

Backfiring in Bad Times: When Rent Control Keeps Rent Too High

Michael Boutros

Geneviève Vallé*

April 2024

Abstract

Rent control, intended to benefit renters by capping rent increases, may disincentivize landlords from lowering rents during temporary negative demand shocks because they are unable to quickly increase rent afterward. To test this prediction, I use a unique combination of exogenous variation in rent control policy in Toronto and a negative demand shock induced by the COVID-19 pandemic. In line with theory, rent per square foot decreased by 1.7% for rent controlled units and 4.7% for exempt units. Using a model of differentiated demand, I construct a counterfactual exercise and estimate that in the absence of rent control, rent would have decreased by 8.3% for rent-controlled units and 8.1% for exempt units.

*Boutros: Financial Stability Department, Bank of Canada (michael.boutros@bankofcanada.ca). Vallé: Financial Stability Department, Bank of Canada (gvallee@bankofcanada.ca). The presented views are those of the authors, not necessarily of the Bank of Canada. Version 2024.4.1.

1 Introduction

Housing affordability is at the forefront of policy discussions, with much of the focus on real estate prices and their relationship to monetary policy and business cycles. Policies intended to increase affordability through mortgage forgiveness or subsidizing home ownership are important tools for affordability. However, many households rent their residences, and the majority of renters who are low-income may never become homeowners. For these lifelong renters, shelter is their primary expense, and fluctuations in housing costs have first-order impacts on welfare.

Rent control keeps rent stable and predictable. Without rent control, tenants are exposed to uninsurable increases in rent that might force them to relocate. The benefits of rent control as insurance accrue to tenants who have made the largest investments in location-specific capital, such as social networks or proximity to school or work, and therefore face the largest non-pecuniary costs to relocating in search of lower rent. However, a large literature documents the inefficiencies of these contracts for tenants, landlords, and neighbourhoods (Diamond, McQuade and Qian, 2019). Advocates of rent control conclude that the benefits of rent control for incumbent tenants outweigh the inefficiencies imposed on others. In this paper, I demonstrate an important caveat in rent control as insurance: although designed to keep rent stable when it would otherwise increase, rent control also keeps rent stable when it would otherwise decrease. This cost is borne completely by tenants, both incumbents and movers, and is especially high for those with local capital, the primary group which rent control is intended to benefit.

I document this caveat using a unique natural experiment which unfolded during the COVID-19 pandemic in Toronto's rental market. This environment generated an ideal natural experiment for two reasons. First, an unexpected policy change removed rent control for any unit first occupied after November 15, 2018. This policy was implemented and announced by Premier Doug Ford in November 2018, four months after his election in June

2018, and five months after a campaign promise in May 2018 in which Ford stated “I have listened to the people, and I won’t take rent control away from anyone. Period.” Nearly identical units on either side of this cutoff are subject to complete or no rent control, and I exploit this variation to isolate the impact of rent control on the change in rent. Second, the COVID-19 pandemic is an exogenous shock to the demand for rentals in Toronto. Faced with the prospect of work-from-home for an extended period of time, tenants in major metropolitan areas fled their small urban apartments in search of larger spaces, while the supply of rental units largely remained fixed, especially in the primary market of purpose-built rentals.

Overall, I find that rent controlled units were more expensive than similar flexible units during the COVID-19 pandemic. Using proprietary data collected by Urbanation, a third-party market research firm, I estimate that rent per square foot decreased by 1.7% for units subject to rent control, and 4.7% for units exempt from rent control. This estimate is obtained by comparing units occupied for the first time for three months before the policy cutoff date (August to October 2018) to units occupied for the first time three months after the policy cutoff date (December 2018 to February 2019). I demonstrate that units occupied for the first time during these six months and subject to starkly different rent control policies are otherwise similar along observables such as average unit size, average building size, and geographic dispersion.

To understand the economic forces at play, I develop a model of Hotelling demand for differentiated rental units in two sectors. In the Flexible sector, landlords can change rent in each period, while landlords in the Controlled sector can only reset rent when a new tenant arrives. The model’s key prediction is that in response to a negative demand shock and all else equal, rent falls by more in the Flexible sector than the Controlled sector. This occurs because landlords in the Controlled sector understand that while they may freely decrease rent to the efficient level in the current period, they will be unable to fully restore rent to the pre-shock level in the future due to the presence of rent control. For example, in many North American cities, rent control limits increases in rent to around 2%, in line

with historical inflation. If the response to a negative demand shock dictates lowering rent by 10%, then it would then it would take roughly five years of 2% increases to return to the initial level. A similar logic applies to decreasing rent for existing tenants who would prefer temporary price relief as opposed to relocation.

On the other hand, landlords in the Flexible sector can freely decrease rent to the efficient level because they can freely increase rent back to the pre-shock level once the shock dissipates. Therefore, in response to the same negative shock, rent decreases by more for Flexible units than Controlled units. However, a key insight of the model is that despite not being directly subject to rent control, units in the Flexible sector are impacted by the presence of Controlled units. Specifically, units in both the Controlled and Flexible sector are priced differently than units in the counterfactual world in which there is no rent control. Since units in the two sectors are partial substitutes, price competition implies that the Flexible sector inherits a degree of rent control from the Controlled sector, and landlords in the Flexible sector respond to this distortion and also set prices differently than they would in the absence of rent control.

This strategic pricing interaction is why the empirical analysis stops short of estimating the true counterfactual of interest: the degree to which rent would have fallen in the complete absence of rent control. In the presence of competition, the simple difference between rent in the flexible and controlled sectors does not identify the impact of rent control. Instead, using the model, I construct the correct theoretical counterfactual that removes rent control considerations for both sectors. The main result is that rent would have decreased by 8.3% for controlled units and 8.1% for exempt units. The model implies that rent control kept rent significantly higher than it otherwise would have been, both for units directly subject to rent control and for units indirectly inheriting rent control via competition.

I contribute to the broader literature studying the intended and unintended consequences of policies surrounding housing. Although policymakers may be well-intentioned, the little available academic evidence on rent control suggests that it is pervasive and likely coun-

terproductive. Nominal rent rigidity is a well-documented phenomena in the United States and around the world (Genesove, 2003; Suzuki, Asami and Shimizu, 2021). Diamond et al. (2019) study the long-term effects of rent control expansion in San Francisco and find that while rent control benefited tenants in the short run, the lost rental housing supply induced by rent control drove up rent in the long-run. Relatedly, I find that rent control prevents rent from decreasing in response to negative demand shocks, which does not benefit tenants. Autor, Palmer and Pathak (2014) find that unexpected rent de-control in Cambridge increased property prices of both controlled and nearby never-controlled units. My model and empirical analysis further demonstrate the impact of rent control on non-controlled units due to strategic interactions.

This paper also contributes to the literature studying mobility and urban flight during the global pandemic. Glaeser, Gorbach and Redding (2020) and Coven, Gupta and Yao (Forthcoming) study urban flight within the United States and document the mass migration from typically popular urban areas to smaller communities with less social amenities. Building on their work, the focus of this paper is on the change in rent from urban flight. Most closely related to this paper is work in Gupta, Mittal, Peeters and Nieuwerburgh (2022) that documents changes in house prices and rent in major US urban areas. This is consistent with the evidence in this paper regarding the decrease in rent in Toronto and, to an extent, surrounding areas. While their paper focuses on the spread between rent in different areas, the focus of this paper is studying the impact of rent control on changes in rent in Toronto.

In the next section, I develop a similar framework of demand for rental units. I then describe the data used in this paper and provide an overview of the rental market in Toronto. I present the impact of the pandemic-induced negative demand shock on rent and map the observable moments from the data to counterfactuals. I estimate the change in rent absent price controls and examine alternative policies that keep the spirit of rent control but allow for additional flexibility in response to negative demand shocks.

2 Background: Rental Market in Ontario

In this section, I describe in detail the rental market in Ontario before and during the pandemic. This is the empirical setting within which I test the impact of rent control during times of economic downturn. For cross-sectional variation between flexible and rent-controlled units, Toronto is an ideal market because of high-quality data on rentals and an exogenous policy change which segmented nearly identical units into flexible and rent-controlled sectors. The COVID-19 pandemic induced a large and unexpected negative demand shock for rentals, which I exploit to study how rent changed in each sector. Later, I use these estimates to construct counterfactual analyses, guided by the model.

2.1 Data

For aggregate statistics on the rental market in Ontario, I use publicly available data from the Canada Mortgage and Housing Corporation (CMHC). I rely on their Rental Market Survey for data on the primary rental market and their Secondary Rental Market Survey for data on the secondary market. The primary rental market consists of purpose-built apartments that are rented by the building owner to tenants. The secondary market consists of units owned by individuals or individual investors and rented to tenants. These surveys provide data for Canada and many subregions, and I focus my analysis on Toronto. The Primary Market Survey contains yearly data on number of units from 1990 to 2021, and the Secondary Market Survey contains similar data from 2007 to 2021.

The main analysis uses rich microdata from Urbanation, a real estate data company that collects information on high-rise real estate projects from conception to occupation. Urbanation's data team is on-the-ground in Toronto and tracks the announcement of new developments. Their surveyors follow these projects through the pre-construction, construction, and occupation phases, and then continue surveying these buildings to track sales and rentals. Urbanation covers both purpose-built rentals in the primary market and condominiums in

Table 1: Summary Statistics for Key Variables (2004Q1 – 2019Q4)

	Occupancy Year	Total Units	Rent	Rent (psf.)	Size (sf.)
Mean	2008.7	348.0	2,250.44	3.00	781.34
Standard Deviation	8.4	155.1	787.98	0.74	273.82
Minimum	1970	8	1,098	.86	298
25 th Percentile	2006	234	1,789	2.48	601
Median	2011	335	2,100	2.95	720
75 th Percentile	2015	440	2,500	3.47	902
Maximum	2020	994	30,000	12.46	7,010
Observations	54,263	54,263	54,263	54,251	54,262

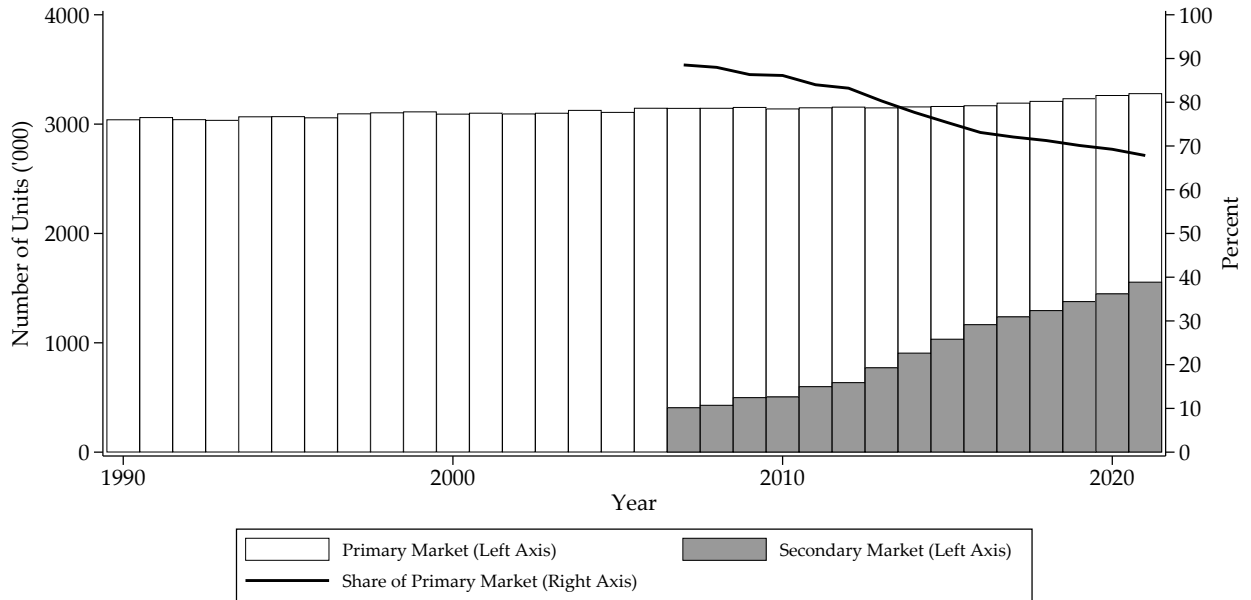
Notes: Rent (psf.) is rent per square foot and size is in square feet.

the secondary market. For each building, Urbanation collects information on location, construction date, occupancy date, and other characteristics such as total suite count. Each quarter, Urbanation records rent and rent per square foot for each unit type in each building, ranging from studio apartments (zero bedrooms) to three or more bedrooms.

2.2 The Rental Market in the Greater Toronto Area

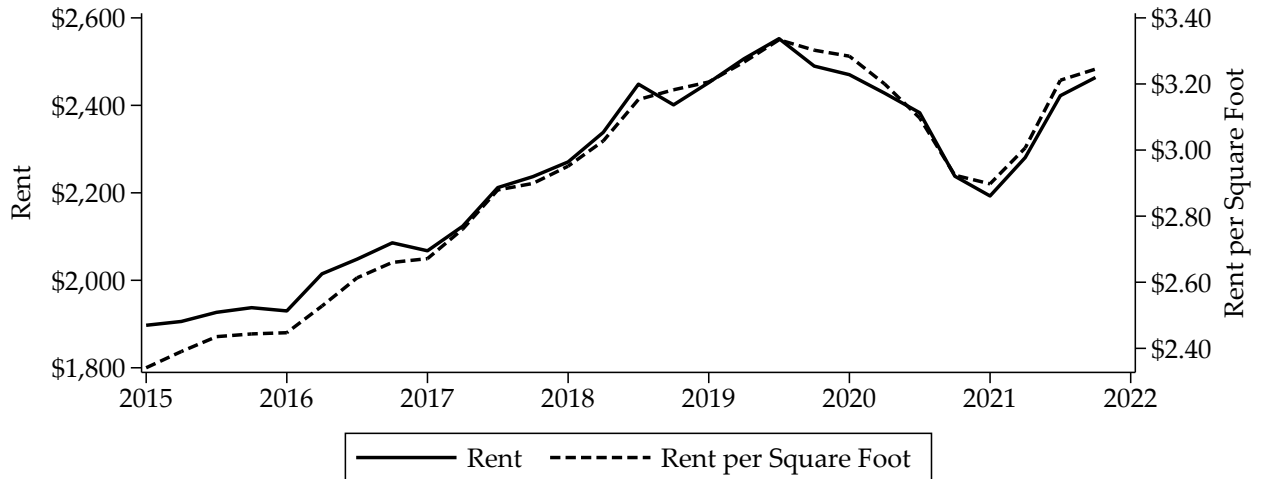
Table 1 contains summary statistics for the key variables in the sample leading up to the pandemic in 2020Q1. On average, buildings in the sample were occupied for the first time in 2009 and have approximately 350 units with an average size of 781 square feet. Average rent prior to the pandemic is roughly \$2,250 per month or \$3.00 per square foot. Figure 1 shows that the rental market in Toronto is dominated by the primary market, but the secondary market has been steadily increasing since data became available in 2007. In that year, the primary market was around 90% of the market, but that share has decreased to just under 70% in 2020. The Urbanation data is available for both the primary and secondary markets. Focusing on rent and rent per square foot, Figure 2 plots these two series from 2015Q1 to 2021Q4. Rent and rent per square foot rise in tandem until 2019Q4, and there is a clear decrease in both series beginning with the pandemic in 2020Q1.

Figure 1: Size and Share of Primary and Secondary Rental Markets



Notes: Number of units in rental market from primary market (e.g., purpose-built rentals) and secondary market (i.e., owner-rented condominiums). Data for secondary market available from 2007 onward. Shares calculated using only primary and secondary market.

Figure 2: Rent and Rent per Square Foot in Toronto



2.3 Exogenous Variation in Rent Control: 2018 Policy Change

In Ontario, there are two key elements to changes in rent: the frequency and the amount. By law, rent for the same unit can change for one of three reasons: a new tenant arrives, the landlord and tenant agree to new capital work and a corresponding change in rent, or the landlord raises rent after 12 months have passed since the last change in rent.

In the third scenario, the Government of Ontario mandates the specific amount by which a landlord can increase rent. The government maintains a website which presents the policy in a way easily digestible by both landlords and tenants. Each year, the government publishes the rent increase guideline, which is based on the Ontario Consumer Price Index and defined as “the maximum a landlord can increase most tenants’ rent during a year without the approval of the Landlord and Tenant Board.” This website also lists a clear exception to the general rent increase guideline: “new buildings, additions to existing buildings and most new basement apartments that are occupied for the first time for residential purposes after November 15, 2018 are exempt from rent control.” This exception is the product of new rent control exemption laws introduced by the Government of Ontario in November 2018. The law created a segmented market in which similar units occupied for the first time on either side of the cutoff date were subject to polar opposite rent controls.

To demonstrate how units are classified as subject to or exempt from rent control, Figure 3 illustrates five different scenarios. All five units are built in 2016 and a marker indicates that the unit was occupied in a given period. Unit 1 is occupied for the first time in 2017 and consistently through 2021. This unit is subject to rent control. Unit 2 is occupied for the first time in 2017, vacant in 2018, and reoccupied by a tenant in 2019. Since the first occupancy is before the policy cutoff, even the new tenancy is subject to rent control. Units 3, 4, and 5 are not subject to rent control. Unit 3 is occupied for the first time in 2019 and continuously through 2021. Unit 4 is occupied for the first time in 2019, vacant in 2020, and reoccupied in 2021. Unit 5 is occupied for the first time in 2021. In all three cases, the first occupancy was

Table 2: Summary Statistics for Flexible and Controlled Units (2004Q1 – 2019Q4)

	Occupancy Year		Total Units	
	Flexible	Controlled	Flexible	Controlled
Mean	2018.8	2008.4	383.6	347.1
Standard Deviation	0.5	8.4	161.3	154.8
Minimum	2018	1970	16	8
25 th Percentile	2019	2005	263	234
Median	2019	2011	363	334
75 th Percentile	2019	2014	506	438
Maximum	2020	2018	697	994
Observations	1,169	53,094	1,169	53,094

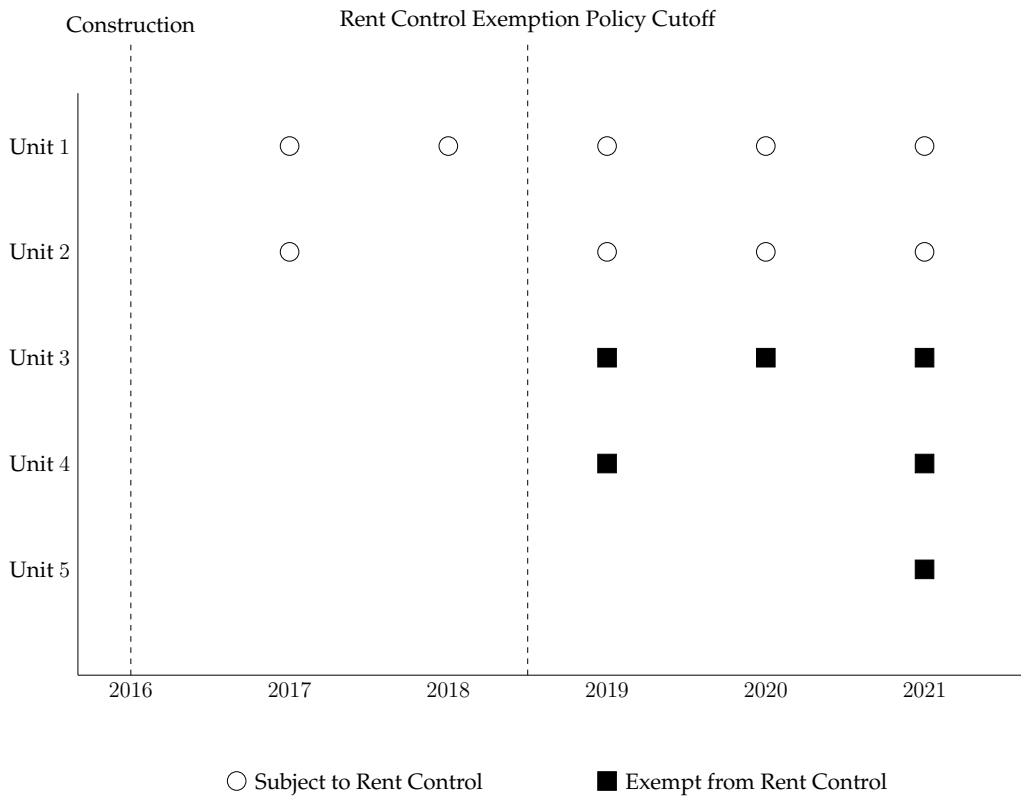
	Rent		Rent per Square Foot		Size (Square Feet)	
	Flexible	Controlled	Flexible	Controlled	Flexible	Controlled
Mean	2563.8	2242.7	3.66	2.98	719.7	782.9
Standard Deviation	856.1	784.7	0.66	.74	251.1	274.2
Minimum	1,098	1,202	1.35	0.86	298	452
25 th Percentile	2,050	1,777	3.21	2.47	563	603
Median	2,332	2,100	3.72	2.94	657	721
75 th Percentile	2,800	2,500	4.08	3.45	838	904
Maximum	11,000	30,000	5.97	12.46	2,551	7,010
Observations	1,169	53,094	1,169	53,082	1,169	53,093

Notes: Statistics for key variables in primary and secondary markets from 2004Q1 to 2019Q4.

after the policy cutoff, so all tenancies in the units are not subject to rent control.

Table 2 presents summary statistics of key variables for units partitioned by rent control exemption status. I refer to units that are subject to rent control laws as being part of the Controlled sector, while units exempt from rent control belong to the Flexible sector. By construction, buildings in the Flexible sector are more recently occupied. Rent for these units is, on average, higher, both overall and per square foot, despite the fact that these units are smaller on average.

Figure 3: Illustration of Rent Control Exemption Eligibility



Notes: Simplified illustration of rent control exemption status after 2018 policy change in Ontario. Illustration of five units that were built in 2016. Policy cutoff is between 2018 and 2019. Units 1 and 2 were occupied for the first time prior to the cutoff, and are therefore not exempt from rent control. Units 3, 4, and 5 were occupied for the first time after the cutoff, and are therefore exempt from rent control.

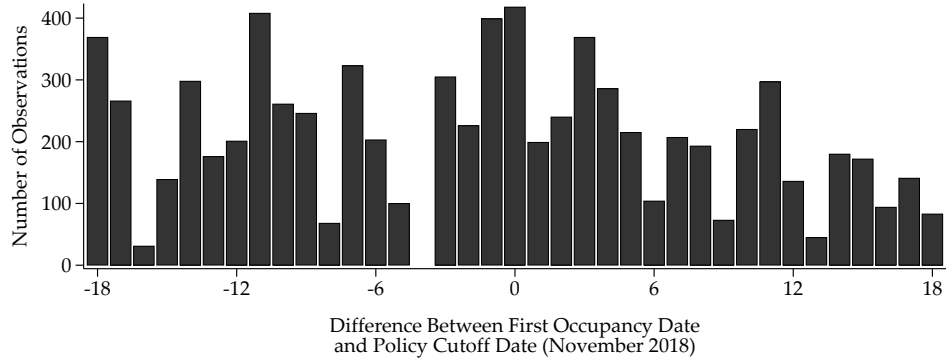
2.3.1 Analysis of Units Just Before and Just After Policy Change

To ensure that the policy change was plausibly exogenous and did not have an immediate impact on the supply of rental units, I analyze key properties of units that were first occupied just before and just after the policy. Figure 4a plots the distribution of observations by number of months from the policy date and demonstrates the ample variation in initial occupancy date, both just before and just after the policy cutoff date in November 2018.

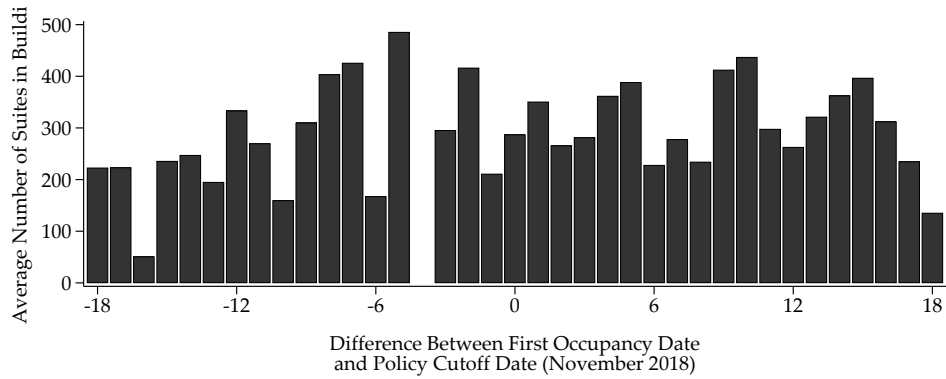
Figures 4b and 4c plot the average number of units and average unit size for all buildings by number of months from the policy cutoff date, and there is no obvious difference between observations from $[-18, -1]$ and $[0, 18]$. To verify this, I perform a balance analysis by regressing an indicator that is equal to 1 for buildings subject to rent control and zero otherwise (i.e., if occupied for the first time after the policy cutoff date) on four covariates: average unit size, average number of units in building, an indicator if the building is located in the core Toronto area, and an indicator if the building is part of the primary rental market.

Table 3 contains estimates from four models with increasing bandwidths around the policy cutoff dates. In the first column, I include all buildings first occupied from August 2018, three months before the policy was introduced, to February 2019, three months after the policy. There is no statistically significant difference between average unit size, average number of suites per building, whether the building is in the Toronto core, and whether the building is in the primary market, and the estimated coefficients are small in magnitude. The same is true when I increase the interval to six months before and after the policy date. As expected, the estimates grow in magnitude and some become marginally significant as the interval increases to one year or 18 months around the policy date, since controlled and flexible buildings are likely different along many other dimensions if their initial occupancy date is also very different.

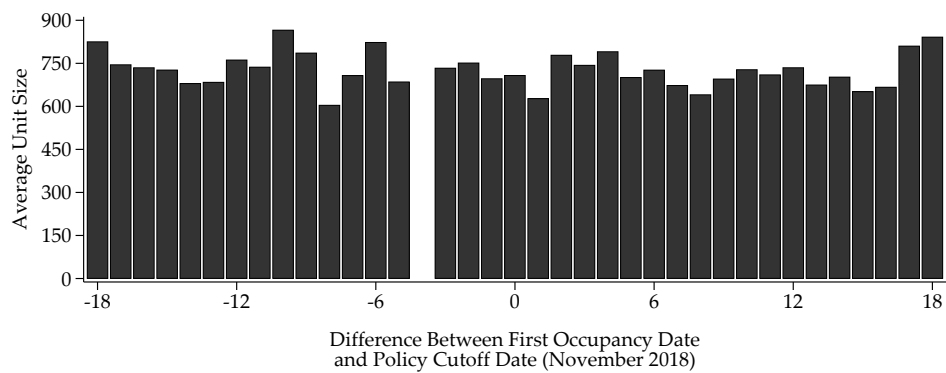
Figure 4: Covariates by Distance from Policy Cutoff Date



(a) Number of Observations



(b) Average Number of Units in Buildings



(c) Average Unit Size

Notes: ...

Table 3: Regression Analysis of Differences in Building Properties by Distance from Policy Cutoff Date

	(1)	(2)	(3)	(4)
	[-3,3]	[-6,6]	[-12,12]	[-18,18]
Size ('00 square feet)	0.004 (0.013)	0.003 (0.011)	0.011 (0.007)	0.010 (0.006)
Suite Count ('00)	0.004 (0.025)	-0.015 (0.021)	0.003 (0.014)	-0.024** (0.011)
Ind.: Toronto Core	-0.025 (0.074)	-0.033 (0.065)	-0.117** (0.048)	-0.058 (0.041)
Ind.: Primary Market	-0.114 (0.092)	-0.104 (0.081)	-0.005 (0.057)	-0.054 (0.052)
<i>N</i>	2,156	3,064	5,696	7,690

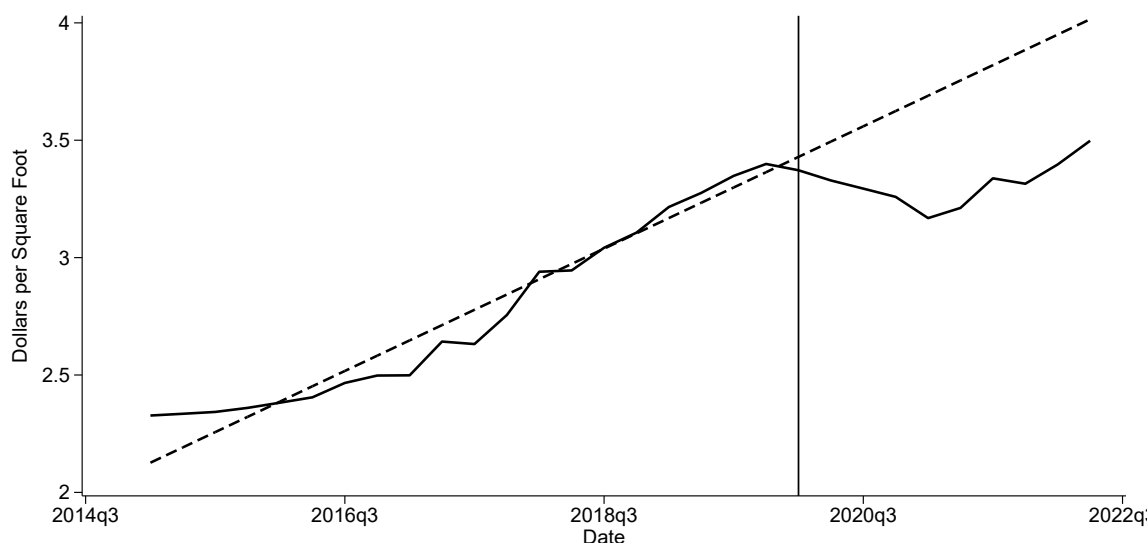
Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2.4 Exogenous Demand Shock: COVID-19 and the Market for Rentals

Figure 5 plots average rent per square foot in the primary market and a linear trendline estimated using data prior to the COVID-19 pandemic in 2020Q1, marked by the vertical black line. The pandemic induced an immediate and large decline in rent per square foot. Although rents stabilized in 2021 and began increasing in 2022, they remain well below their pre-pandemic trend level.

Figure 6 plots the number of new units that were completed in each quarter from 2010Q1 to 2022Q2. The trendline is constructed using all observations prior to the beginning of the pandemic in 2020Q1. Even as the pandemic unfolds, there is no marked increase nor decrease in housing construction for either market. While developers may have anticipated the pandemic would induce structural changes in the employment sector, especially with respect to work-from-home, the design and approval process for new buildings is a timely process which is outside the immediate horizon considered in this paper. On the other hand,

Figure 5: Illustration of Rent per Square Foot and Trendline



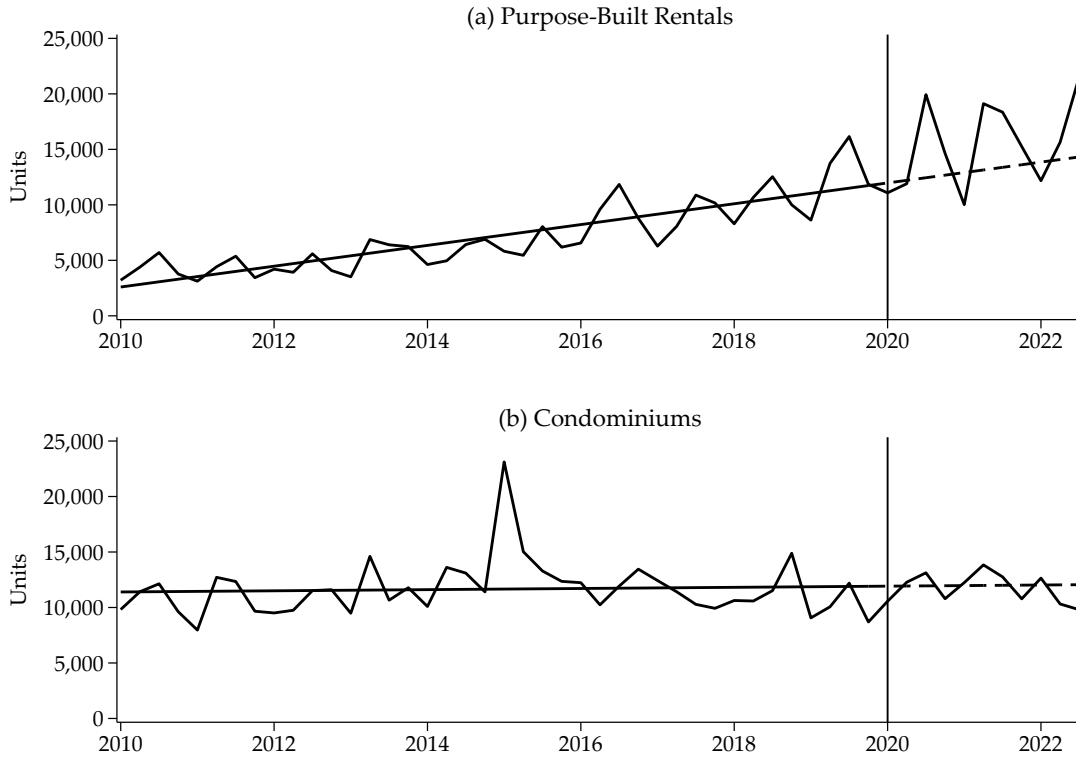
Notes: Solid line is observed rent per square foot. Dashed line is trend calculated using observations prior to the pandemic beginning in 2020Q1 (vertical line).

the absence of a decrease can be explained by Ontario’s designation of construction as an essential sector. Construction projects, especially those which began prior to the pandemic, continued operating mostly as usual during the pandemic. Altogether, these figures suggest that the change in rent was due primarily to a large change in demand, not supply.

3 The Impact of Rent Control on Rent During the Pandemic

In this section, I estimate the impact of the negative demand shock induced by COVID-19 on controlled and flexible units in Toronto. Figure 7 plots average rent per square foot in the Flexible and Controlled Sectors. The Flexible sector contains all buildings with an occupancy date after November 2018, and the Controlled sector contains all other buildings. Given these definitions, data in the Flexible sector begins in 2018Q4. The vertical black line marks the beginning of the COVID-19 pandemic in 2020Q1. In both sectors, there is a clear decrease in rent at the beginning of the pandemic, followed by a flattening and recovery in

Figure 6: New Unit Completions for Purpose-Built Rentals and Condominiums



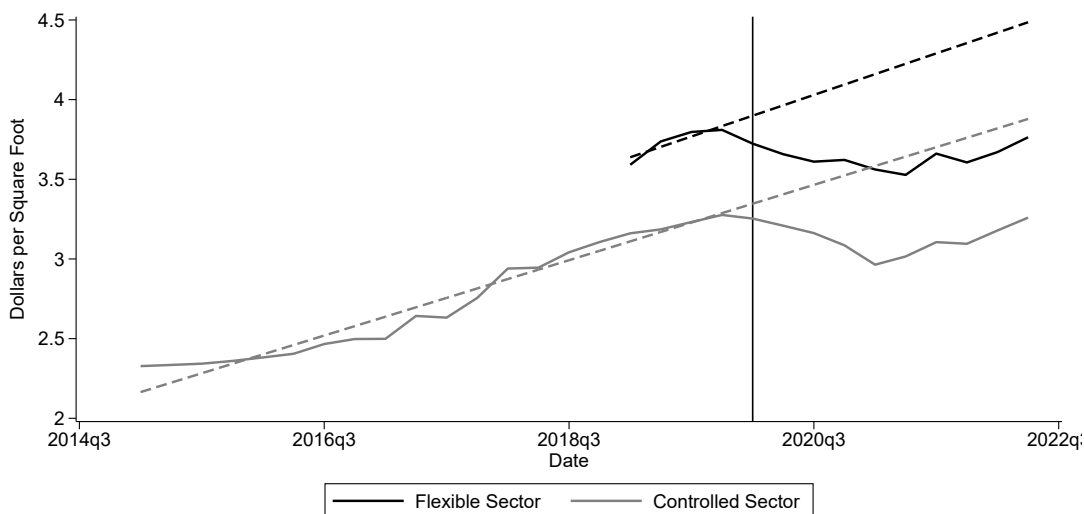
Notes: Construction completions for apartment buildings designated as purpose-built rentals or condominiums from 2010Q1 to 2022Q2. Vertical line at 2020Q1 illustrates beginning of pandemic period. Linear trend constructed using pre-pandemic observations.

the later parts of 2021 and into 2022. I also include a linear trendline estimated using data prior to the COVID-19 pandemic in 2020Q1 in each sector. In both sectors, the trendline indicates the incompleteness of the recovery.

3.1 Specification

In Figure 7, flexible units are more expensive overall than controlled units, which may simply be due to their more recent construction date. To properly study the impact of the demand shock on rent in each sector, I compare units whose first occupancy was just before or

Figure 7: Rent per Square Foot in Flexible and Controlled Sectors



Notes: For each sector, solid line is observed rent per square foot and dashed line is trend calculated using observations before 2020Q1 (vertical line). Controlled sector begins in 2019Q1 because sector was created from rent control exemption policy put in place effective in that period.

just after the policy cutoff date (November 15, 2018), and estimate the following regression:

$$\begin{aligned} \log(\text{Rent per Square Foot}_{it}) = & \beta_0 + \beta_1 \text{Ind.: Rent-Controlled Unit}_i \\ & + \beta_2 \text{Ind.: Pandemic}_t \\ & + \beta_3 \text{Ind.: (Rent-Controlled Unit} \times \text{Pandemic)}_{it} + u_{it} \end{aligned}$$

3.2 Results

Table 4 presents estimates of this specification using four subsamples of the data constructed to include buildings on either side of the policy cutoff date. In each sample, buildings occupied for the first time in November 2018 are excluded, since it is impossible to determine whether they were occupied before or after the exact cutoff date, which landed on November 15, 2018. In the first column, labelled $[-1, 1]$, only buildings occupied for the first time one month before (October 2018) and one month after (December 2018) are in-

cluded and therefore compared in the regression. The indicator for rent-controlled units is close to zero and statistically insignificant, reflecting the fact that these units are nearly identical and command similar rent. As expected, the estimate on the pandemic indicator is negative, and implies that rent during the pandemic was approximately 6 log-points lower than prior to the pandemic. The estimate on the indicator for the interaction term is 3 log-points, implying that rent for controlled units during the pandemic was 3 log-points larger than similar flexible units. Altogether, rent for flexible units decreased during the pandemic by approximately 4.7 log points, while rent for controlled units decreased by only 1.7 ($-4.68 + 3.00 = -1.68$) log points.

Table 4: Differences in Log-Rent per Square Foot Between Controlled and Flexible Units During Pandemic

$100 \times \log(\text{Rent per Square Foot})$ Months Around Policy Cutoff Date:	(1)	(2)	(3)	(4)
	[-3,3]	[-6,6]	[-12,12]	[-18,18]
Ind.: Rent-Controlled Unit	-1.82 (3.12)	-2.43 (2.70)	-6.21*** (1.98)	-6.96*** (1.68)
Ind.: Pandemic	-4.68*** (0.99)	-3.93*** (0.84)	-3.54*** (0.69)	-3.57*** (0.78)
Ind.: Rent Controlled Unit \times Pandemic	3.00** (1.51)	3.29** (1.30)	3.29*** (1.01)	3.42*** (0.99)
N	1,738	2,646	5,271	7,265

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In the second column, buildings occupied up to three months before or three months after the policy date are included in the regression. Now, the difference between rent-controlled and flexible units is slightly larger, since these units are more different, but the estimate is still insignificant. While all units decrease in rent per square foot during the pandemic, rent-controlled units remain 3.3 log-points higher than flexible units. Similar to the narrower window, rent for flexible units decreases by -3.9 log points, while rent for controlled units decreases by only 0.64 log points. Widening the bandwidth to 12 or 18 months before and

after the policy date increases the difference between the two types of units, but it remains the case that during the pandemic, rent for all units decreased, and rent for controlled units was roughly 3 to 4 log-points higher than for flexible units.

To address concerns that the difference between flexible and controlled units is driven by differential change in unit sizes between flexible and controlled units during the pandemic, Table 5 presents estimates of the same specification using total rent. Focusing on buildings occupied for the first time in a three month window around the policy cutoff date, rent is not statistically different between controlled and flexible units. Rent decreased by 8.0 log points for flexible units and 4.32 log points for controlled units. As above, widening the interval around the cutoff date increases the difference in rent between controlled and flexible units, and in every case, rent decreases for both types of units but my more for flexible units.

Table 5: Differences in Log-Rent Between Controlled and Flexible Units During Pandemic

$100 \times \log(\text{Rent})$	(1)	(2)	(3)	(4)
Months Around Policy Cutoff Date:	[-3,3]	[-6,6]	[-12,12]	[-18,18]
Ind.: Rent-Controlled Unit	-4.65 (3.94)	-4.42 (3.00)	-4.66* (2.39)	-5.32*** (1.97)
Ind.: Pandemic	-7.97*** (1.39)	-7.49*** (1.09)	-7.65*** (0.88)	-7.02*** (0.95)
Ind.: Rent Controlled Unit \times Pandemic	3.65* (1.97)	3.95** (1.62)	5.23*** (1.26)	5.48*** (1.21)
N	1,738	2,646	5,271	7,265

Notes: Standard errors in parentheses.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4 Rent Changes Without Rent Controls

In the previous section, I estimated that during the COVID-19 pandemic, rent per square foot decreased by roughly 5 log points for flexible units and only 2 log points for controlled units.

Given the massive impact of the pandemic on the economy, these estimates are relatively small. The model detailed in Section 4.1 rationalizes this finding: in the controlled sector, landlords were reluctant to reduce rent in response to the temporary demand shock due to their inability to quickly return rent to the pre-shock level. The model also elucidates that rent in the flexible sector was also directly impacted by rent controls in the controlled sector, and therefore rent in both sectors decreased by less than it would have in the absence of rent control.

How much would rent have decreased absent rent control? Given the segmented nature of Toronto's rental market into controlled and flexible units, it is straightforward to estimate the change in rent in each sector after the COVID-19 demand shock. However, the simple difference between the change in each sector will not identify the impact of rent controls due to cross-sector price competition.

4.1 A Two-Sector Model of the Market for Rentals

In this section, I build a model of duopolistic competition with differentiated sectors and infrequent price adjustment in one sector. One time period corresponds to one year and the only source of uncertainty is a market-wide demand shock. I model each sector as having a differentiated good to reflect the fact that renters may have preferences over observable differences between rent controlled and flexible buildings. For example, flexible buildings in Ontario are typically newer, while controlled buildings are typically in more established neighbourhoods. The model abstracts from preferences over rent control itself. That is, in what follows, I assume that renters do not select units based on whether or not they are exempt from rent control. This is relevant in the counterfactual analysis because I make all units exempt from rent control and assume demand curves remain the same.

4.1.1 Flexible Sector

Demand for units in the Flexible sector is given by a standard linear demand curve:

$$q_t^F = a_F - b_{FF}p_t^F + \kappa_{FC}b_{FF}p_t^C + z_t,$$

where p_t^F is rent in the Flexible sector at t , p_t^C is rent in the Controlled sector at t , $a_F > 0$ determines differentiated demand for Flexible-sector units, $b_{FF} > 0$ is own-price demand sensitivity in the flexible sector, and $\kappa_{FC} \in [0, 1]$ is relative cross-price sensitivity in the Flexible sector to the price of renting in the Controlled sector. Both sectors are subject to a market-wide AR(1) shock process, z_t , with i.i.d. innovations distributed according to $\sim N(0, \sigma_z)$. In the Flexible sector, the landlord can set the rental price in each period, and therefore solves a static problem,

$$\max_{p_t^F} p_t^F q_t^F = p_t^F (a_F - b_{FF}p_t^F + \kappa_{FC}b_{FF}p_t^C + z_t),$$

with optimal rent given by:

$$p_t^F = \frac{a_F + \kappa_{FC}b_{FF}p_t^C + z_t}{2b_{FF}}. \quad (1)$$

4.1.2 Controlled Sector

As above, demand in the Controlled sector for a given price, p^C , is given by:

$$q_t^C(p^C) = a_C - b_{CC}p^C + \kappa_{CF}b_{CC}p_t^F + z_t,$$

where $a_C > 0$ determines differentiated demand for Control-sector units, $b_{CC} > 0$ is own-price sensitivity, and $\kappa_{CF} \in [0, 1]$ is relative cross-price sensitivity in the Controlled sector to the price of renting in the Flexible sector.

The difference between the two sectors is that landlords in the Controlled sector set rent

only at the initiation of a new tenancy. In the periods that the landlord receives a new tenant, profit maximization is a dynamic problem:

$$\max_{p_t^{C^*}} E_t \sum_{k=0}^{\infty} p_{t+k}^C q_{t+k}^C(p_{t+k}^C).$$

The price the landlord chooses in the current period, $p_t^{C^*}$, is relevant in the current period and all future periods in which no new tenant arrives. Let θ represent the likelihood that an incumbent tenancy ends and a new tenancy begins. Moving all the future terms that include resets (which are not relevant for the optimization) into an auxiliary variable, H , the above expression simplifies to:

$$\begin{aligned} & \max_{p_t^{C^*}} E_t \sum_{k=0}^{\infty} (1 - \theta)^k p_t^{C^*} q_{t+k}^C(p_t^{C^*}) + H \\ & = \max_{p_t^{C^*}} E_t \sum_{k=0}^{\infty} (1 - \theta)^k p_t^{C^*} (a_C - b_{CC} p_t^{C^*} + \kappa_{CF} b_{CC} p_{t+k}^F + z_{t+k}) + H \end{aligned}$$

Taking as given rent in the Flexible sector in the current and future periods, the optimal price for the Controlled landlord is given by:

$$p_t^{C^*} = \frac{E_t \sum_{k=0}^{\infty} \theta (1 - \theta)^k (a_C + \kappa_{CF} b_{CC} p_{t+k}^F + z_{t+k})}{2b_{CC}}. \quad (2)$$

4.1.3 Equilibrium Prices and the Impact of Demand Shocks

The linear demand curves allow a linear model solution in closed form. As a result, the impact of a demand shock on rent is constant for both sectors:

$$\frac{\partial p_t^{C*}}{\partial z_t} = \left(\frac{1}{b_{CC}b_{FF}} \right) \left(\frac{2b_{FF} + \kappa_{CF}b_{CC}}{4 - \kappa_{CF}\kappa_{FC}} \right) \left(\frac{\theta}{1 - (1 - \theta)\rho} \right) \quad (3)$$

$$\frac{\partial p_t^F}{\partial z_t} = \frac{1}{2b_{FF}} + \frac{\kappa_{FC}}{2} \frac{\partial p_t^{C*}}{\partial z_t} \quad (4)$$

These equations summarize the impact of rent control on rent adjustment in light of a demand shock. I note three important insights.

Insight 1: Rent control dampens the effect of demand shocks in the Controlled sector. Rent control forces landlords in the Controlled sector to be forward-looking and anticipate that their changes today have impacts deep into the future. For a given persistence of the demand shock, ρ , the change in rent is smaller because it will stay in place for future periods in which there may be another demand shock.

Insight 2: Rent in the Controlled sector responds more strongly to more permanent demand shocks. In the limit, a fully permanent demand shock is fully priced, regardless of rent control. For less transitory shocks, landlords in the Controlled sector understand that the rent they set in the reset period may outlive the life of the shock, and therefore adjust accordingly to maximize profit until the next time they can change rent.

Insight 3: The Flexible sector inherits a degree of rent control from the Controlled sector. Rent control in the Controlled sector is transmitted to the Flexible sector and governed by the relative cross-price sensitivity in the Flexible market to rent in the Controlled sector, κ_{FC} . When there is no cross-price sensitivity and markets are completely segmented, i.e., $\kappa_{FC} = 0$, then rent control is irrelevant for rent in the Flexible sector. When there is cross-price sensitivity, however, the Flexible landlord understands that their competitor will not fully incorporate the impact of the demand shock, and therefore the Flexible landlord similarly does not fully incorporate the impact of the demand shock.

4.2 Connection to Empirical Results

- connect the three insights above to the empirical results

4.3 Counterfactual: Change in Rent Absent Rent Control

The framework allows me to construct the correct counterfactual changes in rent in the absence of rent control by assuming that a new tenant arrives in every period, i.e., $\theta = 1$. In this case, even the Controlled sector landlord can change rent in each period. I denote these counterfactuals with the following notation:

$$\left(\frac{\partial p_t^{C*}}{\partial z_t} \middle| \theta = 1 \right), \quad \left(\frac{\partial p_t^F}{\partial z_t} \middle| \theta = 1 \right).$$

Using equation 3, the counterfactual change in rent in the Controlled sector if there is no rent control is given by:

$$\left(\frac{\partial p_t^{C*}}{\partial z_t} \middle| \theta = 1 \right) = \left(\frac{1 - (1 - \theta)\rho}{\theta} \right) \frac{\partial p_t^{C*}}{\partial z_t}$$

From equation 2, combined with the counterfactual change in rent in the Controlled sector, the change in the Flexible sector if there is no rent control is given by:

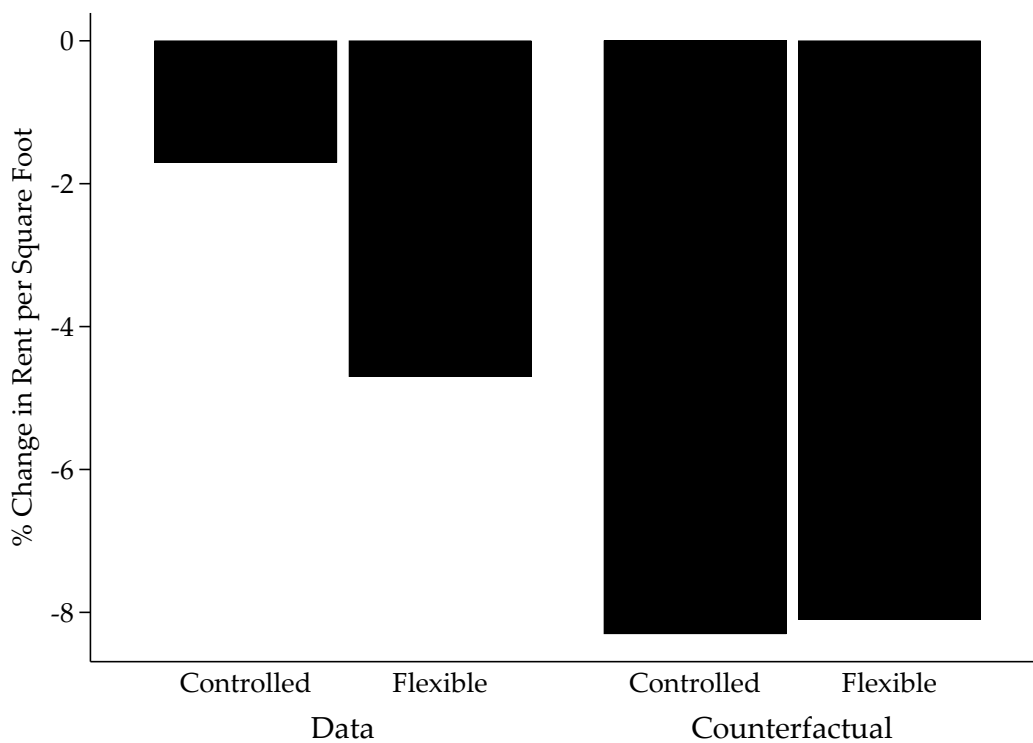
$$\left(\frac{\partial p_t^F}{\partial z_t} \middle| \theta = 1 \right) = \frac{\partial p_t^F}{\partial z_t} + \left[\frac{\kappa_{FC}}{2} \left(\frac{1 - (1 - \theta)\rho - \theta}{\theta} \right) \right] \left(\frac{\partial p_t^{C*}}{\partial z_t} \middle| \theta = 1 \right)$$

The key objects in the counterfactuals are the estimated changes in rent from Section 3, the frequency of new tenant arrivals, θ , the persistence of the demand shock, ρ , and, for the Flexible sector, the cross-price sensitivity, κ_{FC} .

For the frequency of new tenant arrivals, θ , I use the estimate from Han, Ngai and Sheedy (2022) that renters in Toronto remain in their units for approximately three years (Table 3). Using US data, Halket and di Custozza (2015) plot the distribution of rental duration and show a large amount of variation. I set the persistence of the demand shock to 0.25. The

appropriate value for the cross-price sensitivity is hard to pinpoint, but this difficulty is ameliorated by the relative unimportance of this parameter for the counterfactual in the Controlled sector. In the empirical analysis above, the change in rent for buildings first occupied in 2018, just before rent control exemption policy, and in 2019, just after, is stark. It is difficult to believe that the units in these two buildings are not close substitutes, which would imply a value of the cross-price sensitivity closer to unity. At the same time, it may be that newer buildings are indeed uniquely demanded, pushing the cross-price sensitivity closer to zero. For these reasons, and without further evidence on this parameter, I choose a conservative value of 0.50.

Figure 8: Comparison Between Observed Change in Rent and Baseline Counterfactual



Notes: Data reflects observed change in rent during the pandemic for Controlled and Flexible units (see Table ?? for point estimates and standard errors). Counterfactual is model-implied change in rent during the pandemic for each sector under the baseline calibration of setting the persistence of the demand shock to 0.25, the relative cross-sector price sensitivity to 0.50, and the expected tenancy to three years.

Under this baseline parameterization, the model implies that in the absence of rent control, but holding all else equal, rent would have decreased by 8.3% for rent-controlled units and 8.1% for flexible units. Figure 8 plots these counterfactual changes against the observed changes from Table 4. Without rent control, rent for controlled units decreases an additional 6.6 pp, or almost four times the observed change. Rent for flexible units decreases by an additional 3.4 pp, almost double the observed change, highlighting the large extent to which the presence of *any* rent-controlled units impacts flexible units.

4.4 Sensitivity Analysis

In the next three sections, I test the sensitivity of the change in rent in each sector in the absence of rent control for various calibrations of the persistence of the demand shock, the cross-price sensitivity, and the new tenant arrival rate. The tests demonstrate that over a wide range of calibrations, the counterfactual estimates for the change in rent are economically significant in both the Flexible and Controlled sectors. Varying the expected length of the tenancy induces the most variation in the counterfactual scenarios, followed by the persistence of the shock and the cross-price sensitivity.

Persistence of the Demand Shock Panel A of Table 6 presents estimates of the change in rent in the absence of rental control for three calibrations of the persistence of the demand shock. For reference, the first two columns of the table report the same estimates of the change in rent in each sector from Table 4. In each counterfactual, I set the cross-sensitivity of prices between the controlled and Flexible sectors to 0.5 and the tenancy arrival rate to $1/3$, implying an expected duration of three years.

In the first counterfactual, the demand shock is perfectly transitory, i.e., $\rho = 0$. Rent in the Flexible sector decreases by 10.1%, more than double the decrease under rent control, and rent in the Controlled sector decreases by 10.0%, more than fivefold with rent control. Increasing the persistence of the demand shock to 0.25, rent in the Flexible sector decreases

Table 6: Sensitivity Analysis of Counterfactual Changes in Rent

#	Shock Persistence	Cross-Sensitivity	Arrival Rate	Sector	
	ρ	κ_{FC}	θ	Flexible	Controlled
Baseline (Data)	–	–	–	-4.7%	-1.7%
Counterfactual Panel (a): Varying Persistence of Demand Shock					
1	0.00	0.50	1/3	-10.1%	-10.0%
2	0.25	0.50	1/3	-8.1%	-8.3%
3	0.50	0.50	1/3	-6.6%	-6.6%
Counterfactual Panel (b): Varying Cross-Sector Price Sensitivity					
4	0.25	0.25	1/3	-6.5%	-8.3%
5	0.25	0.50	1/3	-8.1%	-8.3%
6	0.25	0.75	1/3	-9.8%	-8.3%
Counterfactual Panel (c): Varying Tenant Arrival Rate					
7	0.25	0.50	0.20	-16.3%	-13.9%
8	0.25	0.50	1/3	-8.1%	-8.3%
9	0.25	0.50	2/3	-5.3%	-4.4%

Notes: Counterfactual changes in rent in the absence of rent control for different calibrations of persistence of the demand shock, cross-sector price sensitivity, and new tenant arrival rate. Data is estimated change in each sector in Table 4.

by 8.1% and in the Controlled sector by 8.3%, and for a persistence of 0.50, rent decreases by 6.6% in both the Flexible and Controlled sectors.

These scenarios illustrate that the counterfactual change in rent is decreasing in the persistence of the shock. As noted above, with rent control, landlords are forced to ignore more transient shocks since their new rent is expected to be in effect for longer than the shock. The counterfactuals show the response of landlords if they can fully internalize the shock into rent. For example, for the observed change to be as large as it is in the data in the case of a perfectly transitory shock, then landlords in the Controlled sector without rent control will decrease their rent by significantly more. As the persistence of the shock increases, the degree of internalization increases, and the counterfactual response decreases (in magnitude).

Cross-Price Sensitivity Panel B of Table 6 presents counterfactuals for three cross-price sensitivities. In the setting studied in this paper, the cross-price sensitivity can be interpreted as the relative size of each sector. The controlled sector is much larger than the flexible sector since the entire stock of buildings first occupied prior to November 15, 2018 are under strict rent control. Given this imbalance, the flexible sector is much more responsive to changes in the controlled sector, reflected in a larger value of the cross-price sensitivity. At the same time, as the degree of sustainability decreases, so too does the cross-price sensitivity.

In all three scenarios, I set the persistence of the demand shock to 0.25 and the tenancy arrival rate to $1/3$, implying an expected duration of three years. In the counterfactual setup, and seen in Equations (3) and (4), the cross-price sensitivity impacts only the demand response for the Flexible sector. In each counterfactual, rent in the Controlled sector decreases by 58 cents. From Equation 3, this change is governed by the parameters for the persistence of the shock and the tenant arrival rate, neither of which vary over the counterfactuals in this section. In this simple framework, Equation 4 shows the linear relationship between the cross-price sensitivity and the two sectoral rents, and this follows directly to the analysis of these counterfactuals.

In the first counterfactual, the cross-price sensitivity is 0.25. Rent in the Flexible sector decreases by 6.5%, roughly 1.8 pp or 40% more than in the data. In the second counterfactual with a cross-price sensitivity of 0.50, rent decreases by 8.7% in the Flexible sector, and in the third counterfactual with a sensitivity of 0.75, by 9.8%. As the Flexible sector inherits more of the rent control from the Controlled sector, the counterfactual change in rent with no rent control increases in magnitude.

New Tenant Arrival Rate Panel C of Table 6 presents three calibrations of the tenant arrival rate, θ , which is the inverse of the expected length of the tenancy. In each scenario, I set the persistence of the demand shock to 0.25 and the cross-price sensitivity to 0.50. The second counterfactual, and the counterfactuals in the last two sections, use the estimate from Han et al. (2022) that the average duration of a renter in Toronto is three years.

The first scenario considers the change in rent if the new tenant arrival rate is $\frac{1}{5}$, corresponding to an expected tenancy length of five years. In this case, in the counterfactual with no rent control, rent per square foot would have decreased by 16.3% in the Flexible sector and 13.9% in the Controlled sector. The intuition for these large changes is similar to above for perfectly transitory shocks. The longer is the expected tenancy, the more the landlord must ignore the demand shock. If under rent control and an expected tenancy of five years the decrease in rent for the Controlled sector is 1.7%, then the counterfactual with no rent control implies a sevenfold decrease in rent since the landlord can fully internalize the shock.

The third scenario considers an expected tenancy of only 1.5 years. Rent decreases by 5.3% in the Flexible sector and 4.4% in the Controlled sector. With this level of tenant turnover, rent control is not very restrictive, and therefore the counterfactual with rent control is not very different than the observed changes in the data.

5 Conclusion

I estimate how much rent decreased in Toronto during the COVID-19 pandemic using detailed building-level data. Exploiting an exogenous policy change that exempted nearly identical units from Ontario's rent control laws, I estimate that rent decreased for controlled units by 1.7% but fell by 4.7% for exempt units. I interpret these results through the lens of a two-sector model of the rental market with rent control in one sector which yields closed-form expressions for the change in rent from a negative demand shock. The Flexible sector decreases rent by relatively more than the Controlled sector since the landlord can freely adjust rent after the shock dissipates. This intuition rationalizes the estimated difference between the Flexible and Controlled sectors.

The model is also necessary to estimate the counterfactual change in rent absent rent control because the Flexible sector inherits a degree of rent control from the Flexible sector

from competing strategically. The key parameters for constructing the counterfactual change in rent absent rent control are the persistence of the demand shock, the cross-sector price sensitivity of demand, and the expected duration of a tenant. I study the counterfactuals for a number of parameter combinations and establish a baseline parameterization using the available empirical evidence.

In the baseline calibration, rent decreases by 64 cents per square foot for Flexible units and 58 cents per square foot for Controlled units. I emphasize that these estimates come from a stylized model that relies on several abstractions from reality. However, these estimates may also understate the true effect of rent control if the demand shock was less persistent or expected durations are longer. Further, using the COVID-19 pandemic as an identified demand shock on housing allows me to study rent control in a partial equilibrium setting and identify a potential short-term detriment to tenants. A full analysis of rent control as it impacts long-term supply and demand of housing requires a full fledged model, such as the one constructed in Favilukis and Nieuwerburgh (2021) or Favilukis, Mabilie and Nieuwerburgh (Forthcoming). Future work will incorporate rent control policies in such a model to fully quantify welfare gains or losses.

References

- Autor, David H., Christopher J. Palmer, and Parag A. Pathak,** "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, June 2014, 122 (3).
- Coven, Joshua, Arpit Gupta, and Iris Yao,** "Urban Flight Seeded the COVID-19 Pandemic Across the United States," *Journal of Urban Economics: Insights*, Forthcoming.
- Diamond, Rebecca, Tim McQuade, and Franklin Qian,** "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, Sept. 2019, 109 (9), 3365–94.
- Favilukis, Jack and Stijn Van Nieuwerburgh,** "Out-of-Town Home Buyers and City Welfare," *Journal of Finance*, Oct. 2021, 76 (5), 2577–2638.
- , **Pierre Mabilie, and Stijn Van Nieuwerburgh,** "Affordable Housing and City Welfare," *The Review of Economic Studies*, Forthcoming.
- Genesove, David,** "The Nominal Rigidity of Apartment Rents," *The Review of Economics and Statistics*, Nov. 2003, 85 (4), 844–853.
- Glaeser, Edward, Caitlin Gorbach, and Stephen J. Redding,** "How Much Does COVID-19 Increase with Mobility? Evidence from New York and Four Other U.S. Cities," *Journal of Urban Economics: Insights*, 2020.
- Gupta, Arpit, Vrinda Mittal, Jonas Peeters, and Stijn Van Nieuwerburgh,** "Flattening the curve: Pandemic-Induced revaluation of urban real estate," *Journal of Financial Economics*, Nov. 2022, 146 (2), 594–636.
- Halket, Jonathan and Matteo Pignatti Morano di Custoza,** "Homeownership and the Scarcity of Rentals," *Journal of Monetary Economics*, Nov. 2015, 76, 107–123.

Han, Lu, L. Rachel Ngai, and Kevin D. Sheedy, "To Own or to Rent? The Effects of Transaction Taxes on Housing Markets," Working Paper 2022.

Suzuki, Masatomo, Yasushi Asami, and Chihiro Shimizu, "Housing rent rigidity under downward pressure: Unit-level longitudinal evidence from Tokyo," *Journal of Housing Economics*, 2021, 52.