Supply, demand, or stickiness?
A causal analysis of the effects of short-term rental activity on residential rents

David Wachsmuth and Cloé St-Hilaire***

* School of Urban Planning, McGill University
** School of Planning, University of Waterloo

September 13, 2024

Abstract

What is the relationship between short-term rental (STR) activity and residential rents? We answer this question with a set of models of STR activity in all Canadian urban areas from 2017-2022. We use the staggered introduction of principal-residence restrictions as a source of exogenous variation in STR activity to measure the impact of STRs on rents in a time-varying difference-in-differences model. We find that STR regulations which reduce the volume of STR activity cause rents to fall in subsequent years, and hence that STR activity has a causal impact on rent levels. We then supplement that approach with direct acyclic graphs for causal hypothesis testing via random-effects eigenvector spatial filtering to decompose STR activity into separate supply, demand, and rental price stickiness channels. We find that commercial STRs and STR prices each independently cause increases in residential rents. This suggests that STRs affect rents in the long-term rental market through a supply channel (by shifting the supply curve leftward as housing units are reallocated from long-term to short-term uses) and a rental price stickiness channel (by increasing landlords' willingness to raise rents aggressively when they have STRs as a plausible fallback). We find weaker support for STRs affecting long-term rents through a demand channel (by shifting the demand curve rightward as residents demand more rental housing in order to offer STRs in their principal residences). The results suggest the ongoing viability of STR regulations as a strategy for housing affordability.

Keywords: *short-term rentals, rents, time-varying difference-in-differences, spatial regression, spatial filtering, causal inference, regulation*

Funding: Social Sciences and Humanities Research Council (Insight Grant 435-2019-0720, "Cities, Short-term Rentals and the Sharing Economy: Housing Impacts, Social Dynamics, and Policy Options")

1. Introduction

Although Airbnb began its corporate life as a platform enabling a disruptive new form of small-scale home-sharing, it has since grown into a giant of the travel industry which has transformed the face of cities around the world. Airbnb and other short-term rental (STR) platforms grew rapidly through the 2010s, and even the Covid pandemic has proven only a temporary speed bump in the development of a professionalized and ubiquitous form of temporary accommodation. Unlike traditional tourist lodgings such as hotels, however, STRs are generally operated out of dwelling units that could otherwise be housing long-term residents. They thus potentially pose zero-sum tradeoffs between the amenity value they pose to tourists and the economic value they generate for their hosts, on the one hand, and the housing needs of local residents, on the other. Acting in part on the fear of these tradeoffs (although often in the absence of reliable evidence), an increasingly large share of municipal governments have imposed restrictions on the operation of STRs.

The key questions are, what is the relationship between short-term rental activity and residential rents, and are governments justified in restricted STR activity in the name of housing affordability? In this paper we combine two independent strategies for measuring the total causal impact of STRs on rent: a time-varying difference-in-differences approach using the staggered introduction of municipal STR regulations across Canadian municipalities as a source of exogenous variation in STR activity, and a set of direct acyclic graphs which allow for hypothesis testing using linear regression. While a growing set of studies have addressed the relationship between STR activity and housing costs, they have mostly done so at narrow spatial and temporal scales, and using relatively coarse measures of STR activity. By contrast, in the following analysis, we take several important steps beyond existing research in the area. First, while much previous research uses a relatively shallow measure of STR activity (often the total number of displayed listings on Airbnb), we use three separate indicators of STR activity (commercial STR listing share of all dwelling units, non-commercial STR listing share of all dwelling units, and nightly STR prices) which allow us to measure separate supply, demand and price-stickiness channels along which STRs could plausibly affect rents. Second, we use a comprehensive country-scaled dataset—all Canadian urban areas—which allows us to assess the impact of STRs on rents across a range of spatial contexts (e.g. large cities versus smaller tourist-oriented communities). Third, our dataset has six years of coverage—2017 to 2022—including before, during and after the Covid pandemic, which gives us access to a wide range of scenarios for our variables of interest. Most previous research on STRs and housing markets was conducted in circumstances where STRs were monotonically increasing; we include such circumstances (2017-2019) but also a period of sharp STR decline (2020-2021) and then a rebound period (2022). We are only aware of two other studies (Ayouba et al., 2020; Lee & Kim, 2023) which have used spatial regression techniques to study the relationship of STRs with housing markets, and those papers concerned a small subset of French cities and a single highly idiosyncratic market—New York City—respectively, in contrast to our comprehensive national data coverage. Finally, we incorporate a set of spatial and temporal effects into our models, using an eigenvector spatial filtering approach to control for spatial autocorrelation in combination with a temporally autoregressive term on the outcome variable and spatiotemporal two-way group effects.

We find that STR regulations which reduce the volume of STR activity cause rents to fall in subsequent years, and hence that STR activity has a causal impact on rent levels. We further find that changes in STR activity causes changes in rents along two of our three hypothesized channels: increases in commercial STRs and STR prices each independently cause increases in residential rents. This suggests that STRs affect prices in the long-term rental market through a supply channel (by shifting the supply curve leftward as housing units are reallocated from long-term to short-term uses) and a price stickiness channel (by increasing landlords' willingness to raise rents aggressively when they have STRs as a plausible fallback). Our results are also directionally consistent with a further effect of non-commercial STRs on residential rents through a demand channel (by shifting the demand curve rightward as residents demand more rental housing in order to offer STRs in their principal residences), but this relationship falls short of statistical significance.

The paper is organized as follows. First, we present a brief literature review of existing research on the relationship between STRs and housing costs. We then provide an empirical overview of the STR market in Canada. Third, we develop the theoretical models and hypotheses to be tested. Fourth, we provide a description of the data and methods used. We then present the findings of our models. We conclude with a discussion of the findings, including the policy implications of our causal estimates. An attached appendix provides additional methodological details, along with diagnostics and robustness checks for all the empirical results.

2. Previous literature

Over the past decade, short-term rental platforms such as Airbnb have emerged as a leading topic of concern for urban economists as well as housing and tourism researchers. Scholars have examined the impacts of STRs on the tourism industry (Guttentag, 2015; Zervas et al., 2017; Cocola-Grant & Tago, 2022), house prices (Sheppard & Udell, 2016; Barron & al., 2021; Todd et al., 2022), neighbourhood crime (Cheung & Yiu, 2023), gentrification (Wachsmuth & Weisler, 2018; Spangler, 2019; Lee & Kim, 2023), vacancy rates (Lie & Xie, 2020; Hill et al., 2023), and rents (Horn & Merante, 2017; Garcia-López et al., 2020; Chang, 2020; Barron et al., 2021; Garay-Tamajón et al., 2022). Scholars have also examined the growing resistance to the Airbnb-led 'touristification' of cities (Morales-Pérez et al., 2022; Marrone & Peterlongo, 2020; Simcock, 2023), and the increasingly vigorous efforts by local and sometimes supra-local governments to enact pro-housing STR regulations (Gurran & Phibbs, 2017; Smigiel, 2020; Hill et al., 2023; Miller, 2024; Wachsmuth & Buglioni, 2024).

2.1. Short-term rentals and housing costs

An increasingly prominent theme of the research on STRs concerns the relationship between STR activity and housing costs. This research has consistently found that greater presence of STRs predicts higher rents, although the scope and level of such impact varies. The most widely cited study is Barron et al.’s (2021) US-wide analysis of Airbnb and housing costs, which found that a one percent increase in Airbnb listings in an area predicted to a 0.018 percent increase in rents in the United States. In Los Angeles, Koster et al. (2021) found that a one percentage point increase in Airbnb listings (or a 0.69 standard deviation increase) was associated with a 4.9% increase in rents. In Taiwan, Chang (2020) determined that a one-standard-deviation increase in Airbnb listings led to a 0.38% increase in rents. In Barcelona, Garcia-López et al. (2019) also found that 54 more active listings in a small neighborhood increased rents by 1.9%. Differentiating their findings by rental unit size in Tel Aviv, Ram and Tchetchik (2022) found that a 1% increase in the number of Airbnb listings was associated with a 1.4% increase in rents for apartments of 4.5 or more rooms, of 1.17% for apartments of 2.5 to 3 rooms, 0.77% for apartments of up to two rooms, and 0.22% for apartments of 3.5 to 4 rooms.

Several studies have examined the effect of changes in STR listing density (i.e. STRs as a proportion of total housing stock or rental housing stock) as opposed to absolute quantity of listings. Horn and Merante (2017) found that a one-standard-deviation increase in Airbnb density was associated with a 0.4% increase in asking rents. Ayouba et al. (2020) found that a one-point increase in the density of commercially-operated Airbnbs led to an increase of average rents by 1.71% and 1.24% in Marseille and Paris, respectively. In London, Shabrina et al. (2022) also examined commercially-operated Airbnbs and determined that a 100% increase in full-time Airbnb listings operated in entire homes led to up to an 8% increase in rental prices per-bedroom per-week, amounting to an average increase of £90 per year.

Studies have found intra-city differences in the relationship between STRs and rents, with Garay-Tamajón et al. (2022) arriving at the conclusion that STRs had the greatest impact on rents in “highly touristified” and “trendy” or “more affluent” neighbourhoods in Spain, where Airbnb listings tended to be concentrated. In Hong Kong, Liang et al. (2022) found that tenants paid 3.61% higher rents in Airbnb-served neighbourhoods. Most analyses of STRs and housing costs have been single-city case studies, although some research has examined larger regions (Rodriguez-Pérez de Arenaza et al., 2022), multiple cities (Ayouba et al., 2020), or entire countries (Barron et al., 2021; Chang, 2020).

A frequent assumption of housing research on STRs is that commercial STRs are responsible for a disproportionate share of negative housing-market impacts, since they represent housing units removed from the long-term market (Combs et al., 2020; Simcock, 2023). This assumption has generally been borne out in quantitative research. Studies have found that commercial STR listings have the strongest impacts on rents (Shabrina et al., 2022; Lee & Kim, 2023), as well as other negative neighbourhood externalities such as noise, littering, and overcrowding (Nieuwland & van Melik, 2020; Celata & Romano, 2022). Lee and Kim (2023) treated Airbnb listings as heterogeneous by differentiating between listings operated by hosts with a single listing and listings operated by hosts with multiple listings, and by differentiating between entire home listings and private and shared rooms. They found that entire-home listings operated by multi-unit hosts were the short-term rentals with the greatest impact on rent, housing value, and gentrification. Ram & Tchetchik (2022) used unit size (e.g., number of bedrooms in the entire home listing) to disaggregate the impact of STRs on rents.

Depending on data availability, studies typically either use asking rents gathered from platforms such as Zillow or sitting rents, which are usually found in government censuses, as an outcome variable. Models typically include neighbourhood attributes as control variables, including neighbourhood categories, tourist attractions, accessibility to jobs and services, income levels, gentrification variables, and distance from city centre.

The type of model most commonly used to examine the relationship between STRs and rents is linear regression with fixed effects, including OLS regression, shift-share regression, and difference-in-differences models. Many models use spatial and temporal fixed effects, but very few studies have attempted to control for spatial dependence beyond simply using fixed effects. The two exceptions we are aware of are Lee and Kim (2023) and Ayouba et al. (2020), who employed a spatial Durbin model and a spatial error model, respectively. Table 1 offers a summary of peer-reviewed research examining the relationship between STR activity and housing costs.

Study	**Model type**	**Outcome variables**	**STR heterogeneity**	**Spatial effect**	**Scale**	**Years**
Horn & Merante (2017)	OLS	Rent Rental housing supply	Size	FE	City (Boston)	1.4
Ayouba et al. (2020)	Spatial error	Rent	Commercial STRs	Spatial error	Multi-city (France)	8
Chang (2020)	OLS	Rent House prices	Listing type	FE	Country (Taiwan)	4.75
Garcia-López et al. (2020)	IV; DiD	Rent House prices Sale price of sold housing	No	FE	City (Barcelona)	3
Barron et al. (2021)	IV	Rent House prices Price-to-rent ratio	No	FE	Country (United States)	5
Benítez-Aurioles & Tussyadiah (2021)	Generalized method of moments	Rent House prices	No	No	City (London)	4
Franco & Santos (2021)	IV; DiD	Rent House prices	Commercial STRs	FE	Country (Portugal)	7
Koster et al. (2021)	Regression discontinuity; DiD	Rent House prices	Listing type	FE	City (Los Angeles)	4
Liang et al. (2022)	DiD	Rent Rent-to-income ratio Unaffordability ratio	No	FE	City (Hong Kong)	5.5
Rodríguez-Pérez de Arenaza et al. (2022)	OLS	Rent	No	No	Region (Andalusia)	0.33
Ram & Tcherchnik (2022)	Generalized method of moments	Rent	Size	No	City (Tel Aviv)	3

Study	Model type	Outcome variables	STR heterogeneity	Spatial effect	Scale	Years
Shabrina et al. (2022)	OLS	Rent	Commercial STRs (Airbnb misuse)	No	City (London)	2
Lee & Kim (2023)	Spatial Durbin	Rent Housing value Poverty propensity	Commercial STRs	Spatial Durbin	City (New York)	4
Duso et al. (2024)	DiD	Rent per square metre Rental housing supply	Commercial STRs	FE	City (Berlin)	4.75

Table 1. Summary of previous research on the relationship between STR activity and housing costs

3. Short-term and long-term rentals overview in Canada

Before presenting detailed model specifications and results, we offer a brief overview of short-term rentals and the rental housing market in Canada. The growth of the country's STR market has occurred against the backdrop of steady and nearly universal increases in housing costs over the last decade. From 2015 to 2023, the average monthly rent in communities of 10,000 people or more in Canada increased from $899 to $1266 in nominal terms—a 4.4% annual increase during a time when inflation averaged 2.8%. Figure 1 shows that this increase was remarkably consistent across different regions in the country.

During the same time frame, short-term rentals grew rapidly, shrunk in the face of the Covid pandemic, and then began to grow again. Figure 2 shows two measures of STR activity across five regions in Canada—Atlantic Canada (the provinces of New Brunswick, Newfoundland and Labrador, Nova Scotia, and Prince Edward Island), British Columbia, Ontario, the prairies (the provinces of Alberta, Manitoba and Saskatchewan), and Québec. The two measures are average daily active listings (the number of listings which are either reserved or available for reservations) and 12-month frequently rented entire-home (FREH) listings (listings active at least 183 nights and reserved at least 90 nights in the past twelve months). FREH listings are a simple measurement of housing units which have been converted to dedicated STRs, since a housing unit which is active as an STR a majority of the time is unlikely to also be the host's principal residence. The figure demonstrates that all regions saw steep and consistent growth in both metrics from 2016 through the beginning of 2020, then depressed activity from mid-2020 through roughly the beginning of 2022, and finally a return to growth from early 2022 onward. As of September 2022, there were 100,270 average active daily listings and 37,620 12-month FREH listings. These numbers represent a 13.9% and 5.6% decline, respectively, from their pre-pandemic peaks. In other words, while active listings and 12-month FREH listings both returned to growth after 2021, they had yet to recover from the pandemic by September 2022.

A final perspective on the evolution of STRs in Canadian communities is found in Figure 3, which shows active STRs as a percentage of all dwelling units in communities of different sizes over time. It demonstrates, first of all, that smaller communities have, on average, a higher concentration of STRs than larger communities. Second, in communities of 50,000 dwellings or

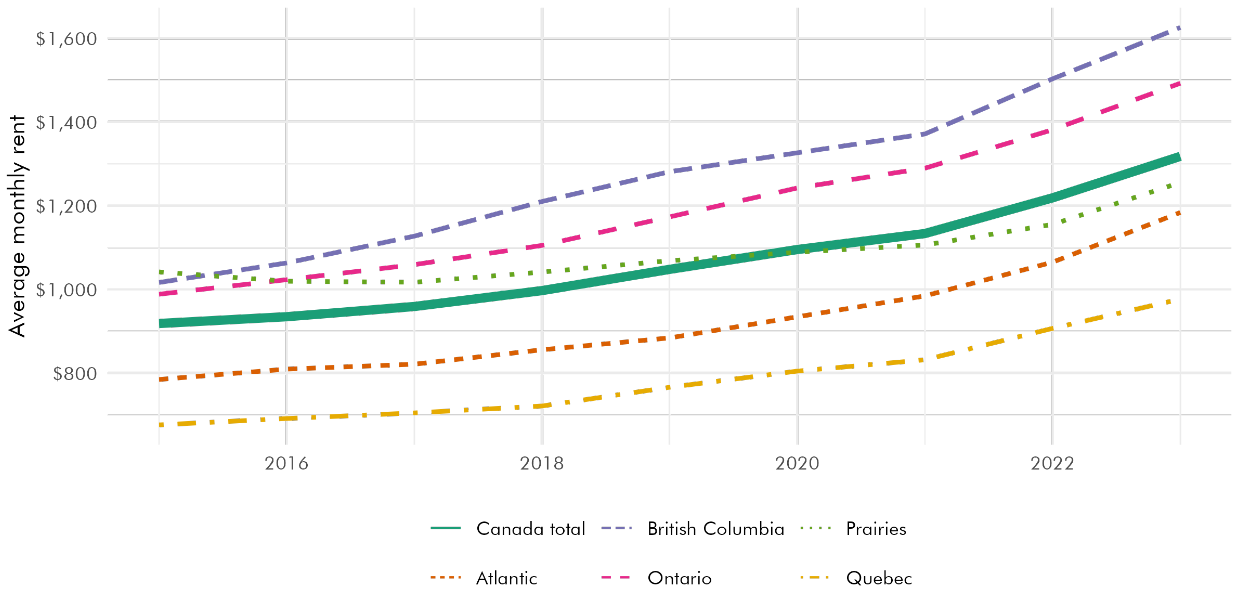

Figure 1. Average monthly rents in Canada, 2015-2023 (2023 dollars)

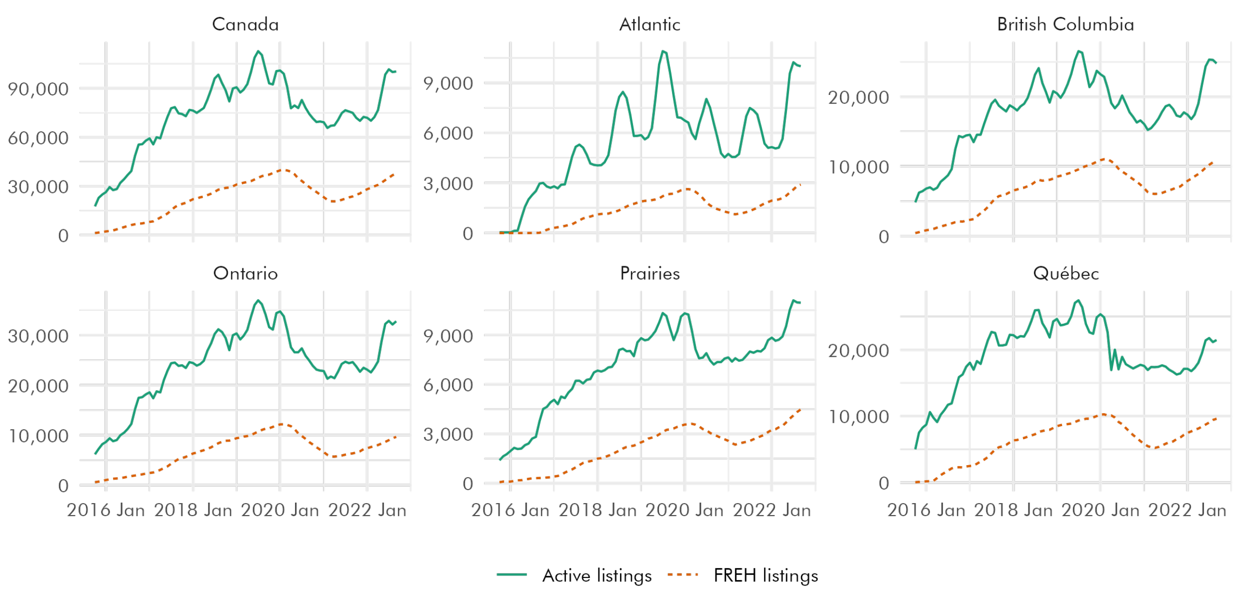

Figure 2. Active listings and frequently rented entire-home (FREH) listings by region of Canada (monthly average)

more, active listings still lagged behind pre-pandemic levels as of September 2022, while in smaller communities active listings were at record highs as of September 2022.

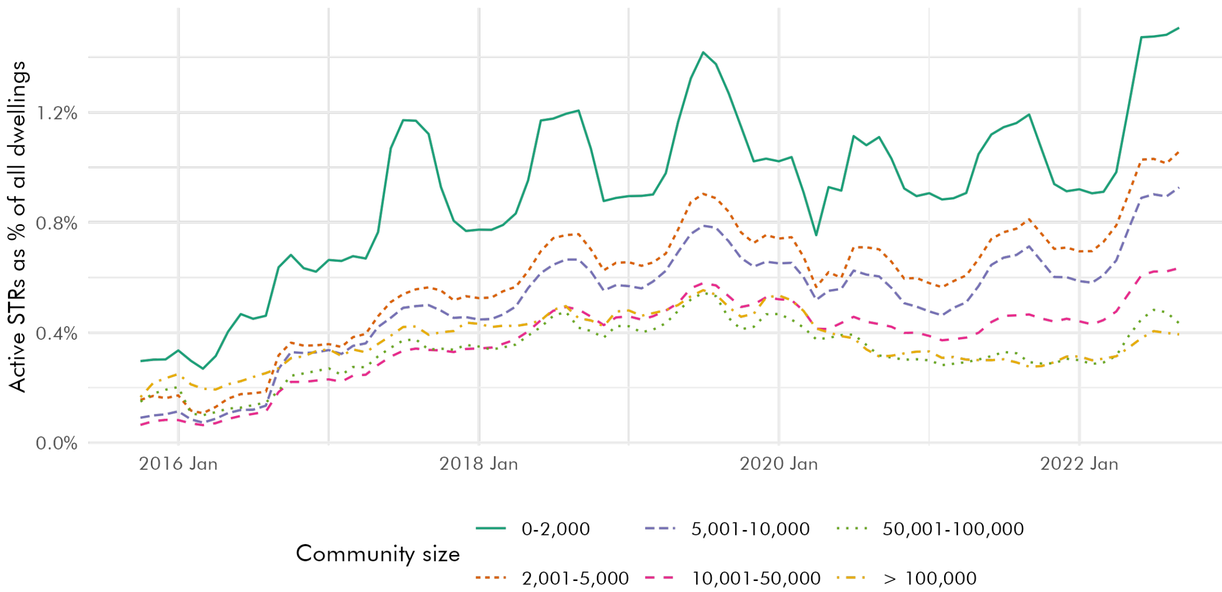

Figure 3. Active short-term rentals as a share of all dwellings by community size in Canada (per-community monthly average)

4. The relationship between short-term-rental activity and residential rents

The short-term segment of the rental market is not new, but the arrival of Airbnb and other online STR platforms has transformed the relationship between short-term and long-term segments, specifically by allowing greater permeability between the two. Specifically, landlords of traditional long-term rental housing have increased opportunities to reallocate their units to the STR market, and tenants have increased opportunities to rent spare capacity in their units (either extra rooms or the entire unit while they are away from home). Here we consider the short-run consequences of these changes, first with the assumption that long-term rentals operate in a competitive equilibrium, and later allowing for price stickiness in the long-term rental market.

Under competitive equilibrium conditions, the decision of landlords whether to allocate their units to the long-term or short-term market can be understood as a function of the place-specific difference between long-term rental prices and short-term rental prices, mediated by the place-specific cost of operating a short-term rental relative to a long-term rental and by idiosyncratic landlord preferences.[1] In equilibrium, long-term rental prices are an outcome of renters' inverse demand function and short-term rental prices are exogenously determined. The growth of online STR platforms such as Airbnb makes it easier to operate an STR, thereby reducing its operating cost relative to the operating cost of a long-term rental, and thus makes some landlords who would otherwise have allocated their units to the long-term market allocate them instead to the

[1] Here we build on Barron et al. (2018), in particular on their identification of decreased STR operating cost as the key STR-induced shock to long-term rental markets.

short-term market on a full-time basis. This outcome represents a leftward shift in the supply curve for long-term rentals: at a given price point, fewer housing units are offered, because some share of them have been reallocated to full-time short-term rentals.

Tenants also have the opportunity to operate short-term rentals by renting out their excess capacity in home-sharing arrangements. Their decision whether to do so is a function of short-term rental prices mediated by the cost of operating an STR and idiosyncratic tenant preferences. (We assume here that tenants do not have the opportunity to rent their excess capacity on the long-term market, but the conclusions remain the same if we relax this assumption.) In equilibrium, tenants demand an amount of housing where the additional rent they are required to pay is equal to their expected gain from home sharing (again mediated by idiosyncratic preferences, which in most cases will result in no home-sharing arrangement at all). The arrival of online STR platforms lowers STR operating costs, and thus makes it more tempting for tenants to rent their excess capacity as an STR. At a given long-term rental price, therefore, tenants will demand more housing. This outcome thus represents a rightward shift in the demand curve for long-term rentals.

Third, we relax the assumption that the long-term rental market operates in a competitive equilibrium, specifically by introducing the possibility of rental price stickiness. In fact, rental price stickiness is a well-known (albeit little-researched) phenomenon of rental housing markets (Genesove, 2003). Gallin and Verbruge (2019) suggest that the key force generating rental price stickiness is that landlords can preempt a tenant's search for a new apartment by offering to renew a lease under its existing terms, and thus that stickiness will be more observed among units operated by small landlords who are more averse to turnover-induced vacancy. In a compatible argument, Wang (2020) argues that higher vacancy rates lead to higher rental price stickiness, because they empower tenants in the bargaining process. When the cost of operating an STR decreases, this should also affect price stickiness, by offering landlords a solution to the aversion to turnover-induced vacancy that Gallin and Verbruge (2019) argue is the key factor explaining stickiness. Landlords could more aggressively pursue rent increases with existing tenants if they have the option to shift their unit to the short-term sector of the rental market in the event that their tenant leaves, and higher STR prices make that option more appealing.

Finally, we consider the effects of regulations which constrain landlords or tenants from freely allocating their units or space within their units to the STR market. Such regulations should be expected to reduce the stock of both dedicated STRs offered by property owners and home-sharing STRs offered by tenants, and thus to shift the long-term rental supply curve right and the demand curve left. Importantly, there is no plausible mechanism by which STR regulations could affect the supply and demand curves for long-term rental housing *except* through their impact on STRs. This means that regulation can serve as a plausible source of exogenous variation in STR activity for the purposes of measuring true causal impacts of STRs on rents.

The hypothesized causal relationships described here are captured in the directed acyclic graph (DAG) in Figure 4. The DAG describes three causal pathways along which the growth of online STR platforms affects long-term rental prices. First, they should shift the supply curve left by causing some housing units to be reallocated to full-time STRs, which we measure by calculating the share of housing units operating as frequently rented entire-home (FREH) STR listings and therefore not housing a long-term resident. Second, the growth of STRs should shift the demand curve right by causing tenants to demand more housing in order to operate home-sharing STRs, which we measure by calculating the share of housing units operating as non-FREH "home sharing" STRs where a long-term resident is present. Third, the growth of STRs should reduce rental price stickiness by reducing the downside risk of landlords demanding large rent increases from existing tenants, which we measure by calculating the average nightly price for a reserved STR. Finally, we should expect STR regulations which restrict STR activities to cause rents in the long-term market to decrease, by reducing the quantity of both dedicated and casual STR listings. We should not expect STR regulations to affect STR prices, however, since the latter are exogenously determined by competition with hotels and other sources of tourism accommodation, and by underlying housing quality. (We return to this point below.) Moreover, STR regulations should act as a source of exogenous variation in STR activity with respect to long-term rents, since there is no plausible channel by which STR regulations could affect long-term rents *except* through their impacts on STR activity. As noted above, this analysis only concerns the short run, and the long-run consequences of the growth of online STR platforms could be quite different. For example, in the long run it is plausible that STR demand would stimulate new residential investment—this is the implication of Bekkerman et al.'s (2023) finding that residential permits decline in the wake of the imposition of STR regulations.

5. Data

The models in this paper test the causal effect of STR activity on rents in Canadian urban areas. To implement these models, we rely on four data sources: rental market data from the Canada Mortgage and Housing Corporation, short-term rental data from Airdna, additional sociodemographic data from the Canadian Census, and a hand-compiled dataset of STR regulations in Canada. The data sources are as follows:

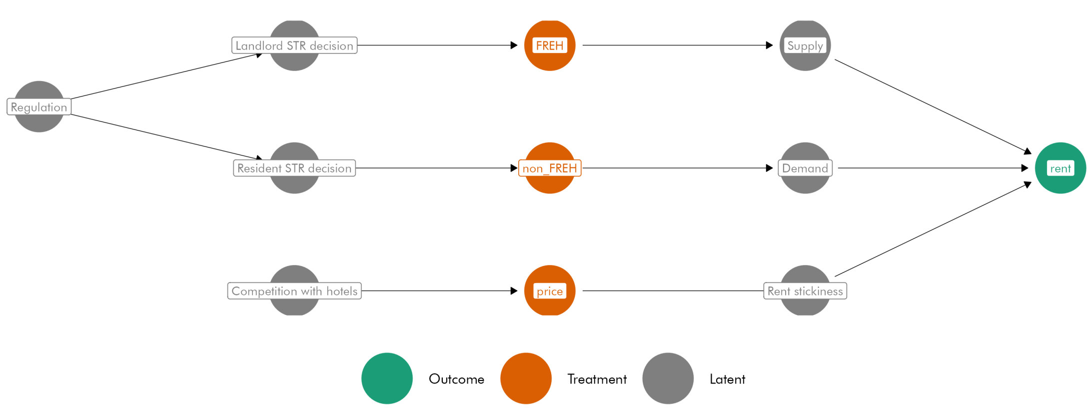

Figure 4. DAG showing hypothesized causal pathways between STR activity, STR regulation, and rent

5.1. Housing data from the Canada Mortgage and Housing Corporation

To measure our outcome variables—average rent and year-over-year change in average rent—we rely on data from the Canada Mortgage and Housing Corporation (CMHC)'s annual Rental Market Survey. This survey is conducted in October of each year in all urban areas with populations of 10,000 or more, and it targets purpose-built rental dwellings of three or more rental units. We use the 2015-2023 Rental Market Surveys with data aggregated at the CMHC neighbourhood level. Neighbourhoods are CMHC-defined collections of census tracts for which data availability is high; these geographies represent the best tradeoff between data granularity and data completeness, since at the census-tract scale there are many missing values in CMHC's rent data. In smaller municipalities CMHC does not define neighbourhoods; in those cases we use municipality-level values. We exclude all neighbourhoods with four or more missing values for total rental units or seven or more missing values for average rent among the nine time periods. To address non-normality in average rent, we log-transform the variable. There are no negative or zero values which could complicate a log transformation.

5.2. STR activity data from Airdna

To measure STR activity, we built a dataset from web-scraped data about Airbnb short-term rental listings gathered by the consulting firm Airdna. This dataset includes information about every STR listing on the Airbnb platform which was active in Canada between October 1, 2016 and September 30, 2022—six complete years. The data includes "structural" information such as the listing type (entire home, private room, shared room or hotel room) and the approximate location of the listing. Airdna collects this information through frequent web scrapes of the public Airbnb and Vrbo websites. The data also includes estimates of listing activity (was the listing reserved, available, or blocked, and what was the nightly price?), which Airdna produces by applying a machine-learning model to the publicly available calendar information of each listing. We process this raw data through a cleaning pipeline using the R *strr* software package (Wachsmuth, 2021). We create three treatment variables from the cleaned data: FREH listings measured as a percentage of dwelling units, non-FREH listings measured as a percentage of dwelling units, and the average nightly price of STR reservations.

5.3. Control variables from CMHC and the Canadian Census

Third, we introduce a set of control variables taken from the 2021 Canadian Census and from the CMHC Rental Market Survey. The former are taken at the dissemination area scale— a set of adjacent blocks with an average population between 400 and 700 persons, which is the smallest scale at which Statistics Canada releases census data. These are log-transformed average total household income, the log-transformed percentage of total employment which is in the entertainment and accommodation sectors, and the log-transformed percentage of occupied private dwellings which are apartments. From the CMHC Rental Market Survey we retrieve the log-transformed and one-year-lagged vacancy rate in purpose-built rental units.

5.4. Hand-compiled STR regulation database

Finally, we developed a comprehensive dataset of STR principal-residence restrictions in every neighbourhood covered by CMHC rent data. This dataset indicates whether any portion of a neighbourhood is subject to a principal-residence restriction, and, if so, the date on which the restriction was implemented. A neighbourhood is considered treated in a given year if there was a principal-residence restriction in place by January 1 of that year, and untreated otherwise. (Since rents are measured in October and the STR variables measured in September, in all cases we allow for at least nine months for treatment effects to occur.) Although we collected information on STR regulations for all neighbourhoods in our main dataset, we only include the 798 neighbourhoods located in provinces which had at least one instance of treatment in the study period (British Columbia, Ontario, New Brunswick, and Quebec), in order to avoid violations of the parallel trends assumption, as discussed below. There were 38 cities which implemented STR principal-residence requirements on or before January 1, 2023, and they collectively account for 314 of the 798 neighbourhoods in the dataset. More details on this dataset are provided in the appendix.

5.5. Dataset assembly and summary statistics

We assembled the final dataset as follows. For each of the 1,080 CMHC neighbourhoods in our dataset, we gathered the average rent and the vacancy rate for each year from 2015-2023, which yielded 9,720 observations. For each of these observations between 2017 and 2022, we aggregated STR data for the September which preceded the survey, and calculated three measures. The first is frequently rented entire-home (FREH) listings, which we calculate by identifying all the listings which were a) an entire-home listing (as opposed to a private or shared room), b) available or reserved a majority of the nights in the prior twelve months, i.e. 183 nights, and c) reserved at least 90 nights in the prior twelve months. The second is non-FREH listings, which we calculate by identifying all listings which were active in September (i.e. had at least one night where they were either reserved or available for a reservation) but which were not an FREH listing based on the definition described above. The third is price: the average nightly price charged by listings for a reserved night in September. We aggregated these measures per neighbourhood and per year, by matching the latitude and longitude of the STR listing with the neighbourhood containing them. Then we joined a set of covariates from the Canadian Census, by using areally weighted interpolation on variables collected at the dissemination area level to estimate the corresponding values at the CMHC neighbourhood level. We constructed our three basic treatment variables (*FREH*, *non_FREH* and *price*) for each neighbourhood by, respectively, dividing the count of FREH listings by the total number of housing units, dividing the count of non-FREH listings by the total number of housing units, and dividing the sum of STR revenue by the total number of reserved nights. To create the final point-in-time version of these variables (*FREH_log*, *non_FREH_log* and *price_log*), we applied a log transformation to these variables after shifting zero values to the minimum non-zero value in the dataset. We likewise applied a log transformation to average rent to create our point-in-time outcome variable *rent_log*. We calculate year-over-year change versions of our outcome (*rent_change*) and treatment

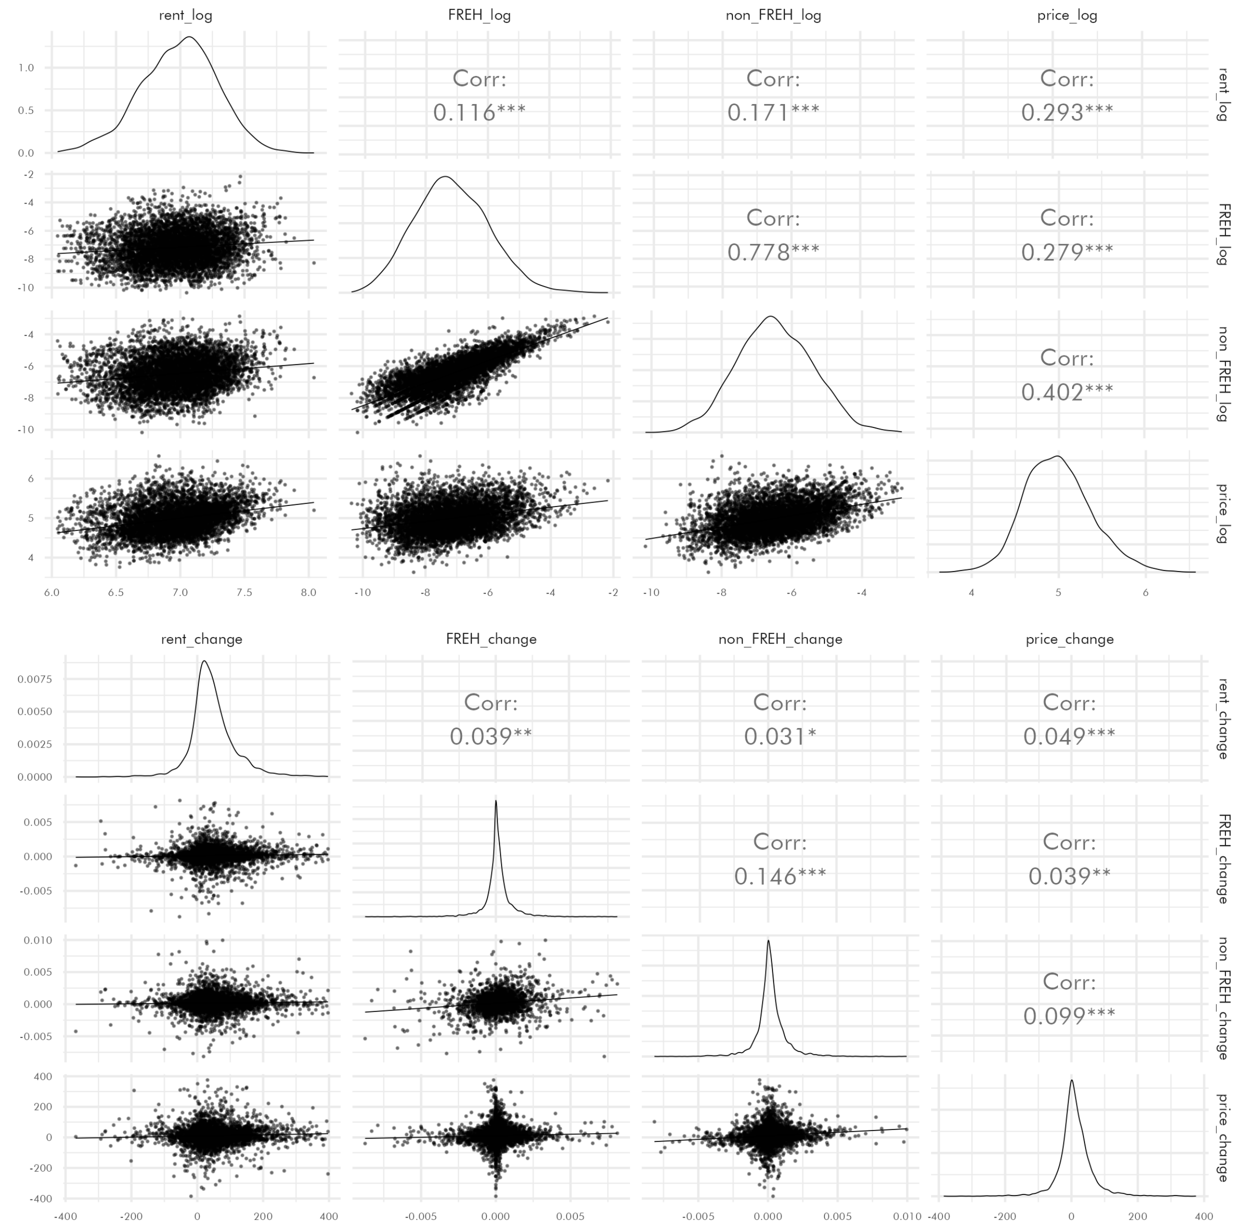

Figure 5. Bivariate relationships between rent_log, FREH_log, non_FREH_log, *and* price_change *(top panel), and between* rent_change*,* FREH_change*,* non_FREH_change*, and* price_change *(bottom panel)*

(*FREH_change*, *non_FREH_change*, and *price_change*) variables using the pre-log-transformed versions. Finally, to enable balanced panel regressions, we use imputation to fill in missing values for rents and vacancy rates. Our method of imputation identifies, for each neighbourhood i in time t, the six nearest neighbours K, and calculates

$$y_t = mean\left(\frac{rent_{i,t}}{rent_{K,t}}\right)$$

Variable	Description	Source	N	Mean	SD	Range
rent_log	Average monthly rent of purpose-built rentals (log)	CMHC	9,210	6.93	0.301	5.93 – 8.04
FREH_log	FREH listings as % of dwellings (log)	Airdna	8,640	-8.64	2.44	-12.1 – -2.18
non_FREH_log	Non-FREH listings as % of dwellings (log)	Airdna	8,640	-7.28	1.85	-11.1 – -2.86
price_log	Average nightly STR price (log)	Airdna	8,640	4.69	0.769	3.05 – 7.20
rent_change	Y.o.y. change in rent	CMHC	7,987	48.3	81.3	-605 – 717
FREH_change	Y.o.y. change in FREH listings as % of dwellings	Airdna	4,320	0.000	0.001	-0.013 – 0.040
non_FREH_change	Y.o.y. change in non-FREH listings as % of dwellings	Airdna	4,320	0.000	0.001	-0.018 – 0.027
price_change	Y.o.y. change in average nightly STR price	Airdna	4,320	16.8	64.3	-965 – 1,115
rent_lag_log	Lagged average monthly rent (log)	CMHC	8,243	6.90	0.292	5.93 – 8.04
vacancy_lag_log	Lagged purpose-built rental vacancy rate (log)	CMHC	7,056	-4.05	1.17	-6.91 – -1.21
income_log	Average total household income (log)	Census	9,720	11.5	0.281	10.7 – 12.8
apart_log	% of occupied private dwelling units which are apartments (log)	Census	9,720	-1.35	0.820	-4.46 – 0.007
tourism_log	% of employment in entertainment and accommodation (log)	Census	9,720	-2.59	0.311	-3.62 – -1.27

Table 2. Descriptive statistics of model variables

for the years in which no observations are missing. We then fit a linear model $y = \beta_0 + \beta_1 year$ for each neighbourhood, and use that model to predict missing observations. (We follow the same process for the vacancy rate.) Our main model uses this imputed data set, but the results of the same regression with no imputed values are described in the robustness checks section of our appendix. A summary of the dataset is given in Table 2. All variables are standardized prior to the regressions being run; the summary statistics in Table 2 pertain to the pre-standardized values. Figure 5 shows a correlation matrix with scatterplots for the three treatment variables alongside the outcome variable, in their point-in-time and year-over-year forms. In both cases, all three bivariate relationships of interest show weak to moderate positive correlations. The code necessary to run the models and obtain all the results in the paper is available under an MIT license at https://anonymous.4open.science/r/cmhc-rent-2023-8FF5/.

6. Identification strategy

Let $Y_{i,t}$ be the average rent in a neighbourhood i at time t and let $STR_{i,t}$ be a place-and-time-specific measure of STR activity. Based on the analysis in section 4 above, we posit the following causal relationship between these variables:

$$Y_{i,t} = \alpha + \beta STR_{i,t} + \gamma X_{i,t} + \epsilon_{i,t} \quad (1)$$

where $X_{i,t}$ is a vector of observed time-varying variables at the neighbourhood level, and $\varepsilon_{i,t}$ contains additional unobserved factors which may affect $Y_{i,t}$. To identify the impacts of STR activity on residential rents we employ a two-part identification strategy to generate complementary estimates. The first part is a time-varying difference-in-differences (DiD) approach which exploits the widespread (although staggered) adoption of STR regulations in Canadian municipalities. This gives us an exogenous source of variation in STR activity which we use to generate an estimate of the impact of STR activity on residential rents and to test the causal assumptions developed in section 4 above. However, while the DiD approach allows us to identify the overall causal effects of STR activity on rents and to rule out any endogeneity in that relationship, it does not allow us to decompose those effects into the three separate channels along which we hypothesize that STR activities can affect rents (supply, demand, and rental price stickiness). We therefore supplement the main DiD model with an additional structural causal model which allows us to identify adjustment sets of conditioning variables to measure the total causal effect of our treatment variables and also yields testable implications to establish the external validity of the causal claims embedded in the model. By implementing this model using a random-effects eigenvector spatial filtering regression with controls for temporal autocorrelation and group-wise fixed effects, we are able to identify accurate parameter estimates for the true causal effect of three separate STR activity treatment variables on residential rents—subject to there being no endogeneity stemming from omitted variables in the model. The combination of these two approaches—a DiD approach which establishes an exogenous source of variation in our treatment variable, and a structural causal model approach which allows us to decompose that treatment variable into separate, theoretically relevant channels—gives us high confidence that we can retrieve the true causal relationship between STR activity and residential rents; we discuss each of them in turn.

6.1. Time-varying difference-in-differences model

As Gibbons & Overman (2012) have argued, identifying true causal relationships requires the identification of an exogenous source of variation in the the treatment variable. We find such a source in the introduction of STR principal-residence restrictions by municipal governments. While there is a wide variety of STR regulations which have been implemented in Canada (like in other countries) at the local and subnational scales, one of the most common is a principal-residence restriction. Generally speaking, this policy requires STR operators to live in the housing unit that they are offering as an STR, which rules out offering dedicated STRs at all. (Although sometimes principal-residence restrictions have an exemption for an accessory dwelling unit or secondary suite.) While the precise details of the principal-residence restrictions vary from municipality to municipality, if our basic causal model shown above in Figure 4 is correct, we should expect to see the following. Principal-residence restrictions will directly reduce the supply of FREH listings, and, as a knock-on effect of making operating an STR more burdensome (e.g. by requiring proof of residency and acquisition of an operating permit), will also indirectly reduce the supply of non-FREH listings. These effects should exert downward pressure on rents through the supply and demand channels, causing them to decline relative to a counterfactual scenario where the restrictions had not been introduced. The price stickiness channel, which is affected by STR nightly prices rather than overall quantities of STR listings or revenue, should not be affected by principal-residence restrictions, since STR prices are exogenously determined by tourism demand, competition with hotels and other forms of tourist accommodation, and underlying housing quality.

Importantly, there is no plausible channel along which STR regulations could affect rents except through their impact on STR activity. So, if the introduction of STR regulations is exogenous of previous rent trajectories (a crucial condition which we test in the next section), then measuring the effect of STR principal-residence regulations on rents gives us an exogenous source of variation in STR activity. This will allow us to generate parameter estimates for the causal impact of STR activity on rents, and also test the assumptions of our causal model, since we have assumed in that model that only FREH and non-FREH listings (and not STR prices) are affected by regulations.

We explore the effects of STR regulations with a difference-in-differences (DiD) model, using the implementation of STR principal-residence restrictions as the treatment, and rent as the outcome. Since Canadian municipalities have implemented principal-residence restrictions at different time periods, a traditional DiD research design cannot be used, since the latter assumes a consistent treatment period for all cases which receive treatment. Instead, we use Callaway and Sant'Anna's (2021) recently developed procedure for estimating DiD causal effect parameters across multiple treatment periods. An outcome Y for a unit i at period t is given by:

$$Y_{i,t} = Y_{i,t}(0) + \sum_{g=2}^{T} (Y_{i,t}(g) - Y_{i,t}(0)) * G_{i,g} \tag{2}$$

Here $Y_{i,t}(0)$ is the potential outcome if the unit were to remain untreated throughout all time periods, and then $Y_{i,t}(g)$ is the potential outcome if the unit were to be treated at time g, summed across all possible time periods 2, …, T. Lastly, $G_{i,g}$ is a binary variable which is true if a unit i was first treated in period g and false otherwise. This specification gives a single potential outcome path for each unit, which is a function of the time period when the unit was treated (and is simply $Y_{i,t}(0)$ if the unit is never treated).

This multi-treatment-period DiD model allows for a a similarly multi-period generalization of the standard DiD causal parameter average treatment effect of the treated (ATT). Callaway and Sant'Anna (2021) call their casual parameter the "group-time average treatment effect", ATT(g,t), and they define it as the expected value of the difference between treatment and no treatment for members of a group g at a time period t, as follows:

$$ATT(g,t) = \mathbb{E}[Y_t(g) - Y_t(0) \,|\, G_g = 1] \tag{3}$$

We implement the model in equation (2) using *rent_log* (log-transformed average rent) as the outcome variable and principal-residence regulation implementation as the treatment to test the main contention of hypothesis 4 (that STR regulations will reduce rents). We then implement the same model with the STR variables (*FREH_log*, *non_FREH_log* and *price_log*) to test the validity of the causal model (which predicts that STR principal-residence regulations will cause

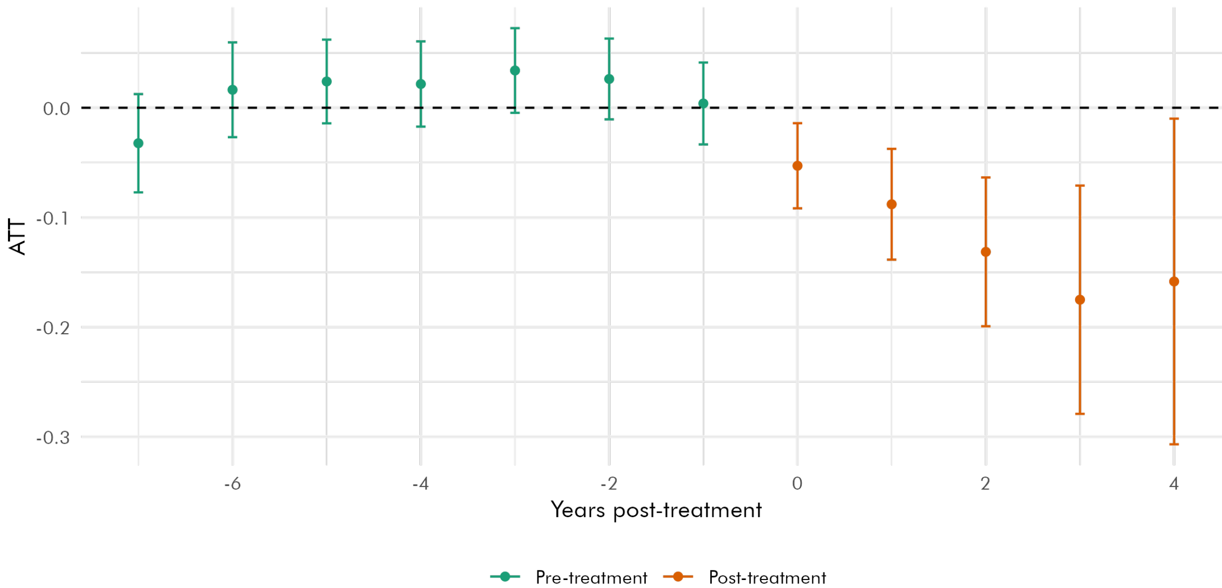

Figure 6. Average treatment effect of the treated of STR principal-residence regulations on rent_log *by length of exposure (point estimates with 95% confidence intervals)*

FREH and non-FREH listings to decline but will not have an effect on STR prices). The regressions run with STR variables as outcomes use an anticipation period of one, meaning that we allow for some pre-knowledge of the impending treatment on outcomes in the year when the regulations are being enacted but are not yet in force.

6.2. Parallel trends assumption test

For the DiD approach to measure actual causal impact, the treatment needs to be applied effectively randomly to individuals, in the sense that the causal pathway between the treatment and the outcome is unidirectional. If individual's previous outcome variable trajectory influences whether they receive the treatment, then the estimates of treatment effects generated by a DiD regression will be biased. Our choice of STR regulations as a treatment raises the prospect that this treatment is not properly exogenous from the outcome variable of average rents, since it is possible that municipalities choose to implement STR restrictions in response to increasing rents. If this is the case, then we are unable to use STR restrictions as a source of exogenous variation in STR activity.

The key test that establishes exogeneity of the treatment from the outcome is the parallel trends assumption: the assumption that, in the absence of treatment, the difference between the treated and non-treated groups is constant over time. This assumption is most commonly tested using an event-study regression—a regression on the outcome variable incorporating both leads and lags of the treatment variable which can identify when treatment effects begin to manifest. However, event-study regressions fail in the face of heterogeneous treatment effects which imply selective treatment timing (Sun and Abraham, 2021). Selective treatment timing is a plausible concern in the present study; for example, municipalities facing larger potential benefits from STR regulation may choose to opt into these regulations sooner than other municipalities. Callaway and Sant'Anna (2021) offer an alternative method for testing the parallel trends assumption which exploits their group-time ATT concept, which is to aggregate average treatment effects by length of exposure. Similar to an event-study regression, for the parallel trends assumption to hold we should expect to see zero treatment effects for negative exposure lengths, and non-zero treatment effects for non-negative exposure lengths.

Figure 6 displays average treatment effects by exposure length for our *rent_log* DiD model, and demonstrates that the parallel trends assumption holds: at 95% confidence intervals, ATTs are statistically indistinguishable from zero for all pre-treatment (negative) exposure lengths, and are negative for all post-treatment (non-negative) exposure lengths. Put differently, we are able to conclude with confidence that municipalities implementing STR principal-residence restrictions do not have different pre-treatment average rent trajectories in comparison to municipalities that did not implement such restrictions. This finding may seem in tension with the high likelihood that municipalities do in fact implement principal-residence restrictions because they hope such restrictions will reduce housing costs. But a plausible explanation is that increasing housing costs are a necessary but not sufficient condition for cities to implement STR restrictions, and since rent has been increasing rapidly throughout Canada (see Figure 1, above), this condition was satisfied in effectively every municipality in the study period. The remaining sufficient conditions could include rapid growth in STRs, the relative strength of pro-STR or anti-STR lobbying groups, the importance of tourism to the local economy, or other factors. Regardless of what they are, figure 6 demonstrates that these factors are exogenous to residential rents, and thus that our time-varying DiD regression can accurately capture the causal impact of STR regulations on rents.

6.3. Structural causal model

The time-varying DiD model in equation 2 allows us to exploit a source of exogenous variation in STR activity to retrieve the true causal relationship between STR activity and rents, but it does not allow us to distinguish between different aspects of that activity. In particular, our basic causal model described in section 4 and Figure 4, above, decomposes STR activity into three dimensions of interest: FREH listings, which are hypothesized to affect long-term rents by shifting the supply curve for long-term rentals, non-FREH listings, which are hypothesized to affect long-term rents by shifting the demand curve for long-term rentals, and STR prices, which are hypothesized to affect long-term rents by decreasing rental price stickiness for long-term rentals.

We therefore supplement the DiD model in equation 2 with a spatial panel regression model which decomposes the STR treatment variable into separate supply, demand and rental price stickiness channels. While this model cannot itself prove a causal relationship between STR activity and long-term rents (since it does not have a source of exogenous variation in the treatment variable), it allows additional insight into the causal pathways along which STR activity effects long-term rents, conditional on the DiD model in equation 2 establishing the overall causal relationship.

We therefore begin by formalizing the intuitions presented in section 4 and Figure 4, above, into a structural causal model of price determination in the rental housing market which can be implemented with spatial panel regression. Our structural causal model builds on equation 1. Although our outcome variable of interest is the level of average rents, this variable is non-stationary (as Figure 1 above demonstrates), and thus unsuitable for analyzing using panel regression techniques. Therefore, as our outcome variable we instead use the first difference—year-to-year change in average rents—which is stationary. We thus arrive at the following hypothesized causal relationship:

$$rent\_change_{i,t} = \alpha + \beta STR\_change_{i,t} + \gamma X_{i,t} + \epsilon_{i,t} \tag{4}$$

We now consider the effects of changes in STR activity on changes in rents along the three separate pathways identified in section 4: supply, demand, and price stickiness. We consider each of these factors separately along with its potential interaction with STR activity, then consolidate the factors into a model of rent price determination which generates testable causal hypotheses.

First, we formalize the argument from section 4 that growth in commercial STRs will shift the long-term rental supply curve to the left as follows. First, we assume that rental housing supply can be decomposed into two sectors (a long-term rental sector and a short-term rental sector) and that landlords are able to freely allocate their rental units between these two sectors. At time *t-1*, landlords take stock of current housing market and STR factors: the prevailing neighbourhood rent (*rent_lag_log)*, the prevailing short-term rental price (*price_lag_log*), the vacancy rate (*vacancy_lag_log*), and the presence of STR regulations (an unobserved variable in this model). Independent of rent and STR price levels, the presence of STR regulations should make landlords less able (or not able at all) to allocate their units to the STR market. On the basis of these three factors (rent, STR prices, and STR regulations) and taking into account their idiosyncratic preferences, landlords make the decision to allocate their units to the short-term or long-term markets, and they implement those decisions over the course of the next year such that by time *t* we can measure the aggregated outcome as the difference between FREH listings in *t-1* and *t* (in both cases measured as a share of total dwelling units), which we henceforth refer to as *FREH_change*. *FREH_change* is thus determined by a combination of the quantity of FREH listings in time t-1 (*FREH_lag_log*) and the outcome of the landlord decision described above. We assume, finally, that rental housing supply is determined by *FREH_change*, the addition or subtraction of new purpose-built rental units (*universe_change*), and an additional set of unobserved factors which are causally unrelated to the relationship between STR activity and rents (*U1*). Because our analysis concerns the short run, we assume that additions or subtractions in purpose-built rental housing supply are exogenous; over the longer run they would be expected to respond to changes in rent levels.

Second, we formalize the argument that growth in non-commercial home-sharing STRs will shift the long-term rental demand curve to the right as follows. Consider the operation of rental housing demand when that demand can be decomposed into the standard household demand for residential services and also demand for additional residential space to operate a non-dedicated STR out of a resident's principal residence. Residents' decision whether to offer a non-FREH STR out of their principal residence as the outcome of four factors: the prevailing neighbourhood rent (*rent_lag_log)*, the prevailing short-term rental price (*price_lag_log*), the share of dwelling units located in apartment buildings (*apart_log*), and the presence of STR regulations (an unobserved variable in this model). Higher rents and higher STR prices should both make residents more likely to operate casual STRs—the former by increasing the salience of additional income and the latter by making the operation of the STR more lucrative per unit of effort necessary to operate it—while apartment buildings are arguably less viable for home-sharing STRs than they are for commercial STRs. Although STR regulations generally target commercial STRs and allow home sharing uses with few or no restrictions, the presence of any regulatory regime arguably raises the barrier to entry on operating an STR, and thus might reduce the quantity of non-commercial STR activity. *non_FREH_change* is thus determined by a combination of the quantity of non-FREH listings in time t-1 (*non_FREH_lag_log*) and the outcome of the resident decision described above. Finally, we decompose rental housing demand into three factors: *non_FREH_change*, the average total household income in a neighbourhood (*income_log*), and a set of unobserved factors which are causally unrelated to the relationship between STR activity and rents (*U2*).

Third, we formalize the argument that the presence of STRs should reduce price stickiness in long-term rents. If landlords have the option to shift their unit to the short-term sector of the rental market in the event that their tenant leaves, they can more aggressively pursue rent increases with existing tenants. The viability of this shift will be a function of the price the landlord could receive from offering the unit as an STR in relation to the rent they could receive on the long-term market. Previous research has found STR pricing to be correlated with a range of listing-level, host-level, and neighbourhood-level factors (Deboosere et al., 2019), but these can mostly be separated into factors which describe the underlying quality and suitability of the housing stock (e.g. the prevalence of entire-home listings, the number of bedrooms, and the average neighbourhood income) and factors which describe the consequences of STR operators responding to competition with hotels and B&Bs in the face of greater or lesser exogenous demand for tourism accommodation (e.g. the prevalence of "superhosts" and the prevalence of multi-listing hosts). We thus model *price_change* as the outcome of pricing in the previous time period (*price_lag_log*), tourist demand (proxied by the proportion of local employment which is located in the entertainment and accommodation sectors, *tourism_log*), competition with hotels (unobserved in our model, but plausibly unconnected to any other variables in our model), and *rent_lag_log*, which captures housing quality. Finally, we can thus decompose rental price stickiness into six factors: rent in the previous time period (*rent_lag_log*, which also allows for temporal autocorrelation in the outcome variable), the vacancy rate (*vacancy_lag_log*), the landlord mix (which is unobserved but which we proxy with *apart_log*, the proportion of occupied private dwelling units which are in apartment buildings), change in the average nightly price of STRs in the previous time period (*price_change*), the change in rental units (*universe_change*, which introduces new units into the market that didn't previously have rents set, and thus cannot have sticky prices), and a set of unobserved factors which are causally unrelated to the relationship between STR activity and rents (*U3*).

In sum, we capture the supply, demand and price-stickiness impacts of STR activity on rent through a decomposition of STR activity into, respectively, FREH listings, non-FREH listings, and STR prices. This gives us an updated model:

$$rent\_change_{i,t} = \alpha + \beta FREH\_change_{i,t} + \gamma non\_FREH\_change_{i,t} + \delta price\_change_{i,t} + \eta X_{i,t} + \epsilon_{i,t} \tag{5}$$

In Figure 7 we formalize the model into a directed acyclic graph (DAG), which also includes some potential additional confounding unobserved variables (*U4*, *U5* and *U6*). Although DAGs are not yet commonly used in econometric research (Imbens, 2020), they are increasingly ubiquitous in other domains of the social sciences where researchers wish to identify causal effects in non-randomized studies, and offer a powerful framework for identifying: 1) confounding effects which interfere with the measurement of causal effects, 2) covariates which allow for statistical adjustment to control for bias, and 3) additional testable implications which are not directly related to the measurement of causal effects but allow the evaluation of whether the assumptions encoded in the model are empirically viable (Thoemmes et al., 2018). For illustrative purposes, Figure 7 displays the DAG for the treatment variable *FREH_change*; the figure is reproduced for the other two treatment variables in the appendix.

To measure the total causal effect of *FREH_change* (or *non_FREH_change* or *price_change*) on *rent_change*, we need to block any so-called "back-door" paths between the treatment and outcome variable. Formal analysis of the DAG using the procedure in Textor et al. (2016) allows us to identify the minimal sufficient adjustment sets of variables which accomplishes this task. Adjustment sets must be identified separately for each treatment-outcome variable pair, since a single model generally cannot measure the total causal impact of multiple treatment variables simultaneously (Westreich and Greenland, 2013). We identify three sets for each of the treatment variables; however, there is one common set:

- *FREH_change*, *non_FREH_change*, *price_change*, *rent_lag_log*, *vacancy_lag_log*, *apart_log*, *income_log*

A regression model which includes this set of variables allows for the simultaneous measurement of the total causal effect of the three treatment variables *FREH_change*, *non_FREH_change*, and *price_change* on the outcome variable *rent_change*. In other words, if the structural causal model in Figure 7 has correctly identified the true causal relationships determining rent price changes, then a model with the common adjustment set will allow us to measure the total causal effect of each of our treatment variables.

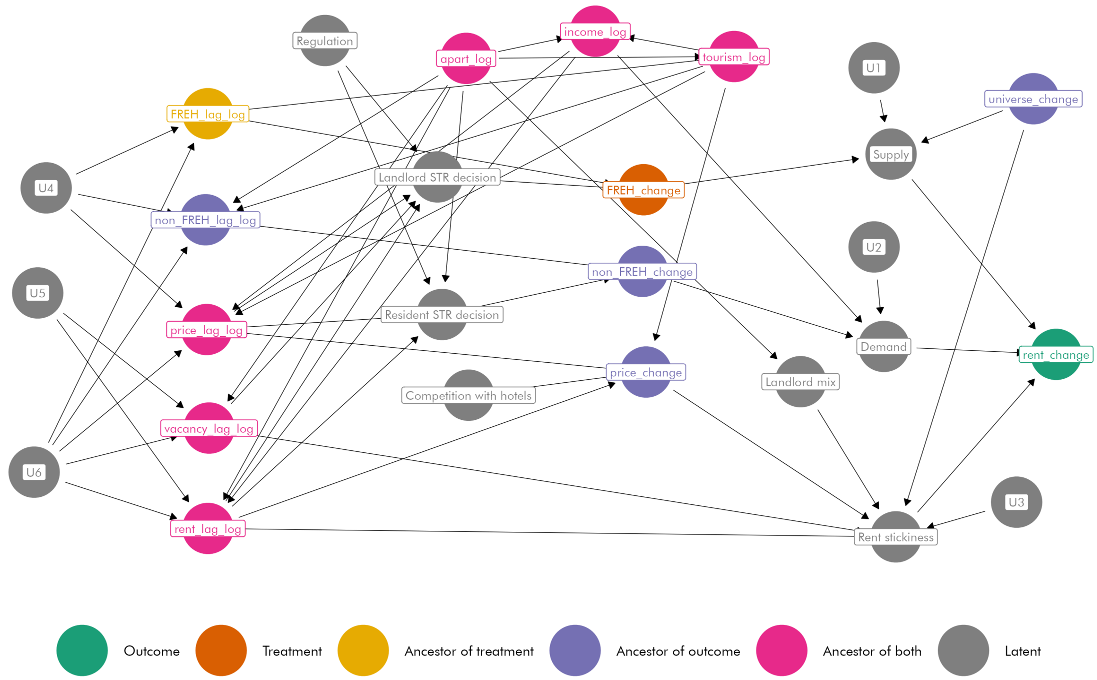

Figure 7. DAG showing hypothesized causal pathways between FREH_change *and* rent_change

6.4. Conditional independences in structural causal model

A necessary step in establishing the validity of a structural causal model formalized in a DAG is conditional independence testing. Causal effects can only be inferred correctly if the underlying structural model accurately represents the processes under investigation, and while the causal relationships cannot be directly tested, these relationships imply a set of conditional independences among variables, which offer testable implications. As an example, if the DAG in Figure 5 is a correct model of the true causal relationships between the variables contained within it, then *income_log* and *vacancy_lag_log* should be independent of each other, after controlling for *apart_log*. If the conditional independences implied by the DAG are confirmed with empirical data, this establishes confidence that the causal relationships are correctly specified. Following the procedure described in Ankan et al. (2021), we conduct tests for each conditional independence involving two or fewer conditioning variables, by calculating OLS regressions for each variable set. (Full results are reported in the appendix.) Figure 8 demonstrates the outcome of the 18 independence tests which involve two or fewer conditioning variables. The figure displays Pearson correlation coefficients, and in 12 of the 18 cases the 95% confidence interval for the coefficient overlaps zero, so we are unable to reject the null hypothesis that the variables are independent. In the remaining six cases, the correlation coefficients are low (never higher than 0.28), and generally concern variable combinations where autoregressive effects were not fully modelled in our structural causal model because we only included a single time period of lagged variables. We are therefore able to conclude that our structural causal model is consistent with the empirical data it is modelling, and thus that we are

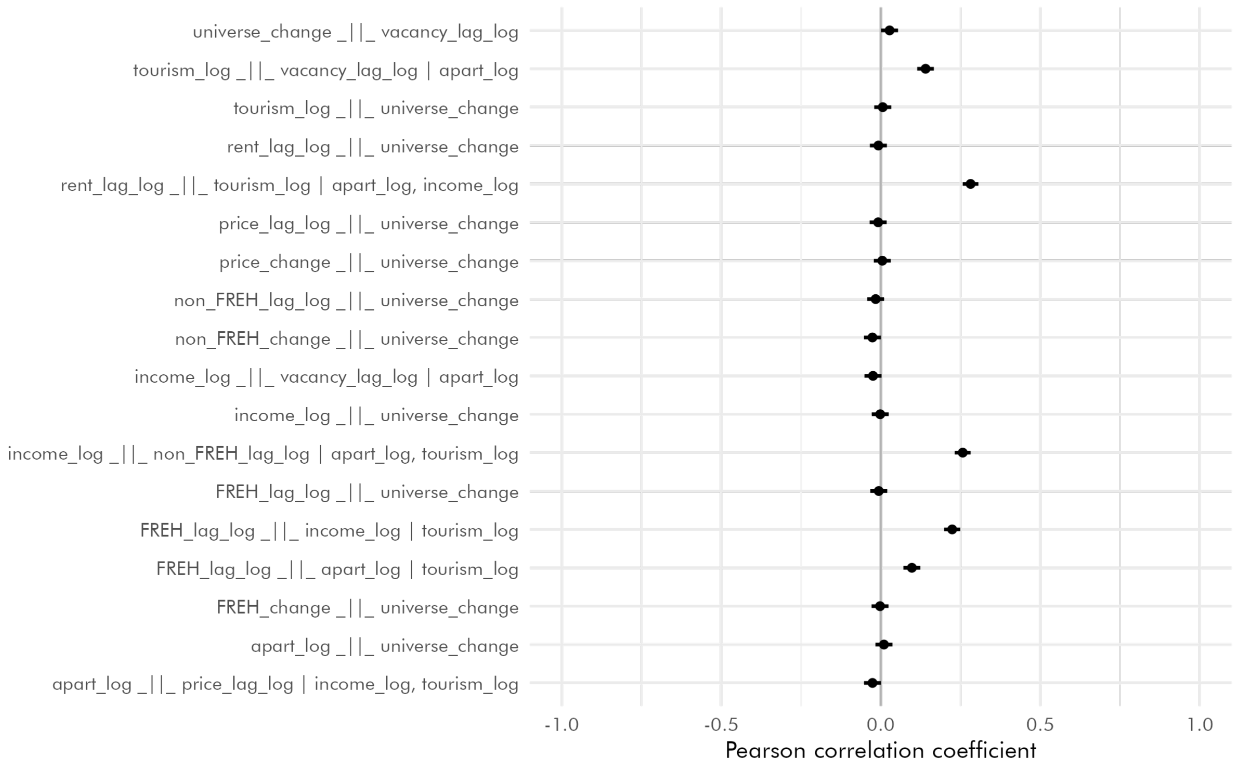

Figure 8. Conditional independences involving two or fewer conditioning variables implied by structural causal model

able to make causal inferences based on the hypothesized causal relationships embedded in the model. Importantly, this conclusion only holds for the variables present in the model; conditional independence analysis does *not* allow us to address the possibility of omitted variables. However, the DiD model in section 6.1 allows us to establish the overall causal relationship between STR activity and long-term rents in a manner which rules out endogeneity, given that the parallel trends assumption was satisfied in section 6.2.

6.5. Random-effects eigenvector spatial filtering panel model

We implement the structural causal model with a random-effects eigenvector spatial filtering (RE-ESF) panel regression. This is an extension of ordinary-least-squares linear regression that explicitly addresses spatial dependence. Although spatial fixed effects are commonly used in econometric research to address spatial dependence, Anselin and Arribas-Bel (2013) demonstrate that they are unreliable in addressing the common issue of non-group-wise dependence. To more robustly account for spatial dependence, therefore, we use an RE-ESF model (Murakami & Griffith, 2015). Eigenvector spatial filtering is a procedure to detect and then address spatial structure in linear regression residuals, and has been shown to outperform more traditional spatial-error models (Sun et al., 2021). It works by taking a spatial weights matrix which represents the spatial structure of observations in a dataset and identifying the eigenvectors which maximize the Moran's I statistic—a commonly used measure of global spatial dependence. These eigenvectors are then added as explanatory variables in the regression model, and thereby account for spatial autocorrelation. The RE-ESF model extends this procedure through the use of random effects to address confounding between treatment variables and latent spatial structure (Hughes & Haran, 2013); we use the RE-ESF variant developed by Murakami & Griffith (2015) which addresses potential oversensitivity to the specific spatial weights matrix chosen by endogenously estimating the correct scale of spatial dependency.

The RE-ESF model for a set of neighbourhoods *i* at time *t* is specified as follows:

$$y_{i,t} = \sum_{k=1}^{7} x_{i,t,k}\beta_k + f_{MC}\left(g_{I(0)}\right) + \sum_{h=0}^{1} \gamma\left(g_{I(h)}\right) + \epsilon_{i,t}, \epsilon_{i,t} \sim N\left(0,\sigma^2\right) \quad (6)$$

Here y is the average year-over-year change in rent in a neighbourhood i between times $t-1$ and t. x_1 , x_2 and x_3 are the the treatment variables (*FREH_change*, *non_FREH_change*, and *price_change*), and x_4 , …, x_7 are the additional 4 set of neighbourhood-level covariates given in the common adjustment set (*rent_lag_log*, *vacancy_lag_log*, *apart_log*, and *income_log*). β are parameters to be estimated, and β_1, β_2 and β_3 give the total causal effect of the three treatment variables. In order to control for unobserved factors at the neighbourhood or region-by-year level (with regions defined as Census Metropolitan Area [CMA] laboursheds), we include $g_{I(0)}$ and $g_{I(1)}$, which are grouping variables at these two levels. $f_{MC}\left(g_{I(0)}\right)$ is a function of the Moran coefficient matrix (derived from the spatial weights matrix) applying spatially dependent group effects at the neighbourhood level, and $\gamma\left(g_{I(h)}\right)$ is the spatially independent group random effects for group h. ϵ is a place- and time-specific error term. We implement this using the *spmoran* package in R (Murakami, 2017).

7. Results

7.1. Difference-in-differences model

Table 3 displays the parameter estimates for aggregated group-time ATTs for a set of difference-in-difference models with *rent_log*, *FREH_log*, *non_FREH_log*, and *price_log* as the outcome variables, and the implementation of STR principal-residence restrictions as the treatment. The results demonstrate that STR regulations have a strongly negative causal impact on rent levels: on average, neighbourhoods located in municipalities which implement principal-residence restrictions have logged average rents which are 0.096 standard deviations lower than they would be if those neighbourhoods did not become subject to such restrictions. The effect of STR regulations on FREH and non-FREH listing levels is even more pronounced: the decline in the log-transformed version of those variables for treated neighbourhoods relative to untreated counterfactuals are 0.251 standard deviations and 0.207 standard deviations, respectively. These results confirm the basic causal account offered in section 4, above: STR activity causes rents in to increase, and STR regulations which decrease STR activity correspondingly cause rents to decrease.

Outcome variable	Average ATT(g,t)	Standard error	Confidence interval (95%)	
rent_log	-0.096***	(0.015)	-0.127	-0.066
FREH_log	-0251***	(0.030)	-0.310	-0.192
non_FREH_log	-0.207***	(0.021)	-0.247	-0.166
price_log	-0.034	(0.034)	-0.102	0.033

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001; Average ATT(g,t): weighted average of group-time average treatment effects of the treated

Table 3. Parameter estimates for aggregated group-time ATTs for four DiD models with STR principal-residence restrictions as treatment

Figure 9 shows average treatment effects by time period, and demonstrates increasing efficacy of STR principal-residence mandates in reducing rents relative to a non-treatment counterfactual. This is a finding which is consistent with the well-documented difficulties municipalities have faced in applying STR regulations in a binding manner (Boeing et al., 2021; Nieuwland & Van Melik, 2020; Wachsmuth & Buglioni, 2024); it is plausible that, over time, principal-residence restrictions have become better implemented, and hence have had stronger ameliorative effects on rents. This interpretation is strengthened by the FREH and non-FREH estimates in Figure 9, which similarly show improvements over time. The figure also includes a variant of the main models which delays the treatment date for the City of Vancouver's STR principal-residence regulations by one year, from 2019 to 2020. Vancouver was the Canadian pioneer in aggressive municipal STR regulation, and Vancouver neighbourhoods dominate the 2019 treatment group in our dataset. If we assume that Vancouver's regulations were slow to take hold, the average treatment effects by time period become considerably more even. (Full details for this model variant as well as several other variants are presented in the appendix.)

Furthermore, these results allow us to partially confirm the plausibility of our structural causal model presented above in Figure 7. We assumed that STR regulations would directly reduce the ability of landlords to offer dedicated FREH STRs, and would indirectly discourage residents from offering non-FREH STRs because of regulatory burden, but would not have an effect on STR prices, since we assume the latter are exogenously determined by tourism demand, hotel competition, and housing quality. As Table 3 demonstrates, the aggregated group-time ATT estimates for *FREH_log* and *non_FREH_log* are both clearly negative, suggesting that principal-residence restrictions are effective in reducing the supply of both dedicated and casual STRs. By contrast, the ATT estimate for *price_log* cannot be distinguished from zero, suggesting that principal-residence restrictions do not have an effect on STR prices.

7.2 RE-ESF model

The DiD model results reported in section 7.1 demonstrate a causal relationship between STR principal-residence requirements and rent levels which is only plausibly mediated by change in STR activity, and further demonstrate a causal relationship between these regulatory requirements and two dimensions of STR activity (commercial FREH listings and casual non-

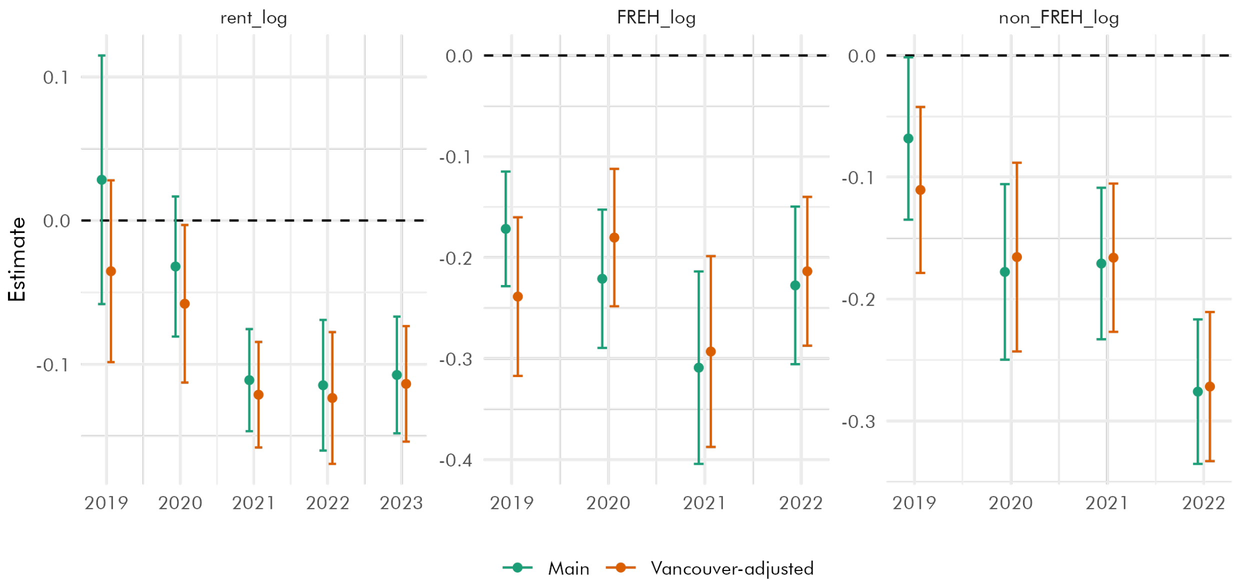

Figure 9. Average treatment effects by time period for the main rent_log, FREH_log *and* non_FREH_log *models, and variants with Vancouver treatment date delayed by one year (point estimates with 95% confidence intervals)*

FREH listings) which represent two possible causal pathways between STR activity and long-term rents: one where STR activity shifts the supply curve for long-term rentals, and one where it shifts the demand curve . However, the DiD models do not allow adjudicating between these two pathways; for example, it could be true that the causal impact of STR regulations on rents is only mediated through the impact of regulations on commercial FREH listings, despite the fact that regulations also cause non-commercial listings to decline). The models also do not allow evaluating our third hypothesized causal pathway between STR activity and long-term rents—rental price stickiness—because the models found that STR regulations did not affect the relevant STR variable (nightly price). Accordingly, we now turn to the results of our RE-ESF regression, which uses *rent_change* as an outcome variable with *FREH_change*, *non_FREH_change* and *price_change* as treatment variables. As discussed above, this regression on its own is vulnerable to endogeneity in the treatment variables. However, the DiD model introduced an exogenous source of variation in the treatment, and thus was able to establish the overall causal relationship between STR activity and long-term rents. We are therefore able to use the RE-ESF model to attempt a more fine-grained analysis of the causal pathways along which STR activity affects long-term rents.

Table 4 presents the RE-ESF regression results. The model demonstrates independent, positive effects of the FREH, non-FREH and STR price treatment variables on the outcome variable *rent_change*. Higher rates of change of any of our measures of STR activity in a neighbourhood predict higher rates of change of rents in the neighbourhood, and this effect holds true independent of both spatial and temporal autocorrelation. However, only the results for *FREH_change* are strongly significant ($p < 0.01$); the results for *price_change* are weakly significant ($p < 0.1$), while the result for *non_FREH_change* displays the hypothesized direction of effect, but with a 90% confidence interval which crosses zero. This result suggests that STR activity has a causal effect on rents primarily through the supply and price-stickiness channels, with a possible weak effect through the demand channel. First of all, commercial FREH STR listings remove housing from the long-term rental market, and thus shift the supply curve left. Second, higher STR prices decrease rent price stickiness by decreasing the risks to landlords of demanding higher rent increases, thus leading to faster rent growth than would have otherwise been expected. Third, non-FREH STR listings possibly increase the amount of rental housing demanded by tenants, and thus may shift the demand curve right, although our model offers less support for this proposition.

Model results demonstrate strong spatiotemporal effects. The temporally autoregressive term *rent_lag_log* has by far the highest parameter estimate among included variables, while the Moran's I term (scaled between 0 and 1) for the spatial effects is 0.535. This indicates a moderately high level of spatial dependence (Griffith, 2003), and validates the decision to use ESF to capture this dependence. Successfully capturing both temporal and spatial autocorrelation

	Estimate	Standard error	Confidence interval (95%)	
(Intercept)	0.000	(0.021)	-0.041	0.041
FREH_change	**0.035****	**(0.013)**	**0.009**	**0.061**
non_FREH_change	**0.018**	**(0.013)**	**-0.008**	**0.044**
price_change	**0.022+**	**(0.013)**	**-0.003**	**0.047**
rent_lag_log	-0.111***	(0.024)	-0.159	-0.063
vacancy_lag_log	-0.041**	(0.015)	-0.070	-0.012
apart_log	0.090***	(0.017)	0.057	0.124
income_log	0.112***	(0.018)	0.076	0.147
Spatial effects (residuals): standard deviation	0.357			
Spatial effects (residuals): scaled Moran's I	0.535			
Random group effects (neighbourhood): standard deviation	0.000			
Random group effects (region-by-year): standard deviation	0.264			
Number of observations	5,400			
Adjusted R2 (conditional)	0.190			
Restricted log likelihood	-7,358			
AIC	14,743			
BIC	14,829			

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; AIC = Akaike information criterion; BIC = Bayesian information criterion

Table 3. RE-ESF regression results for common adjustment set

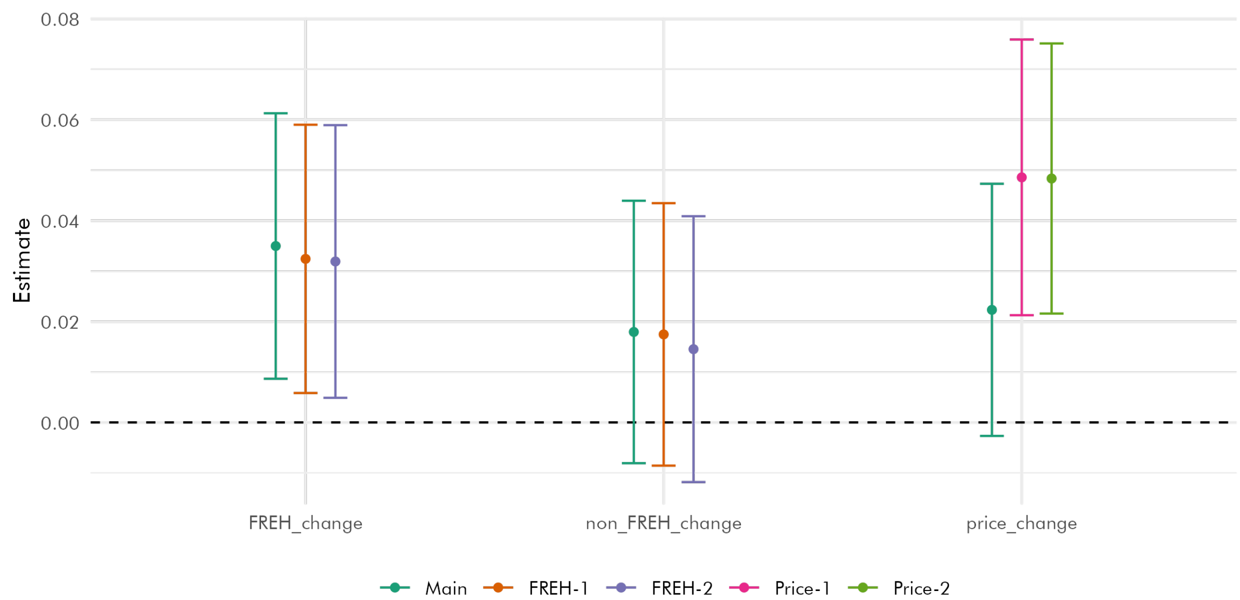

Figure 10. Total causal effect estimates of STR variables on rent_change *for all minimal adjustment sets (point estimates with 95% confidence intervals)*

in the model design increases our confidence that the measured parameter estimates for the three STR activity variables correctly reflect the total causal effect of each of these variables on the neighbourhood change in rent.

Although the results in Table 4 only pertain to a single model variant which includes a common adjustment set of variables for our three treatment variables, formal analysis of the DAG yielded four additional minimal adjustment sets, two of which allow for the direct measurement of total causal impact of *FREH_change* and *non_FREH_change*, and two of which allow for the direct measurement of total causal impact of *price_change*. If our structural causal model is properly specified, any adjustment set should yield the same parameter estimates for the three treatment variables. Figure 10 shows the parameter estimates (with 95% confidence intervals) taken from models run on each of the adjustment sets. (Full model results are reported in the appendix.) The figure makes clear that the parameter estimates for *FREH_change* and *non_FREH_change* are highly reliable: all three adjustment sets produce effectively identical estimates. The estimates for *price_change* are less clear-cut: both alternative adjustment sets produce much higher estimates than the common adjustment set. While the 95% confidence intervals of the *price_change* estimates all overlap (and thus it is possible that the estimates are in fact consistent), it is nevertheless more likely that there is some model misspecification leading to inconsistent results. This inconsistency does not change the qualitative interpretation of the findings, however; under any adjustment set the estimated total causal effect of *price_change* on *rent_change* is positive.

In addition to the conditional independence tests already reported, regression diagnostics and robustness checks (described in the appendix) all support the idea that the RE-ESF model adequately represents the underlying data, and that the treatment variables each have a positive causal impact on year-over-year changes in average rents in Canadian neighbourhoods. In every model variant, the three treatment variables all retain strongly positive and statistically significant total causal effect estimates. While there is still a substantial amount of variation in the outcome variable not explained by the model (adjusted conditional R^2 is 0.190), the purpose of the model is not to account for all variation in *rent_change*, but rather to accurately measure the total causal effect of our treatment variables on *rent_change.* Our conclusion is that, in the case of Canadian urban areas from 2017-2022, our hypothesized causal relationships between STR activity and rents are partially substantiated. That is, changes in dedicated STRs (supply effect) and STR prices (price-stickiness effect) are each independently and positively associated with changes in the level of average rent in a neighbourhood. Casual STRs (demand effect) are possibly independently and positively associated with changes in the level of average rent in a neighbourhood, but this finding cannot be unambiguously substantiated.

7.3. Comparative causal-effect estimates between DiD and RE-ESF models

Our DiD and RE-ESF models offer two complementary strategies for estimating the total causal effects of STR activity on rents. Although the estimates cannot be compared across the entire range of outcomes, since the former uses standardized log of rent as an outcome variable and the latter uses standardized change in rent instead, their effects at mean rent values can be compared.

The aggregated group-time ATTs of STR regulations for the period 2017-2022 (i.e. excluding 2023, when we have rent data but no STR data) are -0.093 for *rent_log*, -0.251 for *FREH_log*, and -0.207 for *non_FREH* (all expressed in standard deviations of the standardized versions of the variables). Put differently, an exogenously caused shock of logged FREH and non-FREH listings of -0.251 and -0.207 standard deviations, respectively, caused an average -0.093 standard-deviation decline in logged average rents, across the entirety of our DiD dataset. At the mean value for average rent in the dataset of $1021, this is a $28.2 decrease. The RE-ESF model, by contrast, estimates that this decline would have only produced a -0.012 standard-deviation decline in average rents, which at the mean value in the dataset would be a $4.1 decrease. In other words, the estimated total causal impact of FREH and non-FREH listings (the two treatment variables which can be estimated from both models) is nearly seven times as high in the DiD model, which controls for any potential endogeneity in the treatment variables, as it is in the RE-ESF model, where endogeneity remains a possibility.

The implication is that any omitted variables in the RE-ESF model are correlated differently with our STR treatment variables and our rent change outcome variable, thus creating negative bias in our measurement of total causal effects. One possible culprit is that our RE-ESF model only includes a single period of temporally lagged variables, and there could be longer-term forms of simultaneity bias at work, whereby higher rents caused by STR activity causes subsequent STR activity to be lower as landlords reallocate units to the long-term market. However, since the evidence from our DiD model suggests that the true causal effect of STR activity on rents is underestimated by the RE-ESF model, we are able to conclude that the parameter estimates generated by the latter represent a conservative lower bound.

8. Discussion

The models provide strong support for the idea that variation in STR activity in Canadian urban areas causes variation in residential rents. We now discuss some of the implications of these findings: the total estimated impact of STRs on rents and policy-relevant counterfactuals.

The top panel of Figure 11 shows the year-to-year change in rent paid in October each year which is accounted for by year-to-year variation in STR activity in the RE-ESF model. (We arrive at these figures by creating a counterfactual scenario for each year where *FREH_change, non_FREH_change and price_change* were all zero for each neighbourhood. We also control for the underlying change in total tenant population in the calculations.) The bottom panel of Figure 11 presents the same results, but expressed cumulatively since 2017. (Here each counterfactual scenario is that *FREH*, *non_FREH* and *price* remained at 2017 levels in each subsequent year.) The figure tells a clear story about the Covid pandemic. In the years prior to the pandemic, an average of 2.6% of the annual increase in total rent paid is caused for by neighbourhood-level variation in STR activity. In 2020, by contrast, this number dropped to -5.5%. In other words, the model suggests that, in a counterfactual scenario where STR activity did not change in those years, rents would have increased by an average of 5.5% more than they actually did. The cumulative figures are similar: the total amount of rent change from a baseline of 2017 that was caused by year-to-year variation in STR activity according to the RE-ESF model was 4.3% before the onset of the pandemic. Thanks to the decline in STR activity during the pandemic, rents increased less than they otherwise would have, so that, by 2022, a counterfactual scenario where STR levels had remained at their 2017 would have seen a 2.8% lower increase in total rents paid in Canadian urban areas.

To contextualize these results, in October 2022, tenants in Canadian urban areas paid approximately \$8.8 billion in monthly rent. (We arrive at this figure by assuming that the average rent for all tenants is the same as the average rent for tenants in purpose-built rentals, since our rent data only covers the latter.) The previous month, hosts on Airbnb in those same urban areas earned \$213 million. Airbnb host revenue was 2.4% the size of total long-term rental revenue. To be clear, this figure is an underestimate of total short-term rental revenue: all non-Airbnb STR platforms (and their associated revenue) are excluded from our dataset. However, as an order-of-magnitude estimate, the notion that a residential land use responsible for approximately 2.4% of total residential rental revenue would have a meaningful impact on rents in the remaining 97.6% of the market is plausible on its face, to say nothing of the supply effects of tens of thousands of rental units being withdrawn from the long-term market.

A complementary estimate of the total impact of STRs on rents in Canada can be obtained from the DiD model. Among the 309 neighbourhoods subject to principal-residence restrictions, in the year after the regulations took effect, rents were on average \$24 lower than they would have been in the absence of the regulations—a 1.7% difference. This rent decrease was driven by a 16.4% decrease in FREH listings compared to the no-regulation counterfactual, and an 7.5% decrease in non-FREH listings. The effects of principal-residence restrictions continue to grow in the years following treatment (as indicated in Figure 6, above). As a result, in 2022 neighbourhoods with

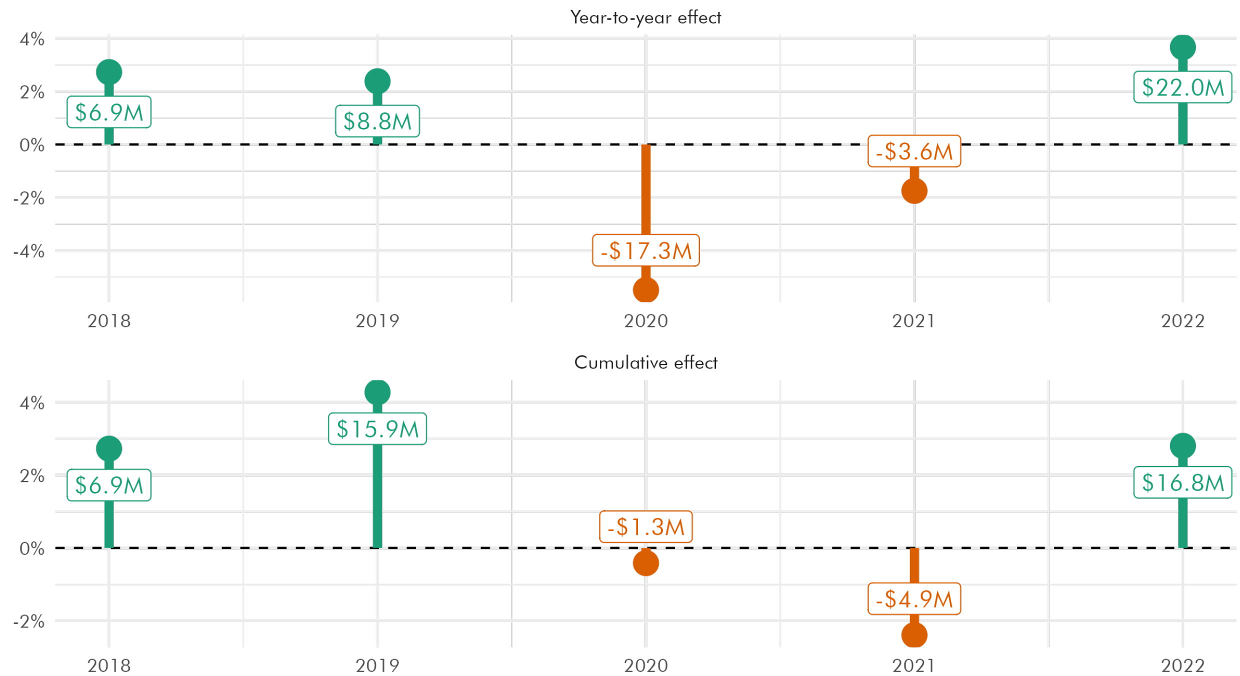

Figure 11. Year-to-year and cumulative estimated effects of STRs on change in total rent paid in Canadian urban areas. Y-axis position is the percentage of total rent change caused by STR variation in the RE-ESF model, and the label is the amount of monthly rent change caused by STR variation in the model.

principal-residence restrictions had rents which were $55 (3.5%) less than they would have been without those restrictions in place, with 25.7% fewer FREH listings and 24.4% fewer non-FREH listings. By contrast, among the neighbourhoods without principal-residence restrictions in 2022, rents could plausibly have been $19 (1.6%) lower than they actually were if these places had had such restrictions in place, based on the average treatment effect experienced in the first year of implementation by jurisdictions which did implement principal-residence restrictions. This would have amounted to $74.5 million in monthly rent payments saved in these areas—an amount which would be expected to increase in subsequent years.

A final consideration arising from the research concerns our strategy of decomposing the STR effect on rents into supply (*FREH_change*), demand (*non_FREH_change*) and price-stickiness (*price_change*) channels. This allows us to apply the results of the models to policy considerations in a more direct way than might otherwise be possible. Table 5 shows, for each of the major regions of Canada, the estimated total causal effect of the different components of STR activity on 2022 rent change. It indicates, first of all, that even despite the substantial number of municipalities which had already imposed principal-residence restrictions on STR operators, that FREH listings continue to be a major source of upward pressure on rents, and that jurisdictions that have not yet implemented such restrictions could expect to see meaningful housing affordability benefits from doing so, since the FREH contribution to rent increases is now concentrated among cities lacking restrictions. Second, it indicates that non-FREH listings are not a meaningful source of upward pressure on rents in most parts of the country, and thus that specifically restricting home sharing STRs would not generally be expected to be an effective pro-affordability housing policy. Finally, it indicates that rising STR prices have become the second most important STR contributor to rising rents in Canadian cities, likely via the price-stickiness channel we have identified. This channel is less amenable to direct regulation. However, although this possibility was not explicitly captured in our structural causal model, it is likely that the actual impact of STR prices on rent price stickiness would diminish in a scenario where non-principal-residence STRs were broadly banned, since landlords would not easily be able to avail themselves of the STR alternative if their tenants refused a rent increase. In sum, the evidence presented here supports the idea that principal-residence restrictions remain the most promising avenue for pro-housing-affordability regulation of STR markets.

Region	Total rent increase	Estimated FREH contribution	Estimated non-FREH contribution	Estimated price contribution
Prairies	$60.0M	$1.6M (2.7%)	$0.7M (1.2%)	$0.6M (1.0%)
British Columbia	$130.3M	$2.6M (2.0%)	$1.0M (0.7%)	$1.3M (1.0%)
Atlantic	$27.0M	$0.9M (3.3%)	$0.3M (1.2%)	$0.4M (1.5%)
Ontario	$234.1M	$2.8M (1.2%)	$0.9M (0.4%)	$3.6M (1.6%)
Quebec	$149.0M	$3.0M (2.0%)	$0.3M (0.2%)	$2.0M (1.3%)

Table 5. Regional estimates of the amount of total 2022 rent increases caused by changes in FREH listings, non-FREH listings, and STR prices

9. Conclusion

Over the last decade, an increasingly loud chorus of housing activists, community groups, and local governments have sounded the alarm about the negative effects of short-term rentals on housing availability and affordability. Are these concerns warranted? Our findings suggest that they are. In this paper we have measured the relationship between STR activity and residential rents in all Canadian urban areas from 2017-2022, leveraging a pair of complementary causal models and a series of data and methodological advances which allow us greater traction into this issue than previous research. Using a time-varying difference-in-differences regression, we were able to exploit a source of exogenous variation in STR activity (the introduction of STR principal-residence restrictions) to establish that STR activity causes long-term rents to increase. Using a structural causal model whose formal properties were validated using conditional independence testing, we further modelled the hypothesis that STR activity causes rents to change via separate supply, demand, and price-stickiness channels. A random-effects eigenvector spatial filtering regression found varying levels of support for the components of the hypothesis: higher rates of change of FREH listings and STR prices each independently and unambiguously cause higher rates of residential rent change, while higher rates of change of non-FREH listings have a directionally consistent effect, but under the threshold for statistical significance. The implication is that STRs affect rental prices in the long-term rental market by shifting the supply curve left (as housing units are shifted from long-term to short-term uses by commercial STR operators) and by decreasing price stickiness (as landlords are more willing to raise long-term rents if they have viable STR options as a fallback strategy). It is also possible, although unambiguously demonstrated by the results, that STRs affect rental prices by shifting the demand curve right (as households demand more housing in order to run home-sharing STRs out of their principal residences).

Our results suggest that the recent trend in North American cities towards imposing principal-residence restrictions on STR operators has already been effective at exerting downward pressure on rents, and that such restrictions remain the most promising avenue for pro-housing-affordability regulation of STR markets in the future.

References

Ankan, A., Wortel, I. M. N., & Textor, J. (2021). Testing graphical causal models using the R package "dagitty". *Current Protocols*, 1, e45. 10.1002/cpz1.45

Anselin, L., & Arribas-Bel, D. (2013). Spatial fixed effects and spatial dependence in a single cross-section. *Papers in Regional Science*, 92(1), 3-18.

Ayouba, K., Breuillé, A., Grivault, C., & Le Gallo, J. (2020). Does Airbnb disrupt the private rental market? An empirical analysis for French cities. *International Regional Science Review*, 43(1-2), 76-104. 10.1177/0160017618821428

Barron, K., Kung, E., & Proserpio, D. (2021). The effect of home-sharing on house prices and rents: Evidence from Airbnb. *Marketing Science*, 40(1), 23–47.

Callaway, B. & Sant'Anna, P.H.C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200-230. 10.1016/j.jeconom.2020.12.001

Canada Mortgage and Housing Corporation. (2023). Rental market survey. [Provided under End User Licence Agreement: This information is reproduced and distributed on an "as is" basis with the permission of CMHC, with no representation or warranty as the quality or accuracy of the information]

Celata, F., & Romano, A. (2022). Overtourism and online short-term rental platforms in Italian cities. *Journal of Sustainable Tourism*, 30(5), 1020-1039. 10.1080/09669582.2020.1788568

Chang, H. (2020). Does the room sharing business model disrupt housing markets? Empirical evidence of Airbnb in Taiwan. *Journal of Housing Economics*, 49(101706). https://doi.org/10.1016/j.jhe.2020.101706

Cheung, K. S., & Yiu, C. Y. (2023). The paradox of Airbnb, crime and house prices: A reconciliation. *Tourism Economics*, 29(5), 1412-1418. https://doi.org/10.1177/13548166221102808

Combs, J., Kerrigan, D., & Wachsmuth, D. (2020). Short-term rentals in Canada: Uneven growth, uneven impacts. *Canadian Journal of Urban Research*, 29(1), 119-134. https://cjur.uwinnipeg.ca/index.php/cjur/article/view/274

Deboosere, R., Kerrigan, D. J., Wachsmuth, D., & El-Geneidy, A. (2019). Location, location and professionalization: a multilevel hedonic analysis of Airbnb listing prices and revenue. *Regional Studies, Regional Science*, 6(1), 143-156.

Duso, T., Michelsen, C., Schäfer, M., & Tran, K. D. (2024). Airbnb and rental markets: Evidence from Berlin. *Regional Science and Urban Economics*, 106, 104007.

Franco, S. F., & Santos, C. D. (2021). The impact of Airbnb on residential property values and rents: Evidence from Portugal. *Regional Science and Urban Economics*, 88, 103667.

Garay-Tamajón, L., Lladós-Masllorens, J., Meseguer-Artola, A., & Morales-Pérez, S. (2022). Analyzing the influence of short-term rental platforms on housing affordability in global urban destination neighborhoods. *Tourism and Hospitality Research*, 22(4), 444-461. https://doi.org/10.1177/14673584211057568

Garcia-López, M., Jofre-Monseny, J., Martínez-Mazza, R., & Segú, M. (2020). Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona. *Journal of Urban Economics*, 119(103278). https://doi.org/10.1016/j.jue.2020.103278

Gibbons, S., & Overman, H. G. (2012). Mostly pointless spatial econometrics?. J*ournal of Regional Science*, 52(2), 172-191.

Griffith, D. A. (2003). *Spatial Autocorrelation and Spatial Filtering: Advances in Spatial Science*. Springer, Berlin, Heidelberg.

Gurran, N., & Phibbs, P. (2017). When tourists move in: How should urban planners respond to Airbnb? *Journal of the American Planning Association*, 83(1), 80-92. 10.1080/01944363.2016.1249011

Guttentag, D. (2015). Airbnb: disruptive innovation and the rise of an informal tourism accommodation sector. *Current Issues in Tourism*, 18(12), 1192-1217. 10.1080/13683500.2013.827159

Hill, R. J., Pfeifer, N., Steurer, M. (2023). The Airbnb rent premium and the crowding-out of long-term rentals. *Journal of Housing Economics*, 61(101935), 1-13. https://doi.org/10.1016/j.jhe.2023.101935

Horn, K., & Merante, M. (2017). Is home sharing driving up rents? Evidence from Airbnb in Boston. *Journal of Housing Economics*, 38, 14-24. https://doi.org/10.1016/j.jhe.2017.08.002

Hughes J., & Haran, M. (2013). Dimension reduction and alleviation of confounding for spatial generalized linear mixed models. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, 75(1), 139–159.

Imbens, G. W. (2020). Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, 58(4), 1129-1179.

Koster, H., van Ommeren, J., & Volhausen, N. (2021). Short-term rentals and the housing market: Quasi-experimental evidence from Airbnb in Los Angeles. *Journal of Urban Economics*, 124(103356). https://doi.org/10.1016/j.jue.2021.103356

Lee, S., & Kim, H. (2023). Four shades of Airbnb and its impacts on locals: A spatiotemporal analysis of Airbnb, rent, housing prices, and gentrification. *Tourism Management Perspectives*, 49(101192). https://doi.org/10.1016/j.tmp.2023.101192

Liang, C., Yeung, M. C. H., & Au, A. K. M. (2022). The impact of Airbnb on housing affordability: Evidence from Hong Kong. *Environment and Planning B: Urban Analytics and City Science*, 49(3), 1048-1066. https://doi.org/10.1177/23998083211043123

Marrone, M., & Peterlongo, G. (2020). Where platforms meet infrastructures: Digital platforms, urban resistance and the ambivalence of the city in the Italian case of Bologna. Work Organisation, *Labour & Globalisation*, 14(1), 119-135. 10.13169/workorgalaboglob.14.1.0119

Miller, C. W. (1 April 2024). The great Airbnb crackdown. *Macleans*. Online: https://macleans.ca/longforms/the-great-airbnb-crackdown/

Morales-Pérez, S., Garay, L., & Wilson, J. (2022). Airbnb's contribution to socio-spatial inequalities and geographies of resistance in Barcelona. *Tourism Geographies*, 24(6-7), 978-1001. 10.1080/14616688.2020.1795712

Murakami, D. (2017). spmoran (ver. 0.2.0): An R package for Moran eigenvector-based scalable spatial additive mixed modeling. https://cran.r-project.org/web/packages/spmoran/spmoran.pdf

Murakami, D., & Griffith, D. A. (2015). Random effects specifications in eigenvector spatial filtering: a simulation study. *Journal of Geographical Systems*, 17, 311-331.

Nieuwland, S., & van Melik, R. (2020). Regulating Airbnb: How cities deal with perceived negative externalities of short-term rentals. *Current Issues in Tourism*, 23(7), 811-825. 10.1080/13683500.2018.1504899

Ram, Y., & Tchetchik, A. (2022). Complementary or competitive? Interrelationships between hotels, Airbnb and housing in Tel Aviv, Israel. *Current Issues in Tourism*, 25(22), 3579-3590. 10.1080/13683500.2021.1978954

Rodríguez-Pérez de Arenaza, D., Hierro, L. A., & Patiño, D. (2022). Airbnb, sun-and-beach tourism and residential rental prices. The case of the coast of Andalusia (Spain). *Current Issues in Tourism*, 25(20), 3261-3278. 10.1080/13683500.2019.1705768

Shabrina, Z., Arcaute, E., & Batty, M. (2022). Airbnb and its potential impact on the London housing market. *Urban Studies*, 59(1), 197-221. https://doi.org/10.1177/0042098020970865

Simcock, T. (2023). Home or hotel? A contemporary challenge in the use of housing stock. *Housing Studies*, 38(9), 1760-1776. 10.1080/02673037.2021.1988063

Smigiel, C. (2020). Why did it not work? Reflections on regulating Airbnb and the complexity and agency of platform capitalism. *Geographica Helvetica*, 75(3), 253-257. https://doi.org/10.5194/gh-75-253-2020

Spangler, I. (2019). Hidden value in the platform's platform: Airbnb, displacement, and the un-homing spatialities of emotional labour. *Transactions of the Institute of British Geographers*, 45(3), 575-588. https://doi.org/10.1111/tran.12367

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175-199.

Sun, Y., Wang, S., Xie, J., & Hu, X. (2021). Modeling local-scale violent crime rate: A comparison of eigenvector spatial filtering models and conventional spatial regression models. *The Professional Geographer*, 73(2), 312-321.

Textor, J., van der Zander, B., Gilthorpe, M.K., Liskiewicz, M., & Ellison, G.T.H. (2016). Robust causal inference using directed acyclic graphs: the R package 'dagitty'. *International Journal of Epidemiology*, 45(6), 1887-1894.

Thoemmes, F., Rosseel, Y., & Textor, J. (2018). Local fit evaluation of structural equation models using graphical criteria. *Psychological Methods*, 23(1), 27–41.

Wachsmuth, D., & Weisler, A. (2018). Airbnb and the rent gap: Gentrification through the sharing economy. *Environment & Planning A: Economy and Space*, 50(6), 1147-1170. https://doi.org/10.1177/0308518X18778038

Wachsmuth, D. (2021). strr. https://github.com/UPGo-McGill/strr

Wachsmuth, D., & Buglioni, B. (2024). Neither housing nor hotel: The emergence of "medium-term rentals" in post-Covid Canadian cities. *Canadian Planning and Policy*, 2024(1), 68–89. https://doi.org/10.24908/cpp-apc.v2024i1.16935

Westreich, D., & Greenland, S. (2013). The table 2 fallacy: presenting and interpreting confounder and modifier coefficients. *American Journal of Epidemiology*, 177(4), 292-298.

Zervas, G., Proserpio, D., & Byers, J. W. (2017). The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry. *Journal of Marketing Research*, 54(5), 687-705. https://doi.org/10.1509/jmr.15.0204

Appendix. Additional methodological details, robustness checks and diagnostics

Principal-residence regulations dataset

The treatment in our difference-in-differences model is the enactment of a principal-residence requirement for an STR—i.e. the requirement that the dwelling unit being rented on an STR platform is the principal residence of the STR host. We specifically define treatment in a given year for a neighbourhood as follows: was any portion of the neighbourhood subject to an in-force principal-residence requirement as of January 1? If a jurisdiction has passed a law restricting STRs to a host's principal residence but the law had not taken effect as of January 1, the neighbourhood is classified as untreated.

Short-term rentals restrictions are typically regulated through zoning bylaws, often also with some form of registration such as a business license or a STR operator license. We considered a municipality to have enacted a STR-specific principal-residence restriction if a bylaw amendment was passed that specifically targeted short-term rentals and provided for some form of licensing to monitor STRs. For example, a municipality that has never allowed any form of short-term renting except bed & breakfast establishments (which they define as a principal-residence tourist accommodation business), and which has not updated its bylaw since the introduction of Airbnb would not be considered to have a principal-residence restriction. A municipality that outright bans STRs but does not monitor STRs nor levies fines against illegal STRs on its territory would also not be considered to have a principal-residence restriction either.

Table A1 lists all the municipalities covered by our main dataset (i.e. cities of at least 10,000 inhabitants) which adopted a principal-residence restriction on or before January 1, 2023. Because STR rules are a borough-level responsibility in Montreal (boroughs are a sub-municipal government in Montreal), the Montreal entries additionally indicate the borough which adopted the regulation.

Difference-in-differences robustness checks

The main version of our DiD dataset has two significant restrictions. First, it only includes neighbourhoods in the four provinces where any municipality had enacted principal-residence restrictions on or before January 1, 2023: British Columbia, Ontario, New Brunswick, and Quebec. These four provinces have 77.1% of the total dwelling units in Canada as of the 2021 census, so in practice this means that we include the large majority of neighbourhoods even with this restriction. As we discuss below, including neighbourhoods in provinces without any principal-residence restrictions leads to a violation of the parallel trends assumption. Second, our main dataset excludes neighbourhoods that had a treatment date before January 1, 2017, since there are only three such neighbourhoods out of the 1,080 total neighbourhoods, and this quantity is insufficient to calculate reliable group-time ATTs. However, we tested a variant of the DiD with all neighbourhoods included as a robustness check (model *all)*; we also tested two other

Municipality	Province	Date	Municipality	Province	Date
Baie-Comeau	Quebec	2022-09-01	Nelson	British Columbia	2017-01-01
Brampton	Ontario	2022-09-30	North Bay	Ontario	2023-01-01
Brant	Ontario	2022-03-31	Oakville	Ontario	2018-11-01
Brantford	Ontario	2022-02-22	Oshawa	Ontario	2020-09-30
Brossard	Quebec	2022-09-01	Ottawa	Ontario	2022-04-01
Burnaby	British Columbia	2021-06-01	Ottawa	Quebec	2022-04-01
Chilliwack	British Columbia	2021-10-19	Penetanguishene	Ontario	2022-09-14
Fredericton	New Brunswick	2021-03-22	Pitt Meadows	British Columbia	2011-10-04
Georgina	Ontario	2019-10-09	Prince Edward County	Ontario	2022-09-20
La Pêche	Quebec	2021-10-05	Québec	Quebec	2019-09-04
London	Ontario	2022-10-01	Rimouski	Quebec	2021-03-11
Magog	Quebec	2021-05-03	Sainte-Agathe-des-Monts	Quebec	2022-12-20
Milton	Ontario	2022-07-15	Sarnia	Ontario	2020-02-10
Mirabel	Quebec	2021-11-01	Sherbrooke	Quebec	2022-03-25
Mississauga	Ontario	2021-01-19	Squamish	British Columbia	2021-04-01
Montréal (Hochelaga-Maisonneuve)	Quebec	2016-12-31	St. Catharines	Ontario	2022-04-01
Montréal (Ville-Marie)	Quebec	2018-06-12	Summerland	British Columbia	2022-09-06
Montréal (Plateau-Mont-Royal)	Quebec	2019-01-01	Thompson-Nicola J	British Columbia	2012-01-01
Montréal (Sud-Ouest)	Quebec	2019-02-11	Toronto	Ontario	2021-01-01
Montréal (Saint-Laurent)	Quebec	2019-08-01	Vancouver	British Columbia	2018-08-31
Montréal (Rosemont/La Petite-Patrie)	Quebec	2020-02-10	Victoria	British Columbia	2018-03-08
Nanaimo	British Columbia	2022-02-07	Windsor	Ontario	2022-09-06

Table A1. STR principal-residence requirements by location and date of implementation

variants: one with no imputed rent values and hence a number of missing observations (model *no_imp*) and one with the treatment date for the City of Vancouver shifted one year from August 31, 2018 to August 31, 2019 to account for the slow roll out of effective enforcement in that municipality (model *van*).

Table A2 shows the aggregated group-time ATTs for all model variants; while the key *rent_log* estimate is lower for the variant with all neighbourhoods included, the qualitative interpretation of the causal relationship between treatment and both rents and STR activity is consistent across all variants.

Model	***rent_log***	***FREH_log***	***non_FREH_log***	***price_log***
main	-0.096***	-0.251***	-0.207***	-0.034
	(0.015)	(0.032)	(0.021)	(0.036)
no_imp	-0.108***	-0.265***	-0.211***	-0.019
	(0.019)	(0.030)	(0.025)	(0.035)
all	-0.037**	-0.265***	-0.238***	-0.029
	(0.015)	(0.034)	(0.021)	(0.032)
van	-0.111***	-0.240***	-0.210***	0.004
	(0.016)	(0.030)	(0.023)	(0.034)

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Table A2. Parameter estimates for aggregated group-time ATTs for all variants of four DiD models with STR principal-residence restrictions as treatment

Difference-in-differences diagnostics

The key assumption required for a DiD model to be valid is the parallel trends assumption, i.e. that, in the absence of treatment, the difference between the treatment and control groups would be constant over time. In a conventional DiD approach, an event-study regression is generally used to test this assumption. In the time-variant DiD approach we have used, the parallel trends assumption can instead be tested with dynamically aggregated group-time treatment effects, where an average effect is computed for each length of treatment exposure (Callaway & Sant'Anna, 2021). If the parallel trends assumption holds, then treatment effects should be zero for all negative exposure lengths (i.e. treated cases have no treatment effect prior to treatment). Figure 6, above, showed that our main dataset fulfills the parallel trends assumption. Figure A1 shows dynamically aggregated group-time ATTs for all four model variants, and shows that the parallel trends assumption is confirmed for the *no_imp* and *van* variants, but is disproven for the *all* variant, since in the latter case most of the average treatment effects for negative exposure lengths are statistically distinguishable from zero. The implication is that treated jurisdictions were *not* on the same trajectory as untreated jurisdictions prior to their treatment, and this fact justifies the decision to exclude observations outside of the provinces where principal-residence restrictions were implemented.

Structural causal model

Figure 7, above, shows a DAG of our structural causal model of *rent_change* with the causal path connecting *FREH_change* to *rent_change* identified. Figure A2 shows variants of the DAG in Figure 7 but with *non_FREH_change* and *price_change* as the identified treatment variables.

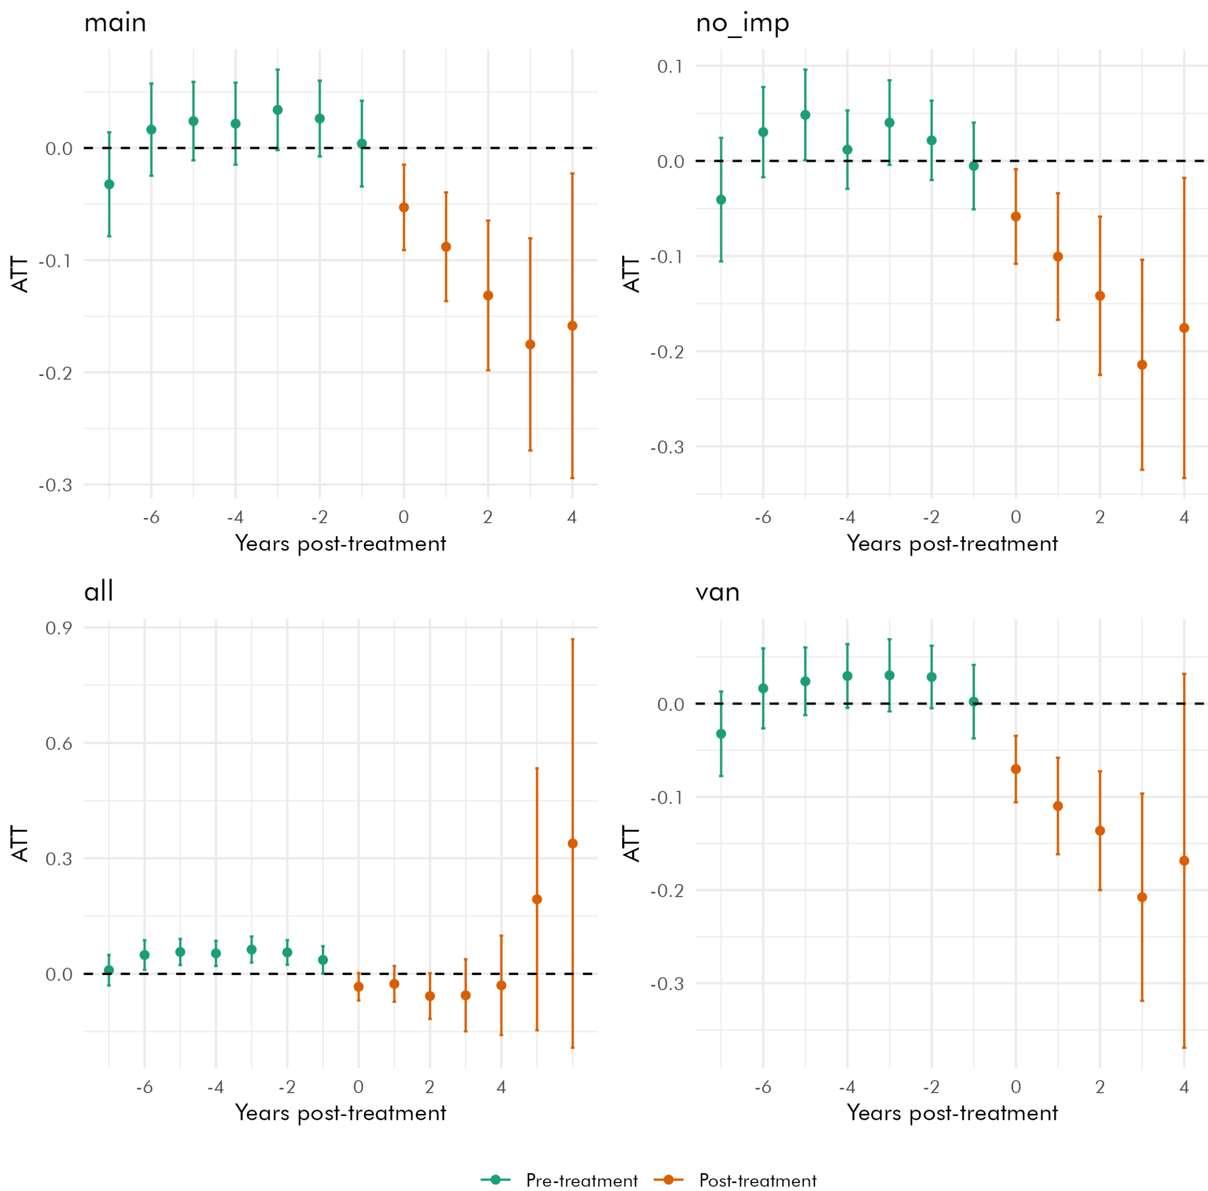

Figure A7. Average treatment effect of STR principal-residence regulations on rent_log *by length of exposure for all model variants (point estimates with 95% confidence intervals)*

These three DAGs were formally analyzed to generate the minimal adjustment sets for measuring the total causal effect of the three treatment variables, and the underlying structural causal model was analyzed to determine the implied conditional independences which could be empirically tested to determine the validity of the model.

Conditional independence tests were conducted using linear regression following the procedure in Ankan et al. (2021). For a condition $X \perp Y \mid Z$, the residuals for $Z \sim X$ and $Z \sim Y$ are calculated, and the Pearson correlation coefficient of the residuals is calculated. The null

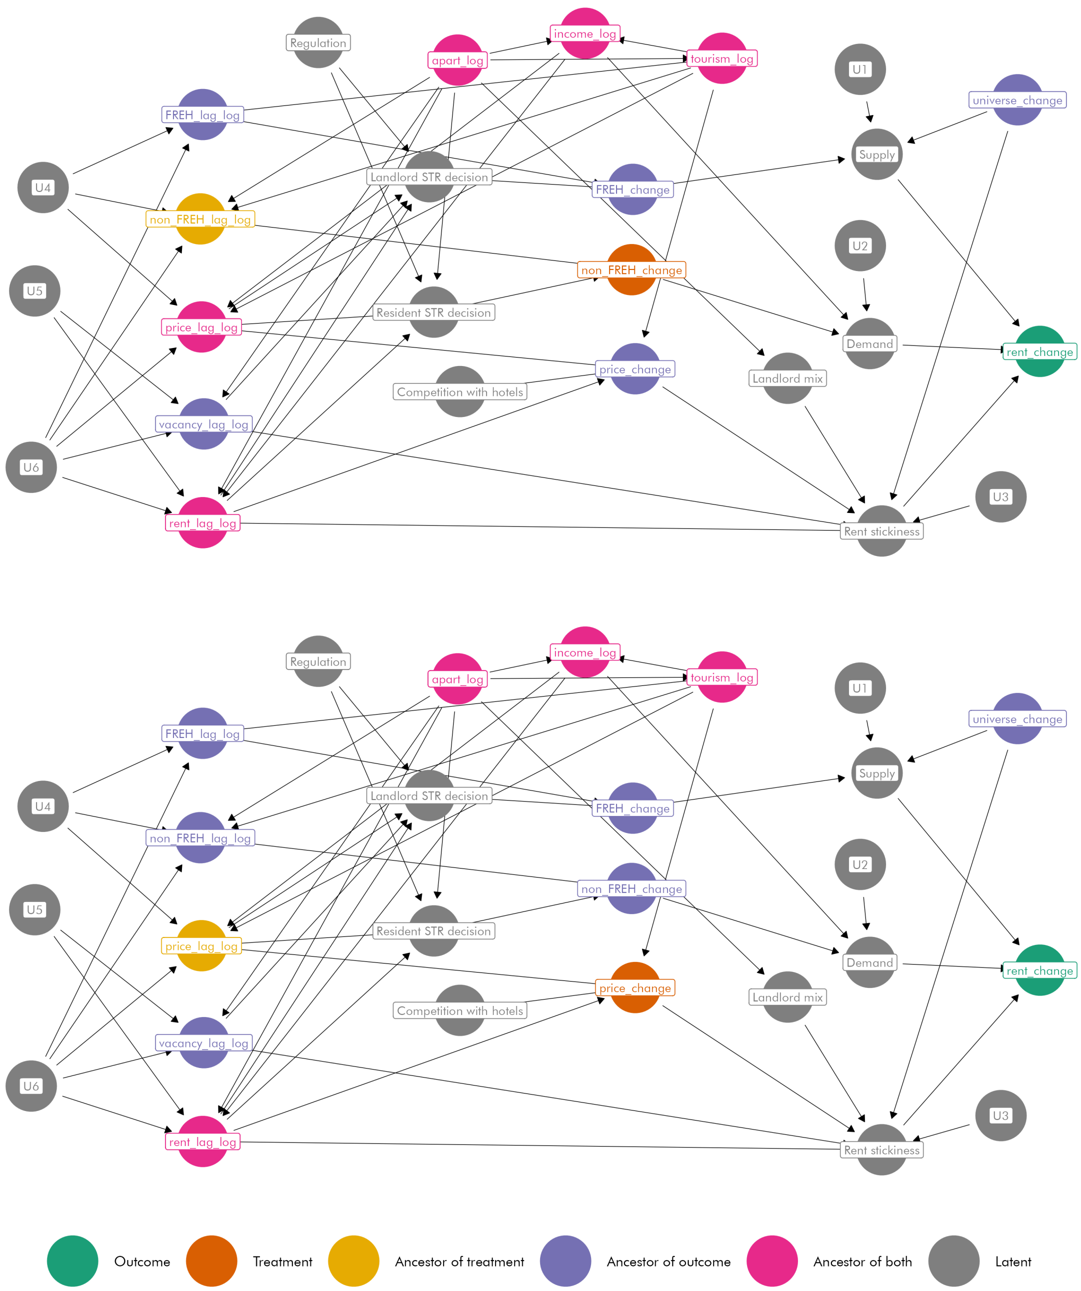

Figure A2. DAG showing hypothesized causal pathways between non_FREH_change *and* rent_change *(top panel) and between* price_change *and* rent_change *(bottom panel)*

hypothesis is that the correlation coefficient is zero, and hence can be rejected if the confidence interval of the estimate does not overlap zero. Unlike most statistical tests, however, a "success" implies failing to reject the null hypothesis, since we seek to establish that the conditional independences implied by the hypothesized causal relationships in the structural model are empirically observable. Table A3 displays the detailed results for all the conditional

Variables		Conditioning variables	Estimate	Confidence interval (95%)	
FREH_change	*universe_change*	-	-0.002	-0.029	0.024
FREH_lag_log	*apart_log*	*tourism_log*	0.098***	0.071	0.124
FREH_lag_log	*income_log*	*tourism_log*	0.224***	0.198	0.249
FREH_lag_log	*universe_change*	-	-0.006	-0.033	0.020
apart_log	*price_lag_log*	*income_log, tourism_log*	-0.026+	-0.053	0.001
apart_log	*universe_change*	-	0.010	-0.017	0.037
income_log	*non_FREH_lag_log*	*apart_log, tourism_log*	0.257***	0.232	0.282
income_log	*universe_change*	-	-0.002	-0.028	0.025
income_log	*vacancy_lag_log*	*apart_log*	-0.024+	-0.051	0.002
non_FREH_change	*universe_change*	-	-0.026+	-0.053	0.001
non_FREH_lag_log	*universe_change*	-	-0.016	-0.043	0.011
price_change	*universe_change*	-	0.005	-0.022	0.031
price_lag_log	*universe_change*	-	-0.008	-0.035	0.018
rent_lag_log	*tourism_log*	*apart_log, income_log*	0.282***	0.257	0.306
rent_lag_log	*universe_change*	-	-0.007	-0.034	0.019
tourism_log	*universe_change*	-	0.006	-0.020	0.033
tourism_log	*vacancy_lag_log*	*apart_log*	0.141***	0.114	0.167
universe_change	*vacancy_lag_log*	-	0.028*	0.001	0.054

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A3. Conditional independence tests for structural causal model.

independence tests, and demonstrates that most variable combinations demonstrate strict independence, and the ones that do not have low correlations. (This is the data underlying Figure 8.)

Robustness checks for rent_change model

Table A4 displays the results for a series of RE-ESF models run on variants of our main model. The variants are as follows:

- Main: The main model discussed in the paper.

- FREH-1: A variant run on the same dataset as the main model, but with a different minimal adjustment set which should be sufficient to measure the total causal impact of *FREH_change* and *non_FREH_change* on *rent_change*, using the following covariates: *FREH_change*, *non_FREH_change*, *rent_lag_log*, *price_lag_log*, *vacancy_lag_log*, *apart_log*, *income_log*, *tourism_log*, *price_lag_dummy*.

	Main	FREH-1	FREH-2	Price-1	Price-2	Non-Gauss	No-imp	No-vac
(Intercept)	0.000 (0.021)	0.001 (0.022)	-0.001 (0.023)	-0.003 (0.023)	-0.002 (0.022)	0.000 (0.022)	-0.003 (0.021)	0.001 (0.021)
FREH_change	0.035** (0.013)	0.032* (0.014)	0.032* (0.014)	-	-	0.028* (0.013)	0.027+ (0.015)	0.035* (0.014)
non_FREH_change	0.018 (0.013)	0.017 (0.013)	0.015 (0.013)	-	-	0.016 (0.013)	0.028+ (0.015)	0.027* (0.014)
price_change	0.022+ (0.013)	-	-	0.049*** (0.014)	0.048*** (0.014)	0.023+ (0.013)	0.009 (0.014)	0.017 (0.013)
rent_lag_log	-0.111*** (0.024)	-0.115*** (0.024)	-0.049* (0.022)	-0.123*** (0.024)	-0.058** (0.022)	-0.136*** (0.024)	-0.070** (0.026)	-0.119*** (0.025)
vacancy_lag_log	-0.041** (0.015)	-0.040** (0.015)	-0.038* (0.015)	-0.039** (0.015)	-	-0.029* (0.015)	-0.078*** (0.018)	-
apart_log	0.090*** (0.017)	0.093*** (0.018)	0.043** (0.016)	0.096*** (0.018)	-	0.096*** (0.017)	0.077*** (0.019)	0.060*** (0.017)
income_log	0.112*** (0.018)	0.103*** (0.020)	-	0.099*** (0.018)	-	0.112*** (0.018)	0.085*** (0.020)	0.100*** (0.019)
price_lag_log	-	0.042* (0.020)	0.063** (0.021)	0.077*** (0.022)	0.079*** (0.020)	-	-	-
tourism_log	-	-0.004 (0.018)	-	-	-0.008 (0.015)	-	-	-
price_lag_dummy	-	-0.011 (0.083)	0.060 (0.105)	0.031 (0.105)	0.035 (0.081)	-	-	-
FREH_lag_log	-	-	0.003 (0.029)	0.007 (0.028)	-	-	-	-
non_FREH_lag_log	-	-	-0.016 (0.028)	-0.023 (0.027)	-	-	-	-
FREH_lag_dummy	-	-	0.031 (0.062)	0.052 (0.062)	-	-	-	-
non_FREH_lag_du mmy	-	-	-0.099 (0.093)	-0.090 (0.092)	-	-	-	-
Spat. eff. (SD)	0.357	0.351	0.342	0.344	0.354	0.413	0.299	0.344
Spat. eff. (Moran's I)	0.535	0.534	0.513	0.529	0.494	0.524	0.523	0.430
RE (nbhd)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RE (region-by-year)	0.264	0.265	0.262	0.263	0.263	0.300	0.245	0.261
N. obs.	5,400	5,400	5,400	5,400	5,400	5,400	4,291	5,043
Adj. R2 (cond.)	0.190	0.191	0.185	0.191	0.187	0.225	0.188	0.185
R. Log. Lik.	-7,358	-7,360	-7,378	-7,360	-7,365	-6,406	-5,849	-6,885
AIC	14,743	14,750	14,790	14,754	14,752	12,851	11,723	13,793
BIC	14,829	14,849	14,902	14,866	14,824	12,976	11,806	13,872

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; AIC = Akaike information criterion; BIC = Bayesian information criterion

Table A4. Additional variants for rent_change model

- FREH-2: A variant run on the same dataset as the main model, but with a different minimal adjustment set which should be sufficient to measure the total causal impact of *FREH_change* and *non_FREH_change* on *rent_change*, using the following covariates: *FREH_change*, *non_FREH_change*, *rent_lag_log*, *FREH_lag_log*, *non_FREH_lag_log*, *price_lag_log*, *vacancy_lag_log*, *apart_log*, *FREH_lag_dummy*, *non_FREH_lag_dummy*, *price_lag_dummy*.

- Price-1: A variant run on the same dataset as the main model, but with a different minimal adjustment set which should be sufficient to measure the total causal impact of *price_change* on *rent_change*, using the following covariates: *price_change*, *rent_lag_log*, *FREH_lag_log*, *non_FREH_lag_log*, *price_lag_log*, *vacancy_lag_log*, *apart_log*, *income_log*, *FREH_lag_dummy*, *non_FREH_lag_dummy*, *price_lag_dummy*.

- Price-2: A variant run on the same dataset as the main model, but with a different minimal adjustment set which should be sufficient to measure the total causal impact of *price_change* on *rent_change*, using the following covariates: *price_change*, *rent_lag_log*, *price_lag_log*, *tourism_log*, *price_lag_dummy*.

- Non-Gauss: A variant run on the same dataset as the main model, and with the same adjustment set, but with the outcome variable transformed using two iterations of SAL ("sinh-arcsinh and affine linear") transformations (Murakami et al., 2021) to address potential non-normality of residuals in the main model.

- No-imp: A variant run with the same adjustment set as the main model, but on a version of the dataset with no imputation of missing values—hence, an unbalanced panel.

- No-vac: A variant run with the adjustment set from the main model but with *vacancy_lag_log* dropped in order to avoid the large number of observations with missing values for *vacancy_lag_log*, and run on a version of the dataset with no imputation of missing values—hence, an unbalanced panel.

Many of these models include lagged point-in-time versions of our treatment variables. To address non-normalcy in these variables, we apply a log transformation to each. Zero values were shifted to the smallest non-zero value present in the variable prior to the log transformation, and additional *FREH_lag_dummy*, *non_FREH_lag_dummy*, and *price_lag_dummy* boolean dummy variables were calculated which are TRUE if the non-log value was zero and FALSE otherwise.

The results in Table A4 demonstrate consistency between model variants. With the exception of the discrepancy between the three adjustment sets for *price_change* noted in the main paper, point estimates for total causal effects are close to each other, and all point estimates for a given treatment variable lie within the 95% confidence interval of all other point estimates. Figure A3 shows the point estimates for the total causal effects across all model variants.

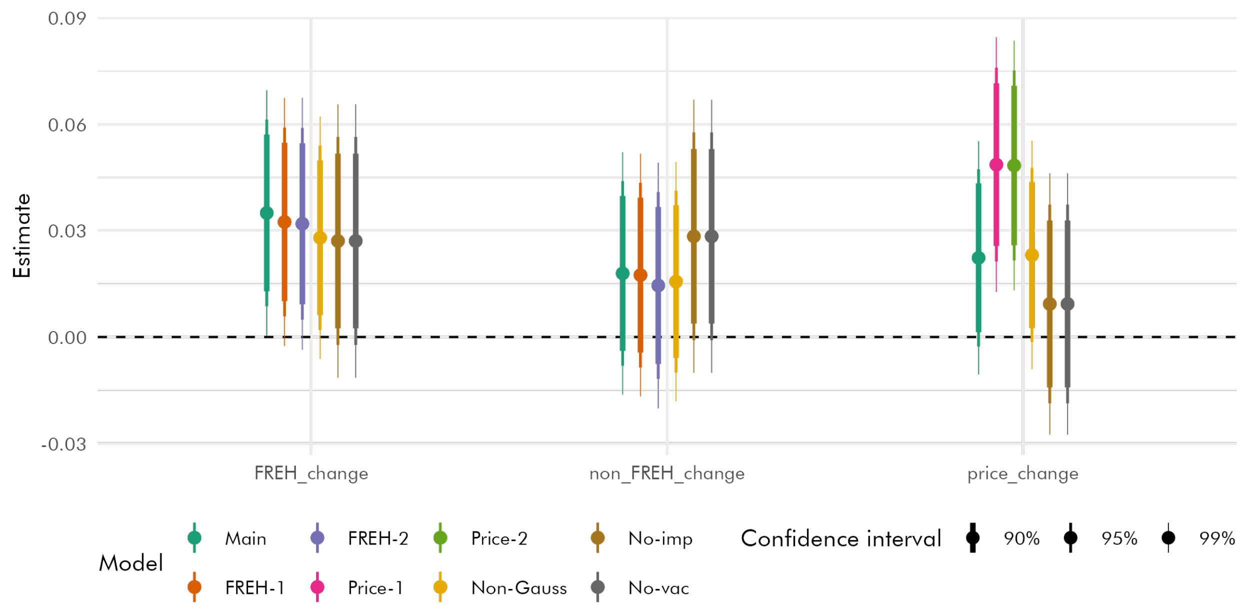

Figure A3. Total causal effect estimates of STR variables on rent_change *for all model variants (point estimates with 95% confidence intervals)*

Model variant with spatially and non-spatially varying coefficients

We account for spatial dependence in our model by implementing random-effects eigenvector spatial filtering, an approach which extracts latent spatial structure in the dataset and models it as random effects. We also evaluated a variant of RE-ESF which estimates spatially and non-spatially varying coefficients (S&NVC) for each variable. In this model, regression coefficients for each variable are allowed to vary spatially, and the estimates are robustified through the addition of a (non-spatially varying) non-linear function to each spatially varying coefficient. The regression equations for the S&NVC model are as follows:

$$y_{i,t} = \sum_{k=1}^{K} x_{i,t,k}\beta_{i,t,k} + f_{MC}\left(g_{I(0)}\right) + \sum_{h=1}^{H} \gamma\left(g_{I(h)}\right) + \epsilon_{i,t}, \epsilon_{i,t} \sim N\left(0,\sigma^2\right) \quad \text{(A1)}$$

$$\beta_{i,t,k} = b_k + f_{MC,k}\left(g_{i,t(0)}\right) + f\left(x_{i,t,k}\right) \quad \text{(A2)}$$

Here $\beta_{i,t,k}$ is the sum of a constant mean b_k, a spatially varying component $f_{MC,k}\left(g_{i,t(0)}\right)$, and a non-spatially varying component $f\left(x_{i,t,k}\right)$, $g_{I(0)}$, $g_{I(1)}$, …, $g_{I(H)}$ are grouping variables, $f_{MC}\left(g_{I(0)}\right)$ is a function of the Moran coefficient matrix applying spatially dependent group effects, and $\gamma\left(g_{I(h)}\right)$ is the spatially independent group effects for group h.

The SNVC approach has been demonstrated to outperform spatial regressions which do not include non-spatially varying coefficients, particularly with respect to the problems of spurious correlation among coefficients and oversmoothing (Murakami & Griffith, 2023). We implemented the S&NVC model using the spmoran package (Murakami, 2017), and relied on iterative Bayesian information criterion (BIC) minimization to decide whether to calculate spatially varying or constant coefficients for each variable. Because BIC minimization did not select spatially varying coefficients for any variable, the result of the SNVC model was effectively identical to our main RE-ESF model. However, we also ran a variant where spatially varying coefficients were forced for the three treatment variables and for the autoregressive term *rent_lag_log*. The results of this model are shown in Table A5, and they are broadly consistent with our main and variant RE-ESF models: all three treatment variables have strong and significant total causal effects.

Diagnostics for *rent_change* model

Standard model diagnostics suggest that the RE-ESF *rent_change* model generally represents the underlying data. Figure A4 shows a residuals-versus-fitted-values plot, and demonstrates no heteroscedasticity. The left panel of Figure A5 shows a normal Q-Q plot which reveals moderate deviation from the expected relationship for high and low values, and hence at least some non-normality in the residuals. For this reason we also ran a non-Gaussian variant of the model (the Non-Gauss model described in Table A4, which has two iterations of SAL transformations to the outcome variable), and the resulting normal Q-Q plot is shown in the right panel of Figure A5. The residuals of the non-Gaussian version of the model are somewhat closer to a normal distribution; however, since the qualitative interpretation of the non-Gaussian and main model variants is the same, and the former is much more difficult to interpret due to the SAL transformation, we have opted to report the main model results in the paper.

To address the possibility of multicollinearity, particularly between the three STR variables included in the models, we calculated variance inflation factors for the treatment variables in an OLS variant of the main model. These were all under 2, suggesting no issues with multicollinearity.

Figure A6 shows the model residuals for the main RE-ESF models in the largest six urban regions in Canada. Although there is no framework for hypothesis testing of spatial dependence in model residuals for an RE-ESF model, visual inspection of Figure A6 does not suggest any meaningful non-randomness to residual distribution. This is expected, given that the RE-ESF model extracts latent spatial structure and converts it into random-effect variables.

Figure A7 plots the outcome variable (rent_change) as well as the model residuals for the main RE-ESF model against year, and demonstrates that both the outcome variable and the residuals are stationary. *Rent_change* values and model residuals remain tightly centred around zero in all years.

	(Intercept)	FREH_ change	non_FREH _change	price_ change	rent_lag_ log	vacancy_ lag_log	apart_log	income_ log
Coefficient estimates								
Min.	-0.190	0.020	0.018	0.024	-0.097	-0.463	0.103	0.109
1st Qu.	-0.104	0.023	0.018	0.024	-0.090	-0.159	0.103	0.109
Median	0.008	0.028	0.018	0.024	-0.068	-0.153	0.103	0.109
Mean	0.046	0.032	0.018	0.024	-0.071	-0.149	0.103	0.109
3rd Qu.	0.241	0.044	0.018	0.024	-0.056	-0.150	0.103	0.109
Max.	0.314	0.049	0.018	0.024	-0.044	0.036	0.103	0.109
Statistical significance								
Not sig.	2,050	2,155	5,400	0	910	455	0	0
10%	20	590	0	5,400	295	6	0	0
5%	215	415	0	0	1,865	95	0	0
1%	3,115	2,240	0	0	2,330	4,844	5400	5400
Variance parameters								
Spatial effects (coefficients on x)								
Std. dev.	0.301	0.025	0.000	0.000	0.045	0	0	0
Scaled Moran's I	0.740	1.000	0.930	0.986	0.872	-	-	-
Non-spatial effects (coefficients on x)								
random_SD	-	0	0	0	0	0.031	0	0
Group effects (std. dev.)		Nbhd: 0.000		Region-by-year: 0.275				
Error statistics								
	Stat							
Resid. SE	0.904							
Adj. R2 (cond.)	0.181							
R.Log.Lik.	-7,444							
AIC	14,932							
BIC	15,077							

Table A5. S&NVC regression results for RE-ESF variant with forced spatially coefficients for treatment variable

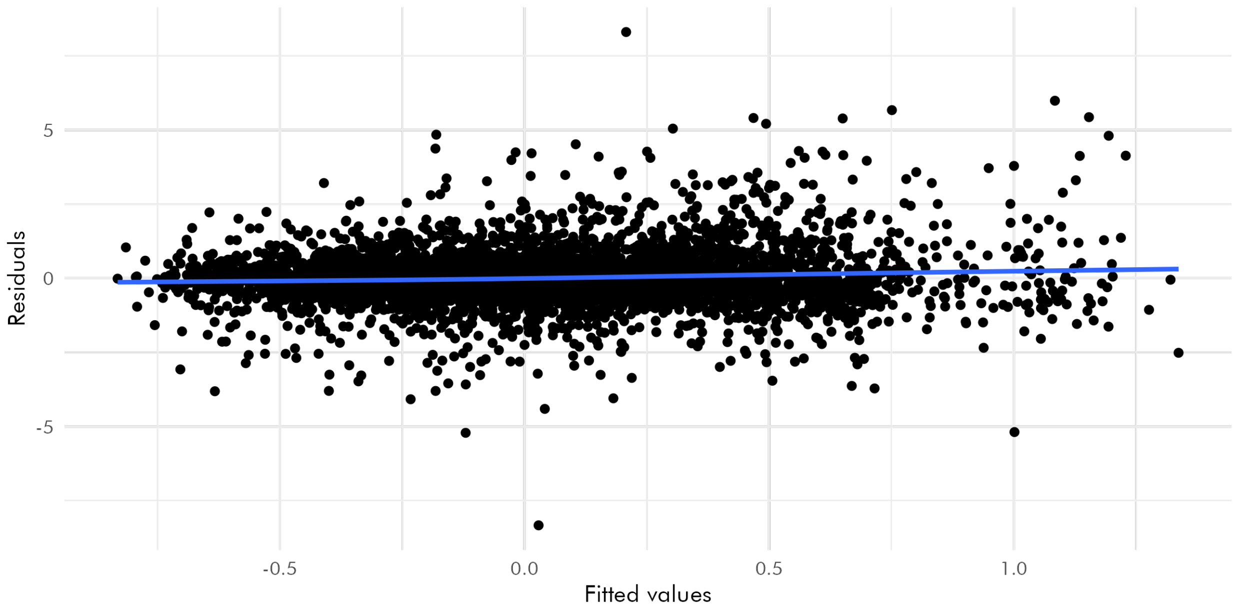

Figure A4. Residuals-versus-fitted-values plot

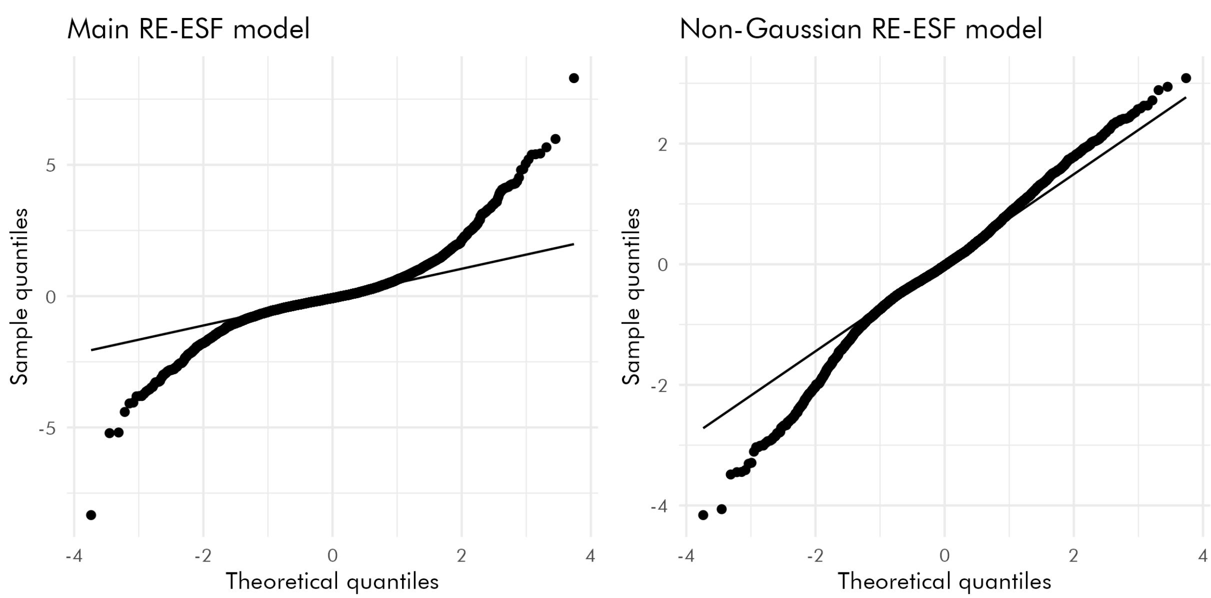

Figure A5. Normal Q-Q plots for main and non-Gaussians versions of the RE-ESF model

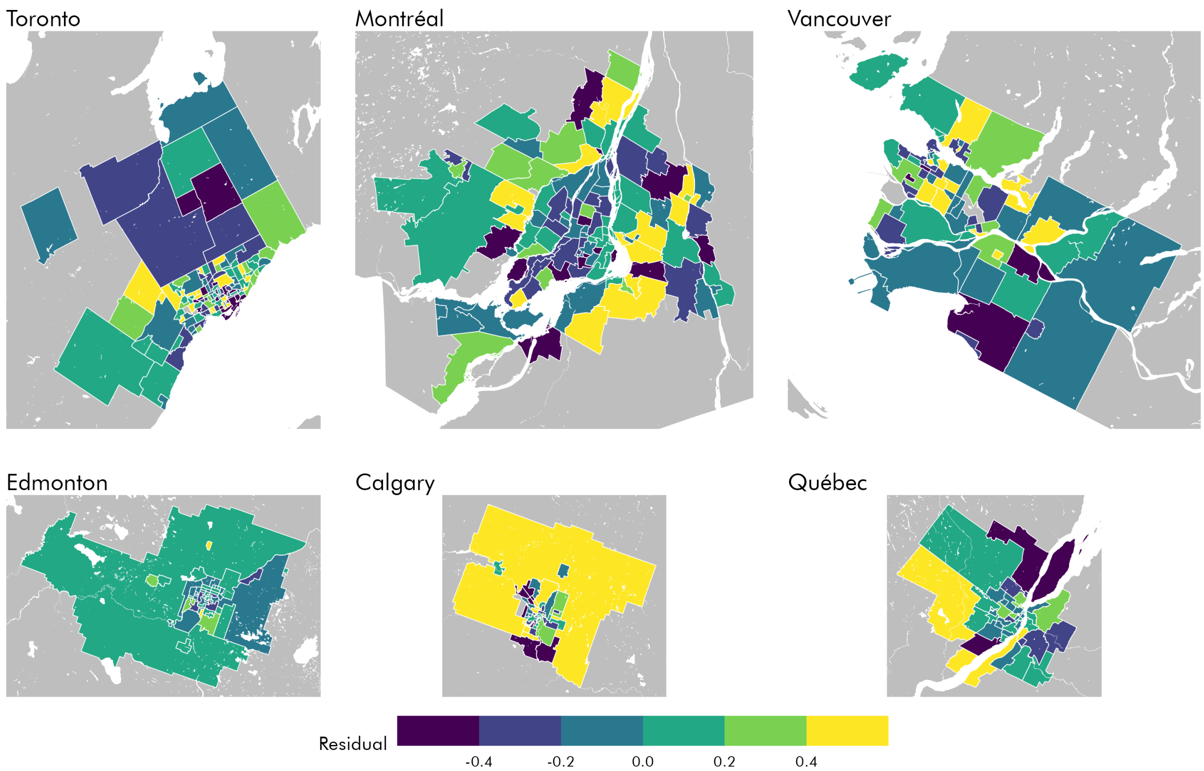

Figure A6. Main RE-ESF model residuals in the six largest urban regions in Canada

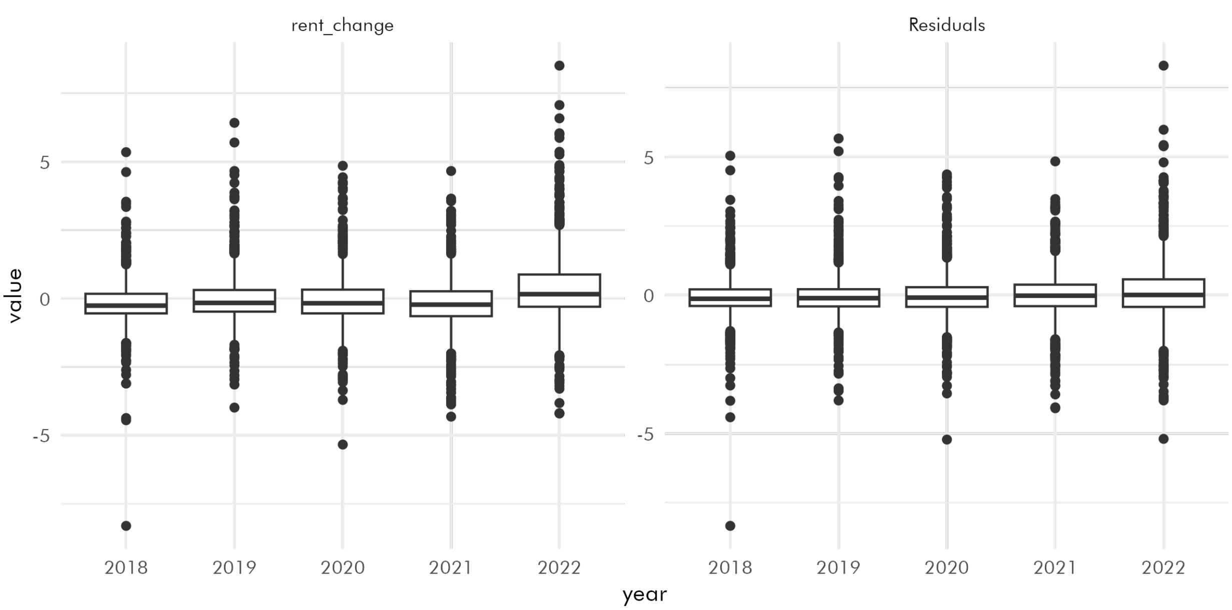

*Figure A7. Outcome variable (*rent_change*) and RE-ESF model residuals by year, demonstrating stationarity in both cases*

Additional appendix references

Murakami, D., Kajita, M., Kajita, S., & Matsui, T. (2021). Compositionally-warped additive mixed modeling for a wide variety of non-Gaussian data. *Spatial Statistics*, 43, 100520.

Murakami, D., & Griffith, D. A. (2023). Balancing spatial and non-spatial variation in varying coefficient modeling: A remedy for spurious correlation. *Geographical Analysis*, 55.1, 31-55.