (iii) the change-in-change (CiC) estimator proposed by Athey and Imbens (2006).

Relative to OLS, the matching estimator allows us to control for the possibility that observationally different local markets in the treated and control groups experienced differential trends. We implement this by estimating a propensity score function that controls for demographics and market-structure characteristics of each FSA. Recall that we selected our treated and control markets based on the pre-merger number of lenders. The propensity-score estimator can therefore be viewed as a way of improving on this matching strategy solely based on market-structure.

The CiC estimator is appealing in our context because it explicitly accounts for heterogenous treatment effects. We use it to estimate the counterfactual distribution of negotiated margins that would have been observed in the treated areas absent the merger. In particular, assuming that the reduced-form pricing functions defined in equation (1) is monotonic in u_i, and that the distribution of u_i in the treated and control groups is time-invariant, Athey and Imbens (2006) showed that the counterfactual price distribution can be computed from three empirical distribution functions (EDF)

$$\hat{F}^{c}_{1,1}(m_i) = \hat{F}_{1,0}\Big(\hat{F}^{-1}_{0,0}\big(\hat{F}_{0,1}(m_i)\big)\Big), \tag{3}$$

where the superscript c identifies the counterfactual distribution, and $\hat{F}_{G,T}(m)$ is the EDF of negotiated margins in the subpopulation (G, T). Intuitively, we obtain the counterfactual distribution by transforming the observed negotiated margin distribution in the treatment group at time 0 $\big($i.e., $F_{1,0}(m)\big)$ to mimic the changes observed in the control group.

In addition to reporting the average effect, we will use this estimated distribution to measure the change in residual rate dispersion, as well as the quantile treatment effect of the merger for each percentile $q_i : \alpha(q_i) = \hat{F}^{-1}_{1,1}(q_i) - \hat{F}^{c^{-1}}_{1,1}(q_i)$. We present the distribution results in Section IIIB.

In online Appendix B, we discuss in greater detail the identification assumptions, and provide a careful description of each estimator.

A. *Average Impact of the Merger*

In Table 3 we present the average effect of the merger on margins and discounting, estimated by OLS, matching, and the CiC estimator. Results are presented for both "baseline" and "trends" specifications.[17] Our baseline specification includes controls that describe the financial and demographic characteristics of the contract, and the identity of lenders.[18] In an effort to further control for confounding factors that could bias our estimate, we also use a richer set of variables that control for heterogenous observable trends across our control and treatment groups.

[17] In online Appendix D we report OLS estimates of the merger effect, along with the marginal effect of other covariates to provide the interested reader with some information on how these are priced. For a more detailed discussion of the pricing of these contracts see Allen, Clark, and Houde (2014c).

[18] The exact set of control variables is listed in footnote 16.

Table 3—Average Effect of the Merger on Margins and Discounts

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

Notes: Standard errors clustered at the FSA level are in parentheses. The dependent variable in columns 1 and 2 is the transaction rate minus bond rate, in columns 3 and 4 is an indicator variable for the transaction rate within ten basis points of the posted rate, which we take to imply zero discount. All specifications include borrower characteristics and bank characteristics as well as FSA and week fixed effects. The trend specifications include province trends as well as the borrower covariates interacted with the "after-merger" dummy. The matching estimator is calculated using the propensity score with four (4) nearest neighbors. Standard errors and hypothesis tests are calculated by bootstrapping the original sample 1,000 times. See Huynh, Jacho-Chávez, and Voia (2011) for analysis of the bootstrap performance in the context of the CIC estimator.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

We do so by interacting every consumer and contract characteristics in $\mathbf{X}_i$ with T_i, our "post-merger" dummy variable. We also include in $\mathbf{X}_i^{\text{trend}}$ province-level linear and quadratic trends. In addition, the OLS specifications also control for FSA fixed effects.

The linear DiD specifications are estimated by OLS

$$\text{(4)} \qquad \mathrm{Y}_i = \overline{\alpha} I_i + \boldsymbol{\beta}'\mathbf{X}_i + \mu_i^{\text{week}} + e_i,$$

where Y corresponds to either margins, or the zero-discount indicator.

To control for the same covariates in the matching and CiC estimators we proceed in two stages. First, we estimate a negotiated margin term m_i by OLS from equation (2). Then, we use m_i as our outcome variable to estimate the ATE of the merger using the matching and CiC estimators.[19] The standard errors of the final estimates are corrected by bootstrapping the original sample 1,000 times.

Columns 1 and 2 present the ATE of the merger on margins in the baseline and trends cases, respectively. Adding the trend variables increases the point estimate from about 6 bps to 7 bps in the OLS specification. This increase is not uniform

[19] We only use the CiC estimator to study margins. A similar estimator is also available to analyze the zero-discount probability, but would yield bound (rather than point) estimates. This would complicate significantly our two-stage procedure to control of observed characteristics of borrowers.

across the three methods of estimation. The matching estimator yields slightly higher estimates of the ATE, but the baseline estimate is higher than the one with trends (i.e., 7.39 versus 6.92 bps). The average effect estimated with the CiC estimator is the lowest, between 5.7 and 6.6 bps.

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

Overall, the three methods produce remarkably similar estimates, ranging from 5.7 to 7.39 for margins, and from 4.77 to 6.7 percentage points for the probability of not receiving a discount. The estimates are also precisely estimated in most specifications; only the matching ATEs with trends have a p-value greater than 1 percent. We can therefore reject the null hypothesis that the merger did not change market power, at least for the subsample of neighborhoods with five to eight lenders pre-merger. This implies that residual dispersion is not solely driven by risk- or cost-based pricing. Recall that this is the standard model assumed in the finance literature to explain the observed dispersion of lending rates (e.g., Edelberg 2006; and Einav, Jenkins, and Levin 2012).

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

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

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

Robustness Analysis.—We have analyzed the robustness of the average merger effect on margins with respect to the choice of controls, the size of local neighborhoods, and different event windows. Results are available in online Appendix C.

First, we consider alternative explanatory variables controlling for location-specific trends. Our most conservative estimate is equal to 5 bps (p-value < 5 percent), which is well within the confidence interval bounds of the estimates reported in Table 3. This specification identifies the ATE solely using within province/year variation (i.e., adding province/year fixed effects). We also consider two alternative control group definitions: (i) locations where only bank A is present, or (ii) locations where none of the two merging parties are present. Both definitions yield results that are nearly identical in magnitude to those in Table 3, albeit less precisely estimated.

Next, we show that the average impact is not implausibly sensitive to the size of the choice set. The ATE is larger and more precisely estimated around a 5 km radius, and smoothly goes toward a zero impact as we increase or decrease the threshold toward 7 km or 3 km. This pattern is consistent with the fact that using radii that are too large or too small exacerbates the measurement error problem in our treatment variable, which biases our results toward zero.

Finally, we consider different event windows, and report the results of a falsification exercise in which we move the merger date by $+/-$ 6 months. On the first point, the results suggest that the effect of the merger is strongest within six months of the announcement (between 6.8 and 9.7 bps), and diminishes over time (between 2 and 4.4 bps after 18 months), possibly due to the closing of duplicate branches. The second exercise confirms that rates did in fact increase suddenly around the actual merger date. We fail to find any evidence of rate increases six months before or after the merger. This suggests that our estimates are not confounded with the presence of unobserved events differentially affecting the treated neighborhoods within a year of the merger.

B. *Distributional Impact of the Merger*

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

Figure 2 plots the counterfactual (long-dashed) and observed (solid) empirical distribution functions (EDFs) for consumers with $G_i = 1$ and $T_i = 1$. The short-dashed lines represent bootstrapped 95 percent confidence intervals. For about three-quarters of the range, the counterfactual distribution is shifted to the left, confirming that most consumers pay higher interest rates as a result of the merger. The figure also reveals that the two distributions are nearly identical for consumers located in the top percentiles, suggesting that consumers paying higher rates were unaffected by the merger. Therefore, the increase in average rates is entirely due to the fact that consumers who, without the merger, were paying relatively low rates, experienced a significant decline in their ability to negotiate discounts. The same pattern is observed if we use a specification with trend, or if we vary the dimension of the choice sets.[20]

[20] These results are available upon request.

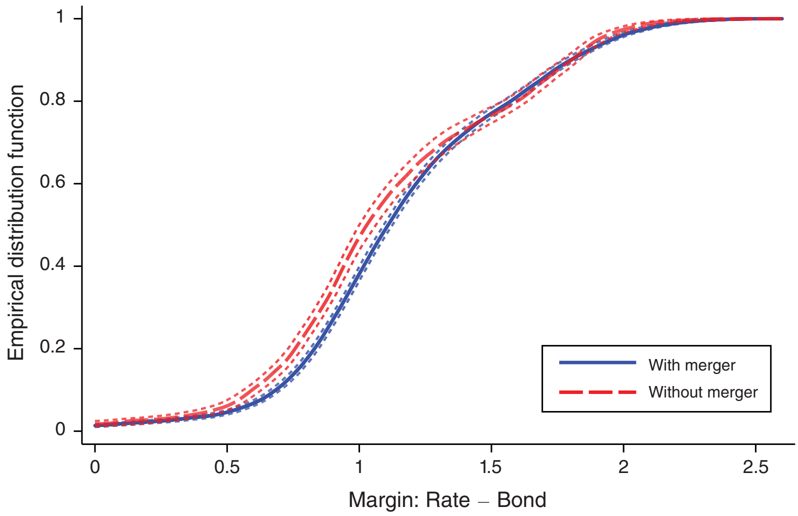

FIGURE 2. EMPIRICAL DISTRIBUTION OF NEGOTIATED MARGINS WITH AND WITHOUT THE MERGER

TABLE 4—DISTRIBUTIONAL EFFECT OF THE MERGER ON NEGOTIATED MARGINS

	Baseline			With trend		
	Est.	95% CI		Est.	95% CI	
Quantile effects						
q_5	0.0766	0.0123	0.157	0.0782	0.0080	0.18
q_{10}	0.0914	0.0505	0.13	0.0844	0.0379	0.138
q_{25}	0.0762	0.0454	0.0994	0.0767	0.0397	0.119
q_{50}	0.0842	0.0433	0.114	0.0874	0.0446	0.128
q_{75}	−0.0008	−0.0621	0.0686	0.0185	−0.0473	0.0885
q_{90}	−0.0042	−0.0485	0.0383	0.0069	−0.0393	0.0591
q_{95}	0.056	0.0053	0.0949	0.0513	−0.0125	0.0973
Dispersion effects						
ΔSD	−0.0287	−0.0587	−0.0066	−0.0256	−0.056	−0.0034
ΔCV	−0.0462	−0.0766	−0.0211	−0.0455	−0.0792	−0.0184
$\Delta q_{75} - q_{25}$	−0.077	−0.131	−0.0099	−0.0582	−0.114	−0.0044
$\Delta q_{90} - q_{10}$	−0.0957	−0.153	−0.0457	−0.0776	−0.134	−0.0261
$H_0 : F^c_{1,1} = F_{1,1}$						
KS	4.7***			4.84***		

Notes: The dependent variable is negotiated margins. ΔSD and ΔCV measure the changes in the standard deviations and coefficient of variation respectively. Confidence intervals were calculated by bootstrapping the sample 1,000 times. The KS statistic is a test of the equality of the observed and counterfactual empirical distribution functions.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 4 summarizes the effect of the merger on various moments of the margin distributions. At the bottom of the table, the Kolmogorov-Smirnov (KS) statistic is a measure of the difference between two EDFs, and tests the null hypothesis of equality of the rate distributions with and without the merger. As suggested by Figure 2, the test statistic easily rejects this null hypothesis. The counterfactual rate distribution without the merger therefore first-order stochastically dominates the post-merger rate distribution observed in the data.

The first seven rows of Table 4 calculate the difference between the two EDFs, and list the quantile treatment effect (QTE) of the merger. Consistent with the figure, we find that the effect of the merger ranges from around 9.1 bps for borrowers at the tenth quantile to zero at the ninetieth quantile. We cannot reject the null hypothesis that the merger had no effect on negotiated rates for consumers beyond the thirtieth percentile. The QTE is similar once we control for trends (column 4).

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

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

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

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

In contrast, we measure the impact of losing one competitor on residual price dispersion in a comparable group of local markets. By exploiting the variation induced by the merger, our approach does not require the use of an instrumental variable, nor does it compare firms operating under different market structures.[21] We therefore think that our method is more credible and could be used in other contexts as well to study the link between competition and dispersion.

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

IV. Mergers When Prices are Negotiated

In this section, we interpret our empirical results through the lens of a search and price negotiation model. We assume that the transaction rate is the sum of a deterministic function of borrower characteristics and the bond rate (i.e., $X_i'\beta$ in equation (2), and negotiated margin m_i. Our interest is in modeling the determination of negotiated margins.

Our objective is to provide a simple and transparent framework that can replicate both qualitatively and quantitatively the effects of the merger. In order to do so, we develop and estimate a model of search and price negotiation that is consistent with the identification assumptions of the CiC estimator described above.[22]

Finally, we use the estimated model to perform a counterfactual experiment in which we lower the search cost of consumers. We use the result of this exercise to quantify the importance of search frictions in explaining the dispersion and level of prices in the market, as well as to predict the effect of mergers in markets with lower search costs.

A. *Model*

We consider an environment with $n + 1$ lenders, in which negotiation takes place over three stages. In the first, consumers receive a take-it-or-leave-it offer m_i^0 from one lender.[23] In the second, if the initial offer is rejected, consumers put forth a search effort e_i at cost κ_i to gather additional quotes. The number of quotes that consumers obtain is stochastically determined by their search effort: with probability $s(e_i)$ they are randomly matched with n banks, and with probability $1 - s(e_i)$ they receive 2 quotes. In the final stage, competition takes place, and consumers choose the lowest price offer. We characterize each stage sequentially, starting with the final competition game.[24]

Competition Stage.—As in Woodward and Hall (2012) and Allen, Clark, and Houde (2014a), we model the competition stage as an English auction between at most n lenders. Alternatively, one can think of this stage as a Bertrand game between a random number of lenders, where the randomness is created by consumers' search effort. This modeling strategy differs from the standard search model often used in I.O. (e.g., Varian 1980), which is based on a price-posting assumption. However, it is a common way of introducing negotiation in on-the-job-search environments (e.g., Postel-Vinay and Robin 2002).

[22] The CiC assumptions are described in greater detail in online Appendix B.

[23] The assumption that firms have all the bargaining power is not important. The predictions of the model hold in a Nash bargaining model in which borrowers and lenders share the transaction surplus.

[24] The timing and assumptions of the model abstract from several features of the mortgage market. For instance, we assume that the initial lender does not participate in the competition stage in order to simplify the pricing problem, and that consumers get a minimum of two quotes in order to generate a finite price even for consumers with high search costs. Alternatively we could have explained the reduced-form results by adding heterogeneous reservation values, for instance due to the posted rate or the value of renting a house. We chose not to include these features, since they would require more than one source of unobserved heterogeneity, and would therefore prevent us from estimating the model using solely the results from the retrospective merger analysis. In Allen, Clark, and Houde (2014a), we relax some of these assumptions by estimating a model of search and negotiation that exploits additional data on the search and switching decisions of consumers, and that incorporates multiple sources of residual price dispersion.

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

The game has a unique dominant-strategy equilibrium: banks are willing to offer up to their privately observed cost $c + \epsilon_j$, and the most efficient bank wins the contract by offering a rate equal to the second lowest cost, $c_{(2)} = c + \epsilon_{(2)}$, where $\epsilon_{(2)}$ is the second order statistic of $\{\epsilon_1, \ldots, \epsilon_k\}$. Let $E(m^* | \tilde{n}) = c + E(\epsilon_{(2)} | \tilde{n})$ denote the expected second-stage transaction price, where $\tilde{n}$ is the number of quotes generated (either two or n). Given our matching assumption, the gain from searching is summarized by the expected cost difference between obtaining 2 or n quotes: $\Delta(n) = E(\epsilon_{(2)} | 2) - E(\epsilon_{(2)} | n) > 0$. This function is increasing in n.

Search Effort Stage.—At this stage, consumers incur a search cost κ_i to gather additional quotes. The search cost has a fixed and variable component:

$$\kappa_i = u_i(1 + e_i) + \eta_i. \tag{5}$$

We assume that the marginal cost of effort u_i is publicly observed by both parties, and that consumers privately observe the sunk cost η_i of gathering extra quotes.

We use a simple matching technology to describe the number of quotes that consumers receive. Search effort e increases the probability $s(e)$ of receiving the maximum number of quotes n. Consumers putting forth zero effort are automatically matched with two lenders, and we use a Pareto distribution to characterize the matching probability function

$$s(e) = 1 - (1 + e)^{-\lambda}. \tag{6}$$

Consumers choose an optimal level of search effort to minimize the sum of the expected transaction price and the search cost. Importantly, this effort level depends only on public information, since η_i is a fixed-cost. A consumer facing n possible options chooses an optimal effort level that minimizes the net borrowing cost

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

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

The search decision depends on the reserve value of consumers net of the private-valued search cost: $r(u_i, n) - \eta_i$. In particular, if m_i^0 is the initial offer received in the first stage, the search probability is given by the following expression:

$$H(m_i^0 | u_i, n) = \Pr\big(\eta_i < m_i^0 - r(u_i, n)\big). \tag{8}$$

Notice that since $\eta_i \geq 0$, the search probability is equal to zero if the initial quote is lower or equal to the reservation value: $m_i^0 \leq r(u_i, n)$.

Initial Offer.—In the first stage, the initial lender has two choices. It can offer the consumer's reservation value in order to prevent search altogether, or it can choose m_i^0 to maximize its expected profits:

$$\max_{m_i^0 \geq r(u_i, n)} \; (m_i^0 - c)\big[1 - H(m_i^0 \,|\, u_i, n)\big]. \tag{9}$$

The first-order condition to this problem is analogous to the monopolist pricing rule

$$m_i^0 - c = \frac{1 - H(m_i^0 \,|\, u_i, n)}{h(m_i^0 \,|\, u_i, n)},$$

where $h(u_i, n) = \partial H(m_i^0 \,|\, u_i, n)/\partial m_i^0$. The initial offer is therefore equal to the cost of lending plus a constant markup corresponding to the inverse of the hazard rate of losing the consumer, determined by the shape of search cost distribution. If η_i is exponentially distributed, this markup corresponds to the average private-value search cost in the population.

The optimal initial offer is the maximum of the solution to equation (9) and the reservation value of consumers. If the interior solution is less than $r(u_i, n)$, the lender can increase its initial offer without increasing the probability of losing the consumer. Intuitively, the lender will find it optimal to prevent search entirely for consumers with high observable search costs (i.e., high u_i). In contrast, it will offer a relatively high price to consumers with good negotiation skills (i.e., low u_i), at the risk of losing these consumers to competing lenders. This prediction is very similar to the results found by Chatterjee and Lee (1998), in which the value of the next best alternative is privately known to the buyer.

In the empirical analysis below, we consider the limiting case in which information is symmetric (i.e., $\eta_i = 0$). In this case, the initial lender offers the consumer's reservation value with probability one, and this offer is automatically accepted. Therefore, in equilibrium, the distribution of transaction margins corresponds to the distribution of reservation values: $m_i = m_i^0 = r(u_i, n)$.

We make this assumption for two reasons. First, as we discuss below, this specification yields clear analytical predictions with respect to the impact of competition on the dispersion of rates. Second, the model implies the existence of a reduced-form pricing equation that matches the assumptions of the CiC estimator; namely that the equilibrium price distribution can be mapped into a time-invariant distribution of consumer unobserved heterogeneity.

In contrast, a model with asymmetric information produces a distribution of prices that is a function of (at least) two unobserved heterogeneity terms: (i) the search cost u_i, and (ii) the competition outcome $\epsilon_{(2)}$. In Section IVC we explain how the asymmetric information case, while no longer satisfying the identifying assumptions of the CiC estimator, nonetheless provides similar results with respect to the relationship between market-power and prices, and accurately reproduces the reduced-form results.

Predictions of the Symmetric Information Model.—The symmetric information model implies that the bargaining leverage of consumers depends on the threat of gathering additional quotes if negotiation fails. As in Bester (1988) and Wolinsky (1987), firms observe the reservation value of each consumer, and tailor their offers such that search does not take place in equilibrium. This modeling strategy is also common in the recent empirical bargaining literature, in which negotiation is assumed not to fail (e.g., Crawford and Yurukoglu 2012; Grennan 2013; and Gowrisankaran, Nevo, and Town 2013).[25]

This leads to a reduced-form pricing function that has the following properties. It (i) is monotonically increasing in u_i, (ii) is weakly decreasing in n, and (iii) features an increasing marginal effect of u_i with respect to n, $\partial^2 m_i / \partial u_i \partial n \geq 0$. These results are presented in online Appendix E.

The first property implies that the pricing function satisfies the monotonicity assumption used to estimate the counterfactual price distribution (see online Appendix E). Therefore, as in Hortaçsu and Syverson (2004) and Hong and Shum (2006), the distribution of search costs can be recovered from the observed post-merger negotiated margin distribution

$$F_{1,1}(m) = \Pr(c + \pi(u_i, n) < m) = H_1\big(\pi^{-1}(m - c, n)\big). \tag{10}$$

The second implies that mergers raise prices for all but those consumers who in equilibrium expect not to search: $\alpha(u_i) = \pi(u_i, n-1) - \pi(u_i, n) > 0$ for all $u_i < \overline{u}(n)$. Therefore, assuming that some consumers expect to exert positive search effort in equilibrium, the average effect of mergers on prices is positive. The presence of a corner solution at the top of the distribution can also generate the type of heterogenous treatment effects that we document in Table 4.

The third property also predicts heterogeneous effects, even absent of the presence of a corner solution. In particular, it implies that the price change due to the merger is decreasing in u

$$\begin{aligned} \alpha(u) &= \pi(u, n-1) - \pi(u, n) \\ &\geq \pi(u', n-1) - \pi(u', n) = \alpha(u'), \quad \text{for all} \quad u' \geq u. \end{aligned} \tag{11}$$

Intuitively, relative to high u_i consumers, low search cost consumers are more likely to receive n quotes, and are therefore more affected by a reduction in $\Delta(n)$. This prediction is consistent with the observed non-linear effect of the merger: consumers at the top of the distribution were not impacted by the merger, while borrowers below the seventieth percentile experienced rate increases between 7 and 9 bps.

The increasing difference property also implies that mergers decrease the amount of residual price dispersion, consistent with our empirical results. In particular,

[25] This assumption is also common in the on-the-job-search literature (e.g., Postel-Vinay and Robin 2002 and Dey and Flinn 2005). In the labor market context, firms and workers observe the current and alternative match values, and bargain over the surplus split. Job separation occurs in equilibrium only if the current match is lower than the match value of the outside alternative. We rule out this possibility in our context by assuming that consumers qualify for a loan at every lender.

equation (11) directly implies that mergers decrease the interquartile range of negotiated margins

$$\alpha(u_{25}) = \pi(u_{25}, n-1) - \pi(u_{25}, n) \geq \pi(u_{75}, n-1) - \pi(u_{75}, n) = \alpha(u_{75})$$

$$0 \geq \mathrm{IQR}(n-1) - \mathrm{IQR}(n),$$

where u_q denotes the qth percentile of the marginal search cost distribution, and IQR denotes the interquartile range.

B. *Model Estimation*

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

Notice also that the gain from searching is a function of the maximum number of quotes pre-merger (i.e., n). We set n to 5, which corresponds to the lower-bound in the actual number of lenders in our estimated sample. While this number is somewhat arbitrary, it roughly corresponds to the maximum number of quotes consumers report in surveys of shopping behavior.[26]

To understand the scale of the model parameters, recall that we use the negotiated margin, defined in equation (2), as the dependent variable to estimate $\hat{F}^c_{1,1}$. Let m_i^I and m_i^N denote the observed (with merger) and counterfactual (without merger) transaction margins for consumer i. The lending cost parameter c therefore measures the average lending cost above the bond rate for consumers with representative characteristics.

We consider a nested fixed-point estimator that minimizes the distance between moments of the predicted and estimated distributional effects of the merger (i.e., model predictions versus CiC results), conditional on the restriction that the search cost distribution is consistent with the observed margin distribution. In practice we impose these restrictions at S random percentiles. This leads to the following GMM estimator:

$$J = \min_{\boldsymbol{\theta}=\{c,\gamma,\sigma_\epsilon\}} (\mathbf{M}(\boldsymbol{\theta}) - \hat{\mathbf{M}})^T \boldsymbol{\Omega}^{-1} (\mathbf{M}(\boldsymbol{\theta}) - \hat{\mathbf{M}}) \tag{12}$$

$$\text{s.t.} \quad u_{q_i} = \pi^{-1}(m_i^I - c, n - 1 | \boldsymbol{\theta}), \forall i = 1, \ldots, S,$$

where $q_i \sim U[0,1]$ denotes a simulated percentile, $m_i^I = \hat{F}^{-1}_{1,1}(q_i)$ is the inverse of the post-intervention margin distribution, and $\mathbf{M}(\boldsymbol{\theta})$ and $\mathbf{M}$ are two $L \times 1$ vectors

[26] See, for instance, the "Residential mortgage survey" conducted by Clayton Research.

describing respectively the predicted and estimated moments associated with the distributional effects of the merger (i.e., $\alpha(u_i)$). The moments are calculated by Monte Carlo simulation using $S = 100$ independent draws. In the empirical application, we use seven moments: the mean and standard deviation of the merger effect (i.e., $\overline{\alpha}$ and $\mathrm{sd}(\alpha_i)$), and the treatment effect at five percentiles of the distribution of u_i (i.e., α_{10}, α_{25}, α_{50}, α_{75}, α_{90}). Finally, the weighting matrix $\boldsymbol{\Omega}$ is constructed using the variance-covariance matrix of the empirical moments, obtained from the same bootstrapping procedure described in Section IIIA.

Identification.—Conditional on $(\gamma, \sigma_\epsilon)$, the distribution of search costs is non-parametrically identified up to an additional scale restriction. This is because the minimum search cost is not separately identified from the intercept c, which both determine the level of prices pre-merger. We therefore impose a lower bound on the search cost distribution, $u_0 = 0$. This is equivalent to assuming that consumers receiving the largest discounts in our sample would have obtained the maximum number of quotes with probability one.[27]

In addition, the scale of the search cost at the top of the distribution is identified by the merger effect at the top of the price distribution. In the limit, consumers in the very top percentiles of the distribution are unaffected by the merger, and pay a price that is additive in the search cost of effort. See equation (7) evaluated at $s(e) = 0$.

Finally, σ_ϵ and γ are identified directly from the effect of the merger at different percentiles of the distribution. First, σ_ϵ determines the size of the pure market-power effect of the merger: $\Delta(n-1) - \Delta(n)$. Under the assumption that consumers at the bottom of the price distribution have low search costs, σ_ϵ is proportional to the merger effect in the lowest percentiles.

The parameter γ, on the other hand, determines the elasticity of search effort with respect to n. If consumers are insensitive to n, the effect of the merger will be homogenous across the different percentiles. Alternatively, if consumers drastically change their search effort in response to a change in n, the model will predict a steeply declining merger effect across u_i's. Since u_i maps directly into the percentiles of the residual price distribution, this parameter is identified by the importance of heterogeneity in effect of the merger.

Estimation Results and Goodness of Fit.—Table 5 presents the model estimation results. Panel A shows the estimated parameters and the value of the GMM objective function. The standard errors are calculated using the bootstrapped price distributions estimated in Section IIIB, holding fixed the S simulated percentiles. Panel B summarizes the model predictions at the estimated parameters (columns 1–3), and compares the predicted and estimated quantile treatment effects of the merger (columns 4 and 5).

The average transaction margin can be calculated as the sum of the estimated average profit, 43 bps, and the estimated lending cost $\hat{c} = 0.714$, which can be interpreted as the average lending cost over the five-year bond rate. Scaling up by the bond rate observed over our sample period, as in equation (2), leads to predicted markups ranging between 5 percent and 10 percent.

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

TABLE 5—MODEL ESTIMATION RESULTS

	c	γ	σ_ε	n	J-stat
Panel A. Estimation results					
Parameters	0.714	0.59	0.8	5	13.1
	(0.033)	(0.096)	(0.13)	—	(7.7)

	Marginal search-cost u_i	Profit margins		Merger effect $\alpha(u_i)$	
		Pre	Post	Predicted	Estimated
Panel B. Simulation results					
Average	0.272	0.382	0.433	0.0512	0.0585
	(0.053)	(0.044)	(0.036)	(0.011)	(0.015)
SD	0.23	0.334	0.308	0.0266	0.0319
	(0.018)	(0.0051)	(0.004)	(0.0051)	(0.0073)
25th percentile	0.0973	0.118	0.191	0.0731	0.0664
	(0.037)	(0.046)	(0.035)	(0.013)	(0.015)
Median	0.211	0.34	0.394	0.0536	0.0841
	(0.058)	(0.044)	(0.036)	(0.011)	(0.019)
75th percentile	0.431	0.659	0.685	0.0258	0.0235
	(0.076)	(0.043)	(0.038)	(0.012)	(0.035)

Notes: Bootstrap standard errors are in parentheses. Number of bootstrap replications = 1,000. Number of simulated prices used in the estimation = 100. The top and bottom 5 percent are eliminated from the estimated price distributions. Parameter definitions: c = lending cost, γ = matching probability parameter, σ_ε = cost heterogeneity, n = maximum number of quotes pre-merger (fixed).

Columns 2 and 3 of panel B show that profit margins vary widely across consumers. Because banks are able to perfectly discriminate, profit margins are less than 19.1 bps below the twenty-fifth percentile of the post-merger distribution, and more than 68 bps above the seventy-fifth. Undoing the merger would have led to a marked reduction in profit margins, but mostly among consumers with low search costs (e.g., 7.3 bps at the twenty-fifth percentile and only 2.6 bps at the seventy-fifth). Similarly, consistent with our empirical results, the merger significantly reduced the dispersion of profit margins across borrowers. Our simulation results show that the coefficient of variation went from 87 percent before the merger, to 71 percent after (i.e., $0.334/0.382$ versus $0.308/0.433$).[28]

The estimated parameter σ_ϵ determines the expected gain from searching. At the estimated value, the gain from obtaining five quotes instead of two corresponds to $\Delta(n) = 53.3$ bps. Taken literally, this parameter also implies important idiosyncratic cost differences across lenders: the range of ϵ_i corresponds to $2 \times 0.8 = 1.6$ bps. However, as we discussed above, heterogeneity across lenders in the competition stage can arise from differentiation and cost heterogeneity without changing the model formulation.

The matching probability parameter γ determines the marginal effect of search effort on the probability of receiving the maximum number of quotes. Since the

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

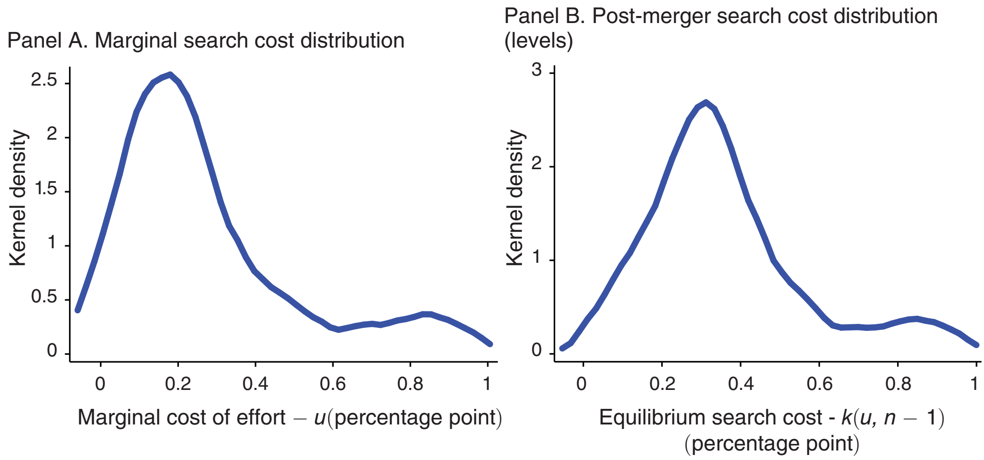

Figure 3. Estimated Search Cost Distributions

estimated marginal cost of effort u_i varies greatly across consumers (standard deviation is 0.23), this marginal product is highly dispersed. Evaluated at the average cost 0.272, a 1 percent increase in effort leads to a 1.7 percent increase in the probability of receiving 5 quotes. In comparison, the search effort elasticity is close to one at the ninetieth percentile of the search cost distribution, and above 6 at the tenth.

The goodness of fit of the model can be studied by comparing the predicted and estimated effect of the merger (columns 4 and 5). As we discussed in Section IIIB, the data suggest a fairly flat merger effect among consumers in the middle and low percentiles of the price distribution. In contrast, the model predicts a monotonically decreasing merger effect, which tends to overestimate its effect at the bottom of the distribution, while underestimating it at the median. As a result, according to the J-statistic, the model is rejected at the 5 percent level.

These differences arise because the estimated effect is fairly noisy at the top of the distribution, as illustrated by the quantile treatment effect beyond the ninetieth percentile (see Table 4). In contrast, the model predicts a strictly monotonically decreasing effect: If the merger has a zero effect at percentile q, the effect must also be zero for all percentiles $q' \geq q$. As a result, the GMM estimator best fits the counterfactual price distribution by predicting relatively few consumers investing zero search effort (i.e., $u_i > \overline{u}$), leading to a smooth decline in the merger effect as a function of u_i.

Finally, panels A and B of Figure 3 plot the density of the marginal search cost and search cost level post-merger, respectively. As the table suggested, the distribution of consumers' marginal cost of effort u_i is highly dispersed, and skewed to the right. For the modal consumer the marginal cost corresponds to about 17 bps, while the average is equal to 27.2 bps.

The distribution shown in Figure 3 panel B suggests that search cost levels are also large. Post-merger, the average cost is equal to 35 bps, with a standard deviation of nearly 20 bps. To calculate the corresponding dollar value, we calculate the implicit increase in monthly payments. For an average loan size of $152,000 amortized over 25 years, an increase in the mortgage rate from 7 percent to 7.35 percent would lead to monthly payments that are $33 larger. Discounted over a 60 months period using

Table 6—Counterfactual and Decomposition Analysis

	Observed search cost distribution			Low search cost distribution		$\Delta\alpha$	Δm
	$\alpha(u)$	$\frac{\alpha(u)}{\alpha(0)}$	$m(u\|n)$	$\alpha(u')$	$m(u'\|n)$	(4)–(1)	(5)–(3)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Distribution							
$u_{5\%}$	9.46	0.89	58.56	9.86	53.93	0.40	−4.63
$u_{25\%}$	7.31	0.69	83.21	8.43	70.30	1.13	−12.91
$u_{50\%}$	5.36	0.50	105.44	7.14	85.07	1.78	−20.37
$u_{75\%}$	2.58	0.24	137.32	5.29	106.24	2.72	−31.08
$u_{95\%}$	0.00	0.00	181.42	2.74	135.46	2.74	−45.96
Average	5.12	0.48	109.62	6.90	87.83	1.78	−21.79
CV	51.96		30.49	28.06	25.24	−23.90	−5.25

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

a 5 percent annual discount rate the average search cost level corresponds to an implicit sunk cost of \$1,760.

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

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

Given our estimate of σ_ϵ, the pure market power effect of the merger corresponds to a 10.66 bps increase.[29] This is the rate increase experienced by consumers expecting to receive the maximum number of quotes. Since consumers have heterogeneous search costs, this represents an upper bound for the effect of the merger on transaction rates. Column 2 in Table 6 reports the ratio of predicted rate increases to the increase of a consumer with zero search cost $(\alpha(0))$. Consumers in the fifth percentile of the search cost distribution experience a price increase that represents 90 percent of the actual market power increase. In contrast, consumers with search costs larger than

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

the seventy-fifth percentile experience a price increase that is less than 30 percent of the upper bound. Overall, the model predicts that the ATE underestimates the market power increase caused by the merger by 50 percent.

The relationship between the retrospective estimate of the merger impact and the actual increase in market power crucially depends on the search cost distribution. In columns 4 and 5 we simulate the equilibrium price distribution and merger effects for an alternative lower search cost distribution. Such a decrease might come about if, for example, the development of online mortgage shopping platforms lowers the cost of search for all consumers. We simulate the effect of a decrease by half for every u. Columns 6 and 7 calculate the difference between the estimated pre-merger price levels and merger effects, and the counterfactual ones.

The simulated change in the pre-merger price distribution (column 7) shows that a 50 percent reduction in search cost would lead to large reductions in the level and dispersion of mortgage rates; on average negotiated margins would fall nearly 20 percent, while the dispersion of prices would fall by 15 percent. These changes mostly come from the fact that consumers at the top of the price distribution are now able to negotiate substantially lower rates; more than 45 bps at the very top.

While this suggests that consumers benefit greatly from a reduction in search costs, we also estimate that the price effect of mergers in low search cost environments would be much more important. On average, the effect of the merger is 32 percent larger in a low search cost market. The effect of the merger is also much more homogenous across consumers, as illustrated by the near 50 percent reduction in the coefficient of variation of the merger effect (last row). Once again, this reduction reflects the fact that consumers at the top of the distribution are now much more affected by the merger, since their outside option depends more on n.

Overall, our results suggest that policymakers evaluating mergers in markets with negotiated prices should worry, not only, about the pure market-power effect of mergers, measured here by the number of lenders and importance of heterogeneity (i.e., σ_ϵ), but also about the distribution of search frictions. In an environment with high search costs, firms enjoy large profits from being able to discriminate across consumers, despite the fact that mergers have little impact on average prices. This implies that researchers and policymakers studying the effect of approved mergers retrospectively should look for heterogenous effects, and concentrate their attention on the effect of mergers among lower percentiles of the price distribution to get a better understanding of the importance of market power.

C. *Discussion of the Symmetric Information Assumption*

As we discussed in Section IVA, when information is asymmetric a fraction of consumers will reject the initial offer and gather multiple quotes. This prediction is consistent with the fact that slightly more than 50 percent of consumers are observed to visit multiple lenders before accepting an offer.

In general, when a fraction of consumers searches, mergers affect both the negotiated price level and the incentives for consumers to search by lowering the search probability. This leads to three sources of price increases resulting from the merger: (i) increase in the auction price for searchers, (ii) positive average difference between the initial quote and the auction price for consumers who go from searching to non-searching, and (iii) increase in the initial quote for non-searchers. The net effect of mergers on prices is therefore positive on average, but can lead to more or less dispersion depending the relative magnitude of the market-power and price discrimination effects (i.e., initial quote). In other words, the extent to which the ATE over- or underestimates market power, and the direction in which dispersion changes, are both model and parameter specific.

Recall, that the presence of multiple unobserved attributes violates the assumption of scalar unobserved heterogeneity necessary to apply the CiC estimator. In other words, the reduced form pricing equation associated with the asymmetric information model is not consistent with the counterfactual price distribution estimated above. Nonetheless, we can simulate the model, and compare its predictions with the distributional effects of the merger. We present the results of this exercise in online Appendix F.

The asymmetric information model successfully reproduces the key distributional effects of the merger: (i) decrease in dispersion, (ii) zero effect at the top of the price distribution, and (iii) positive price increase for consumers at the bottom and middle of the distribution. In addition, the model predicts that roughly 55 percent of the population gathers more than one quote. Finally, the model predicts that the ATE of the merger underestimates the increase in market power by 24 percent. While the magnitude of this prediction is lower than what we find with complete information (i.e., 50 percent), it does not qualitatively change our conclusions.

V. Conclusion

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

We find that the average effect of the merger yields a statistically significant increase in interest rates. This finding, however, masks important heterogeneity. Some borrowers pay significantly higher rates following the merger, while others are barely affected. The evidence we present suggests that much of the heterogeneity in rates can be explained by differences in search costs and negotiation ability. Borrowers at the top of the price distribution, those with high search costs/bargaining abilities, are not affected by the merger, while those lower in the price distribution are affected. Together these results imply that price dispersion falls as a result of the merger.

Our findings on the heterogeneous effect of the merger imply that competition appears to have no impact on rates for consumers who are the most adversely affected by price discrimination; that is borrowers at the top part of residual price distribution. These results have important implications for the design of mortgage market policies. If the objective is to support vulnerable borrowers (those paying the highest rates), policies designed to increase competition, or to prevent increases in concentration, such as restrictions on merger activity or prevention of bank failure may be ineffectual. Instead, what is required are policies designed to help borrowers search for and negotiate better terms.

A 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.

Mortgage brokers can also lower search costs. Brokers in Canada have fiduciary duty, and can search over many banks for the best quote. Impediments to this search process, for example, lenders excluding brokers, can decrease the bargaining power of borrowers by limiting the choice set. Policies that limit exclusion should lower search costs via the broker channel.

REFERENCES

Allen, Jason, Robert Clark, and Jean-François Houde. 2014a. "Search Frictions and Market Power in Negotiated Price Markets." National Bureau of Economic Research Working Paper 19883.

Allen, Jason, Robert Clark, and Jean-François Houde. 2014b. "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.104.10.3365.

Allen, Jason, Robert Clark, and Jean-François Houde. 2014c. "Price Dispersion in Mortgage Markets." *Journal of Industrial Economics* 62(3): 377–416.

Ashenfelter, Orley C., Daniel S. Hosken, and Matthew C. Weinberg. 2009. "Generating Evidence to Guide Merger Enforcement." *Competition Policy International* 5 (1): 67–85.

Ashenfelter, Orley C., Daniel S. Hosken, and Matthew C. Weinberg. 2013. "The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag-Whirlpool." *American Economic Journal: Economic Policy* 5 (1): 239–61.

Athey, Susan, and Guido W. Imbens. 2006. "Identification and Inference in Nonlinear Difference-in-Difference Models." *Econometrica* 74 (2): 431–97.

Berger, Allen N., Rebecca S. Demsetz, and Philip E. Strahan. 1999. "The Consolidation of the Financial Services Industry: Causes, Consequences, and Implications for the Future." *Journal of Banking and Finance* 23 (2–4): 135–94.

Bester, Helmut. 1988. "Bargaining, Search Costs and Equilibrium Price Distributions." *Review of Economic Studies* 55 (2): 201–14.

Borenstein, Severin. 1985. "Price Discrimination in Free-Entry Markets." *RAND Journal of Economics* 16 (3): 380–97.

Borenstein, Severin, and Nancy L. Rose. 1994. "Competition and Price Dispersion in the U.S. Airline Industry." *Journal of Political Economy* 102 (4): 653–83.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer. 2006. "$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions." *American Economic Review* 96 (4): 1253–70.

Chatterjee, Kalyan, and Ching Chyi Lee. 1998. "Bargaining and Search with Incomplete Information about Outside Options." *Games and Economic Behavior* 22 (2): 203–37.

Clayton Research. 2003. "FIRM Residential Mortgage Survey June 2003."

Crawford, Gregory S., and Ali Yurukoglu. 2012. "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review* 102 (2): 643–85.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry." *American Economic Review* 102 (2): 1161–85.

Dey, Matthew S., and Christopher J. Flinn. 2005. "An Equilibrium Model of Health Insurance Provision and Wage Determination." *Econometrica* 73 (2): 571–627.

Edelberg, Wendy. 2006. "Risk-Based Pricing of Interest Rates for Consumer Loans." *Journal of Monetary Economics* 53 (8): 2283–98.

Einav, Liran, Mark Jenkins, and Jonathan Levin. 2012. "Contract Pricing in Consumer Credit Markets." *Econometrica* 80 (4): 1387–432.

Focarelli, Dario, and Fabio Panetta. 2003. "Are Mergers Beneficial to Consumers? Evidence from the Market for Bank Deposits." *American Economic Review* 93 (4): 1152–72.

Gerardi, Kristopher S., and Adam Hale Shapiro. 2009. "Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry." *Journal of Political Economy* 117 (1): 1–37.

Goldberg, Pinelopi Koujianou. 1996. "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey." *Journal of Political Economy* 104 (3): 622–54.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2013. "Mergers When Prices are Negotiated: Evidence from the Hospital Industry." National Bureau of Economic Research Working Paper 18875.

Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review* 103 (1): 145–77.

Hastings, Justine S. 2004. "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California." *American Economic Review* 94 (1): 317–28.

Heckman, James J., Hidehiko Ichimura, and Petra E. Todd. 1997. "Matching as an Econometric Evaluations Estimator: Evidence from Evaluating a Job Training Programme." *Review of Economic Studies* 64 (4): 605–54.

Hendel, Igal, Aviv Nevo, and François Ortalo-Magne. 2009. "The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison.com." *American Economic Review* 99 (5): 1878–98.

Holmes, Thomas J. 1989. "The Effects of Third-Degree Price Discrimination in Oligopoly." *American Economic Review* 79 (1): 244–50.

Hong, Han, and Matthew Shum. 2006. "Using Price Distributions to Estimate Search Costs." *RAND Journal of Economics* 37 (2): 257–75.

Hortaçsu, Ali, and Chad 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* 119 (2): 403–56.

Hortaçsu, Ali, and Chad Syverson. 2007. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." *Journal of Political Economy* 115 (2): 250–301.

Houde, Jean-François. 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review* 102 (5): 2147–82.

Huynh, Kim P., David T. Jacho-Chávez, and Marcel C. Voia. 2011. "Nonlinear Difference-in-Difference Treatment Effect Estimation: A Distributional Analysis." In *Advances in Econometrics: Missing Data Methods: Cross-sectional Methods and Applications*. Vol. 27A, edited by David M. Drukker, 247–68. Bingley, UK: Emerald Group Publishing Limited.

KPMG LLP. 2008. *Canada Mortgage Bonds Program Evaluation*. Canada Mortgage and Housing Corporation Report. http://www.cmhc.ca/en/hoficlincl/in/camobo/upload/CMB-Evaluation-Jun08.pdf (accessed June 2011).

Janssen, Maarten C. W., and José Luis Moraga-González. 2004. "Strategic Pricing, Consumer Search and the Number of Firms." *Review of Economic Studies* 71 (4): 1089–118.

Langer, Ashley. 2012. "Demographic Preferences and Price Discrimination in New Vehicle Sales." Unpublished.

Lee, Jinkook, and Jeanne M. Hogarth. 2000. "Consumer Information Search for Home Mortgages: Who, What, How Much, and What Else?" *Financial Services Review* 9 (3): 277–93.

Lewis, Matthew S. 2008. "Price Dispersion and Competition with Differentiated Sellers." *Journal of Industrial Economics* 56 (3): 654–78.

Lewis, Matthew S., and Kevin E. Pflum. Forthcoming. "Diagnosing Hospital System Bargaining Power in Managed Care Networks." *American Economic Journal: Economic Policy*.

Morton, Fiona Scott, Florian Zettelmeyer, and Jorge 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 (1): 65–92.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70 (6): 2295–350.

Prager, Robin A., and Timothy H. Hannan. 1998. "Do Substantial Horizontal Mergers Generate Significant Price Effects? Evidence from the Banking Industry." *Journal of Industrial Economics* 46 (4): 433–52.

Sapienza, Paola. 2002. "The Effects of Banking Mergers on Loan Contracts." *Journal of Finance* 57 (1): 329–67.

Sorensen, Alan T. 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *Journal of Political Economy* 108 (4): 833–50.

Stigler, G. J. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3): 213–25.

Stole, Lars A. 2007. *"Price Discrimination and Competition."* In *Handbook of Industrial Organization*. Vol. 3, edited by Mark Armstrong and Robert H. Porter, 2221–99. Amsterdam: Elsevier B.V.

Town, Robert, and Gregory Vistnes. 2001. "Hospital Competition in HMO Networks." *Journal of Health Economics* 20 (5): 733–53.

Varian, Hal. 1980. "A Model of Sales." *American Economic Review* 70 (4): 651–59.

Wildenbeest, Matthijs R. 2011. "An Empirical Model of Search with Vertically Differentiated Products." *RAND Journal of Economics* 42 (4): 729–57.

Wolinsky, Asher. 1987. "Matching, Search, and Bargaining." *Journal of Economic Theory* 42 (2): 311–33.

Woodward, Susan E., and Robert E. Hall. 2012. "Diagnosing Consumer Confusion and Sub-optimal Shopping Effort: Theory and Mortgage-Market Evidence." *American Economic Review* 102 (7): 3249–76.