10.14754/CEU.2022.03

Essays on Housing and Media Markets

by

Gábor Révész

Sumbitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy at
Central European University

Supervisor: Ádám Szeidl

Budapest, Hungary

10.14754/CEU.2022.03

10.14754/CEU.2022.03

10.14754/CEU.2022.03

CENTRAL EUROPEAN UNIVERSITY
DEPARTMENT OF ECONOMICS AND BUSINESS

The undersigned hereby certify that they have read and recommend to the Department of Economics and Business for acceptance a thesis entitled **"Essays on Housing and Media Markets"** by Gábor Révész.
Dated: February 17, 2022

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
Chair of the Thesis Committee:

DocuSigned by:
Mátyás László
F608CAA1BA114DB...
László Mátyás

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
Advisor:

DocuSigned by:
Ádám Szeidl
35A428469E444FB...
Ádám Szeidl

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
Internal Examiner:

DocuSigned by:
Marc Kaufmann
86FABDB831EC4EB...
Marc Kaufmann

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
External Examiner:

DocuSigned by:
Simeon Schudy
CA6B1D979DBD4CB...
Simeon Schudy

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
External Member:

DocuSigned by:
Róbert Somogyi
8BC326F5126D4C2...
Róbert Somogyi

10.14754/CEU.2022.03

Author: Gábor Révész

Title: Essays on Housing and Media Markets

Degree: Ph.D.

Dated: February 17, 2022

Hereby I testify that this thesis contains no material accepted for any other degree in any other institution and it contains no materia previously written and/or published by another person except when appropriate acknowledgment is made.

Gábor Révész

Gábor Révész

10.14754/CEU.2022.03

Abstract

The thesis consists of three solo-authored chapters. The first chapter discusses the connection between market conditions and the prevalence of round prices on the housing market. The second chapter evaluates a policy aimed at curbing the spread of short-term rental services in Budapest, Hungary. Finally, the third chapter discusses how the behaviour of a profit-maximizing media outlet can establish a connection between the beliefs and welfare of politically heterogeneous agents. The content of the individual chapters are summarized by the following abstracts.

Chapter 1: Round prices and market power on the housing market

Round prices are disproportionally common on the housing market, and they become even more common as demand gets stronger. To explain this finding, I show that in a simple model of uncertain sale left-digit biased sellers will be more likely to post round price quotes if they gain market power, that is, when demand is higher for the item they want to sell. Intuitively, sellers trade away the monetary benefits of favorable market conditions to utility coming from their preference for round numbers. The model also predicts that depending which party (the buyer or the seller) is biased changes the type of prices that are disproportionally more common in optimum. Data on housing transactions between 2008 and 2017 from Budapest, Hungary is consistent with my model with left-digit biased sellers.

Chapter 2: Short-term rentals and house prices: evidence from a policy change

I study the effect of a local policy change that unexpectedly introduced a de facto licensing fee for new short-term rental units in a district popular with Airbnb hosts. I find strong deterrence effects on new entrants to the short-term rental market. However, local house prices fell only marginally and temporarily after the introduction of the policy: I find the largest, statistically borderline significant effects for larger and more expensive apartments. Using balance sheet and online customer review data, I cannot detect any effect on local business outcomes.

Chapter 3: Media-mediated cross-party effects of political beliefs

In a simple model of politically heterogeneous news consumers with preference for like-minded (pro-attitudinal) slant and a profit-maximizing news media, prior beliefs of left-leaning consumers will affect the posterior beliefs and average welfare of right-leaning customers. This relationship is mediated through the reporting and pricing decision of the news media. I identify two key mechanisms. First, as left-leaning

10.14754/CEU.2022.03

consumers move further to the left, the news media may gain monopoly power over rightists (and thereby extracting all of their surplus) or it may stop serving them altogether. Second, if the taste for like-minded news is high enough, the media will follow the leftward move of left-leaning customer by issuing progressively more left-leaning reporting that rightist consumer will not trust, hence right-leaning news consumers will be more wrong on average.

10.14754/CEU.2022.03

Acknowledgements

First and foremost, I would like to thank my advisor, Ádám Szeidl for the help I got from him during my stay at CEU. Someone once told me that good advisors are like midwives. While they cannot do the work themselves (which would be either impossible and/or defeat the whole purpose of the exercise), they carefully guide you through a long process of producing something new and sometimes unexpected and in every situation they have something encouraging and constructive to say that helps you not to lose sight of the final goal. During my studies, I could always rely on Ádám for careful and level-headed advice that pushed me towards structure, understanding and results. I would like to think that I learned at least a little from it.

In addition to my advisor, lot of other people helped with my thesis. I would like to express my thanks for the examiners, Marc Kaufmann, Simeon Schudy and Róbert Somogyi who provided detailed and valuable feedback on the draft of the three chapters. Early versions of these papers benefitted greatly from discussions with Botond Kőszegi and Ádám Zawadowski. Seminar participants at CEU also gave lots of useful comments and suggestions during the course of years.

Staff members at the Department of Economics (later, Department of Economics and Business) provided valuable administrative help and tried to shield us from the unavoidable bureaucratic complexities as much as possible.

I would like to thank all of my PhD colleagues, but especially Miklós Farkas, Paweł Bukowski, Péter Zsohár, Gergely Hajdú and András Kollarik. They not only kept up the good spirit inside and outside the PhD lab, but I also internalized a lot of knowledge from them without even noticing. I wish Covid and the move to Vienna had not prematurely put an end to it.

Finally, I want to say a big thank you my friends and my family, without whose support I would not been able to complete this thesis and lots of other things in life.

10.14754/CEU.2022.03

10.14754/CEU.2022.03

Contents

10.14754/CEU.2022.03

10.14754/CEU.2022.03

List of Figures

10.14754/CEU.2022.03

10.14754/CEU.2022.03

List of Tables

10.14754/CEU.2022.03

Chapter 1

Round prices and market power on the housing market

1.1 Introduction

Round numbers have long been understood to play a prominent role in markets. It is well documented that on some markets, round prices (e.g. \$1) and prices just below round numbers (e.g. \$0.99) are disproportionally more common than other prices (e.g. \$0.72). So far, most of the attention has been focused on the documentation of this phenomenon in static market settings. However, we have a poorer understanding on how the preferences that can give rise to bunching of prices around round numbers respond to changing economic incentives.

In this paper I study how changes in market power affects the distribution of prices around round numbers in a real-world high stake setting. More specifically, I look at how increased demand on the housing market (a "boom") can lead to a greater share of transactions with round final prices. My proposed mechanism works through sellers having a preference for selling their house above a round target value – or, in other words, they strongly dislike selling just below short of a round price. In addition to replicating the correlation between the relative frequency of round prices and the business-cycle, my model also generates additional predictions: first, on how the missing mass around round prices change due to a demand shock and second, that listings will round prices will take a longer time to sell. I find empirical support for these predictions in my data.

In section 2, I start with a simple model of left-digit biased sellers who have a preference for selling their houses at or over a fixed, round price, leading to an excess mass in the price distribution at round prices. Intuitively, for the seller the difference between e.g. \$198,000 and \$200,000 feels in some sense larger than the difference between e.g. \$200,000 and \$202,000 – in other words, the same two-thousand-dollar difference just below \$200,000 is given disproportionally more attention than just above the round number.

Sellers can choose the price they post their house on the market. This decision comes with a fundamental tradeoff: while a high price benefits the seller conditional on being successful (both through leaving more money in the pocket of the seller and the potential extra utility from not selling it just below a focal, round price), it can lead to a delay in selling, which is assumed to be costly.

The main, novel finding in this setting that the excess mass at round prices becomes larger as demand picks up and sellers gain more market power. Intuitively, a stronger demand for housing makes houses not only more expensive, but also easier to sell, so the psychological gain that sellers experience when selling at a round price is going to be greater in expectation. Therefore if the seller can be certain enough that the listing will eventually sell, the discrete jump coming from "rounding up" the price will be worth more in expectation – if she can charge a high price, she might as well round it up a bit. While a positive correlation between the share of round endings and the strength of local demand has been documented already at least once (Pope et al., 2015), to my knowledge my paper is the first to provide a mechanism for this reduced form result.

There is no theoretical reason why I should assume that out of the parties it is the seller and not the buyer who is left-digit biased. Therefore in the second half of section 2 I modify my setup to one where it is the buyer who values round prices disproportionally more. It turns out that while some of the comparative static remains the same, a model with biased buyers predicts markedly different patterns in the distribution of price endings.

In section 3, I take my model to the data. In the first half of the empirical section using data from residential real estate transactions between 2008 and 2016 from Budapest, Hungary, I document a substantial rise in the share of round prices over the business cycle. I find that as demand picks up after the fallout of the financial crisis, transactions are around 10 percentage point more likely to be realized on a round final price, which means that are roughly a quarter more round prices during the boom than in the bust. This change is to some extent attributable to the fact that there are more round numbers among higher prices, regardless of the strength of demand. However, about one fifth of the increase in the prevalence of round prices are not explained by this nominal effect.

In the second half of the empirical section, I more explicitly connect the predictions of my model with data. The model is able to rationalize my first empirical finding, the increasing share of round prices over the business cycle – in a way that is robust to a variety of the definition of markets and market demand. Additionally, these empirical results suggest that in the setting I study, sellers are much more likely to be biased than buyers. Finally, the model also suggests that in booms buyers round prices up more often than they round them down, which would lead to relatively more missing mass below round numbers than above. This prediction is also corroborated by the data.

My paper directly contributes to the substantial body of evidence documenting empirical patterns which are consistent with economic agents seemingly having special preferences for round numbers. For example, buyers shopping for used cars prefer models with odometers below a round number (Lacetera et al., 2012), runners exert extra effort to finish a marathon just under four hours (Allen et al., 2017), investors prefer round stock prices (Kandel et al., 2001) and drivers choose the amount of gas they buy so that the total value of their purchase is a round number (Lynn et al., 2013).

While these are choices which involve a typically small fraction of the agent's wealth, this behavior persists for economically more significant decisions too. This is evidenced by the literature that studies how market participants with non-standard preferences behave in a high-stake setting, particularly in the residential real estate market. This body of research goes back at least as early as Genesove and Mayer (2001), documenting loss aversion of house sellers in the US housing market or even to Ashenfelter and Genesove (1992), who found a significant premium in prices for condo units sold in an auction overt identical units sold face-to-face. Price quotes just short of round numbers have been documented to be frequently chosen by sellers, possibly to attract more buyers and achieve higher prices in auctions. (Repetto and Solís, 2019).

In addition to studying decision making in a larger stake environment with non-standard utility functions, I also contribute to the broader literature on how a change in the market structure interacts with such preferences, leading to otherwise hard-to-explain changes in the distribution of some market outcomes, mostly prices. Dube et al. (2018) show that employee left-digit biased for wages and employer optimization frictions can explain bunching in the hourly wage distribution and how the size and missing mass can be used to estimate the market power of employers. Backus et al. (2019) document how round prices can serve as a communication tool in online bargaining situations. Shlain (2021) estimates a model of left-biased supermarket consumers using scanner data, showing that even though retail chains react to the bias of their customers when pricing their products, they do it at an inefficiently low level.

Thematically, the closest to my paper is by Pope et al. (2015), who found that house prices divisible by $50,000 are disproportionally more common than other prices and they tentatively documented the positive correlation between the strength of demand and frequency of round final prices, but their focus was the prevalence of round numbers in the cross section and not their evolution over time a response to market condition. Additionally, in a recent paper Leib et al. (2021) looked at a closely related setting, but they emphasize buyer decisions in the process. They argue that precise (nonround) posted prices serve as stronger "anchors" in the negotiating process than round ones, leading to a heterogeneous effect of round prices settings: in a buyer's market (when buyers offer prices below the seller's quote), the seller would want a final sale price as close as possible to the initial quoted price, so it makes sense to "anchor" the offer of the buyer. The direction of this hypothesized mechanism flips in a seller's market: if buyers offer *more* than the seller quotes, it is worth anchoring the buyers as little as possible by quoting a round price.

In contrast to their paper I present and solve a formal model where the main mechanism can work either through the behavioral bias of the seller or the buyer: unlike many papers, I take no ex ante stand on which of the two parties is biased[1]. In fact, using the empirical predictions of my model I am able to make claims about which bias distribution is more consistent with the patterns I observe in my data. In addition to the alternative, simpler behavioral assumption, my paper is different as it explicitly incorporates this behavior into a supply and demand setting and it delivers implications which are consistent with some patterns I observe in the data.

This modelling choice allows me to use transactional data, as opposed to the richer datasets utilized

[1]Given that I examine the effects of left-digit bias I deal with settings where it is unclear in the beginning which part is (more) biased. In a setup with loss-averse agents, it would be logical to focus on the behaviour of the sellers, since they have a natural reference point (the price they had payed for the house when they bought it).

by other contributions on this topic: in my model I do not rely on repeated transaction or the history of bargaining to derive implications. For this, I need to assume, somewhat implicitly, that the final, transactional prices are what matter in my setting, or, rather, they matter to a such a great degree that empirically they are not overwhelmed by other mechanisms.

In line with the literature I will assume that agents are left-digit biased and they make their decisions based on the linear combination of the actual price and its floor – instead of relying on the actual price alone. Which party is biased (that is, which of them puts positive weight on the floor of the price) will turn out to have an important implication on the pattern of prices we observe in optimum. If the seller is biased , it will give rise to bunching at round prices – in contrast, if it is the buyer who is the biased party, bunching will appear at prices that are right below round prices. Additionally, I use data that spans the business cycle which makes me able to document and tentatively explain how round prices become more common in the boom periods.

I start by presenting a simple model of left-digit biased price setting with uncertain sales, solve for the equilibrium price and derive comparative statics, primarily with respect to the strength of demand. Then I replicate some known results about the high prevalence of round prices on the housing market using data from Budapest, Hungary between 2008 and 2017 and I present evidence on how demand changed over this period. Finally, I show how patterns in the data are consistent with these comparative statics derived from the model.

1.2 Model

1.2.1 Setup

There is one seller (she) who wants to sell an item (a house) to a buyer (he). The valuation of the buyer, v is the buyer's private information and it is distributed uniformly on $[0, x]$

The timing is simple: the seller posts a price (makes a take-it-or-leave-it offer) p without knowing the valuation of the buyer. The buyer can only decide between accepting or declining the offer. If the buyer accepts, payments are realized, if she declines, both parties earn 0.

In line with with the literature on round prices (e.g Shlain (2021)), I will make a distinction between the *actual* and the *perceived* price in my model. While the actual price p is the amount of money that changes hands in the transaction (should it occur), it enters the utility of the agents *exclusively* through the perceived price $\hat{p}$, which is defined as

$$\hat{p} = (1-\theta)p + \theta\lfloor p \rfloor, \tag{1.1}$$

where $\lfloor p \rfloor$ is the floor of p and $0 \leq \theta \leq 1$. The parameter θ is the preference parameter attached to the left digit: the higher is θ, the more important the integer part of the price is for the agent, and consequently, the less role the non-integer part of the price will make in decision making. This formulation nests the rational setting with θ being set to 0. I will discuss the implications of the this setup in the next subsection.

10.14754/CEU.2022.03

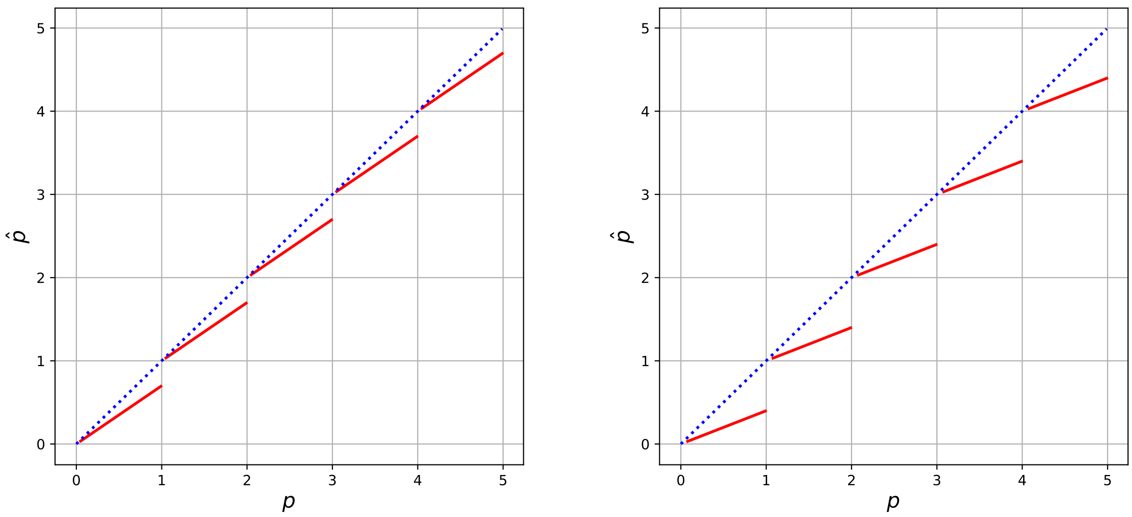

(a) $\theta = 0.3$

(b) $\theta = 0.6$

Figure 1.1: Perceived price $\hat{p}$ as a function of actual price p. (Red solid segments denote the perceived price, the blue dotted line is the 45 degree line.

Payoffs are given as follows. The buyer decides between earning 0 by declining the seller's offer or getting $v - \hat{p}_B = v - ((1-\theta_B)p + \theta_B\lfloor p \rfloor)$ where θ_B is the taste parameter of the buyer. The seller either has utility $\hat{p}_S$ with probability $\frac{1}{x}(x - (1-\theta_S)p - \theta_S\lfloor p \rfloor)$ or 0 otherwise. I use Bayesian Nash Equilibrium to solve this game. [2]

1.2.2 Discussion of modelling assumptions

First I briefly discuss the consequences of the way the perceived price $\hat{p}$ is modelled. The relationship between $\hat{p}$ and p is strictly increasing: something sold at a higher actual price will always feel more expensive. However, the relationship will not be strictly monotonic: an increase of 1 cent in the actual price in general will not directly translate to a 1 cent increase in the perceived price. It will feel less than 1 ($1-\theta$, to be exact) if p is not very close to an integer: e.g. a change from 4.45 to 4.46 increases the perceived price by $(1-\theta)\cdot 0.01$. In contrast, the change from 4.99 to 5 is perceived by the agent as a change of $5 - 4.99(1-\theta) - 4\theta = 0.01 + 0.99\theta$. In other words, the agent with $\theta > 0$ undervalues the effect of price changes if this does not come with the change of the digits left of the decimal point (compared to the non-biased $\theta = 0$ agent), but overvalues it if it does. This makes the graph of the function $\hat{p}(p)$ to have bumps at integers and linear segments with a slope $(1-\theta) < 1$ between integers. Both of these deviations from the $\theta = 0$ case increase in θ as it can be seen in Figure 1: a larger θ leads to flatter segment between integers and hence a larger jump in utility right at integers.

[2] Even though the game is sequential, using PBE as the equilibrium concept does not add much to the discussion: this comes from the very simple structure of the game but mostly from the fact that (unlike in standard signalling games) it is the uninformed party who moves first. This creates a situation where even though there is uncertainty about one player's type (the buyer's valuation), there is no meaningful learning about it, basically eliminating off-the-equilibrium nodes.

In addition, I make the simplifying assumption that the seller fully knows the unique taste parameter of the buyer. I do this for two reasons. First, in this setup I can abstract away from the question of how people think of other's tastes. While it theoretically possible that a seller would project her own taste for round numbers on to the buyer, I have no strong motivation for this additional assumption and if θ is considered a simple taste parameter, it is a fairly standard to assume that, just like in most other economic models, the preferences of the players are common knowledge.

Plus, there is some evidence that people are quite sophisticated when they have to reason about the degree of other people's biases: Fedyk (2021) finds that experiment participants can predict the level of present bias much better for others than for themselves. While results about present bias are not guaranteed to carry over to other settings (especially given that $\beta - \delta$ discounting can lead to time-inconsistent choices in a way that left-digit bias does not, therefore the choices made by a left-digit biased agent should not surprise neither her nor an outside observer), it is consistent with a setup where the seller understands her own tastes and those of the buyer too.

Secondly, by having one buyer in a one-period game I do not have to address the intertemporal choice problem that would arise in e.g. a dynamic game where in every period t the seller would meet a new buyer with taste parameter θ^t_B drawn from some known distribution. This allows me to concentrate on the comparative static that I will bring to the data in the empirical section.

1.2.3 Results

In this section I present the results of two special cases of my model: first, the case with a non-behavioral buyer ($\theta_B = 0$) and a seller who can have an arbitrary degree of left-digit bias ($0 \leq \theta_S \leq 1$); next, the exact opposite of the first case (non-behavioral seller, behavioral buyer).

In both cases, I will proceed in a similar manner: first, I argue that a certain type of prices (integers or prices ending in .99) is going to be disproportionally more common, next, I characterize the set of parameters under which such prices are optimal and finally I provide comparative statics of the structure of prices with respect to the model parameters.

Seller bias

I start with the case when $\theta_S > 0, \theta_B = 0$. The expected utility of the seller (without the normalizing constant) simplifies to

$$(x - p) \cdot (\theta \lfloor p \rfloor + (1 - \theta)p) \tag{1.2}$$

No matter the bias, there is clearly a tradeoff between (1) the probability that price set by the seller is lower than the buyer's valuation, hence the sellers sells, and (2) the utility the seller enjoys in that case.

In this case, however, the seller will attach a relatively smaller weight to the right digits of the price, making him more willing to give up some of the price to increase the probability of sale. This is possible because if such a marginal decrease in price takes place strictly between two integers, leaving the floor term unchanged, its marginal effect on the second term of the expected utility is scaled by $1 - \theta$, while increasing the first term in the same way as it would do for an unbiased seller – since the probability of sale is determined by the unbiased buyer.

These two characteristics explain the overall shape of the graph of the utility function in Figure 1.3. Broadly, they look like a concave parabola which captures the tradeoff between the higher price and the higher probability of sale, which carries through from the case of standard preferences with $\theta = 0$. In addition to that, we find jumps at every integer value of the price, representing the discontinuous utility change that comes from the seller's left-digit bias. [3]

Since I will want to concentrate on the comparative static with respect to the maximum valuation x,[4] I will find the following approach useful. To characterize the optimal pricing decision of the seller, first I define the interval $(x_L(k), x_H(k))$ which contains the x's for which it will be optimal to charge k (where k a nonnegative integer). After that it will be sufficient to determine the optimal price between $x_H(k)$ and $x_L(k+1)$, if such prices exist.

This approach will define the optimal p for every combination of the parameters and will give a tractable way to analyze the comparative statics of the model. Most importantly, I will call a certain price p *more likely* if the width of the interval $(x_L(p), x_H(p))$ increases in some parameter.

Proposition 1 (Optimal pricing if only the seller is biased.)**.** *If $\theta_S > 0$ and $\theta_B = 0$, the following is true:*

1. *The optimal price is at most distance 1 from the unconstrained optimum.*
2. *The optimal price is a weakly increasing function of the maximum valuation x.*
3. *As x goes up, integer prices become more likely.*

Proof. All proofs in the appendix. □

The first two statements, in addition to being necessary to state the third one, imply that rounding is in some sense a local phenomenon: the price quoted by a $\theta > 0$ seller quantitatively do not differ wildly from the prices an unbiased seller would offer and the common intuition from the standard setting that higher demand (x) leads to higher prices carries through.

The third statement establishes a comparative statics on the likelihood of round (integer) prices. Figure 1.2 gives an illustration of this result for a specific set of parameters. We have already seen some of intuition of why round prices can be common for a fixed set of parameters: quoting an integer k rather than $k - \epsilon$ will decrease the probability that the buyer accepts only marginally, but it increase the seller's utility in a discrete way. If this jump is high enough, the seller will find it optimal to "round prices up": even though an unbiased seller would find it optimal to charge something slightly below k, with $\theta > 0$ the seller will quote k instead. Similarly, there will be parameters under which an unbiased seller would set a price that is somewhat higher than some k, however with $\theta > 0$ it will be relatively cheap to attract

[3]The size of the jump measured in utility terms is $\lim_{\epsilon \to +0} U(k) - U(k - \epsilon) = (x - k)\theta$. A marginally small price increase that pushes the seller just above an integer k delivers an extra θ utils if she sells, since the term $\theta \lfloor k \rfloor$ in the utility function goes up by $\theta \cdot 1$. This happens with (the unscaled) probability of $x - k$.

[4]I will assume that when demand goes up, the distribution of valuations will shift to right rather than it expanding. This makes it easier to think about the effect of demand changes, as is this case the mass of buyers between two (internal) prices will stay constant.

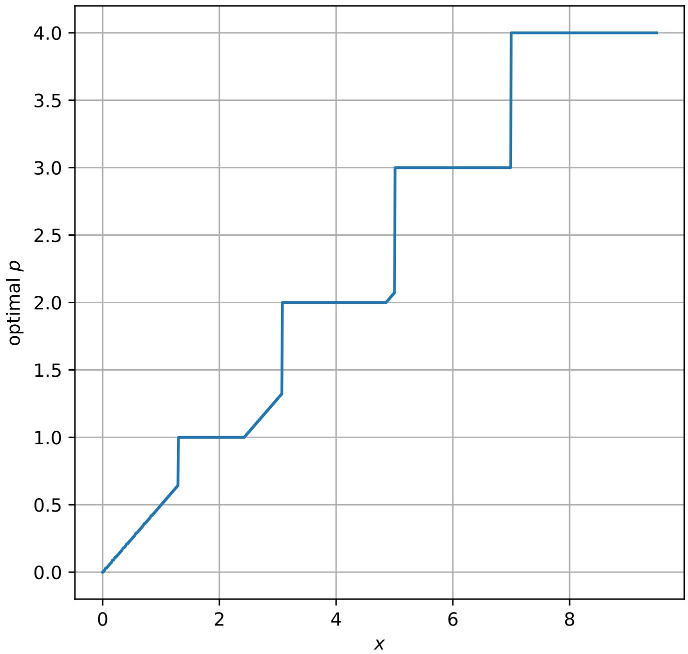

Figure 1.2: The optimal price as a function of x ($\theta = 0.3$)

the buyer by setting a somewhat lower price. These two forces lead to round (integer) prices being more common optima in the parameter space.

Moving on to the intuition of the comparative static, it is useful to rephrase the previous argument through the question how the seller weights the relative importance selling with a higher probability versus selling at higher price.

Having $\lfloor p \rfloor$ with a positive weight in the utility function ($\theta > 0$) modifies it in such a way that the relative importance of the marginal utility of price and the marginal utility of a more certain sale going to depend on the position of the price relative to the closest integers. If p is just above $\lfloor p \rfloor$, a marginal increase in the price would be less appealing to a biased seller, as she would weight the extra money coming from a higher price less than an unbiased seller would. However, she would correctly perceive the impact that it has on the buyer's probability to accept. In this case, her left-digit bias would make her to relatively *overweight* the demand impact of the price hike and *underweight* the marginal expected utility of money.

In contrast, if she is considering the effects of a small price increase from $k - \epsilon$ to k, she would relatively overweight the utility change coming from the price change compared to the utility change coming from the somewhat smaller probability that the buyer accepts her offer (again compared to an unbiased seller). In other words, the jump from 4.9 to 5 will feel very appealing, leading her to underweight the impact this pricing decision has on demand. This intuition holds with x fixed.

Now I move on to the effect of a change in x on the prevalence of round prices. If x goes up , the incentive to sell increases with it: whatever the seller is offering becomes more valuable. That makes the seller more willing to marginally reduce the price and enhance her chances of selling. This effect of a demand increase, however, is *increasing in the bias* θ. As we have seen above, reducing the price from e.g. 5.1 to 5.0 is "cheap" for a biased seller, who weights such marginal changes with a factor of $1-\theta$ when she is thinking about the money she gets if the buyer buys. If demand goes up, which makes her already anxious not to lose the buyer, her relative overweighting of the demand impact of her price cut will make her round down even more.

Technically, this will lead to $x_H(k) - 2k$ increasing in x. This means that the set of x's that lead to an optimal integer price k through rounding down (that is, an unbiased seller would charge something between k and $k+1$, but a biased seller charges k) expands. [5]

This describes the effect of a demand shift on the seller rounding down (and consequently on $x_H(k)$). How does the same shift in demand change how the seller rounds up (and consequently, $x_L(k)$)? It turns out that $2k - x_L(k)$, the set of x's for which the biased seller rounds up to k, is increasing in x too. I offer a partial intuition for this result. For this, let's decompose the change in utility coming from a price change:

$$U(p_2) - U(p_1) = Pr(p_2)\hat{p}_2 - Pr(p_1)\hat{p}_2 = (Pr(p_2) - Pr(p_1))\hat{p_1} + Pr(p_2)(\hat{p}_2 - \hat{p}_1) \tag{1.3}$$

Now let's look at the incentives to round price up by setting $p_2 = k$ and $p_1 = k - \Delta$:

$$\begin{aligned} &(Pr(k) - Pr(k-\Delta))k + Pr(k)(k - \theta(k-1) - (1-\theta)(k-\Delta)) \\ = \quad & \underbrace{-\Delta \cdot k}_{\text{lost on the marginal buyer}} \quad + \quad \underbrace{(x-k)(\theta + (1-\theta)\Delta)}_{\text{gained on inframarginal demand}} \end{aligned}$$

When the seller considers to round the price up to the nearest integer, her profits are affected by two forces: on one hand, she sells with a somewhat lower probability (she loses some profit on the marginal buyer, who will now not buy at a higher price); on the other hand, she makes more if she eventually sells (a gain on the inframarginal buyer). These two forces exactly offset each other whenever the seller is indifferent between charging a higher, round price and lower non-round price. It turns out that the second force is greater whenever demand is stronger (that is, when x is higher).

The partial intuition is as follows. Consider a seller who under some parameter values was indifferent between quoting 4.9 and rounding up to 5. Now, x increases marginally. This change shifts the distribution of buyer valuations slightly to the right, but, as it is very small, does not make it profitable to the seller to increase her price above 5. How did the incentives for rounding up change? The expected loss on the marginal buyer stayed the same: increasing the price by 0.1 leads to the same decrease in the probability of sale – the buyer loses the same "mass" of demand, as the density of buyer valuations shifted to the right. However, the potential gain on the inframarginal buyer went up: a price of 5 will be accepted by the buyer with a slightly higher probability due to a higher x, making rounding up slightly more attractive. This reasoning is partial because it considers only small changes in x and does not account for jumps in the value of the integer candidate price (e.g. it can not explain how the buyer thinks if the optimal integer candidate changes from 5 to 6 due a to a change in x). Furthermore, it can be shown that $\frac{\partial}{\partial k}(x_H(k) - 2k)$ for meaningfully large k's is smaller than $\frac{\partial}{\partial k}(2k - x_L(k))$, which means

[5] Remember, for $x = 2k$ the all sellers charges k and for $x \in (2k, x_H(k))$, the (biased) seller also charges k by definition.

10.14754/CEU.2022.03

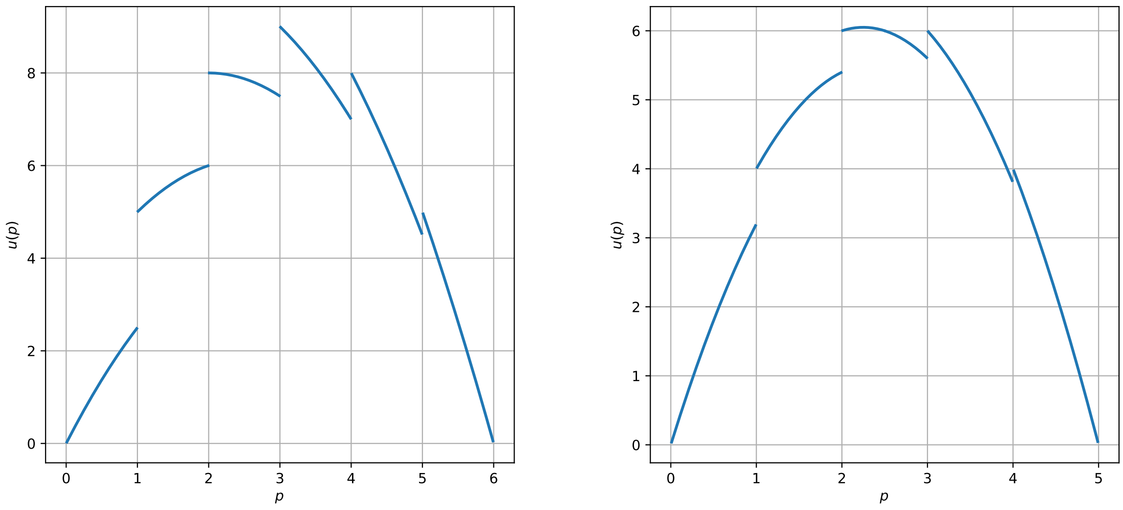

(a) round optimal price, $\theta = 0.5$

(b) internal optimal price, $\theta = 0.2$

Figure 1.3: Utility functions for different values of θ_S (seller bias case).

that interval on which a biased seller rounds up increases faster in k than the interval over which she rounds down – leading to, on average, higher prices from an unbiased seller's point of view.

Buyer bias

I continue with case when only $\theta_B > 0$. Under this assumption the expected utility of the seller simplifies to

$$(x - (1-\theta)p - \theta\lfloor p \rfloor)p. \tag{1.4}$$

Now it is the first term in the utility function – the probability of sale – the one which is discontinuous in the price (as opposed to the previous case when it was the conditional utility of the buyer). This comes from the fact that a change of price from 4.99 to 5.00 will loom larger for the buyer than e.g. a change from 5.01 to 5.00, leading to a discrete drop of the probability that he will buy the object that the seller offers. This discontinuity will carry over to the utility function (Figure 1.4). While the utility function looks broadly similar to that in the case of the biased seller, discussed in the previous subsection, there is an important point of difference too, which will affect the type of prices quoted in equilibrium. Since now the biased party is the buyer and he prefers low prices – as opposed to the previous case, when it was the seller who obviously liked high prices – the jump from e.g. 4.99 to 5.00 will in general lead to a discrete drop in seller profits, so the seller, instead of choosing integer prices often, will pick prices that are just below integers. [6]

[6] In the discussion I will refer them to as "prices just below an integer" or "prices ending in .99". I find this approach useful, since the utility function is not continuous at integers k (rather, it is continuous only at the intervals $[k, k+1)$), which, because these sets are half-open, would lead to somewhat unintuitive outcomes like having no equilibria over the set of all reals. Therefore I treat these cases as if there existed a discrete grid of prices with a well-defined maximum – real-world prices are in fact constrained by the unit of account used.

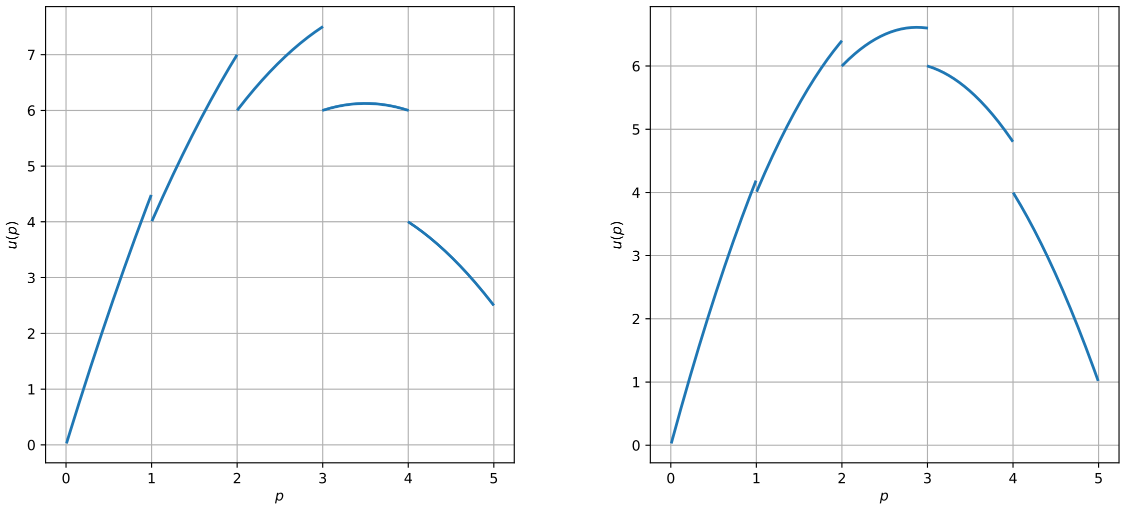

(a) round optimal price, $\theta = 0.5$ (b) internal optimal price, $\theta = 0.2$

Figure 1.4: Utility functions for different values of θ_B (buyer bias case).

Some features of optimal pricing is given by the following proposition.

Proposition 2 (Pricing if the buyer is biased)**.** *If $\theta_B > 0$ and $\theta_S = 0$, the following is true:*

1. *The optimal price is at most distance 1 from the unconstrained optimum*
2. *The optimal price is increasing in the demand (x).*
3. *Higher prices are more likely to be the largest quotable price that is strictly less some integer.*

This proposition is very similar to Proposition 1 in structure and in content, the main difference is that the third claim is no longer about integers, but about prices that are just below integers.

To elaborate, the third statement claims two things. First, as we have seen before, if it is the buyer who is biased, the type of prices that will be more common in equilibrium are the ones that "end in .99". Secondly, it establishes a comparative static with respect to x: as demand goes up, prices ending in .99 will be more common and ultimately will crowd out every other type of price.

The intuition is similar to the previous case. Just above an integer price (say, at 5.1) it will be very appealing to reduce it somewhat (in this case, to 4.99), since it comes with a discrete jump in demand. If prices are high, the buyer is worth more to the seller (as he generates more revenue), so price reductions from $k + 0.1$ to $k - 0.01$ will become more valuable. The difference from the previous case is that now it is not the seller's underweighting of money that makes such decisions profitable, rather it is the buyer's left-digit bias that leads to a disproportionate demand effect. The intuition for the expanding lower bound is very similar to the buyer-bias case.

Somewhat imprecisely, however, I will assume that the difference between such a maximal price p "ending in .99" and the nearest integer (effectively the ceiling of p, $\lceil p \rceil$) is negligible compared to p, which allows me to simplify the interpretation of my results.

10.14754/CEU.2022.03

To sum up, both models predict a disproportionate share of prices that are either (a) round or (b) just under a round price – depending on which party is biased. Furthermore, both models predict that as buyer valuations go up, the share of such prices (round or below round) go up with it. Finally, biased sellers tend to round up more than they round down on average as demand increases.

1.3 Data

1.3.1 Data description and context

My data consists of residential housing transaction from Budapest between 2008 and 2016 sold by a major real estate agency during this period. For every listing I observe its final price, its precise address, the day the listing went online and when it was sold. In addition to that, a wide range of hedonic characteristics are also available (such as the number of rooms, size of the property, the presence of a balcony or a garage etc.). Unfortunately, I only observe the final price and I have no information on the potential bargaining process that could have taken place.

Since in this paper I focus on the effect changing demand, first I look at how market conditions evolved through this period. Not unconnected to the aftermath of the global financial crisis, after 2008 prices and the number of transactions dropped and stayed low up to 2013, when they started to rise markedly.

As evidenced by Figure 1.5, the price of one square-meter has decreased by approximately 10 log points between 2008 (the start of the financial crisis) and 2013 (the trough of the recession in the housing market), while it has grown by almost 40 log points between 2013 and and 2016 (the final year in my dataset). This increase translates to an approximately 60% growth in this three-year period.

The "hotness" of the market may also be captured by the time it takes to sell a listing. If housing supply is relatively stable in the short run, more buyers searching for a suitable house might not only mean higher prices, but faster sales: since there are more buyers looking for houses, matches are created with a higher probability (Díaz and Jerez, 2013). Broadly speaking this is what the right panel of Figure 1.5 shows: as demand kept dropping from 2008 to 2013, it took almost twice as long to sell the median listing, while after 2013 it went to back to levels comparable in 2008. This time pattern is more pronounced for the average as opposed to the median selling time: the average increased by a factor 2.5 between the peak and the trough.

Secondly, both the price and the time-to-sell data points to 2013 as the turning point. This remains true even though mean selling peaks in 2014. The reason for that is somewhat mechanical: as demand grew in 2014, listings which have been on the market for a really long time also started to sell. Only after this "backlog" was cleared out could the average of days-on-market start to fall.

The data I use might not be representative of the market as houses sold through an agency can be systematically different to houses sold through other channels. To address this problem, I compare my observations with data collected by the national office for statistics in figure 1.6.

First I plot the the average price level of transactions in Budapest in the administrative data collected by the statistical office and in my data. In contrast to Figure 1.5, the first panel of Figure 1.6 features average transaction prices, not the average price of a square meter, since only the former is published

10.14754/CEU.2022.03

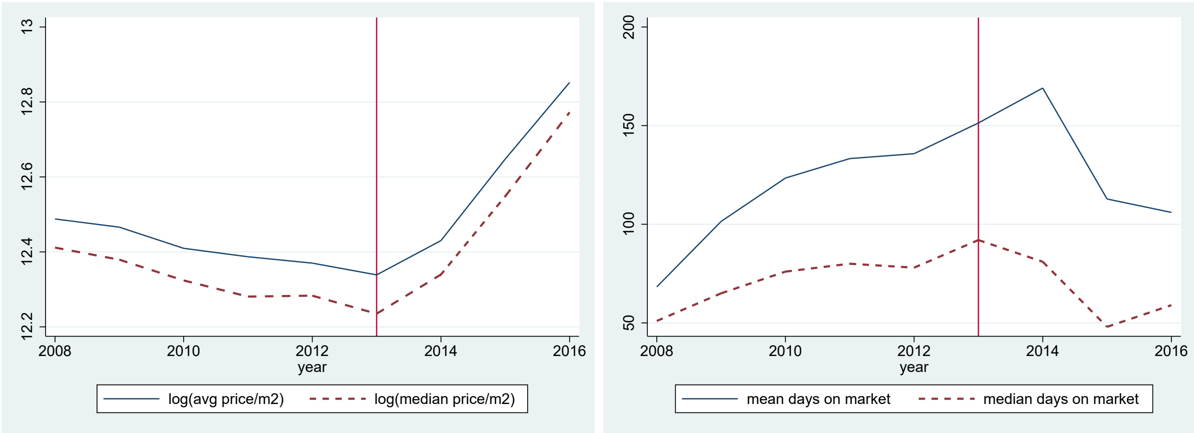

(a) Price

(b) Days on market

Figure 1.5: Market outcomes over time, 2008–2016

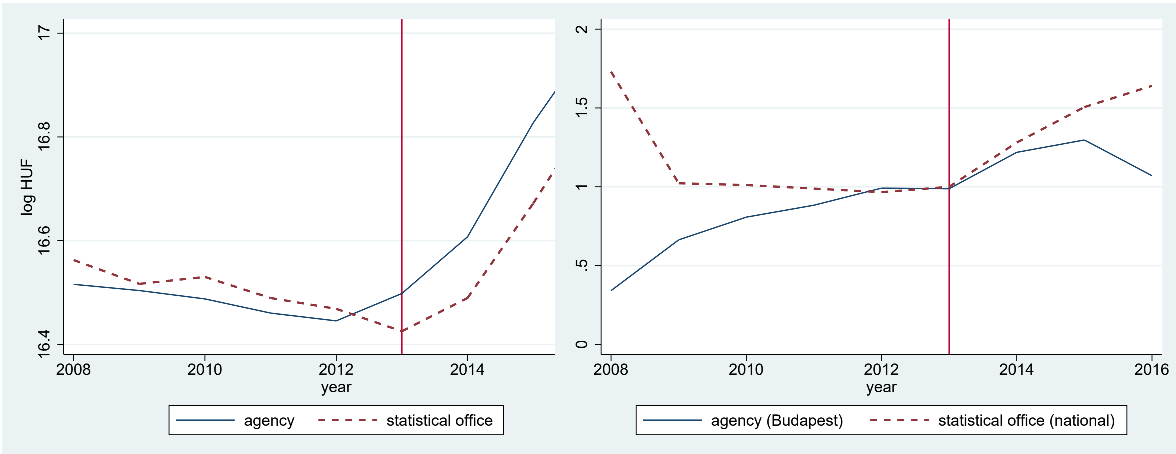

Figure 1.6: Houses sold through agencies vs recorded in the national statistical office. (average national log price and normalized transaction count)

by the statistical office. Before 2013 the average sale price in the datasets track each other very closely. Starting in 2013, a gap appears between the two series: houses sold by the agency tend to be more expensive than the average house on the market. This might be at least partially due to the fact that agency-sold houses seem to be larger in size: while the average square meter price of an agency flat reaches the trough in 2013, the average total price of an agency flat is already considerably higher in 2013 than in 2012. This gap persists and stays stable after 2013.

In the second panel of Figure 1.6, I compare transaction volumes between the two datasets. Official transaction statistics however is available only at the national level, so I normalize both series by their respective levels in 2013. Firstly, 2013 is an important structural break not only in prices, but transaction volumes. Secondly, the agency data features much fewer transactions in towards the endpoints of my timeframe. These missing observations are particularly visible in 2008 and 2016, the first and last year of my sample period. This is likely to be due to data collection issues: looking at the distribution of months in these two years reveals that there are particularly few observations in the first months of 2008 and the last months of 2016.

Even though all three variables I have looked at are clearly endogenous (price, quantity, time to sell), together they tell a consistent story of the evolution of the market over time. Since in my analysis I will focus on the effect of changing demand on various market outcomes, I will split my sample into to periods, the first from 2008 to 2013, (*low demand period*), and the second one, from 2014 to 2016 (*high demand period*).

While the language I use to refer to these two periods suggests that they are different primarily because of different levels of demand, one alternative mechanism that could lead to higher house prices is high levels of inflation. First, as the previous graphs showed, not only prices, but transaction numbers soared after 2013 and it took less time on average to sell a listing, which is suggestive of real changes on the market (as opposed to purely nominal ones). Secondly, consumer prices happened to grow slower in the second, high-demand period than in the first on, characterized by low demand: the average annual CPI in the first period was around 4.9%, while in the second period it averaged around 0.5%. Taken together, these findings suggest that the increase in house prices was most likely to have been caused by an increase in demand and not by an economy-wide increase in prices.

As I have described above, I use change in demand over time as a source of variation in my empirical analysis. While arguably this was a macroeconomic phenomenon connected to the great financial crisis of 2008 and its aftermath and had plausibly little to do with buyers' or sellers' preferences for round numbers, this is obviously not the cleanest of identifications. The reason I use this measure of demand in my analysis, apart from data availability issues, is that, since the size of the effect I am after is small, I need large variation in demand for my empirical exercise and I have not been able to identify a context with sufficiently large and plausibly exogenous variation in demand.

Another feature of my data that may seem like a limitation is the lack of information on the buyer or the seller. This makes conducting analyses like that of Ross and Zhou (2021) impossible: they used repeated house sales and prior seller behavior, such as the exact value of the mortgage that the sellers took out when they purchased the property they are now selling, to obtain individual-level measures of preference for round numbers. However, since my modelling framework does not feature loss-averse agents, I do not explicitly need this information to derive predictions about the distribution of price endings under different conditions.

1.3.2 Stylized facts

Next, I present evidence that there is some excess mass present at round prices.

For some of my results I will assign prices to bins based on their endings. In line with my hypothesis of left-digit bias, I truncate prices to the hundred thousands: if a price is not divisible by 100 000 forints, I round it down to the nearest hundred thousand. Almost 90% of the prices in my data are divisible by 100 000, while only one third of the prices are divisible by one million, so this level of truncation seems like a reasonable one.[7]

Round number are common and get more common over time. Figure 1.7 presents evidence

[7] 100,000 forints hovered roughly between 350 and 400 US dollars during the years in my sample.

10.14754/CEU.2022.03

of the prevalence of round endings. If prices were distributed randomly across bins, then all bins would

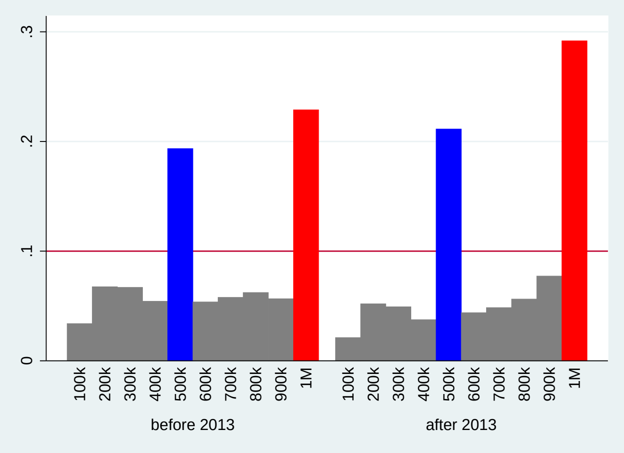

Figure 1.7: Distribution of prices by endings, before and after 2013

contain about one tenth of the observations. What we see instead is the dominance of prices ending in half a million and million forints: they represent around half of the transactions. Furthermore the share of transactions divisible by one million increased markedly from the low demand period to the high demand one: 22.9% of the final prices were divisible by one million before 2013 , which went up to 29.2% in the period between 2014 and 2016. This difference between the periods of low and high demand is statistically significant too, confirmed by a transaction level regression (see appendix)

Round prices are often high, but they get more common in almost every decile. Next, I zoom in on individual transactions and try to decouple the effects of rising prices and increasing demand. This distinction might be necessary if prices which are "high" (in some sense) are rounded with a greater probability. This can happen either because of convenience or for behavioral reasons: e.g. the same fixed difference may loom smaller when compared to a higher final price. Furthermore, buyers and sellers at the same *nominal* price might be potentially different over the business cycle: for example, if rich buyers buy relatively expensive houses and at the same time they care less about rounding, comparing listing based on their nominal prices would lead to comparing a relatively expensive listing (bought by a rich buyer) in a low demand period with a relatively cheap listing (bought by a less rich buyer) in the high demand period. This would understate the change in round transactions over time.

Figure 1.8 shows the share of round transactions by (annual) deciles. First, since both graphs have a positive slope, relatively more expensive listings have a greater tendency to sell at a round price. The expected difference along the price distribution is quite sizeable: going from the second to the ninth decile almost triples the probability that a listing will fetch a round price. More importantly, in almost every decile the share of round transactions are greater in 2015 than in 2012, pointing at a more or less uniform increase in the prevalence of round prices.

In the following step, as before, I zoom in on individual level transactions. In the following regression, I compare listings which are in the same *relative* position in the price distribution in a given year, by including the term $\ln(p_{it}/\bar{p}_t)$, the relative position of the listing compared to the annual median price in that year ($\bar{p}_t$):

10.14754/CEU.2022.03

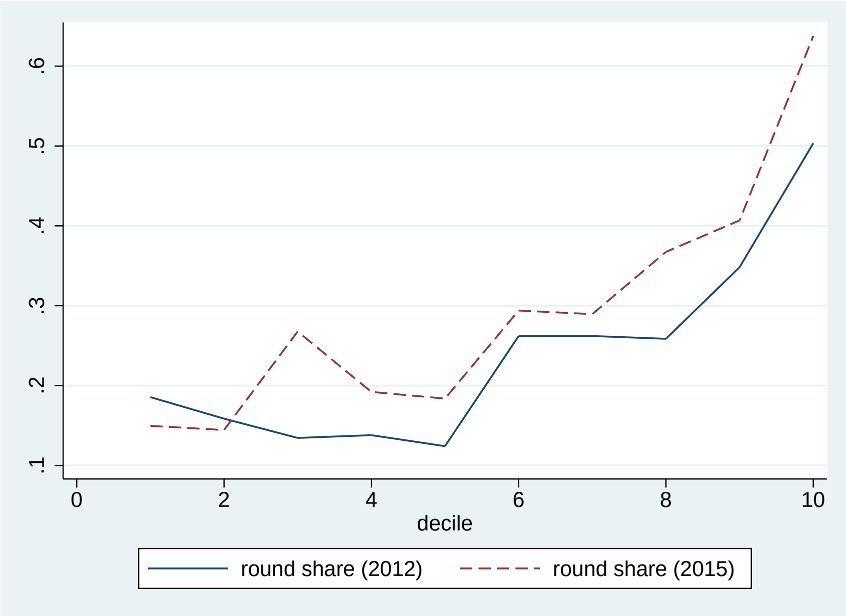

Figure 1.8: Share of round prices by deciles.

$$round_{it} = \beta_0 + \beta_1 \ln(p_{it}/\bar{p}_t) + \text{year dummies} + \epsilon_{it}. \tag{1.5}$$

For this approach I also include a full set of year dummies, instead of a binary before-after variable. Figure 1.9 shows the year-dummy coefficient estimates and their standard errors.[8]

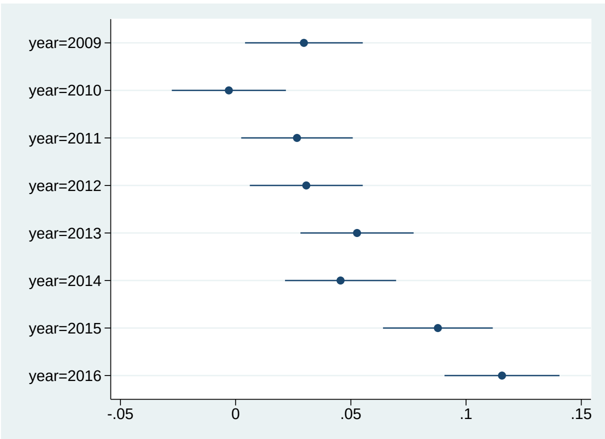

Figure 1.9: Year-dummy coefficient estimates and their standard errors with $\bar{p}_t$ = yearly median in year t. Omitted (reference) year: 2008.

The estimated value of β_1 is positive and strongly significant (not shown in the graphs). This means that listings that were more expensive compared to the annual median had a higher probability of being sold at a round price – in every year.

Furthermore, keeping the relative position of a listing fixed, the probability of it being sold at a round price is somewhat higher after 2010, and more importantly, it is much higher after 2013. As it is apparent

[8] Note that by adding year dummies to the equation, the estimate of β_1 is going to be identical, regardless whether we compare listings with the same relative or absolute position in the distribution.

from the graph, a listing that were in a given relative position in the price distribution in 2016 had about 11% higher change of being sold at a round number than a listing at the same relative position in 2008.

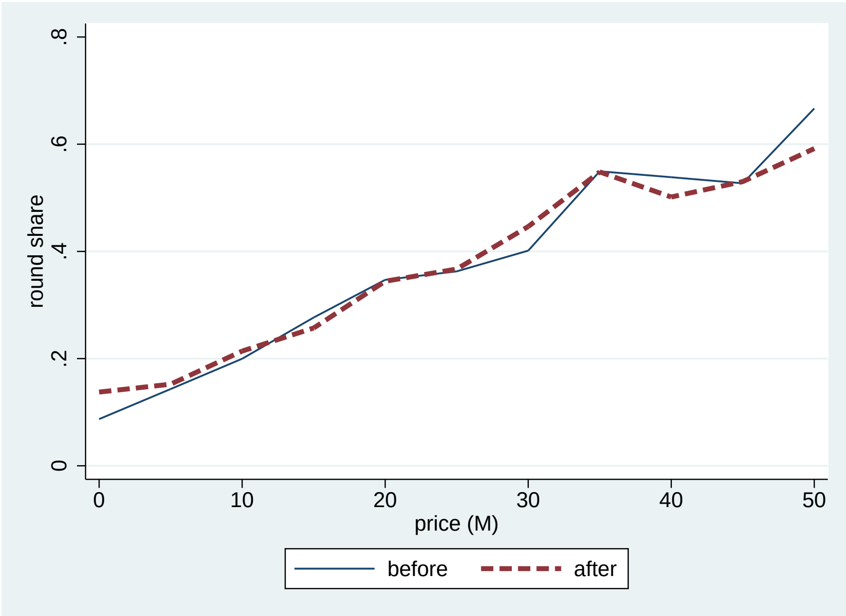

Figure 1.10: Share of round prices in 5 million forint bins for transactions under 50 million, pre- and post 2013.

Decomposition. Write the conditional probability of listing i in group $j \in \{B, A\}$ being round ($y_i = 1$) as

$$y_i^j = X_i^j \beta^j + \epsilon_i^j.$$

Then the average difference between the before (B) and the after (A) group can be decomposed as

$$\bar{y}^A - \bar{y}^B = (\bar{X}^A - \bar{X}^B)\beta^B + \bar{X}^B(\beta^A - \beta^B) + (\bar{X}^A - \bar{X}^B)(\beta^A - \beta^B)$$

The first term of decomposition is usually called the *endowment effect*: it captures the share of the gap which is due to difference in the X's between the two groups. The second term measures the difference due to *different coefficients* in the two groups, while the remaining *interaction* terms captures the possible correlation between β and X across groups.

In my data the raw difference between the after and before group is around 0.063: the share of round transactions are 6.3 percentage points larger after 2013 than before. This corresponds to the difference between the columns in Figure 3. Changing endowments account for 4.9 percentage points out this gap: this is the part of the increase that can be attributed to the fact the higher prices tend to be round more often. In contrast, the rest, 1.2 percentage points, about 19% of the is due to an increase in probability of roundness in every price bucket.

To sum up, I have found that round prices are disproportionally common in my data. Secondly, they became even more common from 2014 onward, when the demand picked up and prices and transaction numbers increased. While this difference is not significant if I compare listings at the same nominal price, this effect seems to be driven by more expensive houses – that is, round prices have indeed become more

high demand	0.289***
	(0.004)
low demand	0.226***
	(0.004)
difference	0.062***
	(0.005)
endowments	0.049***
	(0.003)
coefficients	0.012**
	(0.006)
interaction	0.001
	(0.003)
N	27262

Table 1.1: Decomposition of the increase of round shares. Dependent variable: dummy for round. Includes controls for hedonic characteristics and log of size.

common if I focus on the bottom two-thirds of the market. Furthermore this differential effect is strongly there if I compare listings which are in the same relative position.

1.3.3 Evidence

In this section I present empirical patterns which are consistent with the corollaries of the model.

Which party is biased matters for price endings. It has already been established in the literature that some price endings are disproportionally more common in some markets (e.g. round numbers on the housing market and prices ending in .99 in supermarkets). Given the predominance of round prices in my data, and the fact that the share of round numbers go up with demand, I argue that this is consistent with biased sellers.

The share of disproportionally common price endings reacts to market conditions. In addition to the well-known fact that some price endings are more common, the model outlined in the previous section predicts that in markets where demand is higher, sellers are more likely to set round prices. This is consistent with the empirical results in section 2: as demand gets stronger after 2013, the share of round prices increased. I also observe a slight uptick in relative share of prices falling in the 900 thousand bin, which is consistent with a moderate degree of buyer bias.

In addition to this, I present further evidence using an alternative, more disaggregated measure of local market conditions. Instead of calendar time only (pre-2013 vs post-2013, as in section 2), I use the average number of days-on-market of neighboring listings as a proxy for local demand. I expect this variable to capture local demand for two reasons. First, a large part of housing search is local (as documented by Piazzesi et al. (2020)), so I expect changes in local demand conditions to be contained in the neighborhood. Secondly, these neighboring listings are expected to be similar in other, non-observable characteristics, so changing demand for non-observables will lead to locally correlated demand shocks.

To calculate this measure, I simply take the average of days it takes to sell a listing by year, district and main building type (single family unit, condominium and pre-fab concrete blocks). Then I regress a dummy for round price on this variable while controlling for log-price and adding fixed effects. In this setting, less time to sell would be associated with a hotter market, therefore more round numbers. Indeed this is what we see in Table 1.2

Outcome: 1 if round, 0 otherwise.	(1)	(2)	(3)
log(local average day-on-market)	-0.0105	-0.0921***	-0.0426***
	(0.0072)	(0.0077)	(0.0127)
Observations	22877	22877	22877
year FE	no	no	yes
type of building FE	no	yes	yes

Standard errors in parentheses

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

Table 1.2: Relationship between the final price being round and local average of time-to-sell. Outcome: 1 if price is round. All regressions include the log of relative price, $\ln(p_i/\bar{p}_t)$ as a control (not shown).

The missing mass comes from below round numbers. The model also predicts that in a market which is better for sellers, sellers will adjust their prices *upwards* more frequently (as $\frac{\partial}{\partial k}(2k - x_L(k)) > \frac{\partial}{\partial k}(x_H(k) - 2k)$ for large enough k's) Intuitively, a better market for sellers allows them to make choices which satisfy their biases more.

As a corroborative evidence for this, I examine how the relative mass of some price endings have changed over time. Table 1.3 shows results for the diff-in-diff

$$mass_i = \beta_0 + \beta_1 D_i + \beta_1 \text{post-2013}_i \cdot D_i + \beta_2 \text{post-2013}_i + \gamma' \mathbf{x_i} + \epsilon_i, \tag{1.6}$$

where observation i the relative mass of prices that fall in histogram bin and D is a dummy for the price ending of interest. I look at differential change in three price categories: just below the round price (prices ending in 900 000, column 1), the round price itself (prices ending in one million, column 2) and just above the round price (prices ending in 100 000, column 3).

$D = 1$, if price ends in:	900k	1M	100k
post-2013 · D	-0.0791**	0.0867**	-0.0300
	(0.0350)	(0.0407)	(0.0314)
D	-0.1844***	0.2195***	-0.1506***
	(0.0246)	(0.0296)	(0.0225)
post-2013	-0.3616***	-0.4483***	-0.3616***
	(0.0310)	(0.0794)	(0.0876)
N	1902	1902	1902

Table 1.3: Differential change in price masses pre- and post 2013. Observation: relative bin count in a histogram before/after. All regressions include price polynomials and price dummies, pre- and post-2013. (coefficients not shown)

The second row (the estimated price category coefficients) reproduce the known stylized fact that round prices are more common before 2013 than prices with neighboring endings. The parameters of interest are in the first row: they show that round prices became disproportionally more common after 2013, while prices ending in 900 thousands became somewhat less frequent, which is consistent with the prediction that upward price adjustment are more common during the boom.

1.4 Conclusion

In this paper I document empirical facts using data on housing transactions from Hungary between 2008 and 2015 which are consistent with sellers being left-digit biased. While left-biasedness has ample documentation, some of the findings I uncover here are novel: I show a strong and economically significant connection between market condition (the business cycle) and the share of prices with round endings. These empirical patterns are robust to the exact definition of the business cycle and to the definition of roundness. I argue that these stylized facts in the data are consistent with a simple model where sellers are left-digit biased (as opposed to buyers).

A natural step forward would be to characterize the market equilibrium in a general setup where both sellers and buyers are left-digit biased, possibly at different levels. This has the potential not only to offer a richer set of dynamics but this way the model could be linked closer the empirics.

While in my application I had good measures for the strength of demand, this line of work has the potential to be used in settings for measuring change in demand conditions where such proxies are not available. For example, in an online auction the econometrician typically does not observe demand shifters for a given product, either because they are hard to define (e.g., they are very specific to the object being auctioned and at a given time there are lot of different objects pn sale) or they are not available to the econometrician (e.g. visitor counts are kept private by the auction platform). In such situations the framework of this paper may help to produce variables that are correlated with measures of demand – relying solely on the observed distribution of prices.

10.14754/CEU.2022.03

10.14754/CEU.2022.03

Chapter 2

Short-term rentals and house prices: evidence from a policy change

2.1 Introduction

Short term rental platforms, such as Airbnb have expanded rapidly in the past decade. There exists a perception that this process has inflicted some negative externalities on the neighbourhoods where it gained the most ground. This led to the introduction of regulations that aimed at restricting home sharing in several cities. However, we know relatively little about the true effect of such policies.

In this paper I study how a policy that made new Airbnb hosts liable to a one-time fee upon entry affected local house prices and business outcomes. To my knowledge, this type of policy response to the spread of Airbnb is quite unique, as most regulators reacted by introducing quantity-restricting measures.

The testing ground of this paper is Budapest, Hungary, where local authorities, organized on the district level have a substantial power in regulating local businesses active in tourism/accommodation. Effective as of 2018, a central district which was particularly heavily affected by Airbnb passed a motion which made new Airbnb host liable for an upfront de facto licensing fee, whose exact sum was dependant on the size of the apartment they were planning to put on the short term rental market. The fee affected new entrants only and it curbed the growth of Airbnb supply in treated group. The policy created a geographical discontinuity on the border of this district (District 6) and the neighbouring districts and a discontinuity in terms of size. I use this discontinuity as a source of identification in my paper to measure the effect of this policy on various local outcomes.

My results can be summarized as follows. I find that the policy had a small and statistically borderline significant negative effect on house prices. However, this effect seems to be transitory: while houses in the treatment group cost 2 % less on average in the entire post-treatment period of my sample, zooming on this period yields an estimated effect of -4% for the quarter immediately after the policy change and statistically insignificant results afterwards. The negative effect of prices is stronger for flats which, because of the way the policy was constructed, were liable for a higher licensing fee. This aggregate effect seems to be driven by transaction from the upper end of the price distribution: I detect no effect for the median flat, while the price hit that flats at 75th and 90th percentile take is roughly one and a half times of the average effect.

Next, I turn my attention to the effect of this policy on local businesses. For a more comprehensive view, I use a variety of outcomes from two different sources. First, I collect restaurant reviews from a major review aggregator site, tripadvisor.com, for all bars and restaurants in the treatment and the control group and construct measures of popularity for these establishments. I find that the policy has no effect on the number of reviews bars get. For a different angle, I utilize an administrative dataset to measure employment, profitability and wages payed by local firms active in the accommodation or in the food and beverage sector. I find no effects of the policy on any of these outcomes.

This paper contributes to the small but growing literature on the economic effects of the short-term rental market and the evaluation of policies meant to regulate this market.

Previous studies focused on the connection between the proliferation of short-term rentals and housing market outcomes, such as prices and rents. Horn and Merante (2017) document a correlation between Airbnb penetration and the growh rate of rents in Boston, Barron et al. (2021) use Google Search volume as instrument for short-term rental demand to estimate a causal relationship between Airbnb penetration and prices and rents in several major US metropolitan areas. Garcia-López et al. (2020) perform a similar exercise using shift-share instruments for Barcelona. A key problem these papers have to solve is identification: since the spread of short-term rentals happens endogenously, simply correlating outcomes, such as prices is in general not sufficient to establish causality. Calder-Wang (2020), Hui et al. (2021) and Li and Srinivasan (2019) solve this problem by estimating a structural model of the rental market of New York City and some selected US metro areas, respectively.

Another strand of the literature concentrates on the evaluation of policies restricting short-term rentals. Koster et al. (2019) study the effects of a regulatory change in Los Angeles County on house prices and rents. The setting of paper is the closest to that of Valentin (2021), who, like me, studies a sharp, within-city discontinuity in policy, as opposed to town, municipality or region level implementations of short-term rental regulations. My work, to my knowledge, is the second paper to use this setup.

Apart from replicating their results in a different setting, the contribution of this paper is twofold. First, by introducing a variety of new indicator of local business activity, I am able to show that local businesses are not demonstrably hurt by a policy meant to curb the spread of short-term rentals. Secondly, to my knowledge, I am the first to explore the heterogeneity of such policies by the value of the housing transactions involved. My findings suggest that the policy hurt home sellers (buyers) who happened to own (buy) houses which were considerably more expensive than the median house. This has potentially important distributional consequences, since restrictions of short-term rental supply are often justified as a device to make local housing more affordable: while the policy I am examining did make flats cheaper on average, it predominantly helped buyers of relatively more expensive properties and did nothing to cut sale prices of below-median flats, which may or may not have been the ultimate goal of the restriction.

The rest of the paper is organized as follows. After presenting the data I use in section 2, I begin section 3 by documenting the massive spread of Airbnb in Distict 6 and the neighborhoods surrounding it, which coincided with a rapid appreciation of houses in this area. I present evidence that this relationship is non-causal, since the placement of Airbnb apartments is not random over time.

In section 4, I move on to describe the policy in more detail and present evidence that it was indeed unexpected and effective in reducing the number of new hosts to enter Airbnb. Under a parallel trends assumption, the policy reduced the total stock of Airbnb supply by roughly 9 percent in the treatment group compared to where it would have been in the absence of a policy 12 months after its introduction.

After describing the policy and establishing that it serves as a meaningful "first stage" in my setting, I present my first main sets of results in section 5.1. I estimate DiD regressions that compare the treatment and the control group and I find that apartments in the treatment group cost around 2% less after the introduction of the policy – this result is borderline significant at the 5 % level, depending on the richness of local fixed effects used. To disaggregate this overall effect I look at heterogeneity along three dimensions: time, flat size and price percentiles. I find that the policy decreased the price of larger flats by 6 percent (which is consistent with larger flats being more expensive to license for short-term rental use) only for the first quarter after the introduction of the policy and find no detectable effect afterwards or for smaller flats. Additionally, estimating quantile regressions suggest negative effects for flats at the 75th and 90th percentile in the price distribution – but no effect below that. These last two sets of results (the transitory nature of the price effect and its concentration on the upper end of the price distribution) are novel contributions of this paper.

In section 5.2 I use two other data sources to estimate the impact of the policy on local business. Utilizing balance sheet data, I detect no impact on a rich set of outcomes for business which are primarily active in tourism related industries. Finally, using an online review aggregator platform I detect no drop in the popularity of local bars and restaurants, measured as the number of reviews left in a given month.

2.2 Data

Airbnb listings. I scraped all Airbnb apartments available in Budapest during 2020 November. The site provides information, among others, on the location of the apartment and the year and month the host joined airbnb.com. This allows me to construct a measure of Airbnb penetration that can capture temporal and spatial changes in the number of Airbnb apartments. Since this dataset is basically a cross-section of the apartments which were active when the data were collected, it cannot account for entrants who have already left the platform by then. This problem is exacerbated by the fact that the data were collected during the second wave of the COVID-19 pandemic, which could have led to additional exits from the short-term rental market. Therefore this dataset will be very likely to underestimate the levels of Airbnb penetration, however, under some assumptions that I am going to discuss later, it can be still used to measure the *changes* of Airbnb penetration over time.

House prices. To measure changes in prices, I use the universe of residential real estate transactions between 2008 and 2019, collected by the Hungarian Statistical Office. I omit all those transactions where only a portion of the housing unit was sold. In addition to prices, this dataset contains information on the total floor space of the flat sold and its exact location (both an address and a geocode).

Tripadvisor reviews. Tripadvisor is a platform where users can reviews bar, pubs and restaurants among other establishments of touristic interest. I collect all reviews for all available restaurants, pubs and bars in Budapest for their entire history on the site. This allows me to construct a monthly time series which is a proxy of bar patronage for every establishment in my sample. Based on the language of the review, I can further split this measure to reviews left by ”locals” (reviews written in Hungarian) and by ”tourists” (in any other language).

Balance sheet data. I use the universe of the annual balance sheet statements of all double-bookkeeping Hungarian firms. In my analysis I will focus on firms whose two-digit NACE industry classification is either ”Food and beverage service activities” or ”Accommodation”. I observe the head-quarter of every firm that submits a balance sheet statement, based on which I classify them as treated or non-treated firms.

2.3 Context and descriptive results

Between the the financial crisis in 2008 and the pandemic in 2020, house prices appreciated massively in Budapest: the median price per square meter almost tripled between the trough of 2010q3 and the pre-pandemic peak in 2019q4. As shown in Figure 2.1, this increase coincided with the growth of short term rental services: the stock of Airbnb flats increased on average by 6% in every quarter in the period between 2015 and 2019.

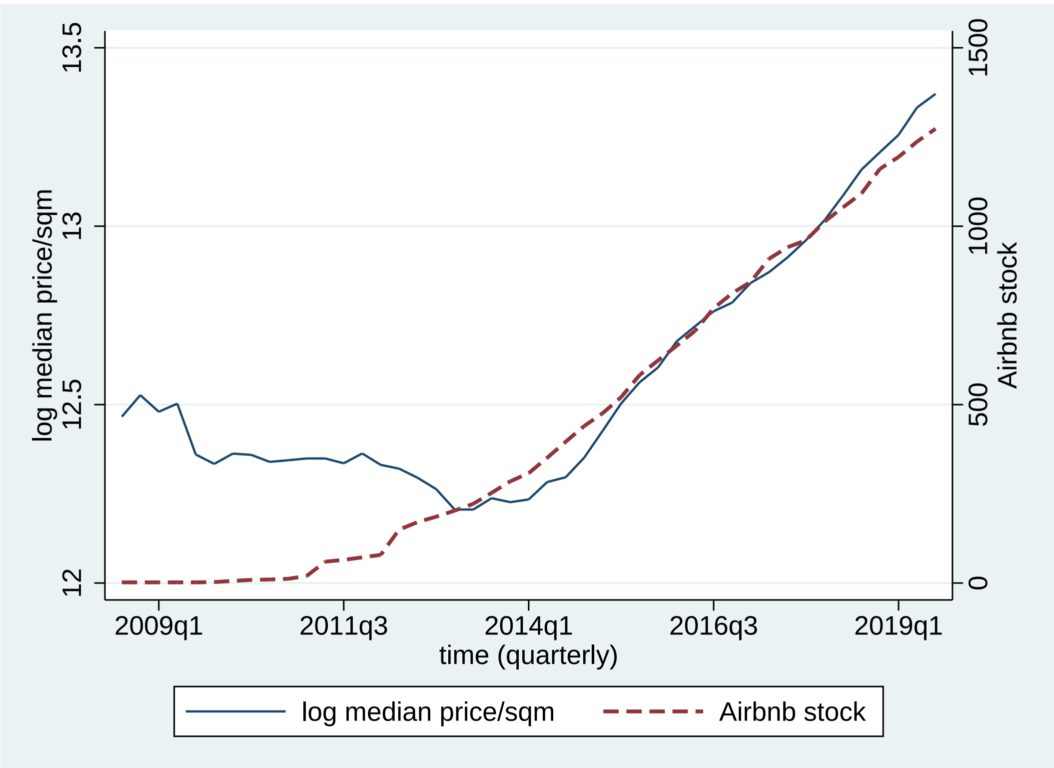

Figure 2.1: The evolution of prices (left axis) and Airbnb supply (right axis) in Budapest, 2009–2019.

To put these figures in context, it would be useful to compare the stock of Airbnb apartments to the total stock of housing. Given that my data almost certainly underestimates the level of Airbnb

10.14754/CEU.2022.03

penetration, I cannot perform these calculations in an exact way. What I do instead is a series of back-of-the-envelope calculations: as I will use the results only to make broad comparisons between cities, some degree of imprecision will be admissible.

In September 2017, Jancsik et al. (2018) identified around 9500 active Airbnb apartments in Budapest, which, given that the 2016 mid-term census found about 908 000 dwellings in Budapest (HCSO, 2016), means that about 1.05% of the housing stock was occupied by Airbnb. In 2015, this measure was about 0.86% in Los Angeles, 1.31% in New York, 2.06% in Barcelona and 2.56% in Paris (Garcia-López et al., 2020), so in terms of Airbnb penetration, while Budapest was not among the most affected cities worldwide, it was comparable to some popular metropolitan areas.

These apartments are typically concentrated in some well-defined neighbourhoods, and the size of the housing stock in the outskirts can potentially distort city-wide Airbnb shares, so, as a next step, I make a rough estimate of Airbnb penetration in District 6, the area which is the focus of the policy evaluation exercise of this paper and was presumably heavily affected by the spread of home sharing. Under the somewhat restrictive assumption that the exit rate of hosts during the pandemic in 2020 is independent of the entry date and location, around 4.3% of the flats in District 6 were available on Airbnb in 2018. [1] The only comparable, neighbourhood-level figures I am aware of are for New York by Calder-Wang (2020), whose findings suggest that the Airbnb penetration of District 6 in Budapest is about two-thirds of that in the most affected neighbourhood in Manhattan. [2]

I zoom in on the previous graph by taking a finer geographical resolution and instead of the entire city I look at the zip code level by estimating the following regression:

$$log(price/m^2)_{izt} = \beta_0 + \beta_1 \log(\text{Airbnb stock})_{zt} + \gamma x_{izt} + a_z + a_t + u_{izt}, \tag{2.1}$$

where i indexes a transaction, z a zip code and t a time period (quarter). The coefficient of interest is β_1, the estimated elasticity of prices with respect to Airbnb supply. As shown in Table 2.1, the connection between these two variables are strongly significant. If I restrict my sample to zip codes with a positive number of Airbnb flats, 1% more Airbnb apartments are associated with 3% higher prices. However, many zip code-quarter observations in the dataset do not have Airbnb units: either because they represent early time periods when Airbnb was not wide-spread, or because they represent zip codes which are not attractive for tourists. Using the log of 1 plus the Airbnb stock on the right-hand side allows me to include these observations, however it introduces some form of attenuation bias: adding a fixed number to the regression would bias the estimated coefficient towards 0. What we see in column 2 instead is a somewhat higher and strongly significant coefficient, suggesting strong selection effects.

Focusing on this, the relationship between contemporaneous Airbnb supply and prices might be en-

[1] For this calculation I take the share of District 6 apartments in my dataset in 2020, multiply this share by 9500, the number of Airbnb flats in 2017 and finally by the the growth rate of the city-wide Airbnb stock between 2017 and 2018. The validity of this calculation clearly depends on the assumptions about the exit rate being exogenous. If hosts outside of District 6 left the platform in relatively higher numbers, this figure would overestimate the true the value. In turn, if hosts who entered late (after 2017) were more likely to leave, I am underestimating Airbnb penetration in Disrict 6.

[2] Airbnb penetration in New York is quite heterogeneous by boroughs: the neighbourhood with the highest share of Airbnb apartments in Brooklyn is Greenpoint & Williamsburg (9.3%), in Manhattan, it is Chelsea, Clinton & Midtown (6.9%), while in the Bronx and in Queens, its 0.4% and 2.8%, respectively.

Dependent variable:	(1)	(2)	(3)
log(price/m^2)	zip codes with Airbnb> 0	all zip codes	zip codes with Airbnb> 0
log(contemp. Airbnb count)	0.0303*		
	(0.0136)		
log(1+contemp. Airbnb count)		0.0342***	0.0572**
		(0.00618)	(0.0176)
Observations	112081	200514	112081
R^2	0.642	0.645	0.642

Standard errors in parentheses. SE's clustered on the zip code level.

Data from 2010-2019. Outcome: log(price/m2). Includes log(size), plus zip and quarter dummies.

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

Table 2.1: Correlation between prices and contemporaneous Airbnb penetration

dogenous: neighbourhoods with more Airbnb units might be more desirable in ways which are not captured by zip code fixed effects, and these omitted variables might be partially responsible for higher prices. One way of testing for this is to check the relationship between prices *before* the entry of Airbnb and Airbnb stock in a much later period. If Airbnb flats were randomly placed (or they were random conditional on the included controls), there would be no connection between these variables. For this, I estimate the following equation:

$$log(price/m^2)_{izt} = \beta_0 + \beta_1 \log(\text{Airbnb stock in 2019})_z + \gamma x_{izt} + a_t + u_{izt}. \tag{2.2}$$

This relationship is identified from cross-sectional variation: I compare similar flats in the same quarter from the period 2010-2019 and see if flats sold in zip codes that end up having more Airbnb flats in 2019 are on average more expensive than flats in zip codes that end up with fewer Airbnb flats in 2019.

Table 2.2 shows that "early" prices and "late" Airbnb stock are indeed positively correlated. Focusing on zip codes that have at least one Airbnb unit in 2019 yield an estimated elasticity of about 0.22: 1 percent more Airbnb in 2019 corresponds to about 22% higher prices *before* the entry of Airbnb to the market.

Since my main control on the right hand side (Airbnb stock in 2019) is fixed for a zip code, I cannot include controls which are fixed at the zip code level, so this elasticity estimate has no causal interpretation.

10.14754/CEU.2022.03

Dependent variable:	(1)	(2)	(3)
log(price/m^2)	zip codes with Airbnb> 0	all zip codes	zip codes with Airbnb> 0
log(Airbnb count in '19)	0.221***		
	(0.000618)		
log(1+Airbnb count in '19)		0.128***	0.242***
		(0.00143)	(0.000678)
Observations	13479	16283	13479
R^2	0.296	0.332	0.296

Standard errors in parentheses. SE's clustered on the zip code level.

Data from 2010. Outcome: log(price/m2). Includes log(size), plus zip and quarter dummies.

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

Table 2.2: Correlation between prices and Airbnb stock in 2019

10.14754/CEU.2022.03

2.4 Empirical strategy

In this section I provide a detailed description of a policy that was meant to restrict the supply of short-term rental and argue that it was successful in reducing the number of new entrants to Airbnb as a platform. After that, I present my design that I use to estimate the effect of this policy on a number of economic outcomes.

2.4.1 Description of the policy

At the end of 2017, the local government of District 6 in Budapest decided to make entry to Airbnb more expensive by introducing an extra fee that new hosts were required to pay when applying for a license to run their own Airbnbs. Nominally, the policy tied the fee to the availability of parking lots: new hosts were required to own a specified number of parking lots per rooms that their guests could use. If hosts had multiple-room units or they wished to get the license for more than one apartment, they were required to be able to produce a parking lot for every room in every apartment. If they did not own the required amount of parking space (or chose not to use it for this purpose), they could opt for paying a one-time fee of 1.2 or 1.5 million HUF, depending on location. [3] Since in District 6 parking space is scarce and generally more expensive than the fee, the policy worked as a de facto licensing fee for the overwhelming majority of the new entrants. To put things in context, the average price of a square meter in the last quarter before the implementation of the policy was about 538 000 HUF, so the licensing fee amounts to the price of around 2.5-3 square meters worth of housing on average.

In my policy evaluation exercise, I will focus on District 6 (as the treatment group). As a control group I pick the neighboring districts 5 and 7, plus those zip codes of district 13 which are located reasonably close to the treatment group. (for a map of these areas, see figures B.2.2 and B.2.5). The treatment and the control groups form a near contiguous mass of flat, densely populated area in the downtown of Budapest: they are not separated by natural boundaries (such as rivers, lakes, hills or valleys) or by man-made structures (like railroad tracks or busy highways) and there are usually no strong and immediate visual clues (such as a change in the age or style of the buildings) when crossing district boundaries. The area I picked for this exercise happens to represent a large chunk of the Airbnb market in Budapest: every year about 70-75% of the total Airbnb stock was located in either the treatment or in the control group (Figure B.1.1). In the appendix, I describe the factors that I weighted when deciding about the exact delineation of control group and I provide robustness checks for some reasonable modifications of the control group.

There are strong reasons to assume that the policy was unexpected. First of all, District 6 was the first to impose such a stringent regulation on Airbnb apartments in Budapest. Secondly, very little preparation time was allowed for prospective new hosts: the motion itself was passed on 21 December 2017 by the district council and it went into effect on 1 January 2018, so even if hosts who had been thinking of converting their apartment into short-term rentals had very few business days to submit their

[3] More central locations carried a higher fee. Given the small difference between these two figures, I will not use this source of heterogeneity anywhere in my paper and I will treat the entirety of District 6 treated in the same way. 1.5 million HUF was worth around 5600 USD in 2018.

application under the earlier, more liberal regulation. Finally, searching local and national media did not turn up any mentions of this topic before the policy took effect.

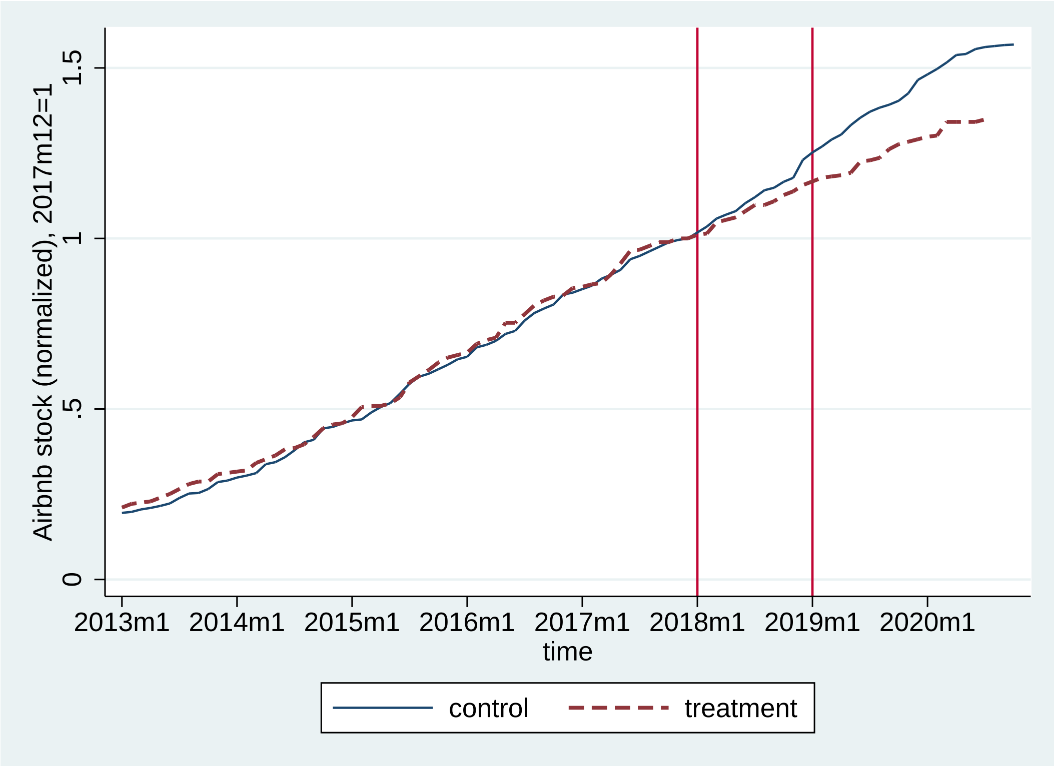

Figure 2.2: The evolution of Airbnb stock (reconstructed from the history of entrants) over time in the treatment and the control group, normalized by their respective mean in the last period before the treatment.

Next, I present some graphical evidence, which supports my claim of the policy having been unexpected and demonstrate that it was effective in reducing the mass of new entrants to the short-term rental market. For every month between 2016 and 2020, figure 2.2 depicts the number of hosts who have already entered by the given month, separately for the treatment group (district 6) and the control group (neighbouring zip codes). Since the control group is much larger, I normalize both times series with their respective values in December 2017, the last month before the policy change.

The evolution of the two normalized graphs is practically indistinguishable from each other pre-treatment. Post-treatment, however, the growth in the treatment group is visibly below that of the control group. Fitting linear trends separately for the four periods (treatment/control $\times$ before/after) confirms that the widening gap between District 6 and the control group is indeed statistically significant. (Table B.1.1).

A very similar policy was introduced in the neighboring district 7 about a year later (effective as of 6 January 2019). New hosts were require to pay a similar licensing fee as in district 6, and the fee was also nominally tied to the lack of parking spaces. This amounted to 2 million forints per 50 square meters of the flat being put on the short term rental market. This policy should theoretically allow for a design with sharper discontinuity, as it uses the total floor area as a basis of the fee, instead of the number of rooms a

10.14754/CEU.2022.03

flat has. This has two potential advantages. Firstly, since this data is available in an independent official housing register, owners are less likely to under-report this number, should they decide to obfuscate the total floor space in order to pay a smaller fee. Secondly, unlike room numbers, I do observe this variable in my dataset, so it would much easier to explore the heterogeneity in the effect of the policy. However, there is some evidence that the policy was to some extent expected by future hosts. Figure 2.3 (which is a reproduction of figure 2.2 for district 7, with district 6 dropped for the control group, as it had been already treated in a slightly different way) suggests an excess mass of new entrants before the policy took effect in January 2019. Therefore in my empirical analysis I omit my post-2019 observations.

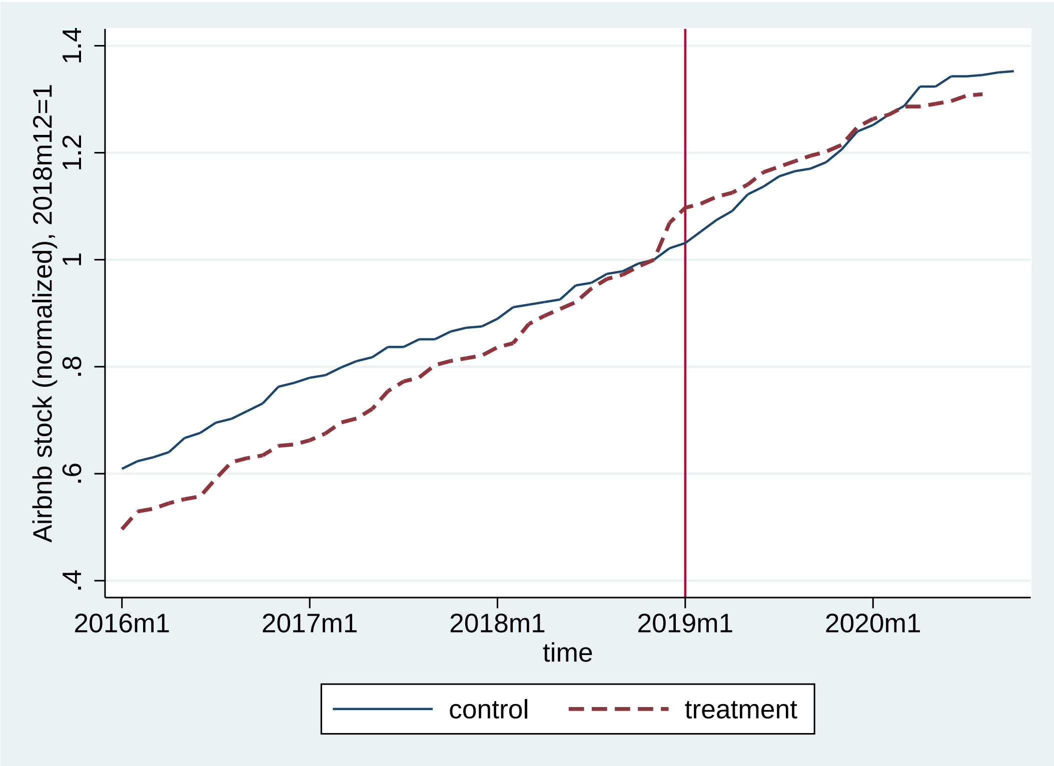

Figure 2.3: The evolution of Airbnb stock (reconstructed from the history of entrants) over time in the treatment (District 7) and the control group, normalized by their respective mean in the last period before the treatment.

These results suggest that the policy seems to have been successful in reducing the growth in the stock of apartments advertised on airbnb.com. This is likely to hold, even though my data on Airbnb flats is clearly is not perfect. As I have discussed before, the data I use is only a proxy of the total Airbnb supply at a given moment and this dataset underestimates the true level of Airbnb penetration. Firstly, even after joining Airbnb, an apartment may have been available only for a fraction of the month or could have been off the market for several months. Since I do not observe the full history of every flat, my data cannot account for these temporary exits/re-entries.

Secondly, as I obtained the data in the fall of 2020, presumably lot of host have left the market due to collapse of tourism brought about by the pandemic. Given the difference-in-difference framework that I have used in this section, this attrition becomes problematic only when it is selective. In this specific case,

10.14754/CEU.2022.03

spurious results could emerge if, out of those host who entered the market after 2018, disproportionally more from District 6 would have left the platform due to the pandemic than post-2018 entrants from the neighbouring districts. I have no strong reasons to assume this, however, I have no way to entirely refute this concern either. Therefore for the results in this section to have causal interpretations, I have to rely on the identification assumption that exit rates were not higher among the post-2018 entrants in the treatment group than in the control group between the end of my sample (end of 2018) and the time of data collection (fall of 2020).

Since in this paper my focus is the effect of the policy on outcome is *not* the total stock of Airbnb apartments, in the following I will move on to an empirical setting where I do not have to deal with these potential measurement issues.

2.4.2 Empirical design

I will use a simple difference-in-difference to framework to estimate the causal impact of the policy. My main specification is

$$y_{it} = \beta_0 + \beta_1'(\textbf{after 2018q1})_{it} + \beta_2'(\textbf{treat x after 2018q1})_{it} + \gamma x_{it} + a_i + u_{it}.$$

Given that the later policy change in District 7 was probably not entirely unexpected, I restrict my sample to 2016–2019. I allow the effect to be time-dependent by including a vector of time dummies and their interaction with the treatment variable.

When a_i's have no natural interpretation as a dummy for a cross-sectional observation (as in the case for housing transactions) I define them to be location-specific to capture the effect of location-specific amenities which are fixed over time. I use two alternative sets of fixed effects: a cruder one where I include a dummy for every zip code, and a finer one, where I overlay a rectangular grid of 250 meter by 250 meter squares on the map and I use the grid where the transaction falls as the location fixed effect. [4] Zip codes have the benefit of being much larger than the grid cells, giving a possibly less noisy measure of the desirability of the neighbourhood. Plus, they are fixed by the postal service, so the definition is less amenable to researcher discretion. However, as it is apparent from the maps (figures B.2.5 and B.2.2), zip codes do not represent equally-sized contiguous, convex areas but have jagged borders and can differ in size greatly. Square grids have none of these shortcomings and can potentially capture fixed neighbourhood characteristics better, giving more power to my tests at the cost of estimating more fixed effects and losing more degrees of freedom.

One possible identification problem is policy-induced demand overflow to neighboring districts, which are included in the control group. If the licensing fees makes an outcome (such as buying property or running a business) less profitable or desirable in the treatment group, those discouraged by the treatment might decide to switch to neighbouring areas, some of which can possibly included in the control group. This would magnify any effect that the policy had on the treated group.

While this may look like a potentially serious threat to identification, there are two reasons why I do not consider this identification threat detrimental to my analysis. First, since most of my results

[4] The results are robust to the choice of the size of the grid cell across a range of plausible values.

will suggest that the policy had no measurable effect on economic outcomes, a mechanism that would inflate the estimated impact of the policy would not threaten the validity of the null results. Secondly, even if we might have good reasons to be suspicious about the exact size and magnitude of negative estimated treatment effects due to the mechanism outlined above, a possible substitution away from the treatment group towards the control suggests the existence of a non-zero policy effect. If the introduction of the fee did not increase demand for flats in the aggregate (that is, in the treatment and control group combined) and it did not attract extra demand to the control group from outside (that is, not from the treatment group), then an increase in demand in the control group must have come to some degree from the treatment group – which is exactly the effect of the policy. So while depending on the exact strength of this channel I may not be able fully identify the size of the effect, I can establish the sign and more importantly the existence of the effect of the policy.

2.5 Results

2.5.1 Price

Since the policy decreases the potential profit from buying and renting out an apartment on Airbnb in District 6, one would expect that the prices would drop after the introduction of the fee.

As the outcome y I use the log of the price per square meter of flat sold and I include the log of total floor space and geographical dummies as controls. First as baseline, I look at the results of a simple before-after comparison of the treatment and the control groups, with varying levels of geographical granularity. As Table 2.3 shows, the results broadly similar, though: the coefficient of the interaction term is around -2 percent, and it is borderline significant at the 5% level, if I use the cruder, zip code level controls. If I include the richer, but potentially noisier set of controls instead, the coefficient of interest decrease in magnitude very slightly and is no more significantly different from zero at the 5% level.

Next, I examine if the effect changes with time. As the policy affected directly buy-to-let buyers only (that is, people who were planning to put their properties on Airbnb after the purchase) who were possibly better informed than other buyers, it could have taken time for less informed buyers to react to drop in demand. Therefore I expect that the price effect of the policy was stronger right after the implementation of the licensing fee and it might have diminished with time. To test for this, I reestimate the the basic specification (equation 2.4.2) by using quarterly dummies (the finest time resolution available in my dataset) instead of a binary before-after variable.

At the same time, I explore heterogeneity along an another dimension. I split the sample between ”big” ($> 50m^2$) and ”small” ($\leq 50m^2$) flats, since I expect that the policy had a bigger effect on larger flats, as the fee was increasing in the number of rooms. Unfortunately, I do not observe the number of rooms directly, so I use size as an imperfect way to classify transactions into flats with two or more rooms (flats larger than $50m^2$) and flats with less than two rooms (every other flat)

Table B.1.2 summarizes the results. No separate interaction dummy is significant after the introduction of the policy in the full sample, even though all of them have the expected negative sign, suggesting

Dependent variable:	(1)	(2)
log(price/m^2)	zip	grid
treated	-0.3170***	0.0583*
	(0.0804)	(0.0287)
	[0.0001]	[0.0445]
after 2018q1	0.2819***	0.2832***
	(0.0074)	(0.0077)
	[0.0000]	[0.0000]
treated x after 2018q1	-0.0251	-0.0212
	(0.0130)	(0.0124)
	[0.0530]	[0.0889]
r2	0.369	0.406
N	15606	15606

Includes location dummies and log(size). SE's clustered on the street x zip and grid level, respectively.

Standard errors in round, p values in square brackets.

Table 2.3: Results from a simple diff-in-diff setting (Equation 2.4.2).

.

a lack of detectable effect which might be due to low power. Focusing on larger flats, I find a significant negative coefficient just after the treatment, which is almost three times greater in absolute value than the point estimate in table 2.3 for the post-treatment period, even though it diminishes after the first quarter. For small flats, I find no significant post-treatment effect, plus no point estimates are negative, which suggests that the average estimated effect I saw in the simpler, before-after setting is mainly driven by transactions (1) shortly after the policy change, and (2) flats which are larger, hence, were, through higher fees, were more heavily affected.

Finally, I explore the heterogeneity of the treatment by price. This is particularly relevant from an applied policy perspective, since one common reason given for restricting short-term rental services is that by driving up prices, they make housing less affordable for locals. I estimate the base specification (with before-after dummies) using quantile regression for the 10th, 25th, 50th, 75h and the 90th quantiles of the price distribution. Tables 2.4 and 2.5 show the results for zip codes and grid fixed effects, respectively. They paint a qualitatively similar picture: the policy has made expensive flats (at the 90th percentile) somewhat cheaper: the per-square-meter price fell by around 3.6 percentage points. While this effect is significant in the specification with zip code fixed effects, it loses its significance using the finer set of grid fixed effects instead of zip codes.

10.14754/CEU.2022.03

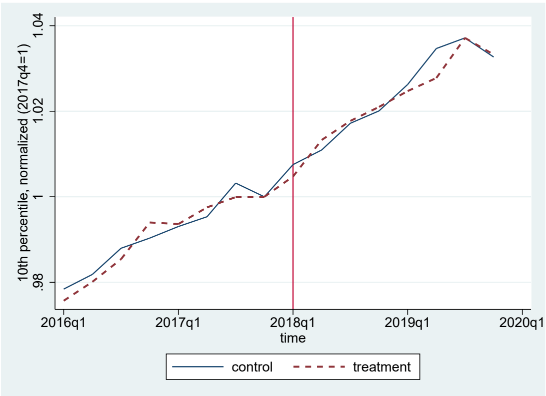

(a) 10th percentile

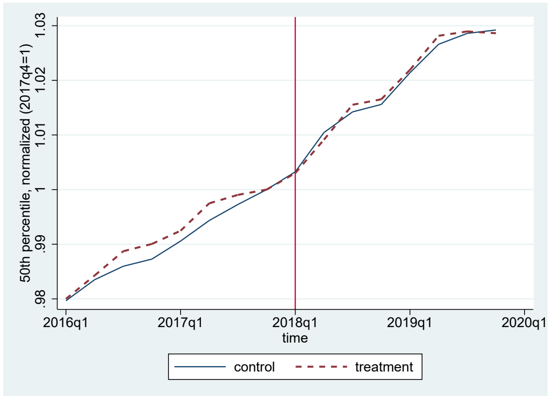

(b) median

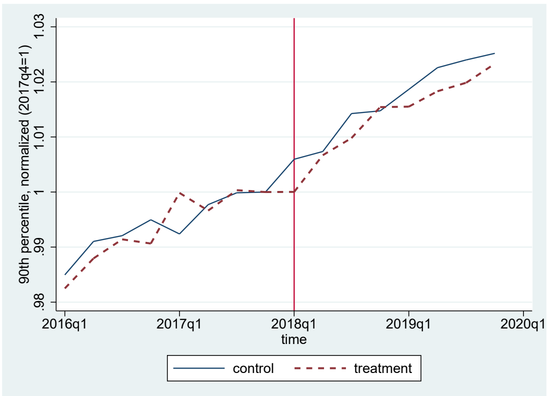

(c) 90th percentile

Figure 2.4: Evolution of log price per square meter quantiles in the treatment and the control group.

Dependent variable:	(1)	(2)	(3)	(4)	(5)
log(price/m^2)	10	25	50	75	90
treated	-0.4612***	-0.4606***	-0.4116***	-0.4110***	-0.3015*
	(0.0529)	(0.0621)	(0.0888)	(0.0975)	(0.1521)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0476]
after 2018q1	0.8267***	0.7548***	0.6031***	0.4556***	0.3849***
	(0.0131)	(0.0113)	(0.0103)	(0.0094)	(0.0100)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	0.0454	0.0223	-0.0266	-0.0262	-0.0376*
	(0.0248)	(0.0167)	(0.0176)	(0.0149)	(0.0162)
	[0.0676]	[0.1804]	[0.1318]	[0.0796]	[0.0203]
r2	0.298	0.307	0.311	0.298	0.282
N	30252	30252	30252	30252	30252

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table 2.4: Quantile regression results, with zip dummies. Standard errors in round, p values in squared brackets.

Dependent variable:	(1)	(2)	(3)	(4)	(5)
log(price/m^2)	10	25	50	75	90
treated	0.1078*	0.1297	0.0865	0.0664	0.0781*
	(0.0525)	(0.0874)	(0.0489)	(0.0358)	(0.0381)
	[0.0401]	[0.1376]	[0.0769]	[0.0638]	[0.0405]
after 2018q1	0.8220***	0.7665***	0.6062***	0.4505***	0.3780***
	(0.0163)	(0.0123)	(0.0104)	(0.0079)	(0.0080)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	0.0367	0.0011	-0.0263	-0.0173	-0.0107
	(0.0266)	(0.0168)	(0.0150)	(0.0127)	(0.0161)
	[0.1686]	[0.9494]	[0.0785]	[0.1724]	[0.5074]
r2	0.312	0.322	0.327	0.312	0.293
N	30252	30252	30252	30252	30252

Includes grid dummies and log(size). SE's clustered on the grid level.

Table 2.5: Quantile regression results, with grid dummies. Standard errors in round, p values in squared brackets.

2.5.2 Local businesses: evidence from online reviews

In this section I will turn to using an alternative data source to quantify the impact of the policy on a subset of local business. I use the number of restaurant reviews from tripadvisor.com as a proxy for restaurant patronage and estimate a simple difference-in-differences model:

$$y_{it} = \beta_0 + \beta_1(\text{after 2018q1})_{it} + \beta_3(\text{T x after 2018q1})_{it} + a_i + u_{it}, \tag{2.3}$$

where y_it is the number of reviews a restaurant receives in a given month. [5]

First I focus on the number of non-Hungarian reviews, as they are are much more likely to be by tourists and tourists are probably more affected by any policy which restricts Airbnb supply.

Dep. variable: number of reviews	(1)	(2)	(3)
after 2018q1	0.827**	0.826**	0.825**
	(0.294)	(0.294)	(0.294)
treated x after 2018q1	0.179	0.178	0.178
	(0.466)	(0.466)	(0.466)
Observations	41801	41801	41014
R^2	0.006	0.010	0.010
sample	all	all	before 2017
time	monthly	quarterly	monthly

Standard errors in parentheses

With restaurant fixed effects. SE's clustered on the restaurant level.

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

Table 2.6: Diff-in-diff results: monthly number of reviews.

Table 2.6 suggests however, that the policy did not have any effect on bar patronage proxied by the number of reviews guest leave and tripadvisor.com. The interaction coefficient is insignificant no matter if I include monthly or quarterly time dummies (columns 1 and 2) or if I drop newly opened establishments from the sample, who might be initially less known, hence the (low) number of reviews they get every month might be a noisier measure of their true popularity (column 3). I repeat this exercise with aggregating the data on the quarterly, instead of the monthly level, with very similar results.

I explore the potential heterogeneity of the treatment by price ranges. The site classifies all establishments into three price bands, "inexpensive" (denoted by "$"), "mid-range" ("$$ - $$$") and "high-end" ("$$$$"). I estimate the same model as in table 2.6 separately for these categories, however, my results remain basically similar. (see tables B.1.8, B.1.9 and B.1.10 in the appendix).

The seemingly stable number of reviews between the treatment and the control group may reflect compositional changes: as district 6 is presumably less frequented by tourists after the treatment, local

[5] I am not the first to use reviews in economics: Dai and Luca (2020) use Yelp reviews to predict hygiene scores of restaurant in San Francisco.

guests might have taken their place. For this, I plot the number of reviews left by locals (identified as having been written in Hungarian) as a share of the total number of reviews that had been submitted on Tripadvisor up to a given month. Figure 2.5 shows the results. Restaurants in district 6, the treatment group had on average relatively more foreign guests (leading to a lower local-to-foreign ratio) than the control group in the before period, while this gap seems to have narrowed after the introduction of the new licensing fee policy at the beginning of 2018. The closing of this gap is statistically significant as shown by the estimates of table B.1.11, however it can be argued that this change had started before the policy change and potentially unrelated to it.

2.5.3 Local businesses: evidence from balance sheet data

It is possible that the policy, in addition to its primary target of local Airbnb hosts, might have had secondary effects on local businesses active in sectors which are the most exposed to tourist demand. In this section, I use a firm-level panel of all Hungarian firms liable to double-entry bookkeeping to detect any possible changes in key measures of firms success. I restrict my sample to businesses which were classified as active either in the ”Accommodation” or in the ”Food and beverage services” 2-digit NACE category. I follow the definition of treatment and control groups from the previous section. Table 2.7 shows the estimation results of a simple diff-in-diff model for various outcomes. The coefficient of interest, treatment x 2018 is never significantly different from zero for either of the outcomes.

It has to be noted however that the lack of significant results might be due data quality issues: since the observations are not establishment, but rather firms, classification into the treatment and the control groups are based on the address on the headquarter of the firm and not on the exact location of the establishment. Multi-establishment firms are also lumped into one observations. While this is unlikely to cause systematic distortion, it can contribute to measurement error, leading to large standard errors.

	(1)	(2)	(3)	(4)	(5)
Dependant variable:	log(persexp)	log(sales '18)	log(pretax profit)	log(pers. exp./emp.)	log(employment)
year 2014	0.057	0.059	0.202**	0.071**	-0.030
	(0.04)	(0.04)	(0.08)	(0.02)	(0.02)
year 2015	0.111**	0.158***	0.389***	0.104***	0.059*
	(0.04)	(0.04)	(0.09)	(0.03)	(0.03)
year 2016	0.306***	0.288***	0.548***	0.221***	0.081**
	(0.04)	(0.05)	(0.09)	(0.03)	(0.03)
year 2017	0.555***	0.469***	0.950***	0.369***	0.111***
	(0.04)	(0.05)	(0.10)	(0.03)	(0.03)
year 2018	0.838***	0.773***	1.556***	0.507***	0.168***
	(0.05)	(0.05)	(0.10)	(0.03)	(0.03)
treat x 2014	0.009	0.108	0.093	-0.086	0.124**
	(0.07)	(0.08)	(0.16)	(0.05)	(0.05)
treat x 2015	0.026	0.089	-0.054	-0.051	0.051
	(0.08)	(0.09)	(0.18)	(0.06)	(0.06)
treat x 2016	0.070	0.121	0.024	0.001	0.070
	(0.08)	(0.10)	(0.18)	(0.06)	(0.06)
treat x 2017	0.044	0.138	0.009	0.013	0.089
	(0.09)	(0.10)	(0.18)	(0.06)	(0.07)
treat x 2018	0.030	0.085	0.048	0.007	0.108
	(0.10)	(0.11)	(0.19)	(0.06)	(0.07)
Constant	8.384***	9.916***	6.971***	7.017***	1.551***
	(0.03)	(0.03)	(0.06)	(0.02)	(0.02)
N	8070	8222	4696	7462	7673
R^2	0.129	0.096	0.174	0.110	0.020

Table 2.7: Accommodation and Food and beverage services combined. Outcomes are logs of (1) total personal expenditure, (2) total sales, in 2018 prices, (3) pretax profit, (4) personal expenditure per employment, (5) number of workers employed.

2.6 Conclusion

I have found that the policy I study did manage to somewhat reduce new entry to the short-term rental market while having only marginal effect on some segments on the housing market and did not inflict any detectable negative impact on local business active in tourism-intensive sectors. Based on this alone,

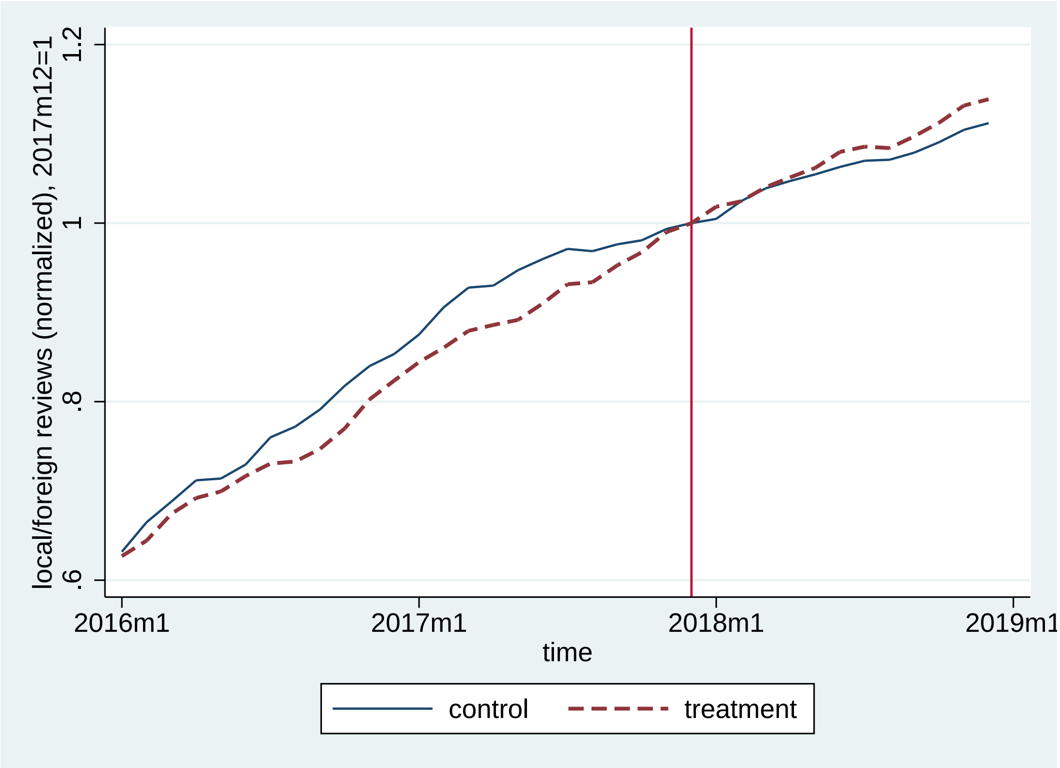

Figure 2.5: Cumulative local/foreign reviews ratio, normalized.

it is probable that the one-off tax did not introduce significant distortions and in this sense it can be considered a good form of taxation.

Whether the policy has achieved its goal is not entirely clear however – mainly because, as discussed in the Identification section, it was never clear what the intentions of the policy maker were. While this policy can prove to be an economically efficient way to raise revenues with the potential added benefit of being politically popular, it might or might not have been effective in e.g. reducing negative externalities inflicted by short-term guest on local communities, perceived or real.

10.14754/CEU.2022.03

10.14754/CEU.2022.03

Chapter 3

Media-mediated cross-party effects of political beliefs

3.1 Introduction

In this paper I will argue that the existence and behavior of a profit maximising news media can establish a link between the beliefs and welfare of politically dissimilar groups. I build a simple model with politically heterogeneous news consumers who, in addition to valuing precise information, enjoy reporting that slants towards their prior beliefs. I find that if the taste for slant is strong enough, as one group (e.g. liberals) move away from the center toward the extremes, the beliefs of the other group (e.g. conservatives) becomes less correct on average. In contrast, if neither side enjoys like-minded reporting too much, a moderate shift of liberals to the left might improve the welfare of conservatives.

My proposed mechanism works through two channels. The more conventional one is the pricing channel that affects consumer welfare: given their different priors, consumers have different valuation for news and a profit maximizing media organization might lower prices to gain market share – which will leave some surplus at high-valuation consumers. For example, if liberals are very liberal and conservatives are slightly off-center, a media monopolist under some conditions will find it optimal to charge low prices and serve both groups with high-quality (truthful) reporting. In my model, less biased voters value truthful reporting more, so in this setting liberals being far from the center increases the welfare of conservative consumers: they get high-quality news for less. In other words, if liberal beliefs moved a bit more toward to the left, conservatives would benefit from it in the form of lower prices and constantly high news quality.

The second, potentially more novel channel works through consumer beliefs and its key insight is the following. As the media tries to cater to one political group by slanting progressively more towards them, it loses some of its credibility in the eyes of the other group, which makes the other group less responsive to the reporting of the media. This makes them put more weight to their own, incorrect prior. For example, if U.S Democrats moved more to the left, the media serving them (e.g the New York Times) would follow suit by increasing its liberal slant. Republican readers would be hurt by that because, while they would still buy the NYT for a while, they would correctly understand that the paper is less informative about the world and they would trust it less, even if it happens to be correct. This would make them always more conservative than they would be if liberals were more centrist. In my model, this would lead to Republicans holding beliefs that turn out to be incorrect more often.[1] To my knowledge, this paper is the first to describe this mechanism formally.

My paper connects to the wider field of the (political) economy of media markets (for reviews, see Gentzkow et al. (2015) or Strömberg (2015)). I explicitly concentrate on how the political slant of the news media affects political beliefs (for empirical contributions see Durante and Knight (2012) and DellaVigna and Kaplan (2007) or for a review Prior (2013)), but I assume, that the media only cares about maximizing profit (and not e.g. getting a candidate elected) and it gets its revenue solely from its subscribers (and not from e.g. government bribes as in Besley and Prat (2006)). From a methodological standpoint my contribution has a lot in common with papers that model news consumers as having explicit preferences for political slant (e.g. as in e.g. Mullainathan and Shleifer (2005) or Bernhardt et al. (2008)), as opposed to papers who assume that news consumers are rational (Gentzkow and Shapiro, 2006). Finally, as my model makes predictions on the evolution of political beliefs of voters of different political views after they are exposed to media reporting, my work naturally connects to the recent literature on the effect of media consumption on political polarization in economics and political science (for a recent review, see Zhuravskaya et al. (2020)).

The rest of the paper is organized as follows. In section 2. I start with the basic setup of the model. Next, I lay out the the most important modelling assumptions, their implications and how they relate to the assumptions and results in the literature. Section 3 contains the main theoretical results of the paper, while section 4 concludes.

3.2 Model

3.2.1 Setup

There is one perfectly informed firm (the news media) who sells news to consumers on the state of the world ω (Left or Right) for price P.[2] While both states of the world are equally likely, consumers are heterogeneous in their priors: the prior about the probability of $\omega = R$ is p_R for right-biased consumers and p_L for the left-biased consumers, with $p_R \geq 1/2 \geq p_L$. Consumers are Bayesian. Both groups have a mass of 0.5. To rule out some knife-edge cases and simplify the discussion of my results, I assume that one group is always strictly farther from the objective probabilities than the other: $p_R \neq 1 - p_L$.

I use the notion of "distortion bias": a left-biased media will report news $n = l$ in state R with probability b: $P(n = l|\omega = R) = b \geq 0$ and $P(n = l|\omega = L) = 1$. For such a reporting strategy, a reporting that says "right" is fully revealing of the underlying state, but a "left" reporting is somewhat informative too.

[1] For this I assume that reality is "centrist in expectation": conservative and liberals views have an equal probability of being true in a given issue.

[2] For my modelling framework I borrow heavily from Gentzkow et al. (2015)

Consumers get utility 1 from matching the state correctly and 0 otherwise (instrumental utility) and listening to news which confirms their prior (behavioral utility). They also pay the nonnegative price P that the media sets if they decide to purchase news.

The timeline is simple: the media announces its reporting strategy b and price P, ω is realized (observed by the media but not by the consumers), consumers decide to buy news or not and finally payoffs are realized.

The expected instrumental utility of a leftist consumer is

$$u(b, p_L) = (1 - p_L) \cdot 1 + p_L(b \cdot 0 + (1 - b) \cdot 1). \tag{3.1}$$

The first term is the payoff if $\omega = L$: in this case the media reports left and a leftist consumer will guess (correctly) L, making 1. The second term is the expected payoff when $\omega = R$. With probability $1 - b$ the media tells the truth by sending a signal r, which the consumer follows, resulting in the payoff of 1. With probability b the media reports distortedly l, which the consumer follows and makes zero, since the truth was R. (For this, note that a left-biased media always goes left if the media reports left: even if the media is totally uninformative, the consumer still thinks that the left state is more likely).

The action of the right-biased consumer depends on the reporting strategy of the media and the consumer's bias. If the bias of the media is low enough and/or the right-leaning consumer's prior is not too far away from the truth of 1/2, then it is possible that $Pr(\omega = L|n = l; p_R) > \frac{1}{2}$, even rightist consumers trust the media in the sense that after it reports left they follow suit (equivalently, $b < \frac{1}{p_R} - 1$). Since in this case the actions of the rightist and leftist consumers will be the same, the rightist consumer will have a very similar instrumental utility $(1 - p_R) + p_R(1 - b)$. If the previous condition does not hold, then the rightist consumers will not hold the reporting informative as it does not change their actions – therefore in equilibrium they won't buy any of it.

The second source of consumer utility is the enjoyment that consumers derive from listening to news that slants towards their prior. For a left-leaning reporting strategy $b \geq 0$ and for some prior p, I define this as

$$\phi(b; p) = -\alpha \cdot b \cdot \left(p - \frac{1}{2}\right), \quad (\alpha > 0) \tag{3.2}$$

The value of this function is positive for leftist consumers (as their prior about the probability that $\omega = R$ is less than one-half) and negative for rightist: the sign of ϕ is determined by the term $\left(p - \frac{1}{2}\right)$. The absolute value of ϕ also increases as p gets farther away from one-half: the more biased the consumer is the more (dis)utility he derives from the slant in itself (for a hypothetical unbiased consumer $\phi(b; \frac{1}{2}) = 0$). Finally, ϕ is increasing in b: the more biased the media is, the more (dis)utility the consumer gets from listening to it.

Consumers can choose not to purchase any news, in which case they will follow their priors when they guess ω, leading to instrumental (and total) utility $1 - p_L$ for the leftist and p_R for the rightist consumer.

A leftist consumer buys the news for P if

$$u(b, p_L) - \phi(b; p_L) - P > 1 - p_L, \tag{3.3}$$

resulting in

$$WTP_L = p_L(1-b) + \alpha b\left(\frac{1}{2} - p_L\right) = p_L + \alpha b\left(\frac{1}{2} - p_L - \frac{p_L}{\alpha}\right) \tag{3.4}$$

The effect of the slant of the media (b) and consumer priors (p_L) on the WTP depends on parameter values in natural ways. A more slanted reporting (higher b) is more valuable for a consumer who prefers like-minded news (α high) and is more biased (low p_L). However, more bias hurts leftist consumers if they are close to the middle (p_L close to $\frac{1}{2}$) and do not mind reading news that goes against their priors too much (low α).

Similarly, rightist consumers will pay at most

$$WTP_R = max\left[1 - p_R - b\left(p_R + \alpha\left(p_R - \frac{1}{2}\right)\right), 0\right] \tag{3.5}$$

for a reporting with left-slant of b. Now b has a clearly negative effect on the consumer valuation, so is a stronger bias to the right (higher p_R). Consumers who are very biased and/or value like-minded news very much may not be willing to purchase any reporting – in an extreme case ($p_R = 1$), even an honest one for free.

3.2.2 Modeling assumptions

Consumers are different from the rational benchmark in two respects: first, they have subjective priors which are different from an objective priors, second, they have preferences over the strategy of the media.

The most important and somewhat debatable assumption I make is about the explicit preference of consumers towards reporting that they know, regardless of its actual content, will be slanted towards their already existing political beliefs (and consequently, having preference against reporting that has counter-attitudinal slant). While this assumption has been made in the literature previously, it mostly lacked microfoundation and was based introspective plausibility. While I am not able to give a full microfoundation for my modelling choice I offer two types of findings that support my decision to move away from a fully rational case.

First, in a standard framework with rational consumers and different (political) priors people prefer news outlets that slant towards them because, given their possibly incorrect priors, they can rationally think that like-minded (pro-attitudinal) media is more trustworthy or better informed than a media outlet with counter-attitudinal reporting: if there is uncertainty about the precision of a report, the closer it is to a Bayesian consumer's prior, the higher valuation the consumer will attach to it. This would imply that people are willing to pay more for like-minded news and they would rate like-minded news outlets more trustworthy. These predictions are empirically also validated (for the first, see Gentzkow and Shapiro (2010), for the second, see e.g. Barnidge et al. (2020)). However, some findings in the political science literature indicate that forceful exposure to counter-attitudinal news (in an experiment) changes the views of at least some participants, usually strengthening them in their priors and making them more extreme (see for example Levendusky (2013)). If the only reason why people avoid listening to counter-attitudinal news is because they find it lacking in information content, being forced consume such reporting should not change their beliefs (and definitely not away from the center), as they would consider such reporting fundamentally irrelevant. The fact that it does not seem to be the case empirically points towards the need for some sort of non-standard consumer behavior.

Secondly, at it has been documented by others (e.g. Levendusky (2013), Leeper and Slothuus (2014)) there is at least a superficial connection between motivated reasoning and preference for like-minded news. If people want to believe for some reason that they are skilled, intelligent or generous, they might avoid information that would prove them wrong, even though that incorrect belief might hurt them in other ways. If for some reason people attach similar value to believing that a liberal or a conservative view is right in some issue, it might also make them avoid information sources that in expectation tells them otherwise. My modelling choice offers a very reduced form way to achieve this effect. However, it is important to point out that the way I model taste for pro-attitudinal slant is different from motivated beliefs. The crucial difference is that while agents with motivated beliefs collect and evaluate information in order to sustain a preferred belief, in my model consumers are fully Bayesian in the sense that they understand that slanted news is less informative and they – correctly – put a low weight on a biased report: however, they still might want to consume as it, as, by assumption, consumer value pro-attitudinal slant in itself.

Another crucial assumption is that consumers have preferences over the reporting strategy of the media, and not the actual news they end up consuming. From a technical perspective, this is a modelling shortcut: it simplifies the consumer's problem as they do not have to form beliefs about the content of the reporting they are considering to buy, hence they are never "surprised" by the reporting they end up buying. However, it is not a completely unrealistic assumption in the sense that consumers usually have a good idea of the partisan lean of most news outlets and their decision whether to consumer the news they produce is based on prior knowledge about the news outlet, rather than the actual reporting.

I also assume that the media chooses its reporting strategy without knowing the true state of the world and it can commit to this announcement. Again, this simplifies the model considerably since the optimal reporting strategy will not convey any information about the truth to consumers – which would clearly affect how they value the reporting in expectation. Similarly, this is not entirely unrealistic either, as the bias of media outlets is typically stable or changes very slowly over time, and, more importantly, does not usually change from issue to issue. Admittedly, there are cases when partisan news consumers are negatively surprised by the editorial decision their favored media outlet takes – those cases fall outside of the applicability of this model.

Additionally, I consider the case of only one news outlet. At first glance, this is a very restrictive assumption. However, I find it useful for two reasons. First, models with two consumer types and two producers can very easily end up delivering "polarized" equilibria in the sense that in equilibrium the two producers serve their respective market segment with (partially of totally) slanted news. By restricting myself to one media outlet, I make the the tradeoff the news producers faces between market size and valuation starker and by turning off the competitive aspect of media behaviour I am able to show that even under this restrictive assumption there is an important sense in which political polarization can occur in equilibrium: under some conditions, as consumers on one end of political spectrum radicalize, it can push the other group towards their respective extreme too.

Another reasoning for this setup is that it can more explicitly speak to a setting where the informed agent in the game is not a media outlet but a set on pundits or opinion leaders that inform the general public on issues that plausibly require expertise (for example, foreign or environmental policy, public health etc). While the objective function of the opinion leadership is very likely to be different to that of a profit-maximizing media company that I assume, my setup at least acknowledges the importance of this scenario and can hopefully illustrate some mechanisms that are likely, but not obviously not guaranteed to survive in a setting that models this setup carefully.

In a similar fashion, I assume that the media produces only one kind of reporting and it cannot price differentiate. As before, this modelling choice allows me to concentrate on the novel mechanisms in my discussion and helps to avoid the very straightforward equilibria when the two types consume two different reports that happen to cater maximally to their tastes. While it is possible to find counterexamples for this setup in the real world[3], the literature on media capture (e.g. Enikolopov and Petrova (2015)) suggests that this kind of behavior is empirically less relevant than the kind of political media capture whose goal is the political persuasion of its viewers.

While it is not central to the results on the optimal behavior of the media, I assume that consumers are Bayesian. This implies that an unbiased news source can (and under some parametric conditions, will) make consumers learn the truth – the only reason why this does not always happen is that they dislike being informed by the "wrong" kind of reporting (that is, reporting that is not biased towards their prior), and in response the media will, in general, distort its reporting to cater to the bias of some of its buyers. This implies that, while this is not featured in the model, should a free source of verifiable unbiased information be available, there would be at least some customers who would actively seek that out and switch to it, ultimately correcting their biased beliefs.

3.3 Results

3.3.1 Optimal media behavior

The following proposition characterizes the optimal behavior of the news provider.

Proposition 3. *For constants $\bar{\alpha}_0$ and $\bar{\alpha}$ which are functions of the underlying parameters and defined in the proof, the optimal action of the media is given as follows.*

1. *It reports the truth ($b = 0$) and*
   (a) *sells to rightist consumers only and charges $P = 1 - p_R$ if $\alpha < \frac{p_L}{\frac{1}{2} - p_L}$ and $p_L \geq 2(1 - p_R)$,*
   (b) *sells to leftist consumers only and charges $P = p_L$ if $\alpha < \frac{p_L}{\frac{1}{2} - p_L}$ and $p_L \leq \frac{1}{2}(1 - p_R)$,*
   (c) *sells to both types if and*
      i. *charges $P = p_L$ if $\alpha < \frac{p_L}{\frac{1}{2} - p_L}$ and $\frac{1}{2}(1 - p_R) < p_L < 1 - p_R$,*
      ii. *charges $P = 1 - p_R$ if $\alpha < \frac{p_L}{\frac{1}{2} - p_L}$ and $1 - p_R < p_L < 2(1 - p_R)$.*

[3] During the Brexit referendum capaign in the United Kingdom, *The Times* endorsed "Remain", while *The Sun* endorsed "Leave", even though both of them were owned by Rupert Murdoch.

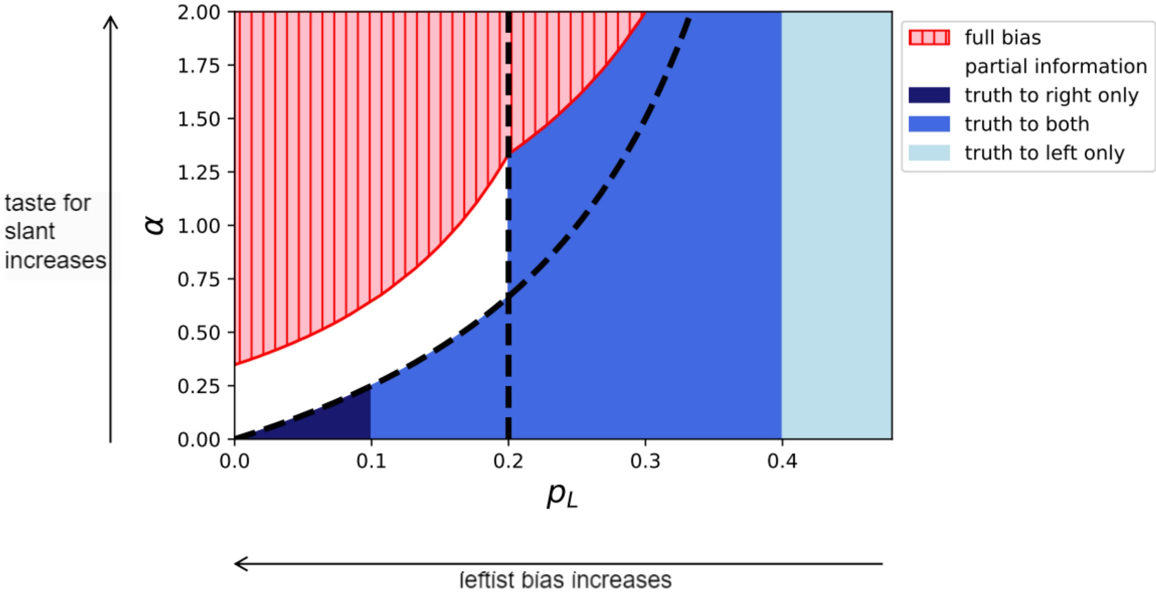

Figure 3.1: Equilibria for various values of leftist priors (p_L) and weight on psychological utility (α), with $p_R = 0.8$. Taste for slant increase from bottom towards top. Leftist bias increases from right to left.

2. *It sets the bias to* $0 < b^*(p_L, p_R, \alpha) < 1$ *and sells to both types for* $P = p_L + \frac{1-p_L-p_R}{(\alpha+1)(p_R-p_L)}\left(\frac{\alpha}{2} - p_L(1+\alpha)\right)$, *if* $p_L < 1 - p_R$ *and* $\bar{\alpha} > \alpha > \frac{p_L}{\frac{1}{2}-p_L}$.

3. *It sets the bias to* $b = 1$, *charges* $P = \alpha\left(\frac{1}{2} - p_L\right)$ *and sells to left-leaning types if* $\alpha > \bar{\alpha}$ *and* $p_L < 1 - p_R$ *or if* $\alpha > \bar{\alpha}_0$ *and* $p_L > 1 - p_R$.

Proof. In the appendix. □

To build intuition for this result, I will focus on a numerical example in Figure 3.1. This figure depicts the optimal level of bias and the types that buy in equilibrium for some combinations of p_L and α, with the value of rightist priors p_R fixed.

For an overview of some of the possible equilibria, I will fix the taste for slant at $\alpha = \frac{1}{2}$ and gradually decrease leftist prior from $p_L = \frac{1}{2}$ to 0, that is, I start out with leftist consumers having fully correct priors and I look at the equilibrium reporting strategy as I gradually consider more and more biased leftist consumers with lower and lower values of p_L. Apart from the basic tradeoff that bias is valued differently by two types, the media has two margins along which it can optimize: it can choose the slant of the news b (the "internal" margin) and, by adjusting the price, the types it sells to (the "external" margin). The relative importance of these margins will be pinned down by the exact value of the parameters.

For $p_L \approx \frac{1}{2}$ (at the right side of the figure) the internal margin dominates: the optimal decision here is to report the truth by setting $b = 0$ and charge such a high price that only leftist consumers buy. In this range, left-leaning consumers value truthful reporting so highly, that for the media it will be optimal to sell only to them and extract all their surplus. To see this, note that in general, less biased consumers value truthful reporting more: if $b = 0$, the total utility simplifies to the instrumental utility, so the type with priors closer to the truth of $\frac{1}{2}$ will value the information more as it helps him to correct his default guess of ω (his type) to the truth (the reporting sent by the media) with a higher subjective probability.

As p_L decreases along the horizontal line pinned down by $\alpha = 0.5$, leftist consumers become relatively more biased compared to rightist consumers, and at one point the media switches to a new optimal strategy which involves serving both types. This is a direct consequence of leftist valuations getting smaller: given that more biased leftist types value truthful reporting less, below a certain level of p_L the external margin becomes dominant and the media will find it profitable to sell to both types, albeit at a lower price.

For even lower values of p_L, the media will eventually abandon its zero-slant policy and will set a positive b in a way that keeps both types buying. Intuitively, in this case, leftist care relatively little about the truth and the media will have to find a better way to keep their valuations high – and it will do it through introducing slant. At this point, however, slant remains strictly partial ($b < 1$) as the external margin is still strong: leftist preference for slant is not high enough that the media would want to sell only to them, so it has to give some value to rightists by providing them with somewhat informative reporting.

Finally, as p_L gets very low, leftist consumers become very biased which means that they value informative reporting very little. At this point the internal margin starts to dominate and serving only leftist consumers with fully slanted news ($b = 1$) becomes the most profitable strategy. Intuitively, high levels of left slant, induced by low levels of p_L hurts right-leaning consumers, meaning they are willing to pay little for such news, so at one point the media would prefer to go all they way and serve leftists only, but now with full slant – as in this range there are no rightists to serve and leftists like slant, there will be no reason not to slant the news fully.

After this brief overview of some features of the optimal strategies, in the following, I will discuss the structure of Figure 1 in more detail and briefly explain how it connects to the media's optimal choice of slant and price.

The two dashed lines divide the figure into four quadrants. In the two quadrants below the curved dashed line, α is relatively low, hence demand for slant is 0 – even leftist consumers value truthful reporting more than any other news with positive slant. This makes the news provider to always report truthfully ($b = 0$ in equilibrium) in this region. Note how the slope of this curve is strictly positive. This comes from the fact that an increasing p_L (a more correct prior) leads to a higher valuation of the truth: a less biased consumer would need a higher α to make him indifferent between honest news and news with a positive slant.

Which types end up buying under the curve depends on the relative sizes of the priors: if leftists are far out in the political spectrum (they have a low prior p_L about $Pr(\omega = R)$), they do not attach too much instrumental utility to learning to truth (as they are fairly certain that $\omega = L$), so the media is left with selling only to rightists, extracting all the surplus from them. As we have seen, as leftists move closer to the center, they gradually understand better the uncertainty they face with respect to the state of the world, so they value honest reporting more. In this region, the media chooses the price in such a way that both types will buy, leaving some surplus at the less biased type. Finally, if left-leaning consumers are very close to the political center, that is, their priors are almost correct, their valuation is so high that selling only to them becomes the dominant strategy: their willingness to pay is so high that it offsets the loss of profits coming from the smaller share of buyers.

The other division, represented by the vertical dashed line at $1-p_R = 0.2$ (with $p_R = 0.8$ fixed) shows the relative biases of the two types: left of this line in the figure, left-leaning consumers are more biased than right-leaning ones – that is, their prior is farther away from one-half than that of the rightist.

If demand for slant is zero (under the curved dashed line), whichever consumer is more biased will value honest reporting less. So if the monopolist will sell to both types (in the middle of this interval), it will have to price the no more than the minimum of the two valuations, which in turn will be determined by the willingness to pay of the more biased the type. Therefore the price will be different on the two sides of this segment, not only because the parameter values are different, but because the type with the lower valuation will be different too.

Next, I turn to the case when the valuation of leftist consumers is increasing in the reporting bias. Moving on the upper half of the figure, it represents parameter combinations under which positive slant is possibly a non-dominated strategy.

In this case, leftist consumers enjoy bias as their valuation increases in b. This creates incentives for the media to provide a strictly positive amount of bias in at least some optima. As we have seen, for α high enough, it will provide totally uninformative reporting (by always reporting left) and sell only to left-leanings customers. One way to understand why such a situation is guaranteed to emerge is to note that $WTP_L(1)$, the leftist valuation of uninformative news is increasing in α without bounds. In contrast, the profits from selling an honest reporting is clearly bounded from above by $\frac{1}{2}$: a consumer would never pay more than that to avoid a loss of 1 that happens with a subjective probability less than one half. Consequently, over a certain threshold of α, the profits from selling $b = 1$ news to leftists will dominate all other options. This basic intuition holds, regardless of which party is more biased.

While the implication that the value of uninformative reporting for a sufficiently biased consumer is theoretically unbounded, may not sound entirely realistic (in the real world, consumers do not spend all of their wealth on e.g. political reporting), it is best understood as the *relative* weight of the behavioral part of the utility function relative to the instrumental part: this way, the case of $\alpha \to \infty$ corresponds to a case when learning the truth is increasingly less important compared to the utility gained from paying attention to the right kind of media.[4]

Still staying in the top half of the figure, if leftists are *less* biased than rightists (top right quadrant), the structure of the optimum is relatively simple. If α is below the threshold described in the previous paragraph, the media will sell unbiased reporting to both types at a price equal to the valuation of the

[4]Another way of thinking about price in this setting is the following. In most real world setting, it is not possible to buy single-issue reporting alone, rather, it is bundled together with other content: traditional TV channels sell political reporting along with non-political entertainment programs, online news sites feature advertisements that some consumers dislike etc. An increasing price can be interpreted as the deterioration of the nonpolitical components: in this setting it would imply that as consumers become more partisan, they will tolerate less exciting sport broadcasts on TV and more advertisement online if they are served with the right amount of political slant.

more biased group, as before.

If, on the other hand, leftists are *more* biased than rightists (top left quadrant), it becomes feasible even in an optimum for the media to distort only partially by setting the bias equal to some $b^* \in (0,1)$ where the level of bias, b^* is picked in a way that equates the valuation of left and right-leaning consumers. Intuitively, this has the potential advantage for the media to extract all the surplus from both types (since both types have positive valuations).

It is important to point out that this was not possible when leftist were the less biased party – left-leaning consumers valued news more for *every* level of bias: they valued it more for $b = 0$ (as less biased consumers value the truth more) and their valued biased reporting progressively even more (as their valuation was increasing in b). Therefore there was no level of bias that would equate the valuation of the two groups. Very simply put, leftist were too good customers for the media not to cater exclusively to their tastes. However, now that leftists are relatively more biased, such a level b^* does exist: in this case, leftists do not value the information content of the news that much (as they think that they are right with a relatively high probability), but this will be offset by the fact they like biased reporting. [5]

To finish, I discuss some comparative static of the optimal level of partial slant, b^*.

It is clearly decreasing in p_L: that is, the more biased leftist consumers are, the less informative is the news that the media sends in equilibrium (as long as partial distortion is optimal). The intuition is almost trivial: by choosing b^*, the media caters to the tastes of a left-leaning consumer who values slant, so in equilibrium a slightly more biased consumer will get served slightly more biased news.

However, b^* turns out to be decreasing in α: that is, the higher the weight that consumers place on the behavioural part of the utility function (value slant), the less the media distorts (again, as longs as $0 < b < 1$). While this might be counterintutive at first, the reason behind this is the following. Conditional on choosing the strategy $b^* \in (0,1)$, the monopolist is effectively constrained by the fact that its reporting strategy has to equate the valuations between the two types. If it did not do that, the media could trivially change the slant and improve on the payoff it gets. However, if leftist are more biased than rightists (we are in the top left quadrant), an increase in α makes leftists more sensitive to changes in bias than rightist[6]. Intuitively, leftists in this range care relatively little about the truth and relatively lot about the slant already so an even higher α means just more weight on the term with the slant in it.

Therefore, if the monopolist picks b^* for some parameters and α goes up, this change will benefit leftists more than rightist lose. To keep both groups buying, the monopolist has to redistribute the positive net gains of a higher α by slightly cutting back on the slant, which hurts leftists somewhat but benefits rightists and thereby restores the equilibrium.[7]

[5]The fact that in this case (top right quadrant) switches to truthtelling is a consequence of the simplifying modelling assumptions: in a more richer model where the monopolist is allowed to produce right-slanted reporting it is conjectured that it would do so in this region.

[6]That is, $\left|\frac{\partial^2}{\partial\alpha\partial b}WTP_L\right| > \left|\frac{\partial^2}{\partial\alpha\partial b}WTP_R\right|$

[7]Graphically, in the top-right panel of Figure C.1.1 an increase in α corresponds to the L curve getting much steeper, with the R getting only somewhat flatter, which leads to the intersection, b^*, moving to the left.

3.3.2 Implications on beliefs and welfare

In this section I will examine how the average belief and welfare evolves for the two types. I will split my discussion in two: I start with presenting results for a setting when the media always reports truthfully, next I repeat the exercise for the rest of the parameter space, where positive slant can be optimal.

I will highlight two important comparative statics as separate corollaries. The first one will describe an important across-type comparative static of welfare for low values of α. Then, I present a second corollary which, for medium values of α establishes a relationship between the posterior of one type and the prior of the other type.

Comparative static when bias is a net bad

First, I look at the case when both types dislike bias – in Figure 3.1, this corresponded to the area under the dashed curve, with values of relatively low α. In this case, as it was established previously, the only kind of news that the media sells is honest reporting. This implies that consumers either buy the news and correct their beliefs to the truth (which is on average 0.5) or they do not buy it, leaving their priors unchanged.

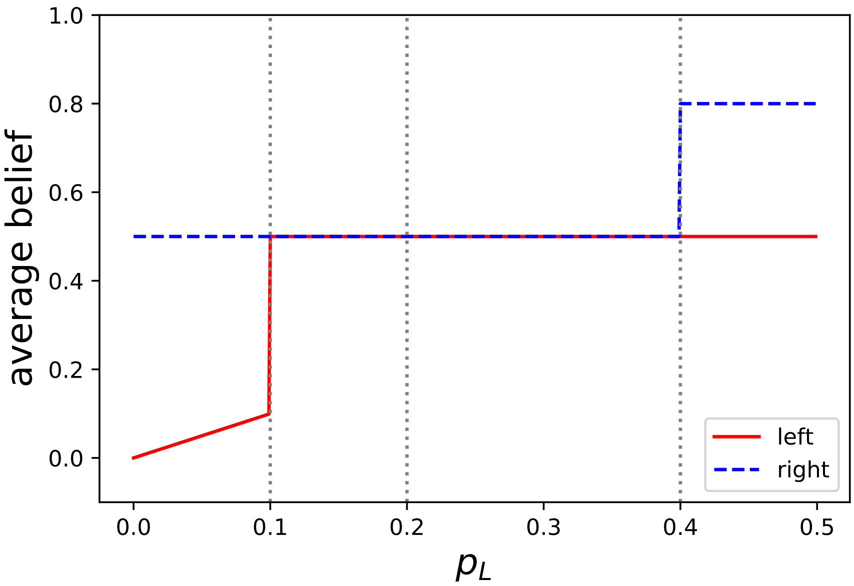

Figure 3.2: The evolution of average beliefs as function of p_L. ($p_R = 0.8$) Red denotes left, blue denotes right beliefs. Dashed lines correspond to $(1 - p_R)/2$, $1 - p_R$ and $2(1 - p_R)$.

As a reminder, for very low values of p_L (very biased leftist consumer) the media sells only to right-leaning buyers, whose beliefs are going to be correct. Next, for medium values of p_L, both types buy the undistorted news, leading to correct belief in the entire market. Finally, if left-leaning consumers are close to the truth (p_L is close to $\frac{1}{2}$), the media prices in such a way that only they buy the reporting, leaving the beliefs of right-leaning consumers unchanged. Note that this last segment does not exist for $p_R < \frac{3}{4}$: if rightists are not too biased in their priors, they will always buy the undistorted news, leading to correct beliefs for both groups above the first threshold $\frac{1-p_R}{2}$.

A change in p_R has an additional effect of changing the set of p_L's for which the media will offer honest reporting to *both* groups, thereby correcting beliefs. If right-leaning consumers become more biased (p_R goes up), the highest leftist belief which is not corrected ($\frac{1-p_R}{2}$) goes down, leading to more leftist but less rightist buying news. While the second effect is somewhat trivial (more biased consumers will value even the truth less), the first effect comes from leftist consumers around the lower threshold becoming relatively more attractive for the media – given that rightist now value news less. While in my model I only have two types, one leaning left and another one to the right, somewhat imprecisely one can say that rightist consumers moving further to the right hurt centrist leftists (whose priors are close to one half) and may even benefit more radical leftists.

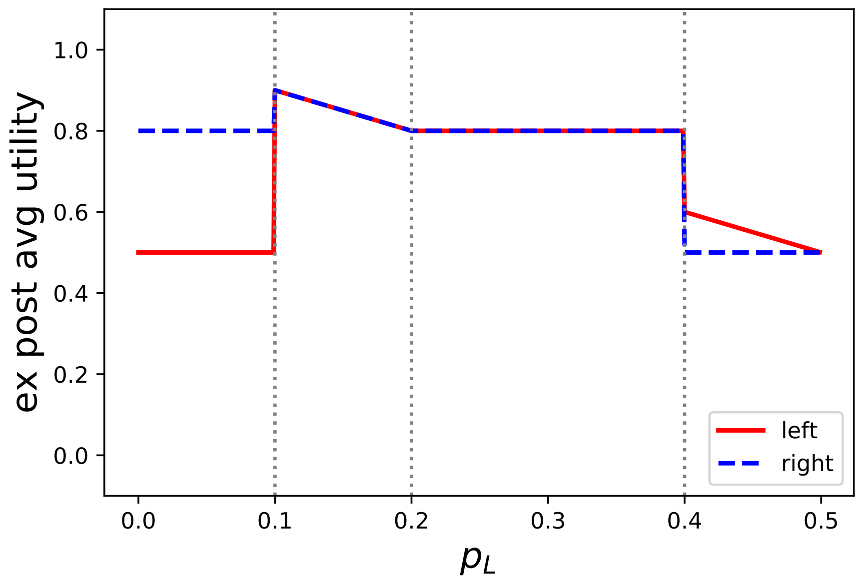

Figure 3.3: The evolution of ex post average utility as function of p_L. ($p_R = 0.8$) Red denotes left, blue denotes right EU. Dashed lines correspond to $(1 - p_R)/2$, $1 - p_R$ and $2(1 - p_R)$.

Next, I plot the ex post average utilities for various parameter values in Figure 3.3. To calculate these, I use objective probabilities (0.5 and 0.5 for $P(\omega = L)$ and $P(\omega = R)$) as opposed to subjective priors. This is an important distinction, since using priors may bias the expected utility upwards. For example for $p_L = 0$, leftist consumers subjectively expect 1: they are certain that the truth is L which makes them guess left and decline to buy any form of news, while the objective expected utility for that interval is $\frac{1}{2}$. Similar considerations hold for rightist consumers in the rightmost interval ($p_L > 0.4$): they do not buy news and from an external point of view they overestimate the payoff of guessing right.

This figure motivates the following corollary:

Corollary 1. *There is an interval for leftist beliefs where rightist welfare in increasing in p_L: on that interval rightist consumers are on average better off as leftists move closer to the truth.*

This results reiterates a key welfare implication of the equilibrium that has been identified in the previous section: for medium levels of p_L the presence of the other group benefits the less biased types. Since in this region the media sells unbiased reporting to both groups, it has to price it at the minimum of the two valuations. For right-leaning consumers this shows up a discrete jump around the first threshold (0.1 in the figure): rightist consumers gain as extreme leftists move a bit towards the centre. This way, the existence of biased leftist consumers benefits the other type.

Comparative static when bias is not a net bad

Next, I move on to the case of leftist valuation increasing in the bias. As a reminder, this is the area over the dashed curve in Figure 3.1. As we have seen, under these conditions, in addition to honest and noninformative reporting, there exists a partially revealing strategy b^*, which will be the focus of this section.[8]

The following corollary establishes a connection between leftist priors and rightist posteriors in this case.

Corollary 2. *If b^* is optimal, average rightist beliefs are weakly decreasing in p_L: as leftist consumers become more biased, rightists will be less correct on average.*

To calculate the average belief in this case, for any combination of priors $p \in \{p_L; p_R\}$, define $B(m; p)$ as the posterior belief that $\omega = R$ as a function of the message $m \in \{l, r\}$ (l and r denote the events that the reporting is a "left" or "right" message, respectively.)

Consequently, the average belief will be given as

$$\bar{B} = P(l) \cdot B(l; p) + P(r) \cdot B(r; p) = \frac{1}{2}(1 + b^*) \cdot \frac{pb^*}{1 - p + pb^*} + \frac{1}{2}(1 - b^*) \cdot 1, \tag{3.6}$$

as under a reporting strategy b^*, the media will send $\frac{1}{2}(1 + b^*)$ left and $\frac{1}{2}(1 - b^*)$ right reports in expectation.

The corollary states that $\frac{\partial}{\partial p_L}\bar{B} < 0$ for right-leaning consumers. This can be directly checked by observing that for rightists

$$\frac{\partial}{\partial p_L}\bar{B} = \underbrace{\frac{\partial \bar{B}}{\partial b^*}}_{+} \cdot \underbrace{\frac{\partial b^*}{\partial p_L}}_{-}. \tag{3.7}$$

Clearly, a lower b^* (more informative reporting) is associated with average beliefs which are closer to the correct value of one half: when more information is bought, beliefs are going to more correct.[9] This explains the positive sign of the first partial derivative: if the reporting becomes less informative, consumers will update less and rightist posteriors will increase. The corollary connects this statement to the comparative static of b^* with respect to p_L outlined at the end of section 3.1: if partial distortion is optimal, a decrease in p_L (leftists getting more biased) increases b^* (less information is revealed). Hence in equilibrium beliefs will be less correct on average.

While this is corollary is a direct consequence of two already stated results, it leads to seemingly counterintuitive consequences – for example, if a left-leaning media moves towards the right (it becomes

[8] The equilibrium when $b = 1$ does not offer any interesting comparative statics.

[9] The derivative of average beliefs with respect to b^* is $\frac{1}{2}\frac{(1-p)(2p-1)}{(b^*p-p+1)^2}$, which is positive for all $1 > p_R > \frac{1}{2}$ and negative for $0 < p_L < \frac{1}{2}$.

somewhat more centrist by reducing its bias from $b^* + \epsilon$ to $b^* > 0$), right-leaning consumers will move to the *left* in expectation. The mechanism for that is that while the media will send more r and less l messages with its bias reduced, l messages will be more informative and this will make the rightists move towards the center on average.[10] Note that in this model, consumers are fully Bayesian so I make no extra assumptions to arrive at this kind of behavior.

Figure 3.4 offers a graphical representation of this phenomenon: it depicts average beliefs as function of p_L for different values of α. The top panel features a middle value of α, which implies that under a certain threshold of p_L, the media will distort partially by setting its bias to b^*. This level of distortion decreases as leftists get less biased, which is reflected in beliefs gradually getting closer to one half. Over this threshold the media provides unbiased news, thereby correcting beliefs completely.

This panel (Figure 3.4(a)) illustrates the key mechanism of the model: a change in leftist priors affects rightist posteriors through media behavior. As leftists become more biased (p_L decreases), the media also moves to the left (as $\frac{\partial b^*}{\partial p_L} < 0$). A more biased media is, almost by definition, produces less informative reporting. Rightist customers, perceiving this correctly, give (on average) less weight to what the left-biased media has to say. As a consequence, they update less, so their posteriors tend to remain closer to their (right-leaning) priors.

This scenario is depicted in the middle of Figure 3.4(a). Consider a decrease in p_L (leftists getting more biased). As soon as they become the more biased party, the truthtelling equilibrium breaks down. On the figure this happens at the dotted line where $p_L = 0.2$ (as p_R was set to $1 - p_L$). Below that threshold, the average beliefs of the two group start to diverge from the objectively correct one-half in a smooth fashion: as the media starts to distort beliefs by setting b^* to be positive, consumers rationally put relatively more weights on their priors, leading to liberals (conservative) staying more liberal (conservative) on average. (This is indeed a smooth process as at the threshold $b^* = 0$.) At one point with lefitsts getting very biased (p_L being very low), the imperfect information revelation strategy will be suboptimal to the media and it will switch to catering to leftist tastes fully by setting $b = 1$. Under this regime, posteriors will be exactly equal to the priors as only type buys the news and it is completely uninformative anyways.

To contrast this logic with that with that of the previous section (when α, the taste for like-minded reporting was low): in the previous case the mechanism operated through an extensive margin – since the media only sold perfectly correct information, some consumers decided to forgo buying news altogether if their priors were far away from the truth (as it made them value honest news less). This was a symmetric phenomenon: extreme leftists were as likely to do that as extreme rightist. Here, under the same parameter values the media finds it more profitable to keep serving leftist (since now the net effect of bias on their valuation is positive), even if this means gradually distorting its reporting (changes on the internal margin). This is the new, internal margin at which rightist beliefs change differently: they do not avoid the news, however they rationally trust it less as the slant of the media increases.

In line with the results of Figue 3.1, partial revelation does not have to be an equilibrium for combi-

[10]To see this, split the total derivative of $\bar{B}$ to $P(r)' + P(l) \cdot P(\omega = R|l)'$ and $P(l)' \cdot P(\omega = R|l)$. The first term pushes right-leaning consumers towards the center, the second term pulls them towards the extreme. It can be shown that the first one is always greater. A similar argument holds for left-leaning consumers.

10.14754/CEU.2022.03

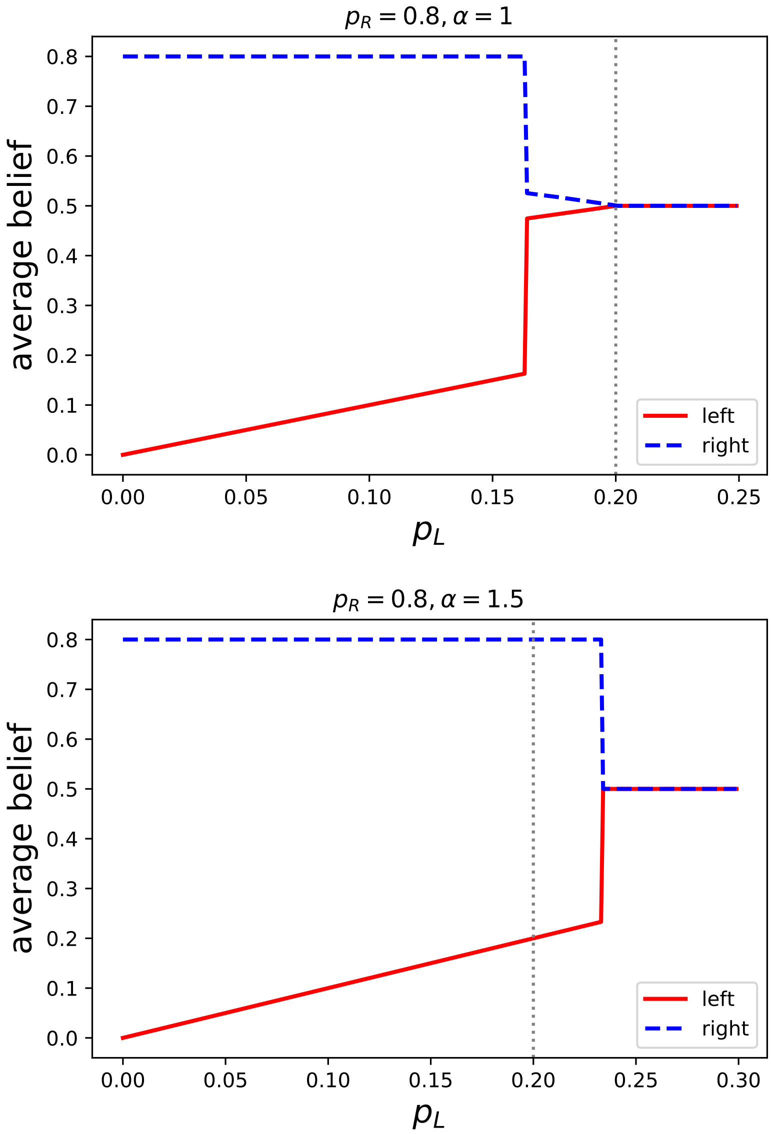

Figure 3.4: Average beliefs for different values of α. Red denotes left, blue denotes right beliefs. Dashed line is at $p_L = 1 - p_R$, the point where the relative size of biases change. Note how the horizontal axis is different between the figures due to the requirement that $\alpha > \frac{p_L}{\frac{1}{2} - p_L}$.

nation of parameters. The bottom panel depicts the case with higher α. This changes the structure of the optimum somewhat, as selling completely biased news to leftist becomes more profitable and for low values of p_L (that is, for more bias), indeed this is will be the best option for the media. This naturally leaves beliefs unchanged. Only for less biased leftist will the media switch to the truthtelling equilibrium and correcting both types. In this setting, there is no space for partial revelation of the truth: if the relative weight of behavioral utility is high enough the media will either tell the truth to both groups or will totally cater to the bias of one group.

For the sake of completeness, in figure 3.5 I plot the expected ex post average utilities for the same parameter values as a function of leftist priors.

Clearly, left-leaning consumers will always follow the media: if the message is r, they will know certainly that $\omega = R$ and $m = l$ will never convince a left-biased consumer that ω is not L. For $b = b^*$, which is the interesting case, the average ex post utility for leftists is therefore

$$\frac{1}{2} \cdot 1 + \frac{1}{2} b^* \cdot 0 + \frac{1}{2}(1 - b^*) \cdot 1 + \phi(b^*, p_L). \tag{3.8}$$

For right-leaning consumers a similar relationship holds.

The main intuition that lower leftist bias (higher p_L's) makes the media to sell unbiased news carries through. This results in a jump of average utility around the threshold when this happens. At this point, both types enjoy the same expected utility: their beliefs will be correct (therefore their behavioral utility will be zero) and they pay the same price.

On the other end (if p_L is very low) the media will sell completely uninformative news that only left-leaning consumers buy. Rightists will always guess "right" and since they do not buy news, they will make 0.5 on average. Left-leaning consumers however will pay a positive price for uninformative reporting which explains what their average utility is less than those of rightists.

For very high levels of α these two scenarios describe the expected utility fully (bottom panel of Figure 3.5). As it was shown above, for an intermediate taste for confirmatory news (top panel) there exists an intermediate range of leftist bias for which partial distortion is the optimal strategy of the media. This delivers high utility for leftist consumers: they gain information and they get to consume reporting that slants towards them. However, right-leaning consumers also benefit: they also get informative reporting and even though they do not like leftist slant, they get it at a relatively low price – so on average their utility might be even slightly greater than in the case of the truthtelling media.

3.4 Conclusion

I have demonstrated how the existence of a media that both sides of the electorate can consume will under some conditions will link the beliefs and welfare of the two parties in specific ways depending on the overall preference for political slant.

While in my model consumers exhibit preferences that can generate a rich set of behavior, the market for news is admittedly a very simple one: there is only one perfectly informed media outlet that provides information on only one issue with a fixed direction of slant. Doing away with these assumptions may give

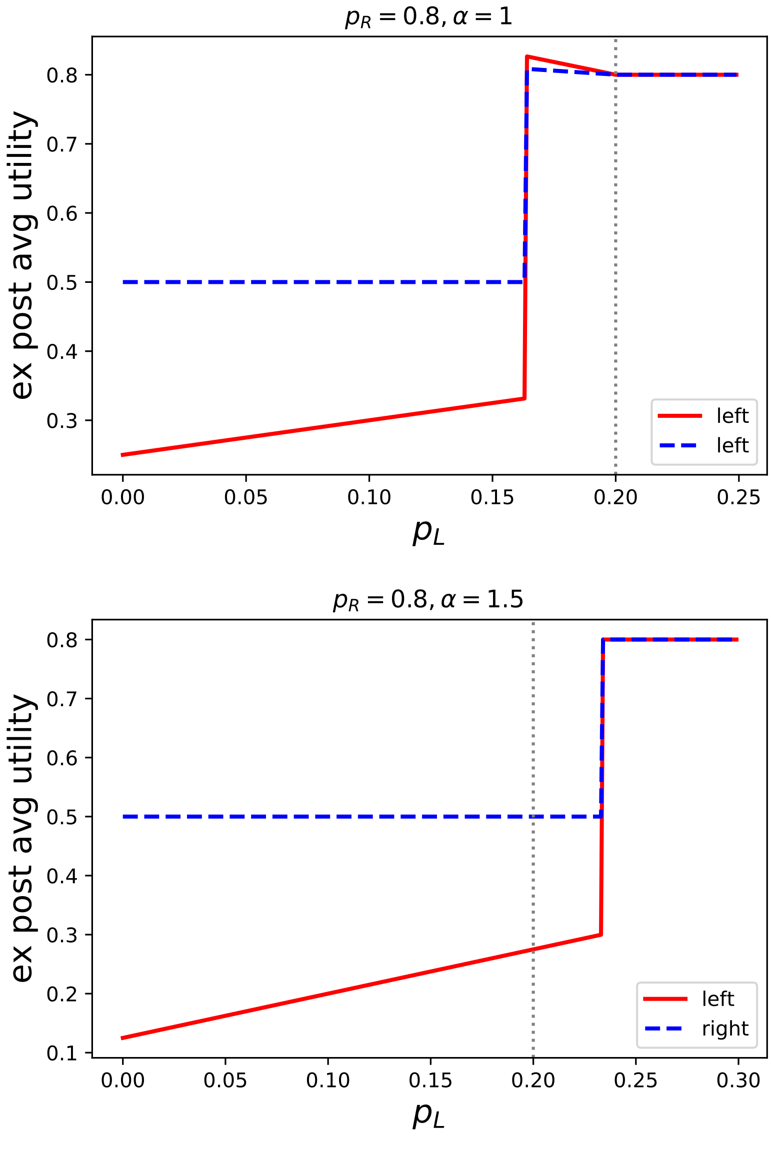

Figure 3.5: Ex post average utilities for different values of α. Red denotes left, blue denotes right EUs. Dashed line is at $p_L = 1 - p_R$, the point where the relative size of biases change. Note how the horizontal axis is different between the figures due to the requirement that $\alpha > \frac{p_L}{\frac{1}{2} - p_L}$.

a more realistic set of predictions. In particular, introducing a second media outlet may shrink the set of parameters that give rise to the partially revealing equilibrium. In this sense, this model can be thought of a first stage of a game where under some conditions a second media outlet can enter the market.

Bibliography

Allen, Eric J., Patricia M. Dechow, Devin G. Pope, and George Wu, "Reference-Dependent Preferences: Evidence from Marathon Runners," *Management Science*, 2017, *63* (6), 1657–1672.

Ashenfelter, Orley and David Genesove, "Testing for Price Anomalies in Real-Estate Auctions," *The American Economic Review*, 1992, *82* (2), 501–505.

Backus, Matthew, Thomas Blake, and Steven Tadelis, "On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining," *Journal of Political Economy*, 2019, *127* (4), 1599–1628.

Barnidge, Matthew, Albert C Gunther, Jinha Kim, Yangsun Hong, Mallory Perryman, Swee Kiat Tay, and Sandra Knisely, "Politically motivated selective exposure and perceived media bias," *Communication Research*, 2020, *47* (1), 82–103.

Barron, Kyle, Edward Kung, and Davide Proserpio, "The Effect of Home-Sharing on House Prices and Rents: Evidence from Airbnb," *Marketing Science*, January 2021, *40* (1), 23–47.

Bernhardt, Dan, Stefan Krasa, and Mattias Polborn, "Political polarization and the electoral effects of media bias," *Journal of Public Economics*, 2008, *92* (5), 1092–1104.

Besley, Timothy and Andrea Prat, "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability," *American Economic Review*, June 2006, *96* (3), 720–736.

Calder-Wang, Sophie, "The Distributional Impact of the Sharing Economy on the Housing Market," *Working paper*, 2020.

Dai, Weijia and Michael Luca, "Digitizing Disclosure: The Case of Restaurant Hygiene Scores," *American Economic Journal: Microeconomics*, May 2020, *12* (2), 41–59.

DellaVigna, Stefano and Ethan Kaplan, "The Fox News Effect: Media Bias and Voting*," *The Quarterly Journal of Economics*, 08 2007, *122* (3), 1187–1234.

Dube, Arindrajit, Alan Manning, and Suresh Naidu, "Monopsony and Employer Mis-optimization Explain Why Wages Bunch at Round Numbers," Working Paper 24991, National Bureau of Economic Research September 2018.

Durante, Ruben and Brian Knight, "Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy," *Journal of the European Economic Association*, 2012, *10* (3), 451–481.

Díaz, Antonia and Belén Jerez, "House prices, sales, and time on the market: a search-theoretic framework," *International Economic Review*, 2013, *54* (3), 837–872.

Enikolopov, Ruben and Maria Petrova, "Chapter 17 - Media Capture: Empirical Evidence," in Simon P. Anderson, Joel Waldfogel, and David Strömberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 687–700.

Fedyk, Anastassia, "Asymmetric Naivete: Beliefs About Self-Control," *Working paper*, 2021.

Garcia-López, Miquel-Angel, Jordi Jofre-Monseny, Rodrigo Martínez-Mazza, and Mariona Segú, "Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona," *Journal of Urban Economics*, 2020, *119*, 103278.

Genesove, David and Christopher Mayer, "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *The Quarterly Journal of Economics*, 2001, *116* (4), 1233–1260.

Gentzkow, Matthew and Jesse M Shapiro, "Media bias and reputation," *Journal of political Economy*, 2006, *114* (2), 280–316.

_ and _ , "What drives media slant? Evidence from US daily newspapers," *Econometrica*, 2010, *78* (1), 35–71.

_ , Jesse M. Shapiro, and Daniel F. Stone, "Chapter 14 - Media Bias in the Marketplace: Theory," in Simon P. Anderson, Joel Waldfogel, and David Strömberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 623–645.

HCSO, "Microcensus data by Budapest districts," *Hungarian Central Statistics Office*, 2016.

Horn, Keren and Mark Merante, "Is home sharing driving up rents? Evidence from Airbnb in Boston," *Journal of Housing Economics*, 2017, *38*, 14–24.

Hui, Li, Yijin Kim, and Kannan Srinivasan, "Market Shifts in the Sharing Economy: The Impact of Airbnb on Housing Rentals," *Working paper*, 2021.

Jancsik, András, Gábor Michalkó, and Márta Csernyik, "The evolution of Airbnb and hotel market in Budapest (in Hungarian)," *Hungarian Economic Review*, March 2018, *65* (3), 259–286.

Kandel, Shmuel, Oded Sarig, and Avi Wohl, "Do investors prefer round stock prices? Evidence from Israeli IPO auctions," *Journal of Banking & Finance*, 2001, *25* (8), 1543 – 1551.

Koster, Hans R.A., Jos van Ommeren, and Nicolas Volkhausen, "Short-term rentals and the housing market: Quasi-experimental evidence from Airbnb in Los Angeles," *CEPR Working paper*, 2019.

Lacetera, Nicola, Devin G. Pope, and Justin R. Sydnor, "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, May 2012, *102* (5), 2206–36.

Leeper, Thomas J and Rune Slothuus, "Political parties, motivated reasoning, and public opinion formation," *Political Psychology*, 2014, *35*, 129–156.

Leib, Margarita, Nils C. Köbis, Marc Francke, Shaul Shalvi, and Marieke Roskes, "Precision in a seller's market: Round asking prices lead to higher counteroffers and selling prices," *Management Science*, February 2021, *67* (2), 1048–1055.

Levendusky, Matthew S, "Why do partisan media polarize viewers?," *American Journal of Political Science*, 2013, *57* (3), 611–623.

Li, Hui and Kannan Srinivasan, "Competitive Dynamics in the Sharing Economy: An Analysis in the Context of Airbnb and Hotels," *Marketing Science*, 2019, *38* (3), 365–391.

Lynn, Michael, Sean Masaki Flynn, and Chelsea Helion, "Do consumers prefer round prices? Evidence from pay-what-you-want decisions and self-pumped gasoline purchases," *Journal of Economic Psychology*, 2013, *36* (C), 96–102.

Mullainathan, Sendhil and Andrei Shleifer, "The Market for News," *American Economic Review*, September 2005, *95* (4), 1031–1053.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel, "Segmented Housing Search," *American Economic Review*, March 2020, *110* (3), 720–59.

Pope, Devin G., Jaren C. Pope, and Justin R. Sydnor, "Focal points and bargaining in housing markets," *Games and Economic Behavior*, 2015, *93*, 89 – 107.

Prior, Markus, "Media and Political Polarization," *Annual Review of Political Science*, 2013, *16* (1), 101–127.

Repetto, Luca and Alex Solís, "The Price of Inattention: Evidence from the Swedish Housing Market," *Journal of the European Economic Association*, November 2019, *18* (6), 3261–3304.

Ross, Stephen L. and Tingyu Zhou, "Loss Aversion in Housing Sales Prices: Evidence from Focal Point Bias," NBER Working Papers 28796, National Bureau of Economic Research, Inc May 2021.

Shlain, Avner S., "More than a Penny's Worth: Left-Digit Bias and Firm Pricing," *The Review of Economic Studies*, 2021, *forthcoming.*

Strömberg, David, "Media and Politics," *Annual Review of Economics*, 2015, *7* (1), 173–205.

Valentin, Maxence, "Regulating short-term rental housing: Evidence from New Orleans," *Real Estate Economics*, 2021, *49* (1), 152–186.

Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov, "Political Effects of the Internet and Social Media," *Annual Review of Economics*, 2020, *12* (1), 415–438.

10.14754/CEU.2022.03

10.14754/CEU.2022.03

Appendix A

Appendix for chapter 1

A.1 Summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
price	27293	1.80E+07	1.65E+07	1000000	3.85E+08
size	27293	62.7013	44.18117	7	1932
nr rooms	27044	1.677636	0.935909	1	14
nr bath	26776	1.094264	0.349901	1	7
has balcony	27293	0.392372	0.488288	0	1
has garage	27293	0.091782	0.288723	0	1
lift	27293	0.483787	0.499746	0	1
build year	26497	1956.979	35.90874	1778	2018

Table A.1.1: Summary statistics for some main variables.

		price	size	nr rooms	nr bath	has balcony	has garage	lift	build year
before '13	mean	14 700 000	59.56	1.61	1.07	0.42	0.10	0.52	1 959.10
	sd	11 800 000	34.59	0.86	0.31	0.49	0.29	0.50	34.95
after '13	mean	20 800 000	65.30	1.73	1.11	0.37	0.09	0.46	1 955.32
	sd	19 200 000	50.63	0.99	0.38	0.48	0.28	0.50	36.56
total	mean	18 000 000	62.70	1.68	1.09	0.39	0.09	0.48	1 956.98
	sd	16 500 000	44.18	0.94	0.35	0.49	0.29	0.50	35.91

Table A.1.2: Differences of some hedonic characteristics between the low and and high demand period.

A.1.1 Change in mass: transaction level regression

$$round_i = \beta_0 + \beta_1(\text{after 2013})_i + \gamma X_i + \epsilon_i \tag{A.1}$$

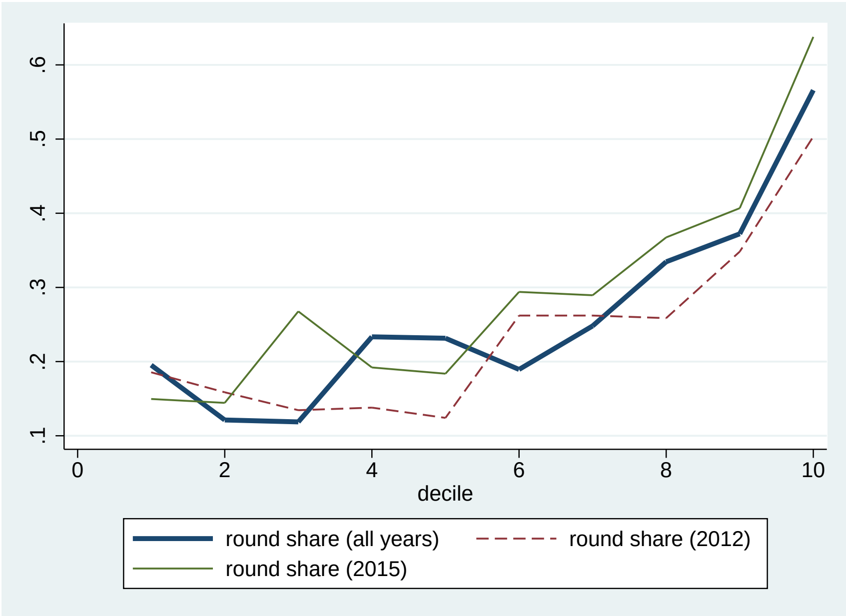

Figure A.1.1: Share of round prices by deciles.

The coefficient of the after 2013 dummy in the first column of Table A.1.3 corresponds to the roughly 6 percentage point difference between the two red columns of Figure 1.7.

	all price	all price	price$< 20M$	price$< 15M$
after 2013	0.0604***	0.0011	0.0100*	0.0152**
	(0.0055)	(0.0055)	(0.0059)	(0.0064)
ln_price		0.1890***	0.0717***	0.0273***
		(0.0046)	(0.0080)	(0.0103)
Observations	25721	25721	18591	14533
R^2	0.005	0.071	0.005	0.001

Table A.1.3: Transaction-level decomposition results.

Even though this regression accounts for the changes in means, it only compares the means before and after 2013. However, this effect might be due to the fact that the entire price distribution shifted to the right during this period, and if "high" and "low" nominal prices are systematically different (for example, there are more round prices among high prices) this alone can be and underlying factor that drives the results. To control for that, I include log price as a control in columns 2 of table A.1.3: this would allow me to compare transactions at or around the same nominal price before and after 2013. In this specification the after 2013 dummy is no longer significant. From this result alone however I cannot conclude that it is the expensive listings that drive the results: in the specifications where I restrict the sample to transactions below 20 and 15 million forints (columns 3 and 4, respectively) the after variable is significant again, suggesting that round prices became more common in the lower segment of the market,

10.14754/CEU.2022.03

even after controlling for the price level. Note that these cutoffs, while somewhat arbitrary, are not exactly low: they are well above the median price both before and after 2013.

A.2 Changes by building type

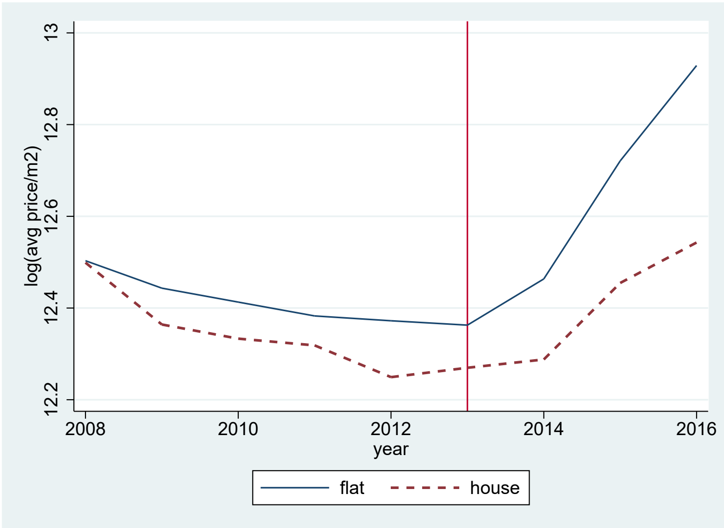

Figure A.2.1: Price change for single-family houses and flats.

Figure A.2.2: Change in the share of round prices over time by for house and flats. (Normalized by the year 2013)

10.14754/CEU.2022.03

A.3 Proofs

A.3.1 Proof for Proposition 1 (optimal pricing if the seller is biased)

Proof. 1. First, I show that the optimal price is nondecreasing in x. The utility can be written as

$$U(p;x) = (x - (1-\theta)p - \theta\lfloor p \rfloor)p \equiv x \cdot p + f(p),$$

where $f(p)$ does not depend on x. Suppose for some x_0, p_0 is the price that maximizes U. For $x_1 > x_0$, can a $p_1 < p_0$ be optimal? No, because at (x_1, p_1) the utility becomes

$$x_1 \cdot p_1 + f(p_1) = (x_0 \cdot p_1 + f(p_1)) + (x_1 - x_0)p_1.$$

The first term is equivalent to $U(p_1; x_0)$, which is less than or equal to $U(p_0; x_0)$ by definition. The second term can also be strictly increased by choosing p_0 over p_1 (as $x_1 - x_0 > 0$ and $p_1 < p_0$).

So for an $x_1 > x_0$ the optimal price p_1 has to be at least as great as p_0.

2. If $\theta = 0$, then optimal price is $\frac{x}{2}$, so if $x = 2k$, the optimal price is still k.

At $p = k$, the incentive of the seller to raise prices a little (not to exceed $k+1$) is given by

$$\frac{\partial}{\partial p}((x-p)\cdot(\theta k + (1-\theta)p)) = (x-2k)(1-\theta) - \theta k \tag{A.2}$$

Equating it with 0 and solving for the x yields

$$x_H(k) = 2k + \frac{\theta}{1-\theta}k \tag{A.3}$$

Clearly, $x_H(k) < 2k+2$ has to hold (for an interior optimum), since we know that for $x = 2k+2$ the optimal price is $k+1$. This means that $\frac{\theta}{1-\theta}k < 2$, which, using the fact that $x = 2k$, is equivalent to

$$x < \frac{1-\theta}{\theta} \tag{A.4}$$

If this condition does not hold (the maximum valuation is high enough), the seller will set only round prices.

As for $x_L(k)$, at $p = k$ the seller compares the utilities from charging an optimal interior price $p \in [k-1, k)$ which comes from an FOC and the utility of staying at $p = k$.

The internal optimal price for the FOC is

$$\frac{x}{2} + \frac{\theta}{2(1-\theta)}(k-1), \tag{A.5}$$

substituting it in the utility function $(1-\theta)(x-p)p + \theta(x-p)\lfloor p \rfloor$ and using that $k-1 \leq p \leq k$ should hold gives

$$(1-\theta)\left(\frac{x}{2} - A\right)\left(\frac{x}{2} + A\right) + \theta\left(\frac{x}{2} - A\right)(k-1), \tag{A.6}$$

where $A = \frac{\theta/2}{1-\theta}(k-1)$.

The utility of $p = k$ is

$$(x-k)k \tag{A.7}$$

Equating the two utilities and solving for x gives candidates for $x_L(k)$

$$\frac{k(2-\theta)+\theta \pm 2\sqrt{k\theta}}{1-\theta} \tag{A.8}$$

Since this is a candidate for a lower bound, $x_H - x_L \geq 0$ must hold. However, comparing $x_H(k) = k\left(2+\frac{\theta}{1-\theta}\right) = \frac{k(2-\theta)}{1-\theta}$ and the greater root clearly shows that only the smaller root can possibly be $x_L(k)$, giving

$$x_L(k) = \frac{k(2-\theta)+\theta - 2\sqrt{k\theta}}{1-\theta}. \tag{A.9}$$

Comparative static:

$$x_H - x_L = \frac{2\sqrt{k\theta}-\theta}{1-\theta}, \quad \frac{\partial}{\partial k}(x_H - x_L) > 0, \quad \frac{\partial}{\partial \theta}(x_H - x_L) = (1-\theta)\sqrt{\frac{k}{\theta}} + (2\sqrt{k\theta}-1) > 0$$

and

$$\frac{\partial}{\partial k}(x_H(k) - 2k) = \frac{\theta}{1-\theta} > 0, \quad \frac{\partial}{\partial k}(2k - x_L(k)) = \frac{1}{1-\theta}\left(\sqrt{\frac{\theta}{k}} - \theta\right) > 0,$$

where the last inequality follows from the fact that $\sqrt{\frac{\theta}{k}} - \theta > 0 \iff \frac{1}{\theta} > k$, which always holds whenever at least some nonround prices are charged (that is, when according to equation A.4, $x < \frac{1-\theta}{\theta} \iff k < \frac{1}{\theta} \cdot \frac{1}{2}(1-\theta)$)

□

A.3.2 Proof for Proposition 2 (optimal pricing if the buyer is biased)

Proof. 1. As in the proof of Proposition 1, I start with showing that the optimal price is nondecreasing in x. Redefining

$$U(p;x) = x \cdot p - p((1-\theta)p + \theta\lfloor p \rfloor) \equiv x \cdot p + f(p),$$

and using the argument in the proof Proposition 1 gives the desired result.

2. Call the price that maximizes $U(p;\theta=0)$ the unconstrained optimum. I show that if the unconstrained optimum is a positive integer k, then for any positive $1 > \theta$ the optimal price is $k - \epsilon$.

   In general, the unconstrained optimum is $\frac{x}{2}$, so if it is equal to some k, then $x = 2k$.

   Define

   $$\bar{U}(k) = \lim_{\epsilon \to 0} U(k-\epsilon) = \lim_{\epsilon \to 0}(x - (1-\theta)(k-\epsilon) - \theta\lfloor k-\epsilon \rfloor)(k-\epsilon) = (x - (1-\theta)k - \theta(k-1))k$$

   Clearly, $\bar{U}(k) > U(k)$.

   First I show that on the interval $(k-1,k)$ the utility is strictly increasing in price. Here $U'(p) = x - (2(1-\theta)p - \theta(k-1))$. Plugging in $x = 2k$ yields

   $$2k + \theta(k-1) - 2(1-\theta)p,$$

   which on this interval is strictly greater than $2k + \theta(k-1) - 2(1-\theta)k = \theta(2k-1) > 0$. So every internal price p is dominated by some $p' > p$ in $(k-1,k)$.

This generalizes to all positive integers below k on the interval $[k-\Delta-1, k-\Delta)$, $U'(p) > 0$, for all integer Δ's with $1 \leq \Delta \leq k$. This means that there is no interior optimum candidate below k and since the the utility is piecewise increasing, it is enough to compare utilities $\bar{U}(n)$ with $0 \leq n \leq k$, among which $\bar{U}(k)$ is the largest (which can be seen by calculating $\bar{U}(k) - \bar{U}(k-\Delta) = 2\Delta\theta + 2k(1-\theta) + 1 > 0$).

Next, I will argue that there are no optimum candidates above k, that is, in the interiors of the intervals $[k+\Delta, k+\Delta+1)$. Let's examine the marginal utility of a price increase on the interval $[k+\Delta, k+\Delta+1)$, where Δ is an integer in $[0, k]$:

$$U'(p) = 2k + \theta(k+\Delta) - 2(1-\theta)p.$$

If on some interval U' is strictly positive or negative, it will be sufficient to examine the utility at the integer endpoint and the limit $\bar{U}$ at the outer bound of interval, respectively, as they clearly dominate any price in the interior of the interval.

If $U' < 0$ on some interval $[k+\Delta, k+\Delta+1)$, we have to compare $\bar{U}(k)$ and $U(k+\Delta)$: $\bar{U}(k) > U(k) > U(k+\Delta)$ as $\bar{U}(k) > U(k)$ for any positive integer and at the integers the utility collapses to $(x-p)p$, which is maximized at k.

If $U' > 0$ on some interval $[k+\Delta, k+\Delta+1)$, we have to compare $\bar{U}(k)$ and $\bar{U}(k+\Delta+1)$. $\bar{U}(k) > \bar{U}(k+\Delta+1) \iff \Delta^2 + 2\Delta + 1 > \Delta\theta + \theta$, which holds term by term.

If $U'(p) = 0$ for some interior price, the FOC holds and potentially we can have an optimum that dominates $k-\epsilon$. Solving for this candidate price yields $p_{foc} = k + \frac{\theta}{2(1-\theta)}(3k+\Delta)$. Comparing this with the conjectured price $k-\epsilon$:

$$U(p_{foc}) - \bar{U}(k) = \frac{\theta}{4(1-\theta)}(-3\Delta^2\theta - 6\Delta k\theta - 4\Delta k - 3k^2\theta + 4k\theta + 4k)$$

$-3\Delta^2\theta - 3k^2\theta$ is clearly negative, the rest in the brackets is negative if

$$6k\Delta\theta > 4k(\Delta - \theta - 1)$$

For $\theta \geq 2/3$, this always holds. For $\theta < 2/3$, it holds only for $\Delta > 0$. But it is easy to check that for $\Delta = 0$ the inequality is always satisfied, irrespective of θ.

To conclude, p_{foc} is always dominated by $k-\epsilon$, therefore the unique optimum is $k-\epsilon$.

3. Given that for some even x the optimal price is $\frac{x}{2} - \epsilon$, it follows that for $x < x' < x+1$ the optimal price p' has to be in $[\frac{x}{2} - \epsilon, \frac{x}{2} + 1 - \epsilon]$. As the optimal price is weakly monotonic in x, $p' \geq \frac{x}{2} - \epsilon$, which sets the lower bound of this interval. The upper bound comes from the fact that for $x+2$ the optimal price is $\frac{x}{2} + 1 - \epsilon$ and obviously $x+2 > x$.

   To put differently, for any positive x, the constrained price is in

$$\left[\lfloor\frac{x}{2}\rfloor - \epsilon, \lfloor\frac{x}{2}\rfloor + 1 - \epsilon\right]$$

4. Next, I characterize the schedule of optimal prices as a function of x. Let $[x_L(k), x_H(k)]$ be the set of x's such that the seller will find it optimal to set the integer price k. Clearly, $x = 2k$ is in this set, but $2k-2$ and $2k+2$ are not.

I start with the assumption than when x just below (above) the interval $[x_L(k), x_H(k)]$ the seller does not jump all the way down to $k-1$ (up to $k+1$) but rather switches to a price somewhere between $k-1$ and k (k and $k+1$), because of the monotonicity of the optimal price in x. Later I will give a (necessary) condition for this assumption to hold.

Start with x_L. Consider $k-1 < p < k$. If the FOC holds, $U(p)$ is maximized at $p^* = \frac{x-(k+1)\theta}{2(1-\theta)}$. Solving for the x that equates $U(p^*)$ and $\bar{U}(k)$ yields

$$x_L(k) = k(2-\theta) - \theta$$

This is the candidate for $x_L(k)$.

For x_H, consider $k < p < k+1$. If the FOC holds, $U(p)$ is maximized at $p^* = \frac{x-k\theta}{2(1-\theta)}$. Solving the quadratic equation $U(p^*) = \bar{U}(k)$ for x yields $k(2-\theta) \pm 2\sqrt{k\theta(1-\theta)}$. Since $x_H \geq x_L$ has to be true and $2\sqrt{k\theta(1-\theta)} > \theta$ for $k \geq 1$ and $\theta \in (0,1)$, I pick

$$x_H(k) = k(2-\theta) + 2\sqrt{k\theta(1-\theta)}.$$

For the assumption that there are some non-integer prices between k and $k+1$ to hold,

$$x_H(k) < x_L(k+1)$$

has to be satisfied. This is equivalent to

$$k < \frac{1-\theta}{\theta}$$

(Comparing $x_H(k) < 2k+2$ and $x_L(k) > 2k-2$ are not useful here).

Comparative static:

$$x_H - x_L = 2\sqrt{k\theta(1-\theta)} + \theta, \quad \frac{\partial}{\partial k}(x_H - x_L) > 0, \quad \frac{\partial}{\partial \theta}(x_H - x_L) = 1 + \sqrt{\frac{k}{\theta(1-\theta)}}(1-2\theta)$$

and

$$\frac{\partial}{\partial k}(2k - x_L(k)) = \theta > 0, \quad \frac{\partial}{\partial k}(x_H(k) - 2k) = \sqrt{\frac{\theta(1-\theta)}{k}} - \theta > 0,$$

where the last line holds whenever at least some nonround prices are quoted ($k < (1-\theta)/\theta$).

It easy to check that $\frac{\partial}{\partial k}(2k - x_L(k)) > \left|\frac{\partial}{\partial k}(x_H(k) - 2k)\right|$ always holds.

□

10.14754/CEU.2022.03

10.14754/CEU.2022.03

Appendix B

Appendix for chapter 2

B.1 Additional results for the main specification

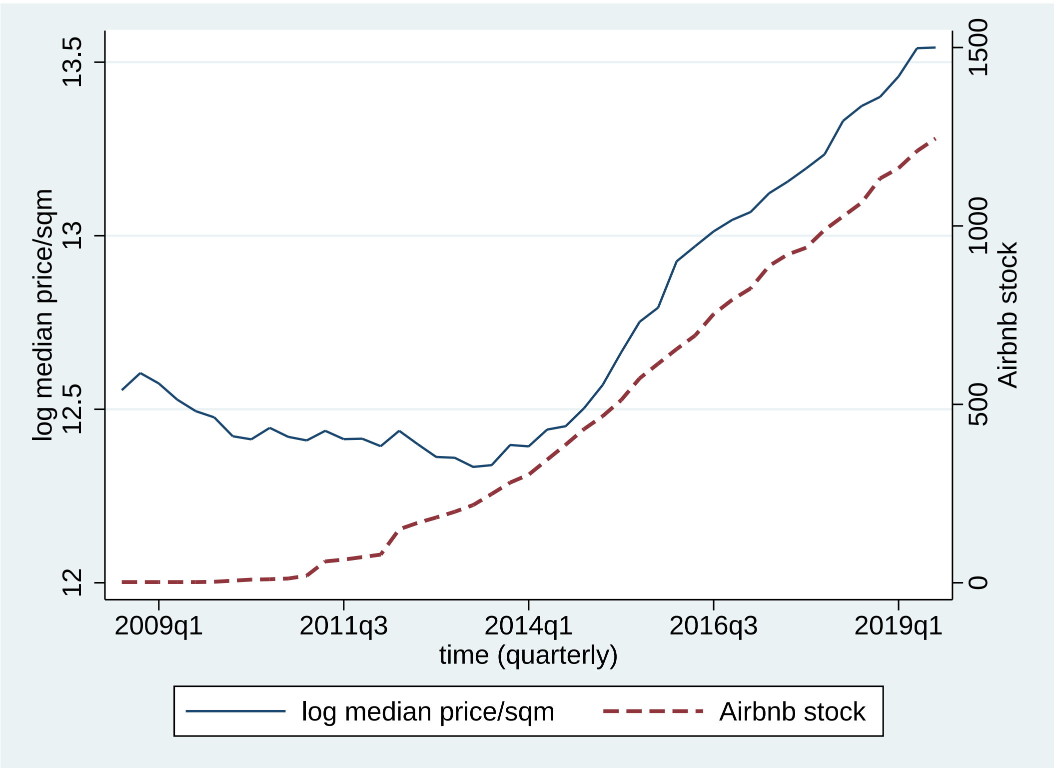

Figure B.1.1: The evolution of prices (left axis) and Airbnb supply (right axis) in the treatment and control group.

10.14754/CEU.2022.03

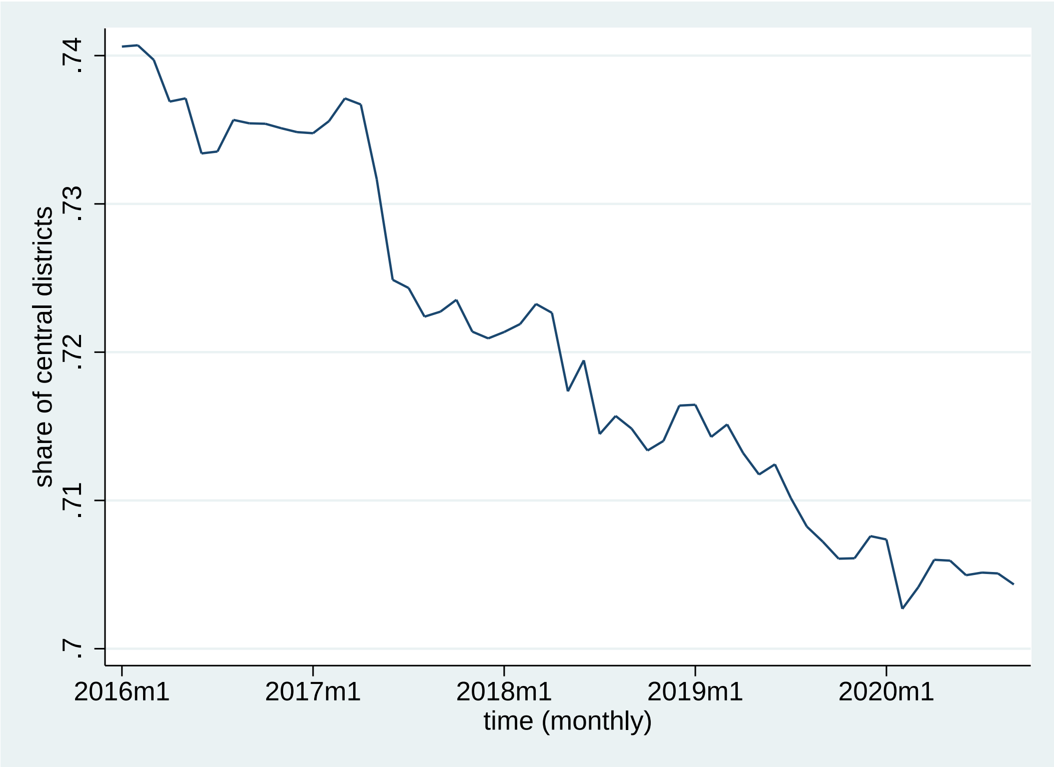

Figure B.1.2: The share of central districts among the total stock of entrants, Airbnb data.

	(1)	
	Airbnb stock (normalized), 2017m12=1	
treat x after x trend	-0.00139***	(-7.52)
after x trend	0.000738***	(4.34)
treat x time	-0.000521***	(-3.79)
time	0.0149***	(56.63)
treat	0.0325***	(6.27)
Constant	0.127***	(12.03)
Observations	185	
R^2	0.996	

t statistics in parentheses

Driscoll-Kraay standard errors

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

Table B.1.1: Regression result for the evolution of the cumulative number of Airbnb entrants for the treatment and the control group with linear trends.

10.14754/CEU.2022.03

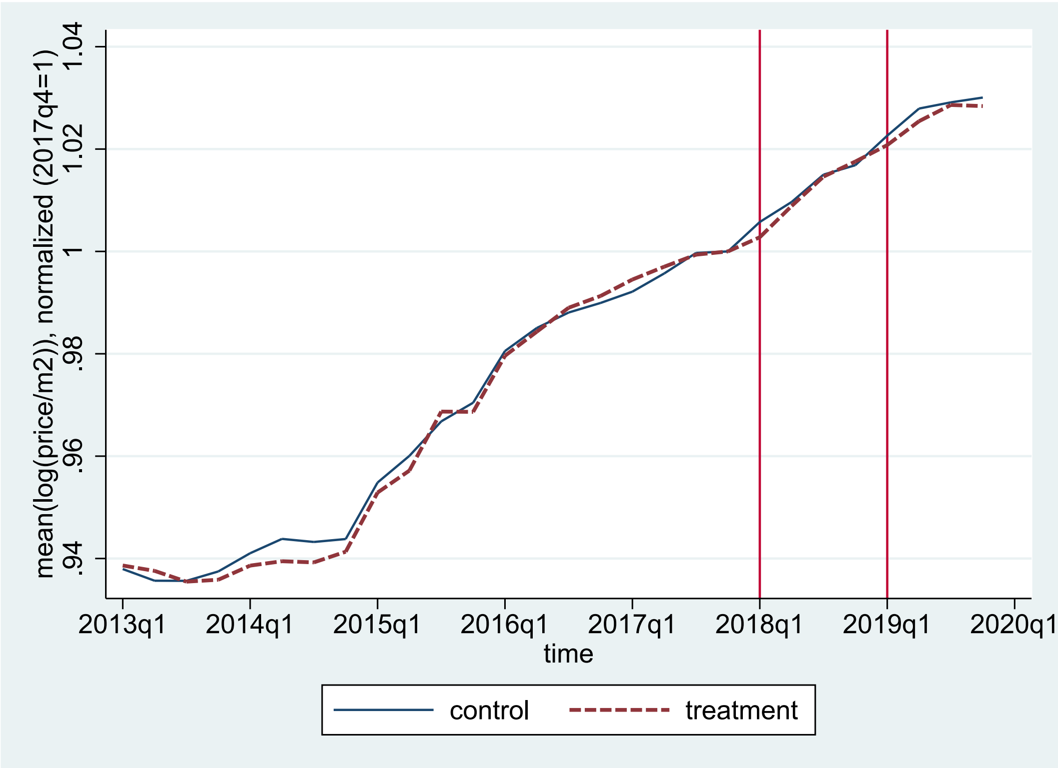

Figure B.1.3: The evolution of the average log price per square meter in the treatment/control group, normalized.

10.14754/CEU.2022.03

	all	small	large
treated	-0.3184***	-0.4576***	-0.1174*
	(0.0641)	(0.1338)	(0.0591)
	[0.0000]	[0.0007]	[0.0478]
treat x 16Q2	-0.0037	-0.0110	0.0004
	(0.0283)	(0.0352)	(0.0435)
	[0.8957]	[0.7538]	[0.9920]
treat x 16Q3	0.0262	0.0456	0.0057
	(0.0325)	(0.0439)	(0.0360)
	[0.4206]	[0.2992]	[0.8743]
treat x 16Q4	0.0309	0.0086	0.0554
	(0.0287)	(0.0354)	(0.0368)
	[0.2808]	[0.8077]	[0.1328]
treat x 17Q1	0.0158	0.0082	0.0314
	(0.0342)	(0.0437)	(0.0405)
	[0.6451]	[0.8517]	[0.4391]
treat x 17Q2	0.0211	0.0548	-0.0234
	(0.0301)	(0.0363)	(0.0387)
	[0.4844]	[0.1314]	[0.5456]
treat x 17Q3	0.0126	0.0425	-0.0159
	(0.0302)	(0.0403)	(0.0369)
	[0.6771]	[0.2922]	[0.6679]
treat x 17Q4	0.0076	0.0439	-0.0372
	(0.0306)	(0.0371)	(0.0390)
	[0.8030]	[0.2379]	[0.3400]
treat x 18Q1	-0.0272	-0.0058	-0.0551
	(0.0321)	(0.0375)	(0.0428)
	[0.3972]	[0.8772]	[0.1990]
treat x 18Q2	-0.0148	-0.0024	-0.0306
	(0.0267)	(0.0334)	(0.0372)
	[0.5794]	[0.9431]	[0.4116]
treat x 18Q3	0.0043	-0.0066	0.0090
	(0.0269)	(0.0325)	(0.0384)
	[0.8731]	[0.8382]	[0.8140]
treat x 18Q4	0.0042	0.0032	0.0082
	(0.0285)	(0.0371)	(0.0383)
	[0.8840]	[0.9316]	[0.8311]
r2	0.422	0.455	0.387
N	15606	8120	7486

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table B.1.2: Quarterly diff-in-diff estimates, by size ("small": < 50, "large": ≥ 50 sq.meter)

10.14754/CEU.2022.03

	(1)	(2)	(3)	(4)	(5)
	log(persexp)	log(sales '18)	log(pretax profit)	log(pers. exp./emp.)	log(employment)
year 2014	0.118	0.088	0.129	0.157**	0.009
	(0.09)	(0.09)	(0.21)	(0.05)	(0.05)
year 2015	0.186	0.231	0.556*	0.116	0.145
	(0.11)	(0.12)	(0.22)	(0.07)	(0.08)
year 2016	0.390**	0.292*	0.815***	0.185**	0.251**
	(0.12)	(0.14)	(0.23)	(0.07)	(0.09)
year 2017	0.514***	0.460**	1.119***	0.299***	0.175*
	(0.13)	(0.14)	(0.24)	(0.08)	(0.09)
year 2018	0.828***	0.827***	1.543***	0.451***	0.258**
	(0.13)	(0.14)	(0.25)	(0.08)	(0.09)
treat x 2014	-0.154	0.189	0.392	-0.282*	0.116
	(0.15)	(0.15)	(0.40)	(0.11)	(0.10)
treat x 2015	-0.020	0.229	0.034	-0.082	0.062
	(0.18)	(0.20)	(0.43)	(0.13)	(0.16)
treat x 2016	-0.005	0.311	-0.493	-0.007	-0.021
	(0.19)	(0.21)	(0.46)	(0.13)	(0.16)
treat x 2017	0.009	0.307	-0.184	-0.061	0.162
	(0.20)	(0.24)	(0.43)	(0.11)	(0.17)
treat x 2018	0.066	0.234	-0.286	-0.165	0.285
	(0.20)	(0.23)	(0.40)	(0.11)	(0.16)
Constant	8.744***	10.273***	7.904***	7.314***	1.693***
	(0.07)	(0.08)	(0.16)	(0.04)	(0.05)
N	1208.000	1278.000	838.000	1094.000	1119.000
r2	0.118	0.107	0.160	0.083	0.060

Table B.1.3: Effect on local business, NACE I55 only (Accomodation)

	(1) log(persexp)	(2) log(sales '18)	(3) log(pretax profit)	(4) log(pers. exp./emp.)	(5) log(employment)
year 2014	0.044	0.053	0.214*	0.056*	-0.038
	(0.04)	(0.04)	(0.08)	(0.03)	(0.02)
year 2015	0.097*	0.143**	0.347***	0.102***	0.043
	(0.04)	(0.05)	(0.09)	(0.03)	(0.03)
year 2016	0.290***	0.285***	0.490***	0.228***	0.049
	(0.04)	(0.05)	(0.10)	(0.03)	(0.03)
year 2017	0.562***	0.468***	0.910***	0.384***	0.098**
	(0.05)	(0.05)	(0.11)	(0.03)	(0.03)
year 2018	0.838***	0.758***	1.553***	0.519***	0.150***
	(0.05)	(0.05)	(0.11)	(0.03)	(0.03)
treat x 2014	0.038	0.095	0.031	-0.056	0.128*
	(0.08)	(0.09)	(0.18)	(0.05)	(0.05)
treat x 2015	0.036	0.066	-0.054	-0.046	0.052
	(0.09)	(0.10)	(0.19)	(0.07)	(0.07)
treat x 2016	0.083	0.087	0.138	0.001	0.087
	(0.09)	(0.11)	(0.19)	(0.06)	(0.07)
treat x 2017	0.049	0.109	0.059	0.019	0.078
	(0.10)	(0.11)	(0.20)	(0.07)	(0.08)
treat x 2018	0.027	0.064	0.119	0.030	0.082
	(0.11)	(0.12)	(0.21)	(0.07)	(0.08)
Constant	8.321***	9.853***	6.770***	6.965***	1.528***
	(0.03)	(0.03)	(0.07)	(0.02)	(0.02)
N	6862.000	6944.000	3858.000	6368.000	6554.000
r2	0.131	0.094	0.180	0.116	0.017

Table B.1.4: Effect on local business, NACE I56 only (Food and beverage activities)

	(1)	(2)	(3)	(4)	(5)
	log(1+persexp)	log(1+sales '18)	log(1+pretax profit)	log(pers. exp./emp.)	employment
year 2014	-0.152*	-0.121	0.110	0.071**	0.096
	(0.06)	(0.08)	(0.10)	(0.02)	(0.25)
year 2015	-0.552***	-0.372***	0.355***	0.104***	2.129***
	(0.09)	(0.10)	(0.11)	(0.03)	(0.55)
year 2016	0.022	0.065	0.556***	0.221***	3.002**
	(0.08)	(0.10)	(0.11)	(0.03)	(0.95)
year 2017	0.365***	0.407***	1.002***	0.369***	4.464***
	(0.09)	(0.11)	(0.12)	(0.03)	(1.27)
year 2018	0.807***	0.908***	1.753***	0.507***	5.481**
	(0.09)	(0.12)	(0.12)	(0.03)	(1.77)
treat x 2014	-0.153	-0.035	-0.009	-0.086	1.654*
	(0.16)	(0.20)	(0.21)	(0.05)	(0.66)
treat x 2015	0.179	0.341	0.061	-0.051	0.791
	(0.18)	(0.21)	(0.20)	(0.06)	(0.90)
treat x 2016	0.183	0.257	0.183	0.001	0.456
	(0.18)	(0.21)	(0.20)	(0.06)	(1.15)
treat x 2017	0.005	-0.042	0.087	0.013	-0.955
	(0.18)	(0.22)	(0.21)	(0.06)	(1.40)
treat x 2018	0.073	0.268	0.005	0.007	-0.910
	(0.19)	(0.23)	(0.22)	(0.06)	(1.90)
Constant	7.439***	8.656***	6.284***	7.017***	12.016***
	(0.05)	(0.06)	(0.07)	(0.02)	(0.61)
N	9361.000	9545.000	5150.000	7462.000	7673.000
r2	0.066	0.043	0.148	0.110	0.018

Table B.1.5: Local businesses, NACE I55 and I56 combined. Outcomes are defined as $\log(1 + x)$.

10.14754/CEU.2022.03

	(1)	(2)	(3)	(4)	(5)
	log(1+persexp)	log(1+sales '18)	log(1+pretax profit)	log(pers. exp./emp.)	employment
year 2014	-0.076	-0.084	-0.133	0.157**	-0.720
	(0.13)	(0.18)	(0.26)	(0.05)	(0.83)
year 2015	-0.660**	-0.384	0.217	0.116	3.489*
	(0.22)	(0.25)	(0.27)	(0.07)	(1.58)
year 2016	-0.080	-0.272	0.551*	0.185**	3.903
	(0.21)	(0.27)	(0.27)	(0.07)	(2.00)
year 2017	0.383	0.290	1.066***	0.299***	5.790*
	(0.21)	(0.25)	(0.28)	(0.08)	(2.53)
year 2018	0.704**	0.861**	1.650***	0.451***	6.541*
	(0.26)	(0.29)	(0.31)	(0.08)	(2.65)
treat x 2014	0.050	0.896*	0.208	-0.282*	8.244*
	(0.35)	(0.43)	(0.51)	(0.11)	(4.14)
treat x 2015	0.638	1.249*	0.222	-0.082	3.859
	(0.48)	(0.53)	(0.53)	(0.13)	(4.76)
treat x 2016	0.600	1.378**	-0.178	-0.007	3.695
	(0.44)	(0.42)	(0.50)	(0.13)	(4.14)
treat x 2017	0.094	0.623	-0.451	-0.061	1.102
	(0.42)	(0.45)	(0.49)	(0.11)	(4.25)
treat x 2018	0.501	1.426*	-0.358	-0.165	1.152
	(0.49)	(0.61)	(0.51)	(0.11)	(4.43)
Constant	7.256***	8.651***	7.399***	7.314***	33.543***
	(0.12)	(0.15)	(0.19)	(0.04)	(1.10)
N	1492.000	1524.000	906.000	1094.000	1119.000
r2	0.071	0.066	0.135	0.083	0.057

Table B.1.6: Effect on local business, NACE I55 only Outcomes are defined as $\log(1 + x)$.

	(1)	(2)	(3)	(4)	(5)
	log(1+persexp)	log(1+sales '18)	log(1+pretax profit)	log(pers. exp./emp.)	employment
year 2014	-0.167*	-0.129	0.159	0.056*	0.236
	(0.07)	(0.09)	(0.11)	(0.03)	(0.25)
year 2015	-0.532***	-0.372***	0.383***	0.102***	1.887**
	(0.10)	(0.11)	(0.11)	(0.03)	(0.58)
year 2016	0.041	0.128	0.559***	0.228***	2.832**
	(0.09)	(0.11)	(0.12)	(0.03)	(1.07)
year 2017	0.359***	0.429***	0.987***	0.384***	4.214**
	(0.10)	(0.12)	(0.13)	(0.03)	(1.43)
year 2018	0.828***	0.916***	1.773***	0.519***	5.287**
	(0.10)	(0.13)	(0.13)	(0.03)	(2.05)
treat x 2014	-0.186	-0.193	-0.044	-0.056	0.625
	(0.18)	(0.22)	(0.22)	(0.05)	(0.39)
treat x 2015	0.099	0.181	0.051	-0.046	0.366
	(0.20)	(0.23)	(0.21)	(0.07)	(0.74)
treat x 2016	0.107	0.046	0.270	0.001	-0.019
	(0.19)	(0.23)	(0.22)	(0.06)	(1.16)
treat x 2017	-0.011	-0.173	0.214	0.019	-1.213
	(0.20)	(0.24)	(0.23)	(0.07)	(1.48)
treat x 2018	-0.007	0.061	0.101	0.030	-1.193
	(0.21)	(0.25)	(0.24)	(0.07)	(2.11)
Constant	7.475***	8.659***	6.043***	6.965***	8.346***
	(0.05)	(0.07)	(0.08)	(0.02)	(0.69)
N	7869.000	8021.000	4244.000	6368.000	6554.000
r2	0.066	0.042	0.152	0.116	0.015

Table B.1.7: Effect on local business, NACE I56 only Outcomes are defined as $\log(1 + x)$.

	(1)	(2)	(3)
after 2018q1	3.605	3.612	3.612
	(3.142)	(3.148)	(3.156)
treated x after 2018q1	-5.494	-5.507	-5.503
	(3.366)	(3.355)	(3.363)
Observations	1602	1602	1586
R^2	0.019	0.023	0.023

Standard errors in parentheses

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

Table B.1.8: Diff-in-diff: number of reviews, high-end price category

	(1)	(2)	(3)
after 2018q1	0.480	0.478	0.477
	(0.364)	(0.364)	(0.364)
treated x after 2018q1	0.491	0.493	0.493
	(0.602)	(0.603)	(0.603)
Observations	24224	24224	23855
R^2	0.005	0.011	0.011

Standard errors in parentheses

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

Table B.1.9: Diff-in-diff: number of reviews, mid-range price category

	(1)	(2)	(3)
after 2018q1	1.193*	1.189*	1.191*
	(0.479)	(0.479)	(0.479)
treated x after 2018q1	0.141	0.139	0.139
	(0.811)	(0.812)	(0.812)
Observations	13047	13047	12699
R^2	0.011	0.016	0.016

Standard errors in parentheses

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

Table B.1.10: Diff-in-diff: number of reviews, inexpensive price category

	local/foreign reviews (normalized), 2017m12=100	
treated	0.00827	(0.783)
time trend	1.964***	(0.0588)
treated x time trend	-0.0518*	(0.0211)
time trend x after 2018q1	-0.358***	(0.0439)
after 2018q1 x treated x time trend	0.0730***	(0.0139)
Constant	-9.002***	(2.416)
Observations	190	
R^2	0.971	

Standard errors in parentheses

Driscoll-Kraay standard errors

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

Table B.1.11: Diff-in-diff regressions with linear time trends for the treated and the control group separately. (Normalization is changed to 100 to save on decimal places)

B.2 Choice of the control group and robustness checks

While the treatment group is clearly defined, the choice of the control group is not trivial, since it has to be large enough that it has enough observations and any possible local substitution effects between the the two groups are "washed out" by more distant areas. However, it cannot be very large as controlling for local characteristics (e.g. local desirability and local demand shocks) can be challenging, even in the presence of rich sets of location-specific fixed effects.

To define the control group, I could rely only on zip codes (see Figure B.2.1). They have the advantage that they are of meaningful size (e.g. district 6, the treatment group is a relatively small district, but it contains 9 zip codes), they are an exclusive and exhaustive partition of administrative districts (no zip code belongs to two or more districts – this is important as the treatment is defined on the district level) and they are readily available in my data.

However their shape is somewhat arbitrary: they are far from convex and consequently it would hard to argue that they capture local characteristics very well in every district. Some of them span a huge distance (e.g. 1138, the northernmost zip code in Figure B.2.1) which can belie underlying differences in local characteristics and degree of substitutability (e.g. the southern half of zip code 1138 includes relatively high prestige locations and flats constructed pre-WW2 but its northern half contains post-WW2 pre-fabricated buildings that are perceived to carry somewhat lower status).

Therefore for the main specification I define the control group as the union of

1. the entirety of neighboring districts 5 and 7,
2. the parts of neighboring district 13 that fall "close" to the treatment group (south of the line in red in the map at Figure B.2.5).

While slicing district 13 in half might seem arbitrary, the boundary (Dózsa György út and Dráva utca) is a natural continuation of the boundary between districts 6, 7 and the neighboring district 14 and serves as a busy route for traffic. Additionally, in district 13 the street Dráva utca represents somewhat of a boundary between buildings constructed before and after World War 2 (south and north from Dráva utca, respectively).

I omit the neighboring district 14 from the control group altogether, even though it borders directly on district 6. The reason for that is while they are indeed geographically close, the zip code that is actually close to the treatment group, 1146 is sparsely populated and its residential neighbourhood is quite far from district 6, as it mostly contains a public park and a major railway station (which accounts for its larger than average size).

To lend some empirical support to these decisions, I calculate a transaction-weighted distance measure for every zip code in Table B.2.1. First, I define the weighted center of the treatment group by taking the average coordinate of all transactions in district 6 in the pre-treatment year. Then I calculate the average distance (and its standard deviation) from this coordinate for every zip code in the same year. This measure has the advantage that it is weighted by transactions so for example it can uncover the fact that while zip code 1146 borders on the treatment group, the average transaction in 1146 falls much farther from district 6's center than in any other zip code of e.g. district 5 – which even has zip codes that share no common border with the treatment group.

In addition to that, I present robustness checks with alternative definitions of the control group for the models with transaction prices on the left hand side. The alternative definitions will be

1. no slicing of district 13: control group will include only those zip codes that they are fully south of the boundary in red

2. omitting southernmost two zip codes of district 5 that do not share a border with the treatment group.

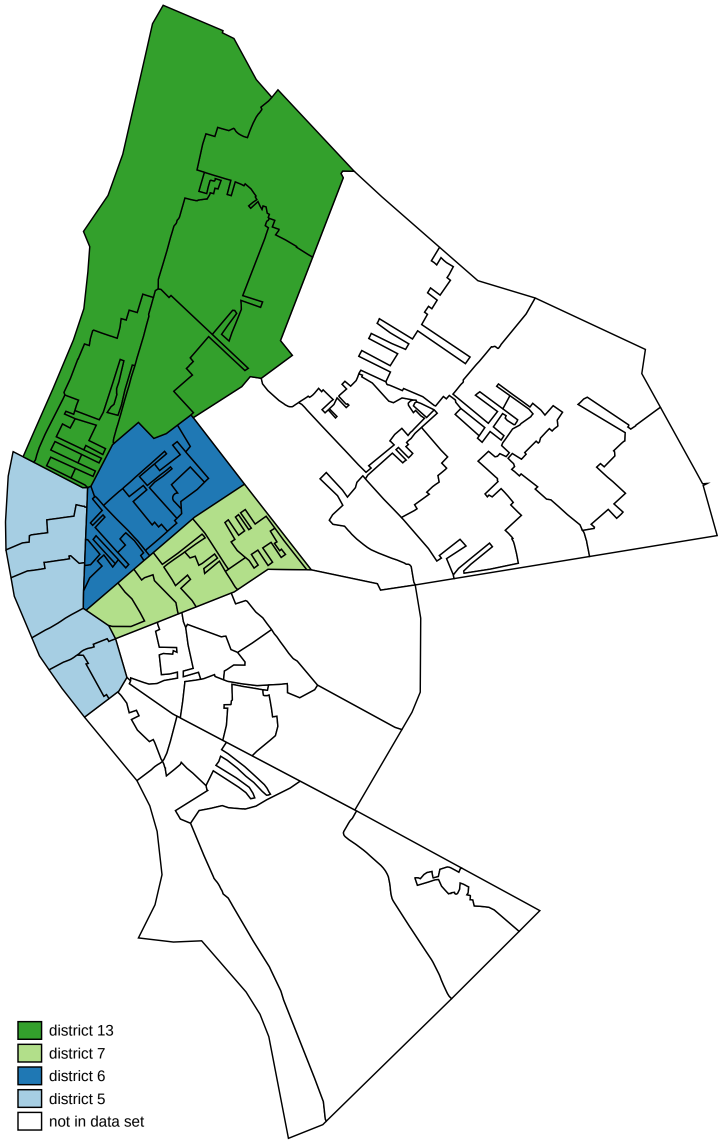

Figure B.2.1: Zip code division of the districts of interest.

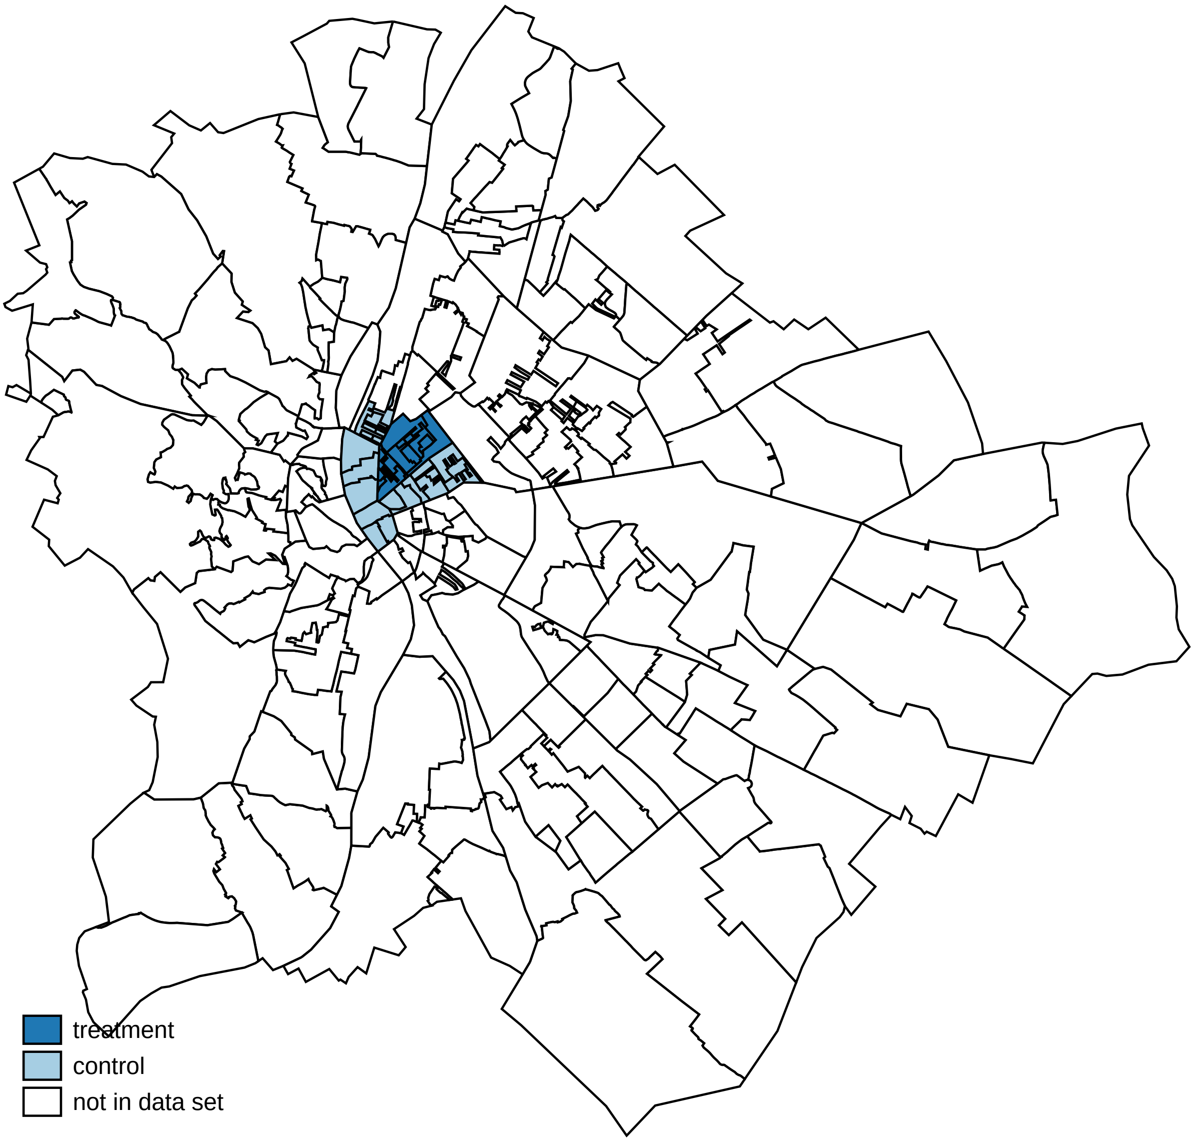

Figure B.2.2: Treatment and control groups, entire Budapest.

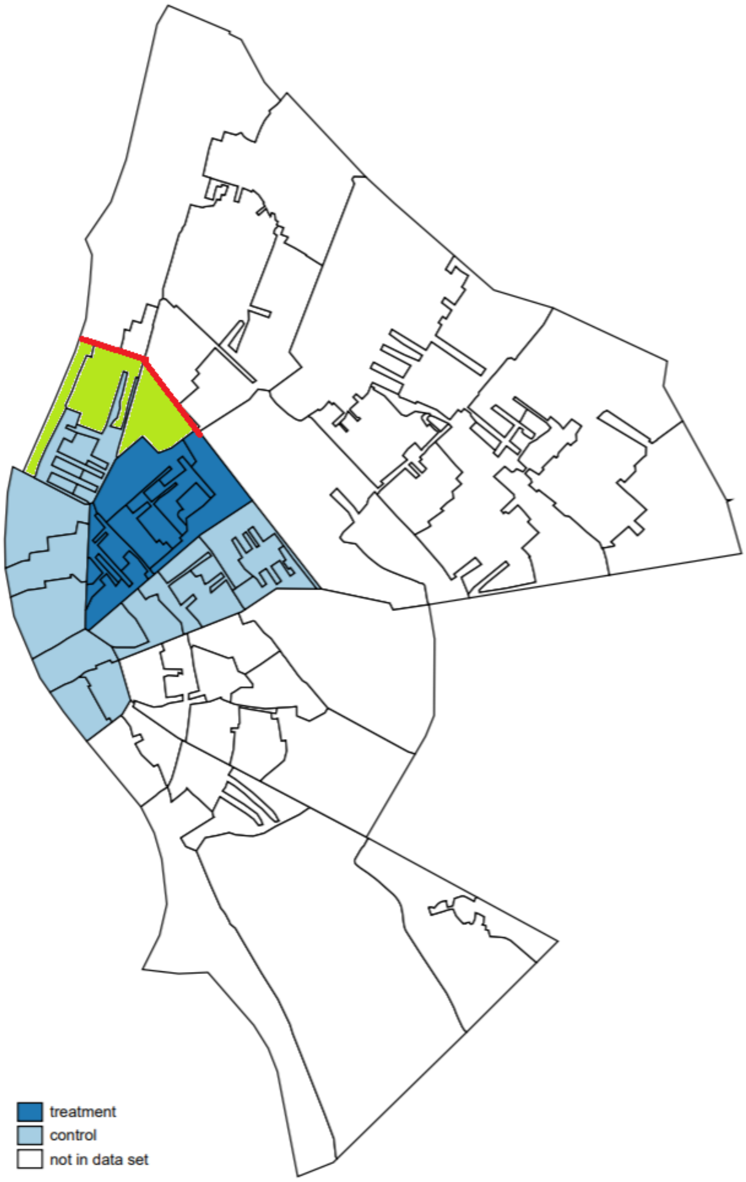

Figure B.2.3: Treatment and control groups, base specification.

10.14754/CEU.2022.03

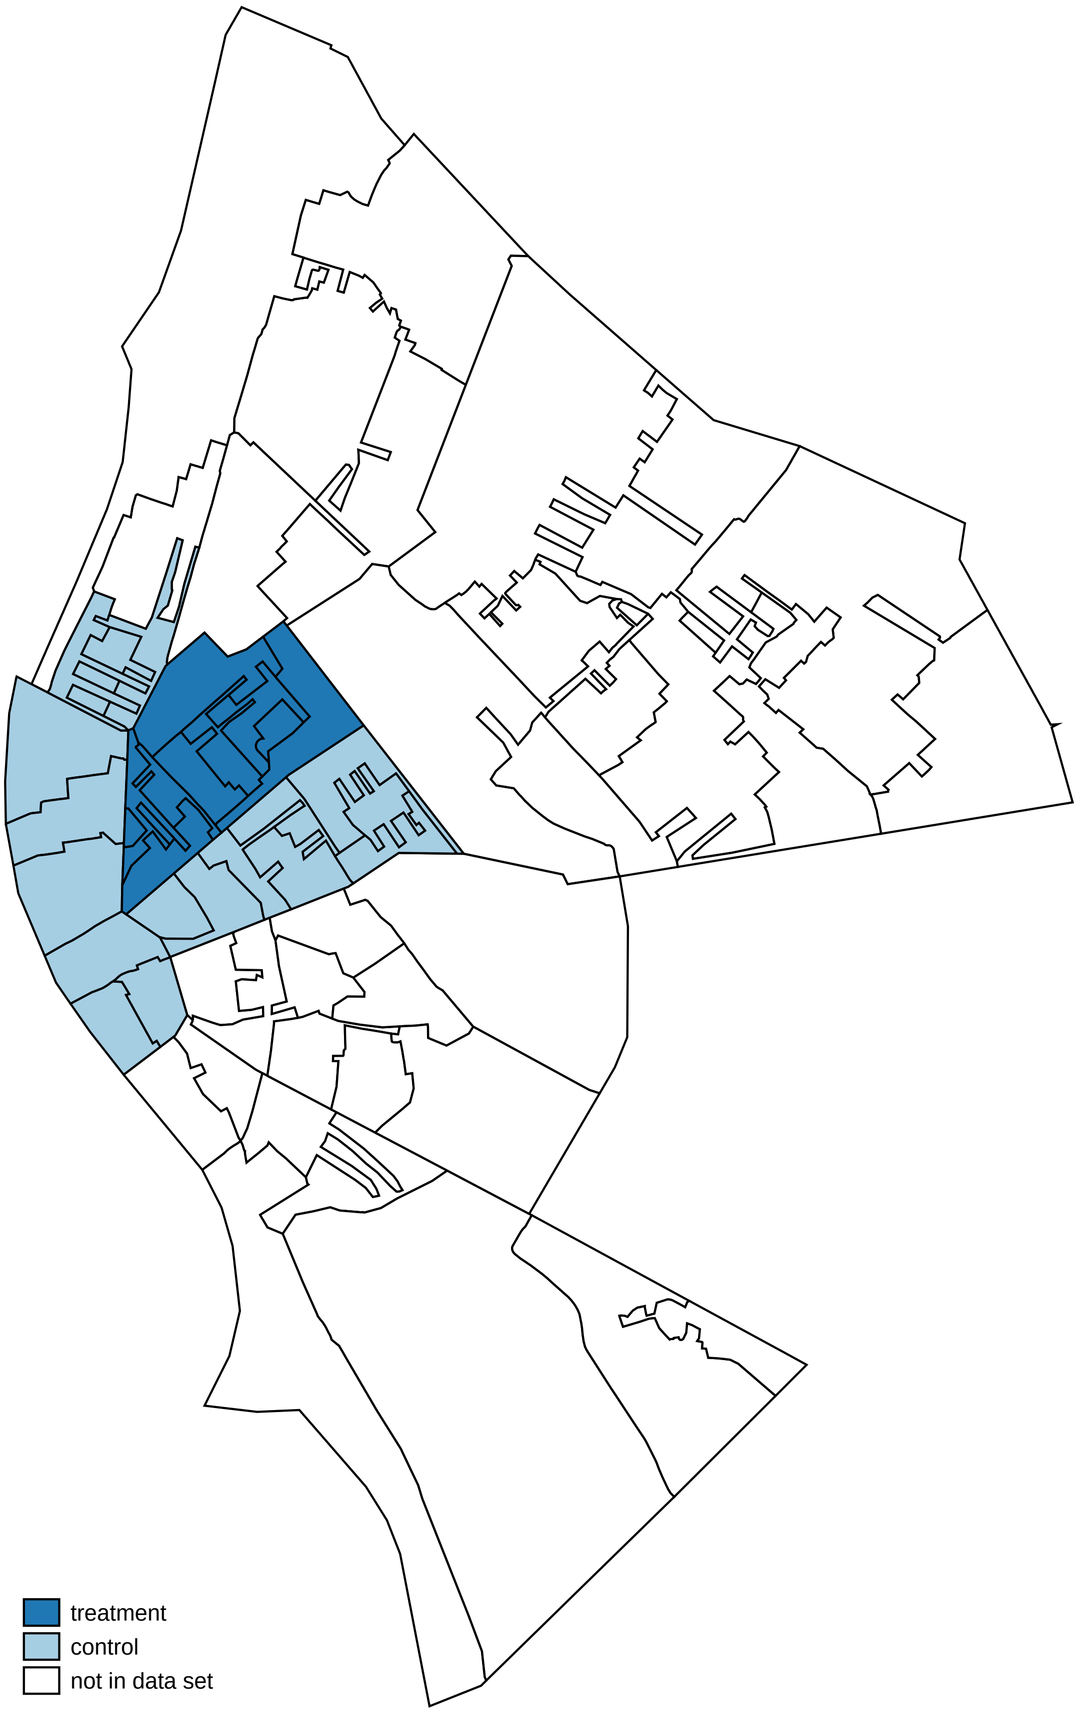

Figure B.2.4: Treatment and control groups, first alternative specification (no slicing of zip codes in district 13.

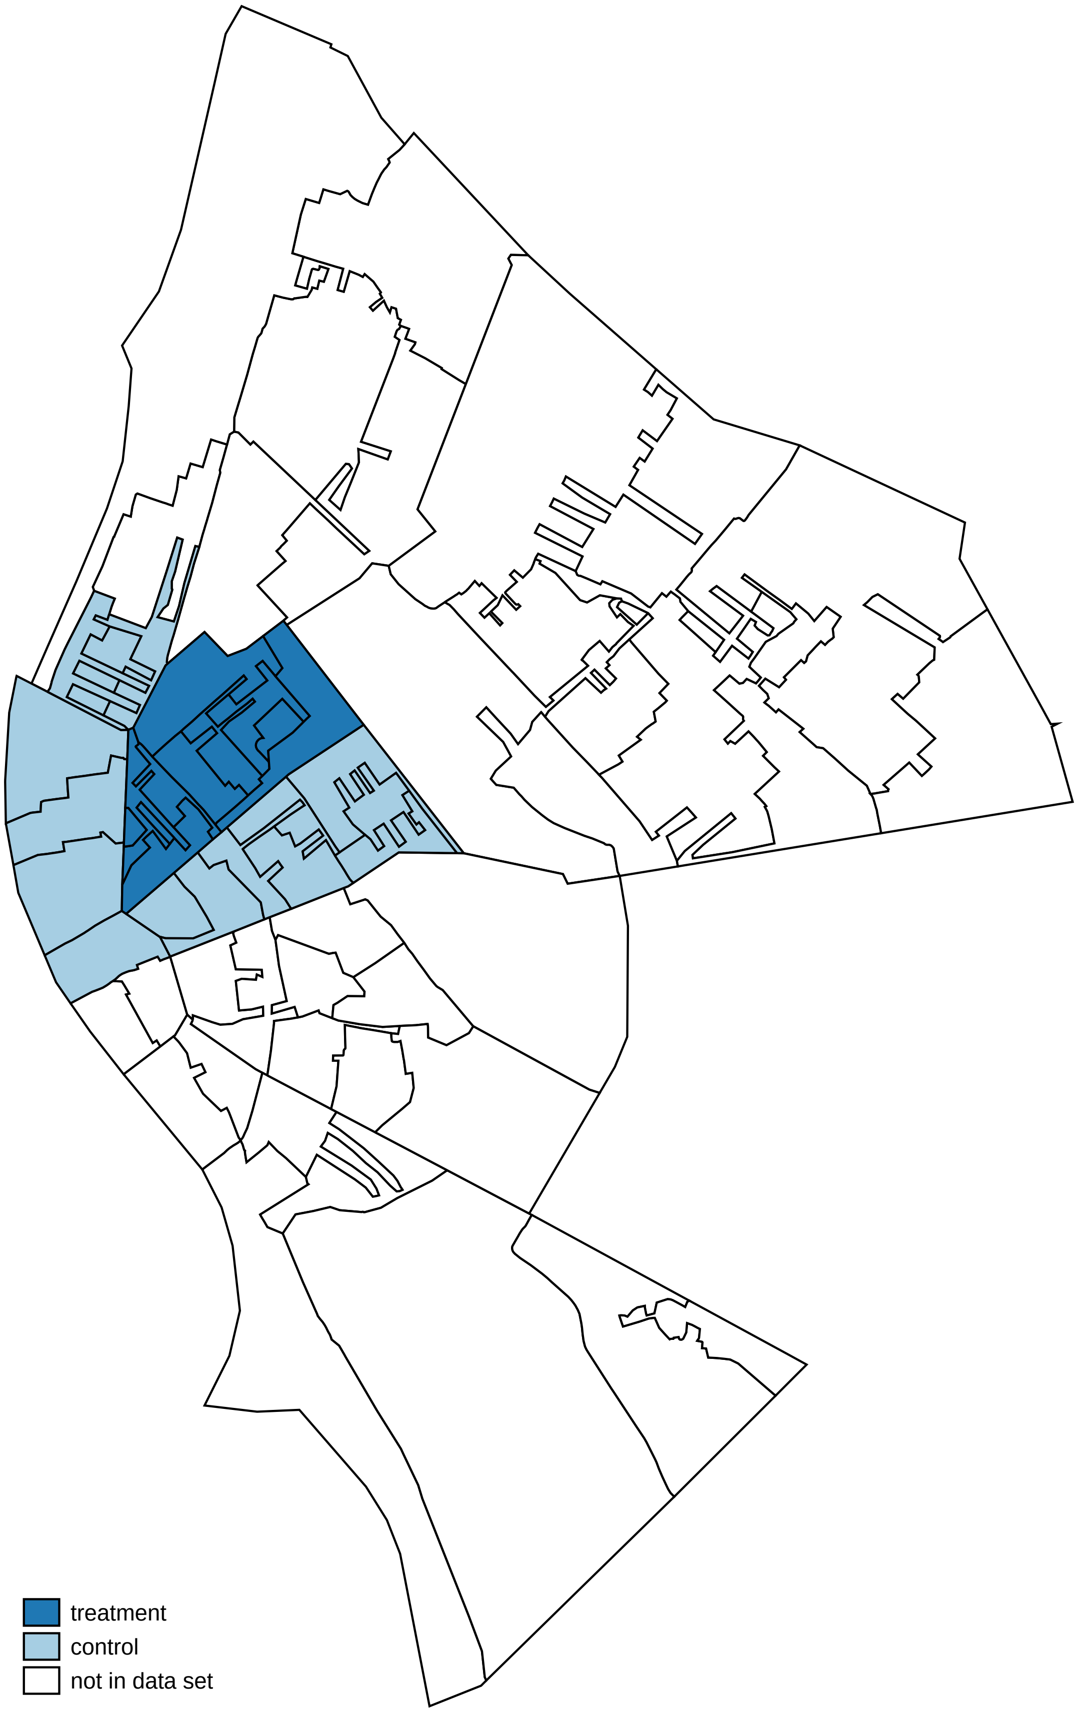

Figure B.2.5: Treatment and control groups, second alternative specification (no slicing of zip codes in district 13, southernmost zip codes in district 5 dropped.

District	zip code	Mean	SD
District 5, control	1051	1300	189
	1052	1638	144
	1053	1861	179
	1054	1037	256
	1055	1094	200
	1056	2116	151
District 6, Treatment	1061	733	278
	1062	609	273
	1063	521	148
	1064	345	158
	1065	672	166
	1066	486	143
	1067	241	111
	1068	768	278
	1069	654	267
District 7, control	1071	1234	293
	1072	1104	183
	1073	817	182
	1074	1051	241
	1075	1180	134
	1076	1642	182
	1077	902	280
	1078	1423	161
District 13, not in control	1131	4244	417
District 13, control	1132	1131	302
District 13, partially in control	1133	1860	424
District 13, partially in control	1134	1767	476
District 13, not in control	1135	2745	466
District 13, control	1136	1189	183
District 13, control	1137	1250	276
District 13, partially in control	1138	3818	924
District 13, not in control	1139	3251	556
District 14, not in control	1141	5200	678
	1142	3754	543
	1143	3007	401
	1144	6190	529
	1145	3194	350
	1146	2414	604
	1147	4250	488
	1148	4676	349
	1149	3872	450
Total		1553	1127

Table B.2.1: Average distance of transactions from the center of the treatment group, pre-treatment year, in meters.

B.2.1 Alternative control group 1: no slicing in district 13

	all	50>	50≤
treated	-0.3174***	-0.4500**	-0.1266**
	(0.0802)	(0.1482)	(0.0389)
	[0.0001]	[0.0026]	[0.0012]
after 2018q1	0.2811***	0.2849***	0.2806***
	(0.0079)	(0.0100)	(0.0102)
	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	-0.0244	-0.0333*	-0.0170
	(0.0132)	(0.0155)	(0.0183)
	[0.0659]	[0.0319]	[0.3528]
r2	0.375	0.411	0.327
N	15039	7981	7058

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table B.2.2: Alternative control group 1: Before-after diff-in-diff results, by size (zip code fixed effects)

	all	50>	50≤
treated	0.0657*	0.1016***	0.0138
	(0.0291)	(0.0211)	(0.0473)
	[0.0258]	[0.0000]	[0.7703]
after 2018q1	0.2811***	0.2816***	0.2854***
	(0.0081)	(0.0086)	(0.0121)
	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	-0.0191	-0.0234	-0.0211
	(0.0127)	(0.0148)	(0.0197)
	[0.1326]	[0.1173]	[0.2874]
r2	0.410	0.462	0.360
N	15039	7981	7058

Includes grid dummies and log(size). SE's clustered on the grid level.

Table B.2.3: Alternative control group 1: Before-after diff-in-diff results, by size (grid fixed effects)

	all	50>	50≤
treated	-0.3178***	-0.4554***	-0.1182*
	(0.0642)	(0.1333)	(0.0590)
	[0.0000]	[0.0007]	[0.0459]
treat x 16Q2	-0.0108	-0.0186	-0.0025
	(0.0285)	(0.0355)	(0.0435)
	[0.7044]	[0.5997]	[0.9534]
treat x 16Q3	0.0288	0.0516	0.0063
	(0.0326)	(0.0439)	(0.0362)
	[0.3769]	[0.2407]	[0.8620]
treat x 16Q4	0.0235	0.0007	0.0489
	(0.0288)	(0.0353)	(0.0369)
	[0.4147]	[0.9835]	[0.1863]
treat x 17Q1	0.0202	0.0164	0.0333
	(0.0342)	(0.0437)	(0.0408)
	[0.5545]	[0.7083]	[0.4143]
treat x 17Q2	0.0191	0.0547	-0.0291
	(0.0303)	(0.0365)	(0.0392)
	[0.5289]	[0.1343]	[0.4594]
treat x 17Q3	0.0111	0.0363	-0.0120
	(0.0304)	(0.0403)	(0.0373)
	[0.7149]	[0.3679]	[0.7483]
treat x 17Q4	0.0090	0.0384	-0.0302
	(0.0308)	(0.0370)	(0.0396)
	[0.7707]	[0.3006]	[0.4457]
treat x 18Q1	-0.0267	-0.0073	-0.0542
	(0.0323)	(0.0377)	(0.0430)
	[0.4086]	[0.8463]	[0.2087]
treat x 18Q2	-0.0190	-0.0106	-0.0296
	(0.0269)	(0.0334)	(0.0377)
	[0.4800]	[0.7506]	[0.4324]
treat x 18Q3	0.0053	-0.0052	0.0097
	(0.0273)	(0.0331)	(0.0387)
	[0.8451]	[0.8749]	[0.8019]
treat x 18Q4	0.0048	-0.0012	0.0154
	(0.0288)	(0.0372)	(0.0389)
	[0.8686]	[0.9735]	[0.6930]
r2	0.427	0.460	0.388
N	15039	7981	7058

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table B.2.4: Alternative control group 1: Quarterly diff-in-diff results, by size (zip code fixed effects)

	(1)	(2)	(3)	(4)	(5)
	10	25	50	75	90
treated	-0.3968***	-0.3958***	-0.3636***	-0.3543***	-0.2551
	(0.0435)	(0.0618)	(0.0863)	(0.1021)	(0.1531)
	[0.0000]	[0.0000]	[0.0000]	[0.0005]	[0.0957]
after 2018q1	0.8267***	0.7548***	0.6031***	0.4556***	0.3849***
	(0.0131)	(0.0114)	(0.0104)	(0.0094)	(0.0100)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	0.0453	0.0223	-0.0266	-0.0262	-0.0376*
	(0.0250)	(0.0167)	(0.0177)	(0.0149)	(0.0162)
	[0.0694]	[0.1815]	[0.1329]	[0.0791]	[0.0201]
r2	0.298	0.307	0.311	0.298	0.282
N	30252	30252	30252	30252	30252

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table B.2.5: Alternative control group 1: Quantile regression estimates (zip code fixed effects)

	(1)	(2)	(3)	(4)	(5)
	10	25	50	75	90
treated	0.1078*	0.1297	0.0865	0.0664	0.0781*
	(0.0525)	(0.0874)	(0.0489)	(0.0362)	(0.0381)
	[0.0401]	[0.1377]	[0.0769]	[0.0663]	[0.0405]
after 2018q1	0.8220***	0.7665***	0.6062***	0.4505***	0.3780***
	(0.0164)	(0.0122)	(0.0105)	(0.0079)	(0.0081)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	0.0367	0.0011	-0.0259	-0.0173	-0.0107
	(0.0268)	(0.0166)	(0.0150)	(0.0127)	(0.0161)
	[0.1709]	[0.9489]	[0.0827]	[0.1716]	[0.5077]
r2	0.312	0.322	0.327	0.312	0.294
N	30252	30252	30252	30252	30252

Includes grid dummies and log(size). SE's clustered on the grid level.

Table B.2.6: Alternative control group 1: Quantile regression estimates (grid fixed effects)

B.2.2 Alternative control group 2: dropping two distant zip codes from district 5

	(1)	(2)	(3)
	all	50>	50≤
treated	-0.3164***	-0.4480**	-0.1250**
	(0.0800)	(0.1470)	(0.0388)
	[0.0001]	[0.0025]	[0.0014]
after 2018q1	0.2855***	0.2875***	0.2865***
	(0.0080)	(0.0104)	(0.0104)
	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	-0.0288*	-0.0359*	-0.0230
	(0.0133)	(0.0157)	(0.0184)
	[0.0308]	[0.0226]	[0.2134]
r2	0.369	0.404	0.321
N	14320	7515	6805

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table B.2.7: Alternative control group 2: Before-after diff-in-diff results, by size (zip code fixed effects)

10.14754/CEU.2022.03

	(1)	(2)	(3)
	all	50>	50≤
treated	0.0674*	0.1030***	0.0153
	(0.0292)	(0.0212)	(0.0473)
	[0.0224]	[0.0000]	[0.7464]
after 2018q1	0.2855***	0.2843***	0.2911***
	(0.0076)	(0.0083)	(0.0121)
	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	-0.0235	-0.0261	-0.0267
	(0.0124)	(0.0147)	(0.0197)
	[0.0593]	[0.0792]	[0.1771]
r2	0.406	0.457	0.355
N	14320	7515	6805

Includes grid dummies and log(size). SE's clustered on the grid level

Table B.2.8: Alternative control group 2: Before-after diff-in-diff results, by size (grid fixed effects)

	all	50>	50≤
treated	-0.3161***	-0.4559***	-0.1121
	(0.0643)	(0.1323)	(0.0591)
	[0.0000]	[0.0006]	[0.0587]
treat x 16Q2	-0.0119	-0.0150	-0.0091
	(0.0287)	(0.0361)	(0.0436)
	[0.6792]	[0.6779]	[0.8341]
treat x 16Q3	0.0301	0.0591	0.0016
	(0.0330)	(0.0447)	(0.0366)
	[0.3619]	[0.1871]	[0.9653]
treat x 16Q4	0.0195	-0.0036	0.0426
	(0.0290)	(0.0359)	(0.0370)
	[0.5028]	[0.9193]	[0.2500]
treat x 17Q1	0.0212	0.0190	0.0314
	(0.0343)	(0.0441)	(0.0411)
	[0.5363]	[0.6663]	[0.4453]
treat x 17Q2	0.0175	0.0566	-0.0343
	(0.0306)	(0.0371)	(0.0395)
	[0.5681]	[0.1285]	[0.3867]
treat x 17Q3	0.0096	0.0415	-0.0206
	(0.0307)	(0.0409)	(0.0376)
	[0.7552]	[0.3108]	[0.5839]
treat x 17Q4	0.0079	0.0404	-0.0338
	(0.0313)	(0.0381)	(0.0400)
	[0.8012]	[0.2902]	[0.3983]
treat x 18Q1	-0.0297	-0.0042	-0.0637
	(0.0324)	(0.0382)	(0.0431)
	[0.3591]	[0.9129]	[0.1403]
treat x 18Q2	-0.0185	-0.0047	-0.0357
	(0.0273)	(0.0340)	(0.0381)
	[0.4996]	[0.8907]	[0.3491]
treat x 18Q3	-0.0041	-0.0075	-0.0058
	(0.0275)	(0.0338)	(0.0389)
	[0.8827]	[0.8241]	[0.8812]
treat x 18Q4	-0.0036	-0.0062	0.0030
	(0.0293)	(0.0379)	(0.0392)
	[0.9034]	[0.8691]	[0.9396]
r2	0.423	0.453	0.385
N	14320	7515	6805

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Standard errors in round, p values in squared brackets.

Table B.2.9: Alternative control group 2: Quarterly diff-in-diff results, by size (zip code fixed effects)

	(1)	(2)	(3)	(4)	(5)
	10	25	50	75	90
treated	-0.6359***	-0.4757***	-0.1297	-0.2961***	-0.4364***
	(0.0854)	(0.0943)	(0.0734)	(0.0592)	(0.0818)
	[0.0000]	[0.0000]	[0.0771]	[0.0000]	[0.0000]
after 2018q1	0.8330***	0.7613***	0.6087***	0.4572***	0.3881***
	(0.0135)	(0.0116)	(0.0103)	(0.0095)	(0.0100)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	0.0397	0.0151	-0.0339	-0.0258	-0.0433**
	(0.0253)	(0.0170)	(0.0175)	(0.0154)	(0.0155)
	[0.1162]	[0.3733]	[0.0526]	[0.0935]	[0.0053]
r2	0.292	0.302	0.305	0.292	0.275
N	28766	28766	28766	28766	28766

Includes zip dummies and log(size). SE's clustered on the street x zip level.

Table B.2.10: Alternative control group 2: Quantile regression estimates (zip code effects)

	(1)	(2)	(3)	(4)	(5)
	10	25	50	75	90
treated	0.1069*	0.1318	0.0856	0.0686	0.0856*
	(0.0539)	(0.0877)	(0.0462)	(0.0356)	(0.0395)
	[0.0474]	[0.1329]	[0.0642]	[0.0540]	[0.0301]
after 2018q1	0.8316***	0.7725***	0.6116***	0.4515***	0.3818***
	(0.0178)	(0.0124)	(0.0104)	(0.0080)	(0.0080)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
treated x after 2018q1	0.0299	-0.0037	-0.0334*	-0.0152	-0.0137
	(0.0269)	(0.0169)	(0.0150)	(0.0122)	(0.0161)
	[0.2658]	[0.8257]	[0.0262]	[0.2131]	[0.3960]
r2	0.308	0.318	0.323	0.307	0.288
N	28766	28766	28766	28766	28766

Includes grid dummies and log(size). SE's clustered on the grid level.

Table B.2.11: Alternative control group 2: Quantile regression estimates (grid fixed effects)

10.14754/CEU.2022.03

Appendix C

Appendix for chapter 3

C.1 Proof of the Proposition

Proof. I start with characterizing the strategies which, depending on the exact values of the parameters, are not always dominated by some other strategies. After that, I move on to finding the profit-maximizing strategy for every possible parameter set.

Possible reporting strategies. Since buyer valuations are linear in b, the media has three kinds of possibly non-dominated reporting strategies (see Figure C.1.1):

1. complete bias ($b = 1$). In this case, it can only sell to leftitsts (as $WTP_R(1) < 0$). The price is $WTP_L(1)$ and the profits are

$$\pi^{b=1}_{\text{left only}} = \frac{1}{2} WTP_L(1) = \frac{1}{2}\alpha\left(\frac{1}{2} - p_L\right). \tag{C.1}$$

   This will be a feasible option only if $\frac{\partial}{\partial b} WTP_L > 0 \iff \alpha > \frac{p_L}{\frac{1}{2} - p_L}$, otherwise, honest reporting and selling to leftists only will strictly dominate this strategy.

2. honest reporting ($b = 0$). Unless the rightist consumers are completely biased, the agency can potentially sell them its honest reporting for a low enough price: $WTP_R(0, \theta) > 0$ unless $p_R = 1$.

   In general, the media will compare the profits of serving unbiased news to only left-leaning consumers ($\pi_L = \frac{1}{2} WTP_L(0, p_L) = \frac{1}{2} p_L$), only right-leaning consumers ($\pi_R = \frac{1}{2} WTP_R(0, p_R) = \frac{1}{2}(1 - p_R)$) and both segments ($\pi_B = min\{WTP_R(0, p_R), WTP_L(0, p_L)\} = min\{p_L, 1 - p_R\}$).

3. partial distortion ($b = b^*$). In this case, the media will distort with $b^* \in (0, 1)$ and serve both markets (selling to one type with $b^* \in (0, 1)$ will be dominated). Given that consumer valuations are linear in b, it will chose a reporting strategy b^* such that $WTP_L(b^*) = WTP_R(b^*)$:

$$b^* = \frac{1 - p_L - p_R}{(\alpha + 1)(p_R - p_L)} \tag{C.2}$$

   resulting in profits

$$\pi^* = WTP_L(b^*) = WTP(b^*) = p_L + \frac{1 - p_L - p_R}{(\alpha + 1)(p_R - p_L)} \left(\frac{\alpha}{2} - p_L(1 + \alpha)\right) \tag{C.3}$$

Depending on whether WTP_L is increasing or decreasing in b and the relative size of the biases, the following comparisons will be relevant for the media (see the corresponding subfigures in Figure C.1.1).

(a) *WTP is increasing and rightists are more biased (Figure C.1.1a).*

The media would set $b = 1$ and serve exclusively leftists only if

$$\frac{1}{2}WTP_L(1) > 1 - p_R \iff \bar{\alpha}_0 \equiv \frac{2(1-p_R)}{\frac{1}{2} - p_L} < \alpha. \tag{C.4}$$

Otherwise, it will sell undistorted news to both types at $1 - p_R$ (selling honest reporting to only type is dominated by selling uninformative reporting to leftists only).

(b) *WTP is increasing and leftists are more biased* In this region, $\alpha > \frac{p_L}{1-p_L}$ (WTP_L is increasing in b) and $p_L < 1 - p_R$ (leftists are more biased).

I start with showing that the optimal level of partial bias b^* dominates truthtelling $b = 0$ in this region. In general, the firm can sell to (a) leftists only, (b) rightist only or (c) both groups.

Clearly, selling truth to leftists only and charging p_L is dominated by selling truth to rightist only and charging $1 - p_R$, as in this region $p_L < 1 - p_R$. Therefore the firm pick the best option conditional on truthtelling, which is either $\frac{1}{2}(1 - p_R)$ by selling to rightists only or p_L by selling to both types and compare it with profits of b^*:

$$\pi^* > max\{p_L, \frac{1}{2}(1-p_R)\}. \tag{C.5}$$

If $p_L > \frac{1}{2}(1 - p_R)$, we need to verify that

$$p_L + \frac{1-p_L-p_R}{(\alpha+1)(p_R-p_L)}\Big(\frac{\alpha}{2} - p_L(1+\alpha)\Big) > p_L. \tag{C.6}$$

This holds, given that $\frac{1-p_L-p_R}{(\alpha+1)(p_R-p_L)} > 0$ as $p_L < 1 - p_R$ in this region and the value in the bracket being positive is equivalent with $\alpha > \frac{p_L}{\frac{1}{2}-p_L}$, which also holds in this region.

If $\frac{1}{2}(1 - p_R) \geq p_L$, however, we need to verify that

$$p_L + \frac{1-p_L-p_R}{(\alpha+1)(p_R-p_L)}\Big(\frac{\alpha}{2} - p_L(1+\alpha)\Big) > \frac{1}{2}(1-p_R). \tag{C.7}$$

It can be checked that the left hand side of this inequality is increasing in p_L, so if it holds for $p_L = 0$, then it holds for any positive p_L. Substituting $p_L = 0$, rearranging and using that fact that $\alpha > \frac{p_L}{\frac{1}{2}-p_L}$ again gives the desired result.

To finish, I have to establish the condition for $b = 1$ dominating b^*. This is given implicitly by condition

$$\pi^*(\bar{\alpha}) < \bar{\alpha}(\frac{1}{2} - p_L) \tag{C.8}$$

(c) *WTP is decreasing and rightists are more biased (Figure C.1.1c).* Selling unbiased news to leftists clearly dominates selling them fully distorted news. Therefore the media will set $b = 0$ and sell to both groups if

$$1 - p_R > \frac{1}{2}p_L \iff 2(1-p_R) > p_L \tag{C.9}$$

for $1 - p_R$. Otherwise, it will serve only leftists for p_L.

10.14754/CEU.2022.03

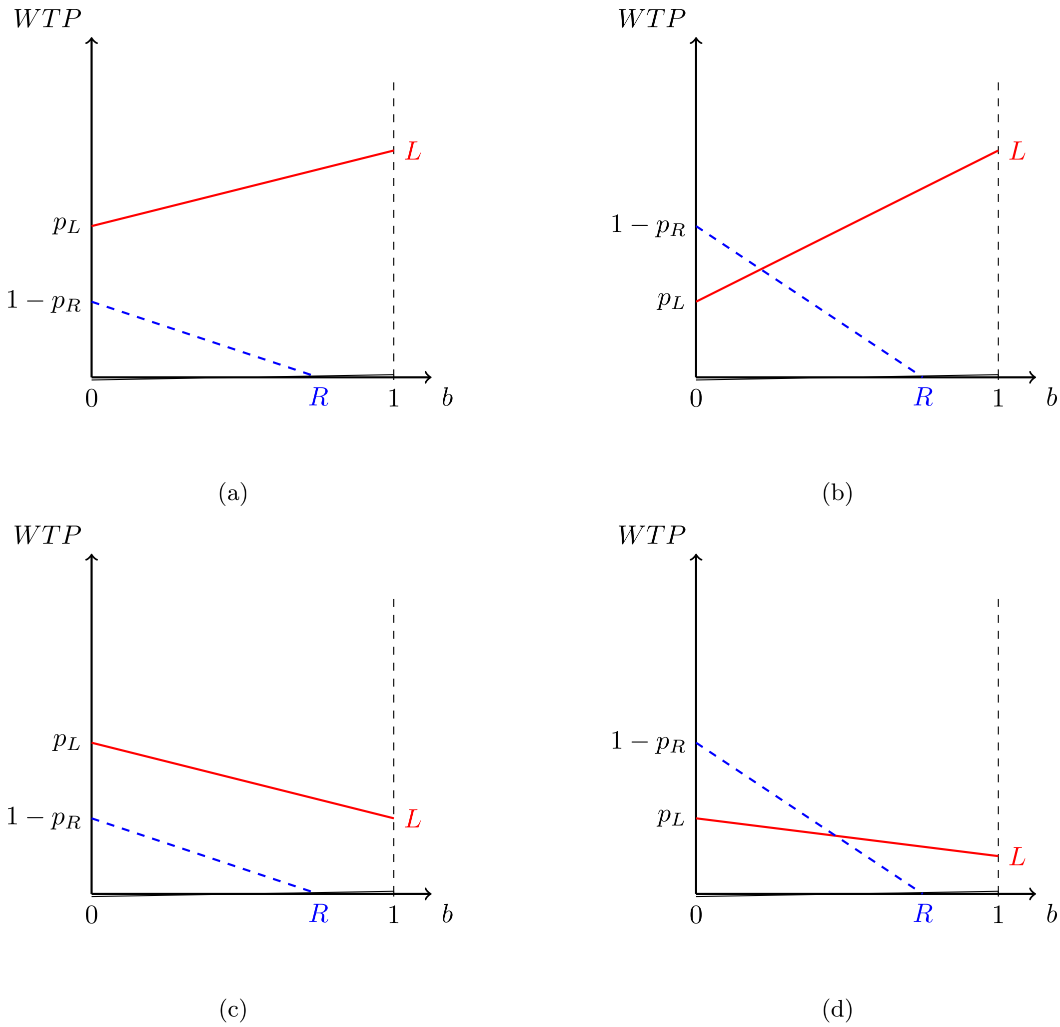

Figure C.1.1: Willingness to pay as a function of reporting bias.

(d) *WTP is decreasing and leftists are more biased (Figure C.1.1d).* Zero bias dominates $b = 1$ and serving rightists at $1 - p_R$ dominates $b = b^*$. As above, the media will maximize its profits given $b = 0$.

It will sell to both groups for p_L if

$$p_L > \frac{1}{2}(1 - p_R) \tag{C.10}$$

otherwise, it will serve only rightists by charging them $1 - p_R$.

□