The Effect of Rent Controls along the 'Excess' Price Distribution

Joan Monras, UPF-BSE-CREI, and CEPR
Jose G. Montalvo, UPF-BSE-IPEG*

This version: February 2025

Abstract

Catalonia enacted a rental cap policy that affected only some municipalities and, within those, only units with prices above their "reference" price. We show that, as intended, the policy led to a reduction in rental prices, but with price increases at the bottom and price declines at the top of the distribution. The policy also affected supply, with exit at the top which was not compensated by entry at the bottom. We show that a model with quality differences in rental units rationalizes the empirical facts and allows us to compute the welfare consequences of an (increasingly popular) policy that was designed as a short-term remedy to a tight rental market.

Keywords: *rental markets, rent control, housing supply*
JEL classification numbers: D4, R21, R28, R31

*We are grateful to the Agència de l'Habitatge de Catalunya (AHC), the INCASOL and the Colegio de Registradores de España for providing access to their data. We would also like to thank Blaise Melly, Gregor Jarosch, Ezra Oberfield, Monika Piazzesi, and Martin Schneider for insightful discussions. We acknowledge the generous support of the Fundación BBVA (Scientific Research Projects in Social Science, 2022 Call) and the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in RD (CEX2019-000915-S). Monras acknowledges the generous support of the Ramon y Cajal grant (AEI - RYC-2017-23682) and the Fundacion BBVA Beca Leonardo (LEO24-1-12346) for its generous support. Daniele Alimonti, Celeste Liu, Pablo de Llanos, and Schuyler Louie provided excellent research support.

1 Introduction

The increase in rental prices has reduced access to housing in many large cities and prompted much political debate. One increasingly popular measure has been the (re)introduction of various types of rent controls. This type of policy has a long history. The Increase of Rent and Mortgage Interest Act of 1915 introduced rent controls in the United Kingdom, something that was subsequently replicated by several countries around Europe. Over the last 40 years, some U.S. cities, including New York City, Washington, San Francisco, Boston, and Los Angeles, have also introduced various forms of rent control. Most commonly, the United States and many other countries have used rent controls which limit the rate of rental price growth within tenancies, although in some instances, rent controls also affect rental price growth across tenancies. With increasingly detailed micro data, some European countries have recently implemented rent controls that establish a reference price for each unit in the rental market.[1]

There is also a long history of studying the effects of rent controls among economists. Since Friedman & Stigler (1946), much economic research has used *theoretical* arguments to emphasize the potential negative consequences of keeping rents below market prices. As we discuss in more detail later, arguments against rent controls include the fear of distortions in the housing rental and sales markets, mismatches between types of housing units and types of tenants, and inefficient levels of consumption of housing services. On the plus side, rent caps might be an effective tool to preserve a neighborhood's social capital and to weaken landlords' bargaining positions, redistributing surplus towards tenants. This is why rental cap policies are sometimes advocated as a way to help low income families, who are disproportionately represented among renters.

In contrast, the *empirical* literature on the effects of rental price controls has been scarce. As noted in Diamond et al. (2019), "we have little well-identified empirical evidence evaluating how introducing local rent controls affects tenants, landlords, and the broader housing market." The reason for this lack of evidence is twofold. First, microlevel data on house prices and rents are mostly available for recent periods, whereas sources of potentially exogenous variation are available for earlier times. Second, as already alluded, there are various ways in which rental cap policies are implemented; hence, we have only a partial understanding of how certain types of rental cap policies can affect the market. Recent empirical evidence in the United States has evaluated the removal of rental price growth caps, as in the case of Diamond et al. (2019), but less is known about the effects of introducing them, and even less about introducing them by computing unit-specific reference prices.

In this paper, we analyze the effects of a rent control regulation that was introduced in Catalonia, Spain, in September 2020. This policy established a maximum rental price (reference price) for each housing unit in the

1. In contrast, a number of U.S. states forbid rent control laws entirely, as described in https://www.lawdistrict.com/articles/rent-control-laws-by-state, last accessed in December 2024.

market, which depended on the prevailing rental price of (at least) its 25 closest neighboring units, according to a methodology that we describe in more detail below. The policy also limited, with some exceptions, price increases relative to the previous contract, both within and across tenancies. The policy was meant to be a temporary solution to slow down rental price growth while developing other longer-run policies. It was operative for a bit over a year, and has since been the example that the Spanish government has used to implement similar short-term solutions to rapid rental price growth.[2] The policy affected some municipalities, but not others, providing us with multiple sources of variation, as it allows us to compare various groups of treated and non treated municipalities, as well as housing units with prices above and below their reference price within each group of municipalities. Moreover, unlike other studies, we have detailed housing unit microdata on rental housing in Catalonia. Specifically, we have information about each rental contract signed, its start date, its end date (in cases where the contract terminated), the rental price, the unit-specific rental index, and some unit characteristics, such as size, age of the building, and exact geographic location.

Using these data, we report several empirical results. First, we document that the policy led to rental price declines. On average, rental prices dropped by approximately 5% in treated municipalities relative to non-treated ones. The estimates are precise, and we find no evidence of systematic differences in the relative trends leading to the policy change – even when we compute pre-trends using more than 50 months of data. We obtain these results using both nonparametric approaches that rely on comparing histograms of "excess prices" – i.e., the price relative to the unit-specific reference price set by the policy – and those that rely on standard difference-in-differences comparisons. Interestingly, we also find a strong convergence of prices toward the reference price in treated municipalities. While price increases are uniform across the entire excess price distribution in non-treated municipalities, we find a strong gradient in price changes along this distribution in treated municipalities. "Low-price" units, defined as units with prices below the reference price, experience price increases above those seen in control municipalities, while "high-price" units experience strong price declines.[3] Studying the effects of rental cap policies along the entire excess price distribution is the most novel empirical contribution of this paper, something that can be done because we have information on the rental price index in *both* treatment and control municipalities that the policy makers established for each rental unit.

As we discuss in more detail later, the (mostly theoretical) literature has argued that price caps can result in changes in the supply of units to the rental market. These changes might be driven by both a change in the composition of the units and a change in the overall number of units available in the market. We document that

2. When introduced, the policy was supposed to last for a maximum of 5 years. It was ruled unconstitutional, and, hence, removed, 18 months after its introduction.

3. We apply some of the insights in Autor et al. (2015) and Cengiz et al. (2019), who study the effect of minimum wages along the wage distribution, to study the effect of rental price caps along the rental price distribution.

the overall supply of units in the market declined by around 10%, mostly explained by a large decrease in units available at prices above the reference price, which is not compensated by the increase in units below. These estimates allow us to compute an elasticity of rental unit supply, which we define as $\frac{\Delta \ln(\text{New rental contracts})}{\Delta \ln(\text{Rental Price})}$, of around 2.[4]

From the perspective of a standard perfect competition model of the rental market with homogeneous housing it is, in principle, straightforward to evaluate the welfare consequences of rental cap policies. In this standard model, as long as the supply of rental units is not perfectly inelastic, rental caps lead to a reduction in rental prices and reduced welfare among owners. Rent caps also generate a gap between supply and demand. All renters who participate in the market, however, are better off. Overall, the perfect competition framework also predicts a deadweight loss which, from a policy perspective, can be viewed as the cost of redistributing welfare from landlords to renters.

Much of the empirical evidence that we document, however, does not fit well with this perfect competition framework. For instance, in perfect competition models with homogeneous housing, there is only one rental price. Hence, the standard model is not able to reproduce price dispersion in equilibrium nor the heterogeneous effects of the policy along the excess price distribution, two central empirical features. Instead, we argue that the facts that we document are best understood through the lens of a perfect competition model with quality differences. In our model, we think of quality differences as different types of housing. Some rental units are of high-quality and some of low-quality. We assume that potential renters substitute across these types of housing (and potentially an outside option) only imperfectly. This captures the idea that renters may see two (far apart) housing units facing the sea in Barcelona as perfect substitutes, whereas, substituting toward a unit in the basement may be more difficult, even within the same building.[5] We assume landlords own a unit that is of either high- or low-quality. Given the different prices for quality, landlords decide whether to supply the unit into the market or not. The model generates price dispersion in equilibrium. High-quality units, to the extend

4. We discuss in detail below why we think that our measure of supply in the rental market accurately reflects changes in the total supply of housing units in the rental market. For instance, we document that there does not seem to be a change in the time elapsed between contracts, suggesting that there is a fraction of rental units that leave the rental market, and hence, that new contracts signed is a good approximation of changes in the overall supply of housing in the rental market. To provide more credibility to our estimates we also use an alternative measure of housing supply, namely the likelihood that one of the units over the entire sample period is rented at a given point in time. A likely mechanism through which rentals left the market is by staying empty for a while, by being sold and used by its buyers, or by entering in short-term contracts – a contract with a maximum duration of 11 months intended for short stays –, which are likely harder to detect in our data.

5. The crucial innovation is to study how the demand for quality differences shapes the market, rather than how quality differences affect or determine the supply of rentals, as studied, for example in Frankena (1975). Our model is related to models where part of the market is under a price cap, while part is uncontrolled. The difference is, again, how we model this through imperfect substitution across markets from the view point of potential renters.

that are valued more in utility terms and that supply is not significantly above that of low-quality units, are going to be more expensive.

We use our model to study how a price cap that is binding for high quality units affects the market equilibrium. The model predicts that the price of high-quality units will decline. At the prevailing new price there are some landlords of high-quality units that leave the market. With the cap, the supply of high quality housing does not satisfy the demand. Some of this demand substitutes toward low-quality units and some to the outside option. As a result of this increased demand for low-quality units, low-quality prices increase. Whether the total size of the rental market (i.e. adding high- and low-quality units) decreases or not depends on whether there is an outside option.[6] We argue that this model, when we define high- and low- quality units as units with prices above and below their reference price, delivers predictions that are completely in line with our empirical findings.

The model implies that, if we abstract from misallocation and rent-seeking considerations (Glaeser & Luttmer, 2003; Bulow & Klemperer, 2012), we can approximate the welfare gain for renters in high-quality units (which is equal to the loss of the landlords of those units) by estimating the percentage change difference in high-quality prices in treated municipalities (relative to control) and multiply this by the size (in Euros) of the market. Similarly, we can approximate the welfare cost of the policy for renters in low-quality units (or gain of the landlords of those units) by estimating the percentage change difference in low-quality prices scaling, in this case, by the size of the low-quality market. Finally, we can also calculate the deadweight loss of the policy by multiplying the percentage change in high-quality prices and quantities, and the size of the high-quality market, all this multiplied by one half. Using our data we estimate that the policy lead to a welfare equivalent gain of around 300 million Euros per year for renters of units with prices (initially) above their reference price, a loss of around 75 million Euros per year for renters in units below their reference price, and a deadweight loss of about 36 million Euros per year. These exercises provide a new insight to the existing literature. Redistribution effects of rental cap policies may not only go from landlords to renters (as predicted in standard models), but rather from landlords of high-quality units *and* renters in low-quality units, to landlords of low-quality units *and* renters in high quality ones. This insight highlights the importance of evaluating rental cap policies by studying what happens to prices and supply along the entire distribution, new in this paper.

In principle, assuming perfect competition in our model with quality differences may be too restrictive. Perfect competition assumes both that landlords do not internalize the demand for their units (i.e., that they take prices as given instead of setting prices optimally) and that there is perfect information (or more broadly,

6. In the model, to preserve simplicity, we assume that the rental housing supply elasticity is the same in the low- and high-quality markets. In a more general model, the total size of the market would also depend on the relative shape of these two rental housing supply elasticities.

no search frictions). Often times, rental cap policies are justified as a way to limit market power. Rental cap policies may also be effective in solving search frictions, as argued by Galenianos & Kircher (2009) in the case of labor markets. For these reasons, it is worth investigating whether models that allow for market power or search frictions alone can explain our empirical results, and whether incorporating these elements makes our model with quality differences more realistic or leads to fundamentally different conclusions. If the rental market in Catalonia was characterized by market power and quality differences, a rental cap policy, which would be likely more binding for high-quality units, would lower prices and *expand* supply among high quality units. This is contrary to what we see in the data. Similarly, if all the price dispersion that we see in the data was due to search frictions, rather than quality differences, then, after the policy, we would not observe price dispersion below the cap, and we would observe entry rather than exit above the cap, something that, again, is not what we observe in the data. For these and other reasons that we discuss in detail below, we argue that our perfect competition model with quality differences, while intentionally simple, captures the essential forces that help us rationalize the empirical evidence and study the welfare consequences of the policy.

Overall, these results provide empirical support for the notion that rental price caps are effective at reducing rental prices but also affect supply and welfare, potentially in richer ways than what standard frameworks suggests.

Related literature

There is a rather large theoretical literature that has identified several channels through which rental caps can lead to unintended consequences. For instance, rent controls can lead to mismatches between tenants and rental housing units. Once a tenant has obtained a rent-controlled apartment, the household might decide to reside in that dwelling longer than would have been optimal based on their needs for housing services, a point made by Suen (1989), Glaeser & Luttmer (2003) and Bulow & Klemperer (2012) among others, which has consequences for the whole housing market, as emphasized in Basu & Emerson (2000).[7] Similarly, if rental prices are below market prices, renters might choose to consume excessive amounts of housing (Olsen, 1972; Gyourko & Linneman, 1989). Rental price caps can also cause rental housing to deteriorate. Landlords might decide not to invest in maintenance because they cannot recoup the investment when rental price growth is limited (Sims, 2007; Downs, 1988; Moon & Stotsky, 1993). Rent controls can also affect the relative size of the rental and home-ownership markets (Fetter, 2016). On the positive side, rent control also offers benefits for incumbent tenants and potentially the neighborhood. For example, longer-term tenants might develop specific social capital, such as a network of friends and family or proximity to a job or a local school, which may result in

7. One example of these mismatches is households with adult children who might decide to stay in family-size apartments, and, hence, displace families with young children into crowded accommodation and small studios.

positive spillovers for the neighborhood. Similarly, rent controls may offer insurance for risk averse renters facing variability on future rental prices (Favilukis et al., 2023). All the aforementioned work studies mechanisms that likely operate over the longer-run. Relative to these papers, our focus is to study empirically the short-run effects of a policy that established a unit-specific price index with the intention to temporarily limit rental price growth.

In the short-term, the theoretical literature argues that rental caps potentially lower prices but may also affect housing supply. We contribute to this theoretical literature by exploring the short-term effects of the introduction of rental caps in a perfect competition model where potential renters substitute imperfectly between high- and low-quality units. We show that allowing for quality differences leads to richer empirical predictions that match what we observe in the data. We argue that our model provides a novel insight: welfare redistribution of rental cap policies may go, not just from landlords to renters, but also from certain type of landlords and renters to some other type of landlords and renters. To some extent our model is closely related to Couture et al. (2023), although the focus of that paper is very different than ours. Couture et al. (2023) develop a model with heterogeneous quality across neighborhoods, and demand-driven imperfect substitutability across these neighborhoods, to explain why households of various income groups sort across space.[8] Couture et al. (2023) do not use their framework to analyze the effect of a rental cap. However, their model would predict effects comparable to the ones we find if rental caps were placed only in the highest quality neighborhood. Our model is also related to the work by Landvoigt et al. (2015) and Piazzesi et al. (2020). Like them, we consider that the housing market is better understood as a collection of submarkets. Landvoigt et al. (2015) and Piazzesi et al. (2020), however, focus on the homeownership market, and use their framework to study why within metropolitan areas we see price change heterogeneity. Their analysis relies on search frictions rather than perfect competition. We incorporate (a version) their framework into ours in the appendix and explore how our conclusions would change.

The increasing availability of reliable microdata at the level of individual rental units has allowed researchers to study rent control policies empirically in some detail. In a seminal contribution, Sims (2007) analyzed the end of rent control in the Boston metropolitan area using data from the American Housing Survey. He found that rent control induced landlords to remove units from the rental market and lowered rents. In addition, Sims (2007) showed that there was no effect on construction and a reduction of maintenance in controlled units, and that the spillover effect reduced the rental price even of noncontrolled units. In a similar vein, Autor et al. (2014) analyzed the impact of the unexpected elimination of rent control in Cambridge, Massachusetts, finding that the market value of properties no longer subject to rent control increased 45%. In addition to the direct

8. We simplify their framework in that we do not link heterogeneous income levels to residential choices, although this is something we could incorporate.

effect of removing rent caps, Autor et al. (2014) showed that removing rent control had significant, indirect effects on neighboring properties, also increasing their value. Unlike our paper, these two papers analyze the removal of rental price controls, rather than its introduction. Their focus is also on the long-term consequences of price controls, rather than the shorter-term impact on the rental market. Moreover, we focus our study on what happens to prices in different parts of the excess price distribution.

Diamond et al. (2019) discusses the consequences of an expansion of rent control on tenants, landlords, and the overall San Francisco housing market. In 1979, San Francisco imposed a rent control on buildings with five or more apartment units. This policy included regulated increases indexed to the Consumer Price Index (CPI) within tenancies, but the regulation did not apply between tenancies. New and smaller buildings were exempt from control. However, in 1994, the exemption for buildings with fewer than five apartments was eliminated. Diamond et al. (2019) showed that, while rent control appears to help ordinary tenants in the short term, it decreases affordability in the long term. Rental units are sold to very high-income buyers, increasing gentrification. We differ from Diamond et al. (2019) in that we analyze a rental cap policy that introduces a reference price for each unit in the market and we study, both empirically and theoretically, the heterogeneous effects of this policy along the entire distribution of rental units in the market.

Our paper also contributes to the recent literature on rent control in Europe. Breidenbach et al. (2022) uses detailed rental price microdata and a triple differences event study to show that the initial success of rent caps in reducing rental prices vanished after one year. Some of the empirical findings in our paper resonate with theirs; however, we provide a more complete study of how the policy affects the rental market along the entire price distribution, and we provide a novel theoretical framework which allows us to estimate the welfare consequences of the policy. Mense et al. (2019) and Mense et al. (2023) also analyze the introduction of what can be considered a rental cap in Germany (applied both within and between tenancies), but where the policy does not apply to new construction. Their emphasis is, as a result, slightly different. Mense et al. (2019) show that land prices grow faster in municipalities where rent caps are implemented than in the control group, while Mense et al. (2023) document a positive spillover effect of rent control on units not subject to rent caps within municipalities. Relative to these two papers, we provide evidence on heterogeneous effects and a novel framework to analyze these types of policies. In our case, moreover, we do not detect spillovers in rental prices *across* municipalities. It is also worth noting that, relative to these two papers, we use data on actual market rents rather than rents posted on internet platforms, which, we argue, provides more accurate information, especially in periods of distress.[9] In a contemporaneous paper, Jofre-Monseny et al. (2022) study the same policy that we analyze in our paper but they use a purely empirical approach to study the effects of the policy

9. Although the broad trend of the movement of offered and market rents over time is similar, the changes of these two prices can diverge, especially around special events or cyclical changes.

on *average* outcomes. We argue that not studying what happens along the entire distribution misses crucial insights.

To summarize, in this paper we contribute to the literature in two ways. First, we provide well-identified short run evidence on the effect of rent controls along the entire distribution of rental units. Second, we provide a novel framework, which relies on quality differences, to interpret the empirical evidence and to estimate the welfare consequences of this type of policy. Our analysis is possible because we obtained data on the reference price for each rental unit, both in treatment and control municipalities, information which is missing in previous studies. Our results show that redistribution effects of rental cap policies are richer than suggested in prior work.

2 Rent control in Catalonia

The Spanish housing market is characterized by a large proportion of homeowners. In 2011, homeowners represented 79% of the market, while renters at market price amounted to 11%.[10] After the financial crisis, the demand for rental housing increased significantly in a very short period of time. In 2018, tenants at market price represented 16% of the market, while homeowners decreased to 75%. Several factors explain this change. First, banks reduced lending and increased the requirements for obtaining a loan (Carozzi, 2020). Second, young people likely realized that obtaining a mortgage to buy a house before obtaining financial stability could lead to financial distress and a high risk of foreclosure. Third, most of the population growth in a number of cities was primarily driven by international immigrants, who are, in general, less likely to buy. Other factors, such as rental platforms and tourist pressure, might have also played a role, at least in some specific neighborhoods (Garcia-Lopez et al., 2020; Almagro & Dominguez-Iino, 2022). This large increase in the demand for rental housing was not met with a sufficient increase in supply, and rental prices increased, especially in some municipalities.

To contain the increase in prices and following the experiences of some European cities, such as Berlin and Paris, the Parliament of Catalonia approved, quite unexpectedly, a system of rent ceilings that was implemented by mid September 2020.[11] The policy was meant to be a temporary solution to the increase in rental prices in a selected set of municipalities.

The Catalan rental price index in fact began as an informational tool to provide a reference to renters of the minimum and maximum rental price per square meter for flats, considering both their location and specific characteristics. Following the (unexpected) enactment of *Llei 11/2020* by the Catalan Parliament on September

10. These two numbers do not total 100% because some tenants have reduced rents and some units are rented free of charge.

11. Act 11/2020 of September 18, 2020. The law passed unexpectedly on September 8, 2020, after the center-right party Junts per Catalunya, which traditionally opposed these type of policies, voted in favor. See link to news article and link to the law.

18, 2020, a plain version of this index became a regulatory benchmark in areas declared as having a "stressed" housing market. For a municipality to qualify as "stressed", two conditions had to be satisfied: (a) rental prices must have had to increase by more than 20% between 2014 and 2019, and (b) the area had to be part of the Barcelona Metropolitan Area or had to have a population exceeding 20,000 inhabitants. This resulted in 61 "treated" municipalities.[12]

The rental price index was based on data from the Generalitat de Catalunya's Catalan Land Institute (*INCASOL*), the Catalan agency responsible for rental deposit records, and was computed by the Agencia de l'Habitatge de Catalunya (*AHC*). To calculate the reference price for each flat, the system created a comparison group for each unit by selecting rental contracts from properties within a radius starting at 50 meters and increasing incrementally (up to 1,050 meters) until a minimum of 25 contracts were included. Additionally, the comparison group was stratified by the unit's size: comparison contracts for units larger than 90 m^2 included units between 91 and 250 m^2, while comparison contracts for smaller units were matched to contracts within ± 10 m^2 of their size.

For each unit, the reference price was then computed as the average price per square meter from the unit-specific comparison group. A range of prices (minimum and maximum) was also computed by adjusting the reference price based on specific characteristics of the unit under consideration such as maintenance condition, availability of elevator, parking space, year of construction, furnishings, floor level, and energy certification. Prior to the policy, the intervals generated by these adjustments were meant to provide information on rental prices for each individual property.[13]

After the approval of the law, reference prices became the highest price that the rental price per square meter could reach, at least on paper. Apart from the rent cap, the Catalan law also limited rental price increases, both within and across tenancies.[14] The law, however, allowed for some exceptions to tenants that had performed major investments on the unit and to new construction, which could charge the maximum price

12. 1. Badalona 2. Barberà del Vallès 3. Barcelona 4. Blanes 5. Calafell 6. Castellar del Vallès 7. Castelldefels 8. Cerdanyola del Vallès 9. Cornellà de Llobregat 10. Esplugues de Llobregat 11. Figueres 12. Gavà 13. Girona 14. Granollers 15. L'Hospitalet de Llobregat 16. Igualada 17. Lleida 18. Manlleu 19. Manresa 20. Martorell 21. El Masnou 22. Mataró 23. Molins de Rei 24. Mollet 25. Montcada i Reixac 26. Montgat 27. Olesa de Montserrat 28. Olot 29. Palafrugell 30. Pallejà 31. Pineda 32. El Prat de Llobregat 33. Premià de Mar 34. Reus 35. Ripollet 36. Rubí 37. Sabadell 38. Salou 39. Salt 40. Sant Adrià de Besòs 41. Sant Andreu de la Barca 42. Sant Boi de Llobregat 43. Sant Cugat del Vallès 44. Sant Feliu de Guíxols 45. Sant Feliu de Llobregat 46. Sant Joan Despí 47. Sant Just Desvern 48. Sant Pere de Ribes 49. Sant Vicenç dels Horts 50. Santa Coloma de Gramenet 51. Santa Perpètua de Mogoda 52. Sitges 53. Tarragona 54. Terrassa 55. Tortosa 56. El Vendrell 57. Vic 58. Viladecans 59. Vilafranca del Penedès 60. Vilanova i la Geltrú 61. Vilassar de Mar.

13. This methodology is explained in detail in the following link: https://portaljuridic.gencat.cat/ca/document-del-pjur/?documentId=881998.

14. The Spanish regulation already limited rent increases within tenancies using CPI inflation, something that was extended to also included new tenants.

within the aforementioned interval, as long as it was less than 5% higher than the reference price. In case the contracts were not complying with the new norm tenants could use the Agencia Catalana del Consum to initiate a complain. As far as we know, though, there was no active and independent effort on the part of the administration to track whether the law was followed.

As mentioned, the main goal of the policy was to (temporarily) slow down rental price growth in a number of municipalities deemed to be "stressed". The law was meant to be in place for a maximum of 5 years. There were also discussions on longer-run type policies, such expanding the supply of housing, whose implementation required additional time and resources. Hence, this paper can be understood as an evaluation of these short-run type of policies, implemented in periods of distress, that leverage rich microdata to compute unit specific price ceilings. Almost four years after the implementation of this policy, the Spanish government passed a law inspired in the Catalan experience. It is important to notice that the adoption of this policy for the whole country was partly justified by the apparent success of reducing average rents in Catalonia.

3 Data

Our data cover more than half a million rental contracts signed from the beginning of 2016 until June 2021. In Catalonia, there is a legal obligation to report rental contracts to the Generalitat de Catalunya's Catalan Land Institute (INCASOL) in the form of a bail. Our data come from the AHC, which took the information available in INCASOL and cleaned it to compute the reference prices.[15] As mentioned in the previous subsection, the AHC is the public agency that was responsible for setting the methodology and performing the calculation of the reference price for each rental unit and as such, it had to make sure that the information was accurate and reliable since there were potentially legal consequences for mistakes in the computation of the reference price. The AHC checked for mistakes and inconsistencies in the original data,[16] imputed the adequate measure of the size of the houses and obtained the necessary information to compute reference prices, such as the geographical coordinates of each residential unit.[17]

The final dataset contains information on the closing price of each rental contract, the start date of the contract, the end date (if it had already ended), the area in square meters of the unit, and the reference price per square meter. The AHC computed the reference price for multiple periods prior to the implementation of

15. Jofre-Monseny et al. (2022) used first, aggregated data, and then data directly from INCASOL. Hence, they could not use the reference price calculated by the AHC which is very important for our analysis.

16. For instance, INCASOL data contains many rentals of less than 100 euros that are actually an increase in the bail and not a new contract, something that is taken care of in the AHC data that we use.

17. Some times there are small inconsistencies in the size of units in the INCASOL, given that there are two measures to compute surface: usable surface and the built surface. We use the former since that is the source for the reference price per square meter.

the policy and for both treated and control municipalities.[18]

One key feature of these data is that we have information about rental market prices, rather than the asking price or an imperfect measure of the agreed price – which are the typical measures from web scraped or rental agency type data sets. Hence, we can measure potential price effects more accurately. This difference, for example, is important with respect to the research that evaluated the impact of rent control in Germany, which used posted rents (Mense et al., 2019, 2023). These data are also more accurate than the online price data used sometimes in other research that are widely available through scrapping. As seen in Figure A.1, in the Appendix, prices posted by online intermediaries are much more volatile, with faster declines in downturns and faster growth in upturns.[19]

As mentioned above, our data cover two types of municipalities. We denote as the rental contracts of units in municipalities that were subject to the rent cap "treated", and the ones that were not subject to the rent control cap, as "non-treated". It is worth emphasizing that only contracts signed after September 22, 2020, were subject to the new regulation. Figure I shows a map of the distribution of the municipalities subject to rent caps. As anticipated before, many of the treated municipalities are in the metropolitan area of Barcelona, although some are in other areas of Catalonia.

Treated and non-treated municipalities are markedly different in a number of dimensions, beyond their proximity to Barcelona. Table I reports a number of summary statistics. Treated municipalities tend to be larger, more expensive, and denser. Since our empirical design depends on comparisons in the trends across municipalities, systematic differences in the level of outcomes of interest do not necessarily pose a threat to our identification strategy. The major thread to our identification is systematic differences in the trends of the outcome variables of interest. To mitigate these concerns we show data many months prior to the policy change, as we discuss in what follows.

Finally, it is worth noting that the ownership of rental units in Spain is remarkably dispersed. The Bank of Spain estimates that only around 5% of all the units in the rental market are owned by *large* corporations and investment funds, and that only 8.1% of the stock of rental houses is owned by private companies, both small and large, a number that increases to 11.1% in Catalonia. The picture for ownership structure of rental houses is not much different if we look at particular cities. For instance, as we document in Table II, ownership of rental properties is not very concentrated in Barcelona, the main city of Catalonia, and, in fact, the largest concentration is due to the public sector and NGOs, not private corporations. Hence, contrary to what happens in other countries, the fraction of private owners who have many units in a market is small. We use this

18. After the rent caps were declared unconstitutional, INCASOL stopped sharing the data with the AHC since there was no need to update the reference prices. This means that we cannot update our data.

19. Online prices correspond to Idealista, which is the leading online real estate intermediary in Catalonia.

information to justify the introduction of a perfect competition model as our baseline framework which helps us to rationalize the evidence that we present in what follows.

4 Empirical Evidence

In this section, we evaluate the impact of the policy change that took effect in late September 2020. We analyze various dimensions through which the policy affected the rental housing market. First, we study how the policy affected rental prices using a number of empirical strategies. Second, we study whether the policy also affected the supply of housing, using two alternative measures, one in the main text and the other in the appendix.

4.1 The effect of the policy on rental prices

4.1.1 Nonparametric evidence

We start our analysis of how the policy change affected rental prices by providing nonparametric evidence. As we explained before, the rental cap implied that rents could not exceed the reference price computed for each particular unit. Hence, we can define the *excess price* as:

$$\text{Excess price}_{i(m),t} = \ln\left(\frac{p_{i(m),t}}{\overline{p}_{i(m),t}}\right)$$

where $p_{i(m),t}$ denotes the rental price per square meter in unit i, located in municipality m, at time t, and where $\overline{p}_{i(m),t}$ is the price cap calculated for this housing unit based on the surrounding units in that time period – this is computed directly by the AHC, as explained above.[20] The excess price measures the percentage deviation of the actual rental price from the cap. That is, if the rental price is at the cap, then this measure is equal to 0. Any negative number means that the rental price is below the cap. Hence, if the excess price is equal to -0.1, it means that the rental price is approximately 10% below the rental cap.

It is worth noting that the reference price is also calculated for the control area. Using this information we can construct excess price histograms for both treated and non treated areas around the implementation of the policy. Figure II shows four histograms around the months of the policy change and around the *same* months in the previous year. Specifically, the graphs on the left show histograms for September and October 2019, when the policy was still not in place, and the graphs on the right show September and October 2020 – the

20. Note that this reference price is fixed after the implementation of the policy in treated municipalities. Before the policy it was updated only infrequently.

first month with the new rental cap policy. Panel A includes the contracts in the treated municipalities, while Panel B shows the distributions of rental prices for the contracts in the control areas.

Comparing figures on the left and on the right indicates that the control municipalities do not show much difference in rental prices before and after the policy change. However, in treated areas there is a clear movement of the distribution of excess rents to the left. It is also interesting that, after the policy change, a significant proportion of the rental prices are above the corresponding reference price, which seems to suggest that landlords found ways, both legal and illegal, to charge renters a price above the reference price – a point we return to below, and that we investigate in detail in Appendix A.4.

While the histograms presented in this section are indicative of the effects of the policy, it might be useful to quantify the results using various regression models. We turn to this in what follows.

4.1.2 Two-way fixed effects estimates

A second way to look at how the policy affected rental prices is to compare treated and non treated municipalities, following a standard two-way fixed effects or difference-in-differences approach. These types of research designs crucially depend on whether the trends in the outcomes of interest between the treated and control groups are parallel before the treatment. This condition is usually checked visually by running dynamic difference-in-differences specifications such as the following one – ideally with multiple time periods prior to the treatment:

$$y_{i(m),t} = \delta_m + \lambda_t + \sum_k \beta_k \mathbf{1}_{t=k} \times \textit{Rent Cap}_m + \varepsilon_{i(m),t} \tag{1}$$

where $y_{i(m),t}$ is an outcome of interest, such as the excess price defined earlier, and where $\textit{Rent Cap}_m$ is a dummy variable taking the value of 1 if municipality m ever faces a rent cap. The coefficients β_k estimate the differential effect of the intention to treat at various points in time, including pre-treatment periods when the treatment should have no effect. In this and subsequent equations, i indicates units, m municipalities, and t time – usually at a monthly frequency. In some specifications we aggregate the outcome of interest at the municipality level.

Figure III plots these β_k coefficients for three different outcomes of interest for each month around the policy change, with a window of more than 50 months prior to the change and 9 afterward. The three outcome variables that we consider are the excess price (i.e., the log ratio of the rental price divided by the reference price), the price per square meter, and the proportion of rental prices in a municipality above the reference price. In each of these three cases, we consider the baseline specification (1) and a second specification that includes flexible province trends – i.e. non parametric trends common for a larger geographic units (provinces)

that group various municipalities.[21] We also include in some specifications unit fixed effects, since we have a number of units with multiple contracts, something that does not change any of the results reported in what follows.

Using the baseline specification, we see in Panel A that the average excess price drops immediately after the introduction of the rent cap. It is also worth emphasizing that all of the coefficients prior to the policy change fluctuate around zero and are never significantly different from zero. We get very similar patterns when we include provincial trends. If anything, including province-specific flexible time trends makes the trends prior to the policy change even more parallel. When we include these flexible province trends, the drop in prices after the implementation of the rental cap is of the same magnitude.

Panel B of Figure III shows the results for the rental price per square meter. In this case, we observe a slight upward trend in the price per square meter in treated relative to control municipalities. This is what we would have expected from a policy that tries to reduce housing prices in places where rental prices are increasing. Still, the differential trend is small, and even smaller when we control flexibly for province trends. In contrast, the drop after the policy change is sharp and substantial. Comparing these estimates to the ones in panel A suggests that a benefit of using excess price rather than price per square meter as the outcome variable is that the computation of the reference price effectively removes the differential trends between treatment and control.

Panel C of Figure III shows yet another way to measure whether the policy was effective at reducing rental prices. In this case, the dependent variable is the fraction of rentals that are above reference prices in each municipality. Again, trends leading to the policy change are remarkably parallel between treatment and control municipalities, and the effect of the policy is clearly visible.

We can quantify the results observed in Figure III by running the following difference-in-differences specification:

$$y_{i(m),t} = \delta_m + \lambda_t + \beta Rent\ Cap_{m,t} + \varepsilon_{i(m),t} \tag{2}$$

where Rent $\text{Cap}_{m,t}$ is a dummy variable taking the value of 1 if municipality m has a rent cap in place at time t. β measures the average (intention to treat) effect of the policy.

Table III shows the results using the same three outcome variables: excess prices, price per square meter, and the fraction of prices above the reference price. Across the three panels, column 1 shows the baseline specification, which includes only municipality and month fixed effects. Column 2 includes province times

21. We have also estimated models with "comarca" fixed effects. "Comarques" are smaller geographic units than provinces. Results are almost identical.

month fixed effects, which allows for each of the four Catalan provinces to evolve flexibly over time.[22] Column 3 includes rental unit fixed effects. Hence, this column identifies the effect of the policy by comparing the outcome of interest for the same unit before and after the policy change. Columns 4 and 5 repeat the specifications in columns 1 and 3 but with only municipalities that have fewer than 40,000 inhabitants. This specification should mitigate, to some extent, concerns about the treatment and control groups being markedly different, as explained above. Finally, column 6 repeats the estimation of column 2 but using quarterly instead of monthly fixed effects.[23] Standard errors reported in Table III and in the rest of the paper are always clustered at the municipality level to account for serial and group correlation.

Panel A presents the results for the outcome excess price. All of the specifications show that rent caps reduced excess rental prices. Excess prices are estimated to decrease by between 4% and 8%. Several things are worth noting. First, when we include flexible province trends, the point estimate is only slightly lower than in our baseline specification, suggesting that differential trends across provinces are not particularly pronounced. The point estimate is somewhat smaller (in absolute terms) when restricting the sample to comparable municipalities (at least in terms of size). This outcome can be explained by somewhat differential trends between Barcelona and the other smaller municipalities in our sample, or, most likely, by a smaller effect of the policy in smaller municipalities. It is interesting to explore how point estimates change once we include unit fixed effects. When we compare columns 1 and 2 to column 3, we see that the point estimate increases somewhat in absolute terms. This outcome indicates that the policy change seems to have a disproportionate effect on the units that stay in the market.

Panel B shows the same specifications as Panel A but using rental prices as the outcome variable.[24] Estimates fluctuate between -0.08 and -0.037.[25] Moving from column 1 to column 2 shows that differential province-specific trends may be a larger concern when studying rental prices than deviations of those from the reference price, as one could expect from the fact that the reference price removes part of the differential trend. Moving to column 3, we again observe that the point estimate is substantially larger when conditioning on unit fixed effects.

Panel C shows another way to look at the data, akin to panel C of Figure III. For each municipality-month

22. We can also include flexible county-specific flexible trends and the results are almost identical. We prefer province flexible trends because there are counties that consist of only one municipality.

23. Appendix A.2 shows additional graphs for quarterly frequency.

24. It is worth mentioning that we have a number of units in control for which reference prices are not computed, hence, we have a smaller number of observations in Panel A than in Panel B.

25. The results are also robust to using the estimator proposed in de Chaisemartin & D'Haultfoeuille (2020) (see Table A.1 in the Appendix), which delivers a main estimate of -4.3%, and other estimators that allow for dynamic effects, such as Callaway & Sant'Anna (2021). In this case, the estimation is -4.8%. Using the recently proposed synthetic diff-in-diffs estimator (Arkhangelsky et al., 2021), the estimation is -4.1%. The average treatment effect estimate using propensity score matching is a bit higher, around 8.7% as we show in Table A.5 in the Appendix.

pair, in this panel we compute the fraction of housing units that is above the reference price. The estimates corroborate the evidence presented in Panels A and B. The decline in the fraction of units with prices above the reference price, which was around 50% before the policy, fluctuates between 13 and 22 percentage points – both statistically significantly different from 0 and from 50% –, depending on the types of comparisons in each specification. Hence, again, the policy was followed to some extent but not fully: non-compliance with the policy is estimated to be at least 56% ($(50-22)/50$).

In Appendix A.4, we investigate whether non-compliance was systematically related to the unit characteristics that we have at our disposal (surface, location characteristics, and age of the building, see Table A.2), and whether the share of non-compliance in a municipality is systematically correlated with municipality characteristics (see Table A.3). Within municipalities, non-compliance was slightly higher among mid-size units in central locations, as can be seen very clearly in Figure A.3. Across municipalities, non-compliance is not systematically correlated with municipality level characteristics, except for a slight negative correlation with population size.

In Appendix A.5, we explore whether the estimated effects are heterogeneous across municipalities. Figure A.4 plots the histogram of the point estimate for each individual municipality relative to the control group. It shows that all the estimates are distributed around the average estimated reported in this section. Table A.4 shows that there is some statistically detectable but *quantitatively* small differences in the estimates that are correlated with some municipality characteristics.

4.2 The effect of rent caps on the supply of rental units

4.2.1 Overall supply of rental units

We start our empirical investigation of the effect of the policy on the overall housing supply by running an ordinary least squares (OLS) and a Poisson model (PPML), using a pseudo-maximum likelihood algorithm for multiple high-dimensional fixed effects, following equations (1) and (2). Our dependent variable is the (ln) number of new contracts signed in each municipality each month.[26] Hence, this measure of volume is a flow measure. The main reason for using this variable rather than the stock of rentals in the market is that we have good data on the new contracts signed, but we lack data on the initial stock in the market at the beginning of our sample. Hence, we have many more pre-policy periods in the number of new contracts relative to the stock of rental units in the market, something that is important to check for differential pre-trends. Appendix A.6 shows an alternative measure of housing supply based on the probability that the units in our data have a

26. It is worth noting that when we use OLS, we miss roughly 50% of the observations, namely, the municipality - month pairs with zero new contracts. PPML estimates take into account this potential concern by including the zeros in the estimation. Montalvo (1997) discussed the use of pseudo-maximum estimation procedures for dynamic count data models.

tenant in each month.

In Figure IV, we show the evolution of the volume of new contracts for rentals, both in total and below and above the reference prices. The number of new contracts fluctuates around zero during the months prior to the policy. This is so when we allow for municipality linear time trends, otherwise there is a small upward trend in treatment municipalities relative to control. After the policy is implemented, however, we observe a very substantial drop of around 20%.[27] The graph that shows the evolution of rental contracts above reference prices declines by substantially more than this 20%. In contrast, the series for the new contracts below reference prices experiences a substantial increase after the policy change in treated relative to control municipalities. This increase is smaller, and hence does not compensate, for the drop in the volume of new contracts above the reference price, explaining the marked drop in the graph that shows the effects on the overall volume of new contracts.

We quantify these estimates in Table IV, using OLS (in Panel A) and PPML (in Panel B). The effect of rent caps on the number of new contracts is negative and statistically significant. The reduction is large for contracts above the reference price. There is also a significant increase in contracts below the reference price, but this effect does not compensate for the decrease in contracts above it. The estimates are stable to the inclusion of flexible province trends and to restricting the estimation sample to municipalities with fewer than 40,000 inhabitants. When comparing across panels, we see that taking the 0s into account (Panel B) reduces the point estimates somewhat, but the estimates are still precise and always statistically different than zero. The estimates shown in Table IV confirm that the number of new contracts decreased by around 10% to 20%. In the Appendix we explore an alternative measure of the supply of rental units, see Figure A.5. For this figure we compute the share of units that are rented at each point in time, and we compare this share across treated and control municipalities. This measure more closely reflects changes in the stock of units in the market, but has the limitation that we can use a smaller window around the treatment.[28] We find that the probability of being rented decreases by about two percentage points following the implementation of the policy.

Using the results reported so far, we can obtain estimates of the elasticity of new contracts with respect to price changes. We show these results in Table V. Panel A of Table V displays the "naive" least square estimates, for various specifications and subperiods. In general, however, when we regress the (log) number

27. Controlling for this small upward trend and using AHC rather than INCASOL data are the main differences between our approach and the one by Jofre-Monseny et al. (2022). As argued in the data section, INCASOL data is more noisy, which likely makes estimates less precise, making it harder to reject the null of no effect.

28. The window for this exercise is smaller because we need some months to stabilize the overall number of units in our data, due to the left censoring that emerges from the fact that many units are under contracts that started before our data begins. Another benefit of this measure is that it is computed at the housing unit level, which allows us to estimate the effect of the policy change using propensity score matching. We show these results in table A.5 in the Appendix.

of new contracts on (log) rental prices, we obtain small point estimates. The small estimates likely reflect the standard endogeneity problem when estimating demand or supply curves. For example, when demand increases, the supply may react and increase as well, resulting in small equilibrium price changes. Our setting, however, allows us to estimate the supply of rental units by exploiting the policy change. That is, the policy changes the price of housing. By tracking how this change affects the overall number of new contracts, we can estimate the rental market housing supply curve.

Panel B of Table V presents the instrumental variables (IV) estimates. Given the results shown so far, it should not be surprising that we obtain a positive estimate when using the policy change as an IV strategy. In the baseline specification with the full sample, a decrease of 1% in rental prices implies a decrease of around 2% in the supply of rental housing. We obtain a similar point estimate when we allow for flexible province-specific time trends (column 2) and when we restrict the sample to municipalities with fewer than 40,000 inhabitants (column 3). In Columns 4 to 6, we present estimates of this key elasticity when removing the lockdown period from the sample, as is also done in Table IV. Point estimates are, if anything, slightly larger than in the full sample, reaching an estimate of approximately three.

In Appendix A.5, we again explore whether these effects are heterogeneous across municipalities. Figure A.4 plots the histogram of the point estimate of the effect of the policy on volume of new contracts for each municipality. It shows that all the estimates are distributed around the average estimate reported in this section. Table A.4 shows that there are some detectable but quantitatively small differences across municipalities' characteristics in the estimated effects.

Overall, the evidence presented in this section suggests that the supply of housing units reacts to price changes in the market.

4.2.2 The effect along the distribution of excess price

While the previous subsections suggest that the supply of rental units is upward sloping, they also hint at the idea that the overall distribution of housing units in the market might have changed with the rental cap policy. We investigate this point in detail in what follows.

For that, we group all the contracts in our data by sextiles of excess price within treated and control municipalities, once we have removed municipality and time fixed effects. This creates six *equally* sized groups of housing units for both the treated and control municipalities as a function of its price relative to the reference price. The municipality and time fixed effects make these excess prices fully comparable across time and space.[29]

29. In practice, as can be seen in Figure A.10, it doesn't matter much whether we condition on time and location fixed effects, because we are essentially normalizing the rental prices by the reference price, hence excess prices are comparable across time and space. This approach is inspired in the work on minimum wages by Autor et al. (2015) and Cengiz et al. (2019).

Once we have these six groups, or sextiles, which summarize the entire distribution of excess prices, we plot in Figure V the change in prices and quantities, before and after the policy change, for each of these six groups. The price change results are shown in the graph on the left of the figure, Panel A. The results show that prices in the control municipalities increased uniformly along the entire distribution, by around 7%. In contrast, there is a clear gradient in rental price changes along the distribution in the treatment municipalities. Prices decline at the top of the distribution by as much as 4% in absolute terms and by close to 10% when we compare them to the control group. In contrast, prices increase above the control group at the bottom of the distribution.

In the right-hand graph of Figure V, in Panel A, we investigate whether the policy induced changes in the overall number of monthly contracts along the excess price distribution. The figure suggests a movement towards higher priced units in control municipalities. This movement is also observable in treatment municipalities up until the third sextile. After this, i.e., at the top of the distribution, we see a clear decline in the number of new contracts signed. Since all the groups are roughly similar, we can add up the different estimates to check that the overall supply in treatment municipalities declined relative to control ones. This graph shows very clearly the missing mass of new contracts that we estimated in the previous section above and below the reference price (see Figure IV and Table IV). These results indicate that high-price units leave the rental market in treated relative to control municipalities.

Panel B of Figure V repeats the exact same exercise as panel A, but assumes that the policy took place one year prior to the true implementation of the policy (and using the same post-treatment period length, but one year earlier). Hence, it provides placebo estimates that should reassure that the effects estimated in Panel A are a consequence of the policy and are not driven by pre-existing trends in different parts of the distribution. Both the left and the right hand side graph show that price and quantity changes were similar in treatment and control municipalities along the entire excess price distribution prior to the implementation of the policy.

Overall, the evidence in this section suggests that the introduction of reference prices leads to heterogeneous effects along the excess price distribution, both in prices and supply. It is worth noting, perhaps, that any methodology that computes unit-specific reference prices will lead to an excess price distribution, and potentially effects like the ones reported in this and the previous sections.

4.2.3 Spillovers

Arguably, the most important result in Mense et al. (2023) is the evidence of spillovers from units facing price caps to those that did not. In our context, there may be two type of spillovers. On the one hand, there may be spillovers from treated to non-treated units *within* treatment municipalities, as we have already seen. Indeed, the effects along the entire excess price distribution can be interpreted as a form of spillover, something that the model (see Section 5) makes even more explicit. On the other hand, there may also be spillovers *across*

municipalities, something that we have implicitly assumed away in our empirical strategy. We check this second source of spillovers in this section. We do so using two strategies. First, we examine whether there is a gap in excess prices before and after the reform in rental units in treated municipalities that share a border with non treated ones. Figure A.11 shows that there is no gap in excess prices at the border prior to the policy reform. After the policy, a gap opens for units on each side of the border. With spillovers from treated to non treated municipalities, one would expect the gap around the border to be larger than suggested by the difference-in-differences estimates (as long as demand moves to non treated municipalities). We estimate a gap of similar magnitude than with our difference-in-differences estimates, which indicates an absence of spillovers. Our second strategy is to compare the price and volume changes for units in non treated municipalities but close to treated ones, relative to non treated municipalities that are far from the treatment. Figure A.12 provides evidence which is, again, consistent with no spillovers between treatment and control.[30]

5 Model

In the previous section we have presented empirical evidence documenting how the policy affected rental prices and supply. We have shown that both prices and supply declined with the policy. We have also documented rich patterns along the entire excess price distribution of both prices and supply. Relative to the control group, low-price units' prices, which we defined as being to the left of the excess price distribution, increased with the policy, while high-price units' prices declined. At the same time, we documented that the supply of rental units declined mainly due to high price units exiting the rental market (or not entering as they were prior to the policy).[31]

In this section, we introduce a model that rationalizes these empirical results and allow us to evaluate the welfare consequences of the policy. In a standard perfect competition model with homogeneous housing, introducing a binding price cap has several consequences. A cap on prices generates a gap between the amount of units that potential tenants would demand and how many units landlords would be willing to offer in the rental market. As a result, at the price cap, some potential renters would want to rent, but the supply does not satisfy this demand. Moreover, as long as the supply of rental units is not perfectly inelastic, there is a deadweight loss associated with the policy. Hence, a rental price cap policy can be understood as a trade-off between efficiency

30. In Figure A.6 in the Appendix, we investigate whether there is any heterogeneity in our results for units where we observe multiple contracts and for units where we only observe one contract in our data. We show no meaningful difference in the estimates, in line with the idea that there was no scope for spillover effects in the case of our policy given that reference prices were applied to all units in treated municipalities.

31. We also documented some degree of non-compliance. We abstract from this result in the model, although it can be easily incorporated by assuming that some landlords of high quality units comply with the cap, and some do not.

and redistribution from landlords to renters (who stay in the rental market), as shown graphically in Appendix B.[32] These type of policies can be justified in terms of equity, since homeowners tend to be richer than renters.

While the standard perfect competition model with homogeneous housing may provide a useful benchmark, there are several aspects of the evidence presented above that do not fit well into the framework and justify expanding this standard model. First, in the data we see substantial heterogeneity in prices across units in the market, whereas in the standard framework there is a unique price. This is a well known feature in many data sets which is robust to conditioning on a large host of observable characteristics (see for example Bierdel et al. (2022)), or, in our case, when we deflate prices by rental prices of nearby units, as we do with our measure of excess price. Second, the results indicate interesting effects of the policy along the entire distribution, both in terms of price and supply changes. A simple homogeneous housing framework with perfect competition is ill-suited to speak to these findings since, again, there is a unique price in the market.

For these and potentially other reasons, it may be useful to depart from the simple benchmark. In this section, we extend the standard perfect competition model by considering quality differences. This assumes that the observed price heterogeneity in the data is due to heterogeneity in unit characteristics that is valued (and perfectly observed) in the market, and that units with different quality are imperfectly substitutable from the view point of renters. We use this model to quantify the welfare consequences of the policy change.

The perfect competition model with heterogeneous housing makes two, potentially restrictive, assumptions. On the one hand, the model assumes perfect competition, i.e. that renters and landlords in the model take prices as given. While the ownership structure of housing units is not very concentrated in Catalonia, as documented in Table II, it may still be the case that preferences for particular units are very specific, in which case, landlords could potentially internalize the demand for their units. On the other hand, there may be search frictions that make reality depart from the perfect competition and perfect information framework. Perhaps renters do not fully know the price and characteristics of all the units in the market and this can be exploited by landlords to set prices differently than how prices are set with full information. We explore these two possibilities formally in Appendix C and D and we summarize the findings in subsection 5.7, below.

5.1 Demand for rental units

We assume that there are two types of rental units, high- and low-quality, and an outside option, such as buying or moving to another economy. Potential tenants decide among high- and low-quality units and the outside option based on the utility they obtain from each option. This can be represented as:

32. We show a graphical representation of the perfect competition model in Appendix B. The graph shows the initial equilibrium $\{p^0, Q^0\}$ and how a cap on price moves the equilibrium to $\{\underline{p}, \underline{Q}\}$, leaving an unsatisfied demand from $\underline{Q}$ to Q^D.

(3) $$\ln u^q - \ln p^q + \varepsilon_{iq}$$

where u^q denotes the utility of living in a unit of quality q, p^q denotes its price, and ε_{iq} is an individual (indexed by i) idiosyncratic taste shock for quality q drawn from an extreme value distribution. Under this assumption, by the law of large numbers, the number of potential tenants that demand quality q, indicated as D^q, is given by the following expression:[33]

$$D^q = \left(\frac{(u^q/p^q)^{1/\lambda}}{(u^l/p^l)^{1/\lambda} + (u^h/p^h)^{1/\lambda} + (b/p^l)^{1/\lambda}} \right) N$$

where N denotes the total number of persons in the economy and b indicates the value of the outside option – denominated in low-price metric. λ governs the rental demand elasticity, i.e. by how much demand changes for a given change in rental prices in market q.

To ease exposition, we keep quality differences to two categories. However, it is worth emphasizing that the model can easily accommodate multiple quality categories. In that case the denominator would be the sum over all these quality categories. The results that we derive in what follows would be almost identical.

5.2 Supply of rental units

Landlords either own a high- or a low-quality unit and put their unit of quality q in the rental market according to an upward sloping supply function:

$$S^q = \left(\frac{p^q}{C} \right)^{1/\gamma} L_q$$

where C is the cost of supplying a unit in the rental market and L_q is the mass of landlords who have a unit of quality q. To simplify, and without loss of generality, we assume $L_q = 1$.[34] Also to simplify, we assume that the rental housing supply elasticity (γ) is the same across quality types. We discuss later the places where this simplification matters for the results.

5.3 Equilibrium

The equilibrium is determined by equating the demand and supply equations in both the high- and low-quality markets:

33. Anderson et al. (1992) provide an early and beautiful introduction to discrete choice models.

34. Note that L_q and C can be used interchangeably. This equation can also be micro-founded using discrete choice theory. In that case, C would also include p_q which makes it a little bit inconvenient for the derivations to follow. This explains why we assume directly an iso-elastic rental supply curve.

$$\left(\frac{(u^l/p^l)^{1/\lambda}}{(u^l/p^l)^{1/\lambda}+(u^h/p^h)^{1/\lambda}+(b/p^l)^{1/\lambda}}\right)N=\left(\frac{p^l}{C}\right)^{1/\gamma};\quad \left(\frac{(u^h/p^h)^{1/\lambda}}{(u^l/p^l)^{1/\lambda}+(u^h/p^h)^{1/\lambda}+(b/p^l)^{1/\lambda}}\right)N=\left(\frac{p^h}{C}\right)^{1/\gamma}$$

By dividing the two equilibrium conditions we obtain that relative prices depend on relative utilities:[35]

$$\left(\frac{p^h}{p^l}\right)=\left(\frac{u^h}{u^l}\right)^{\frac{\gamma}{\lambda+\gamma}}$$

Hence, we obtain that $p^h=\Omega p^l$, with $\Omega\equiv(\frac{u^h}{u^l})^{\frac{\gamma}{\lambda+\gamma}}>1$ (as long as $u^h>u^l$). This means that, in equilibrium, high quality units are more expensive.

We can now substitute for p^h in the equilibrium condition for the low-quality market and obtain:

$$p^l=\left(\frac{u^l}{U}\right)^{\gamma/\rho}N^\gamma C$$

where $U\equiv(u^l)^{\frac{1}{\rho}}+(u^h)^{\frac{1}{\rho}}+(B)^{\frac{1}{\rho}}$, $\rho=\lambda+\gamma$, and $B\equiv b(u^l)^{\frac{-\gamma}{\lambda+\gamma}}$.

With this expression we can also obtain the equilibrium price in the high-quality submarket:

$$p^h=\left(\frac{u^h}{U}\right)^{\gamma/\rho}N^\gamma C$$

and the equilibrium quantities, which are given by:

$$Q^l=\left(\frac{u^l}{U}\right)^{1/\rho}N,\qquad Q^h=\left(\frac{u^h}{U}\right)^{1/\rho}N$$

These four equations determine the equilibrium in the market as function of exogenous parameters (fundamentals and elasticities). They also capture many standard forces. For instance, more people in the market will lead to more housing services and higher prices in equilibrium. Similarly, higher costs for landlords to supply units to the rental market will translate into higher equilibrium prices. Better outside options (i.e. higher B) will translate into lower prices and quantities in the rental market.

5.4 A rental price cap

Suppose a policy maker introduces a price cap (P) that is only binding for high price units. Then, the supply of high-quality units at price P is given by:

35. This result depends, obviously, on the assumption that the supply cost C and the supply elasticity (γ) is the same in the high- and low-quality markets. Heterogeneity in supply elasticities is harder to handle, but heterogeneity in the cost (C) is equivalent to re-expressing utilities in each submarket as a function of costs.

$$S^h = \left(\frac{P}{C}\right)^{1/\gamma}$$

But given this supply and high-quality prices, there will be excess demand. We can denote by $\tilde{p}^h$ the price that buyers would be willing to pay if the supply of units of high quality is determined by the binding price cap P. We can obtain $\tilde{p}^h$ using the following equation:

$$\left(\frac{(u^h/\tilde{p}^h)^{1/\lambda}}{(u^l/p^l)^{1/\lambda} + (u^h/\tilde{p}^h)^{1/\lambda} + (b/p^l)^{1/\lambda}}\right) N = \left(\frac{P}{C}\right)^{1/\gamma}$$

This expression depends on the (endogenous) low-quality prices. Hence, as before we need the equilibrium condition for the low-quality market, which is given by:

$$\left(\frac{(u^l/p^l)^{1/\lambda}}{(u^l/p^l)^{1/\lambda} + (u^h/\tilde{p}^h)^{1/\lambda} + (b/p^l)^{1/\lambda}}\right) N = \left(\frac{p^l}{C}\right)^{1/\gamma}$$

Following the same steps as before we can express the willingness to pay for quality as a function of prices of low-quality units as follows:

$$\tilde{p}^h = \left(\frac{u^h}{u^l}\right)(p^l)^{\frac{\gamma+\lambda}{\gamma}} P^{-\frac{\lambda}{\gamma}}$$

Note that we can always write $P = \eta\tilde{p}^h$, with $\eta < 1$. That is, the rental cap price is set at a fraction η of the high price that potential renters would be willing to pay. In this case we can write:

$$\tilde{p}^h = p^l \left(\frac{u^h}{u^l}\right)^{\frac{\gamma}{\lambda+\gamma}} \eta^{-\frac{\lambda}{\gamma+\lambda}}$$

where this $\tilde{p}^h$ is, again, the willingness to pay for high-quality units when supply is determined by price P. Hence, we have that:

$$\tilde{p}^h = p^l \Omega \eta^{-\frac{\lambda}{\gamma+\lambda}}$$

where Ω is defined as before, i.e. $\Omega \equiv (\frac{u^h}{u^l})^{\frac{\gamma}{\lambda+\gamma}} > 1$.

Using this into the equilibrium condition for the low-quality market, we obtain:

$$p^l = \left(\frac{(u^l)^{\frac{1}{\rho}}}{(u^l)^{\frac{1}{\rho}} + (u^h)^{\frac{1}{\rho}} \eta^{\frac{1}{\gamma+\lambda}} + (B)^{1/\rho}}\right)^{\gamma} N^{\gamma} C$$

Hence, we can express the equilibrium conditions *with* a binding rental price cap in the high-quality submarket with the following four equations:

$$p^l = \Gamma^{\gamma} \left(\frac{u^l}{U}\right)^{\gamma/\rho} N^{\gamma} C; \quad p^{h*} = P; \quad Q^l = \Gamma \left(\frac{u^l}{U}\right)^{1/\rho} N; \quad Q^h = \left(\frac{P}{C}\right)^{1/\gamma}$$

where $\Gamma \equiv \frac{(u^l)^{\frac{1}{\rho}}+(u^h)^{\frac{1}{\rho}}+(B)^{1/\rho}}{(u^l)^{\frac{1}{\rho}}+(u^h)^{\frac{1}{\rho}}\eta^{\frac{1}{\gamma+\lambda}}+(B)^{1/\rho}} > 1$.

Given these equations, it is easy to see that the rental cap decreases prices and quantities in the high-quality market and increases them in the low-quality one, since $P < p_h$ and $\Gamma > 1$. The term η in Γ, which is the gap between the reference price P and the willingness to pay for high-quality units when supply is determined by price P, generates all the difference between this and the unconstrained equilibrium. It is also worth emphasizing that, with the outside option, the increase in the number of units in the low-quality market does not compensate the decline of units in the high-quality market. Without outside option, we could have no changes in the overall number of units in the rental market, with large distributional consequences. This later result also depends on whether the rental housing supply elasticity is the same across quality markets or not. We provide further intuition for this model in the following subsection.

5.5 Graphical representation

The perfect competition model with quality differences can be represented in a set of simple graphs, which are shown in Figure VI. The figure is divided into two panels. Panel A shows the equilibrium in the market when there is no rental cap. Panel B shows the new equilibrium once a rental cap is introduced which is only binding for the high-quality units.

Initially, the equilibrium is at points A and B. At these points, prices for high-quality units are higher than for low-quality ones, reflecting the higher value of high-quality. A reference price P is more likely to be binding in the high-quality market. At that price, as shown in Panel B, fewer owners of high quality units will be willing to supply their unit into the market, hence, the amount of high quality units in the market (Q^h) will be lower with the price cap.

The price cap will generate excess demand for high-quality units. This excess demand can be computed by looking at the location of $\tilde{p}^h$ in panel B, which is the price that renters would be willing to pay at the limited supply. This price helps to determine how much demand moves to the low-quality market, or, potentially, to the outside option. This is shown in Panel B as a modest shift in the demand for low-quality units. This increase in the demand for low-quality units leads to increases in both prices and quantities in the low-quality market.

Hence, these two sets of graphs show the essential forces behind the model. A rental cap is more likely to bind in the high-quality market. The cap restricts supply and lowers prices among high-quality units. Some renters that would have otherwise entered the high-quality submarket move towards low-quality units, and others chose their outside option, i.e. buying, living with the family, or moving somewhere else. The higher demand for low-quality units, in turn, translates into higher prices at the bottom of the rental price distribution. The patterns in the data are well aligned with these predictions.

5.6 Welfare analysis

Following Figure VI, it is relatively simple to compute the welfare consequences of the policy using this framework. For instance, we can compute the welfare gains of renters in the high quality submarket (i.e. those who were and stayed in the high quality submarket), which we label by ΔW^h, using the following expression:

$$\Delta W^h \approx -\Delta V^h \approx -\frac{\Delta p^h}{p^h} p^h Q^h > 0$$

This expression says that the welfare change among high-quality renters is given by the percentage change in high-quality prices multiplied by the size of the high-quality market, measured in value. For instance, according to Figure V, prices in the high-qualtiy market declined by around 7 to 10 percent, and the size of the high quality market in the treatment municipalities is around 3,600 million euros in 2019.[36] Hence, the welfare gain of high-quality renters is equivalent to around 252 to 360 million euros per year (in 2019 Euros). Moreover, as we can see following Figure VI, this welfare change is approximately the opposite of the one experienced by owners of high-quality units (labeled as ΔV^h). In other words, landlords of high quality units loose the welfare equivalent of about 252 to 360 million euros per year.[37]

Similarly, the renters' welfare loss of those initially in the low-quality submarket (ΔW^l) is similar to the welfare gain of landlords in the low-quality submarket (ΔV^l). They can both be approximated by:

$$\Delta W^l \approx -\Delta V^l \approx -\frac{\Delta p^l}{p^l} p^l Q^l < 0$$

This expression says that the welfare change can be approximated by the percentage change in low-quality prices multiplied by the size of the low-quality market. In this case, the data suggest that low-quality prices increased by around 3 or 4 percent, which we can multiply by the size of the low quality market in treated areas which is around 75 million euros per year.

Finally we can also approximate the value of the deadweight loss in the high quality market using the following expression:

$$DWL^h = -\frac{1}{2} \frac{\Delta p^h}{p^h} \frac{\Delta Q^h}{Q^h} p^h Q^h$$

36. We compute the total size of the high- and low-quality market as follows. First, we take the rental prices of 2019, i.e. the year prior to the application of the rental cap policy. We use these units to estimated the high-quality and low quality proportion in each of the controlled municipalities, defined as units above and below the reference price. Next, we obtained from the Incasol the stock of all the houses rented in Catalonia in 2021, and applied the estimated proportion of high-quality and low-quality to each municipality. Finally, we summed the annual rental revenue of each of the segments of the market across municipalities.

37. It is worth noting that we are making approximations here because we omit some of the smaller triangles, and we are a bit loose in computing the size of the high- and low-quality market using either the period prior or post-policy change.

In this case, the expression says that the deadweight loss in the high-quality market is equal to one half of the percentage change in high-price units (i.e. around 7 to 10 percent), multiplied by the percentage change in the size of the high-quality market – measured in housing units –, which is around 25 percent, all this multiplied by the size of the high-quality market (i.e. around 3,600 million). Hence the deadweight loss from the policy is of around 36 million Euros per year.

It is worth noting that these numbers assume a non-random allocation mechanism, by which the renters most willing to pay for quality are the ones allocated to high-quality units once the price cap is in place. We could incorporate other allocation mechanisms following Glaeser & Luttmer (2003). Under these alternative mechanisms, our estimates would be a lower bound, since we would also need to take into account the potential misallocation induced by the policy.[38] It is also worth noting that part of the price dispersion observed in the data may be due to search frictions. If so, we would be overestimating the overall welfare losses from the policy, but we would be estimating adequately the relative ranking across groups of renters, as we discuss in the subsequent section and the appendix, where we also argue that part of the empirical evidence is harder to rationalize with just search frictions.

It is worth highlighting the main conceptual departure of our model relative to other, more standard, frameworks. While from the perspective of other models, rent cap policies redistribute surplus from landlord to renters, potentially with a deadweight loss, our model has more nuanced implications. Some landlords gain from the policy (those who own low-quality units), while others loose, and the same applies to tenants, even those who stay in the market, who do not uniformly benefit from the policy.

5.7 Discussion of alternative models

As mentioned before, there are two assumptions in the perfect competition model presented in the previous subsections that may be viewed as strong assumptions. First, there may be instances where home owners may be able to internalize the demand for rental units, and hence, exert some degree of market power. While this is probably more prevalent in markets where the ownership structure is more concentrated than what we see in the data, see Table II, it may still be the case that landlords have some market power. This is important because in contexts where market power is prevalent rental caps may lead to price declines without affecting supply, similar to what minimum wages do in monopsonistic labor markets.

In Appendix C, we study the predictions of a model with quality differences and market power. The main point we make is that this alternative model generates some predictions that are not in line with the empirical exercise. With quality differences, the rental cap is likely to be more binding in the high-quality market than

38. We could potentially also incorporate rent-seeking forces following Bulow & Klemperer (2012), given their analysis of controlled and uncontrolled markets. We do not incorporate these considerations because these would require additional assumptions.

among low-quality units. Hence, the rental cap is likely to eliminate market power among owners of high-quality units. If the rental cap is binding but not too high, this should lead to price declines among high-price units without affecting the supply of high-quality units in the market. This prediction is inconsistent with the evidence presented above. Moreover, with the rental cap policy, high-quality units become more attractive, and hence demand should move away from low-quality housing towards higher quality one. As a result, prices in low-quality units should also decline, and again, this is in contradiction to what we observe in the data.

A second assumption that may be viewed as too strong is that perfect competition requires perfect information. In other words, assuming that agents take prices as given implicitly means that any price dispersion we observe in the data must be the result of differences in unit characteristics. While the variance in housing prices is smaller when one controls for observable characteristics, see for example Bierdel et al. (2022), there is a substantial amount of residual variation that may either reflect unobservable (to the econometrician) variation in unit characteristics or that may, instead, reflect search frictions. It may be that someone would prefer a particular unit in the market at a certain price, but that she is unaware of the existence of that unit, and hence, settles for a different option.

In Appendix D we investigate the role of search frictions. In particular, directed search models of the labor market have been used to show that wage differences can be sustained in equilibrium, even among identical workers. The main idea in these models is that workers with identical characteristics may look for jobs that pay more but are harder to find, or settle instead for jobs that pay less but are easier to find. It is also optimal for firms to post vacancies that either pay more or less for the same type of workers, because again, it may be to the firm's advantage to pay little and reduce on labor costs, but have a non-filled vacancy for a longer time period. In these models, the key is whether workers are allowed to apply to different types of jobs.

Hence, directed search models can potentially be a useful framework to analyze rental markets where, as we also see in labor markets, there is substantial variation in prices. As such, and to build intuition, we start in Appendix D by studying a directed search model of the rental market with homogeneous housing. We show that in this model a rental cap policy that acts as a reference price may help to collapse the various submarkets into a unique market, and that this may be welfare enhancing. Intuitively, if low-price rental units co-exist in equilibrium with high-price units, with the exact same characteristics, renters may be sending too many applications, increasing the queue length for forming matches. A reference price solves this inefficiency.

A rental price cap in a directed search model with homogeneous housing results in predictions that are to some extend in line with the data. For instance, the model predicts that rental prices will converge to a unique price, consistent with the empirical evidence that high-price units experience price declines while low-price units experience price increases following the introduction of reference prices. Moreover, as long as the rental price cap is not too high, the model also predicts that supply will change, driven primarily by the exit of high-price

units. There is, however, one piece of empirical evidence that does not fit well into the framework. In the data we observe substantial (excess) price dispersion in treated municipalities even after the implementation of the policy. This piece of evidence suggests that it is very likely that at least some of the price variation is due to quality differences.

For this reason, in Appendix D we also study a directed search model with quality differences. We show that the predictions of this more realistic model do not match what we observe in the data. In particular, it is again more likely that rental cap policies are (at least more) binding among high-quality units. Hence, we would expect price convergence, especially among high-price units, in contrast to what we see in the data. Moreover, because a rental price cap is welfare enhancing in markets where it is binding, we should see both supply and demand move towards markets where the policy is binding, i.e. towards high-quality units. This should alleviate price pressures among low-price units, something that we do not see in the data. Finally, we should also see that the time between rental contracts declines and flattens – at least in the top end of the distribution. We investigate whether this is the case in Figures A.7, A.8, and A.9. When comparing treatment and control municipalities, we see that the time in between contracts does not seem to change differentially between treatment and control municipalities, neither on average, nor for high- or low-price units.

Overall, we conclude that the perfect competition framework with quality differences is well equipped to rationalize the empirical results documented in the first part of the paper, while other assumptions seem to lead to predictions that are not supported by the data. This gives us confidence that our welfare calculations are an accurate estimate of the consequences of the policy – as long as we abstract from potential misallocation concerns (Glaeser & Luttmer, 2003).[39]

6 Conclusions

In this paper, we analyze the impact of the rent control enacted in Catalonia in September 2020. We find that the introduction of rental price controls led to a reduction in both prices and the supply of rental units. We estimate that the policy change led to price reductions of around 5%, and that the fraction of units above the reference price, which was around 50% prior to the policy, was reduced by between 13 and 22 percentage points. This evidence suggests that the policy was effective in reducing rental prices, although there also was some degree of non-compliance.

39. Something that we do not explore in this paper are the dynamic effects of the rent control policy implemented in Catalonia. There is a number of papers that make the point that rent controls alter or shape the dynamics in housing markets, see for example Rotberg & Zhang (2024); Basu & Emerson (2000); Favilukis et al. (2023). However, the policy in Catalonia was meant to be temporary, and in fact, it was in place only for about eighteen months. Hence, it is beyond the scope of our paper to explore the effect of this temporary policy on long-run outcomes.

The supply of rental units in the market also changed as a consequence of the policy. We estimate that the overall supply of rental units declined by around 10%. This number masks substantial shifts in the composition of rental units in the market. We document a very pronounced decline in the units above the reference price that was not compensated by an increase in units below.

Finally, we argue that a perfect competition model with quality differences can rationalize the evidence and provides a framework to estimate the welfare consequences of the policy. Rental price caps are likely more binding for high-quality units. With the new prices owners of high-quality units exit the market. With this limited supply, some renters that would otherwise have been in a high-quality unity move their demand to the low-quality market. This, in turn, leads to price increases in low-price units, as observed in the data.

References

Almagro, M., & Dominguez-Iino, T. (2022). Location sorting and endogenous amenities: Evidence from Amsterdam. *mimeo*.

Anderson, S., De Palma, A., & Thisse, J.-F. (1992). Discrete choice theory of product differentiation. *MIT Press*.

Arkhangelsky, D., Athey, S., Hirshberg, D., Imbens, G., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, *111*(12), 4088-4118.

Autor, D., Manning, A., & Smith, C. (2015). The contribution of the minimum wage to u.s. wage inequality over three decades: A reassessment. *American Economic Journal: Applied Economics*.

Autor, D., Palmer, C., & Pathak, A. (2014). Housing market spillovers: evidence from the end of rent controls in Cambridge, Massachusetts. *Journal of Political Economy*, *122*(2), 661-717.

Basu, K., & Emerson, P. (2000). The economics of tenancy rent control. *Economic Journal*.

Bierdel, A., Drenik, H. J., A., & Ottonello, P. (2022). Asymmetric information and capital accumulation. *mimeo*.

Breidenbach, P., Eilers, L., & Fries, J. (2022). Temporal dynamics of rent regulations - the case of german rent control. *Regional Science and Urban Economics*, *92*, 103737-.

Bulow, J., & Klemperer, P. (2012). Regulated prices, rent seeking, and consumer surplus. *Journal of Political Economy*, *120*(1), 160-186.

Callaway, B., & Sant'Anna, P. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*, 200-230.

Carozzi, F. (2020). Credit constraints and the composition of housing sales. farewell to first-time buyers? *Journal of the European Economic Association*.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*.

Couture, V., Gaubert, C., Handbury, J., & Hurst, E. (2023). Income growth and the distributional effects of urban spatial sorting. *Review of Economic Studies*.

de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964-2996.

Diamond, R., McQuade, T., & Quian, F. (2019). The effect of rent control expansion on tenants, landlords and inequality: evidence from San Francisco. *American Economic Review*, *109*(9), 3365-3394.

Downs, A. (1988). Residential rent controls: An evaluation. *DC: Urban Land Institute*.

Favilukis, J., Mabille, P., & Van Nieuwerburgh, S. (2023). Affordable housing and city welfare. *The Review of Economic Studies*, *90*(1).

Fetter, D. (2016). The home front: Rent control and the rapid wartime increase in home ownership. *Journal of Economic History*.

Frankena, M. (1975). Alternative models of rent control. *Urban Studies*.

Friedman, M., & Stigler, G. (1946). *Roofs or ceilings?* Foundation for Economic Education.

Galenianos, M., & Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, *144*(2), 445-471.

Garcia-Lopez, M.-A., Jofre-Monseny, J., Martínez-Mazza, R., & Segú, M. (2020). Do short-term rent platforms affect housing markets? evidence from airbnb in barcelona. *Journal of Urban Economics*, *119*.

Glaeser, E. L., & Luttmer, E. F. P. (2003). The misallocation of housing under rent control. *American Economic Review*, *93*(4), 1027-1046.

Gyourko, J., & Linneman, P. (1989). Equity and efficiency aspects of rent control: An empirical study of new york city. *Journal of Urban Economics*, *26*, 54–74.

Jofre-Monseny, J., Martinez, R., & Segu, M. (2022). Effectiveness and supply effect of high-coverage rent control policies. *IEB Working Paper 2022/02*.

Landvoigt, T., Piazzesi, M., & Schneider, M. (2015). The housing market(s) of san diego. *American Economic Review*.

Mense, A., Michelsen, C., & Kholodilin, K. (2019). The effect of second-generation rent control on land values. *American Economic Review Paper and Proceedings*, *109*(3), 385-388.

Mense, A., Michelsen, C., & Kholodilin, K. (2023). Rent control, market segmentation, and missallocation: Causal evidence from a large-scale policy intervention. *Journal of Urban Economics*.

Montalvo, J. G. (1997). GMM estimation of count panel data models with fixed effects and predetermined instruments. *Journal of Business and Economic Statistics*, *15*(1), 82-89.

Moon, C.-G., & Stotsky, J. G. (1993). The effect of rent control on housing quality change: A longitudinal analysis. *Journal of Political Economy*, *101*(6), 1114-1148.

Olsen, E. (1972). An econometric analysis of rent control. *Journal of Political Economy*, *80*(6).

Piazzesi, M., Schneider, M., & Stroebel, J. (2020). Segmented housing search. *American Economic Review*.

Rotberg, S., & Zhang, L. (2024). The welfare consequences of rent controls. *mimeo*.

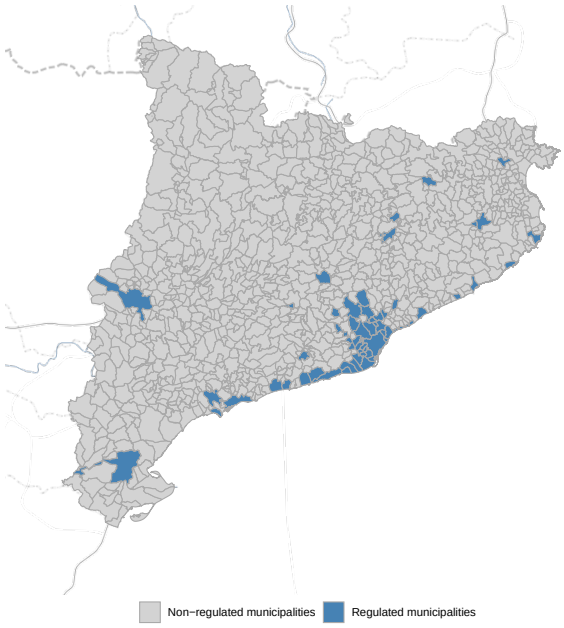
Sims, D. (2007). Our of control: what can we learn from the end of massachusetts rent control? *Journal of Urban Economics*, *61*(1), 129-151.

Suen, W. (1989). Rationing and rent dissipation in the presence of heterogeneous individuals. *Journal of Political Economy*, *97*(6), 1384-94.

Wright, R., Kircher, P., Julien, B., & Guerrieri, V. (2021). Directed search and competitive search: A guided tour. *Journal of Economic Perspectives*, *59*(1), 90-148.

7 Figures

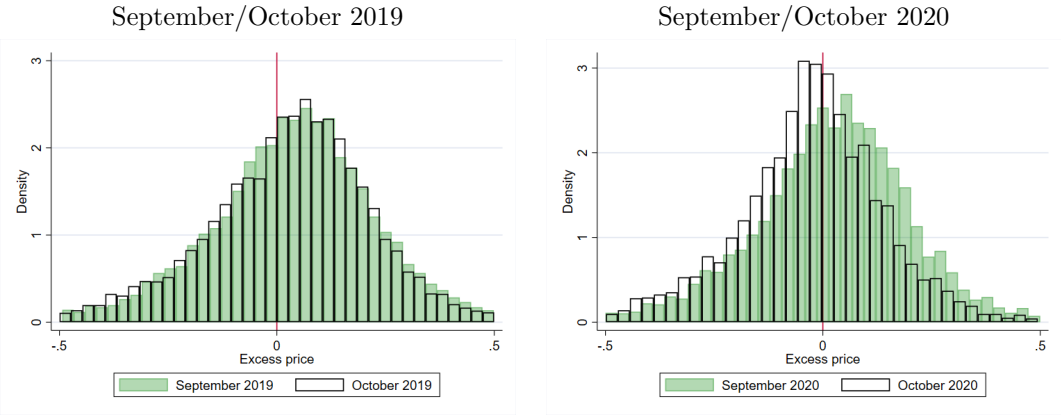
Figure I: Map of municipalities affected by rental price caps



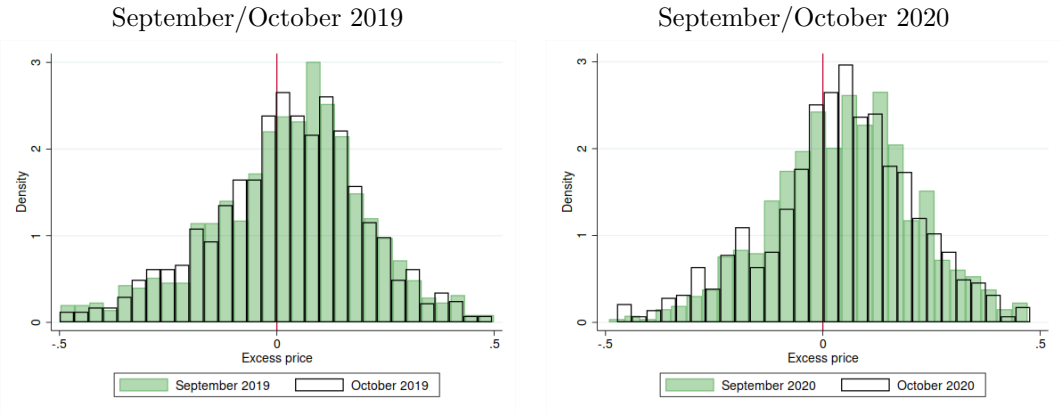
Notes: These maps show treated, in blue, and control municipalities, in gray. Treated municipalities are defined as municipalities that were subject to a rental price cap.

Figure II: Histograms of excess prices around the policy change date versus the previous year

Panel A: Treated municipalities



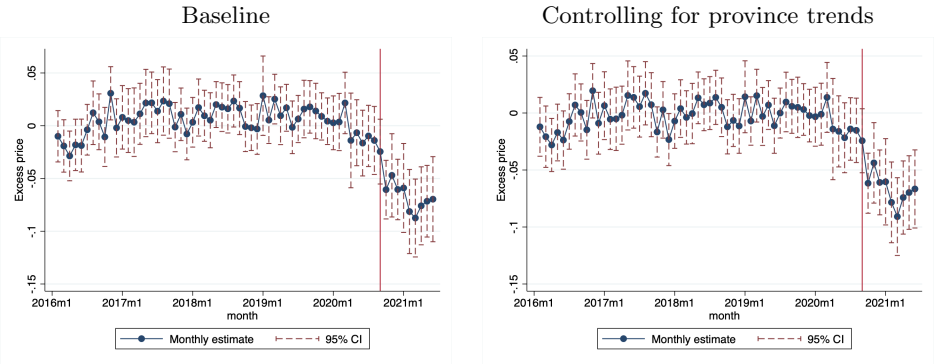
Panel B: Control municipalities



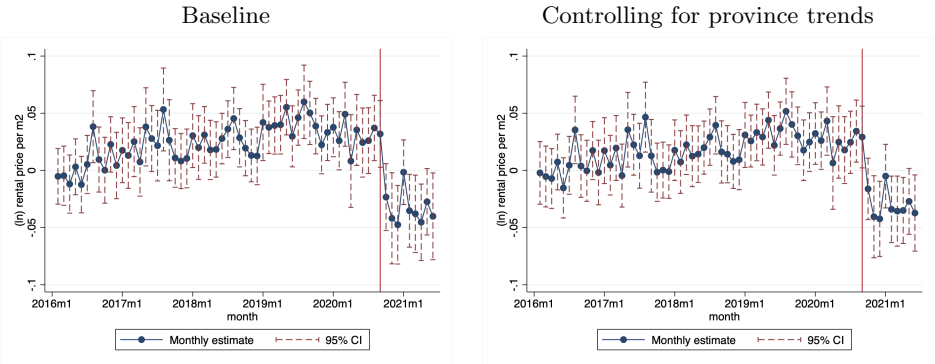
Notes: These graphs show the histograms of Excess price in treatment and control municipalities during selected months around the policy change and in the preceding year. Excess price is defined as the log difference between the rental price and the reference price calculated for each unit.

Figure III: Dynamic estimates of the effect of the policy on rental prices

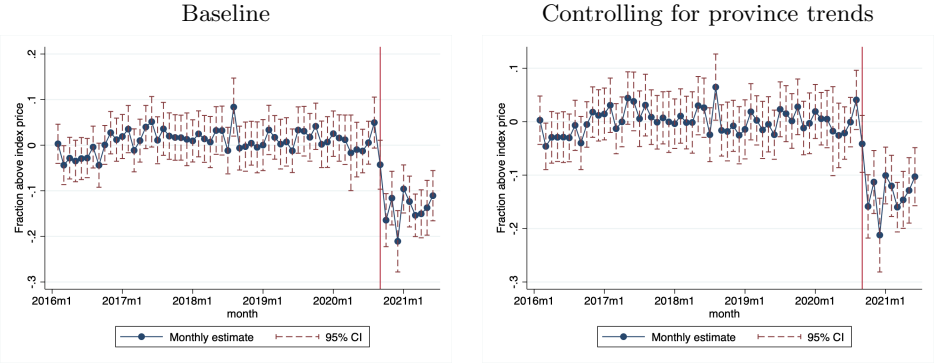
Panel A: Outcome variable: Excess price



Panel B: Outcome variable: (ln) rental price per m2

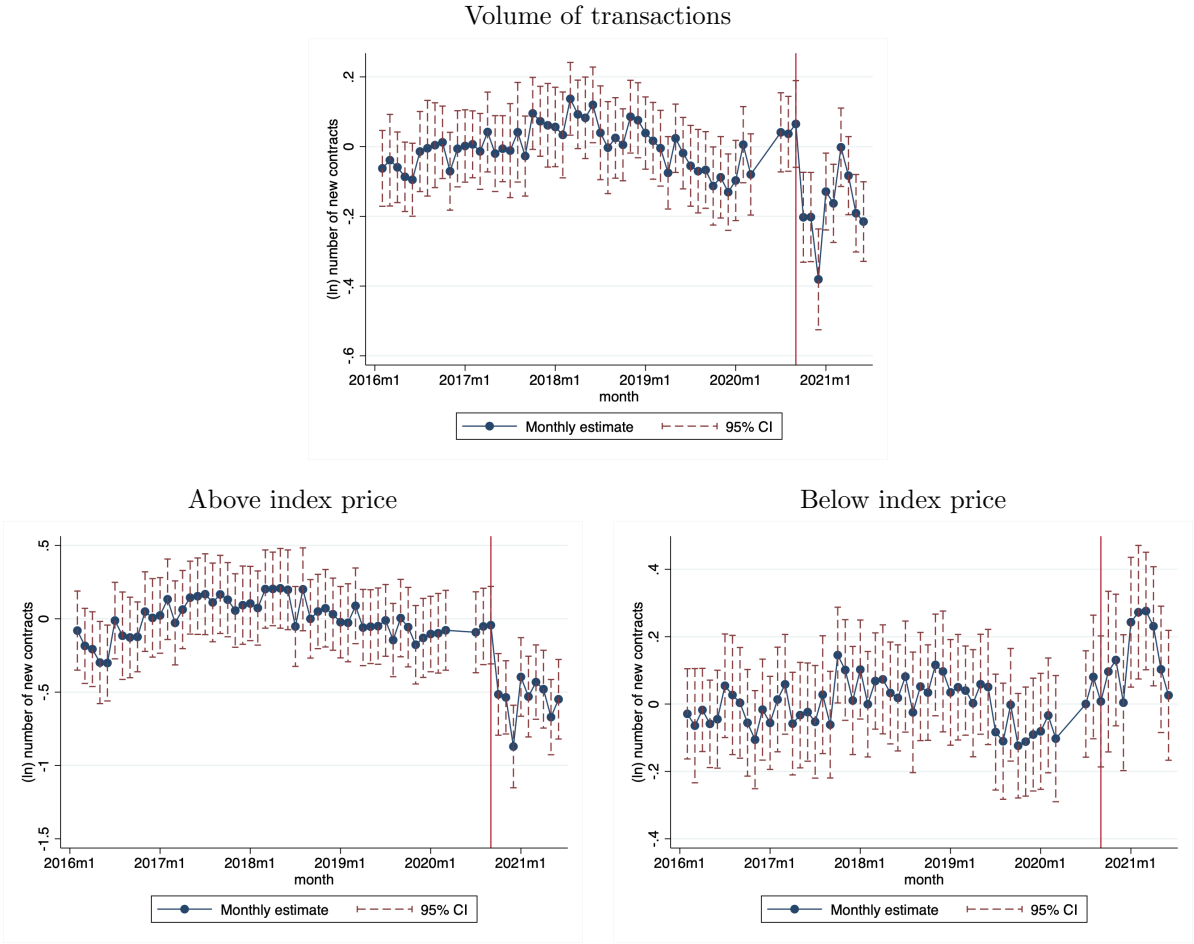


Panel C: Outcome variable: Fraction above the index price



Notes: This figure shows two types of graphs of the effects of the policy on three different outcome variables, as denoted in the titles. The graphs on the left show our baseline specification, while the graphs on the right show our baseline specification controlling for flexible province-specific time trends. Ninety-five percent confidence intervals of robust standard errors clustered at the municipality level are reported.

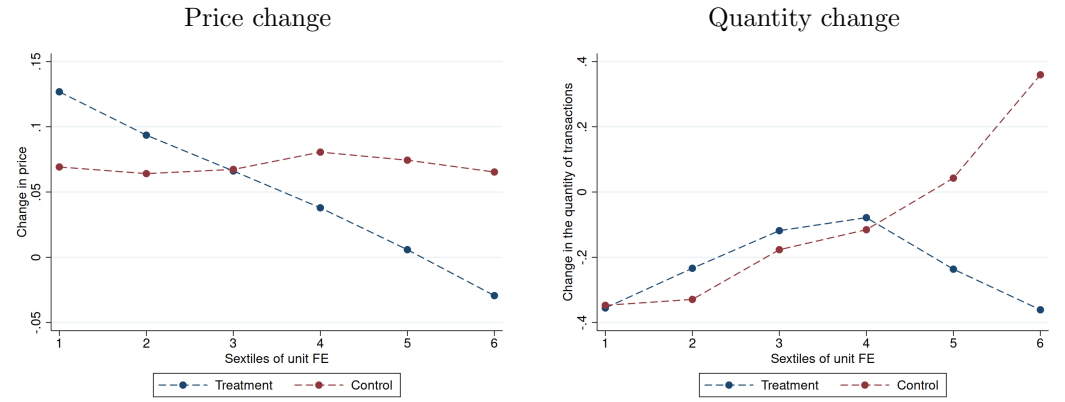
Figure IV: Dynamic estimates of the effect of the policy on the volume of transactions



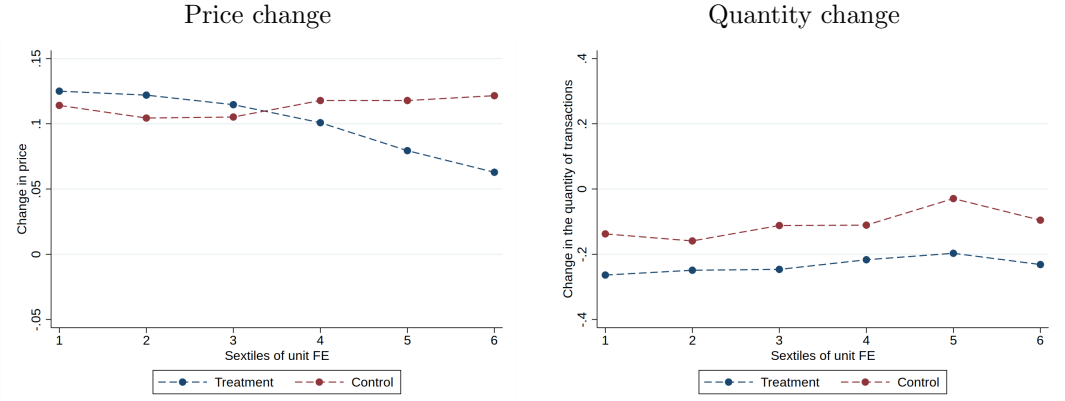
Notes: These figures show graphs of the effect of the policy on the (log) number of the overall number of transactions and the overall (log) number of transactions above and below the reference price in the treatment relative to the control municipalities. These graphs use our baseline specification, controlling by flexible province-specific time trends and allowing for municipality-specific linear time trends estimated prior to the policy. Ninety-five percent confidence intervals of robust standard errors clustered at the municipality level are reported.

Figure V: Price and quantity changes by sextile

Panel A: Price and quantity change around the policy

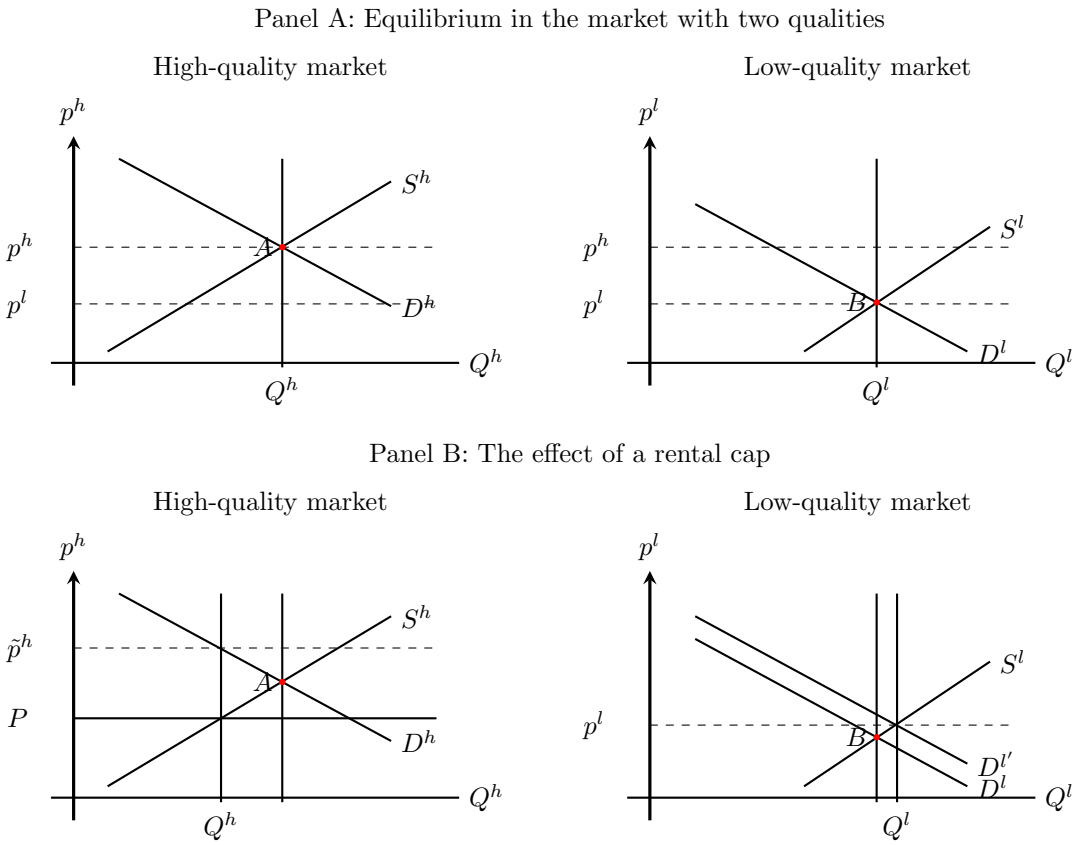


Panel B: Placebo assuming policy change took place in the preceding year



Notes: The figure shows the change in prices and quantity of monthly new contracts between the pre- and the post-period for the entire sample of units along the excess price distribution. Sextiles define six equally sized groups within treatment and control municipalities. Panel A shows the change before and after the policy change, while Panel B shows a placebo that assumes that the policy was implemented one year before its actual implementation.

Figure VI: Perfect competition model with hetergeneous quality



Notes: Panel A of this figure shows the equilibrium in the rental market without two types of units, high- and low-quality. Panel B shows the equilibrium when a price cap at P is introduced that is only binding for the high-quality market.

8 Tables

Table I: Descriptive statistics

	Regulated municipalities			**Non-regulated municipalities**		
	Mean	Std. Dev.	Obs.	Mean	Std. Dev.	Obs.
Rent (€/month)	745.44	355.92	474,477	525.41	236.10	118,758
Square meters	64.55	32.14	474,477	71.82	31.41	118,758
Price per sq.meter	12.20	5.14	474,477	7.75	3.26	118,758

Notes: This table shows summary statistics for treated and non treated municipalities. Sources: INCASOL and AHC.

Table II: Number of rental properties by owner in Barcelona

Number of houses	**Proportion of all houses**	**Number of houses**	**Proportion of all houses**
1 house	33.4 %	25 - 49 houses	6.6 %
2 houses	11.7 %	50 - 99 houses	4.5 %
3 - 5 houses	12.3 %	100 - 149 houses	1.4 %
6 - 10 houses	9.3 %	150 - 199 houses	0.7 %
11 - 15 houses	5.5 %	200 - 299 houses	0.5 %
16 - 24 houses	6.2 %	300 or more	7.0 %

Notes: This table shows the proportion of the number of rental units by owner in the city of Barcelona. Source: Barcelona Metropolitan Housing Observatory.

Table III: Effect of rent cap on rental prices

Panel A: Excess price						
	(1)	(2)	(3)	(4)	(5)	(6)
Effect of Rent Cap	-0.076	-0.068	-0.082	-0.042	-0.067	-0.067
	(0.016)	(0.012)	(0.015)	(0.010)	(0.012)	(0.012)
Observations	523844	523844	257461	112763	50857	523844
Panel B: (ln) Price						
	(1)	(2)	(3)	(4)	(5)	(6)
Effect of Rent Cap	-0.058	-0.048	-0.080	-0.037	-0.062	-0.048
	(0.014)	(0.010)	(0.014)	(0.011)	(0.012)	(0.010)
Observations	591568	591568	286899	164337	72193	591568
Panel C: Fraction above reference price						
	(1)	(2)	(3)	(4)	(5)	(6)
Effect of Rent Cap	-0.146	-0.135	-0.186	-0.160	-0.216	-0.135
	(0.012)	(0.011)	(0.015)	(0.025)	(0.033)	(0.017)
Observations	106149	106149	26957	59264	8911	4138
Specification:						
Fixed effects	yes	yes	yes	yes	yes	yes
Province trends	no	yes	yes	no	no	yes
Unit FE	no	no	yes	no	yes	yes
Condition	no	no	no	<40k	<40k	no
Quarterly freq	no	no	no	no	no	yes

Notes: This table shows the difference-in-differences estimates of the effect of the policy on Excess price – defined as the percentage difference between the contract price and the reference price –, the rental price, and the fraction of contracts above the reference price. Column (1) is our baseline specification. Column (2) adds province × month fixed effects. Column (3) adds unit fixed effects. Column (4) limits our baseline specification to municipalities with fewer than 40,000 inhabitants. Column (5) adds unit fixed effects and limits the sample to municipalities with fewer than 40,000 inhabitants. Column (6) uses the same specification as (3), but considers the data at quarterly frequency. In Panel C, we compute the fraction of contracts above the reference price in each municipality conditioning in columns (3) and (5), on the sample of units with multiple contracts and in columns (4) and (5) on the sample of municipalities with fewer than 40,000 inhabitants. Robust standard errors clustered at the municipality level are reported.

Table IV: Effect of rent cap on number of new contracts signed

Panel A: (ln) Number of new contracts, OLS estimates

	All				Above				Below			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Effect of Rent Cap	-0.207	-0.181	-0.196	-0.093	-0.577	-0.545	-0.546	-0.449	0.166	0.146	0.171	0.233
	(0.021)	(0.023)	(0.030)	(0.025)	(0.042)	(0.042)	(0.052)	(0.048)	(0.038)	(0.038)	(0.047)	(0.049)
Observations	26176	26176	24273	26176	26176	26176	24273	26176	26176	26176	24273	11617

Panel B: (ln) Number of new contracts, PPML estimates

	All				Above				Below			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Effect of Rent Cap	-0.111	-0.106	-0.069	-0.086	-0.401	-0.343	-0.286	-0.359	0.238	0.213	0.259	0.210
	(0.018)	(0.021)	(0.031)	(0.021)	(0.044)	(0.035)	(0.048)	(0.036)	(0.044)	(0.041)	(0.053)	(0.043)
Observations	56007	56007	53545	19558	56007	56007	53545	18669	56007	56007	53545	18669
Specification:												
Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Province trends	no	yes	no	yes	no	yes	no	yes	no	yes	no	yes
Condition	no	no	<40k	no	no	no	<40k	no	no	no	<40k	no
Quarterly freq	no	no	no	yes	no	no	no	yes	no	no	no	yes

Notes: This table estimates the effect of the policy change on the volume of new contracts at the monthly-municipality level. Columns (1), (5), and (9) are our baseline specifications. Columns (2), (6), and (10) add province × month fixed effects. Columns (3), (7), and (11) limit our baseline specifications to municipalities with fewer than 40,000 people. Columns (4), (8), and (12) consider the data at quarterly frequency. All the regressions exclude the lockdown period. Fixed effects include municipality, time fixed effects, and municipality specific linear time trends. Robust standard errors clustered at the municipality level are reported.

Table V: Estimates of the elasticity of new rentals in the market

Panel A: (ln) Number of new contracts, OLS estimates

	Full Sample				No lockdown			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(ln) Rental price	0.068	0.065	0.068	0.064	0.068	0.064	0.067	0.066
	(0.012)	(0.013)	(0.012)	(0.021)	(0.013)	(0.013)	(0.013)	(0.022)
Observations	27302	27302	25303	12168	26176	26176	24273	11617

Panel B: (ln) Number of new contracts, IV estimates

	Full Sample				No lockdown			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(ln) Rental price	1.987	2.046	2.139	1.605	3.069	3.195	3.110	2.206
	(0.395)	(0.517)	(0.599)	(0.512)	(0.459)	(0.613)	(0.646)	(0.571)
Observations	27302	27302	25303	12168	26176	26176	24273	11617
FS F-stat	74.572	45.646	43.819	33.478	73.136	44.729	43.543	34.646
Specification:								
Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes
Province trends	no	yes	no	yes	no	yes	no	yes
Condition	no	no	<40k	no	no	no	<40k	no
Quarterly freq	no	no	no	yes	no	no	no	yes

Notes: This table uses the policy change to estimate the elasticity of new contracts to price changes. Columns (1) and (5) are our baseline specifications. Columns (2) and (6) add province × month fixed effects. Columns (3) and (7) limit our baseline specifications to municipalities with fewer than 40,000 people. Columns (4) and (8) consider the data at quarterly frequency. Fixed effects include municipality, time fixed effects, and municipality specific linear time trends. Panel A presents OLS estimates, while panel B presents IV estimates based on the policy change. Robust standard errors clustered at the municipality level are reported.

APPENDIX FOR ONLINE PUBLICATION

Table of contents of the paper and the appendix

1 Introduction **1**

2 Rent control in Catalonia **8**

3 Data **10**

4 Empirical Evidence **12**

4.1 The effect of the policy on rental prices . . . 12

4.1.1 Nonparametric evidence . . . 12

4.1.2 Two-way fixed effects estimates . . . 13

4.2 The effect of rent caps on the supply of rental units . . . 16

4.2.1 Overall supply of rental units . . . 16

4.2.2 The effect along the distribution of excess price . . . 18

4.2.3 Spillovers . . . 19

5 Model **20**

5.1 Demand for rental units . . . 21

5.2 Supply of rental units . . . 22

5.3 Equilibrium . . . 22

5.4 A rental price cap . . . 23

5.5 Graphical representation . . . 25

5.6 Welfare analysis . . . 26

5.7 Discussion of alternative models . . . 27

6 Conclusions **29**

7 Figures **33**

8 Tables **39**

A Empirical Appendix **45**

A.1 Comparison of different price measures . . . 45

A.2 Quarterly frequency . . . 45

A.3 Heterogeneity-robust estimates . . . 47
A.4 Non-compliance . . . 48
A.5 Heterogeneity . . . 51
A.6 Alternative measure of housing supply . . . 53
A.7 Propensity score matching . . . 54
A.8 Units with one versus multiple contracts . . . 55
A.9 Contract duration . . . 56
A.10 Price and quantity changes along the excess price distribution, robustness . . . 58
A.11 Spillovers figures . . . 59

B Perfect competition with homogeneous housing **61**
B.1 Graphical representation . . . 61

C Market power **61**
C.1 Demand . . . 61
C.2 Supply . . . 61
C.3 Equilibrium . . . 62
C.4 Rental Cap . . . 62

D Search Frictions **63**
D.1 Model setup with homogeneous housing . . . 63
D.2 Multiple submarkets . . . 63
D.3 Free entry . . . 65
D.4 Welfare implications of a reference price rule . . . 66
D.5 Quality differences . . . 68
D.6 Graphical representation . . . 69
D.7 Formal proof of Proposition 1 . . . 72

A Empirical Appendix

A.1 Comparison of different price measures

Figure A.1: Comparison between online posted prices and market prices

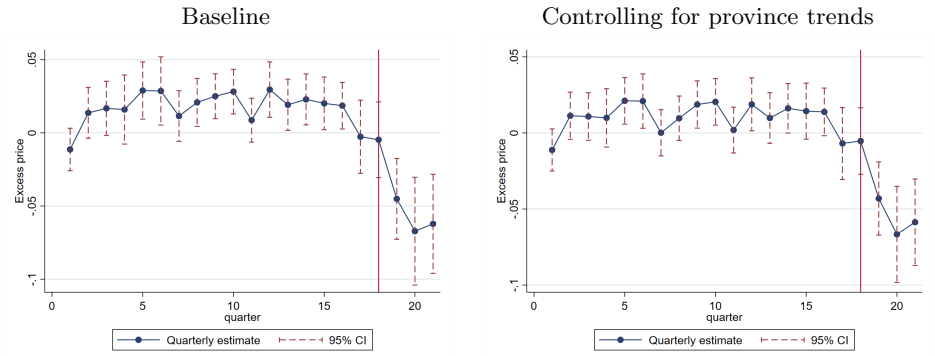


Notes: The top graph shows the level of rental market prices and online posted prices. The bottom graph shows the growth rate of rental market prices and online posted prices. Sources: INCASOL and Idealista.

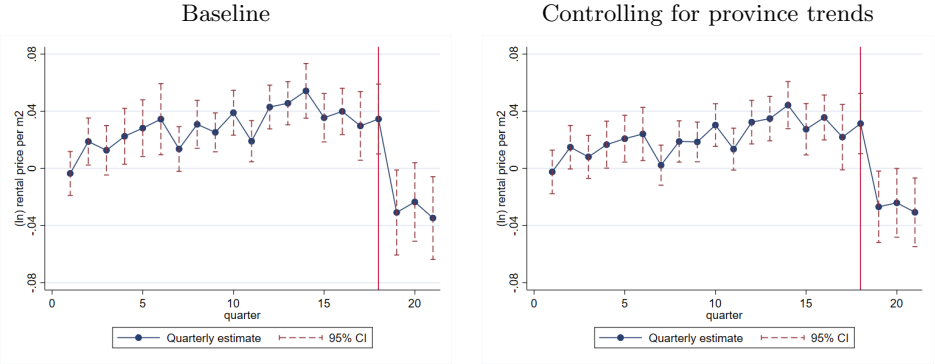
A.2 Quarterly frequency

Figure A.2: Dynamic estimates of the effect of the policy - quarterly frequency

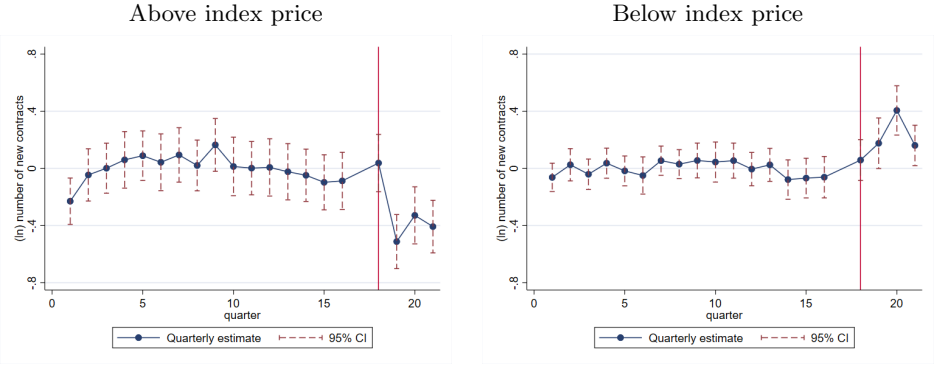
Panel A: Outcome variable: Excess price



Panel B: Outcome variable: (ln) rental price per m2



Panel C: Volume of transactions



Notes: This figure repeats some of the main graphs shown in Figures III and IV using data grouped at the quarterly frequency. Details on each of these graphs are provided in the notes to Figures III and IV.

A.3 Heterogeneity-robust estimates

Table A.1: Heterogeneity-robust estimates of first stage

	Excess price				(ln) rental price				Fraction above excess price			
	β	s.e.	β	s.e.	β	s.e.	β	s.e.	β	s.e.	β	s.e.
t-5			-.0093	.024			.011	.034			-.064	.053
t-4			-.0075	.015			-.012	.021			.0052	.034
t-3			.0082	.01			.016	.012			.0092	.031
t-2			.0097	.011			.0091	.013			.076	.031
t-1			-.022	.013			-.03	.016			-.11	.047
Main estimate	-.051	.019	0		-.043	.028	0		-.13	.043	0	
t+1			-.038	.013			-.032	.021			-.12	.04
t+2			-.022	.011			-.019	.02			-.07	.037
t+3			-.045	.013			-.058	.022			-.16	.045
t+4			-.04	.012			-.0076	.019			-.073	.04
t+5			-.059	.013			-.039	.019			-.1	.033
t+6			-.062	.015			-.047	.017			-.13	.04

Notes: This table reports the same results reported in Table III but using the estimator proposed in de Chaisemartin & D'Haultfoeuille (2020).

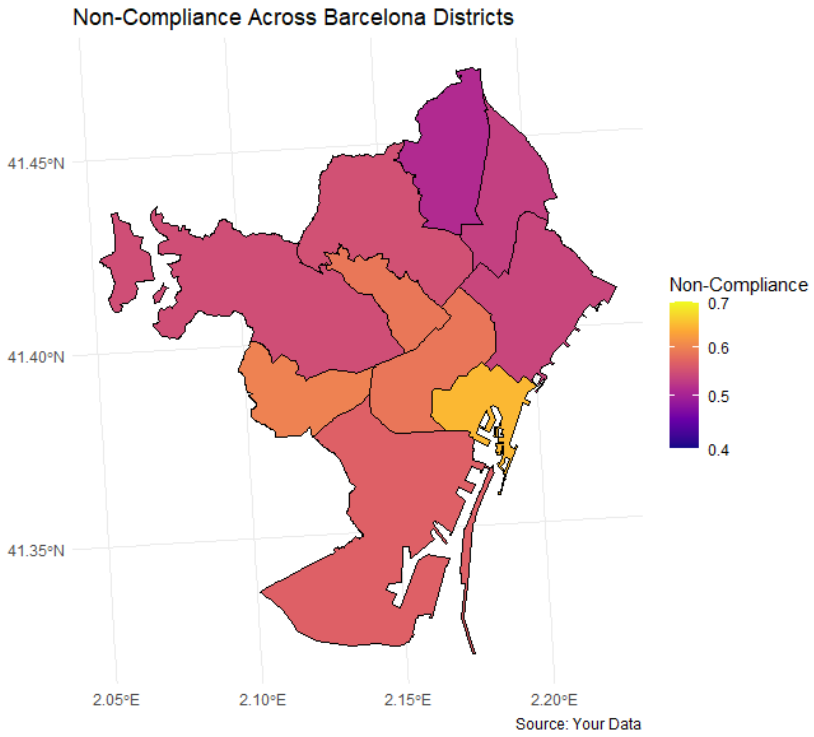
A.4 Non-compliance

Table A.2: Non-compliance, unit level, within municipality variation

	Non-compliance prob					
	(1)	(2)	(3)	(4)	(5)	(6)
Dummy units between 50 to 100m2	0.111	0.112	0.142	0.117	0.128	0.133
	(0.018)	(0.016)	(0.015)	(0.019)	(0.020)	(0.024)
Dummy units between 100 to 150m2	-0.158	-0.160	-0.116	-0.105	-0.089	-0.151
	(0.047)	(0.049)	(0.054)	(0.044)	(0.046)	(0.080)
(ln) dist to non-treated municipality		0.030	0.016		0.011	0.027
		(0.019)	(0.010)		(0.009)	(0.013)
(ln) net income pc, tract			-0.075		0.031	0.147
			(0.075)		(0.087)	(0.057)
(ln) population, tract			0.026		0.089	0.144
			(0.019)		(0.025)	(0.059)
(ln) share over 65, tract			0.044		-0.166	-0.206
			(0.046)		(0.044)	(0.063)
(ln) avg household size, tract			-0.750		-0.682	-0.684
			(0.094)		(0.120)	(0.187)
(ln) share native, tract			-0.427		-0.055	-0.166
			(0.055)		(0.111)	(0.103)
Dummy built 1900-36				-0.008	0.031	0.035
				(0.024)	(0.018)	(0.043)
Dummy built 1936-75				-0.128	-0.044	-0.029
				(0.029)	(0.021)	(0.041)
Dummy built 1975-2000				-0.350	-0.273	-0.208
				(0.025)	(0.028)	(0.048)
Dummy built post 2000				-0.735	-0.683	-0.822
				(0.057)	(0.042)	(0.055)
Observations	50087	50087	49961	40356	40236	13626

Notes: The dependent variable is the probability of the unit having a contracted rental prices above the reference price. The independent variables are unit-specific level characteristics from our baseline saple: surface in squared meters and distance to the closest treated municipality; district level characteristics (column 3); and unit level characteristics from the catastro (column 4), which include year of construction category fixed effects. Column 5 combines information on district level characteristics and unit characteristics. Column 6 restricts the regression to Barcelona. All regressions include month and municipality fixed effects. Only units in treated municipalities in the post-treatment period are included. Columns 1 to 5 report standard errors clustered at the municipality level. Column 5 reports robust standard errors.

Figure A.3: Share of units that did not comply with the policy within Barcelona



Notes: This maps shows the share of units that charged rental prices above the reference price in each district in Barcelona after the policy change.

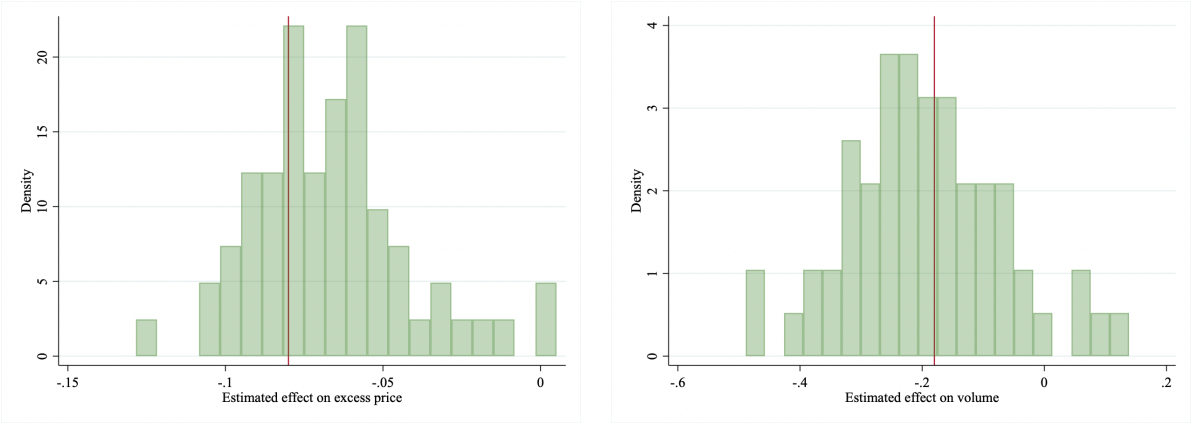
Table A.3: Non-compliance, municipal level

	Non-compliance share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(ln) income pc	-0.089	-0.046	-0.041	-0.078	-0.061	-0.026	-0.034	-0.031
	(0.043)	(0.042)	(0.044)	(0.166)	(0.048)	(0.049)	(0.054)	(0.110)
(ln) density	-0.016	0.000	-0.022	-0.039	-0.002	0.007	-0.002	-0.001
	(0.009)	(0.009)	(0.015)	(0.031)	(0.009)	(0.009)	(0.014)	(0.018)
Percent under 18	0.018	0.015	0.034	0.025	0.021	0.019	0.019	0.013
	(0.007)	(0.007)	(0.010)	(0.024)	(0.008)	(0.008)	(0.011)	(0.014)
Percent over 65	0.001	-0.001	0.017	0.034	0.005	0.005	0.008	0.008
	(0.006)	(0.005)	(0.007)	(0.024)	(0.006)	(0.005)	(0.007)	(0.010)
Percent single	0.004	0.002	0.004	-0.003	0.006	0.004	0.005	0.005
	(0.003)	(0.003)	(0.004)	(0.011)	(0.003)	(0.003)	(0.004)	(0.005)
Percent native	0.002	0.002	-0.003	-0.009	0.002	0.002	0.001	-0.002
	(0.002)	(0.001)	(0.002)	(0.007)	(0.002)	(0.002)	(0.002)	(0.003)
(ln) population	-0.016	-0.021	-0.023	-0.082	-0.030	-0.034	-0.036	-0.065
	(0.007)	(0.007)	(0.010)	(0.104)	(0.008)	(0.008)	(0.011)	(0.029)
Observations	61	61	61	31	207	207	207	177
Fixed_Effects	no	Provincia	Comarca	Comarca	no	Provincia	Comarca	Comarca
Treatment_only	yes	yes	yes	yes	no	no	no	no
Sample	all	all	all	<40k	all	all	all	<40k

Notes: The dependent variable is the share of units with contracted rental prices above the reference price. The independent variables are municipal level characteristics.

A.5 Heterogeneity

Figure A.4: Histograms of municipal level estimates of the effect of rental prices on the excess price



Notes: This figure shows the histograms of the estimate of the policy change on excess price and volume for each municipality in the treatment group relative to non-treated municipalities. The vertical red line indicates the average estimate.

Table A.4: Effect of rent cap on rental prices

Panel A: Excess price

	Excess price				
	(1)	(2)	(3)	(4)	(5)
Effect of Rent Cap	-0.050	-0.051	-0.051	-0.064	-0.050
	(0.007)	(0.006)	(0.006)	(0.005)	(0.006)
× (ln) density	-0.016				0.001
	(0.007)				(0.005)
× (ln) low-skill share		0.014			0.134
		(0.002)			(0.051)
× (ln) immigrant share			0.015		-0.131
			(0.002)		(0.051)
× (ln) income				-0.225	-0.075
				(0.017)	(0.078)

Panel B: (ln) Price

	Price				
	(1)	(2)	(3)	(4)	(5)
Effect of Rent Cap	-0.044	-0.047	-0.047	-0.045	-0.046
	(0.008)	(0.007)	(0.007)	(0.004)	(0.007)
× (ln) density	-0.017				-0.000
	(0.007)				(0.006)
× (ln) low-skill share		0.013			0.145
		(0.002)			(0.064)
× (ln) immigrant share			0.014		-0.159
			(0.002)		(0.068)
× (ln) income				-0.180	-0.246
				(0.018)	(0.110)

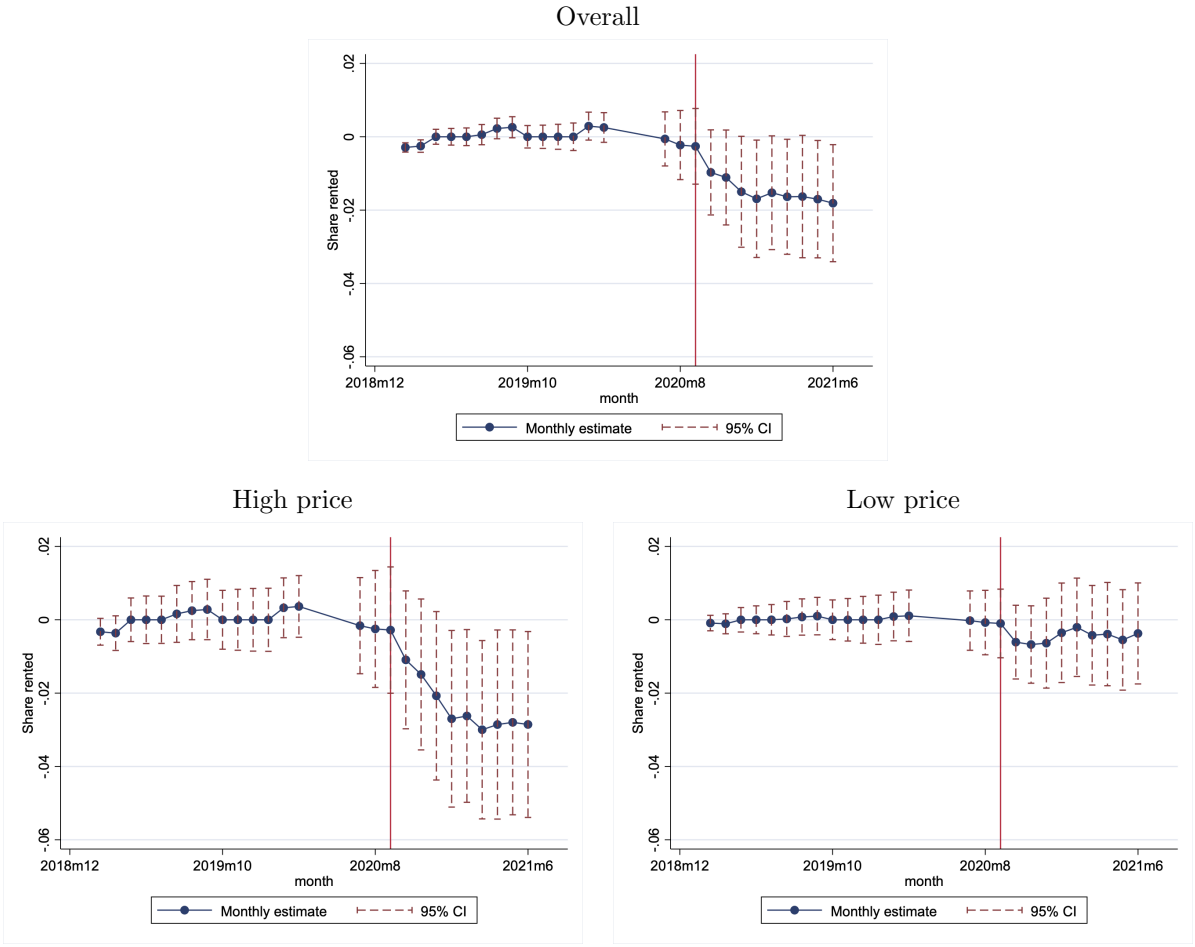
Panel C: Volume

	Volume				
	(1)	(2)	(3)	(4)	(5)
Effect of Rent Cap	-0.273	-0.220	-0.219	-0.252	-0.251
	(0.066)	(0.027)	(0.028)	(0.028)	(0.077)
× (ln) density	0.022				0.001
	(0.015)				(0.020)
× (ln) low-skill share		-0.017			-0.295
		(0.006)			(0.218)
× (ln) immigrant share			-0.018		0.323
			(0.007)		(0.228)
× (ln) income				0.238	0.375
				(0.066)	(0.385)

Notes: This table shows the heterogeneity of our estimates by municipal characteristics.

A.6 Alternative measure of housing supply

Figure A.5: Alternative measure of housing supply



Notes: These graphs show the fraction of units that are rented. The top graph shows the overall fraction. The bottom graphs show the share among high and low prices respectively. We show a smaller window because of the left censoring in not having units whose contracts started prior to when our data is available.

A.7 Propensity score matching

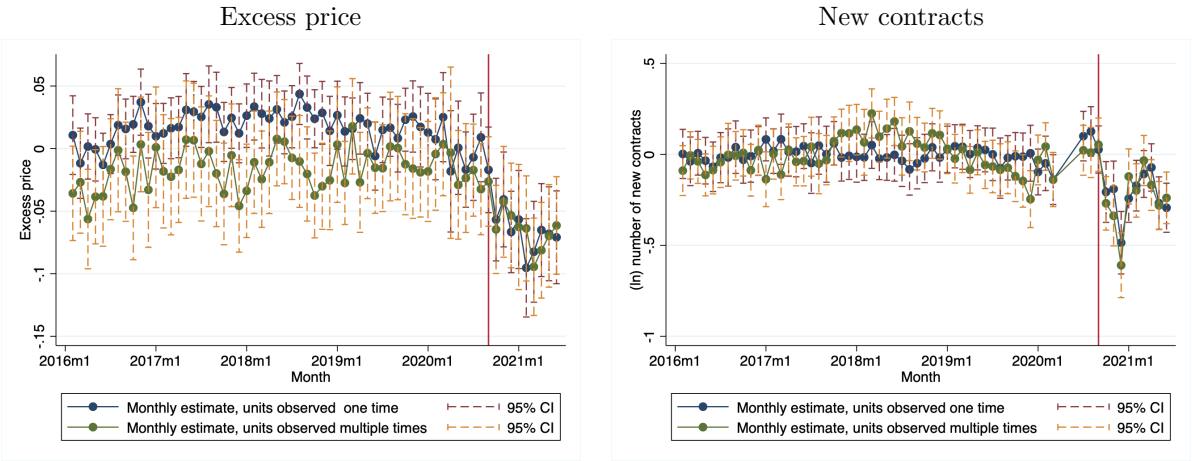
Table A.5: Estimates of the effect of the policy using a propensity score matching estimator

	(1)	(2)	(3)	(4)
	Excess price	(ln) price	Prob. rented	
Effect of the Rent Cap	-0.091	-0.087	-0.020	-0.017
	(0.003)	(0.002)	(0.002)	(0.001)
Matched on	surface	surface	surface	excess price

Notes: This tables shows average treatment effect propensity score estimates of the effect of being on a treated municipality on the change in Excess prices, (ln) prices, and the probability of being rented before and after the implementation of the policy. Columns 1, 2, and 3 match housing units based on square meters. Column 4 matches units based on the excess price distribution. Standard errors in parenthesis.

A.8 Units with one versus multiple contracts

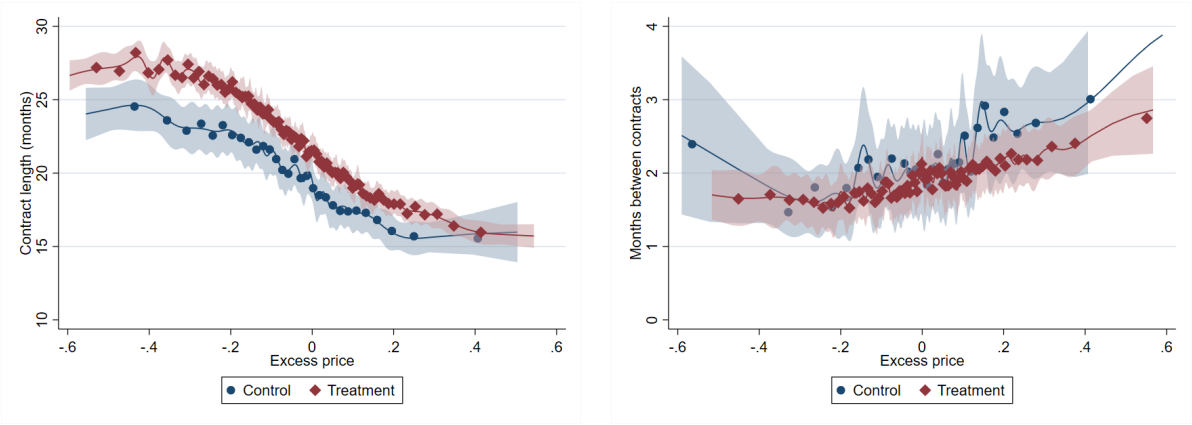
Figure A.6: Price and quantity results by type of unit



Notes: These graphs show the main results on excess prices and new contracts signed separating by units for which we observe multiple contracts in our data and units for which we only observe one contract.

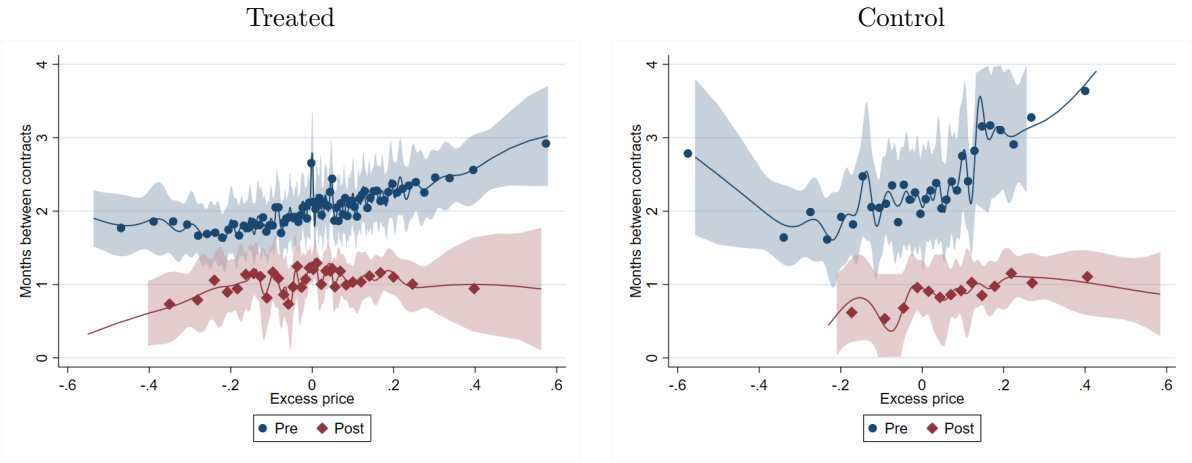
A.9 Contract duration

Figure A.7: Relationship between contract duration and excess price



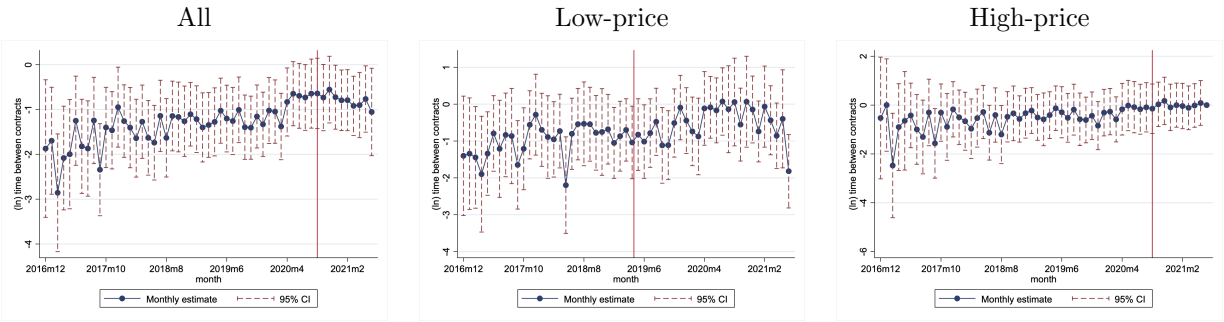
Notes: This figure shows two different binned scatter plots. The data points are grouped into equal-sized bins, and the selected number of bins is optimal in minimizing the (asymptotic) integrated mean-squared error (IMSE). Each data point corresponds to the mean of the two variables within each bin, controlling for municipality. We add a fitted line and a confidence band based on cubic B-splines.

Figure A.8: Relationship between time in-between contracts and excess price



Notes: This figure shows binned scatter plots for time in-between contracts and excess price. The data points are grouped into equal-sized bins, and the selected number of bins is optimal in minimizing the (asymptotic) integrated mean squared error (IMSE). Each datapoint corresponds to the mean of the two variables within each bin, controlling for municipality. We add a fitted line and a confidence band based on cubic B-splines.

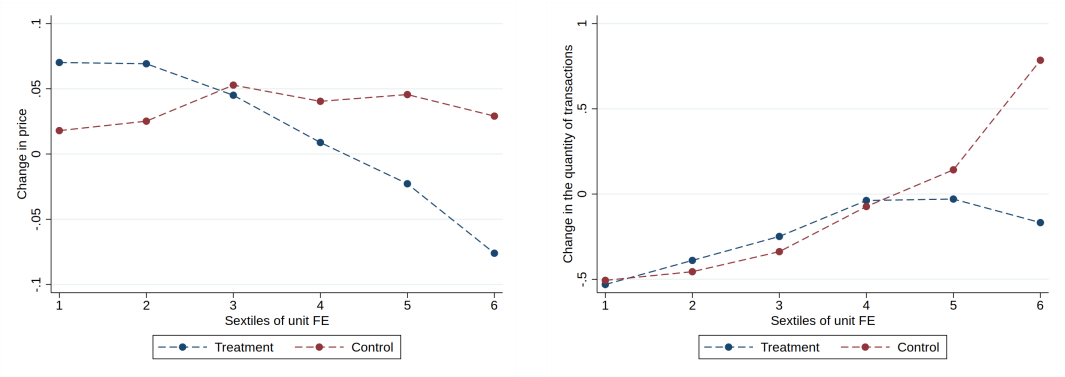
Figure A.9: Relationship between time in-between contracts difference-in-difference



Notes: This figure shows the difference in difference estimates on contract duration comparing treatment and control units. The first graphs shows the estimates for all the units in the market, the second graph for all units with prices below the reference price, and the third graph shows the estimates for units with prices above the reference price.

A.10 Price and quantity changes along the excess price distribution, robustness

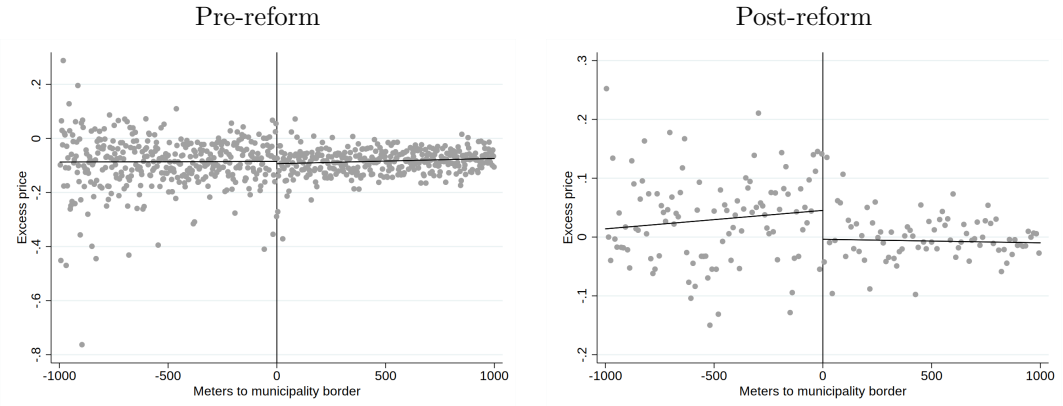
Figure A.10: Price and quantity changes by sextile



Notes: This figure shows the same results as Figure V but where sextiles are computed directly using the excess price, i.e. without removing municipality and month fixed effects.

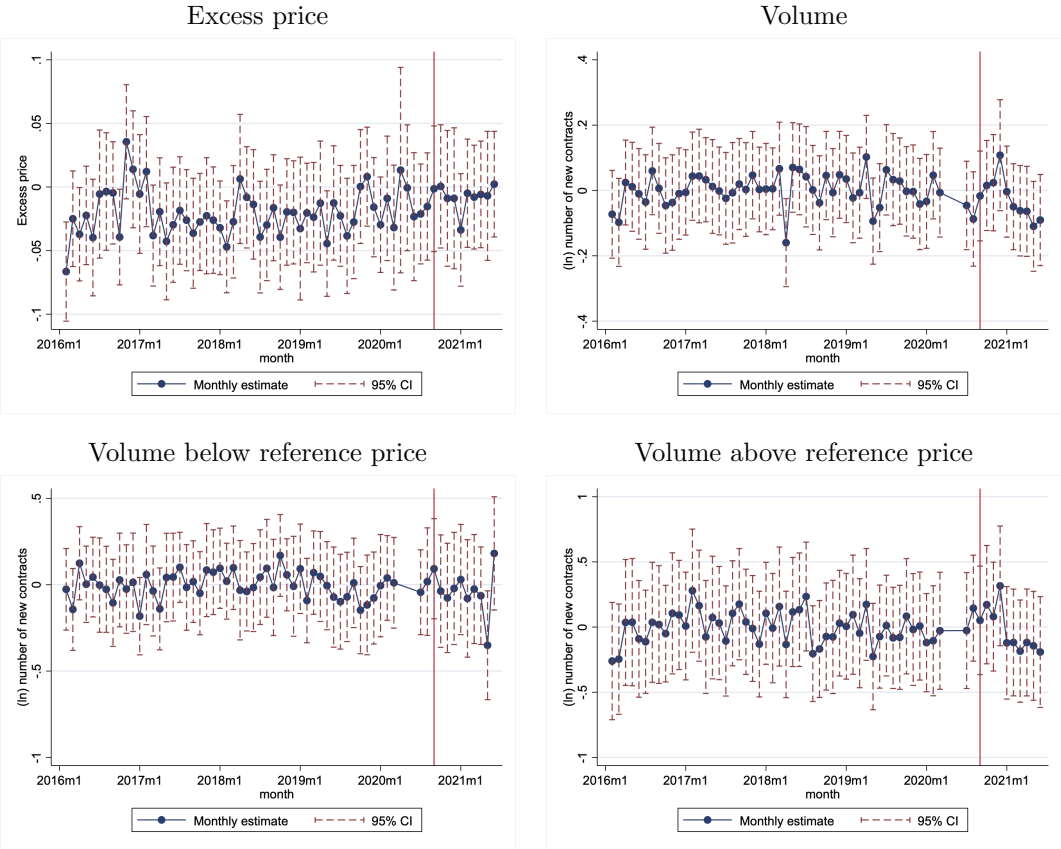
A.11 Spillovers figures

Figure A.11: Border discontinuity



Notes: The graph on the left estimates a discontinuity along the border between non-treated (left) and treated (right) in terms of excess prices before the policy. The graph on the right repeats the exercise after the policy. Controls include province fixed effects.

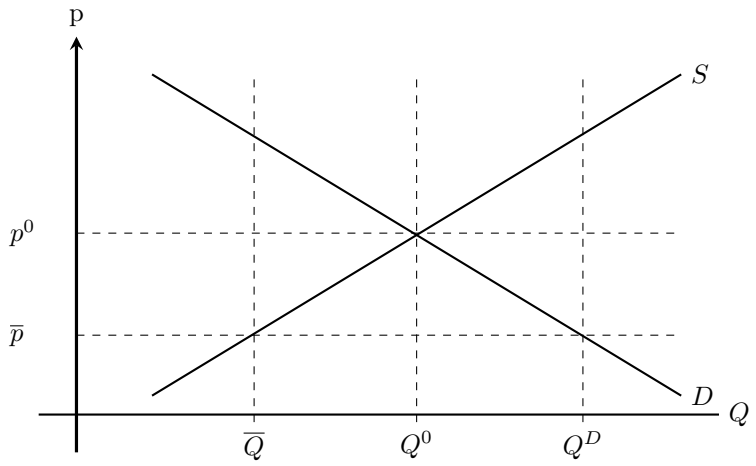
Figure A.12: Spillovers



Notes: The top left graph estimates the effect of the policy change on the excess price on units in non treated municipalities within 1,000 meters of treated ones relative to units in non treated municipalities. The top right graph estimates the effect of the policy on volumes in municipalities with units on average within 1,000 meters of treated municipalities, relative to non treated municipalities. The bottom graphs repeat the exercise on volumes for units with prices above and below the reference price.

B Perfect competition with homogeneous housing

B.1 Graphical representation



C Market power

Often, justifications for rental cap policies hinge on the fact that landlords have some degree of market power. In this appendix, we explore the welfare consequences of a rental cap policy in a model with heterogeneous quality where landlords internalize the demand for their units. We directly assume different types of units based on quality because even with market power, a homogeneous housing model would not generate price dispersion in equilibrium, and this is a prominent feature of the data.

C.1 Demand

The demand for rental units is the same as in the main text. Hence, we have the same expressions as the ones shown in Section 5.1.

C.2 Supply

There are different ways to think about landlord with market power. It is standard to assume that in these cases, landlords will internalize the demand for their units which will lead to prices being a markup over marginal costs, which, in turn, means that charged prices will be a markup over the equilibrium price that would prevail in the absence of market power. Similarly, quantities would be lower than the quantities prevailing in

equilibrium without market power. Instead of deriving explicitly these markup and markdowns, we will assume them directly into the equilibrium.

C.3 Equilibrium

The equilibrium with market power will be similar to the equilibrium without market power described in the main text, but with higher prices and lower quantities. We can express it as:

$$p^l = \mu_p^l (\frac{u^l}{U})^{\gamma/\rho} N^\gamma C$$

$$p^h = \mu_p^h (\frac{u^h}{U})^{\gamma/\rho} N^\gamma C$$

where μ_p^l and μ_p^h are markups larger than 1 and all the other variables are as before. Equilibrium quantities, which are given by:

$$Q^l = \frac{1}{\mu_q^l} (\frac{u^l}{U})^{1/\rho} N$$

$$Q^h = \frac{1}{\mu_q^h} (\frac{u^h}{U})^{1/\rho} N$$

where μ_p^l and μ_p^h are wedges with respect to the perfect competition equilibrium, i.e. larger than 1.

C.4 Rental Cap

Suppose now that a rental price cap is introduced as before. Let's assume that it is only binding in the high-quality submarket. Then, in that case, the rent cap will eliminate the ability of high quality owners to internalize the demand. Hence, the equilibrium will shift to perfect competition allocation in this submarket. This will also lower prices in the high-quality submarket, driving some renter in the low-quality submarket towards the high-quality one. This shift in demand will also lead to lower prices in the lower-quality submarket (at least to the extent that prices are partially tied to demand in the market, something that it is not always true in price setting models).

It is worth noting that this produces predictions that are not in line with the data. In the data, we see that prices of low-price units increase, whereas models with market power would predict the opposite. Similarly, the model would predict expansions in the number of high-quality units in the market, again, contrary to the empirical evidence.

D Search Frictions

Search frictions have the potential to explain price dispersion in equilibrium, as shown by Galenianos & Kircher (2009) and Wright et al. (2021) using a directed search model of the labor market. In this section we show how the model can be adapted to the rental market. We start the analysis by considering homogenous housing. This helps to build the intuition on how a rental cap policy affect allocations and prices in this model.

We show that in this model, a rental price cap that acts as a reference price leads to price convergence. This is in line with the data, except for the fact that even after the policy there is price dispersion below the reference price even in excess prices. In the later part of this section we consider heterogeneous housing in this directed search environment.

D.1 Model setup with homogeneous housing

There are two types of agents in the model: unit owners or landlords and tenants. Following the terminology in the directed search literature, we denote by "sellers" the owners of rental units, and by "buyers" the potential tenants. Sellers post (and commit to) a single price. Posted prices are observed by buyers, who can apply to many rental units. There is a mass N_b of buyers and a mass N_s of sellers. We denote by N the ratio of buyers to sellers, so $N = N_b/N_s$. All buyers and sellers are ex ante homogeneous.

We denote by $\alpha(n)$ the probability that a mass n_s of sellers meets a mass n_b of buyers for a posted rental price p, with $n = n_b/n_s$. Note that this may or may not coincide with N, depending on whether multiple prices are posted in equilibrium. This ratio is also referred to in the literature as market tightness. We denote the value that renters derive from renting a unit by u and the cost that sellers incur when having a unit in the rental market by c. Hence, the (ex post) payoff for buyers to rent at price p is given by $u - p$, while the payoff for sellers is $p - c$.

Matches are determined by a constant return to scale (CRS) matching function $\mu(n_b, n_s)$. We can define the probability that a seller meets a buyer, $\alpha(n)$, using the matching function: $\alpha(n) = \frac{\mu(n_b,n_s)}{n_s} = \mu(n, 1)$. We can also define $\frac{\alpha(n)}{n} = \frac{\mu(n_b,n_s)}{n_b} = \frac{\mu(n,1)}{n}$, which is the probability that a buyer meets a seller. We denote by $\varepsilon(n) = \frac{n\alpha'(n)}{\alpha(n)}$ the elasticity of the meeting probability with respect to the ratio of buyers and sellers.

D.2 Multiple submarkets

We assume that buyers apply to two different types of units.[40] Sellers, instead, can only post one price. We denote by $\underline{p}$ and $\overline{p}$ the prices associated with each of the different submarkets and by $\underline{n}$ and $\overline{n}$ the ratio of buyers

40. More generally, we could consider buyers applying to an arbitrary number of different types of units, but using two types simplifies the model, the exposition, and the intuitions that we describe in what follows.

to sellers in each submarket.[41] Hence, the number of buyers in each submarket is equal to the total number of buyers, i.e. $\underline{n}_b = \bar{n}_b = N_b$, whereas the number of sellers in each submarket is determined in equilibrium.

With multiple submarkets, the value of the buyers can be expressed as:

$$V_b = \frac{\alpha(\underline{n})}{\underline{n}}(u - \underline{p}) + (1 - \frac{\alpha(\underline{n})}{\underline{n}})\frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p}) \tag{4}$$

This expression says that with probability $\frac{\alpha(\underline{n})}{\underline{n}}$, the buyer meets a low-price seller, in which case she accepts the price and derives utility $(u - \underline{p})$. With probability $(1 - \frac{\alpha(\underline{n})}{\underline{n}})$, she does not meet a low-price unit, and hence is forced to consider high-price options. In the high-price market, with probability $\frac{\alpha(\overline{n})}{\overline{n}}$, she meets a high-price seller and derives utility $(u - \overline{p})$. With probability $(1 - \frac{\alpha(\underline{n})}{\underline{n}})(1 - \frac{\alpha(\overline{n})}{\overline{n}})$ she does meets neither a low- nor a high-price seller, and hence obtains her outside option, which we assume to be equal to 0.

With multiple submarkets, sellers need to decide whether to post a high or a low price. We can start the analysis by looking into the optimization problem of sellers considering the high-price submarket. This maximization problem can be written as follows:

$$V_{\overline{s}} = \max_{\overline{n},\overline{p}}\{(\alpha(\overline{n})(1 - \frac{\alpha(\underline{n})}{\underline{n}})(\overline{p} - c))\} \text{ subject to } \frac{\alpha(\underline{n})}{\underline{n}}(u - \underline{p}) + (1 - \frac{\alpha(\underline{n})}{\underline{n}})\frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p}) = V_b$$

This equation says that sellers in the high-price market meet a buyer who did not find a low-price unit with probability $\alpha(\overline{n})(1 - \frac{\alpha(\underline{n})}{\underline{n}})$, in which case sellers obtain a payoff equal to $(\overline{p} - c)$. We also assume that sellers can internalize the demand they face given the price they decide to post.

We can solve this maximization problem by substituting the constraint into the objective function and taking first-order conditions. Doing so, we obtain that:

$$\overline{p} = (1 - \varepsilon(\overline{n}))u + \varepsilon(\overline{n})c \tag{5}$$

Equation (5) is easy to interpret. It says that sellers who decide to post in the high-price submarket, post a price that is a weighted average between their cost of posting and buyers' utility.

Using the equilibrium price, we obtain that the sellers' value when posting in the high-price submarket is given by:

$$V_{\overline{s}} = \alpha(\overline{n})\left(1 - \frac{\alpha(\underline{n})}{\underline{n}}\right)(1 - \varepsilon(\overline{n}))(u - c) \tag{6}$$

We can now turn to sellers considering the low-price submarket. Sellers in the low-price market also need to decide what price to post given the demand they face. This can be written as:

41. Galenianos & Kircher (2009) show that there are as many submarkets as buyer applications.

$$V_{\underline{s}} = \max_{\underline{n},\underline{p}} \{\alpha(\underline{n})(\underline{p} - c)\} \text{ subject to } \frac{\alpha(\underline{n})}{\underline{n}}(u - \underline{p}) + (1 - \frac{\alpha(\underline{n})}{\underline{n}})\frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p}) = V_b$$

This equation says that sellers in the low-price market meet a buyer with probability $\alpha(\underline{n})$, in which case, given that buyers always prefer low prices, they match, and hence, sellers obtain a payoff equal to $(\underline{p} - c)$.

We can also solve this maximization problem by substituting the constraint into the objective function and taking first-order conditions. When we do so, we obtain:

$$\underline{p} = (1 - \varepsilon(\underline{n}))\underline{u} + \varepsilon(\underline{n})c \tag{7}$$

where $\underline{u} = u - \frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p})$.

Equation (7) is also easy to interpret. It says that sellers who decide to post in the low-price submarket post a price that is a weighted average between their cost of posting and buyers' utility, in this case, relative to the one obtained in the high-price submarket.

As before, we can obtain the payoff of posting in the low-price submarket by introducing the equilibrium price into the objective function:

$$V_{\underline{s}} = \alpha(\underline{n})(1 - \varepsilon(\underline{n}))(\underline{u} - c) \tag{8}$$

It is worth noting that $\underline{p} < \overline{p}$, since we can use the fact that $V_{\underline{s}} = V_{\overline{s}}$ to obtain $\frac{\overline{p}}{\underline{p}} = \frac{\alpha(\underline{n})}{\alpha(\overline{n})(1 - \frac{\alpha(\underline{n})}{\underline{n}})}$. The numerator on the right-hand side, which is the probability a seller meets a buyer in the low-price market, needs to be larger than the probability of finding a buyer in the high-price market conditional on not finding one in the low-price market. Otherwise, there would be more entry in the high-price market, which delivers a higher payoff to sellers.

D.3 Free entry

Entry conditions have two margins. First, landlords who participate in the rental market need to be indifferent between posting in the high- and the low-price submarkets. This leads to the following indifference condition, which results from equating Equations (6) and (8):

$$\alpha(\underline{n})(1 - \varepsilon(\underline{n}))(\underline{u} - c) = \alpha(\overline{n})\left(1 - \frac{\alpha(\underline{n})}{\underline{n}}\right)(1 - \varepsilon(\overline{n}))(u - c) \tag{9}$$

The second margin of the entry decision is whether landlords want to participate in the market. One way to model this is to assume that there is an outside option to owners, for instance to sell the unit they have or

to enter short-term rental contracts, with payoff $k_s(N_s)$. We assume that this payoff is a function that depends on the number of sellers (N_s) in the rental market. This captures the idea that, when the rental market is attractive, even owners for whom it may be very costly to offer their unit in the rental market may consider doing so. To capture this, we assume that $k_s(N_s) > 0, \forall N_s \in [0, \infty)$ and, that $k_s(N_s)$ is (weakly) increasing. In this case, the indifference condition between posting in the high- and in the low-price submarkets needs to be equal to the outside option $k_s(N_s)$:

$$\alpha(\underline{n})(1-\varepsilon(\underline{n}))(\underline{u}-c) = \alpha(\overline{n})\left(1-\frac{\alpha(\underline{n})}{\underline{n}}\right)(1-\varepsilon(\overline{n}))(u-c) = k_s(N_s) \tag{10}$$

Equation (10) defines a system of two equations and two unknowns, $\underline{n}$ and $\overline{n}$, that has a unique solution, as shown in Galenianos & Kircher (2009).

D.4 Welfare implications of a reference price rule

In this subsection, we turn to study the implications of introducing a reference price rule. We can prove the following proposition.

PROPOSITION 1.

1. A policy that forces a unique reference price $P \in [c, u]$ results in the following:
   (a) Prices converge to P.
   (b) The policy increases total welfare.
   (c) The policy is ex ante welfare-improving for buyers and sellers if and only if:
   $$(1-\gamma)c + \gamma u < P < (1-\eta)u + \eta c$$
   for $\gamma > 0$ and $\eta > 0$, defined in the appendix D.7.
   (d) Ex post the policy is welfare improving or welfare declining depending on whether the reference price is above or below existing prices in each submarket.
2. A policy that forces a unique price $P \notin [c, u]$ destroys the rental market.
3. Overall supply in the rental market is affected by the policy as follows:
   (a) If $P < c$ then all landlords exit the rental market.
   (b) If $c < P < (1-\gamma)c + \gamma u$ then some landlords exit the rental market.

(c) If $P > (1-\gamma)c + \gamma u$ then more landlords enter the rental market.

Proof.

See Appendix D.7, for a detailed proof. ■

This proposition has several implications. First, when a reference price is introduced, prices converge to the reference price. This is similar to what we observ in the data, except that price dispersion is still prevalent below reference prices after the policy. Second, the proposition says that there is a range of prices, determined by the utility derived from housing and the cost of having a unit in the market, where the policy is (ex ante') welfare improving. This is so because the policy removes the inefficiency created by the fact that renters who end up in high-price units also queue in low-price ones, making matching more costly. Furthermore, there is a range of prices where the policy is welfare improving for both renter and landlords. Proposition 1 characterizes the (smaller) range of prices where this happens.

The proposition shows that, if the reference price P is lower than $(1-\gamma)c + \gamma u$ but still above c, then landlords lose, while if the reference price is higher than $(1-\eta)u + \eta c$ but still below u, renters lose. Hence, there is a range of prices that does no destroy the market and that favors either landlords or renters.

The reference price set by the policymaker has effects on the overall size of the market. A policy that sets the price P in the range $(c, (1-\gamma)c + \gamma u)$, leads to lower entry in the rental market since it lowers the value of having units in the rental market. In that case, high-price matches are the ones more likely to break. A reference price above $(1-\gamma)c + \gamma u$ encourages entry.

All these results can be summarized in a graph, as shown in Figure A.13. This graph provides an example where $\overline{p}$ is above $(1-\eta)u + \eta c$, and where $\underline{p}$ is below $(1-\gamma)c + \gamma u$ (something that does not necessarily need to be the case, but is the most plausible ordering).[42] In this case, we define six regions where the policymaker could set the price P (without destroying the market). A reference price P set in region A would be good for landlords, who value high prices, but would be detrimental to renters (both ex ante and ex post), even for those initially in the high-price submarket. Region B is similar to region A, except that in this region the price P is lower than the prevailing price in the high-price market, so renters matched in the high-price submarket would have incentives to find a new match. Region C is ex ante welfare improving for both landlords and renters. In this region the post-policy average price would be higher than prior to the policy. Region D is like region C except that the policy leads to a lower average price, something that is in line with the data. In this region, however, we would observe new entrants into the rental market, something that we did not observe in the data. Region E is ex ante welfare improving for renters, but ex ante detrimental to landlords. In this region

42. Figure A.14 in the Appendix explores a reversed ordering. The empirical evidence provided in the paper rules out this situation.

we should observe overall exit in the rental market. In this case, however, renters in the low-price units would have incentives to stay in the match, and landlords in the low-price submarket would have strong incentives to stay in the market. This means that we would see matches in the high-price submarket disappear. Hence, this region is in line with all the empirical results discussed above. Finally, Region F is similar to region E, except that even matched landlords in low-price units would lose and hence would have incentives to leave the market, something that we do not observe.

D.5 Quality differences

To gain intuition, we assumed thus far that housing units are homogeneous. This assumption simplifies the exposition of the model and is adequate if reference prices correctly account for any unobserved heterogeneity. In practice, however, there may be quality differences that are hard to take into account when computing reference prices.

As in the main text, we assume that there are two qualities, high and low. We index quality by $q \in \{High, Low\}$. Using this notation, we denote the payoff of living in a high-quality unit by u^h, while the payoff for a low-quality unit is u^l. We use similar notation for all the other variables in the model. We assume $u^h > u^l$ and $c^h > c^l$.

As in the main text, we assume that the utility of an individual i for looking into the market of quality q is given the utility derived from the rental unit and an idiosyncratic individual-specific shock. In equations:

$$\ln V_b^q + \epsilon_{iq} = \ln[\frac{\alpha(\underline{n}^q)}{\underline{n}^q}(u^q - \underline{p}^q) + (1 - \frac{\alpha(\underline{n}^q)}{\underline{n}^q})\frac{\alpha(\overline{n}^q)}{\overline{n}^q}(u^q - \overline{p}^q)] + \epsilon_{iq} \tag{11}$$

We further assume that ϵ is drawn from an extreme value distribution such that a fraction of the buyers search in the high-quality market and the rest in the low-quality market (imagine there are two online platforms where people apply). Then:

$$N^q = \pi^q N = (\frac{V_b^q}{V_b})^{\frac{1}{\lambda}} \tag{12}$$

where $V_b = [(V_b^l)^{\frac{1}{\lambda}} + (V_b^h)^{\frac{1}{\lambda}} + (B)^{\frac{1}{\lambda}}]^{\lambda}$, in a slight abuse of notation relative to the previous sections.

Finally, as before we assume that sellers either have high- or low-quality units supply to the market depending on a free entry condition as before. We also assume that sellers take the number of buyers in their market N^q as given. Under these assumptions, we have the same equations as in the previous section, but all depending on each quality market. On top, we have equation (12) that governs the overall size of each quality market. In particular, there will be four equilibrium prices in the market:

$$\underline{p}^q = (1 - \varepsilon(\underline{n}^q))\underline{u}^q + \varepsilon(\underline{n}^q)c^q \tag{13}$$

and,

$$\overline{p}^q = (1 - \varepsilon(\overline{n})^q)u^q + \varepsilon(\overline{n}^q)c^q \tag{14}$$

for $q \in \{High, Low\}$.

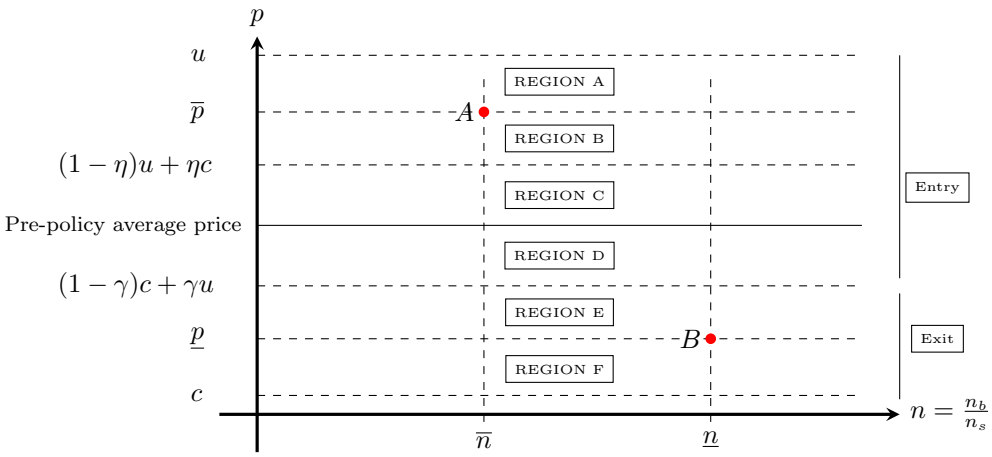
The logic of Proposition 1 applies in this model with quality differences. First, in equilibrium, each quality market can be divided in two submarkets. Depending on the values of u^q and c^q, we may have overlap across markets. That is, it may happen that the high-quality market has uniformly higher prices than the low-quality market, or that the low-price submarket of the high-quality market has in equilibrium lower prices than the high-price submarket of the low-quality market. In equations:

$$\text{Either } \underbrace{\overline{p}^h > \underline{p}^h > \overline{p}^l > \underline{p}^l}_{\text{Case 1}} \text{ or } \underbrace{\overline{p}^h > \overline{p}^l > \underline{p}^h > \underline{p}^l}_{\text{Case 2}}$$

With quality differences, a reference price P that does not distinguish between qualities will have similar effects as the ones described in Proposition 1 *within* each of the submarkets affected by the policy. For instance, if we assume that the policy affects only high quality submarkets, then prices will converge to a reference price among high-quality markets. Moreover, because the policy is welfare enhancing in these submarkets, renters will move towards them. Note that this would generate patterns in the data that are somewhat contrary to what we observe. In particular the policy would shift demand toward high-quality units, and would probably result in more (not less) high quality units in the market.

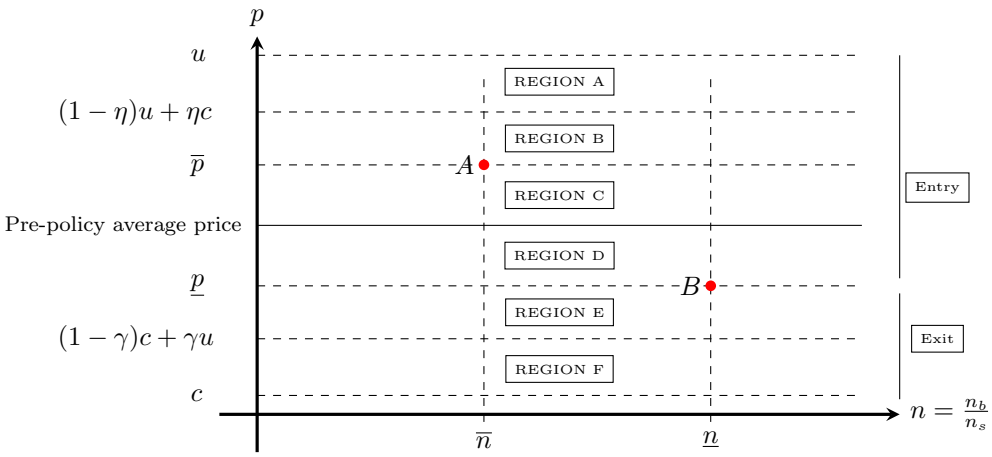
D.6 Graphical representation

Figure A.13: Graphical representation of the model



Notes: This graph shows the original equilibrium with two submarkets, denoted by A and B in a diagram with prices on the y-axis and market tightness in the x-axis. It also depicts the various regions where the policymaker can set the reference price P. Entry and Exit denote whether in those regions the model predicts a higher number of rental contracts or a lower one, after the implementation of a reference price. Region A: Welfare improving for all landlords, detrimental to renters even ex post. Region B: Welfare improving for all landlords, detrimental to renters. Region C: Welfare improving for landlords and renters, higher post-policy average prices. Region D: Welfare improving for landlords and renters, lower post-policy average prices. Region E: Welfare improving for all renters, detrimental to landlords. Region F: Welfare improving for renters, detrimental to landlords even ex post.

Figure A.14: Alternative theory graph



Notes: This graph shows the original equilibrium with two submarkets, denoted by A and B in a diagram with prices on the y-axis and market tightness on the x-axis. It also depicts the various regions where the policymaker can set the reference price P. Entry and Exit denote whether in those regions the model predicts a higher number of rental contracts or a lower one, after the implementation of a reference price. Region A: Welfare improving for all landlords, detrimental to renters even ex post. Region B: Welfare improving for landlords and renters, ex post welfare improving for all landlords. Region C: Welfare improving for landlords and renters, higher post-policy average prices. Region D: Welfare improving for landlords and renters, ex post welfare detrimental for all landlords. Region E: Welfare improving for landlords and renters, even low-price landlords exit. Region F: Welfare improving for all renters, detrimental to landlords.

D.7 Formal proof of Proposition 1

Point a) is by assumption.

To prove point b) we need to show that welfare is higher with one submarket than with two submarkets. For this, with one submarket, when buyers are allowed to submit two applications, we have that total buyer welfare is:

$$N_b V_b = N_b \frac{\mu(N_s, 2N_b)}{N_b}(u - P) = \mu(N_s, 2N_b)(u - P)$$

Total seller welfare is given by

$$N_s V_s = N_s \frac{\mu(N_s, 2N_b)}{N_s}(P - c) = \mu(N_s, 2N_b)(P - c)$$

Hence:

$$W^* = \mu(N_s, 2N_b)(u - P) + \mu(N_s, 2N_b)(P - c) = \mu(N_s, 2N_b)(u - c)$$

In the two-market equilibrium, total buyer welfare is given by

$$\begin{aligned} N_b V_b &= N_b\left(\frac{\mu(\underline{n}_s, N_b)}{N_b}(u - \underline{p}) + \left(1 - \frac{\mu(\underline{n}_s, N_b)}{N_b}\right)\frac{\mu(\bar{n}_s, N_b)}{N_b}(u - \overline{p})\right) \\ &= \mu(\underline{n}_s, N_b)(u - \underline{p}) + \left(1 - \frac{\mu(\underline{n}_s, N_b)}{N_b}\right)\mu(\bar{n}_s, N_b)(u - \overline{p}) \end{aligned}$$

Likewise, total seller welfare is given by

$$N_s V_s = \underline{n}_s \frac{\mu(\underline{n}_s, N_b)}{\underline{n}_s}(\underline{p} - c) + \bar{n}_s \frac{\mu(\bar{n}_s, N_b)}{\bar{n}_s}(\overline{p} - c)$$

Hence:

$$N_s V_s = \mu(\underline{n}_s, N_b)(\underline{p} - c) + \mu(\bar{n}_s, N_b)(\overline{p} - c)$$

Therefore, in the two-market equilibrium, total welfare is given by

$$W = (\mu(\underline{n}_s, N_b) + \mu(\bar{n}_s, N_b))(u - c) - \frac{\mu(\underline{n}_s, N_b)}{N_b}\mu(\bar{n}_s, N_b)(u - \overline{p})$$

We can compare welfare:

$$W^* - W = [\mu(N_s, 2N_b) - (\mu(\underline{n}_s, N_b) + \mu(\bar{n}_s, N_b))](u - c)$$
$$+ \frac{\mu(\underline{n}_s, N_b)}{N_b}\mu(\bar{n}_s, N_b)(u - \overline{p}) =$$
$$= [2\mu(\frac{1}{2}N_s, N_b) - (\mu(\underline{n}_s, N_b) + \mu(\bar{n}_s, N_b))](u - c) +$$
$$+ \frac{\mu(\underline{n}_s, N_b)}{N_b}\mu(\bar{n}_s, N_b)(u - \overline{p})$$

And this is smaller than W^*, since $\mu(\underline{n}_s, N_b) \leq \mu(N_s, N_b)$ and $\mu(\bar{n}_s, N_b) \leq \mu(N_s, N_b)$ and because of the properties of CRS matching functions.

For point c), we have that welfare of buyers is:

$$V_b^* = \frac{\mu(N_s, 2N_b)}{N_b}(u - P)$$

whereas in the one market equilibrium is:

$$V_b = (\frac{\mu(\underline{n}_s, N_b)}{N_b}(u - \underline{p}) + (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{N_b}(u - \overline{p}))$$

Before comparing the two, it is useful to realize that from $\underline{p} = (1 - \varepsilon(\underline{n}))(u - \frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p})) + \varepsilon(\underline{n})c$ and $\overline{p} = (1 - \varepsilon(\overline{n}))u + \varepsilon(\overline{n})c$ we have that:

$$\overline{p} - u = (1 - \varepsilon(\overline{n}))u + \varepsilon(\overline{n})c = \varepsilon(\overline{n})(c - u)$$

and that:[43]

43.

We have that

$$\underline{p} - u = (1 - \varepsilon(\underline{n}))(u - \frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p})) + \varepsilon(\underline{n})c - u$$

hence

$$\underline{p} - u = \varepsilon(\underline{n})(c - u) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p})$$

thus

$$\underline{p} - u = \varepsilon(\underline{n})(c - u) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}\varepsilon(\overline{n})(c - u)$$

and so:

$$\underline{p} - u = (\varepsilon(\underline{n}) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}\varepsilon(\overline{n}))(c - u)$$

hence:

$$\underline{p} - u = (\frac{\varepsilon(\underline{n})}{\varepsilon(\overline{n})} - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}})\varepsilon(\overline{n})(c - u)$$

$$\underline{p} - u = \Omega(\underline{n}, \overline{n})\varepsilon(\overline{n})(c - u)$$

with $\Omega(\underline{n}, \overline{n}) \equiv (\frac{\varepsilon(\underline{n})}{\varepsilon(\overline{n})} - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}) < 1$.

Using this, we can compare buyers' welfare in the one and two submarket equilibria:

$$V_b^* - V_b = \frac{\mu(N_s, 2N_b)}{N_b}(u - P) - \frac{\mu(\underline{n}_s, N_b)}{N_b}(u - \underline{p}) - (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{N_b}(u - \overline{p})$$

Hence $V_b^* - V_b > 0$ if and only if:

$$P < u - \frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}(u - \underline{p}) - (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}(u - \overline{p})$$

Hence, using the expression for prices:

$$P < u - \frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}\Omega(\underline{n}, \overline{n})\varepsilon(\overline{n})(u - c) - (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}\varepsilon(\overline{n})(u - c)$$

Re-arranging we obtain:

$$\begin{aligned} P < u(1 - \frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}\Omega(\underline{n}, \overline{n})\varepsilon(\overline{n}) - (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}\varepsilon(\overline{n})) + \\ + c(\frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}\Omega(\underline{n}, \overline{n})\varepsilon(\overline{n}) + (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}\varepsilon(\overline{n})) \end{aligned}$$

And this defines $\eta \equiv (\frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}\Omega(\underline{n}, \overline{n})\varepsilon(\overline{n}) + (1 - \frac{\mu(\underline{n}_s, N_b)}{N_b})\frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}\varepsilon(\overline{n}))$.

For the second part, we first need to realize that:

$$\overline{p} - c = (1 - \varepsilon(\overline{n}))(u - c)$$

and, after some algebra:[44]

44. From

$$\underline{p} - c = (1 - \varepsilon(\underline{n}))(u - c) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(u - \overline{p})$$

we have:

$$\underline{p} - c = (1 - \varepsilon(\underline{n}))(u - c) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(u - c - (\overline{p} - c))$$

$$\underline{p} - c = (1 - \varepsilon(\underline{n}))(u - c)[1 - \frac{\alpha(\overline{n})}{\overline{n}} - \frac{\alpha(\overline{n})}{\overline{n}}(1 - \varepsilon(\overline{n})]$$

which can be rewritten as:

$$\underline{p} - c = (1 - \varepsilon(\underline{n}))(u - c)\Gamma(\underline{n}, \overline{n})$$

with $\Gamma(\underline{n}, \overline{n}) \equiv [\frac{(1-\varepsilon(\underline{n}))}{(1-\varepsilon(\overline{n}))} - \frac{\alpha(\overline{n})}{\overline{n}}\frac{(1-\varepsilon(\underline{n}))}{(1-\varepsilon(\overline{n}))} - \frac{\alpha(\overline{n})}{\overline{n}}(1 - \varepsilon(\underline{n}))] < 1$.
Now, we have that the welfare of sellers in the one market is:

$$V_s^* = \frac{\mu(N_s, 2N_b)}{N_s}(P - c)$$

whereas in the two submarket equilibrium is:

$$V_s = \frac{\underline{n}_s}{N_s}\frac{\mu(\underline{n}_s, N_b)}{\underline{n}_s}(\underline{p} - c) + \frac{\bar{n}_s}{N_s}\frac{\mu(\bar{n}_s, N_b)}{\bar{n}_s}(\overline{p} - c)$$

Hence

$$V_s^* - V_s = \frac{\mu(N_s, 2N_b)}{N_s}(P - c) - \frac{\underline{n}_s}{N_s}\frac{\mu(\underline{n}_s, N_b)}{\underline{n}_s}(\underline{p} - c) - \frac{\bar{n}_s}{N_s}\frac{\mu(\bar{n}_s, N_b)}{\bar{n}_s}(\overline{p} - c)$$

and this is greater than zero if and only if:

$$\mu(N_s, 2N_b)(P - c) - \mu(\underline{n}_s, N_b)(\underline{p} - c) - \mu(\bar{n}_s, N_b)(\overline{p} - c) > 0$$

Hence, if and only if:

$$P > c + \frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}(\underline{p} - c) + \frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}(\overline{p} - c)$$

Now, using the definitions of prices:

$$P > c + \frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}(1 - \varepsilon(\underline{n}))(u - c)\Gamma(\underline{n}, \overline{n}) + \frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}(1 - \varepsilon(\overline{n}))(u - c)$$

And hence:

$$\underline{p} - c = (1 - \varepsilon(\underline{n}))(u - c) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(u - c) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(\overline{p} - c))$$

Hence:

$$\underline{p} - c = (1 - \varepsilon(\underline{n}))(u - c) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(u - c) - (1 - \varepsilon(\underline{n}))\frac{\alpha(\overline{n})}{\overline{n}}(1 - \varepsilon(\overline{n}))(u - c)$$

$$P > c[1 - \frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}(1 - \varepsilon(\underline{n}))\Gamma(\underline{n}, \overline{n}) - \frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b)}(1 - \varepsilon(\overline{n}))] +$$
$$+ u[\frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}(1 - \varepsilon(\underline{n}))\Gamma(\underline{n}, \overline{n}) + \frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b))}(1 - \varepsilon(\overline{n}))]$$

And this defines $\gamma \equiv [\frac{\mu(\underline{n}_s, N_b)}{\mu(N_s, 2N_b)}(1 - \varepsilon(\underline{n}))\Gamma(\underline{n}, \overline{n}) + \frac{\mu(\bar{n}_s, N_b)}{\mu(N_s, 2N_b))}(1 - \varepsilon(\overline{n}))]$.

The rest of the proposition follows from these price ranges and welfare comparisons.