The Effects of a Targeted Financial Constraint on the Housing Market*

Lu Han† Chandler Lutz‡ Benjamin Sand§ Derek Stacey¶

February 28, 2021

Abstract

We study the housing market implications of financial constraints by exploiting a regulatory change that increases the downpayment requirement for homes that sell for \$1M or more. Using Toronto data, we find that the policy causes excess bunching of homes listed at \$1M, heightened bidding intensity for these homes, but only a muted response in sales. While difficult to reconcile in a frictionless market, these findings are consistent with the implications derived from an equilibrium search model with auctions and financial constraints. Our analysis points to the importance of designing macroprudential policies that recognize the strategic responses of market participants.

JEL classification: R51, R28, D83, D44, C14

Keywords: macroprudential regulation, directed search, financial constraints, bunching estimation, housing policy

*This paper was previously circulated under the title "Do Financial Constraints Cool a Housing Boom?" We gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada. We thank the editor, two anonymous referees and participants at seminars at Arizona State University, Federal Reserve Board, George Washington University, University of Wisconsin-Madison, Stockholm University, McGill University, National University of Singapore, Singapore Management University, University of Alberta, University of Colorado, University of Calgary, University of Toronto, University of Windsor, Queen's University, Copenhagen Business School, the AREUEA National Conference, Atlanta Fed Conference, Bank of Canada, Carleton Macro-Finance Workshop, pre-WFA conference at Whistler, UBC summer symposium, McMaster's AWSOME and the Annual CEA Conference. All errors are our own.

†Rotman School of Managment, University of Toronto. lu.han@rotman.utoronto.ca; Wisconsin School of Business, University of Wisconsin-Madison. lu.han@wisc.edu.

‡Securities and Exchange Commission, lutzc@sec.gov. The Securities and Exchange Commission disclaims responsibility for any private publication or statement of any SEC employee or Commissioner. This article expresses the author's views and does not necessarily reflect those of the Commission, the Commissioners, or other members of the staff.

§Department of Economics, York University, bmsand@yorku.ca.

¶Department of Economics, University of Waterloo, dstacey@uwaterloo.ca.

1 Introduction

This paper examines how financial constraints targeting a specific housing market segment impact house price formation. A growing class of "targeted" policies aim to cool a red-hot housing segment rather than the overall market. In Toronto and Vancouver, a higher down-payment has been required for homes purchased for over \$1M. In New York and London, so-called "mansion taxes" have been imposed on purchases of all homes valued over \$1M (USD) (since 1989) and over £1.5M (GBP) (since 2014), respectively.[1] While varied in their design, these policies impose additional financial constraints on prospective homebuyers in a particular segment relative to those in other segments, which in turn can affect a seller's decision to list a house, their choice of asking price, and the process of final price determination. The central role of financial constraints makes them an appealing macroprudential vehicle for policymakers to intervene in housing markets, often for the purpose of "[ensuring] that shocks from the housing sector do not spill over and threaten economic and financial stability" (IMF 2014).[2] While financial constraints represent a recurring theme in the finance literature, there remains no micro analyses of the links between financial constraints and search behavior among housing market participants. Moreover, policies targeting a particular housing segment have just begun to attract serious attention from economists (e.g., Kopczuk and Munroe 2015). This paper fills the gap in the literature by examining how financial constraints affect price formation in the targeted segment of a frictional housing market. Our empirical methodology exploits a natural experiment arising from a mortgage insurance policy change implemented in Canada in 2012. The interpretation of our results is motivated by a search-theoretic model of sellers competing for financially constrained buyers.

Canada experienced one of the world's largest modern house price booms, with house prices more than doubling between 2000 and 2012. In an effort to cool this unprecedentedly

[1]Million dollar homes are not the mansions they used to be. In Toronto, a \$1M (CAD) house represents the 86th percentile in 2012 but the 58th percentile in 2017. Furthermore, in 2019, a \$1M (USD) home represents the 52nd percentile in San Francisco and the 33rd percentile in Manhattan among homes purchased with mortgages.

[2]Kuttner and Shim (2016) document 94 actions on the loan-to-value ratio and 45 actions on the debt-service-to-income ratio in 60 countries between 1980–2012.

long boom, the government implemented the so-called "million dollar" policy that restricts access to mortgage insurance when the purchase price of a home exceeds one million Canadian dollars ($1M). Note that lenders are required to insure mortgages with loan-to-value (LTV) ratios over 80 percent. As such, the minimum downpayment jumps from 5 to 20 percent of the entire transaction price at a threshold of $1M, increasing the minimum downpayment by $150,000 for million dollar homes. The existence or absence of bunching around the threshold should provide compelling and transparent evidence about how home buyers and sellers respond to a targeted financial constraint.

Understanding the mechanisms that generate bunching requires an equilibrium analysis of a two-sided market. To this end, we preface the empirical work with a search-theoretic model that features financial constraints on the buyer side.[3] Sellers pay a cost to list their house and post an asking price, and buyers allocate themselves across sellers subject to search frictions governed by a many-to-one meeting technology. Prices are determined by an auction mechanism: a house is sold at the asking price when a single buyer arrives; but to the highest bidder when multiple buyers submit offers to purchase the same house. In that sense, our model draws from the competing auctions literature (McAfee 1993, Peters and Severinov 1997, Julien, Kennes, and King 2000, Albrecht, Gautier, and Vroman 2014, Lester, Visschers, and Wolthoff 2015). The distinguishing feature of the model is that the million dollar policy tightens the financial constraints faced by a subset of buyers and limits how much they can bid on a house.[4]

We characterize the pre- and post-policy equilibria and derive a set of empirical predictions. The post-policy equilibrium features a mass of sellers with asking prices at the $1M threshold. These price adjustments can come from either side of $1M. In some circumstances,

[3]Financial constraints and search frictions represent recurring themes in the housing literature. Financial constraints are emphasized in Stein (1995), Lamont and Stein (1999), Ortalo-Magne and Rady (2006), and Favilukis, Ludvigson, and Nieuwerburgh (2017), among others, whereas search frictions play a central role in Wheaton (1990), Williams (1995), Krainer (2001), Genesove and Han (2012), Diaz and Jerez (2013), Head, Lloyd-Ellis, and Sun (2014), and Head, Lloyd-Ellis, and Stacey (2018). The interaction between search and financial frictions is a distinguishing feature of our analysis.

[4]Others have studied auction mechanisms with financially constrained bidders (Che and Gale 1996a,b, 1998; Kotowski 2016), but to our knowledge this is the first paper to consider bidding limits in a model of competing auctions.

sellers lower the asking price from above \$1M in order to attract both constrained and unconstrained buyers to compete for their homes. In other circumstances, sellers increase the asking price from below \$1M to extract a higher payment in bilateral situations. In both cases, the policy generates an excess mass of homes listed at \$1M. As the bunching response passes through to the sales price distribution, however, the effect on sales prices is mitigated by search frictions and bidding wars. For example, even though some sellers lower the asking price to \$1M, the induced competition among constrained and unconstrained buyers creates a heated market just under \$1M that both pushes the sales price above \$1M and leads to shorter time-on-the market.

Ultimately, the magnitude of the impact of the policy on prices is an empirical question. We test the model's predictions using the 2010-2013 housing market transaction data in the Greater Toronto Area, Canada's largest housing market. This market provides an ideal setting for this study for two reasons. First, home sellers in Toronto typically initiate the search process by listing the property and specifying a date on which offers will be considered (often 5-7 days after listing). This institutional practice matches well with our model of competing auctions. Second, the million dollar policy caused two discrete changes in the market: one at the time the policy was implemented, and another at the \$1M threshold. The market thus provides a natural experimental opportunity for examining the price response to targeted financial constraints.

Figure 1 presents the distribution of *listings* (left column) and *sales* (right column) in the segments around the \$1M threshold. Panels A and B display frequency counts of asking prices in each \$5,000 dollar bin during the pre- and post-policy periods, respectively.[5] In both periods, there is a substantial mass of listings right below \$1M, possibly due to a psychological bias associated with the million dollar threshold. Panels C and F net out the time-invariant threshold effects by presenting the difference in the frequency of listings and

[5] The figure shows the raw frequency counts of Toronto homes for one year prior to the July 12th, 2012 policy implementation (the pre-policy period) and one year after the implementation (the post-policy period). The frequency counts were created by sorting the data by either asking or sales price and grouping into \$5,000 dollar bins. The figure is restricted to within \$100,000 dollars of the \$1M policy threshold.

sales between the post- and pre-policy periods, along with confidence interval bars.[6] The results are striking. First, there is a substantial and statistically significant positive jump in the number of *listings* in the \$995,000 – \$999,999 bin, suggesting that the policy induces excess bunching of listings at \$1M. Second, this excess bunching appears to come from both sides of the threshold, as reflected by the reduction in the number of homes listed in bins just to the left and right of the million dollar bin. Finally, the mass of *sales* in the million dollar segment is much less pronounced, and the difference is statistically insignificant. The evidence here, in its most descriptive form, lends support to the key implications of our model and forms the basis for our empirical estimation design.

Despite the appealing first-cut evidence presented in Figure 1, identifying the million dollar policy's impact on asking and sales prices is difficult for several reasons. First, housing composition may have shifted around the time the policy was implemented. As a result, changes to the distributions of prices between pre- and post-policy periods may simply reflect the changing characteristics of houses listed/sold rather than buyers' and sellers' responses to the policy. Second, the implementation of the policy coincided with a number of accompanying government interventions,[7] complicating the challenge of attributing any changes in the price distributions to the million dollar policy.

Our solution relies on a two-stage estimation procedure that examines changes in the price distribution. First, leveraging the richness of our data on house characteristics and using the well-known reweighting approach introduced by DiNardo, Fortin, and Lemieux (1996), we decompose the observed before-after-policy change in the distribution of house prices into: (1) a component that is due to changes in house characteristics; and (2) a component that is due to changes in the price structure. The latter represents the quality-adjusted changes in the distribution of house prices that would have prevailed between

[6] Confidence bars were created by bootstrapping 399 random samples with replacement.

[7] The law that implemented the million dollar policy also reduced the maximum amortization period from 30 years to 25 years for insured mortgages; limited the amount that households can borrow when refinancing to 80 percent (previously 85 percent); lowered the maximum total debt service ratio (all housing expenses, credit card, and car loan payments relative to income) from 45% to 44%; and set a maximum gross debt service ratio (mortgage payments, property taxes, and heating costs relative to income) at 39%. Source: "Harper Government Takes Further Action to Strengthen Canada's Housing Market." *Department of Finance Canada*, June 21, 2012.

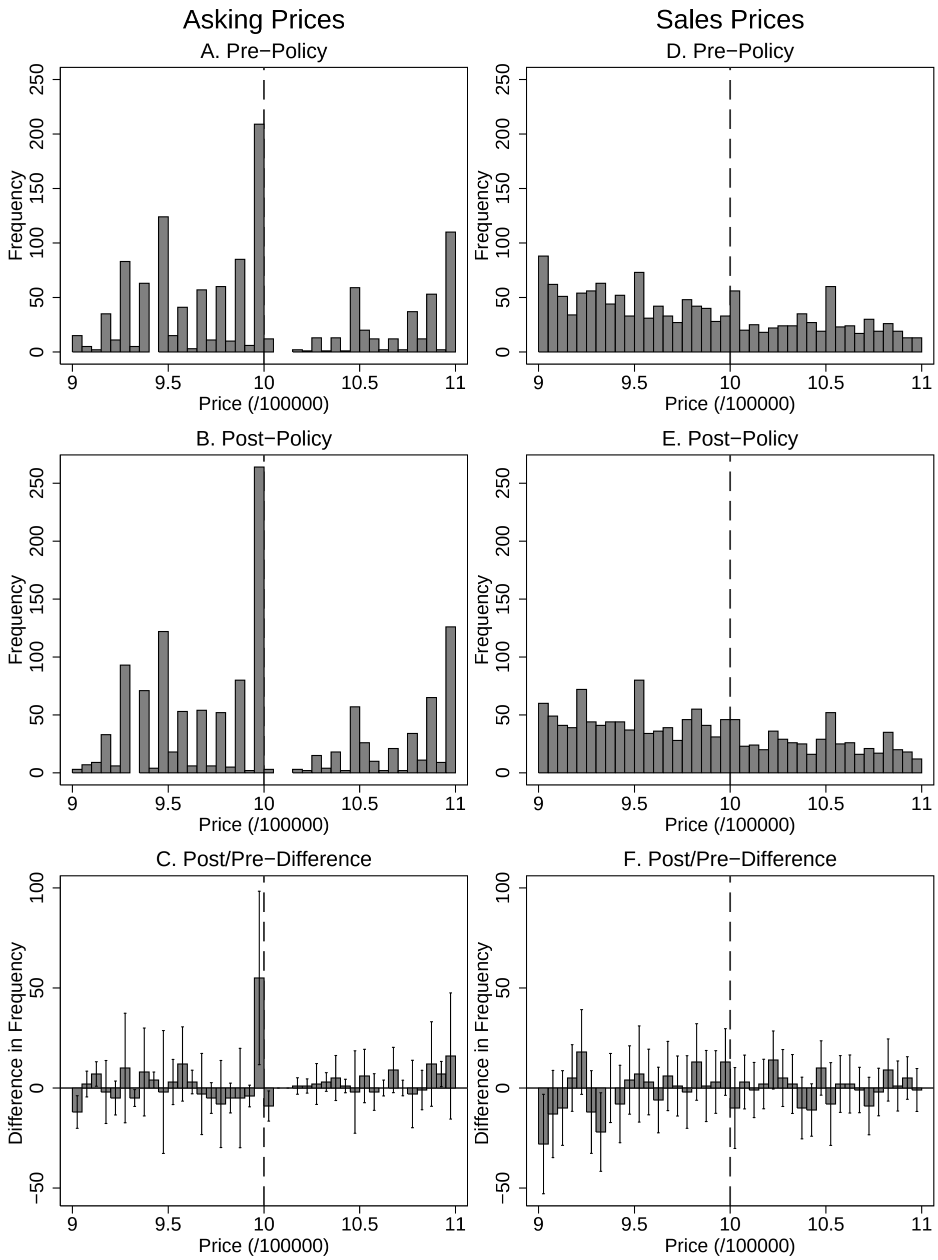

Figure 1
Frequency Counts of Asking and Sales Prices in the Pre- vs. Post-Policy Periods

Notes: The figure uses data on asking and sales prices for the city of Toronto in the year before (pre-period) and after (post-period) the implementation of the million dollar policy. Panels A, B, D, and E show frequency counts for $5,000 bins within $100,000 dollars of the policy threshold for the indicated period. Panels C and F show the difference in the frequency counts for the post- vs pre-periods in each bin. The confidence bars in Panels C and F are constructed via bootstrap for 399 random samples with replacement.

the pre- and post-policy periods if the characteristics of houses remained the same as in the pre-policy period. Next, we measure the effects of the \$1M policy by comparing the actual changes in the quality-adjusted price distributions to counterfactual changes in the quality-adjusted price distributions that would have prevailed in the absence of the policy, separately for asking and sales prices. We employ a recently developed bunching estimation approach (Chetty et al. 2011, Kleven and Waseem 2013).[8] The key idea is to use price segments that are not subject to the policy's threshold effects to form a counterfactual near the \$1M threshold. By comparing the counterfactual changes in distributions with the actual changes in distributions around \$1M, bunching estimation allows us to difference out impacts of contemporaneous factors on house prices, such as other mortgage rule changes and market trends. Further, working with *changes* in price distributions over time allows us to net out any time-invariant threshold price effects unrelated to the policy, such as psychological bias.

Our main findings are the following. In the single-family-housing market, the asking price distribution features large and sharp excess bunching right at the \$1M threshold with corresponding holes both above and below \$1M. In particular, the policy adds 86 homes to *listings* in the million dollar bin (from \$995,000 to \$999,999) in the post-policy year for the city of Toronto, which represents about a 38 percent increase relative to the number of homes that would have been listed in this \$5,000 bin in the absence of the policy. Among these, half would have otherwise been listed below \$995,000; the other half above \$1M. In contrast, the policy adds only about 11 homes to *sales* in the million dollar bin, which is economically small and statistically insignificant. These findings are robust to an extensive set of specification checks, including a counterfactual constructed using only data below \$1M, alternative functional forms, estimation windows and excluded regions, different definitions of pre- and post-policy periods, and allowing for possible spillovers near the \$1M segment. We also find similar patterns in the condominium and townhouse markets.

[8] In the context of real estate, Kopczuk and Munroe (2015) and Slemrod, Weber, and Shan (2017) analyse bunching behavior in sales volume induced by discontinuities in real-estate transfer taxes; Best et al. (2018) exploit variation in interest rates that produce notches in the loan-to-value ratio at various thresholds; and DeFusco and Paciorek (2017) estimate leverage responses to a notch created by the conforming loan limit in the U.S. Our approach differs from these related studies in that we consider a two-sided bunching estimator to accommodate both possibilities explored in our theoretical framework.

The lack of excess bunching in the sales price, together with the sharp bunching in the asking price, suggests that the intended cooling impact of the policy is mitigated by sellers' listing decisions and buyers' bidding behavior. Consistent with this interpretation, we find that housing segments right below the $1M threshold experience a shorter time-on-the-market and a higher incidence of bidding wars, confirming the notion that the million dollar policy created a "red hot" market for homes listed just below $1M.[9]

Together, our findings contribute to a better understanding of policies that use targeted financial constraints to temper a heated market segment. We find that the million dollar policy did not achieve the specific goal of cooling the housing boom in the million dollar segment. This is not because market participants did not respond to the policy. In fact, quite the opposite appears to be true: it is precisely the strategic responses of home sellers and buyers that interact to undermine the intended impact of the policy on sales prices. Our analysis thus points to the importance of designing policies that recognize the endogenous responses of buyers and sellers in terms of listing strategies, search decisions and bidding behavior.

While our main focus is on segments around $1M, we also go beyond the bunching estimation and examine the policy effects in segments further above $1M. An *extensive margin response* would imply that some transactions above $1M did not occur due to the additional financial constraint. An *intensive margin response* would imply depressed prices in at least some segments of the market above $1M. Either of these responses should manifest as a systematic discrepancy between the counterfactual post-policy price distribution and the observed post-policy price distribution for price bins above the $1M threshold. Using a distribution decomposition method, we do not find such discrepancies, which suggests that transactions above $1M are not markedly affected by the policy.

Despite failing to curb house price appreciation, the policy may have nonetheless succeeded in improving the creditworthiness of homebuyers. A key implication of the model is that, when facing the $1M policy, less constrained buyers have an advantage over con-

[9]See "Ottawa's new rules creating 'red hot' market for homes under $999,999." Financial Post. July 3, 2013.

strained ones in multiple offer situations and hence have incentive to participate in the segment below \$1M. Consistent with this, we observe a large share of homebuyers with LTV ratios below 80% even in the segment just below \$1M after the implementation of the policy. Thus, the policy improves borrower creditworthiness in segments above \$1M without compromising borrower creditworthiness in the segment slightly below \$1M via its influence on buyers' and sellers' search and listing behaviors. More broadly, by reallocating million dollar homes from more constrained to less constrained homebuyers, the policy effectively prevents lenders from making riskier loans. As such, our analysis is also related to the recent literature on macroprudential interventions in mortgage markets. Significant contributions have been made towards our understanding of how these policies affect mortgage market outcomes (Allen et al. 2016; DeFusco, Johnson, and Mondragon 2019), as well as financial stability and mortgage market efficiency (Elenev, Landvoigt, and Van Nieuwerburgh 2016; Van Bekkum, Gabarro, and Irani 2017; Elenev, Landvoigt, and Van Nieuwerburgh 2018; Acharya, Berger, and Roman 2018). From that perspective, the effects of the \$1M policy may help reduce the likelihood of mortgage crises and safeguard the stability of the financial system, which would be welfare enhancing in the long run. Studying such benefits by quantifying the effects of the policy on mortgage market outcomes may prove to be an important area for future research.[10]

2 Background

2.1 Mortgage Insurance

Mortgage insurance is an instrument used to transfer mortgage default risk from the lender to the insurer and represents a key component of housing finance in many countries including

[10] As preliminary evidence, we observe that in aggregate Canada exhibited a decrease in the fraction of new mortgage holders with a credit score below 660 after 2012. See Panel A of Figure A1 in Appendix A. We do not examine the policy effects on credit market outcomes for two reasons. First, we do not have micro-level mortgage data. Second, default is not widespread in Canada due to its highly regulated financial system. Panel B of Figure A1 shows the difference in the delinquency rates (defined as overdue on a payment by 90 days or more) between Canada and the U.S. over time. In 2012, the fraction of all mortgages with delinquencies was 7.14 percent in the U.S., but only 0.32 percent in Canada (and 0.23 percent in Toronto).

the United States, the United Kingdom, the Netherlands, Hong Kong, France, and Australia. These countries share two common features with Canada: (i) the need to insure high LTV mortgages, and (ii) the central role of the government in providing such insurance. The combination of these two requirements gives the government the ability to influence the financial constraints faced by homebuyers.

In Canada, all financial institutions regulated by the Office of the Superintendent of Financial Institutions (OSFI) are required to purchase mortgage insurance for any mortgage loan with an LTV above 80 percent. The mortgage insurance market is comprised of the government-owned Canada Mortgage and Housing Corporation (CMHC) as well as two private insurers, Genworth Financial Mortgage Insurance Company Canada and Canada Guaranty. All three institutions benefit from guarantees provided by the Canadian government and therefore are subject to federal regulations through the OFSI.

In practice, while it is possible for buyers to obtain uninsured residential mortgages with a loan-to-value ratio greater than 80 percent from unregulated lenders, we find that private lending accounted for only 4% of all loans in the Greater Toronto Area in 2013 and this sector did not experience any noticeable growth around the million dollar policy period. The reason is that, compared to traditional mortgages from regulated lenders, private mortgages on average have one-fifth duration, over three times higher interest rates, and loan amounts that are one-third of the size. Hence they operate in a small disparate niche corner of the Canadian mortgage market.[11] In addition, anecdotal evidence suggests that it is generally difficult for a borrower to obtain a second mortgage at the time of origination to reduce the downpayment of the primary loan below 20 percent in Canada, making this strategic circumvention of macroprudential regulation less of a concern. The pervasiveness of government-backed mortgage insurance within the housing finance system makes it an appealing macroprudential policy tool for influencing housing finance and housing market outcomes.

[11] See **??** for the statistics reported here and relevant discussions.

2.2 The Million Dollar Policy

Figure 2 plots the national house price indices for Canada and the U.S. reflecting Robert Shiller's observation in 2012 that "what is happening in Canada is kind of a slow-motion version of what happened in the U.S."[12] As home prices in Canada continued to escalate post-financial crisis, the Canadian government became increasingly concerned that rapid price appreciation would eventually lead to a severe housing market correction.[13] To counter the potential risks associated with this house price boom, the Canadian government implemented several rounds of housing market macroprudential regulation, all through changes to the mortgage insurance rules.[14] This paper examines the impact of the so-called "million dollar" policy that prevents regulated lenders from offering mortgage loans with LTV ratios above 80 percent when the purchase price is \$1M or more. The objective of the regulation was to curb house price appreciation in high price segments of the market and at the same time improve borrower creditworthiness. The law was announced on June 21, 2012, and effected July 9, 2012. Anecdotal evidence suggests that the announcement of the policy was largely unexpected by market participants.[15]

3 Theory

To understand how the million dollar policy affects strategies and outcomes in the housing market, we present a two-sided search model that incorporates auction mechanisms and financially constrained buyers. We characterize pre- and post-policy directed search equilibria and derive a set of empirical implications. The purpose of the model is to guide the empirical

[12] "Why a U.S.-style housing nightmare could hit Canada." *CBCNews*. September 21, 2012.

[13] In 2013, Jim Flaherty, Canada's Minister of Finance from February 2006 to March 2014, stated: "We [the Canadian government] have to watch out for bubbles - always - . . . including [in] our own Canadian residential real estate market, which I keep a sharp eye on." Sources: "Jim Flaherty vows to intervene in housing market again if needed." *The Globe and Mail*, November 12, 2013.

[14] These changes included increasing minimum down payment requirements (2008); reducing the maximum amortization period for new mortgage loans (2008, 2011, 2012); reducing the borrowing limit for mortgage refinancing (2010, 2011, 2012); increasing homeowner credit standards (2008, 2010, 2012); and limiting government backed mortgage insurance to homes with a purchase price of less than one million Canadian dollars (2012).

[15] See "High-end mortgage changes seen as return to CMHC's roots." *The Globe and Mail*, June 23, 2012.

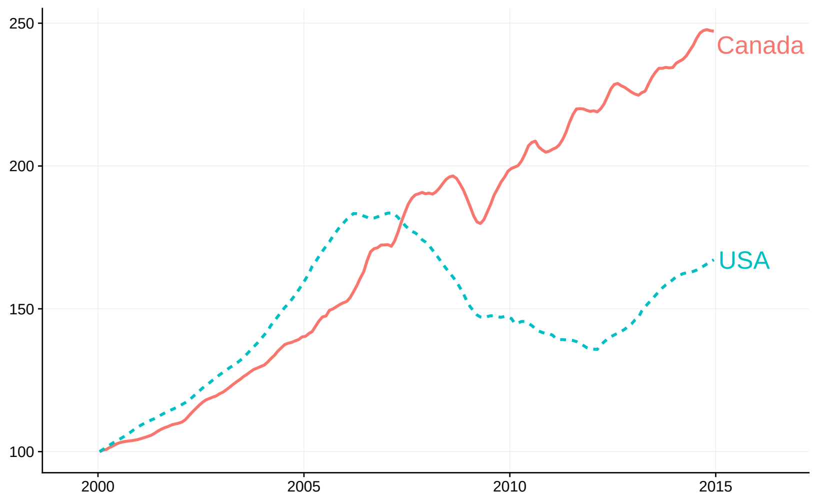

Figure 2
House Price Indices for Canada and the U.S.

Notes: Monthly house price indices from S&P Case-Shiller (US) and Teranet (Canada). All series downloaded from Datastream and are indexed to 100 in 2000. Series ID numbers: USCSHP20F and CNTNHPCMF.

analyses that follow. As such, we present a simple model of directed search with auctions and bidding limits that features heterogeneity only along the financial constraints dimension. The clean and stylized nature of the model allows for a quick understanding of the intuition underlying plausible strategic reactions among buyers and sellers to the implementation of the policy.

3.1 Environment

Agents. There is a fixed measure $\mathcal{B}$ of buyers, and a measure of sellers determined by free entry. Buyers and sellers are risk neutral. Each seller owns one indivisible house, their value of which is normalized to zero. Buyer preferences are identical; a buyer assigns value $v > 0$ to owning the home. No buyer can pay more than some fixed $u \leq v$, which can be viewed as a common income constraint or debt-service constraint.[16]

Million dollar policy. The introduction of the million dollar policy causes some buyers

[16] A non-binding constraint (i.e., $u > v$) would have the same implications as the case where $u = v$ in the analysis that follows.

to become more severely *financially constrained*. Post-policy, a fraction Λ of buyers are unable to pay more than c, where $c < u$. Parameter restrictions $c < u \leq v$ can be interpreted as follows: all buyers may be limited by their budget sets, but some are further financially constrained by a binding wealth constraint such as a minimum downpayment requirement following the implementation of the policy.[17] Buyers with financial constraint c are hereinafter referred to as *constrained* buyers, whereas buyers willing and able to pay up to u are termed *unconstrained*.

Search and matching. The matching process is subject to frictions which we model with an urn-ball meeting technology. Each buyer meets exactly one seller. From the point of view of a seller, the number of buyers she meets is a random variable that follows a Poisson distribution. The probability that a seller meets exactly $k = 0, 1, \ldots$ buyers is

$$\pi(k) = \frac{e^{-\theta}\theta^k}{k!}, \tag{1}$$

where θ is the ratio of buyers to sellers and is often termed *market tightness*. The probability that exactly j out of the k buyers are unconstrained is

$$\phi_k(j) = \binom{k}{j}(1-\lambda)^j \lambda^{k-j}, \tag{2}$$

which is the probability mass function for the binomial distribution with parameters k and $1-\lambda$, where λ is the share of constrained buyers. Search is *directed* by asking prices in the following sense: sellers post a listing containing an asking price, $p \in \mathbb{R}_+$, and buyers direct their search by focusing exclusively on listings with a particular price. As such, θ and λ are endogenous variables specific to the group of buyers and sellers searching for and asking price p.

Price determination. The price is determined in a sealed-bid second-price auction. The seller's asking price, $p \in \mathbb{R}_+$, is interpreted as the binding reserve price. If a single

[17] We model the implied bidding limit rather than the downpayment constraint explicitly. The interpretation is as follows: the discontinuous downpayment requirement at \$1M effected by the policy means that buyers with wealth levels less than \$200,000 must bid less than \$1M.

bidder submits an offer at or above p, he pays only p. In multiple offer situations, the bidder submitting the highest bid at or above p wins the house but pays either the second highest bid or the asking price, whichever is higher. When selecting among bidders with identical offers, suppose the seller picks one of the winning bidders at random with equal probability.

Free entry. The measure of sellers is determined by free entry so that overall market tightness is endogenous. Supply side participation in the market requires payment of a fixed cost x, where $0 < x < c$. It is worthwhile to enter the market as a seller if and only if the expected revenue exceeds the listing cost.

3.2 Equilibrium

3.2.1 The Auction

When a seller meets k buyers, the auction mechanism described above determines a game of incomplete information because bids are sealed and bidding limits are private. In a symmetric Bayesian-Nash equilibrium, it is a dominant strategy for buyers to bid their maximum amount, c or u. When $p > c$ ($p > u$), bidding limits preclude constrained (and unconstrained) buyers from submitting sensible offers.

3.2.2 Expected payoffs

Expected payoffs are computed taking into account the matching probabilities in (1) and (2). These payoffs, however, are markedly different depending on whether the asking price, p, is above or below a buyer's ability to pay. Each case is considered separately in Appendix B.1. In the *submarket* associated with asking price p and characterized by market tightness θ and buyer composition λ, let $V^s(p, \lambda, \theta)$ denote the sellers' expected net payoff. Similarly, let $V^c(p, \lambda, \theta)$ and $V^u(p, \lambda, \theta)$ denote the expected payoffs for constrained and unconstrained buyers.

For example, if the asking price is low enough to elicit bids from both unconstrained and constrained buyers, the seller's expected net payoff is

$$V^s(p \leq c, \lambda, \theta) = -x + \pi(1)p + \sum_{k=2}^{\infty} \pi(k) \left\{ [\phi_k(0) + \phi_k(1)]\, c + \sum_{j=2}^{k} \phi_k(j) u \right\}.$$

Substituting expressions for $\pi(k)$ and $\phi_k(j)$ and recognizing the power series expansion of the exponential function, the closed-form expression is

$$\begin{aligned} V^s(p \leq c, \lambda, \theta) = -x + \theta e^{-\theta} p + \left[1 - e^{-\theta} - \theta e^{-\theta}\right] c \\ + \left[1 - e^{-(1-\lambda)\theta} - (1-\lambda)\theta e^{-(1-\lambda)\theta}\right] (u - c). \end{aligned}$$

The second term reflects the surplus from a transaction if they meet only one buyer. The third and fourth terms reflect the surplus when matched with two or more buyers, where the last term is specifically the additional surplus when two or more bidders are unconstrained.

The expected payoff for a buyer, upon meeting a particular seller, takes into account the possibility that the seller meets other constrained and/or unconstrained buyers as per the probabilities in (1) and (2). The expected payoff for a constrained buyer in this case is

$$V^c(p \leq c, \lambda, \theta) = \pi(0)(v - p) + \sum_{k=1}^{\infty} \pi(k) \phi_k(0) \frac{v - c}{k + 1}$$

and the closed-form expression is

$$V^c(p \leq c, \lambda, \theta) = \frac{e^{-(1-\lambda)\theta} - e^{-\theta}}{\lambda\theta}(v - c) + e^{-\theta}(c - p).$$

The first term is the expected surplus when competing for the house with other constrained bidders; the last term reflects the possibility of being the only buyer. Note that whenever an unconstrained buyer visits the same seller, the constrained buyer is outbid with certainty and loses the opportunity to purchase the house. Finally, the expected payoff for an unconstrained buyer can be similarly derived to obtain

$$
\begin{aligned}
V^u(p \leq c, \lambda, \theta) &= \pi(0)(v-p) + \sum_{k=1}^{\infty} \pi(k) \left[\phi_k(0)(v-c) + \sum_{j=1}^{k} \phi_k(j) \frac{v-u}{j+1} \right] \\
&= \frac{1-e^{-(1-\lambda)\theta}}{(1-\lambda)\theta}(v-u) + e^{-(1-\lambda)\theta}(u-c) + e^{-\theta}(c-p).
\end{aligned}
$$

The first term is the expected surplus when competing for the house with other unconstrained bidders, and the second term is the additional surplus when competing with constrained bidders only. In that scenario, the unconstrained bidder wins the auction by outbidding the other constrained buyers, but pays only c in the second-price auction. The third term represents the additional payoff for a monopsonist. Closed-form solutions for the other cases are derived in Appendix B.1.

3.2.3 Directed Search

Agents perceive that both market tightness and the composition of buyers depend on the asking price. To capture this, suppose agents expect each asking price p to be associated with a particular ratio of buyers to sellers $\theta(p)$ and fraction of constrained buyers $\lambda(p)$. We will refer to the triple $(p, \lambda(p), \theta(p))$ as *submarket* p. When contemplating a change to her asking price, a seller anticipates a corresponding change in the matching probabilities and bidding war intensity via changes in tightness and buyer composition. This is the sense in which search is *directed*. It is convenient to define $V^i(p) = V^i(p, \lambda(p), \theta(p))$ for $i \in \{s, u, c\}$.

Definition 1. *A* directed search equilibrium *(DSE) is a set of asking prices $\mathbb{P} \subset \mathbb{R}_+$; a distribution of sellers σ on $\mathbb{R}_+$ with support $\mathbb{P}$, a function for market tightness $\theta : \mathbb{R}_+ \to \mathbb{R}_+ \cup +\infty$, a function for the composition of buyers $\lambda : \mathbb{R}_+ \to [0, 1]$, and a pair of values $\{\bar{V}^u, \bar{V}^c\}$ such that:*

1. *optimization:*
   - *(i) sellers: $\forall p \in \mathbb{R}_+$, $V^s(p) \leq 0$ (with equality if $p \in \mathbb{P}$);*
   - *(ii) unconstrained buyers: $\forall p \in \mathbb{R}_+$, $V^u(p) \leq \bar{V}^u$ (with equality if $\theta(p) > 0$ and $\lambda(p) < 1$);*

(iii) constrained buyers: $\forall p \in \mathbb{R}_+$, $V^c(p) \leq \bar{V}^c$ *(with equality if* $\theta(p) > 0$ *and* $\lambda(p) > 0$*);*

where $\bar{V}^i = \max_{p \in \mathbb{P}} V^i(p)$ *for* $i \in \{u, c\}$*; and*

2. *market clearing:*

$$\int_{\mathbb{P}} \theta(p)\, d\sigma(p) = \mathcal{B} \quad \textit{and} \quad \int_{\mathbb{P}} \lambda(p)\theta(p)\, d\sigma(p) = \Lambda\mathcal{B}.$$

The definition of a DSE is such that for every $p \in \mathbb{R}_+$, there is a $\theta(p)$ and a $\lambda(p)$. Part 1(i) states that θ is derived from the free entry of sellers for active submarkets (*i.e.*, for all $p \in \mathbb{P}$). Similarly, parts 1(ii) and 1(iii) require that, for active submarkets, λ is derived from the composition of buyers that find it optimal to search in that submarket. For inactive submarkets, parts 1(ii) and 1(iii) further establish that θ and λ are determined by the optimal sorting of buyers so that off-equilibrium beliefs are pinned down by the following requirement: if a small measure of sellers deviate by posting asking price $p \notin \mathbb{P}$, and buyers optimally sort among submarkets $p \cup \mathbb{P}$, then those buyers willing to accept the highest buyer-seller ratio at price p determine both the composition of buyers $\lambda(p)$ and the buyer-seller ratio $\theta(p)$. If neither type of buyer finds asking price p acceptable for any positive buyer-seller ratio, then $\theta(p) = 0$, which is interpreted as no positive measure of buyers willing to search in submarket p. The requirement in part 1(i) that $V^s(p) \leq 0$ for $p \notin \mathbb{P}$ guarantees that no deviation to an off-equilibrium asking price is worthwhile from a seller's perspective. Finally, part 2 of the definition makes certain that all buyers search.

3.2.4 Pre-Policy Directed Search Equilibrium

We first consider the initial setting with identically unconstrained buyers by setting $\Lambda = 0$.[18] Buyers in this environment direct their search by targeting the asking price that maximizes their expected payoff. Because the buyer correctly anticipates the free entry of sellers, the

[18] A DSE when $\Lambda = 0$ is defined according to Definition 1 except that we impose $\lambda(p) = 0$ for all $p \in \mathbb{R}_+$ and ignore condition 1(iii).

search problem can be written

$$\bar{V}^u = \max_{p,\theta} V^u(p, 0, \theta) \quad \text{s.t.} \quad V^s(p, 0, \theta) = 0. \tag{P$_0$}$$

We construct a DSE with a single active submarket with asking price and market tightness determined by the solution to problem P_0, denoted $\{p_0, 0, \theta_0\}$.[19] Given the auction mechanism and the role of the asking price, a strictly positive expected surplus from searching requires $p \leq u$. If the solution is interior it satisfies the following first-order condition and the constraint:

$$x = [1 - e^{-\theta_u^*} - \theta_u^* e^{-\theta_u^*}]v \tag{3}$$

$$\theta_u^* e^{-\theta_u^*} p_u^* = [1 - e^{-\theta_u^*} - \theta_u^* e^{-\theta_u^*}](v - u). \tag{4}$$

If this solution is infeasible because of financial limit u, the solution is instead u and θ_u, where θ_u satisfies the free entry condition $V^s(u, 0, \theta_u) = 0$, or

$$x = [1 - e^{-\theta_u}]u. \tag{5}$$

The solution to problem P_0 can therefore be summarized as $p_0 = \min\{p_u^*, u\}$ and θ_0 satisfying $V^s(p_0, 0, \theta_0) = 0$.

The following proposition provides a partial characterization of the pre-policy DSE constructed using this solution as per the algorithm in Appendix B.2.

Proposition 1. *There is a DSE with* $\mathbb{P} = \{p_0\}$, $\theta(p_0) = \theta_0$, $\sigma(p_0) = \mathcal{B}/\theta_0$ *and* $\bar{V}^u = V^u(p_0, 0, \theta_0)$.

As buyers' ability to pay approaches their willingness to pay (i.e., as $u \to v$), the equi-

[19] The same active submarket can instead be determined by solving the seller's price posting problem and imposing free entry. Specifically, sellers set an asking price to maximize their expected payoff subject to buyers achieving their market value $\bar{V}^u$. The seller's asking price setting problem is therefore

$$\max_{p,\theta} V^s(p, 0, \theta) \quad \text{s.t.} \quad V^u(p, 0, \theta) = \bar{V}^u. \tag{P$'_0$}$$

librium asking price tends to zero (i.e., $p_0 = p_u^* \to 0$), which is the seller's reservation value. This aligns with standard results in the competing auctions literature in the absence of bidding limits (McAfee 1993; Peters and Severinov 1997; Albrecht, Gautier, and Vroman 2014; Lester, Visschers, and Wolthoff 2015). When buyers' bidding strategies are somewhat limited (i.e., $p_0 = p_u^* \leq u < v$), sellers set a higher asking price to capture more of the surplus in a bilateral match. The equilibrium asking price is such that the additional bilateral sales revenue exactly compensates for the unseized portion of the match surplus when two or more buyers submit offers but are unable to bid up to their full valuation. This the economic interpretation of equation (4). When buyers' bidding strategies are too severely restricted (i.e., $p_0 = u < p_u^*$), the seller's choice of asking price is constrained by the limited financial means of prospective buyers. Asking prices in equilibrium are then set to the maximum amount, namely u. In this case, a seller's expected share of the match surplus is diminished, and consequently fewer sellers choose to participate in the market (i.e., $\theta_u > \theta_u^*$).

If $p_0 = p_u^* \leq u$, the equilibrium expected payoff $\bar{V}^u$ is independent of u (in particular, $\bar{V}^u = \theta_u^* e^{-\theta_u^*} v$). As long as the constraint remains relatively mild, a change to buyers' ability to pay, u, will cause the equilibrium asking price to adjust in such a way that market tightness and the expected sales price remain unchanged. This reflects the fact that the financial constraint does not affect the *incentive* to search. When $p_0 = u < p_u^*$, the constraint is sufficiently severe that it affects the *ability* to search in that it shuts down the submarket that would otherwise achieve the mutually desirable trade-off between market tightness and expected price. This feature highlights the distinction between the roles of financial constraints and reservation values, since a change to buyers' willingness to pay, v, would affect the incentive to search, the equilibrium expected payoff, and the equilibrium trade-off between market tightness and expected sales price.

3.2.5 Post-Policy Directed Search Equilibrium

As in the previous section, an active submarket with $p \leq c$ is determined by an optimal search strategy. The search problem of a constrained buyer takes into account the participation of both sellers and unconstrained buyers:

$$\bar{V}^c = \max_{p,\lambda,\theta} V^c(p,\lambda,\theta) \quad \text{s.t.} \quad V^s(p,\lambda,\theta) = 0 \quad \text{and} \quad V^u(p,\lambda,\theta) \geq \bar{V}^u. \tag{P$_1$}$$

Let $\{p_1, \lambda_1, \theta_1\}$ denote the solution to problem P_1 when $\bar{V}^u$ is set equal to the maximized objective of problem P_0. The bidding limit once again limits the set of worthwhile submarkets. In particular, the optimal submarket for constrained buyers must feature an asking price less than or equal to c. If the solution is interior, it satisfies the two constraints with equality and a first-order condition derived in Appendix B.3. This interior solution is denoted $\{p_c^*, \lambda_c^*, \theta_c^*\}$. The corner solution is denoted $\{c, \lambda_c, \theta_c\}$, where λ_c and θ_c satisfy the free entry condition $V^s(c, \lambda_c, \theta_c) = 0$ and an indifference condition for unconstrained buyers $V^u(c, \lambda_c, \theta_c) = \bar{V}^u$. In summary, the solution to problem P_1 is $p_1 = \min\{p_c^*, c\}$ with λ_1 and θ_1 satisfying $V^s(p_1, \lambda_1, \theta_1) = 0$ and $V^u(p_1, \lambda_1, \theta_1) = \bar{V}^u$.

As long as the aggregate share of constrained buyers, Λ, does not exceed λ_1, we can construct an equilibrium with two active submarkets associated with the asking prices obtained by solving problems P_0 and P_1 in the manner described above.

Proposition 2. *Suppose $\Lambda \leq \lambda_1$. There is a DSE with $\mathbb{P} = \{p_0, p_1\}$, $\lambda(p_0) = 0$, $\lambda(p_1) = \lambda_1$, $\theta(p_0) = \theta_0$, $\theta(p_1) = \theta_1$, $\sigma(p_0) = (\lambda_1 - \Lambda)\mathcal{B}/(\lambda_1\theta_0)$, $\sigma(p_1) = \Lambda\mathcal{B}/(\lambda_1\theta_1)$, $\bar{V}^c = V^c(p_1, \lambda_1, \theta_1)$ and $\bar{V}^u = V^u(p_0, 0, \theta_0) = V^u(p_1, \lambda_1, \theta_1)$.*

Intuitively, constrained buyers would prefer to avoid competition from unconstrained buyers because they can out-bid them. For the same reason, some unconstrained buyers prefer to search alongside constrained buyers. The equilibrium search decisions of constrained buyers takes into account the unavoidable competition from unconstrained buyers to achieve the optimal balance between price, market tightness, and the bidding limits of potential auction participants.

The incentive to search alongside constrained buyers in a submarket distorted by a binding financial constraint is increasing in the share of buyers constrained by the policy. If the fraction of constrained buyers is not too high (i.e., $\Lambda < \lambda_1$), the DSE features partial pooling.

That is, only some unconstrained buyers search for homes priced at p_1 while the rest search in submarket p_0.[20] As $\Lambda \to \lambda_1$, it can be shown that $\sigma(p_0) \to 0$ and the DSE converges to one of full pooling, with all buyers and sellers participating in submarket p_1. Finally, if $\Lambda > \lambda_1$, market clearing (part 2 of Definition 1) is incompatible with unconstrained buyer indifference between these two submarkets, which begets the possibility of full pooling with unconstrained buyers strictly preferring to pool with constrained buyers. We restrict attention to settings with $\Lambda \leq \lambda_1$ for the analytical characterization of equilibrium and rely on numerical results for settings with $\Lambda > \lambda_1$.[21]

3.3 Empirical Predictions

This section summarizes the housing market implications of the million dollar policy by comparing the pre- and post-policy directed search equilibria. Since financial constraint c is intended to represent the maximum ability to pay among buyers affected by the million dollar policy, parameter c corresponds to the \$1M threshold and Λ reflects the share of potential buyers with insufficient wealth from which to draw a 20 percent downpayment.[22]

There are four possible cases to consider depending on whether financial constraints u and c lead to corner solutions to problems P_0 and P_1. In this section we focus on the most empirically relevant case where the financial constraint is slack in problem P_0 but binds in problem P_1. In other words, we consider the possibility that pre-existing financial constraints are mild, but that the additional financial constraint imposed by the policy is sufficiently

[20]The partial separation of unconstrained buyers in this case arises because the source of heterogeneity is bidders' *ability to pay* and not their *willingness to pay*. A similar environment with heterogeneous valuations rather than financial means would not necessarily deliver more than one active submarket in equilibrium (Cai, Gautier, and Wolthoff 2017).

[21]We construct fully pooling DSE numerically when $\Lambda > \lambda_1$ by increasing $\bar{V}^u$ above the maximized objective of problem P_0 until the share of constrained buyers in the submarket that solves problem P_1 is exactly Λ. A thorough analysis of such DSE would require abandoning the analytical convenience of block recursivity (i.e., the feature that equilibrium values and optimal strategies do not depend on the overall composition of buyers). We sacrifice completeness for conciseness and convenience by restricting the set of analytical results to settings with $\Lambda \leq \lambda_1$.

[22]Since the million dollar policy effectively imposes a 20 percent downpayment requirement when the purchase price is \$1M or more, c more precisely represents a bidding limit of \$999,999 expressed relative to the seller's reservation value. So as to avoid awkward wording, we hereinafter use the \$1M threshold to refer to the price point *just under* \$1M.

severe. Under this assumption, the equilibrium asking prices are $p_0 = p_u^*$ and $p_1 = c$. There are still two possible subcases, namely (i) $p_u^* \leq c$ and (ii) $p_u^* > c$, which we use to derive several testable predictions that we bring to the data in Section 5.

Prediction 1. *The million dollar policy motivates some sellers to change their asking price to \$1M. This asking price response corresponds to "bunching from below" if $p_0 < p_1$, or "bunching from above" if $p_0 > p_1$.*

As per Propositions 1 and 2, the set of asking prices changes from just $\mathbb{P} = \{p_0\}$ pre-policy to $\mathbb{P} = \{p_0, p_1\}$ post-policy. Following the introduction of the policy, some or all sellers find it optimal to target buyers of both types by asking price $p_1 = c$. The million dollar policy can thus induce a strategic response among sellers in market segments near the newly imposed financial constraint. If $p_0 < p_1$, some sellers who would have otherwise listed below c respond to the policy by increasing their asking price to the threshold. The intuition for *bunching from below* is the following: as buyers become more constrained, the distribution of possible sales prices features fewer extreme prices at the high end. Sellers respond by raising their asking price to effectively truncate the distribution of prices from below. The higher price in a bilateral situation can offset (in expectation) the unseized sales revenue in multiple offer situations arising from the additional financial constraint. Constrained buyers tolerate the higher asking price because they face less severe competition from unconstrained bidders in submarket p_1. If instead $p_0 > p_1$, the policy induces some sellers who would have otherwise listed above c to drop their asking price to exactly equal the threshold. In the case of *bunching from above*, the reduction in asking prices is designed to attract constrained buyers. Because there is pooling of both buyer types in submarket p_1, these sellers may still match with unconstrained buyers and sell for a price above c.

Prediction 2. *Bunching at \$1M in asking prices only partially passes through to the sales price distribution because of search frictions and bidding wars.*

The frictional matching process between buyers and sellers results in some homes failing to sell. With probability $e^{-\theta_1}$, a seller listing a home post-policy at price $p_1 = c$ does not meet even a single buyer. The auction mechanism further reduces the mass of sales relative to listings at price c. With probability $1 - e^{-(1-\lambda_1)\theta_1} - (1-\lambda_1)\theta_1 e^{-(1-\lambda_1)\theta_1}$, competition among unconstrained bidders in submarket p_1 escalates the sales price up to u.

This bidding war effect intensifies (diminishes) in response to the million dollar policy if $p_0 > p_1$ ($p_0 \leq p_1$). This is related to the ratio of unconstrained buyers to sellers and relies on the indifference condition for unconstrained buyers between submarkets p_0 and p_1. If $p_1 < p_0$, the ratio is higher in submarket p_1 (i.e., $\theta_0 < (1-\lambda_1)\theta_1$), which shifts the Poisson distribution that governs the random number of unconstrained buyers meeting each particular seller in the sense of first-order stochastic dominance. The policy therefore increases the probability of multiple offers from unconstrained buyers and the overall share of listed homes selling for u. The intuition for this is that unconstrained buyers enter the pooling submarket until the lower sales price when not competing against other unconstrained bidders (that is, p_1 instead of p_0) is exactly offset by the higher incidence of price escalation, resulting in indifference between the two submarkets. If instead $p_0 < p_1$, the indifference condition for unconstrained buyers implies the opposite, namely $\theta_0 \geq (1-\lambda_1)\theta_1$. In that case, the policy raises asking prices but lowers the probability of multiple offers from unconstrained buyers.

In both cases, the effect of the policy on sales prices via sellers' revised listing strategies (Prediction 1) is partly neutralized by the endogenous change in bidding intensity. We should therefore expect a more dramatic impact of the million dollar policy on asking prices than sales prices.

Prediction 3. *The million dollar policy increases the probability of selling-above-asking and shortens expected time-on-the-market for homes listed below \$1M. This results in discrete jumps in the probability of selling-above-asking (downward) and in expected time-on-the-market (upward) at asking price \$1M.*

At asking price $p_1 = c$, the presence of constrained buyers does not alter the payoff to an unconstrained buyer. This is because, in a second price auction with reserve price exactly equal to constrained buyers' ability to pay, offers from constrained buyers affect neither the probability of winning the auction nor the final sales price when an unconstrained buyer bids u. For submarkets priced above c, these constrained buyers cannot afford to participate. Given that $\bar{V}^u$ is unchanged by the policy, it follows that the ratio of unconstrained buyers to sellers is also unaffected by the policy in any submarket asking c or more. The policy, however, induces the participation of constrained buyers in submarket c and a range of inactive submarkets below c. Submarkets that attract both constrained and unconstrained buyers post-policy feature higher market tightness because the presence of constrained buyers does not deter unconstrained buyers. On the contrary, unconstrained buyers are drawn to these submarkets because they have an advantage when competing bidders face tighter financial constraints. The resulting discontinuous drop in market tightness at asking price c can be understood as discontinuous reductions in both the probability of selling and the probability of receiving multiple offers and hence selling-above-asking. The inverse of the probability of selling in the static model proxies for expected time-on-the-market in a dynamic setting. Prediction 3 therefore summarizes the implications for time-on-the-market. Specifically, the million dollar policy causes homes listed just below \$1M to sell faster, as well as induces a discontinuous increase in average selling time at the threshold.

In Appendix B.5, we illustrate Predictions 1, 2 and 3 by simulating a parameterized version of the model that has been extended to incorporate a form of seller heterogeneity. Specifically, sellers with different reservation values implement different asking price strategies, which permits the characterization of equilibria featuring *bunching from both above and below* simultaneously. Figures B1 and B2 plot the asking and sales price distributions. These simulated distribution functions reveal an excess mass of listings at \$1M from both above and below the threshold (Prediction 1), and a much less pronounced excess mass of sales at \$1M (Prediction 2). Figures B3 and B4 present visualizations of Prediction 3 by plotting expected time-to-sell and the probability of selling-above-asking as functions of the asking price.

Our analysis so far has focused exclusively on housing market outcomes. On the normative side, the model implies that the policy reduces the social surplus derived from housing market activity, as it affects the entry decision of sellers and hence market tightness, distorting the total number of housing market transactions.[23] It is worth noting that the million dollar policy was introduced not only to cool housing markets but also to improve financial stability and mortgage market efficiency. The latter is a central theme in the recent macro-finance literature on macroprudential policies.[24] Although the model is not designed to assess the policy's impact on borrowers' creditworthiness, it nevertheless offers an important insight in this regard. In particular, less constrained buyers have an advantage over constrained ones in multiple offer situations, and as such we would expect post-policy homebuyers to be wealthier and hence more "creditworthy".[25]

Prediction 4. *An unconstrained buyer is more likely to purchase a house than a constrained buyer following the introduction of the million dollar policy.*

By reallocating million dollar homes from financially constrained buyers to less financially constrained buyers, the policy effectively improves borrower creditworthiness and prevents lenders form making more risky loans. A normative argument in favor or against the million dollar policy would weigh these credit market benefits against the distortions introduced in the housing market.

[23]The welfare maximizing level of housing market activity is achieved in the pre-policy DSE, provided the pre-existing financial constraint is slack in problem P_0.

[24]See Jeske, Krueger, and Mitman (2013), Clerc et al. (2015), Elenev, Landvoigt, and Van Nieuwerburgh (2016), Elenev, Landvoigt, and Van Nieuwerburgh (2018) and Begenau (2019).

[25]Only the unconstrained search for and buy homes in segment p_0 of the post-policy DSE. In submarket p_1 of the post-policy DSE, the buying probabilities for constrained and unconstrained buyers are

$$\pi(0)+\sum_{k=1}^{\infty}\pi(k)\phi_k(0)\frac{1}{k+1}=\frac{e^{-(1-\lambda_1)\theta_1}-e^{-\theta_1}}{\lambda_1\theta_1} \quad \text{and} \quad \pi(0)+\sum_{k=1}^{\infty}\pi(k)\sum_{j=0}^{k}\phi_k(j)\frac{1}{j+1}=\frac{1-e^{-(1-\lambda_1)\theta_1}}{(1-\lambda_1)\theta_1}.$$

The probability of success in purchasing a house for unconstrained buyers therefore exceeds that for constrained buyers by

$$\frac{1-e^{-(1-\lambda_1)\theta_1}}{(1-\lambda_1)\theta_1}-\frac{e^{-(1-\lambda_1)\theta_1}-e^{-\theta_1}}{\lambda_1\theta_1}=\frac{1}{\lambda_1}\left[\frac{1-e^{-(1-\lambda_1)\theta_1}}{(1-\lambda_1)\theta_1}-\frac{1-e^{-\theta_1}}{\theta_1}\right]>0.$$

3.4 Caveats about Modeling Assumptions

Further discussion of some features of the model is in order. First, the asking price is assumed to represent a firm commitment to a minimum price, which results in a sales price either above or at the asking price. In practice, sales prices can be above, at, and below asking prices. Embellishing the price determination mechanism may allow for transaction prices below asking prices without compromising the asking price-related implications of the theory.[26] The theory of asking prices advanced in Khezr and Menezes (2018), for example, considers the situation wherein sellers learn their reservation value *after* setting an asking price and observing buyers' interest. As in our setting, transactions at the asking price arise in bilateral meetings; but unlike our model, multilateral meetings can, in some circumstances, result in transactions below the asking price. Alternative price determination mechanisms would add considerably to the analytical complexity of the model. Such extensions, however, would not affect our theoretical results substantively as long as (i) the asking price remains meaningful (in expectation) for price determination in a bilateral match, and (ii) competition among bidders in a multilateral match tends to drive up the sales price. The former is to ensure the directing role of the asking price, which is key for establishing Prediction 1. The latter is to allow for price escalation in multilateral matches so that sellers can list at $1M in response to the policy but may still end up selling for more. This is important for Prediction 2.

Second, entry on the supply side of the market is a common approach to endogenizing housing market tightness in directed search models with auctions (e.g., Albrecht, Gautier, and Vroman, 2016 and Arefeva, 2016). This assumption equates the seller's expected surplus with the listing cost. Keeping instead the measure of sellers constant pre- and post-policy would further reduce the seller's expected payoff and hence sales prices. A third alternative is to allow entry on the demand side, as in Stacey (2016). Buyer entry would be less straightforward in our context given that the demand side of the market is homogeneous pre-policy but heterogeneous thereafter. With post-policy entry decisions on the demand

[26]See Albrecht, Gautier, and Vroman (2016) and Han and Strange (2016) for more sophisticated pricing protocols that can account for sales prices above, at, and below the asking price.

side, buyers would self-select into the market in such a way that the effects of the policy would be mitigated or even non-existent. Suppose for a moment that both types of buyers face entry decisions subject to an entry fee or search cost. Provided there are sufficiently many unconstrained potential market participants, unconstrained buyers would enter the market until they reach indifference about market participation: their expected payoff would equal the participation cost. Because constrained buyers are outbid by unconstrained buyers, the expected payoff for a constrained buyer would be strictly less than the cost of market participation. It follows that constrained buyers would optimally choose not to participate in this segment of the housing market and consequently the post-policy equilibrium would be indistinguishable from the pre-policy equilibrium with identically unconstrained buyers. In contrast, we have shown in the preceding analysis that the policy does affect equilibrium strategies and outcomes when entry decisions are imposed on the supply side of the market.

Finally, the scope of the model shrinks to a narrow segment of the market around \$1M if the parameter values for v, u and c are close to the seller's reservation value, which is normalized to zero. The 20% downpayment constraint also reduces the maximum affordable price for buyers in segments well above \$1M. The model's implications for these segments are the same as the \$1M segment, albeit with a reinterpretation of parameter c. More specifically, the perceived reduced ability to pay among prospective buyers generates incentives for the sellers to adjust their asking prices, and heightened competition among less constrained buyers can further drive sales prices above asking. In these segments well above \$1M, however, the downpayment requirement is continuous in the sales price, which makes it empirically challenging to identify the policy effects on buyers' and sellers' strategies and price formation. In contrast, the policy creates a discrete change in the downpayment requirement at the \$1M threshold. The degree of excess bunching at \$1M in the data, which we turn to next, provides evidence on the extent to which buyers and sellers respond to this targeted financial constraint.

4 Data and Methodology

4.1 Data

Our dataset includes transactions of residential homes in the Greater Toronto Area from January 1st, 2010 to December 31st, 2013. For each transaction, we observe asking price, sales price, days on the market, transaction date, location, as well as detailed housing characteristics. In particular, we define a number of variables to control for house quality. We create indicator variables for whether the house is detached, semi-detached, condominium or townhouse. Houses in our data are coded in 16 different styles. We condense this information into three housing styles (2-story ($\approx 65\%$), bungalow ($\approx 25\%$), other ($\approx 10\%$)), where the style 'other' includes 1-1/2 story, split-level, backsplit, and multi-level. We observe the depth and width of the lot in meters, which we convert to the total size of the lot by taking their product. We create a categorical variable for the number of rooms in a house that has 7 categories, from a minimum of 5 to ≥ 11, and another for the number of bedrooms that has 5 categories from 1 to ≥ 5. We create an indicator for the geographic district of the house listing. For our main sample of the city of Toronto, this district variable identifies 43 districts corresponding to the MLS district code.

We observe the final asking price posted in each listing, but not the changes in the asking price. From a local brokerage office's confidential database, we learned that about 12% of overall listings experienced revisions to the asking price. This number reduces to 2% when it comes to the estimation sample of properties around $1M. For our analysis, we split our data into two mutually exclusive time periods. We define a *post-policy period* from July 15th, 2012, to June 15th, 2013. Our *pre-policy period* is similarly defined as July 15th, 2011, to June 15th, 2012. That is, we choose one year around the policy implementation date, but we omit a month covering the pre-implementation announcement of the policy. For the purposes of assigning a home to the pre- or post-policy period, we use the date the house was listed.[27] We do so because a seller's listing decision depends on the perceived ability to

[27] There are no notable differences in our results (available on request) when we instead assign homes to the pre- or post-policy period based on the date the house sold.

pay among potential buyers, which in turn depends on whether the policy is implemented. We assess the sensitivity of our results to different time windows of the pre- and post-policy periods in a later robustness section.

For the main analysis, we focus on single-family homes in the city of Toronto.[28] Table 1 contains summary statistics. Panel (a), containing information on all districts, includes 22,244 observations in the pre-policy period and 19,061 observations in the post-policy period. The mean sales price in Toronto was $723,396.82 in the pre-policy period and $760,598.15 in the post-policy period, reflecting continued rapid price growth for single family houses (all figures in CAD). Our focus is on homes near the $1M threshold, which corresponds to approximately the 86th percentile of the pre-policy price distribution. There were 1,448 homes sold within $100,000 of $1M in the pre-policy period and 1,423 in the post-policy period. Panel (b) of Table 1 shows summary statistics for the central district only. The central district of Toronto is more expensive than suburban markets in general; in the post-policy period, a $1M home is at the 56th percentile of the sales price distribution in the central district. The central district contains nearly 40 percent of the homes sold within the $0.9M-1.1M price range. In the empirical analysis below, we will examine the policy impact for the city of Toronto, and for central Toronto separately.

4.2 Empirical Methodology

To measure price responses, we use a bunching approach recently developed in the public finance literature (e.g., Saez 2010, Chetty et al. 2011, and Kleven and Waseem 2013). Our theoretical model established that the downpayment discontinuity can create incentives for bunching at the $1M threshold in terms of listings (Prediction 1), but less so in terms of sales (Prediction 2). To test these predictions, we use the price segments which are not subject to the policy's threshold effects to form a valid counterfactual near the $1M threshold. The two

[28] The geographic area of our study includes the city of Toronto and the immediate bordering municipalities of Vaughan, Richmond Hill, and Markham. We do not include the municipalities to the west (Mississauga and Brampton) or east (Pickering) because there are very few million dollar homes. Our main results (available on request) are very similar when we include them.

Table 1
Summary Statistics: City of Toronto

(a) All Districts

		Pre-Policy		Post-Policy	
		Asking	Sales	Asking	Sales
All Houses	Mean	722430.15	723396.82	770836.16	760598.15
	25th Pct	459900.00	465000.00	499000.00	491000.00
	50th Pct	599000.00	605000.00	639000.00	635000.00
	75th Pct	799000.00	807500.00	849000.00	845000.00
	N	22244.00	22244.00	19061.00	19061.00
	Median Duration	10.00	10.00	13.00	13.00
	$1M Percentile	0.87	0.86	0.85	0.84
Houses $0.9–1.0M	N	840.00	934.00	888.00	907.00
	Median Duration	9.00	8.00	13.00	12.00
	Mean Price	964427.90	942427.89	966120.77	946257.88
Houses $1.0–1.1M	N	364.00	514.00	410.00	516.00
	Median Duration	10.00	9.00	13.00	12.00
	Mean Price	1071802.41	1043508.98	1073840.91	1044025.97

(b) Central District

		Pre-Policy		Post-Policy	
		Asking	Sales	Asking	Sales
All Houses	Mean	1082210.56	1087206.62	1172612.53	1153957.65
	25th Pct	649000.00	665000.00	699900.00	718000.00
	50th Pct	849000.00	875000.00	899900.00	925000.00
	75th Pct	1288000.00	1295000.00	1395000.00	1362500.00
	N	4943.00	4943.00	4065.00	4065.00
	Median Duration	9.00	9.00	11.00	11.00
	$1M Percentile	0.64	0.60	0.58	0.56
Houses $0.9–1.0M	N	334.00	363.00	336.00	335.00
	Median Duration	8.00	8.00	8.00	8.00
	Mean Price	966559.71	943206.85	968328.73	945389.72
Houses $1.0–1.1M	N	163.00	228.00	186.00	226.00
	Median Duration	8.00	8.00	10.00	10.00
	Mean Price	1073393.17	1044304.37	1074523.94	1045802.38

Notes: This table displays summary statistics for the city of Toronto for single family homes (attached and detached). The pre-policy period is defined as July 15th, 2011, to June 15th, 2012, and the post-policy period is defined as July 15th, 2012, to June 15th, 2013. The columns labelled Asking refer to asking prices and the columns labelled Sales refer to sales prices. Duration refers to the number of days a home is on the market.

underlying assumptions are that (1) the policy-induced incentives for bunching occur locally in segments near the \$1M threshold, leaving other parts of the price distributions unaffected by threshold consequences; and (2) the counterfactual is smooth and can be estimated using these other parts of the price distributions. In forming the counterfactual, we use a two-step approach: first constructing counterfactual price distributions that would have prevailed if there were no changes in the composition of the housing stock using a common reweighting method; then applying the bunching approach to the difference between each composition-constant post-policy price distribution and the observed pre-policy distribution.

4.2.1 First step: controlling for housing composition

If houses listed or sold in the million dollar segment in the post-policy year differ in terms of quality from those in the previous year, then the difference between price distributions in the two periods could simply reflect the changes in the composition of housing rather than the effect of the policy. We alleviate this concern by leveraging the richness of our data to flexibly control for a set of observed house characteristics to back out a counterfactual distribution of house prices that would have prevailed if the characteristics of houses in the post-policy period were the same as in the pre-policy period.

Let Y_t denote the (asking or sales) price of a house and let X_t denote the characteristics of a house that affect prices at $t = 0$ (the pre-policy period), and $t = 1$ (the post-policy period). The conditional distribution functions $F_{Y_0|X_0}(y|x)$ and $F_{Y_1|X_1}(y|x)$ describe the stochastic assignment of prices to houses with characteristics x in each of the periods. Let $F_{Y\langle 0|0\rangle}$ and $F_{Y\langle 1|1\rangle}$ represent the observed distribution of house prices in each period. We are interested in $F_{Y\langle 1|0\rangle}$, the counterfactual distribution of house prices that would have prevailed in the post-period if the characteristics of the houses in the post-period were as in the pre-period. We can decompose the observed change in the distribution of house prices:

$$\underbrace{F_{Y\langle 1|1\rangle} - F_{Y\langle 0|0\rangle}}_{\Delta_O=\text{Observed}} = \underbrace{\left[F_{Y\langle 1|1\rangle} - F_{Y\langle 1|0\rangle}\right]}_{\Delta_X=\text{Composition}} + \underbrace{\left[F_{Y\langle 1|0\rangle} - F_{Y\langle 0|0\rangle}\right]}_{\Delta_S=\text{Price Structure}}. \tag{6}$$

Since the counterfactual is not observed, it must be estimated. We use a simple reweighting method proposed by DiNardo, Fortin, and Lemieux (1996) based on the following relation:

$$F_{Y\langle 1|0\rangle} = \int F_{Y_1|X_1}(y|x) \cdot \Psi(x) \cdot dF_{X_1}(x)$$

where $\Psi(x) = dF_{X_0}/dF_{X_1}$ is a reweighting factor that can be easily estimated using a logit model (for details, see Fortin, Lemieux, and Firpo 2011). To implement this, we obtain the weighting function by pooling pre- and post-policy data and estimating a logit model where the dependent variable is a pre-policy period dummy. The covariate vector contains indicators for district, month, the number of rooms, the number of bedrooms, whether the house is detached or semi-detached, the lot size and its square, and the housing style (2-story, bungalow, other).[29] The estimated counterfactual distribution is given by $\hat{F}_{Y\langle 1|0\rangle} = \int \hat{F}_{Y_1|X_1}(y|x) \cdot \hat{\Psi}(x) \cdot d\hat{F}_{X_1}(x)$, where $\hat{F}$ denotes a distribution function estimated using grid intervals of $5,000. The result is a reweighted version of the observed price distribution in the post-policy period that can be interpreted as the price distribution that would prevail if the characteristics of homes were the same as in the pre-policy period.

4.2.2 Second step: bunching estimation

With the estimated $\hat{\Delta}_S(y_j) = \hat{F}_{Y\langle 1|0\rangle}(y_j) - \hat{F}_{Y\langle 0|0\rangle}(y_j)$ in hand, we are now ready to estimate the policy effects on asking and sales price using a bunching estimation procedure. This procedure requires separation of the observed $\hat{\Delta}_S(y_j)$ into two parts: the price segments near $1M that are subject to the policy's threshold effects, and the segments that are not. The affected segments are known as the "excluded region" in the bunching literature. Since knowledge of this region is not known *a priori*, it must also be estimated and we develop a procedure below to do so. Once this region around $1M is determined, we use standard methods to estimate the counterfactual difference in distributions by fitting a flexible polyno-

[29] The weighting function is $\Psi(x) = \frac{p(x)}{1-p(x)} \cdot \frac{1-P(t=1)}{P(t=0)}$, where $p(x)$ is the propensity score: i.e., the probability that $t = 0$ given x.

mial to the estimated $\hat{\Delta}_S(y_j)$ outside the excluded region. We use the estimated polynomial to predict or "fill in" the excluded region which forms our counterfactual. Our estimates of the policy effects are derived from the difference between the observed $\hat{\Delta}_S(y_j)$ and the estimated counterfactual.

In particular, consider the equation:

$$\begin{aligned}\hat{\Delta}_S(y_j) = \sum_{i=0}^{p} \beta_i \cdot y_j^i + \beta_A \cdot 1[y_j = \$1\text{M}] + \beta_B \cdot 1[y_j = \$1\text{M} - h] \\ + \sum_{l=1}^{L} \gamma_l \cdot 1[y_j = \$1\text{M} - h \cdot (1+l)] + \sum_{r=1}^{R} \alpha_r \cdot 1[y_j = \$1\text{M} + h \cdot r] + \epsilon_j \qquad (7)\end{aligned}$$

where p is the order of the polynomial, L is the excluded region to the left of the cut-off, R is the excluded region to the right of the cut-off, and h is the bin size.[30]

The total observed jump at the \$1M cut-off is

$$\begin{aligned}\underbrace{\hat{\Delta}_S(\$1\text{M}) - \hat{\Delta}_S(\$1\text{M} - h)}_{\text{Jump at threshold}} = \underbrace{\sum_{i=0}^{p} \hat{\beta}_i \cdot y_{\$1\text{M}}^i - \sum_{i=0}^{p} \hat{\beta}_i \cdot y_{\$1\text{M}-h}^i}_{\text{Counterfactual (C)}} \\ + \underbrace{\underbrace{\hat{\beta}_A}_{\text{Bunching from above (A)}} - \underbrace{\hat{\beta}_B}_{\text{Bunching from below (B)}}}_{\text{Total Policy Response}} \qquad (8)\end{aligned}$$

It is important to note that the interpretation of the total jump at the threshold, as shown in the left-hand-side of equation (8), is not all causal. Since changes in listings and sales between the two periods can be more pronounced in some price segments than others, we should not expect the difference in CDFs to be flat even in the absence of the million dollar policy. In our case, an upward sloping curve is captured by our polynomial estimates as a counterfactual. Specifically, the first two terms on the right-hand-side of equation (8) reflect the counterfactual difference at the \$1M threshold.

[30] Note that there is no residual component in equation (7) since, throughout the excluded region, every bin has its own dummy and the fit is exact. We observe the population of house sales during this time, thus, the error term in (7) reflects specification error in our polynomial fit rather than sampling variation. We discuss the computation of our standard errors of our estimates in more detail below.

After netting out the counterfactual, we are left with $\hat{\beta}_A - \hat{\beta}_B$, which is the policy response we aim to measure. A finding of $\hat{\beta}_A > 0$ is consistent with *bunching from above* since it indicates that sellers that would have otherwise located in bins above \$1M instead locate in the \$1M bin. A finding of $\hat{\beta}_B < 0$, on the other hand, is consistent with *bunching from below* since it indicates that sellers that would otherwise locate below the \$1M bin now move up to locate in the \$1M bin. Both are responses to the million dollar policy.

In the absence of an extensive margin response, the two sources of response described above imply the following two constraints. First, the excess mass in the distribution at \$1M resulting from *bunching from below* should equal the the responses from lower adjacent bins, implying

$$R^B \equiv \beta_B - \sum_{l=1}^{L} \gamma_l \cdot 1[y_j = \$1M - h \cdot (1+l)] = 0. \tag{9}$$

Similarly, for those sellers coming from above the threshold,

$$R^A \equiv \beta_A - \sum_{r=1}^{R} \alpha_r \cdot 1[y_j = \$1M + h \cdot r] = 0. \tag{10}$$

In order to implement our estimator, several decisions must be made about unknown parameters, as is the case for all bunching approaches. In particular, the number of excluded bins to the left, L, and right, R, are unknown, as is the order of the polynomial, p. In addition, we choose to limit our estimation to a range of price bins around the \$1M threshold. We do this because the success of our estimation procedure requires estimation of the counterfactual in the region local to the policy threshold (Kleven 2016). Using data points that are far away from the excluded region to predict values within the excluded region can be sensitive to polynomial choice and implicitly place very high weights on observations far from the threshold (Gelman and Imbens 2014; Lee and Lemieux 2010). Thus, we focus on a more narrow range, or estimation window, W, of house prices around the policy threshold. Since we are fitting polynomial functions, this can be thought of as a bandwidth choice for local polynomial regression with rectangular weights (Imbens and Lemieux 2008). Thus, the parameters we require for estimation of the regression coefficients are (L, R, W, p).

We use a data-driven approach to select these parameters. The procedure we implement is a 5-fold cross-validation procedure, described fully in Appendix C.1. Briefly, we split our individual-level data into 5 equally sized groups and carry out both step 1 and 2 of our estimation procedure using 4 of the groups (i.e., holding out the last group), and then obtain predicted squared residuals from equation (7) for the hold-out group. We repeat this procedure 5 times, holding out a different group each time, and average the predicted squared residuals across each repetition. This is the cross-validated Mean Squared Error (MSE) for a particular choice of (L, R, W, p). We perform a grid search over several values of each parameter, and choose the specification which minimizes the MSE.[31]

4.2.3 Caveats about Empirical Methodology

One legitimate concern is that our bunching estimates pick up threshold effects in pricing that are caused by, for example, marketing convention or psychological bias surrounding \$1M, or other macro forces that affected the housing market at the same time as the implementation of the million dollar policy. Our estimation methodology addresses this concern in two ways. First, we examine the post-policy CDF relative to the CDF in the pre-policy period. Time-invariant threshold price effects unrelated to the policy are therefore differenced out in our estimation. Second, we allow for round number fixed effects to capture potential rounding in the price data. Thus, all estimates reported below include a dummy variable for prices in \$25,000 increments, and another for prices in \$50,000 increments.

Another potential concern is that the million dollar policy is announced in combination

[31] In the literature on bunching estimation, the excluded region is sometimes selected by visual inspection (Saez 2010; Chetty et al. 2011) in combination with an iterative procedure (Kleven and Waseem 2013; DeFusco and Paciorek 2017) that selects the smallest width consistent with adding-up constraints. Often, high-order global polynomials are used in estimation and robustness to alternative polynomial orders are shown. In the closely related regression discontinuity literature, free parameters are sometimes chosen by cross-validation (Lee and Lemieux 2010). In a recent paper by Diamond and Persson (2016), there are many different regions and time periods where bunching occurs, and so visual inspection is impractical. They develop a k-fold cross-validation procedure to choose the width of the manipulation region and polynomial order. Our approach closely follows theirs. In addition, we consider a series of robustness checks to assess the sensitivity of our estimates to the choice of parameters L, R, W, and p. We find that our estimates are quite robust to reasonable deviations from the parameter values selected by our cross-validation procedure.

with three other mortgage rule changes, which may complicate the challenge for identification. However, unlike the million dollar policy, these contemporaneous mortgage rule changes apply to the entire housing market.[32] Their housing market impacts are accounted for in the counterfactual price distribution that would have prevailed in the absence of the million dollar policy. By comparing the actual post-policy distribution of house prices to the counterfactual distribution, the bunching estimation teases out the effect of the million dollar policy from these confounding factors. To address the possibility that buyer and seller strategies evolve dynamically in ways that are not reflected in the counterfactuals, we implement several placebo experiments that use alternative cut-offs in Section 5.1.2. We also present results based on counterfactuals constructed only from data below $1M.

5 Empirical Evidence

The core estimation is presented in Section 5.1 with an aim to test Predictions 1 and 2 by examining asking and sales prices near the $1M threshold. We then explore Prediction 3 in Section 5.2.1 and Prediction 4 in Section 5.2.2. Finally, we extend the analysis to study all segments above $1M in Section 5.3.

5.1 Predictions 1 and 2: Asking Price and Sales Price

5.1.1 Main Results

The main predictions of the model are that the million dollar policy leads to an excess mass of homes listed at the $1M threshold (Prediction 1), which only partially passes through to the sales price distribution (Prediction 2).

Our main analysis focuses on single-family-housing markets. Figures 3 and 4 present graphical results from the first step estimation of asking and sales price distributions based on equation (6). We first discuss asking prices. Panel A of Figure 3 plots the distribution functions for the asking price between $600,000 and $1,400,000 in the pre- and post-policy

[32] See footnote 6 for details.

years. The post-policy CDF lies everywhere below the pre-policy CDF, indicating that all housing market segments experienced a boom. Panel B plots the difference between the two CDFs. If the CDFs were the same pre- and post-policy for a given bin, the difference would show up as a zero in the figure. The displayed difference in CDFs is always below zero, indicating that houses in general are becoming more expensive over time. Following equation (6), we then decompose the difference in CDFs into two components: (i) price difference due to shifting housing characteristics in each segment (Panel C); and (ii) price difference due to changes in sellers' listing strategies (Panel D). The latter is the market response that we aim to measure at the \$1M threshold. As shown in Panel C, the price change caused by shifting housing characteristics is small in magnitude and relatively flat. In contrast, Panel D shows that the price difference caused by sellers' updated listing strategies generally changes smoothly with price, with a relatively large jump at \$1M. Given the minimal composition effect, nearly all of the difference in the observed distribution of asking prices is driven by sellers' listing behavior.

Turning to the sales prices, the top panels of Figure 4 plot the distribution of sales price in the pre- and post-policy years and their differences. The bottom panels of Figure 4 show that after accounting for the composition effect, a jump in the sales price at the \$1M threshold is hardly apparent, supporting the notion that buyers' non-trivial search and bidding activities disentangle sales prices from asking prices, potentially mitigating the overall impact of the policy.

The descriptive findings presented in Figures 3 and 4 are consistent with the model. However, this evidence alone does not distinguish the policy effects from the impact of other contemporaneous macro forces. To isolate the million dollar policy's effects on the price distributions, we now turn to the second step estimation: namely, bunching estimation. We choose to plot the bunching estimates in both the CDFs and the PDFs. While the latter is more standard in the bunching literature, the former allows us to visualize the decomposition of the estimated jump based on equation (8) in a more transparent way.

Figure 5a presents a graphical test of Prediction 1 based on the estimation of equation

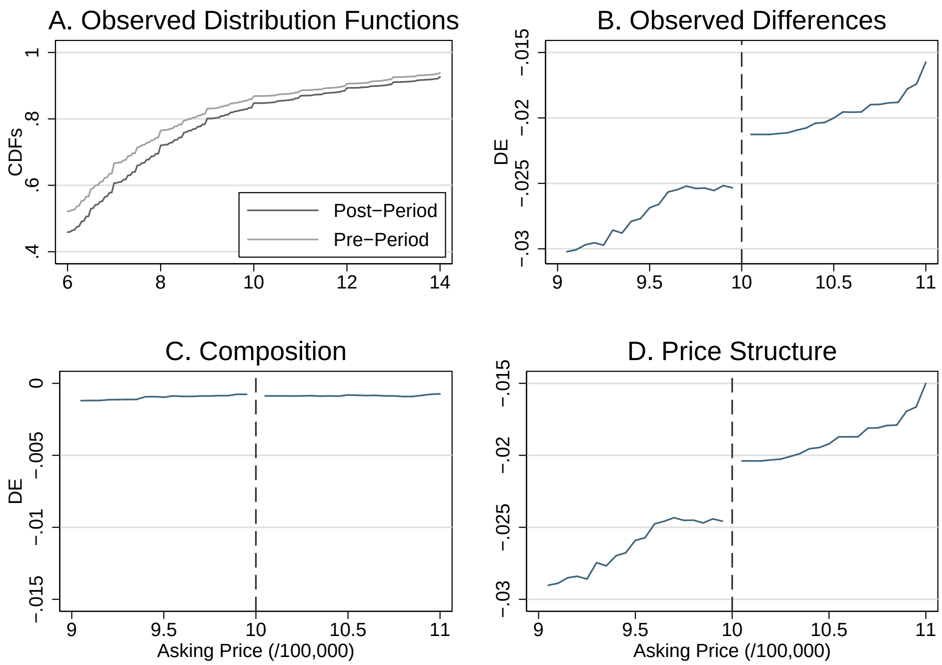

Figure 3
Observed Distribution and Decomposition of Asking Prices

Notes: The figure uses data on asking prices for the city of Toronto in the year before (pre-period) and after (post-period) the implementation of the million dollar policy. Panel A plots the empirical CDF of asking prices for each year. Panels B through D decomposes the difference in the CDFs according to equation (6). Panel B plots the observed difference in the CDFs, Δ_O. Panel C plots the difference in the CDFs due to composition, Δ_X. Panel D plots the difference due to the change in the price structure, Δ_S.

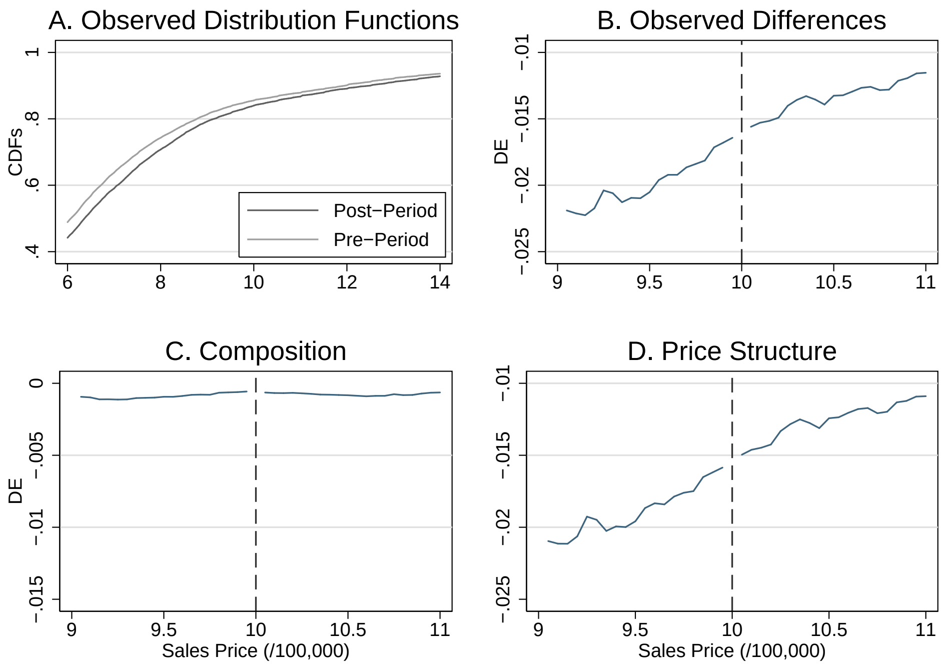

Figure 4
Observed Distribution and Decomposition of Sales Prices

Notes: The figure uses data on sales prices for the city of Toronto in the year before (pre-period) and after (post-period) the implementation of the million dollar policy. Panel A plots the empirical CDF of asking prices for each year. Panels B through D decomposes the difference in the CDFs according to equation (6). Panel B plots the observed difference in the CDFs, Δ_O. Panel C plots the difference in the CDFs due to composition, Δ_X. Panel D plots the difference due to the change in the price structure, Δ_S.

(7). In particular, we plot changes in the CDFs of the asking price, $\hat{\Delta}_S(y_j) = \hat{F}_{Y\langle 1|0\rangle}(y_j) - \hat{F}_{Y\langle 0|0\rangle}(y_j)$, holding housing characteristics constant. The solid line plots the quality-adjusted observed changes, with each dot representing the difference in the CDFs before and after the policy for each \$5,000 price bin indicated on the horizontal axis. The dashed line plots the counterfactual changes in the absence of the policy, while the vertical dashed lines mark the lower and upper limits of the bunching region (\$975,000 and \$1,025,000). Note that the width of the estimation widow (\$100,000 dollars on each side of the threshold), the order of the polynomial (cubic), and the width of the excluded region were chosen based on the cross-validation procedure outlined in Section 4.2.2.

The empirical distribution of asking prices exhibits a sharp discontinuity at the \$1M threshold. After the policy, a total of 0.45 percent of listings were added to the \$1M bin. Following equation (8), Figure 5a decomposes this total jump into three distinct components: 0.06 percent of listings reflect the counterfactual change in the absence of the million dollar policy (marked by C), 0.18 percent of listings bunched from below (marked by B) and 0.20 percent of listings bunched from above (marked by A). Thus, 87% (i.e., (A+B)/(A+B+C)) of the excess bunching in the asking price at the \$1M threshold is attributed to the policy.

Figure 5b further presents a graphical test of Prediction 1 based on the difference in densities. The spike in homes listed at the \$1M is accompanied by dips in homes listed to the right and left of \$1M. The spike reflects the excess mass of homes listed between \$995,000 and \$1M after the implementation of the policy. The dips reflect missing homes that would have been listed at prices further from the \$1M in the absence of the policy.

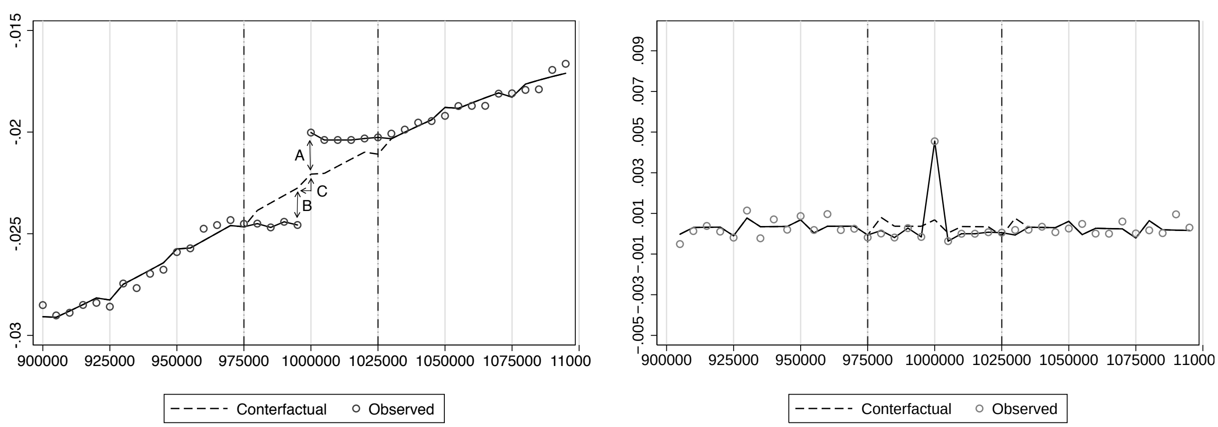

(a) $\hat{\Delta}_S(y_j)$ and counterfactual estimate

(b) $\hat{\Delta}_{S'}(y_j)$ and counterfactual estimate

Figure 5
Visual representation of column (1) in Table 2

Notes: Panel (a) of the figure shows a visual representation of the bunching specification in column (1) of Table 2 which uses data on asking prices for the city of Toronto. The dots indicate the before-after policy differences in the CDFs, $\hat{\Delta}_S(y_j)$. Vertical dashed lines in the figure indicate the excluded region. The solid line is the fitted polynomial from equation (7) outside the excluded region and the fitted dummies within it. The dashed line, formed from predicted values of the polynomial within the excluded region, indicates the counterfactual estimate of the CDF difference that would have prevailed in the absence of the policy. The figure labels correspond to those in equation (8) that decompose the vertical jump at the policy threshold, indicating the magnitude of bunching from above (A) and below (B), and the counterfactual estimate (C). Panel (b) represents the same specification in terms of differences in PDFs.

Column (1) of Table 2 reports our baseline bunching estimates underlying the above graphical presentation. The specification used is chosen by the cross-validation procedure outlined above. Standard errors are calculated via bootstrap.[33] Overall, we find that approximately 86 homes that would have otherwise been listed away from \$1M were shifted to the \$1M bin. While seemingly small, 86 additional homes represent a 38.3 percent increase relative to the number of homes that would have been listed in the million dollar bin in the absence of the policy. Among these additional listings, about half are shifted from below \$995,000; the remaining half come from above \$1M. Both estimates are significant at the five percent level. When viewed through the lens of our search model, price adjustments from both sides of \$1M are quite sensible. On the one hand, the policy induces some sellers who would have otherwise listed homes below \$1M to increase their asking prices towards the \$1M mark. By doing so, these sellers demand a higher price in a bilateral situation to offset any price dampening effect of the policy in multilateral situations. On the other hand, the policy induces other sellers who would have otherwise listed homes above \$1M to lower their asking price to just below the cut-off, attracting both constrained and unconstrained buyers to compete for their homes.

As noted earlier, we do not observe sellers' revisions to asking prices in our main data. With a one-time access to a local brokerage office's confidential database, we find that about 44 houses in the estimation sample (2% of all houses that sold within \$100,000 of \$1M) were originally listed before the policy, pulled off the market, re-listed after the policy and then sold. Restricting attention to these 44 houses, Figure 6 shows that 18 (41%) of them adjusted their asking price to [\$975,000, \$1,000,000]. These adjustments come from both sides, complementing our bunching estimation results. Among the 22 re-listed houses that were originally asking between \$1M and \$1.1M before the policy, 14 (64%) of them reduced their price to just under \$1M after the policy. These asking price revisions are consistent the intuition that some sellers lower their asking price to invite competition from both constrained and unconstrained buyers.

[33]We calculate standard errors for all estimated parameters by bootstrapping both steps 1 and 2 of the estimation procedure. We draw 399 random samples with replacement from the household-level data, and calculate the standard deviation of our estimates for each of these samples.

Table 2
Regression Bunching Estimates: City of Toronto and Central District

	City of Toronto		Central District	
	(1) Asking	(2) Sales	(3) Asking	(4) Sales
Jump at cut-off	0.0045*	0.00094*	0.0094*	0.0032*
	(0.0010)	(0.00042)	(0.0032)	(0.0014)
Total Response	0.0039*	0.00050	0.0068*	0.0028
	(0.0010)	(0.00053)	(0.0031)	(0.0017)
From Below	-0.0018*	0.00017	-0.0049*	-0.0013
	(0.00072)	(0.00060)	(0.0022)	(0.0018)
From Above	0.0020*	0.00067	0.0020	0.0014
	(0.00089)	(0.00079)	(0.0025)	(0.0022)
Observations	41305	41305	9008	9008
Excluded Bins:				
L	4	1	3	1
R	5	2	4	2
Tests of Fit:				
$B - \sum_l^L \beta_B^l$	-.0013	.00017	-.00045	-.0013
	(.0013)	(.0006)	(.0022)	(.0018)
$A - \sum_r^R \beta_A^r$	.0025	-.00036	.0048	-.0012
	(.0016)	(.00057)	(.0041)	(.0018)
Joint p-val.	0.20	0.82	0.52	0.45
Impact:				
Δ Houses at cutoff	85.9	11.1	33.8	13.7
Specifications:				
Poly. Order	3	3	2	2
Window	20	20	25	20
Other	CV Opt.	CV Opt.	CV Opt.	CV Opt.

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto and the central district. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using asking prices (columns 1 and 3) or sales prices (columns 2 and 4). The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

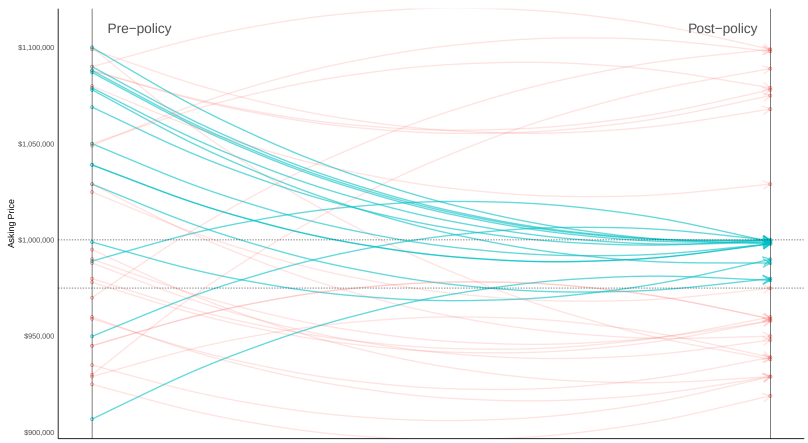

Figure 6
Observed Changes in Asking Price Among Re-Listed Properties

Notes: The figure uses data for a subset of houses that were listed prior to the implementation of the million dollar policy that were withdrawn and listed again after the implementation of the policy. This subset includes only houses that had asking prices in the $900K to $1M range within 180 days of June 9th, 2012. The left vertical bar shows the house's pre-policy asking price and the right vertical bar shows the house's post-policy asking price. Lines in green show pricing behaviour consistent with our bunching analysis; that is, these houses re-listed just below $1M (specifically, the $975K to $1M segment, indicated by dashed horizontal lines).

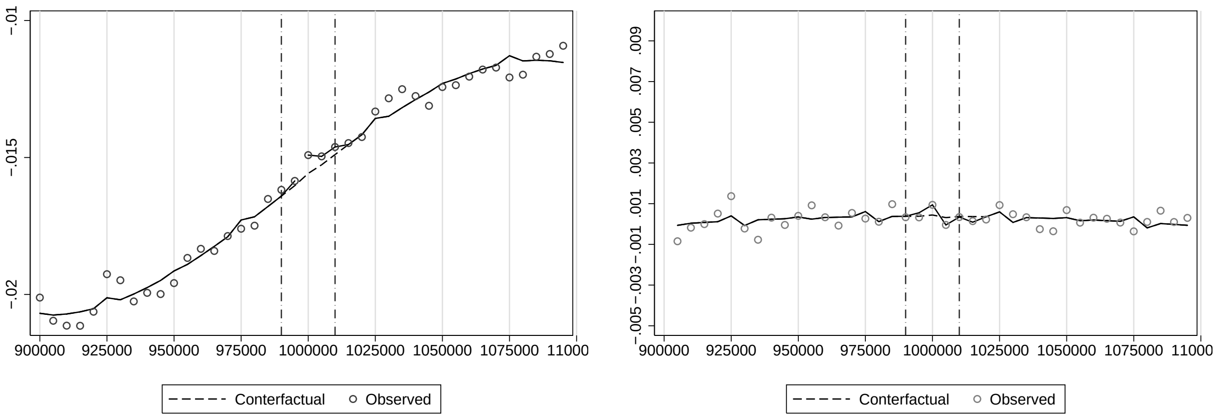

(a) $\hat{\Delta}_S(y_j)$ and counterfactual estimate

(b) $\hat{\Delta}_{S'}(y_j)$ and counterfactual estimate

Figure 7
Visual representation of column (2) of Table 2

Notes: Panel (a) of the figure shows a visual representation of the bunching specification in column (2) of Table 2 which uses data on sales prices for the city of Toronto. The dots indicate the before-after policy differences in the CDFs, $\hat{\Delta}_S(y_j)$. Vertical dashed lines in the figure indicate the excluded region. The solid line is the fitted polynomial from equation (7) outside the excluded region and the fitted dummies within it. The dashed line, formed from predicted values of the polynomial within the excluded region, indicates the counterfactual estimate of the CDF difference that would have prevailed in the absence of the policy. Panel (b) represents the same specification in terms of differences in PDFs.

Turning to Prediction 2, we report the bunching estimates for the sales price in column (2) of Table 2, with a visualization of the estimates shown in Figure 7. Despite sharp excess bunching of asking prices, we do not find evidence of excess bunching of sales prices at the $1M price bin; the estimated total response attributable to the million dollar policy is small and statistically insignificant. This evidence is consistent with Prediction 2, which characterized how the price dampening effect of the policy can be undermined by the strategic search and bidding behavior of market participants in a setting with search frictions and auctions. To test this interpretation, we estimate the policy effect on bidding intensity in Section 5.2.1.

Million dollar homes are concentrated in central Toronto. In columns (3) and (4) of Table 2, we restrict the sample to the central district and repeat the same estimation for asking and sales prices as in columns (1) and (2). Despite the much reduced sample size, the resulting estimates are qualitatively consistent with what we find above for the city of Toronto.

Condominiums and townhouses make up an important sector of the Toronto housing market with 21,768 transactions in the pre-policy period.[34] For this sector, Table 3 shows that the million dollar policy adds 19 listings at \$1M and 12 sales at \$1M, aligning again with Predictions 1 and 2. Quantitatively, the degree of excess bunching is smaller, because there are much fewer million dollar condominiums than houses. While a \$1M home corresponds to the 86th percentile in the single-family-housing market, it corresponds to the 99th percentile in the condominium and townhouse market. Given this, we focus the analysis hereinafter on single-family-homes.

5.1.2 Robustness Checks

In Appendix E.1, we estimate an extensive set of specifications to assess the robustness of our main results. We briefly review these results here and provide a more extensive discussion in the appendix. Our first robustness exercise deals with the concern that the bunching estimates could be altered by plausible policy responses above the \$1M threshold. Suppose the introduction of the policy hindered potential listings or transactions above \$1M. In that case, our counterfactuals estimated by excluding an area around the \$1M threshold would not accurately reflect what would have occurred in the absence of the policy. Note that this is a common issue in the bunching literature (Kopczuk and Munroe 2015; Best et al. 2018; Best and Kleven 2018). As suggested by Kleven (2016), we construct the counterfactual using only data below \$1M under the assumption that the distribution below the threshold is unaffected by the policy, and the results are similar to those presented above.

Our analysis above hinges on the assumption that homes further below \$1M are unaffected by the policy. A second concern, then, is that there are unintended policy consequences in market segments below the threshold. This could occur, for example, if a seller of a below-\$1M home intends to trade-up to an above-\$1M home. The seller may be *constrained* in that the proceeds from the sale of their current home must not compromise their ability to make a 20% downpayment on their next home. To that end, they may alter their listing and selling

[34] See Table D1 in the Appendix.

Table 3
Assessing Robustness to Housing Types

	Condos/Townhouses		All Homes	
	(1) Asking	(2) Sales	(3) Asking	(4) Sales
Jump at cut-off	0.00077	0.00052*	0.0030*	0.00074*
	(0.00040)	(0.00018)	(0.00059)	(0.00025)
Total Response	0.00086*	0.00056*	0.0025*	0.00049
	(0.00040)	(0.00020)	(0.00059)	(0.00030)
From Below	-0.00012	-0.000076	-0.00097*	0.000076
	(0.00027)	(0.00022)	(0.00040)	(0.00034)
From Above	0.00074*	0.00049	0.0016*	0.00057
	(0.00031)	(0.00025)	(0.00048)	(0.00042)
Observations	40025	40025	83058	83058
Excluded Bins:				
L	4	1	4	1
R	5	3	5	2
Tests of Fit:				
$B - \sum_l^L \beta_B^l$	.00016	-.000076	-.00048	.000076
	(.00054)	(.00022)	(.00073)	(.00034)
$A - \sum_r^R \beta_A^r$	.001	.00012	.0022*	-.00021
	(.00054)	(.00029)	(.00084)	(.00032)
Joint p-val.	0.18	0.86	0.036	0.80
Impact:				
Δ Houses at cutoff	18.7	12.2	114.4	22.1
Specifications:				
Poly. Order	3	3	3	3
Window	20	20	20	20

Notes: This table displays the bunching estimates of the million dollar policy for condos/townhouses and all housing types in Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using using either asking prices (columns 1 and 3) or sales prices (columns 2 and 4). The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

strategies, which could affect prices within the estimation window but below the excluded region. Given the rate of house price appreciation, however, this trading-up constraint is unlikely to bind except possibly for sellers that bought their home very recently.[35] In Table E4 of Appendix E.2, we report specifications that exclude sellers that bought their current home within the previous three, four, or five years. The resulting bunching estimates are very similar to those reported in Table 2, alleviating concerns about constrained buyer-sellers in the Toronto market.

To the extent that the policy might have other spillover effects, we rely on our data-driven method for model selection to appropriately determine, among other things, the estimation window and the size of the exclusion region. In Appendix E.1 we perform an extensive set of robustness checks to ensure that our estimates are not overly sensitive to the parameters chosen by our data-driven procedure. In particular, in Tables E1 and E2 we present results based on alternative criteria for parameter selection, with the estimation window widened by $25,000, with a fourth-order polynomial, and with the constraints in equations (9) and (10) imposed. Reassuringly, the bunching estimates are extremely robust, suggesting that our results are not driven by the selection of the size of the estimation window, order of the polynomial, or the width of the excluded region. Appendix E.1 also contains a graphical depiction of a larger set of robustness checks, providing further support in these regards.

Next, we perform two different types of *placebo tests* as additional checks of our identification strategy. First, we designate two years prior to the implementation of the million dollar policy as *placebo years*. There were no changes to policies specifically affecting houses around the $1M threshold during this time, and so we would not expect to find patterns of excess bunching.[36] Second, we designate alternative *placebo thresholds* at prices well below

[35]To see this, consider the possibility that sellers below $1M put down the minimum 5% with a 25-year amortizing loan when purchasing their home initially and are contemplating trading up to an above-$1M Toronto home. With a mortgage rate of 4% and annual house price appreciation of 5% , the seller of a $900K home in 2012 would have accumulated over $316K in home equity, which is enough for a 20% dowpayment on a $1.5M home purchase, provided they owned their home for at least 5 years.

[36]In a similar spirit, in Tables E7 and E8, we compare the last six months of 2011 with the first six months of 2012 in column (6) and the last six months of 2012 and the first six months of 2013 in column (7). The former are two periods before the million dollar policy, the latter are two periods after the policy. As expected, we find no evidence of excess bunching in asking or sales prices around the $1M threshold.

or above the \$1M threshold, and again estimate our baseline specification at each of these points. The idea is straightforward: since the million dollar policy generates a notch in the downpayment required of buyers at precisely \$1M, house prices in market segments well below or well above the \$1M cut-off should not be affected by the policy in a discontinuous manner. Table E3 contains 50 placebo estimates: 24 during the years overlapping the implementation of the policy for alternative price thresholds (the estimates excluding the \$1M threshold), and 26 during the pre-policy years. Of these, only 4 are statistically significant and only 1 is economically large. Taken together, these results support the notion that the bunching results presented in Section 5.1 provide an accurate measure of the threshold effects of the million dollar policy on house prices.

Finally, we assess the robustness of our main results to alternative choices of the pre- and post-policy periods. Our baseline specification groups pre- and post-policy periods by listing date and omits the few weeks following the announcement of the policy but before its implementation. In Appendix E.4, we show that our main results are not sensitive to these choices. We also show that our results are qualitatively robust to narrowing the pre- and post-periods from one year to six or three months.

5.2 Predictions 3 and 4: Bidding Wars

5.2.1 Sales-above-Asking and Time-on-the-Market

Turning to the policy effects on market liquidity, Prediction 3 states that the million dollar policy reduces expected time-on-the-market and increases the incidence of sales-above-asking for homes listed just under \$1M. The policy thus triggers a discontinuous increase in the expected time-on-the-market and a discontinuous decrease in the probability of selling-above-asking right at asking price \$1M. We bring this prediction to the data by employing a regression discontinuity design. The variables of interest are (1) the probability that a house sold above the asking price conditional on being listed at $p \geq y_j^A$; and (2) the probability that a house stayed on the market for more than two weeks conditional on being listed at $p \geq y_j^A$, where two weeks is roughly the median time-on-the-market in the sample. We construct these two variables in three steps.

First, we estimate the complementary CDFs, $\hat{S}_{Y\langle 0|0\rangle}(y_j^A) = 1-\hat{F}_{Y\langle 0|0\rangle}(y_j^A)$ and $\hat{S}_{Y\langle 1|1\rangle}(y_j^A) = 1 - \hat{F}_{Y\langle 1|1\rangle}(y_j^A)$, which represent the probability of a house being listed for at least y_j^A. Holding the distribution of housing characteristics the same as the pre-policy period using the reweighting method described in Section 4.2.1, we then estimate the counterfactual probability $\hat{S}_{Y\langle 1|0\rangle}(y_j^A) = 1 - \hat{F}_{Y\langle 1|0\rangle}(y_j^A)$. Second, we estimate the rescaled complementary CDFs, $RS_{Y\langle 0|0\rangle}(y_j^A, y^S \geq y^A)$ and $RS_{Y\langle 1|1\rangle}(y_j^A, y^S \geq y^A)$, which give the joint probability of a house being listed for at least y_j^A and selling above the asking price. Similarly, we estimate $\hat{RS}_{Y\langle 1|0\rangle}(y_j^A, y^S \geq y^A)$, the counterfactual rescaled complementary CDF, holding the distribution of housing characteristics the same as the pre-policy period. Finally, using the estimated probabilities above and Bayes' rule, we derive the conditional probability that a house is sold above asking conditional on being listed for at least y_j^A in the pre-policy period,

$$\hat{S}_{Y\langle 0|0\rangle}(y^S \geq y^A | y_j^A) = \frac{\hat{RS}_{Y\langle 0|0\rangle}(y_j^A, y^S \geq y^A)}{\hat{S}_{Y\langle 0|0\rangle}(y_j)},$$

and the corresponding counterfactual post-policy conditional probabilities,

$$\hat{S}_{Y\langle 1|0\rangle}(y^S \geq y^A | y_j^A) = \frac{\hat{RS}_{Y\langle 1|0\rangle}(y_j^A, y^S \geq y^A)}{\hat{S}_{Y\langle 1|0\rangle}(y_j)}.$$

Using this three-step procedure, we impute the two variables of interest: (1) the change in the probability of being sold above asking, $\hat{S}_{Y\langle 1|0\rangle}(y^S \geq y^A | y_j^A) - \hat{S}_{Y\langle 0|0\rangle}(y^S \geq y^A | y_j^A)$; and (2) the change in the probability of being on-the-market for more than two weeks, $\hat{S}_{Y\langle 1|0\rangle}(D \geq 14 | y_j^A) - \hat{S}_{Y\langle 0|0\rangle}(D \geq 14 | y_j^A)$. Both are constructed relative to the pre-policy period, conditional on being listed for at least y_j^A and holding the distribution of housing characteristics constant.

We plot each of the two constructed variables above as a function of the asking price, along with third order polynomials which are fit separately to each side of \$1M. In Figure 8a, changes in the probability of being sold above asking exhibit a discrete downward jump at \$1M, with an upward sloping curve to the left of \$1M. In 8b, changes in the probability of staying on the market for more than two weeks exhibit a discrete upward jump at $1M, with a downward sloping curve to the left of $1M. The evident discontinuities at $1M for time-on-the-market and sales-above-asking correspond exactly with Prediction 3, suggesting that the policy's minimal effect on sales prices can be at least partially attributed to heightened competition for homes listed just under $1M. Moreover, Figures 8c and 8d display changes in these two variables during two years prior to the policy. It is clear that changes in sales-above-asking can be represented by a smooth function of the asking price before the policy. Changes in time-on-the-market still exhibit a discrete jump at $1M even during the pre-policy periods, but the jump is statistically insignificant and much smaller in magnitude than during the policy periods. Together, these patterns are congruent with the finding that the policy caused a discrete jump in bidding intensity at the $1M and a heated market right below the $1M threshold.

5.2.2 Reallocation of Million Dollar Homes

Prediction 4 implies that the policy encourages an allocation of million dollar homes that favours less constrained over more constrained homebuyers. With a one-time-access to restricted proprietary mortgage data, we impute the fraction of *constrained buyers* (defined as having an LTV ratio above 80%) around the million dollar segment in our sample market during one year before and one year after the policy. For the segment slightly above $1M, the fraction of constrained buyers is reduced to zero, which is as intended. For the segment slightly below $1M, the fraction of buyers making downpayments of more than 20% remains high, even after the implementation of the policy. This is consistent with the model's implication that less constrained buyers have an incentive to participate in the segment below $1M because they can outbid constrained buyers in multiple offer situations. Thus, achieving the desired mortgage market outcome above $1M did not necessarily compromise the segment just below $1M. Lacking suitable micro-level mortgage data, we leave a formal test of these credit market implications for future research.

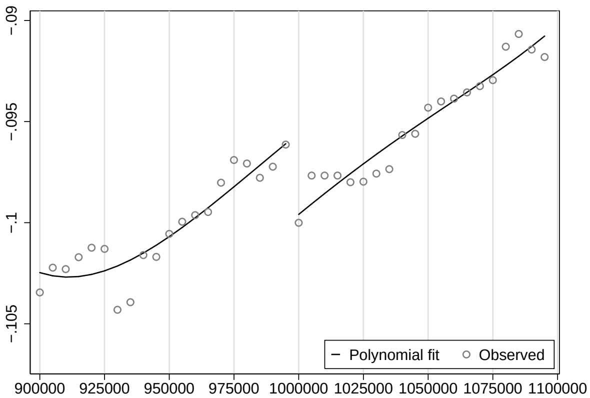

(a) Sales above Asking - Policy Period

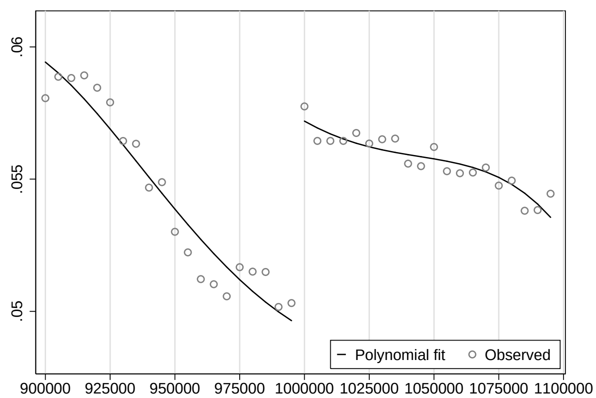

(b) Duration on Market - Policy Period

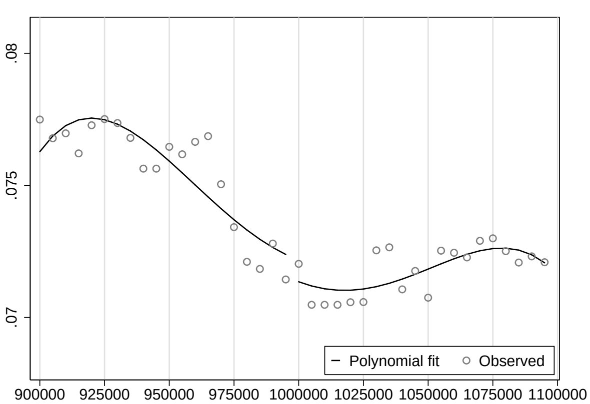

(c) Sales above Asking - Pre-Policy Period

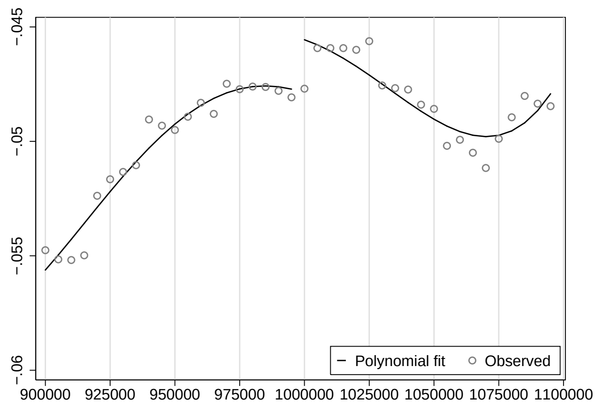

(d) Duration on Market - Pre-Policy Period

Figure 8
Policy Effects on Sales above Asking and Time on the Market

Notes: Panel (a) of the figure plots the change in the probability that a home is sold above asking, conditional on the asking price during the policy period. Panel (b) plots the change in the probability that a home is on the market for a duration longer than two weeks, conditional on the asking price during the policy period. Panels (c) and (d) repeat the analysis in panels (a) and (b), respectively, for the pre-policy period. Each dot represents the observed change in probability, while the solid line plots the predicted values from a second-order polynomial fit separately to either side of \$1M.

5.3 Policy Responses Above \$1M

The million dollar policy may not only affect homes around the \$1M threshold but homes above \$1M as well. Lacking a discrete change in the downpayment requirement in the segments far above \$1M, the bunching approach cannot be used to estimate policy consequences there. In this section, we design an alternative approach to examine the possible price responses above the \$1M threshold. An extensive margin response would occur if some transactions above the \$1M threshold did not transpire due to the additional financial constraint. Fewer transactions means less probability mass in segments above \$1M, and hence relatively more probability mass in segments below \$1M. Consequently, the CDF for post-policy prices would diverge above the counterfactual CDF, with the largest discrepancies at and around the \$1M price threshold. An intensive margin response, on the other hand, would occur if some prices above \$1M are lower than they would have been otherwise, in which case probability mass shifts to lower prices closer to \$1M. The post-policy CDF would again diverge above the counterfactual CDF, and the discrepancies would appear above the \$1M threshold.

We apply the distribution decomposition method proposed by Fortin, Lemieux, and Firpo (2011) to analyze the evolution of the price distributions over the sample period. Following the approach of Chernozhukov, Fernández-Val, and Melly (2013), we decompose the total change in the price distribution between the pre- and post-policy periods into three parts: (1) the effect of changes in house characteristics, (2) the effect of the changes in market conditions, and (3) any residual differences in price distributions. We consider part (3) to encompass all potential policy effects. Each of these can be measured as differences between appropriately chosen counterfactual distributions. We further provide confidence intervals for each component, thereby quantifying their economic and statistical significance. Consider the following decomposition:

$$\underbrace{[F_1 - F_0]}_{\text{observed difference}} = \underbrace{[F_1 - F_c]}_{\text{house characteristics}} + \underbrace{[F_c - \hat{F}_c]}_{\text{price counterfactual}} + \underbrace{[\hat{F}_c - F_0]}_{\text{market conditions}}, \tag{11}$$

where F_1 is the observed post-policy distribution, F_0 is the observed pre-policy distribution, F_c is the estimated composition-adjusted post-policy distribution (i.e., holding constant the pre-policy distribution of housing characteristics using the reweighting method discussed in Section 4.2.1), and $\hat{F}_c$ is a counterfactual composition-adjusted post-policy distribution. For the counterfactual distribution, $\hat{F}_c$, we use the pre-policy distribution, but shifted and rescaled to reflect market forces. In particular, we assume that the overall market trends between the the pre- and post-policy periods are captured by shifting and rescaling the pre-policy distribution along the horizontal axis.[37] Importantly, we use only the shape of the distribution below a cut-off to estimate the counterfactual shifting and rescaling so as to minimize the possibility of interpreting any policy effects as broad market trends. These shape-preserving differences in distributions are labelled "market conditions", whereas the differences due to changes in house characteristics are labelled "house characteristics". We are interested in differences that are not attributed to either the changing composition of homes sold or evolving market conditions, which we label "price counterfactual".[38]

More formally, we form the counterfactual post-policy distribution $\hat{F}_c$ using the pre-policy distribution F_0, but shifted and rescaled:

$$y = \hat{F}_c(p) \equiv F_0\left(\frac{p - \beta_0}{\beta_1}\right), \tag{12}$$

where β_0 is the shift (location) parameter and β_1 is the scale parameter. We want $\hat{F}_c$ to closely approximate the composition-constant post-policy distribution, F_c, for all prices less than some cut-off $\tau <$ \$1M. Inverting the distribution functions to obtain quantile functions yields the following relationship:

$$p = \hat{F}_c^{-1}(y) = \beta_0 + \beta_1 F_0^{-1}(y), \tag{13}$$

To achieve the desired approximation, we estimate β_0 and β_1 by regressing the quantiles of F_c below $\tau =$ \$900K on the corresponding quantiles of F_0 as well as a constant.[39] Denoting

[37] We shift and rescale along the horizontal rather than vertical axis to preserve the boundedness properties of a cumulative distribution function. This is equivalent to selecting a distribution $\hat{F}_c$ from the same *location-scale family* as the pre-policy distribution, F_0.

[38] Note that compared to the decomposition detailed previously in Section 4.2.1, we have further decomposed the "price structure" into "market conditions" and a residual "price couterfactual".

[39] The R-squared value for this regression is 0.9996, meaning the linear transformation of F_0^{-1} closely approximates F_c^{-1} for prices up to \$900K.

these estimates $\hat{\beta}_0$ and $\hat{\beta}_1$, the estimated counterfactual quantile function is

$$p = \hat{F}_c^{-1}(y) = \hat{\beta}_0 + \hat{\beta}_1 F_0^{-1}(y), \tag{14}$$

which can then be inverted to recover the estimated counterfactual post-policy distribution function, $\hat{F}_c$. Finally, we compare $\hat{F}_c$ to F_1 at prices above τ = \$900K to make inferences about policy effects above the threshold.

Our proposed method for disentangling market trends from other potential policy effects therefore relies on two identification assumptions: (1) that market trends in the absence of the policy can be suitably represented by an intercept and slope shift in the pre-policy quantile function (equivalently, a shifting and rescaling of the pre-policy distribution function); and (2) that these parameters can be estimated using only price segments below τ. To assess these assumptions, we apply the same procedure using only pre-policy sample periods. We also apply the same procedure to simulated data that feature no policy response, an extensive margin response, and an intensive margin response to further justify assumptions (1) and (2), and to show that our method can readily detect policy effects above \$1M. To further address assumption (2), we assess robustness to a lower τ cut-off: namely, \$800K.

We first summarize the results of the simulation exercises presented in Appendix F.3 of the supplemental material. In the absence of a policy response, a shape-preserving change in the distribution is well-summarized by a linear transformation applied to the pre-policy quantile function. Moreover, the intercept and slope coefficients are estimated reasonably well using only prices below a cut-off of \$900K.[40] Reassuringly given the absence of a simulated policy response, the price counterfactual is everywhere close to zero. We further simulate an extensive margin response to the policy by dropping 20% of prices above \$1M in the post-policy sample, and an intensive margin response by lowering prices in excess of \$1M by 30% of this excess amount. Using the proposed distribution decomposition method, both extensive and intensive margin responses are immediately apparent.

We present our empirical results in Figure 9, which plots the observed CDFs and their

[40]This still holds true if we lower the cut-off, τ, from \$900K to \$800K.

differences, along with the decomposition. To aid with data visualization, a smoothing algorithm was applied to each curve following Chernozhukov, Fernández-Val, and Melly (2013), and dots corresponding to ventiles of the post-policy price distribution illustrate how many transactions are represented by different segments of each curve. Panel (a) uses the main sample for the city of Toronto, whereas panel (b) focuses on the central district of Toronto. Two patterns emerge in both panels. First, the post-policy price distributions lie everywhere below the pre-policy distributions, reflecting an upward price trend in the Toronto market over time. Second, differences attributed to market conditions, as captured by the intercept and slope coefficients estimated using price data below \$900K and applied to the observed pre-policy distribution, account for nearly all of the observed differences between the pre- and post-policy CDFs in Figure 9. The price counterfactual differences left unexplained by market conditions and house characteristics are thus nearly indistinguishable from zero.[41] In particular, there appears to be no visual evidence of positive discrepancies in price segments around or above \$1M. Had the policy either inhibited sales or dampened prices in segments above \$1M, we would expect the price counterfactual to diverge above zero, as discussed above. Given the standard errors, we are unable to reject the hypothesis that the price counterfactual differences in these price segments are zero. This is true for both the city of Toronto (panel (a)), and the central district (panel (b)). This analysis suggests that the policy effects in the above \$1M segments are minimal.[42] This may not be surprising given that \$1M was at the 86th percentile of the house price distribution in 2012. Homes priced above \$1M thus represent very high-end segments. If buyers in these segments tend to be wealthy, many of them may not be financially constrained by a 20% downpayment requirement, in which case high prices could prevail from bidding competition among these less constrained buyers.

[41] The standard errors used to construct the confidence bands are obtained by bootstrapping our procedure 399 times.

[42] In Appendix F we perform several robustness exercises in generating our price counterfactual differences. In particular, we examine the sensitivity to a lower cut-off of $\tau = \$800$ and narrower estimation windows. The latter exercise should, in principle, minimize potentially confounding market trends. Our results are similar to those reported here. We also present the results as a decomposition of the differences in PDFs rather than CDFs.

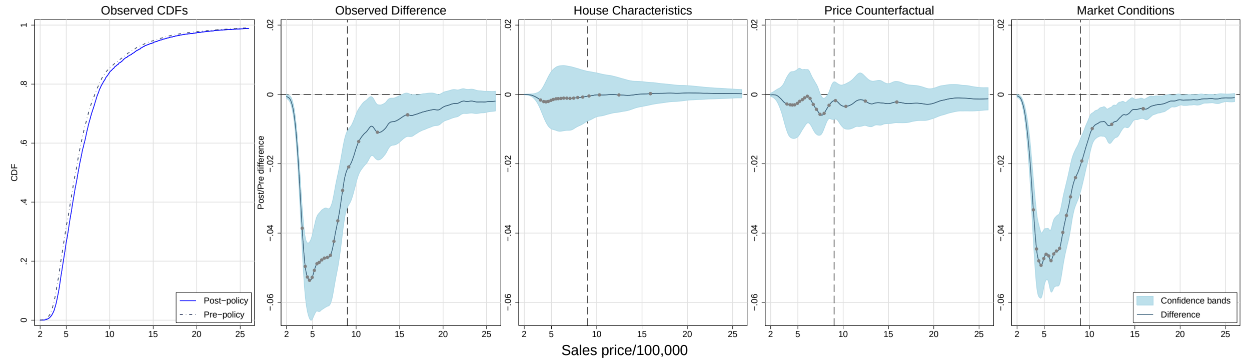

(a) City of Toronto

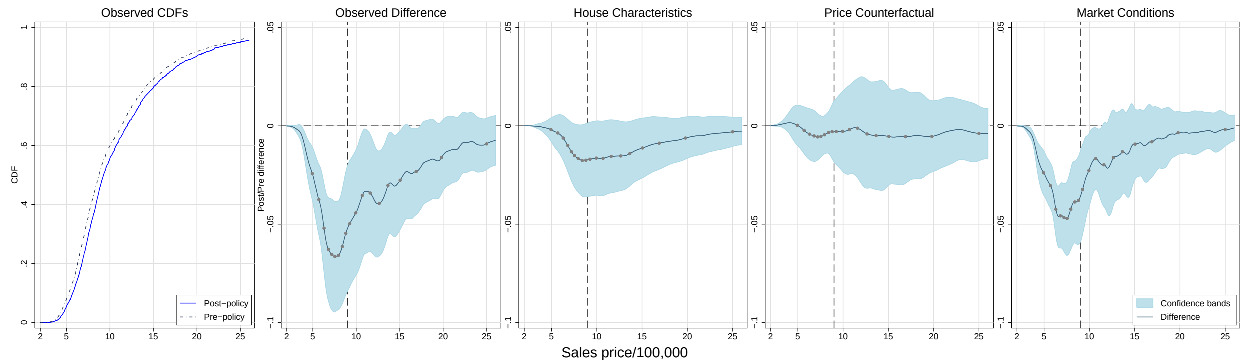

(b) Central District

Figure 9
Examining Policy Responses Above $1M (CDFs)

Notes: Panel (a) of the figure plots the pre- and post-policy sales price distributions, their differences, and the decomposition based on house characteristics, market conditions, and any residual price counterfactual differences for the city of Toronto. Panel (b) represents the same procedure for the central district. The cut-off for estimating market trends is set to $900K. The shaded area represents a 95% confidence interval, obtained via bootstrap.

6 Conclusion

In this paper we assess the impact of a financial constraint on price formation in the targeted segment of a frictional housing market. Our empirical methodology exploits a natural experiment arising from a mortgage insurance policy change that effectively imposes a 20 percent minimum downpayment requirement on homebuyers paying $1M or more. The interpretation of our results is motivated by a search-theoretic model of sellers competing for financially constrained buyers in the $1M segment of the housing market. We model the million dollar policy as a targeted financial constraint affecting a subset of prospective buyers. We show that sellers respond strategically by adjusting their asking prices to $1M, which attracts both constrained and unconstrained buyers. Because of the interactions of search, bidding and listing strategies of buyers and sellers, asking price effects translate into milder sales price effects.

We exploit the policy's $1M threshold to isolate the effects of the policy on prices and other housing market outcomes. Specifically, we implement an estimation procedure that combines a decomposition method with bunching estimation. Using housing market transaction-level data from the city of Toronto, we find that the million dollar policy results in excess bunching at $1M for asking prices but not for sales prices. These results, together with evidence that homes listed just below the $1M threshold sell faster with a higher incidence selling-above-asking, match the intuition derived from the theory. For segments well above $1M, we apply a distribution decomposition approach to uncover potential policy effects. We do not find evidence that the million dollar policy impacted home sales above the $1M threshold.

Overall, we find that the million dollar policy did not achieve the specific goal of cooling the housing boom, but instead heated a narrow segment of the market right below $1M. These findings are difficult to reconcile in a frictionless market, but are fully consistent with an equilibrium model of financial constraints with search frictions and auction mechanisms. Our analysis thus points to the importance of designing macroprudential policies that consider the underlying market microstructure and recognize the strategic responses of market participants.

References

Acharya, V. V., A. N. Berger, and R. Roman. 2018. Lending Implications of U.S. Bank Stress Tests: Costs or Benefits? *The Journal of Financial Intermediation* 34:58–90.

Albrecht, J., P. A. Gautier, and S. Vroman. 2014. Efficient entry in competing auctions. *The American Economic Review* 104:3288–96.

———. 2016. Directed search in the housing market. *Review of Economic Dynamics* 19:218–31.

Allen, J., T. Grieder, B. Peterson, and T. Roberts. 2016. The impact of macroprudential housing finance tools in canada: 2005-2010 .

Arefeva, A. 2016. How auctions amplify house-price fluctuations. *Mimeo* .

Begenau, J. 2019. Capital requirements, risk choice, and liquidity provision in a business-cycle model. *Journal of Financial Economics* .

Berger, D., N. Turner, and E. Zwick. 2020. Stimulating housing markets. *The Journal of Finance* 75:277–321.

Best, M. C., J. Cloyne, E. Ilzetzki, and H. Kleven. 2018. Estimating the elasticity of intertemporal substitution using mortgage notches. Working Paper 24948, National Bureau of Economic Research. doi:10.3386/w24948.

Best, M. C., and H. J. Kleven. 2018. Housing market responses to transaction taxes: Evidence from notches and stimulus in the uk. *The Review of Economic Studies* 85:157–93.

Cai, X., P. A. Gautier, and R. P. Wolthoff. 2017. Search frictions, competing mechanisms and optimal market segmentation. *Journal of Economic Theory* 169:453–73.

Che, Y.-K., and I. Gale. 1996a. Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. *Economics Letters* 50:373–9.

———. 1996b. Financial constraints in auctions: effects and antidotes. In M. Baye, ed., *Advances in Applied Microeconomics, Volume 6*, 97–120. JAI Press Inc.

———. 1998. Standard auctions with financially constrained bidders. *The Review of Economic Studies* 65:1–21.

Chernozhukov, V., I. Fernández-Val, and B. Melly. 2013. Inference on counterfactual distributions. *Econometrica* 81:2205–68.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri. 2011. Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *The Quarterly Journal of Economics* 126:749–804.

Clerc, L., A. Derviz, C. Mendicino, S. Moyen, K. Nikolov, L. Stracca, J. Suarez, and A. Vardoulakis. 2015. Capital regulation in a macroeconomic model with three layers of default .

DeFusco, A., S. Johnson, and J. Mondragon. 2019. Regulating household leverage. *Review of Economic Studies, Forthcoming* .

DeFusco, A. A., and A. Paciorek. 2017. The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy* 9:210–40.

Diamond, R., and P. Persson. 2016. The long-term consequences of teacher discretion in grading of high-stakes tests. Working Paper 22207, National Bureau of Economic Research. doi:10.3386/w22207.

Diaz, A., and B. Jerez. 2013. House prices, sales, and time on the market: A search-theoretic framework. *International Economic Review* 54:837–72.

DiNardo, J., N. M. Fortin, and T. Lemieux. 1996. Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica* 64:1001–44.

Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh. 2016. Phasing out the gses. *Journal of Monetary Economics* 81:111–32.

———. 2018. A macroeconomic model with financially constrained producers and intermediaries .

Favilukis, J., S. C. Ludvigson, and S. V. Nieuwerburgh. 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy* 125:140–223. doi:10.1086/689606.

Fortin, N., T. Lemieux, and S. Firpo. 2011. *Decomposition Methods in Economics*, vol. 4 of *Handbook of Labor Economics*, chap. 1, 1–102. Elsevier.

Gelman, A., and G. Imbens. 2014. Why high-order polynomials should not be used in regression discontinuity designs. Working Paper 20405, National Bureau of Economic Research. doi:10.3386/w20405.

Genesove, D., and L. Han. 2012. Search and matching in the housing market. *Journal of Urban Economics* 72:31 – 45. ISSN 0094-1190. doi:http://dx.doi.org/10.1016/j.jue.2012.01.002.

Han, L., and W. C. Strange. 2016. What is the role of the asking price for a house? *Journal of Urban Economics* 93:115–30.

Head, A., H. Lloyd-Ellis, and D. Stacey. 2018. Heterogeneity, frictional assignment and home-ownership. Working papers, Queen's University, Department of Economics.

Head, A., H. Lloyd-Ellis, and H. Sun. 2014. Search, liquidity, and the dynamics of house prices and construction. *American Economic Review* 104:1172–210.

Imbens, G. W., and T. Lemieux. 2008. Regression discontinuity designs: A guide to practice. *Journal of Econometrics* 142:615 – 635. ISSN 0304-4076. doi:https://doi.org/10.1016/j.jeconom.2007.05.001. The regression discontinuity design: Theory and applications.

IMF. 2014. Managing House Price Booms: The Role of Macroprudential Policies, Remarks at the IIM-Bangalore and IMF Conference on Housing Markets, Financial Stability and Growth, By Ratna Sahay, Deputy Director, Monetary and Capital Markets Department, IMF.

Jeske, K., D. Krueger, and K. Mitman. 2013. Housing, mortgage bailout guarantees and the macro economy. *Journal of Monetary Economics* 60:917–35.

Julien, B., J. Kennes, and I. King. 2000. Bidding for labor. *Review of Economic Dynamics* 3:619–49.

Khezr, P., and F. Menezes. 2018. Auctions with an asking price. *International Journal of Game Theory* 47:1329–50.

Kleven, H. J. 2016. Bunching. *Annual Review of Economics* 8:435–64. doi:10.1146/annurev-economics-080315-015234.

Kleven, H. J., and M. Waseem. 2013. Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan. *The Quarterly Journal of Economics* qjt004.

Kopczuk, W., and D. Munroe. 2015. Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market. *American Economic Journal: Economic Policy* 7:214–57.

Kotowski, M. H. 2016. First-price auctions with budget constraints .

Krainer, J. 2001. A Theory of Liquidity in Residential Real Estate Markets. *Journal of Urban Economics* 49:32–53.

Kuttner, K. N., and I. Shim. 2016. Can non-interest rate policies stabilize housing markets? Evidence from a panel of 57 economies. *Journal of Financial Stability* 26:31–44. doi:10.1016/j.jfs.2016.07.014.

Lamont, O., and J. C. Stein. 1999. Leverage and House-Price Dynamics in U.S. Cities. *RAND Journal of Economics* 30:498–514.

Lee, D. S., and T. Lemieux. 2010. Regression discontinuity designs in economics. *Journal of Economic Literature* 48:281–355.

Lester, B., L. Visschers, and R. Wolthoff. 2015. Meeting technologies and optimal trading mechanisms in competitive search markets. *Journal of Economic Theory* 155:1–15.

McAfee, R. P. 1993. Mechanism design by competing sellers. *Econometrica* 1281–312.

Mian, A., and A. Sufi. 2012. The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. *Quarterly Journal of Economics* 127:1107–42.

Ortalo-Magne, F., and S. Rady. 2006. Housing market dynamics: On the contribution of income shocks and credit constraints. *The Review of Economic Studies* 73:459–85. ISSN 00346527, 1467937X.

Peters, M., and S. Severinov. 1997. Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory* 75:141–79.

Pierce, J. R., and P. K. Schott. 2016. The surprisingly swift decline of us manufacturing employment. *American Economic Review* 106:1632–62.

Saez, E. 2010. Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy* 2:180–212. ISSN 1945-7731. doi:10.1257/pol.2.3.180.

Slemrod, J., C. Weber, and H. Shan. 2017. The behavioral response to housing transfer taxes: Evidence from a notched change in D.C. policy. *Journal of Urban Economics* 100:137–53.

Stacey, D. 2016. Commitment and costly signaling in decentralized markets. *International Economic Review* 57:1507–33.

Stein, J. C. 1995. Prices and trading volume in the housing market: A model with down-payment effects. *The Quarterly Journal of Economics* 110:379–406. ISSN 00335533, 15314650.

Van Bekkum, S., M. Gabarro, and R. M. Irani. 2017. Does a larger menu increase appetite? collateral eligibility and credit supply. *The Review of Financial Studies* 31:943–79.

Wheaton, W. C. 1990. Vacancy, Search, and Prices in a Housing Market Matching Model. *Journal of Political Economy* 98:1270–92.

Williams, J. T. 1995. Pricing Real Assets with Costly Search. *Review of Financial Studies* 8:55–90.

Supplemental Appendices: The Effects of a Targeted Financial Constraint on the Housing Market

A Delinquency and Credit Score Metrics

Figure A1
Delinquency Rates and Origination Credit Scores

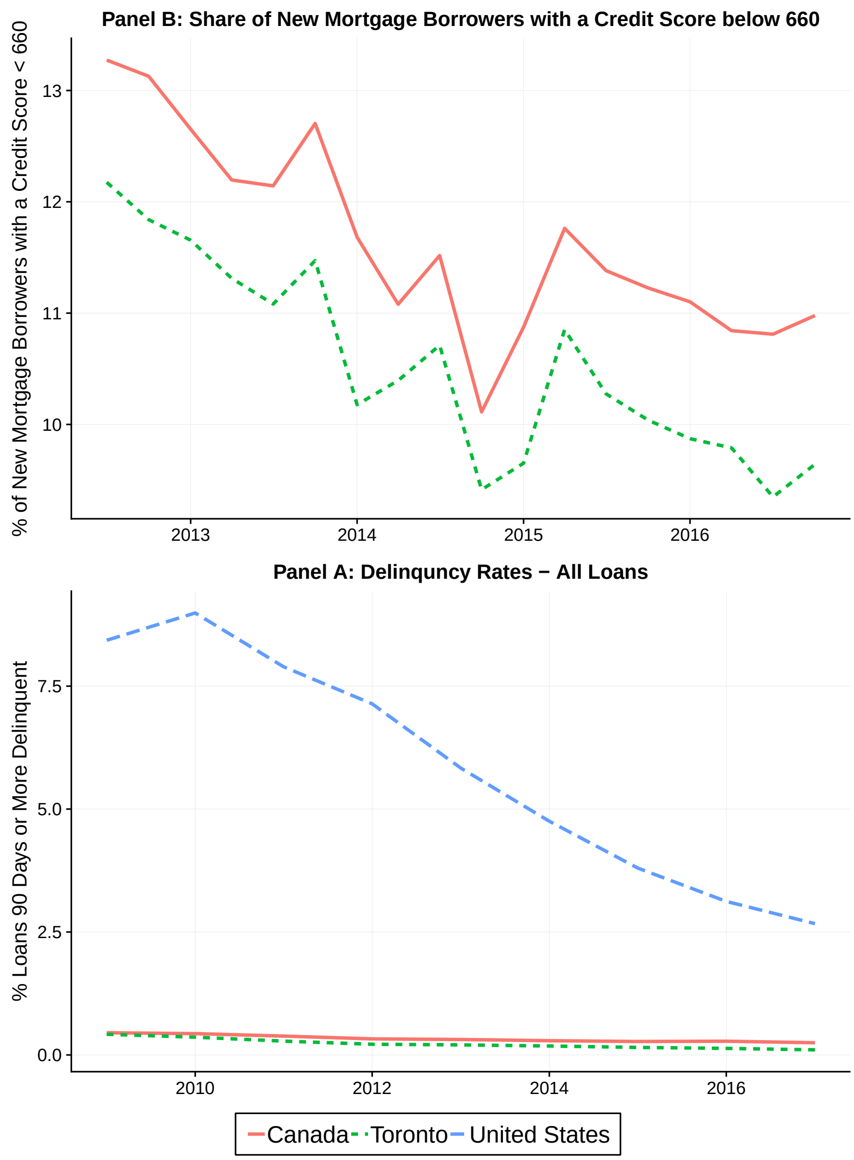

Notes: Data from Canada Mortgage and Housing Corporation (CMHC)

B Theory: Details and Derivations

B.1 Expected Payoffs

Expected payoffs are markedly different depending on whether the asking price, p, is above or below buyers' ability to pay. Consider each scenario separately.

Case I: $p \leq c$. Expected payoffs in this case, denoted $V_I^i(p,\lambda,\theta)$ for $i \in \{s,u,c\}$, are the ones derived in Section 3.2.2.

Case II: $c < p \leq u$. The seller's expected net payoff is

$$V_{II}^s(p,\lambda,\theta) = -x + \sum_{k=1}^{\infty} \pi(k)\phi_k(1)p + \sum_{k=2}^{\infty} \pi(k) \sum_{j=2}^{k} \phi_k(j)u.$$

The closed-form expression is

$$V_{II}^s(p,\lambda,\theta) = -x + (1-\lambda)\theta e^{-(1-\lambda)\theta}p + \left[1 - e^{-(1-\lambda)\theta} - (1-\lambda)\theta e^{-(1-\lambda)\theta}\right] u. \tag{B.1}$$

The second term reflects the surplus from a transaction if she meets exactly one unconstrained buyer; the third term is the surplus when matched with two or more unconstrained buyers.

The unconstrained buyer's expected payoff is

$$V_{II}^u(p,\lambda,\theta) = \pi(0)(v-p) + \sum_{k=1}^{\infty} \pi(k) \left[\phi_k(0)(v-p) + \sum_{j=1}^{k} \phi_k(j)\frac{v-u}{j+1}\right].$$

The closed-form expression is

$$V_{II}^u(p,\lambda,\theta) = \frac{1 - e^{-(1-\lambda)\theta}}{(1-\lambda)\theta}(v-u) + e^{-(1-\lambda)\theta}(u-p). \tag{B.2}$$

The first term is the expected surplus when competing for the house with other unconstrained bidders; the second term reflects additional surplus arising from the possibility of being the exclusive unconstrained buyer.

Since constrained buyers are excluded from the auction, their payoff is zero:

$$V_{II}^c(p,\lambda,\theta) = 0. \tag{B.3}$$

Case III: $p > u$. In this case, all buyers are excluded from the auction. Buyers' payoffs are zero, and the seller's net payoff is simply the value of maintaining ownership of the home (normalized to zero) less the listing cost, x:

$$V_{III}^s(p,\lambda,\theta) = -x, \quad V_{III}^u(p,\lambda,\theta) = 0 \quad \text{and} \quad V_{III}^c(p,\lambda,\theta) = 0. \tag{B.4}$$

Using the expected payoffs in each of the different cases, define the following value functions: for $i \in \{s, u, c\}$,

$$V^i(p, \lambda, \theta) = \begin{cases} V^i_{III}(p, \lambda, \theta) & \text{if } p > u, \\ V^i_{II}(p, \lambda, \theta) & \text{if } c < p \leq u, \\ V^i_I(p, \lambda, \theta) & \text{if } p \leq c. \end{cases} \tag{B.5}$$

B.2 Algorithm for Constructing Pre-Policy DSE

Solution to Problem P_0: Assuming (for the moment) an interior solution, the solution to problem P_0 satisfies the following first-order condition with respect to θ and the free-entry condition:

$$\begin{aligned} x &= [1 - e^{-\theta_u^*} - \theta_u^* e^{-\theta_u^*}]v \\ x &= \theta_u^* e^{-\theta_u^*} p^* + [1 - e^{-\theta_u^*} - \theta_u^* e^{-\theta_u^*}]u, \end{aligned}$$

which combine to yield

$$p_u^* = \frac{[1 - e^{-\theta_u^*} - \theta_u^* e^{-\theta_u^*}](v - u)}{\theta_u^* e^{-\theta_u^*}}. \tag{B.6}$$

Now taking into account the constraint imposed by bidding limit u, the solution is $p_0 = \min\{u, p_u^*\}$ and θ_0 satisfying $V^s(p_0, 0, \theta_0) = 0$.

Algorithm: If $\Lambda = 0$, set $\mathbb{P} = \{p_0\}$, $\theta(p_0) = \theta_0$, $\sigma(p_0) = \mathcal{B}/\theta_0$ and $\bar{V}^u = V^u(p_0, 0, \theta_0)$. For $p \leq u$, set θ to satisfy $\bar{V}^u = V^u(p, 0, \theta(p))$ or, if there is no solution to this equation, set $\theta(p) = 0$. For $p > u$ set $\theta(p) = 0$.

B.3 Algorithm for Constructing Post-Policy DSE

Solution to Problem P_1: Assuming (for the moment) an interior solution, the solution to problem P_1 satisfies the two constraints with equality, $V^s(p_c^*, \lambda_c^*, \theta_c^*) = 0$ and $V^u(p_c^*, \lambda_c^*, \theta_c^*) = \bar{V}^u$, and the following first-order condition.

$$\begin{aligned} e^{-\theta_c^*} p_c^* = &\left(1 - \frac{\left[1 - e^{-\theta_c^*} - \theta_c^* e^{-\theta_c^*}\right] v - x}{(1 - \lambda_c^*)\theta_c^*} \frac{1}{\bar{V}^u - \bar{V}^c}\right) \\ &\times \left(\frac{1 - e^{-(1-\lambda_c^*)\theta_c^*} - (1 - \lambda_c^*)\theta_c^* e^{-(1-\lambda_c^*)\theta_c^*}}{(1 - \lambda_c^*)\theta_c^*}(v - u) + (1 - \lambda_c^*)\lambda_c^*\theta_c^* e^{-(1-\lambda_c^*)\theta_c^*}(u - c)\right) \end{aligned}$$

where $\bar{V}^c = V^c(p_c^*, \lambda_c^*, \theta_c^*)$ and $\bar{V}^u$ is set equal to the maximized objective of problem P_0. Now taking into account the constraint imposed by bidding limit c, the solution is $p_1 = \min\{c, p_c^*\}$ with λ_1 and θ_1 satisfying $V^s(p_1, \lambda_1, \theta_1) = 0$ and $V^u(p_1, \lambda_1, \theta_1) = \bar{V}^u$.

Algorithm: If $0 < \Lambda \leq \lambda_1$, set $\mathbb{P} = \{p_0, p_1\}$, $\lambda(p_0) = 0$, $\theta(p_0) = \theta_0$, $\lambda(p_1) = \lambda_1$, $\theta(p_1) = \theta_1$, $\sigma(p_0) = (\lambda_1 - \Lambda)\mathcal{B}/(\lambda_1\theta_0)$ and $\sigma(p_1) = \Lambda\mathcal{B}/\lambda_1\theta_1$. The equilibrium values are $\bar{V}^u = V^u(p_0, 0, \theta_0) = V^u(p_1, \lambda_1, \theta_1)$ and $\bar{V}^c = V^c(p_1, \lambda_1, \theta_1)$. For $p \leq c$, set λ and θ to satisfy $\bar{V}^u = V^u(p, \lambda(p), \theta(p))$ and $\bar{V}^c = V^c(p, \lambda(p), \theta(p))$. If there is no solution to these equations with $\lambda(p) > 0$, set $\lambda(p) = 0$ and θ to satisfy $\bar{V}^u = V^u(p, 0, \theta(p))$; or if there is no solution to these equations with $\lambda(p) < 1$, set $\lambda(p) = 1$ and θ to satisfy $\bar{V}^c = V^c(p, 1, \theta(p))$. If there is still no solution with $\lambda(p) \in [0, 1]$ and $\theta(p) \geq 0$, set $\lambda(p)$ arbitrarily and set $\theta(p) = 0$. For $p \in (c, u]$, set $\lambda(p) = 0$ and θ to satisfy $\bar{V}^u = V^u(p, 0, \theta(p))$ or, if there is no solution to this equation, set $\theta(p) = 0$. Finally, for $p > u$, set $\lambda(p) = 0$ and $\theta(p) = 0$.

B.4 Omitted Proofs

Proof of Proposition 1. Construct a DSE as per the algorithms in Appendix B.2. Conditions 1(ii) and 2 of Definition 1 hold by construction. Condition 1(i) also holds for all $p > u$ because $V^s(p > u, \lambda, \theta) = -x$. To show that condition 1(i) holds for all $p \leq u$, suppose (FSOC) that there exists $p \leq u$ such that $V^s(p, 0, \theta(p)) > 0$, or

$$\theta(p)e^{-\theta(p)}p + \left[1 - e^{-\theta(p)} - \theta(p)e^{-\theta(p)}\right]u > x. \tag{B.7}$$

There exists $p' < p$ such that $V^s(p', 0, \theta(p)) = 0$, or

$$\theta(p)e^{-\theta(p)}p' + \left[1 - e^{-\theta(p)} - \theta(p)e^{-\theta(p)}\right]u = x.$$

Note, however, that

$$\bar{V}^u = \underbrace{\frac{1 - e^{-\theta(p)}}{\theta(p)}(v - u) + e^{-\theta(p)}(u - p)}_{V^u(p,0,\theta(p))} < \underbrace{\frac{1 - e^{-\theta(p)}}{\theta(p)}(v - u) + e^{-\theta(p)}(u - p')}_{V^u(p',0,\theta(p))}. \tag{B.8}$$

The equality follows by construction since inequality (B.7) requires $\theta(p) > 0$. The inequality follows from the fact that V^u is decreasing in the asking price and $p' < p$. The pair $\{p', \theta(p)\}$ therefore satisfies the constraint set of problem (P_0) and, according to (B.8), achieves a higher value of the objective than $\{p_0, \theta_0\}$: a contradiction. $\square$

Proof of Proposition 2. Construct a DSE as per the algorithm in Appendix B.3. Conditions 1(ii), 1(iii) and 2 of Definition 1 hold by construction. Condition 1(i) also holds for all $p > u$ because $V^s(p > u, \lambda, \theta) = -x$. To show that condition 1(i) holds for all $p \leq u$, suppose (FSOC) that there exists a profitable deviation: either (1) there exists $p \leq u$ such that $\lambda(p) = 0$ and $V^s(p, \lambda(p), \theta(p)) > 0$, or (2) there exists $p \leq c$ such that $\lambda(p) > 0$ and $V^s(p, \lambda(p), \theta(p)) > 0$.

For case (1), the contradiction can be derived in the same manner as in the proof of Proposition 1. For case (2), the profitable deviation under consideration is $V^s(p \leq c, \lambda(p), \theta(p)) > 0$, or

$$\theta e^{-\theta}p + \left[1 - e^{-\theta} - \theta e^{-\theta}\right]c + \left[1 - e^{-(1-\lambda)\theta} - (1-\lambda)\theta e^{-(1-\lambda)\theta}\right](u - c) > x, \tag{B.9}$$

where, for notational convenience, λ and θ refer to $\lambda(p)$ and $\theta(p)$. There exists $p'' < p$ such that $V^s(p'', \lambda, \theta) = 0$, or

$$\theta e^{-\theta} p'' + \left[1 - e^{-\theta} - \theta e^{-\theta}\right] c + \left[1 - e^{-(1-\lambda)\theta} - (1-\lambda)\theta e^{-(1-\lambda)\theta}\right](u - c) = x.$$

Note, however, that

$$\bar{V}^c = \underbrace{\frac{e^{-(1-\lambda)\theta} - e^{-\theta}}{\lambda\theta}(v - c) + e^{-\theta}(c - p)}_{V^c(p,\lambda,\theta)} < \underbrace{\frac{e^{-(1-\lambda)\theta} - e^{-\theta}}{\lambda\theta}(v - c) + e^{-\theta}(c - p'')}_{V^c(p'',\lambda,\theta)} \tag{B.10}$$

The equality follows by construction since inequality (B.9) requires $\theta > 0$ and, by assumption, $\lambda > 0$. The inequality follows from the fact that V^c is decreasing in the asking price and $p'' < p$. Similarly, $\bar{V}^u = V^u(p, \lambda, \theta) < V^u(p'', \lambda, \theta)$. The triple $\{p'', \lambda, \theta\}$ therefore satisfies the constraint set of problem (P_1) and, according to (B.10), achieves a higher value of the objective than $\{p_1, \lambda_1, \theta_1\}$: a contradiction. □

B.5 Numerical Simulation

To illustrate the predictions of the theory, we parameterize and simulate a version of the model that has been extended to incorporate a form of seller heterogeneity. Specifically, we assume that, upon listing their house for sale at cost x, a seller's reservation value is an idiosyncratic random variable that takes one of N possible values, $\{r_1, \ldots, r_N\}$ satisfying $r_1 < \cdots < r_N < u$, with equal probability.[43] The free entry condition on the supply side must now be satisfied in expectation. Modifying the model environment along this dimension does not affect the incentives facing buyers, but the expressions for sellers' expected net payoffs must be modified accordingly. For example, if the asking price is low enough to elicit bids from both unconstrained and constrained buyers, the expected net payoff for a seller with reservation value $r_n < c$ is

$$\begin{aligned} V_n^s(p \le c, \lambda, \theta) &= r_n - x + \pi(1)(p - r_n) + \sum_{k=2}^{\infty} \pi(k) \left\{ [\phi_k(0) + \phi_k(1)]\, c + \sum_{j=2}^{k} \phi_k(j) u - r_n \right\} \\ &= e^{-\theta} r_n - x + \theta e^{-\theta} p + \left[1 - e^{-\theta} - \theta e^{-\theta}\right] c \\ &\quad + \left[1 - e^{-(1-\lambda)\theta} - (1-\lambda)\theta e^{-(1-\lambda)\theta}\right](u - c). \end{aligned}$$

Sellers with different ex post reservation values will implement different asking price strategies, which permits the characterization of equilibria featuring *bunching from both above and below* simultaneously.

For our parameterization we set $c = 1000$, $v = 1100$ and $x = 50$, so that the unit of measurement corresponds to \$1,000 (CAD). We then choose $N = 200$ equally spaced reservation values, $\{r_1, \ldots, r_{200}\} = \{900, \ldots, 1041\}$, and set $u = 1070$ so that the pre-policy

[43] The interpretation is that a seller may not be certain about their value of moving/staying in advance. Their precise reservation value is ascertained at some stage of the listing process.

equilibrium asking prices range from 950 to 1,050. Finally, we calibrate the overall fraction of constrained buyers to $\Lambda = 0.0061$, which means only 0.61 percent of prospective buyers in this segment are constrained by the policy in this particular parametrization of the model.

The simulated distributions of asking prices are plotted in Figure B1a. These distributions have been rescaled to accommodate (unmodeled) asking prices outside of $[950, 1050]$. Specifically, we apply scale factors 0.0300 and 0.0390 to the pre- and post-policy asking price distributions to match the shares of overall listings in this segment of the Toronto market. Next, the pre- and post-policy distributions are anchored at 0.87 and 0.85 at asking price $c = 1000$ to mimic the percentiles of the \$1M home in the Toronto market in the two sample periods (see summary statistics in Table 1, discussed below in Section 4). Figure B1b plots the differences in the simulated distribution functions, along with a counterfactual obtained using the pre-policy asking price distribution, but shifted and rescaled to reflect the post-policy percentile at $c = 1000$ and measure of sellers in the $[950, 1050]$ segment. The discontinuity at the threshold is a visual representation of Prediction 1. Following the introduction of the policy, 25 out of the 200 seller types (12.5 percent of sellers) find it optimal to list at exactly price c. Of these sellers, roughly one third would have otherwise listed further below the threshold, the remaining two thirds would have listed above. Note that the excess mass of listings at policy threshold $c = 1000$ matches our estimates in Section 5. The share of constrained buyers, $\Lambda = 0.0061$, was chosen precisely to mimic this feature of the data.

Figure B2a displays the simulated sales price distributions, and Figure B2b plots their differences. The pre- and post-policy distributions are anchored at 0.86 and 0.84 at sales price $c = 1000$ to again mimic the summary statistics in Table 1. The scale factors applied to the simulated sales price distributions, however, are the same as those applied to the asking price distributions. Notice that the discontinuity in Figure B2b is much less pronounced than the discontinuity in Figure B1b: an unmistakable illustration of Prediction 2. Many of the homes listed with asking price c sell for more than c in a bidding war involving multiple unconstrained buyers.

To illustrate Prediction 3, we plot the expected time to sell (i.e., the reciprocal of the probability of selling) as well as the probability of selling-above-asking, as functions of the asking price, in Figures B3a and B4. The pre- and post-policy differences are plotted in Figures B3b and B4b. As discussed above, the liquidity of homes listed above the threshold is unaffected by the policy. In contrast, homes listed below the threshold post-policy attract both constrained and unconstrained buyers and consequently sell with higher probability and are more likely to sell for more than the asking price. The threshold nature of the policy induces discrete changes in these liquidity measures at asking price $c = 1000$.

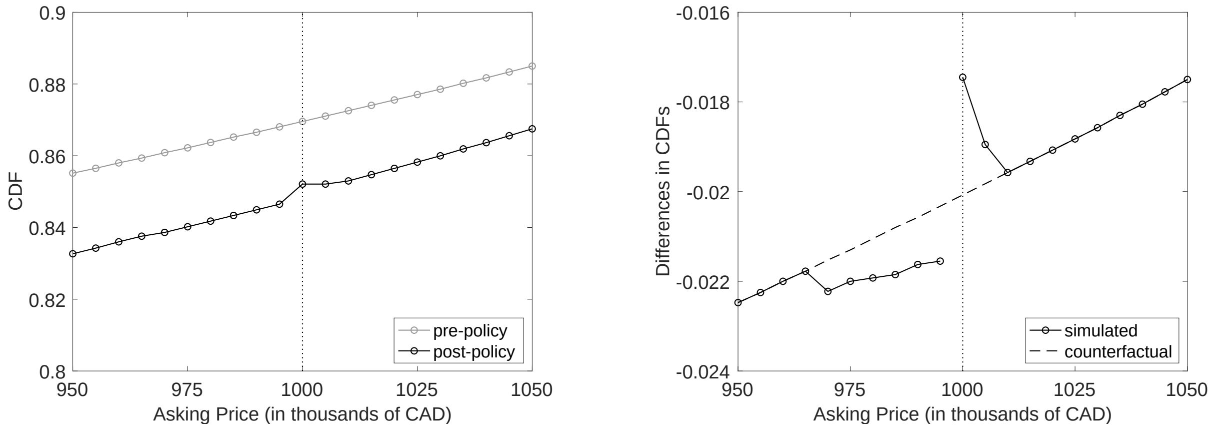

(a) Distribution Functions (b) Differences

Figure B1
Asking Prices

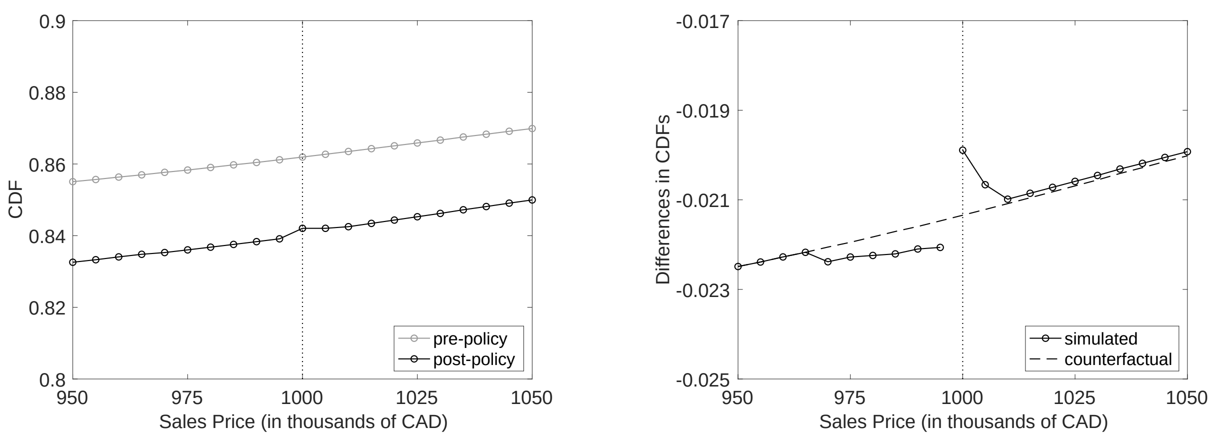

(a) Distribution Functions (b) Differences

Figure B2
Sales Prices

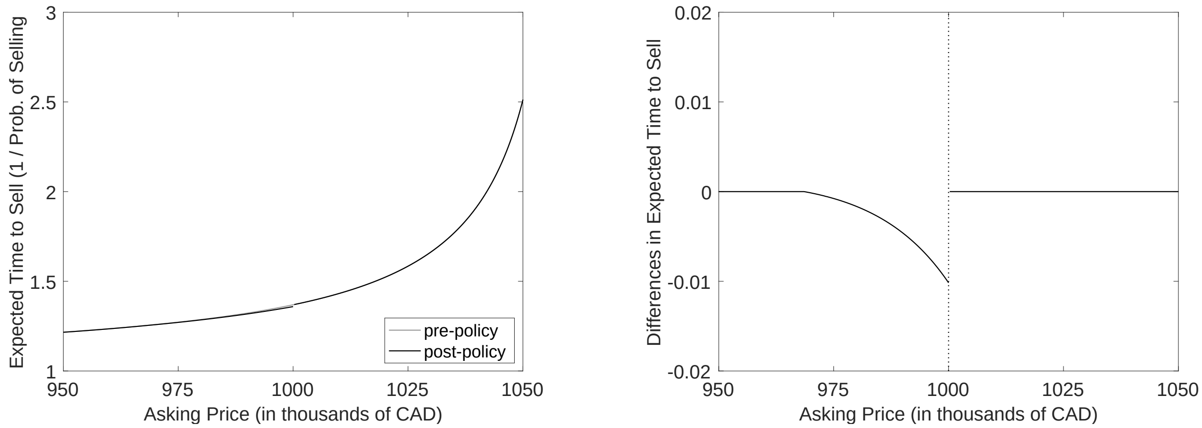

(a) Expected Time-to-Sell (b) Differences

Figure B3
Expected Time-to-Sell (Reciprocal of the Probability of Selling), by Asking Price

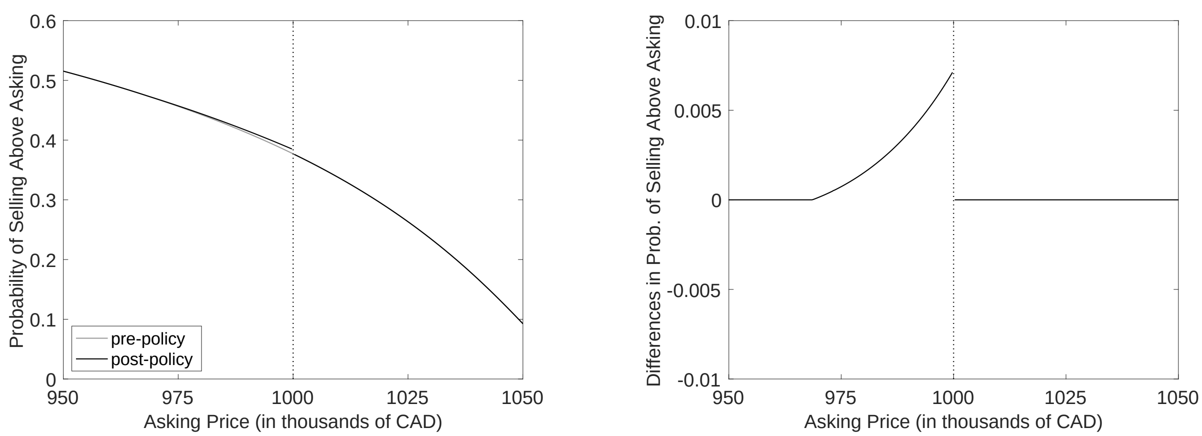

(a) Probability of Selling-Above-Asking (b) Differences

Figure B4
Probability of Selling-Above-Asking, by Asking Price

C Estimation Details

C.1 Cross Validation

We use a 5-fold cross validation procedure to select unknown hyperparameter involved in our estimation procedure. In particular, we aim to select the number of excluded bins to the left, L, and right R, the order of polynomial, p, and the estimation window, W, that determines how many house price bins are used in the estimation procedure. The latter can be thought of as a bandwidth choice for a local polynomial regression with rectangular weights (Imbens and Lemieux 2008). In order to select the quadruple $\theta \equiv \{L, R, p, W\}$ we use a minimum mean squared error criterion.

We begin our procedure by splitting the microdata on houses into 5 groups in a structured way. We cross validate both steps of our estimating procedure, first constructing the reweighted distribution functions and then estimating the bunching regression. Since the construction of the CDFs depend on ordered data, we respect this by sorting the data in increasing order of house price. We construct fold 1 by taking the observations $n_1 \in \{1, k+1, 2k+1, ...\}$, the 2nd fold by taking observations $n_2 \in \{2, k+2, 2k+2, ...\}$, and so on, where $k = 5$ in our implementation.

We estimate both steps of our empirical procedure using observations in $n_1, \ldots, n_4$ and the only the first step of our procedure (the construction of the empirical distribution) using observations in the 5th fold. We iterate the second step of our empirical procedure over a grid of potential hyperparameter values in the set $L, R \in \{1, 2, \ldots, 8\}$, $p \in \{2, 3\}$, and $W \in \{20, 25, 30\}$. For each combination of these values, we fit the bunching estimator on folds $n_1, \ldots, n_4$ and using the estimated coefficients, predict the residuals on n_5. When estimating our bunching estimator, we impose the adding up constraints given in (9) and (10) to assist in regularization. These restrictions are not imposed in our estimation in our main text. We repeat this procedure five times, holding out a different fold each time. For each choice of θ the cross-validation error is:

$$CV(\theta) = \frac{1}{W \cdot 5} \sum_{k}^{5} \sum_{j}^{W} \left(\hat{\Delta}_S(y_j)_{n_k} - \widehat{\hat{\Delta}_S(y_j)}_{\theta, n_{-k}} \right)^2$$

Where $\widehat{\hat{\Delta}_S(y_j)}_{\theta, n_{-k}}$ are fitted values for the k fold from the bunching estimator estimated on folds n_{-k} with parameter values θ, W is the estimation window (the number of observations used in the bunching estimation). We choose as our optimal hyperparameters:

$$\theta_{Opt.} = \text{argmin}_{\theta \in \{\theta_1, \ldots, \theta_V\}} CV(\theta)$$

where V is the total number of combinations of parameter values in $\{L, R, p, W\}$. We also compute the standard error for the cross-validation, letting

$$CV_k(\theta) = \frac{1}{W}\sum_{j}^{W}\left(\hat{\Delta}_S(y_j)_{n_k} - \widehat{\hat{\Delta}_S(y_j)}_{\theta,n_{-k}}\right)^2$$

we compute $SD(\theta) = \sqrt{\mathrm{Var}\left(CV_1(\theta), \ldots, CV_5\right)}$ and $SE = SD(\theta)/\sqrt{5}$ as the standard error of $CV(\theta)$. We use a 'one standard error rule': $CV(\theta) \leq CV(\theta_{opt.}) \pm SE(\theta_{Opt.})$ to find the widest and narrowest excluded region that is within one standard error of the optimum chosen $\theta_{Opt.}$ A graphical representation of this procedure is given in Figure C1. In the Panel (a), the root-mean squared error for each θ is plotted against the width of the excluded region (given by $L + R$) for asking price. Each dot on the figure represents one iteration of our procedure. The solid line gives the 'one standard error' rule. The points chosen by our procedure are labelled as (L, R). For instance, in panel (a), the optimum excluded region is given by $(4, 5)$, the narrowest by $(3, 4)$, and the widest by $(5, 5)$. Panel (b) shows the results for the sales prices. Notice that the CV values are much flatter and that no excluded region $(1, 1)$ is not rejected by the one standard error rule.

`../Figures-dfl-T2-General/Figure-CV_ask.pdf`

(a) Cross-Validation for Asking Prices

`../Figures-dfl-T2-General/Figure-CV_sold.pdf`

(b) Cross-Validation for Sales Prices

Figure C1
Cross-Validation Procedure

Notes: Panel (a) plots the root-mean squared error of each iteration of the cross-validation procedure for a given θ against the width of the excluded region for the city of Toronto using asking prices. Panel (b) plots the root-mean squared error of each iteration of the cross-validation procedure for a given θ against the width of the excluded region for the city of Toronto using sales prices. The chosen width of the excluded region is labelled on each panel.

D Supplemental Summary Statistics

Table D1
Summary Statistics for Condominiums and Townhouses: City of Toronto

		Pre-Policy		Post-Policy	
		Asking	Sales	Asking	Sales
All Condos	Mean	383058.22	377722.79	390153.40	381873.22
	25th Pct	275000.00	270000.00	278999.00	270000.00
	50th Pct	349000.00	343000.00	349900.00	345000.00
	75th Pct	439900.00	437000.00	454900.00	448000.00
	N	21768.00	21768.00	18257.00	18257.00
	Median Duration	20.00	20.00	24.00	24.00
	$1M Percentile	0.99	0.99	0.99	0.99
Condos $0.9–1.0M	N	92.00	89.00	100.00	105.00
	Median Duration	22.50	17.00	18.50	18.00
	Mean Price	963255.20	943872.89	968476.07	949377.90
Condos $1.0–1.1M	N	50.00	71.00	46.00	47.00
	Median Duration	22.50	30.00	19.50	23.00
	Mean Price	1071297.66	1051929.44	1073844.09	1046271.06

Notes: This table displays summary statistics for the city of Toronto for condominiums and townhouses. The pre-policy period is defined as July 15th, 2011, to June 15th, 2012, and the post-policy period is defined as July 15th, 2012, to June 15th, 2013. The columns labeled Asking refer to asking prices and the columns labeled Sales refer to sales prices. Duration refers to the number of days a home is on the market.

E Supplemental Material for the Bunching Estimation

In this section, we estimate an extensive set of specifications to assess the robustness of our main results.

E.1 Alternative Parameterizations: City of Toronto

Robustness to potential extensive margin responses in bunching estimates: Our first empirical exercise deals with the concern that the bunching estimates could be altered by plausible policy responses above the \$1M threshold. Suppose the introduction of the policy hindered potential listings/transactions above \$1M: an issue that is interesting in its own right and separately investigated in Section 5.3. In that case, our counterfactual difference in distributions, estimated by fitting a flexible polynomial through the empirical difference in distributions excluding an area around the \$1M threshold, would not accurately reflect what would have occurred in the absence of the policy. Note that this is a common issue in the bunching literature (Kopczuk and Munroe 2015; Best et al. 2018; Best and Kleven 2018). As suggested by Kleven (2016), we construct the counterfactual using only data below \$1M under the assumption that the distribution below the threshold is unaffected by the policy. These results are presented in column (2) of Tables E1 and E2, where we use the same excluded region and polynomial order as in our main specification (reproduced in column (1) for ease of reference), and extend the estimation window leftward in order to maintain the same number of bins despite only using data below \$1M.[44] These results are similar to the other columns of the tables, providing reassurance of the robustness of our main bunching estimates.

Robustness to alternative parametrizations: To the extent that the policy might have spillover effects below \$1M,[45] we rely on our data-driven method for model selection to appropriately determine, among other things, the estimation window and the size of the exclusion region. Our main estimation uses a 5-fold cross-validation procedure and a grid search over several values of each parameter. The results are reproduced in column (1) of Tables E1 and E2. Columns (3) and (4) use alternative excluded regions. These regions are chosen based on the cross-validation "plus one-standard error rule" outlined in Appendix

[44] In implementing this specification, we add additional constraints to the estimating procedure. In particular, to avoid unstable behavior of polynomial estimates near boundary points, we restrict the coefficients within the excluded region to be negative to the left of \$1M ($\gamma_l < 0$) and positive in the excluded region in the bins above \$1M ($\alpha_r > 0$). The idea behind these constraints are similar to natural splines that place shape restrictions near boundary values. Note that these constraints do not impose the adding up conditions in (9) or (10), but simply restricts the extrapolated polynomial counterfactual to lie between the observed $\hat{\Delta}_S(y_j)$ within the excluded region.

[45] One might be concerned that congestion externalities arising from, for example, decreasing returns to scale in the matching process between buyers and sellers, cause the excess bunching in listings just below \$1M to spill into even lower price segments. Such spillover effects are unlikely, however, as we find no evidence of excess under-listing in price bins below \$975K following the implementation of the policy (results available upon request).

C.1, where we choose the widest and narrowest excluded region specifications whose MSE is no more than one standard error above the minimum MSE obtained from the model in Column (1). Column (3) adds one excluded bin to the left of the threshold, whereas Column (4) subtracts one bin from the left and the right of the excluded region. These two specifications yield nearly identical results. Column (5) extends the estimation window by $25,000. Column (6) includes a fourth-order, rather than third-order polynomial used in the baseline specification. Column (7) imposes the constraints in equations (9) and (10) during estimation. Columns (3) – (7) of Table E2 present these same robustness checks for sales prices. Reassuringly, the bunching estimates are extremely robust, suggesting that our results are not driven by the selection of the size of the estimation window, order of the polynomial, or the width of the excluded region.

We present a larger set of specification checks in Figures E1 and E2 to further characterize the sensitivity of our bunching estimation results to various parameterizations. Rather than picking a few to report, we present a graphical representation of over 1,400 parameter combinations to display the sensitivity of our results to, for example, the width of the estimation window and the size of the excluded regions. As the estimation window gets larger, the polynomial should also become more flexible to accommodate observations further from the threshold. With this in mind, the extensive set of robustness specifications displayed in Figures E1 and E2 reveal once again that estimates of the threshold effects of the policy are remarkably robust.

Tables E1 and E2 correspond to the city of Toronto. Tables E5 and E6 in Appendix E.3 restrict samples to the central district and yield the same conclusions.

Placebo tests: Our next empirical exercise involves two *placebo tests* as additional checks of our identification strategy. We first designate two years prior to the implementation of the million dollar policy as *placebo years*. Specifically, we estimate our baseline specification for asking and sales prices (columns (1) and (2) of Table 2) to compare the distribution of house prices between the period from July 15th, 2011, to June 15th, 2012, and the period from July 15th, 2010, and June 15th, 2011. Between and during these time periods, there were no changes to policies specifically affecting houses around the $1M threshold, and so we would not expect to find patterns of excess bunching. The middle row of Table E3 presents the results. The total observed jump at $1M is 0.0006 for the asking price and 0.0001 for the sales price. Neither are statistically significant, as expected. In a similar spirit, in Tables E7 and E8, we compare the last six months of 2011 with the first six months of 2012 in column (6) and the last six months of 2012 and the first six months of 2013 in column (7). The former are two periods before the million dollar policy, the latter are two periods after the policy. As expected, we find no evidence of excess bunching in asking or sales prices around the $1M threshold.

Second, we designate alternative *placebo thresholds* at prices well below or above the $1M threshold, and again estimate our baseline specification at each of these points. The idea is straightforward: since the million dollar policy generates a notch in the downpayment required of buyers at precisely $1M, house prices in market segments well below or

Table E1
Robustness Checks to Alternative Parameterizations (Asking Price): City of Toronto

	Asking Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0045*	0.0045*	0.0045*	0.0045*	0.0045*	0.0045*	0.0045*
	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)
Total Response	0.0039*	0.0041*	0.0039*	0.0039*	0.0035*	0.0040*	0.0049*
	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0010)	(0.0013)
From Below	-0.0018*	-0.0019*	-0.0018*	-0.0018*	-0.0018*	-0.0024*	-0.0022*
	(0.00072)	(0.00088)	(0.00074)	(0.00069)	(0.00078)	(0.00074)	(0.00087)
From Above	0.0020*	0.0022	0.0021*	0.0021*	0.0017	0.0016	0.0027*
	(0.00089)	(0.0013)	(0.00088)	(0.00087)	(0.00088)	(0.00093)	(0.00096)
Observations	41305	41305	41305	41305	41305	41305	41305
Excluded Bins:							
L	4	4	5	3	4	4	4
R	5	5	5	4	5	5	5
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-.0013	-.0012	-.0011	-.00063	-.0016	-.0023	
	(.0013)	(.00096)	(.0018)	(.00061)	(.0015)	(.0012)	
$A - \sum_r^R \beta_A^r$	.0025	.00049	.0025	.0018	.0032	.00093	
	(.0016)	(.0064)	(.0016)	(.0011)	(.002)	(.0018)	
Joint p-val.	0.20	0.46	0.27	0.19	0.15	0.17	
Impact:							
Δ Houses at cutoff	85.9	91.4	86.0	85.7	77.7	88.9	108.5
Specifications:							
Poly. Order	3	3	3	3	3	4	3
Window	20	20	20	20	25	20	20
Other	CV Opt.	Extensive	CV Wide	CV Narrow			Constrained

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using asking prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

Table E2
Robustness Checks to Alternative Parameterizations (Sales Price): City of Toronto

	Sales Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.00094*	0.00094*	0.00094*	0.00094*	0.00094*	0.00094*	0.00094*
	(0.00042)	(0.00042)	(0.00042)	(0.00042)	(0.00042)	(0.00042)	(0.00042)
Total Response	0.00050	0.00032	0.00035	0.00051	0.00048	0.00054	0.00035
	(0.00053)	(0.00049)	(0.00052)	(0.00055)	(0.00047)	(0.00054)	(0.00045)
From Below	0.00017	-5.3e-38	0.00051	0.00015	0.00016	-0.0000099	0.00013
	(0.00060)	(0.00023)	(0.00077)	(0.00060)	(0.00066)	(0.00051)	(0.00060)
From Above	0.00067	0.00032	0.00086	0.00065	0.00065	0.00053	0.00048
	(0.00079)	(0.00054)	(0.00090)	(0.00077)	(0.00080)	(0.00073)	(0.00058)
Observations	41305	41305	41305	41305	41305	41305	41305
Excluded Bins:							
L	1	1	3	2	1	1	1
R	2	2	8	1	2	2	2
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	.00017	-5.3e-38	.00058	.000058	.00016	-9.9e-06	
	(.0006)	(.00023)	(.0009)	(.00037)	(.00066)	(.00051)	
$A - \sum_r^R \beta_A^r$	-.00036	-.00032	.0043	.00065	-.00027	-.00039	
	(.00057)	(.00035)	(.004)	(.00077)	(.00055)	(.00058)	
Joint p-val.	0.82	0.64	0.52	0.67	0.88	0.74	
Impact:							
Δ Houses at cutoff	11.1	7.03	7.83	11.3	10.8	11.9	7.75
Specifications:							
Poly. Order	3	3	3	3	3	4	3
Window	20	20	20	20	25	20	20
Other	CV Opt.	Extensive	CV Wide	CV Narrow			Constrained

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using sales prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

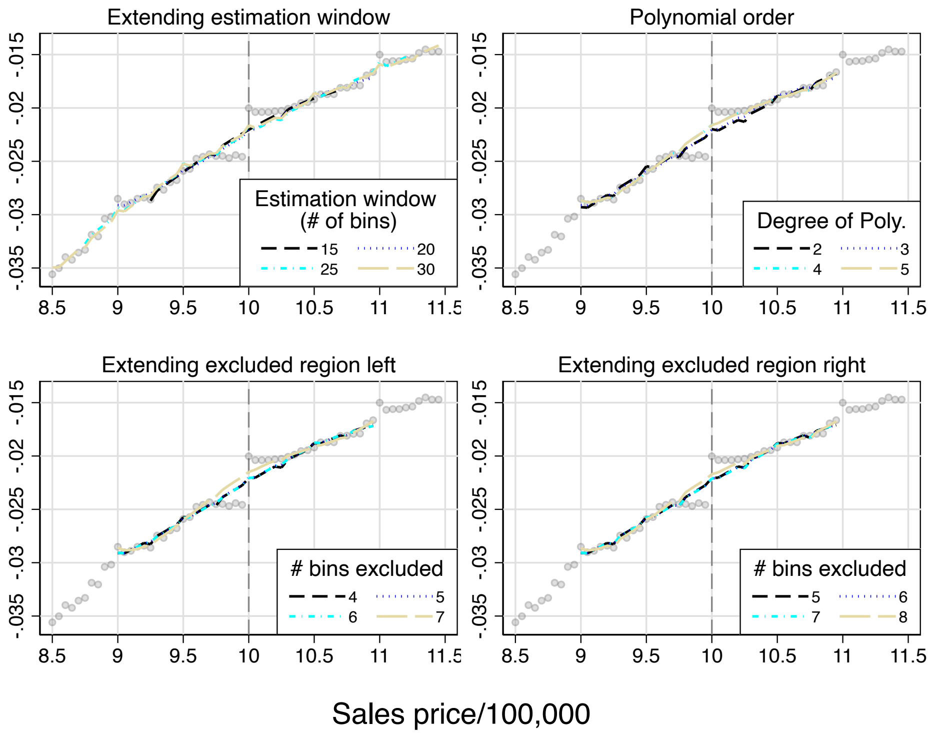

Figure E1
Examining Robustness to Deviations from Baseline

Notes: Each panel of the figure examines robustness to alternative parametrizations of our estimation procedure. Each figure displays our optimal specification (estimation widow = 20, polynomial = 3, excluded left = 4, excluded right = 5) and assesses 3 deviations from this by varying the estimation window (Panel A), the polynomial order (Panel B), the excluded bins on the left (Panel C) and the excluded bins on the right (Panel D).

well above the $1M cut-off should not be affected by the policy in a discontinuous manner. Excess bunching at selected placebo thresholds (multiples of $25,000 between $800,000 and $1,150,000) would thus represent contradictory evidence. Table E3 contains 50 placebo threshold estimates: 24 during the years overlapping the implementation of the policy for alternative price thresholds (the estimates excluding the $1M threshold), and 26 during the pre-policy years. Out of the 50 bunching estimates, only 4 are statistically significant and only 1 is economically large. Most estimates are statistically insignificant and economically small. Taken individually, each estimate alone may not be sufficient to alleviate concerns regarding marketing convention, psychology bias or other threshold factors unrelated to mortgage insurance regulation. But all together, these estimates provide compelling evidence that the bunching results presented in Section 5.1 provide an accurate measure of the threshold effects of the million dollar policy on house prices.

Table E3
Robustness Checks to Alternative Cut-offs and Period

	Post-Policy Difference		Pre-Policy Difference	
	(1) Asking	(2) Sales	(3) Asking	(4) Sales
800000	0.0017	-0.0015*	0.0018	-0.0000097
	(0.0015)	(0.00077)	(0.0014)	(0.00073)
825000	-0.00015	-0.000038	-0.00023	0.00032
	(0.00056)	(0.00077)	(0.00047)	(0.00071)
850000	0.00100	0.00068	0.00088	-0.000087
	(0.0013)	(0.00078)	(0.0012)	(0.00069)
875000	0.00011	-0.000064	0.00030	0.00034
	(0.00045)	(0.00070)	(0.00038)	(0.00063)
900000	0.00037	-0.00059	0.0030*	0.00039
	(0.0013)	(0.00072)	(0.0012)	(0.00059)
925000	-0.00055	0.00099	-0.00016	0.00033
	(0.00042)	(0.00066)	(0.00035)	(0.00059)
950000	-0.00098	-0.00022	-0.0015	-0.00085
	(0.00090)	(0.00062)	(0.00081)	(0.00051)
1000000	0.0039*	0.00051	0.00062	0.00013
	(0.0011)	(0.00054)	(0.00092)	(0.00043)
1050000	-0.0015*	0.00030	-0.00027	-0.000018
	(0.00066)	(0.00044)	(0.00058)	(0.00037)
1075000	-0.00038	-0.00055	0.00011	0.00026
	(0.00023)	(0.00043)	(0.00022)	(0.00041)
1100000	0.0012	-0.00011	0.00088	-0.00045
	(0.00073)	(0.00036)	(0.00066)	(0.00031)
1125000	-0.000066	-0.00062	0.00020	0.00041
	(0.00025)	(0.00036)	(0.00020)	(0.00036)
1150000	-0.00056	0.00055	-0.0012*	-0.00064
	(0.00061)	(0.00040)	(0.00060)	(0.00033)

Notes: This table displays the bunching estimates at various price thresholds for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using either asking prices (columns 1 and 3) or sales prices (columns 2 and 4). Each row of the table shows the total policy component of equation (8) using price thresholds indicated in the left-side panel. The post-policy difference (columns 1 and 2) use data on sales one year before and after the million dollar policy. The pre-policy difference (columns 3 and 4) compares the two years prior to the implementation. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

E.2 Robustness to Exclusion of Short Ownership Spells

Our analysis in Section 5 hinges on the assumption that homes further below \$1M are unaffected by the policy. This is by design of the million dollar policy. One legitimate concern, however, is that the policy may have unintended consequences in market segments below the threshold. For example, suppose a seller of a below-\$1M home intends to trade-up to an above-\$1M home. The seller may be *constrained* in that the proceeds from the sale of their current home must not compromise their ability to make a 20% downpayment on their next home. To that end, there may be an incentive to raise the asking price and then patiently await a buyer with a high willingness to pay. This could affect our estimated counterfactual and hence the validity of the bunching estimates if the revised listing/selling strategy affects prices within the estimation window but below the excluded region.

To address this concern, we consider the possibility that sellers below \$1M put down the minimum 5% with a 25-year amortizing loan when purchasing their home initially and are contemplating trading up to an above-\$1M Toronto home. With an insured mortgage interest rate of 4% and 5% annual house price appreciation, the seller of a \$900K home in 2012 would have accumulated over \$316K in home equity, which is enough for a 20% dowpayment on a \$1.5M home purchase, provided they owned their home for at least 5 years (i.e., purchased the home in 2007 or earlier).[46] In Table E4 of Appendix E.2, we report specifications that exclude sellers that bought their current home within the previous three, four, or five years. Excluding sellers in price segments below \$1M with short ownership spells ensures that the estimated counterfactual is unaffected by a trading-up constraint. As shown in Table E4, the resulting bunching estimates are very similar to those reported in Table 2, alleviating concerns about constrained buyer-sellers in the Toronto market.

[46] A shorter ownership spell implies potentially less home equity. Under similar assumptions, the seller of a \$900K home in 2012 would have accumulated nearly \$268K in home equity if they purchased the home in 2008 (i.e., a 4-year ownership spell), and around \$216K if they purchased in 2009 (i.e., a 3-year ownership spell).

Table E4
Assessing Robustness to Excluding Short Ownership Spells

	3-years		4-years		5-year	
	(1) Asking	(2) Sales	(3) Asking	(4) Sales	(5) Asking	(6) Sales
Jump at cut-off	0.0044*	0.00090	0.0045*	0.00094*	0.0046*	0.0010*
	(0.0011)	(0.00047)	(0.0011)	(0.00046)	(0.0013)	(0.00049)
Total Response	0.0038*	0.00045	0.0038*	0.00062	0.0039*	0.00066
	(0.0011)	(0.00058)	(0.0012)	(0.00057)	(0.0012)	(0.00062)
From Below	-0.0018*	-0.000046	-0.0018*	0.000042	-0.0020*	0.000023
	(0.00082)	(0.00059)	(0.00082)	(0.00063)	(0.00082)	(0.00071)
From Above	0.0020*	0.00040	0.0020*	0.00066	0.0019*	0.00069
	(0.00087)	(0.00079)	(0.00095)	(0.00080)	(0.00093)	(0.00085)
Observations	38053	38053	36281	36281	34360	34360
Excluded Bins:						
L	4	1	4	1	4	1
R	5	2	5	2	5	2
Tests of Fit:						
$B - \sum_l^L \beta_B^l$	-.00085	-.000046	-.00095	.000042	-.00065	.000023
	(.0014)	(.00059)	(.0014)	(.00063)	(.0014)	(.00071)
$A - \sum_r^R \beta_A^r$	.0026	-.00033	.0028	-.00053	.0026	-.00056
	(.0016)	(.0006)	(.0016)	(.00066)	(.0017)	(.00068)
Joint p-val.	0.24	0.81	0.17	0.70	0.32	0.67
Impact:						
Δ Houses at cutoff	77.7	9.14	73.9	12.1	71.0	12.1
Specifications:						
Poly. Order	3	3	3	3	3	3
Window	20	20	20	20	20	20

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto when houses with short ownership spells are removed from the sample. Short ownership spells are defined as houses that where bought and sold within 3 years (columns 1 and 2), 4 years (columns 3 and 4), or 5 years (columns 5 and 6). The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using using either asking prices (columns 1, 3 and 5) or sales prices (columns 2, 4 and 6). The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

Figure E2
Robustness of Excess Bunching Estimates to Various Alternative Parametrizations

Notes: Each panel of the figure examines the robustness of estimates to alternative parametrizations of our estimation procedure. Panel A shows how the parameter Total varies as we alter the polynomial order (displayed in each column), the estimation window around the $1M threshold (the x-axis), and the excluded region (number of bins excluded to the left + the number of bins excluded to the right, and highlighted by color shading). The yellow dots correspond to our optimal excluded region (left = 4, right = 5) and the baseline specification (Column 1 of Table 2 in the main text) is highlighted in the second column and illustrated by the dashed line in other columns. Panel B shows the corresponding t-statistics. In panel C, we show the total fraction of bunching that comes from Above (Above / Total).

E.3 Alternative Specifications: Central District

Table E5
Robustness of Regression Bunching Estimates to Alternative Parametrizations: Central Toronto

	Asking Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0094*	0.0094*	0.0094*	0.0094*	0.0094*	0.0094*	0.0094*
	(0.0032)	(0.0032)	(0.0032)	(0.0032)	(0.0032)	(0.0032)	(0.0032)
Total Response	0.0068*	0.0069*	0.0069*	0.0072*	0.0068*	0.0096*	0.0088*
	(0.0031)	(0.0031)	(0.0031)	(0.0030)	(0.0030)	(0.0043)	(0.0030)
From Below	-0.0049*	-0.0047	-0.0049*	-0.0051*	-0.0049*	-0.0052	-0.0053
	(0.0022)	(0.0024)	(0.0022)	(0.0023)	(0.0022)	(0.0027)	(0.0028)
From Above	0.0020	0.0021	0.0019	0.0021	0.0020	0.0044	0.0035
	(0.0025)	(0.0026)	(0.0024)	(0.0024)	(0.0024)	(0.0031)	(0.0037)
Observations	9008	9008	9008	9008	9008	9008	9008
Excluded Bins:							
L	3	5	3	3	3	3	3
R	4	6	3	4	4	4	4
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-.00045	.0019	-.00051	-.00065	-.00045		-.001
	(.0022)	(.0068)	(.0022)	(.0023)	(.0021)	(.)	
$A - \sum_r^R \beta_A^r$	.0048	.0087	.0031	.0039	.0048		.0023
	(.0041)	(.0087)	(.002)	(.0046)	(.0036)	(3.9e-11)	
Joint p-val.	0.52	0.59	0.32	0.68	0.42	.	0.76
Impact:							
Δ Houses at cutoff	33.8	33.9	33.9	35.7	33.8	47.5	43.6
Specifications:							
Poly. Order	2	2	2	2	3	2	2
Window	25	25	25	30	25	25	25
Other	CV Opt.	CV Wide	CV Narrow			Constrained	Extensive

Notes: This table displays the bunching estimates of the million dollar policy for central Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using asking prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

Table E6
Robustness of Regression Bunching Estimates to Alternative Parametrizations: Central Toronto

	Sales Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0032*	0.0032*	0.0032*	0.0032*	0.0032*	0.0032*	0.0032*
	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0014)
Total Response	0.0028	0.0028	0.0028	0.0027	0.0026	0.0022	0.0034
	(0.0017)	(0.0016)	(0.0017)	(0.0016)	(0.0018)	(0.0016)	(0.0021)
From Below	-0.0013	-0.0015	-0.0014	-0.0019	-0.0013	-0.0015	0
	(0.0018)	(0.0019)	(0.0018)	(0.0020)	(0.0018)	(0.0018)	(0.0011)
From Above	0.0014	0.0013	0.0014	0.00086	0.0013	0.00069	0.0034
	(0.0022)	(0.0023)	(0.0021)	(0.0023)	(0.0023)	(0.0017)	(0.0026)
Observations	9008	9008	9008	9008	9008	9008	9008
Excluded Bins:							
L	1	1	1	1	1	1	1
R	2	4	1	2	2	2	2
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-.0013	-.0015	-.0014	-.0019	-.0013	-.0015	0
	(.0018)	(.0019)	(.0018)	(.002)	(.0018)	(.0018)	(.0011)
$A - \sum_r^R \beta_A^r$	-.0012	-.0018	.0014	-.0011	-.001	0*	-.00062
	(.0018)	(.0048)	(.0021)	(.0017)	(.0017)	(0)	(.0014)
Joint p-val.	0.45	0.70	0.25	0.31	0.46	0.43	0.90
Impact:							
Δ Houses at cutoff	13.7	13.9	13.7	13.5	13.1	10.7	16.7
Specifications:							
Poly. Order	2	2	2	2	3	2	2
Window	20	20	20	25	20	20	20
Other	CV Opt.	CV Wide	CV Narrow			Constrained	Extensive

Notes: This table displays the bunching estimates of the million dollar policy for central Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using sales prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

E.4 Robustness to Alternative Timing Choices

E.4.1 Alternative Timing Choices

We assess the robustness of our main results to alternative choices of the pre- and post-policy periods. Our baseline specification groups pre- and post-policy periods by listing date. In column (1) of Table E7, we instead use the sold date to distinguish the two periods. Our baseline omits the few weeks following the announcement of the policy but before its implementation. In column (2) of Table E7, we use the entire year including the announcement period. This gives results that are very similar to our baseline. Column (3) takes the post-policy period to be the six months following implementation and the pre-policy period to be the six months prior to implementation. The results using this timing choice are slightly greater in magnitude than our baseline results. To further minimize the possible confounding effects of changing market conditions, the pre- and post-policy periods are further reduced to three-month windows in column (4). The estimates become less significant due to a much reduced sample size, but the coefficients remain comparable. One issue with columns (3) and (4) is that they compare two different sets of months within the same year, and so one might worry about seasonality of housing sales. Column (5) addresses this by again taking a six-month window, but using as a pre-policy period the same six calendar months in the prior year. By leaving out the six months between the two periods, we likely also circumvent potential issues related to properties listed before the policy but sold after. Across all these specifications, the bunching estimates are remarkably consistent, highlighting the robustness of the main findings. Table E8 repeats this exercise using sales prices. Again, we find that our main results are not sensitive to the choice of the pre- and post-policy time periods.[47]

[47] Tables E9 and E10 in Appendix E.4.2 present the same specifications as in columns (2) – (7) of Tables E7 and E8, except that pre- and post-policy periods are defined instead by sold date rather than listing date. The results are nearly identical.

Table E7
Robustness Checks to Alternative Time Choices (Asking Price)

	Asking Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0048*	0.0044*	0.0049*	0.0030	0.0057*	0.0011	-0.00047
	(0.0011)	(0.00090)	(0.0016)	(0.0020)	(0.0015)	(0.0010)	(0.00090)
Total Response	0.0040*	0.0037*	0.0051*	0.0037	0.0050*	0.00026	-0.0014
	(0.0011)	(0.00092)	(0.0015)	(0.0021)	(0.0016)	(0.0012)	(0.00081)
From Below	-0.0018*	-0.0018*	-0.0026*	-0.0022	-0.0027*	-0.00018	0.0010
	(0.00084)	(0.00081)	(0.0011)	(0.0013)	(0.00084)	(0.00082)	(0.00085)
From Above	0.0022*	0.0019*	0.0024*	0.0016	0.0023	0.000084	-0.00032
	(0.0010)	(0.00078)	(0.0011)	(0.0017)	(0.0016)	(0.0010)	(0.00075)
Observations	40838	44766	22257	12634	17071	24068	19061
Excluded Bins:							
L	4	4	4	4	4	4	4
R	5	5	5	5	5	5	5
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-.001	-.0014	-.0037	-.0029	-.0013	.0024	.0018
	(.0012)	(.0016)	(.0021)	(.0022)	(.0013)	(.0016)	(.0015)
$A - \sum_r^R \beta_A^r$	.003	.0023	.0027	.001	.0015	-.0012	.00082
	(.002)	(.0013)	(.0023)	(.0025)	(.0023)	(.0015)	(.0017)
Joint p-val.	0.22	0.16	0.16	0.42	0.45	0.32	0.35
Impact:							
ΔHouses at cutoff	88.3	87.6	73.8	29.0	46.4	2.46	-10.1
Specifications:							
Poly. Order	3	3	3	3	3	3	3
Window	20	20	20	20	20	20	20
Timing	Sold Date	Include Announce- ment	6 Months Before and After	3 Months Before and After	Jul.-Dec. 2012 VS. Jul.-Dec. 2011	Jan.-Jun. 2012 VS. Jul.-Dec. 2011	Jul.-Dec. 2012 VS. Jan.-Jun. 2013

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using asking prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

Table E8
Robustness Checks to Alternative Time Choices (Sales Price)

	Sales Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0010*	0.00095*	0.00095	0.000096	0.0011	0.00023	0.000022
	(0.00044)	(0.00042)	(0.00065)	(0.00063)	(0.00066)	(0.00048)	(0.00070)
Total Response	0.00057	0.00035	0.00089	0.00038	0.00024	-0.00053	-0.00054
	(0.00063)	(0.00043)	(0.00064)	(0.00093)	(0.00082)	(0.00050)	(0.00083)
From Below	-0.0000078	0.00013	0.00030	0.000031	0.00011	-0.000060	-0.00043
	(0.00062)	(0.00038)	(0.00084)	(0.0010)	(0.00090)	(0.00048)	(0.0012)
From Above	0.00057	0.00049	0.0012	0.00041	0.00035	-0.00059	-0.00097
	(0.00079)	(0.00047)	(0.0010)	(0.0015)	(0.0014)	(0.00066)	(0.0011)
Observations	40838	44766	22257	12634	17071	24068	19061
Excluded Bins:							
L	1	1	1	1	1	1	1
R	2	2	2	2	2	2	2
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-7.8e-06	.00013	.0003	.000031	.00011	-.00006	-.00043
	(.00062)	(.00038)	(.00084)	(.001)	(.0009)	(.00048)	(.0012)
$A - \sum_r^R \beta_A^r$	-.0003	-.00014	-.00021	.00044	.00044	.00077	.00012
	(.00064)	(.00048)	(.00078)	(.0011)	(.00089)	(.00074)	(.00082)
Joint p-val.	0.90	0.92	0.94	0.88	0.72	0.55	0.93
Impact:							
ΔHouses at cutoff	12.7	8.45	13.0	2.94	2.25	-5.00	-4.00
Specifications:							
Poly. Order	3	3	3	3	3	3	3
Window	20	20	20	20	20	20	20
Timing	Sold Date	Incclude Announcement	6 Months Before and After	3 Months Before and After	Jul.-Dec. 2012 VS. Jul.-Dec. 2011	Jan.-Jun. 2012 VS. Jul.-Dec. 2011	Jul.-Dec. 2012 VS. Jan.-Jun. 2013

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using sales prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

E.4.2 Classifying by Sold Date Instead of Listing Date

Table E9
Assessing Robustness to Alternative Time Choices: Sold Date

	Asking Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0048*	0.0047*	0.0049*	0.0034	0.0062*	0.0015	-0.00086
	(0.0011)	(0.0011)	(0.0015)	(0.0019)	(0.0016)	(0.0012)	(0.0017)
Total Response	0.0040*	0.0041*	0.0054*	0.0038	0.0055*	0.00050	-0.0019
	(0.0011)	(0.0011)	(0.0015)	(0.0020)	(0.0016)	(0.0013)	(0.0017)
From Below	-0.0018*	-0.0017*	-0.0022*	-0.00080	-0.0025*	-0.00048	0.00066
	(0.00074)	(0.00072)	(0.0010)	(0.0014)	(0.0010)	(0.00091)	(0.0011)
From Above	0.0022*	0.0024*	0.0033*	0.0030	0.0030*	0.000021	-0.0012
	(0.00087)	(0.00088)	(0.0012)	(0.0016)	(0.0013)	(0.0011)	(0.0013)
Observations	40838	44545	22128	12467	18280	24175	18675
Excluded Bins:							
L	4	4	4	4	4	4	4
R	5	5	5	5	5	5	5
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-.001	-.001	-.0021	.0012	-.0012	.00091	.00054
	(.0013)	(.0013)	(.0019)	(.0025)	(.0018)	(.0017)	(.0019)
$A - \sum_r^R \beta_A^r$	.003	.0032*	.0046*	.0046	.0038	-.0011	-.0012
	(.0015)	(.0015)	(.0022)	(.0027)	(.0022)	(.0018)	(.0022)
Joint p-val.	0.13	0.091	0.084	0.24	0.22	0.72	0.84
Impact:							
Δ Houses at cutoff	88.3	97.8	75.8	29.8	55.7	5.01	-14.6
Specifications:							
Poly. Order	3	3	3	3	3	3	3
Window	20	20	20	20	20	20	20
Timing	Baseline	Include Announcement	6 Months Before and After	3 Months Before and After	July-Dec. 2012 VS. July-Dec. 2011	Jan.-Jun. 2012 VS. July-Dec. 2011	July.-Dec. 2012 VS. Jan-June 2013

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using asking prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

Table E10

Assessing Robustness to Alternative Time Choices: Sold Date

	Sales Price						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jump at cut-off	0.0010*	0.00099*	0.00057	-0.00026	0.0010	0.00049	0.00014
	(0.00044)	(0.00044)	(0.00059)	(0.00065)	(0.00062)	(0.00047)	(0.00068)
Total Response	0.00057	0.00065	0.00099	0.00098	0.00053	-0.00029	-0.00053
	(0.00056)	(0.00054)	(0.00072)	(0.00086)	(0.00074)	(0.00063)	(0.00079)
From Below	-0.0000078	0.00020	0.00095	0.0017	0.00013	-0.00058	-0.00053
	(0.00059)	(0.00061)	(0.00086)	(0.0011)	(0.00082)	(0.00069)	(0.00088)
From Above	0.00057	0.00085	0.0019	0.0027*	0.00066	-0.00087	-0.0011
	(0.00075)	(0.00076)	(0.0011)	(0.0013)	(0.00099)	(0.00091)	(0.0011)
Observations	40838	44545	22128	12467	18280	24175	18675
Excluded Bins:							
L	1	1	1	1	1	1	1
R	2	2	2	2	2	2	2
Tests of Fit:							
$B - \sum_l^L \beta_B^l$	-7.8e-06	.0002	.00095	.0017	.00013	-.00058	-.00053
	(.00059)	(.00061)	(.00086)	(.0011)	(.00082)	(.00069)	(.00088)
$A - \sum_r^R \beta_A^r$	-.0003	-.00033	-.00099	-.0013	.00026	.0011	.00024
	(.00056)	(.00057)	(.0008)	(.0011)	(.00073)	(.00069)	(.00081)
Joint p-val.	0.85	0.85	0.41	0.28	0.89	0.27	0.83
Impact:							
ΔHouses at cutoff	12.7	15.7	13.8	7.61	5.34	-2.89	-4.17
Specifications:							
Poly. Order	3	3	3	3	3	3	3
Window	20	20	20	20	20	20	20
Timing	Baseline	Include Announcement	6 Months Before and After	3 Months Before and After	July-Dec. 2012 VS. July-Dec. 2011	Jan.-Jun. 2012 VS. July-Dec. 2011	July.-Dec. 2012 VS. Jan-June 2013

Notes: This table displays the bunching estimates of the million dollar policy for the city of Toronto. The dependent variable is $\hat{\Delta}_S(y_j)$ constructed using sales prices. The rows of the table correspond to the components of (8). The first row shows the total jump at the million dollar threshold, the second row shows the total response due to the policy ($\hat{\beta}_A - \hat{\beta}_B$), and the last two rows show the response from above ($\hat{\beta}_A$) and below ($\hat{\beta}_B$) the threshold, respectively. Standard errors, in parentheses, are constructed via bootstrap discussed in the main text. (*) denotes significance at the 5% level.

F Supplemental Material for the Distribution Decomposition

F.1 Robustness of the Price Counterfactual

Figure F1 presents the price counterfactual differences for a variety of alternative specifications. The leftmost plots in panels (a) and (b) of Figure F1 replicate the price counterfactual differences in Figure 9, but with a narrower range on the vertical axis. The adjacent plots display the price counterfactual differences when we restrict the set of transactions to just six months before and after the implementation of the policy. We repeat the exercise again using only three-month pre- and post-policy periods. The differences attributed to market conditions are smaller when we consider shorter pre- and post-policy periods (not shown), but the implied price counterfactual differences remain close to zero and statistically insignificant. To generate the rightmost plots of panels (a) and (b) of Figure F1, labeled "falsification", we perform the same procedure outlined above using only pre-policy transaction data. We use the same two pre-policy years as the placebo tests in Table E3 to create the black line. The superimposed grey lines display the six- and three-month price counterfactual differences again using two pre-policy periods, and the confidence bands correspond to the 6-month sample periods. After accounting for shape-preserving market trends as well any changes in the composition of homes sold, the residual price counterfactual differences are relatively small and not statistically significantly different from zero, as expected. This provides some reassurance that the evolution of a price distribution in a booming housing market can be summarized reasonably well by scale and shift parameters, after controlling for house characteristics.

../Figures-dfl-T2-General/Figure-Price-Cntf-900000-Toronto-CDF.pdf

(a) City of Toronto

../Figures-dfl-T2-General/Figure-Price-Cntf-900000-Central-CDF.pdf

(b) Central District

Figure F1
Price Counterfactual Robustness (CDFs)

Notes: The leftmost plots in panels (a) and (b) replicate the price counterfactual differences in Figure 9. The remaining plots present price counterfactual differences for a variety of alternative specifications. The cut-off for estimating market trends is set to $900K. The shaded area represents a 95% confidence interval, obtained via bootstrap.

F.2 Supplemental Material: Decomposition of PDFs

As an alternative to displaying CDFs and their differences, the main results from Figure 9 can be displayed in terms of PDFs or histograms. Figure F2 plots the PDFs, their differences, and the decomposition for both the city of Toronto and the central district. An extensive margin response would be evident if the price counterfactual differences indicated fewer homes sold at prices above the $1M threshold. An intensive margin response would display fewer homes sold at high prices, but more sales at prices closer to $1M. Neither phenomenon is visually evident in the price counterfactual differences plotted in Figure F2. Figure F3 shows robustness to alternative specifications.

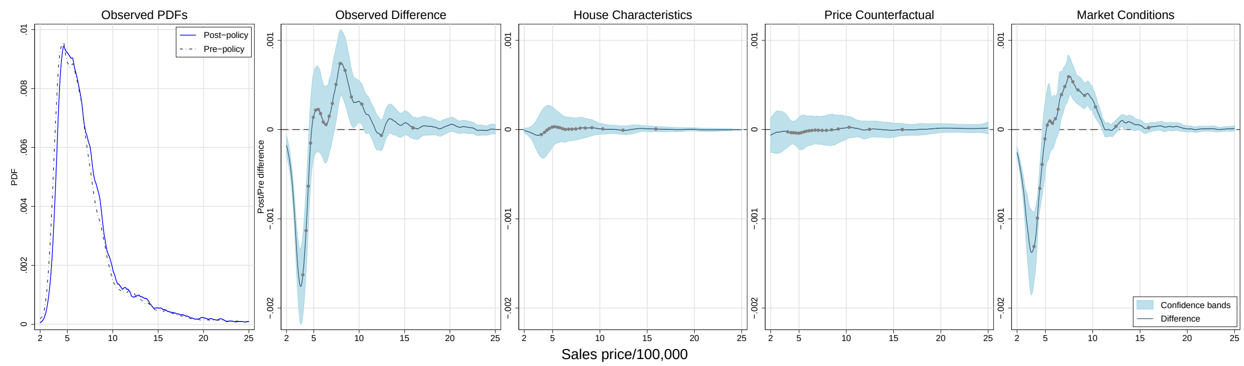

(a) City of Toronto

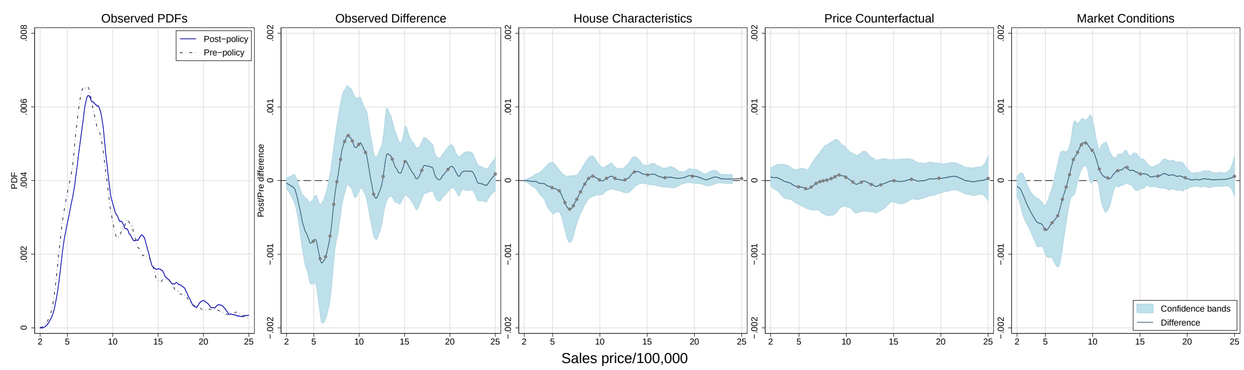

(b) Central District

Figure F2
Examining Policy Responses Above $1M (PDFs)

Notes: Panel (a) of the figure plots the pre- and post-policy sales price PDFs, their differences, and the decomposition based on house characteristics, market conditions, and any residual price counterfactual differences for the city of Toronto. Panel (b) represents the same procedure for the central district. The cut-off for estimating market trends is set to $900K. The shaded area represents a 95% confidence interval, obtained via bootstrap.

../Figures-dfl-T2-General/Figure-Price-Cntf-900000-Toronto-f.pdf

(a) City of Toronto

../Figures-dfl-T2-General/Figure-Price-Cntf-900000-Central-f.pdf

(b) Central District

Figure F3
Price Counterfactual Robustness (PDFs)

Notes: The leftmost plots in panels (a) and (b) replicate the price counterfactual differences in Figure F2. The remaining plots present price counterfactual differences for a variety of alternative specifications. The cut-off for estimating market trends is set to \$900K. The shaded area represents a 95% confidence interval, obtained via bootstrap.

F.3 Distribution Decomposition Simulations

Here, we apply the distribution decomposition methodology to simulated data that feature no policy response, an extensive margin response, and an intensive margin response to further justify assumptions (1) and (2) in Section 5.3.

Two samples of 20,000 log sales prices are generated from the logistic distribution with mean and scale parameters estimated by maximum likelihood to fit the observed sales price distributions. The sample designated "pre-policy" features mean and scale parameters 13.3345 and 0.2489, whereas the "post-policy" distribution sets the mean to 13.3865 and the scale to 0.2434. Figure F4a plots the two CDFs and their differences, which is then decomposed into differences that are attributed to "market conditions" and those that are not (i.e., "price counterfactual" differences).[48] Not surprisingly given the absence of any additional effect(s) beyond the two parameters of the logistic distribution, the price counterfactual differences are everywhere close to zero. For this simulation exercise, the shape-preserving change in the distribution is well-summarized by a linear transformation applied to the pre-policy quantile function. Moreover, the intercept and slope coefficients are estimated reasonably well using only prices below a cut-off of \$900K.

An extensive margin response is then simulated by randomly dropping 20% of prices above \$1M in the post-policy sample. The same procedure is implemented and the results are displayed in Figure F4b. As argued above, the extensive margin response is visually identifiable because the price counterfactual difference diverges above zero near the \$1M threshold. An intensive margin response is then simulated (from the original post-policy sample) by reducing every price above \$1M from p to $\$1\text{M} + 0.7(p - \$1\text{M})$. In other words,

[48]There are no differences attributed to "house characteristics" as the composition of homes is assumed to be unchanged for the purposes of these simulations.

prices in excess of \$1M are lowered by 30% of this excess amount. Figure F4c displays the CDFs and the decomposition of their differences. Once again, the price counterfactual differences diverge above zero, this time slightly above the \$1M threshold. Each of these simulated responses cause the average post-policy house price to drop by only about 2% relative to the original post-policy sample of simulated prices, yet these simulations affirm that evidence of extensive and intensive margin responses can be readily uncovered under the proposed methodology.

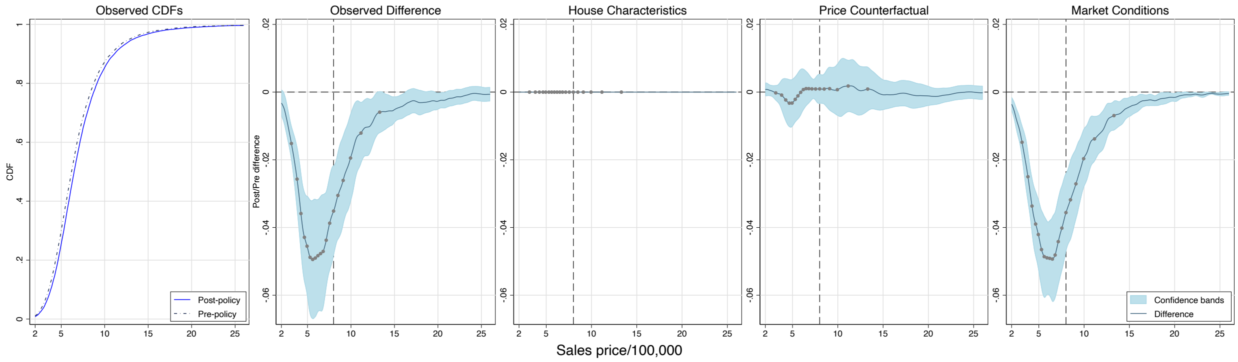

(a) Simulation 1: No Policy Response

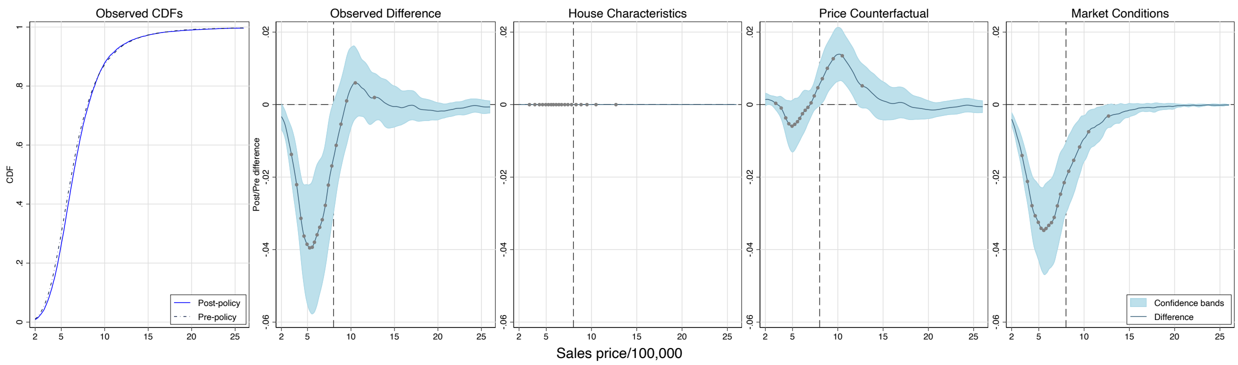

(b) Simulation 2: Extensive Margin Response

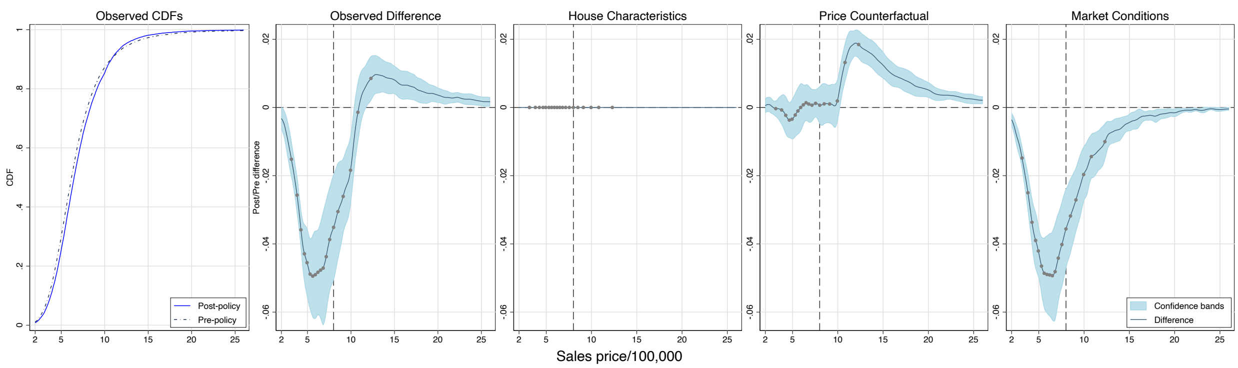

(c) Simulation 3: Intensive Margin Response

Figure F4
Distribution Decomposition Applied to Simulated Data

Notes: This figure displays results from applying our distribution decomposition to simulated data. Panel (a) contains the simulated prices under the assumption that there was neither an extensive nor intensive margin effect. Panel (b) simulates an extensive margin response by dropping 20% of the prices above \$1M. Panel (c) simulates an intensive margin response by reducing every price in excess of \$1M by 30% of this excess amount.

F.4 Robustness to a Lower Cut-off: $800K

As a final robustness check, we consider estimating the intercept and coefficient parameters using only transaction data further below the $1M threshold. More specifically, we lower the cut-off, τ, from $900K to $800K to be sure that the differences in distributions we ascribe to market conditions are not contaminated by unintended consequences of the million dollar policy in price segments further below the policy threshold. Figure F5 in Appendix F.4 displays the results with $\tau = \$800K$. The resulting market conditions and price counterfactual differences are remarkably similar to those in Figure 9, except the confidence bands are larger when the lower cut-off is applied.

../Figures-dfl-T2-houses-GTA2/800000/Figure-Extensive-CDFs-post.pdf

(a) City of Toronto

../Figures-dfl-T2-houses-Central/800000/Figure-Extensive-CDFs-post-central.pdf

(b) Central District

Figure F5
Examining Policy Responses Above $1M: Robustness to a Lower Cut-off

Notes: Panel (a) of the figure plots the pre- and post-policy sales price distributions, their differences, and the decomposition based on house characteristics, market conditions, and any residual price counterfactual differences for the city of Toronto. Panel (b) represents the same procedure for the central district. The cut-off for estimating market trends is set to $800K. The shaded area represents a 95% confidence interval, obtained via bootstrap.

G Difference-in-Differences Results

In this appendix, we consider a difference-in-differences (diff-in-diff) research design that (i) exploits market segmentation by geography to cleanly identify policy effects and (ii) leverages the time dimension of the analysis to assess the parallel pre-trends assumption. We employ two new left-hand-side variables (namely, sales counts and the log first-difference in quarterly house price indices) for which the parallel pre-trends assumption appears to hold. Our diff-in-diff approach builds on a recent literature that exploits ex ante treatment intensity across geographies to study policies implemented at the national level (Mian and Sufi 2012; Pierce and Schott 2016; Berger, Turner, and Zwick 2020). A natural extension of this literature to the \$1M policy yields a continuous treatment intensity proxy based on the share of homes sold over \$1M two years prior to the policy by FSA (three-digit zip code). With this treatment variable in hand, we then estimate a diff-in-diff model using data from central Toronto only, and again for all of Toronto for robustness.[49] The diff-in-diff regression of interest becomes:

$$z_{kt} = \sum_{y \neq 2012Q2} \beta_t \times 1\{y = t\} \times Intensity_k + \sum_{y \neq 2012Q2} \eta_t \times 1\{y = t\} \times \mathbf{X}'_k + \tau_k + \tau_t + \epsilon_{kt}, \quad \text{(G.1)}$$

where z_t is the LHS variable aggregated to the FSA level for each FSA (indexed by k) at time t. The ex ante treatment intensity variable, $Intensity_k$, represents the share of homes sold over \$1M in FSA k two years before the policy. The parameter of interest is the intensity coefficient, β_t, which compares z_t across FSAs that differ in treatment intensity. In other words, it measures the difference in z_t between high- and low-intensity FSAs (first difference), relative to this same difference evaluated in 2012Q2, the quarter before the implementation of the \$1M policy (second difference). To ease the interpretation of the regression estimates, we scale $Intensity_k$ by its standard deviation. With this scaling, β_t measures the change in z_t due to a one standard deviation increase in $Intensity_k$, relative to the change in z in 2012Q2 also associated with a one standard deviation increase in $Intensity_k$.

FSA level controls, $\mathbf{X}_k$, consist of the average lot size and its square measured two years before the policy by FSA. We interact $\mathbf{X}_k$ with time fixed effects so that the impact of average property size on the LHS variable can vary with time. Finally, τ_k and τ_t are FSA and time fixed effects, and robust standard errors are clustered at the FSA level.

The first LHS variable we consider is the number of homes sold (sales counts) by year-quarter. When the LHS variable is sales counts, the identifying assumption is that the trend in sales counts in FSAs with low intensity (fewer ex ante \$1M sales) can provide the counterfactual for the trend in sales counts for high intensity FSAs. As the excluded dummy in equation (G.1) corresponds to 2012Q2, the diff-in-diff parallel pre-trends assumption main-

[49] Homes in the central district constitute the bulk of \$1M homes and are concentrated geographically, which lends itself well to satisfying the diff-in-diff parallel pre-trends assumption. Comparisons of homes across central and suburban Toronto are less appealing within this diff-in-diff framework given the natural urban-suburban differences and the relative scarcity of \$1M homes in suburban Toronto. Be that as it may, our results are similar if we use all of Toronto, as shown below.

tains that the change in sales counts in each year-quarter during the pre-treatment period associated with a change in $Intensity_k$ is not statistically different from this same change evaluated in 2012Q2, the quarter before policy implementation.

Intuitively, if the \$1M policy tempered activity in high-price segments of the market (i.e., an extensive margin response), then we would expect the number of sales in these more expensive FSAs to fall. Figure G1, panel A plots β_t when the LHS variable is sales counts by FSA. Prior to the implementation of the \$1M policy (left of the black-dashed vertical line), there was nearly no difference in the number of homes sold across high- and low-intensity FSAs relative to the intensity coefficient evaluated in 2012Q2, congruent with the diff-in-diff parallel pre-trends assumption. Indeed, the diff-in-diff estimates to the left of the black-dashed vertical line indicate that a one standard deviation increase in $Intensity_k$ corresponds with just an extra 1.57 homes sold on average by year-quarter during the pre-treatment period, relative to the impact of $Intensity_k$ on sales counts in 2012Q2. These pre-treatment diff-in-diff estimates are also all statistically insignificant, where the largest cluster-robust t-statistic associated with these estimates is just 1.29.

Following the implementation of the million dollar policy, there was likewise no statistically significant change in the path of year-quarter sales counts associated with an increase in $Intensity_k$, indicating that the policy did not impact the number of homes sold in high intensity FSAs. In the first quarter following policy implementation (2012Q3), a one standard deviation increase in $Intensity_k$ corresponds with just 4.28 extra homes sold (cluster-robust S.E. = 4.38; t-statistic = 0.98), relative to impact of a one standard deviation increase in $Intensity_k$ in 2012Q2. The remaining point estimates for 2012Q4–2013Q1 are also small and not statistically significant, indicating that the million dollar policy did not impact sales counts across high- and low-intensity FSAs in these later periods.

Next, in panel B of Figure G1, we let the LHS variable be the log-difference in quarterly FSA hedonic house price indices. One of the objectives of the million dollar policy was to curb house price appreciation in high-price segments of the market. If successful in achieving this objective (i.e., an intensive margin response), the million dollar policy would translate into subdued house price growth in FSAs with higher $Intensity_k$. To examine house price growth within our diff-in-diff framework, we first estimate a hedonic regression to adjust FSA year-quarter house prices for property-level characteristics using transaction level data. Then we take the log first-difference of each FSA house price index to use as the LHS variable in equation (G.1). The results in panel B of Figure G1 first indicate that, during the pre-treatment period (left of the black-dashed vertical line), increases in ex ante treatment intensity are uncorrelated with pre-treatment house price growth, after accounting for the relationship between ex ante treatment intensity and house price growth in 2012Q2, consistent with the diff-in-diff parallel pre-trends assumption. Moreover, there is likewise no statistically significant difference in the trend of house price growth between FSAs with differential intensity during the treatment period. Hence, the million dollar policy does not correspond with a reduction in house price growth following the implementation of the policy in neighborhoods rich with \$1M homes.

Finally, Figure G2 displays similar results when we include all of Toronto.

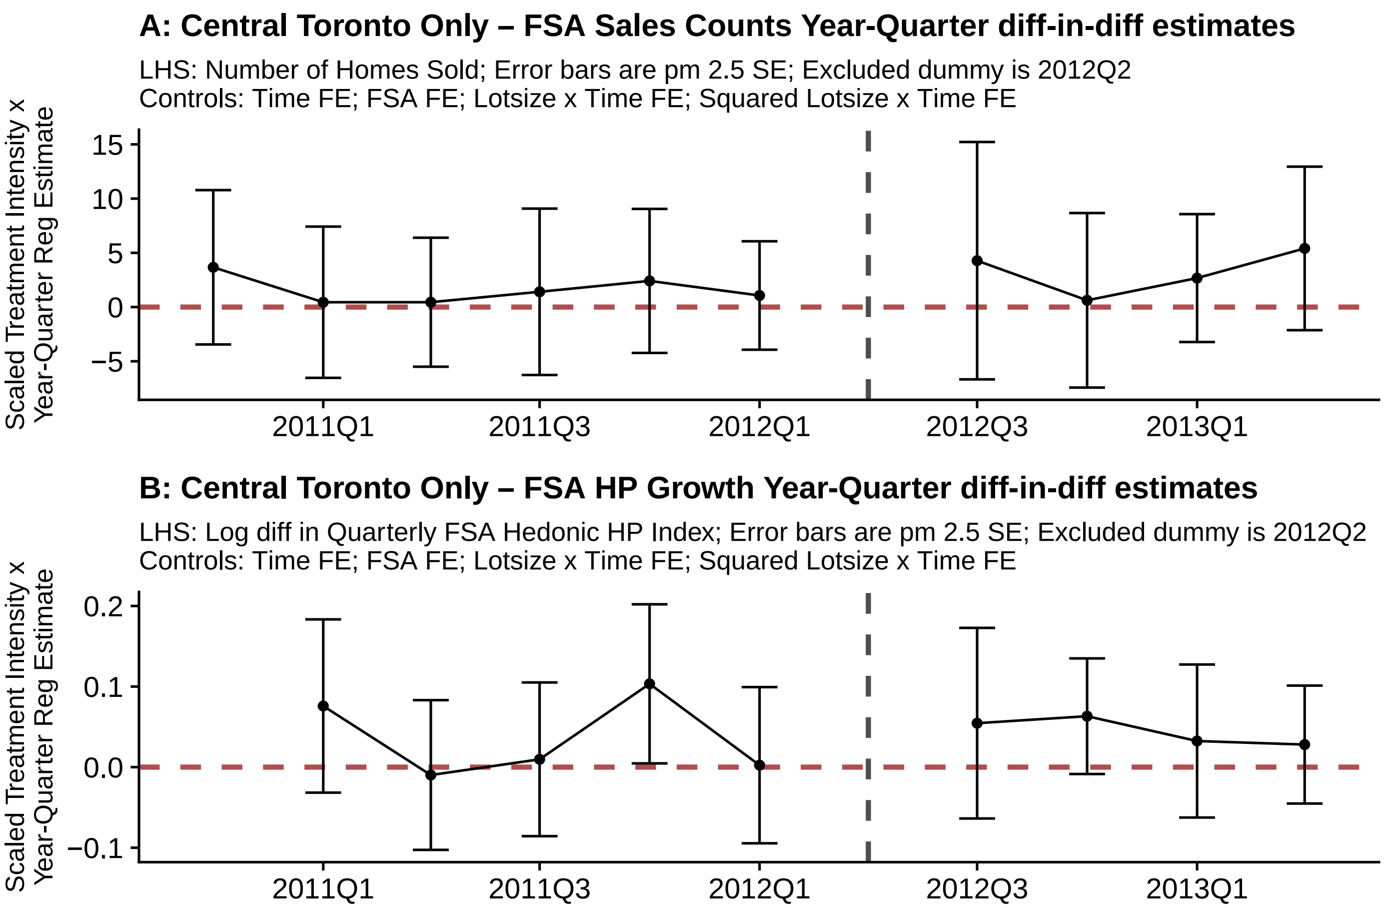

Figure G1
Central Toronto Diff-in-Diff Estimates

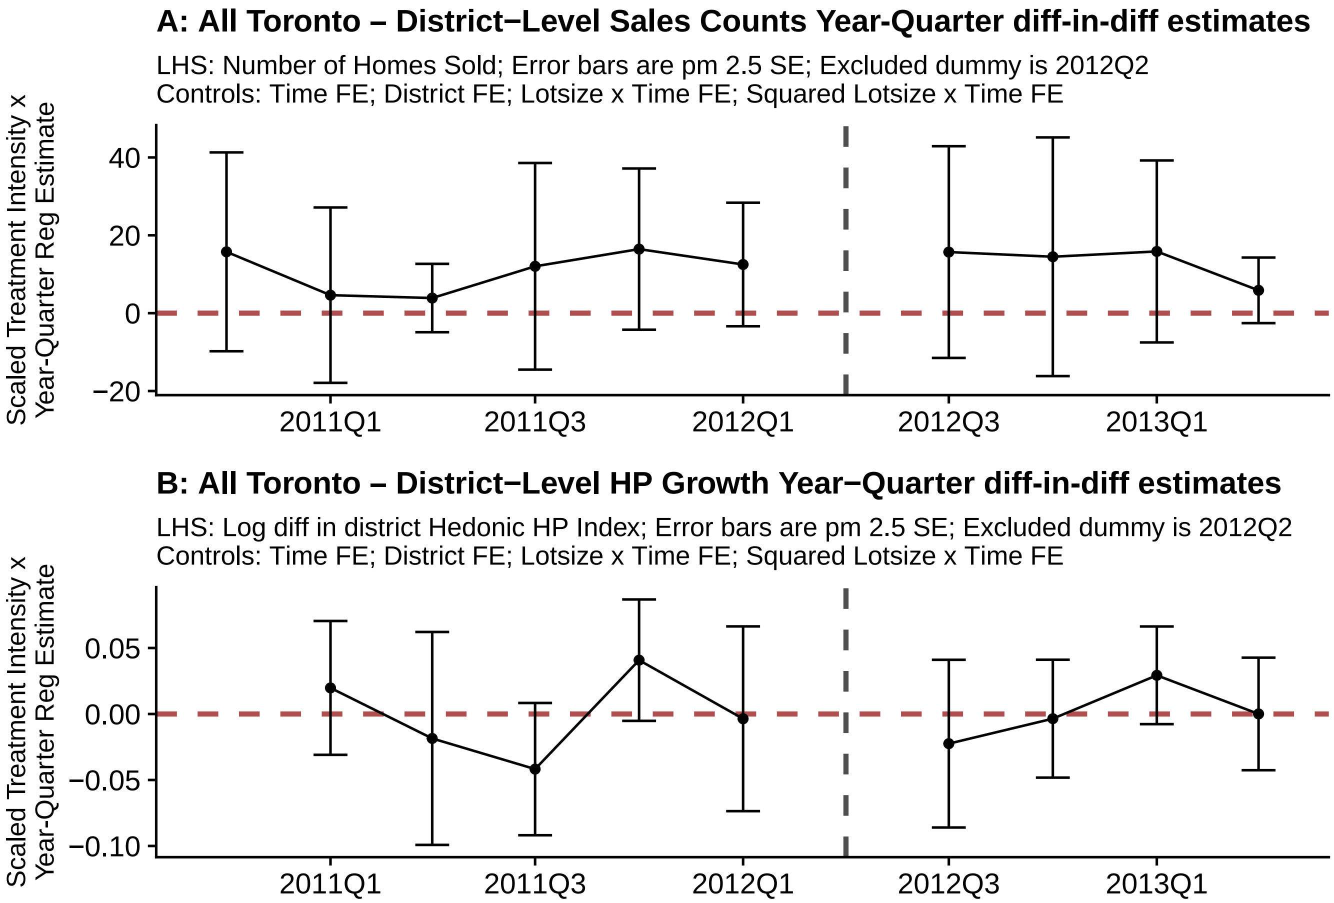

Figure G2
All Toronto Diff-in-Diff Estimates