Essays on Innovation and Markets

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Economics

have examined a dissertation entitled

"Essays on Innovation and Markets"

presented by **Sophie Calder-Wang**

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature ______________________

Typed name: Prof. Ariel Pakes

Signature ______________________

Typed name: Prof. Edward Glaeser

Signature ______________________

Typed name: Prof. Robin Lee

Signature Adi Sunderam

Typed name: Prof. Adi Sunderam

Signature Paul Gompers

Typed name: Prof. Paul Gompers

Date: May 20, 2020

Essays on Innovation and Markets

A dissertation presented

by

Sophie Calder-Wang

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2020

Dissertation Advisors:
Professor Ariel Pakes
Professor Edward Glaeser

Author:
Sophie Calder-Wang

Essays on Innovation and Markets

Abstract

This thesis includes three essays that explore the interaction between innovation and markets. The first essay studies the distributional implications of technological innovation in the context of the rapid rise of the sharing economy. In particular, it evaluates the welfare impact of the home-sharing platform Airbnb on residents of New York City by estimating an integrated model of the housing market. It finds that Airbnb leads to a rise in equilibrium rents that negatively affect high-income, educated, and White renters the most. Moreover, the benefits generated by home-sharing accrue to a small fraction of city residents. The second essay deals with the econometric issue of estimating production functions. It highlights that including second-order conditions as moment inequalities can improve estimation, especially when weak instruments are present. The last essay explores the financing of innovations through venture capital, focusing on the role of gender diversity. It finds that when existing partners have more daughters, the likelihood of hiring women increases. It also suggests that firms with greater gender diversity have better financial performance.

Contents

Abstract iii
Acknowledgments x

Introduction **1**

1 The Distributional Impact of the Sharing Economy on the Housing Market **3**
1.1 Introduction 3
1.1.1 A Stylized Model 12
1.2 Background and Data 16
1.3 Model Description 26
1.3.1 A Model of an Integrated Housing Market 26
1.4 Estimation and Results 31
1.4.1 Estimating the Long-Term Rental Demand 32
1.4.2 Estimating the Short-Term Rental Supply 36
1.4.3 Parameter Estimates 39
1.5 Counterfactual Analysis 43
1.5.1 The Distributional Impact via the Rent Channel 45
1.5.2 The Distributional Impact via the Host Channel 62
1.5.3 The Net Impact on Renters 67
1.5.4 Implications for the Planner 74
1.6 Conclusion 78

2 The Value of Information: Why You Should Add the Second Order Conditions **81**
2.1 Introduction 81
2.2 The Model 82
2.2.1 The Continuous Optimal Choice Problem 83
2.2.2 The Discrete Optimal Choice Problem 86
2.3 Conclusion 90

3 And the Children Shall Lead: Gender Diversity and Performance in Venture Capital **91**
3.1 Introduction . . . 91
3.2 Data Collection . . . 97
3.3 Methodology . . . 108
3.4 Empirical Results . . . 113
3.4.1 Effects on Venture Capital Hiring . . . 113
3.4.2 Effects on Venture Capital Performance . . . 119
3.4.3 Instrumental Variable Regression . . . 124
3.5 Conclusion . . . 140

References **143**

Appendix A Appendix to Chapter 1 **150**
A.1 Supplementary Figures . . . 150
A.2 More on LTR Demand Estimation . . . 154
A.2.1 Model . . . 154
A.2.2 Estimation . . . 155
A.2.3 Simulation Study . . . 157
A.3 More on STR Supply Estimation . . . 160
A.3.1 Problem Formulation . . . 160
A.3.2 Analytical Derivations . . . 161

Appendix B Appendix to Chapter 2 **163**
B.1 Simulation Details . . . 163

Appendix C Appendix to Chapter 3 **164**
C.1 Supplementary Tables and Figures . . . 164

List of Tables

1.1 Breakdown of Airbnb Activity in NYC in 2018 19
1.2 Income and Demographic Characteristics of NYC Renters 24
1.3 Reduced-Form Correlations between Airbnb and Neighborhood Characteristics 25
1.4 Parameter Estimates for the Long-Term Rental Demand Model (Part I) 41
1.5 Parameter Estimates for the Long-Term Rental Demand Model (Part II) . . . 42
1.6 Parameter Estimates for the Short-Term Rental Supply Model 44
1.7 Counterfactual using Actual Airbnb Penetration vs. Uniform Airbnb Penetration 55
1.8 Estimated Own- and Cross- Price Semi-Elasticities 61
1.9 Net Welfare Impact by Household Income 69
1.10 Net Welfare Impact by Household Demographics 70
1.11 Welfare Impact on All Participants 75

3.1 Children Data Collection 100
3.2 Firm Sample Selection 101
3.3 Summary Statistics 104
3.4 Number of Female Hires 105
3.5 Female Hired Ratio by Firm Size 105
3.6 Hiring Patterns Over Time 106
3.7 Industry Patterns 107
3.8 Partner Characteristics by Gender 108
3.9 Hiring Level Regression 115
3.10 Hiring Level Regression (Alternative Measures of Daughters) 118
3.11 Daughter Effect on Performance (Deal-Level Reduced-Form) 121
3.12 Daughter Effect on Performance (Fund-Level Reduced-Form) 123
3.13 Deal-Level Instrumental Variable Regression 127
3.14 Deal-Level First-Stage Regression 128
3.15 Fund-Level Instrumental Variable Regression 130
3.16 Fund-Level First-Stage Regression 131
3.17 Daughter Effect on Entrepreneurs 134
3.18 Daughter Effect on Performance Controlling for Venture Capitalist Gender . 135

3.19 Daughter Effect on Performance Controlling for Individual Venture Capitalist Family Characteristics 137
3.20 Hiring Level Regression (Daughter Age Effects) 139

A.1 Regression on Mean Utility 159
A.2 Estimation of the Willingness to Pay for Amenities 159

C.1 Impact of the First Child's Gender 164
C.2 Robustness: Hiring Level Regression when Children Age Available 165
C.3 Robustness: Hiring Level Regression Excluding Email Respondents 166
C.4 Hiring Level Regression All Partners (Alternative Measures of Daughters) . 167
C.5 Robustness: Deal-Level Instrumental Variable Regression with IPOs Only . . 169
C.6 Robustness: Deal-Level Instrumental Variable Regression with U.S. Deals . . 170
C.7 Robustness: Impact on Career Outcomes (Deal Count) 171
C.8 Robustness: Impact on Career Outcomes (Partner Tenure) 172
C.9 Robustness: Fund-Level Instrumental Variable Regression with U.S. Funds . 173

List of Figures

1.1 An Illustrative Model of an Integrated Housing Market 13
1.3 Growth of Airbnb across the U.S. 18
1.5 Percentage (%) of Housing Units on Airbnb (Having At Least One Reservation in 2018) 21
1.6 Percentage (%) of Rental Housing Units Reallocated (Available on Airbnb for 180+ Days in 2018) 22
1.7 Housing Types Comparison 48
1.9 Welfare Impact (Rent Channel) by Household Size and Race 49
1.11 Welfare Impact (Rent Channel) by Education and Income 50
1.13 Race and Ethnicity across Neighborhoods 52
1.14 Education Attainment across Neighborhoods 53
1.15 The Equilibrium Effects of Supply Restrictions (Successive Steps of Best Response Function) 58
1.16 Distribution of Short-Term Rental Surplus 64
1.18 Distribution of Short-Term Rental Surplus by Income Quintiles 65
1.20 The Net Welfare Impact on Renters 68
1.22 Net Welfare Impact by Neighborhood (Median) 72
1.23 Net Welfare Impact by Neighborhood (Right Tail) 73

2.1 Parameter Identification for the Continuous Problem 85
2.3 The Effects of Instrument Strengths 85
2.5 Impact of Moments from the Second Order Condition 86
2.7 Identification from Moment Inequalities Conditions 87
2.9 Discrete vs. Continuous Problem 88
2.11 The Distribution of the Upper Bounds 89
2.13 Improvements Introduced by the Additional Moment Condition 90

3.1 The Probability of Hiring a Woman 116

A.1 Distribution of Listings by Calendar Availability 150

A.2 The Daily Number of Reservations of Private Rooms Sold on Airbnb in New York City. 151
A.3 The Daily Average Price of Private Rooms sold on Airbnb in New York City. 152
A.4 Building Age of NYC Housing Units 153

C.1 Fund Return Distribution 168
C.2 Randomization Inference: Hiring Level Regression 174
C.3 Randomization Inference: Deal-Level Performance Regression 175

Acknowledgments

I am greatly indebted to my advisors Ariel Pakes, Edward Glaeser, Paul Gompers, Robin Lee, and Adi Sunderam for their invaluable support, guidance, and feedback on my research. I am grateful for the advice and support of Alex Bell, John Campbell, Lauren Cohen, Mark Egan, Chiara Farronato, Sam Hanson, Nir Hak, Yizhou Jin, Myrto Kalouptsidi, Jing Li, Chris Malloy, David Martin, Mark Shepard, Jeremy Stein, Boris Vallee, Ron Yang, and the participants of Harvard's Industrial Organization and Finance research seminars. I thank Kevin Huang, Coco Zhang, and Patrick Sweeney at Harvard Business School for collecting and organizing the data used in the last chapter.

To my parents Xu Xiaoping and Wang Qiang and to my husband Bradley Calder for their encouragement, support, and love

Introduction

The main theme of this dissertation is the interaction between innovation and markets. The first essay provides a direct analysis when technological innovation lowers the cost of production, which brings about changes in the market dynamics and the welfare of market participants. The second essay provides an econometric framework to improve the measurement of production costs by using second-order conditions as moment inequalities. The last essay deals with the financing side of innovative activities, focusing on the impact of gender diversity in venture capital. The three essays are briefly summarized below.

The Distributional Impact of the Sharing Economy on the Housing Market

What is the impact of the sharing economy, pioneered by companies such as Airbnb, on the housing market? I estimate the welfare and distributional impact of Airbnb on the residents of New York City. I develop a model of an integrated housing market, in which a landlord can offer a housing unit for rent either on the traditional long-term rental market or on the newly available short-term rental market. By estimating a structural model of residential choice and linking it to detailed Airbnb usage data, I estimate the effect of such reallocation on the equilibrium rents across different housing types and demographic groups. In addition, to evaluate the gains from direct home-sharing, I estimate a supply system featuring heterogeneous costs. Overall, renters in New York City suffer a loss of \$178mm per annum, as the losses from the rent channel dominate the gains from the host channel. I find that the increased rent burden falls most heavily on high-income, educated, and white renters because they prefer housing and location amenities that are most desirable to tourists. Moreover, there is a divergence between the median and the tail, where a few enterprising low-income households obtain substantial gains from home-sharing. Thus, this paper delivers a nuanced characterization of the winners and losers of the sharing economy, and provides a framework for understanding the consequences of regulating such technological innovations.

The Value of Information

When conducting estimation based on agent optimization, I show that one can improve the performance of the estimator when information such as the second order condition is appropriately incorporated as moment inequality restrictions, especially when there are weak instruments. I run a simulation study to demonstrate the effectiveness of this approach in both continuous and discrete choice problems, and illustrate to empirical researchers how to include the additional moment inequalities in practice.

And the Children Shall Lead: Gender Diversity and Performance in Venture Capital

Given an overall lack of gender diversity in venture capital and entrepreneurship shown in Calder-Wang and Gompers (2017), we ask: What promotes greater gender diversity in hiring? Does diversity lead to better firm performance and higher investment returns? In this paper, using a unique dataset of the gender of venture capital partners' children, we find strong evidence that when existing partners have more daughters, the propensity to hire female partners increases. Moreover, our instrumental variable results suggest that increased gender diversity improves deal and fund performance. Lastly, the effects are primarily driven by the gender of senior partners' children.

Chapter 1

The Distributional Impact of the Sharing Economy on the Housing Market

1.1 Introduction

Economic theory teaches that cost-reducing new technologies should improve welfare. By substantially reducing transaction costs, platform companies such as Airbnb allow existing housing units to be used by short-term visitors in exchange for payment. Such innovation improves the allocation and utilization of the underlying asset. However, it is not necessarily Pareto improving.

As housing supply is constrained in many coastal markets in the United States,[1] there are significant concerns that Airbnb exacerbates housing affordability problems. Many worry that housing units are being reallocated away from the traditional long-term rental market and thereby displacing existing residents. Legal battles continue in places such as New York City where Mayor de Blasio signed legislation to curb Airbnb rentals in late 2018.

[1] The amount of housing construction in New York City has been depressed over the past three decades. In fact, 41% of the homes today were built prior to 1940, and 88% of the homes were built prior to 1990. Only 2.9% of the homes were built since 2010.

However, the law was subsequently blocked by the court amid privacy concerns.

Proponents of Airbnb argue that the additional income that hosts earn from home-sharing, especially in expensive cities, is vital to their livelihood. An important feature of many sharing economy platforms is that services are produced by peers, rather than firms, and therefore can distribute gains directly to individuals. Therefore, the question is empirical: For New York City residents, does the welfare gain from home-sharing offset the welfare loss from increased housing costs? Moreover, how does the welfare impact differ across key demographic characteristics, such as income, education, race, and family structure?

To answer these questions, I specify and estimate a structural model that highlights two key innovations that Airbnb brings to the housing market. First, the long-term rental market and the short-term rental market become integrated on the supply side: An absentee landlord who owns a housing unit can choose between the two markets, whichever yields greater profit. In equilibrium, a fraction of the housing units are reallocated to Airbnb. Since housing supply is inelastic, such reallocation raises prices for long-term renters and decreases their welfare.[2] Second, the utilization of existing housing units increases when renters themselves offer space in their homes to host short-term visitors, especially during times when short-term rental demand is high. The proceeds from such direct home-sharing raise the welfare of these residents.

My structural model flexibly incorporates rich heterogeneity in household preferences as well as heterogeneity in the cost of home-sharing, which allows me to analyze the distributional impact across demographic groups. I model the demand for long-term rentals as a discrete choice problem featuring heterogeneous household preferences over housing attributes (McFadden, 1978; Bayer, Ferreira and Mcmillan, 2007). In addition to the rental price, the demand model captures a set of housing attributes ranging from hedonic attributes (such as the number of bedrooms, year built, type of structure) to neighborhood attributes

[2] In the case of owner-occupiers, they can be thought of as renting from themselves. However, as 67% of the housing units in NYC are renter-occupied, my analysis focuses primarily on the renters and returns to owner-occupiers at the end.

(such as distance to job centers and neighborhood demographics in terms of race, ethnicity, and education). Moreover, it also allows for a horizontal preference over these attributes; For example, it captures the differential preferences over living in a predominantly ethnic neighborhood depending on the ethnicity of the household members. Compared to the traditional Alonso-Mills-Muth analysis of the urban problem (Alonso, 1964; Mills, 1967; Muth, 1969), the discrete choice framework makes explicit the preferences over a wide vector of housing attributes. Given that the entry of Airbnb varies greatly across space and housing types, the ability to incorporate heterogeneity into the long-term rental demand is crucial to evaluating its distributional impact.

I model the supply of short-term rentals by residents as a binary choice problem where a resident decides whether to share her home on a given day with a short-term visitor at the prevailing market price. She makes this decision based on the trade-off between the income she makes and the cost of providing such short-term rental services. The short-term rental supply model allows households to have differential costs based on their demographic characteristics, such as age, education, and family structure. It also allows for differential price sensitivity based on household income. To my knowledge, even though others have estimated the overall distribution of home-sharing costs (Farronato and Fradkin, 2018), this paper is the first to examine how costs of peer production differ across income and demographic characteristics.

To estimate the structural model, I adapt and extend methods from the empirical industrial organization literature. For long-term rental demand, I use the American Community Survey (ACS) Public Use Microdata Sample, whose key benefit is the availability of individual-level data. I observe the full vector of household demographic characteristics together with the full vector of housing attributes chosen by the household, including the location of the home at the neighborhood level. I construct moment conditions that match market shares, as well as the covariance between the housing attributes and the average demographic characteristics of households living in those homes (Berry, Levinsohn and Pakes, 2004). To address the concern that there may be an unobservable housing quality correlated with price, I construct an instrument based on a home's location in the characteristics space, following Bayer, Ferreira and Mcmillan (2007).

To estimate the short-term rental supply, I leverage high-frequency Airbnb transaction data. Although Berry, Levinsohn and Pakes (1995); Nevo (2000, 2001) (BLP) methods are widely used by researchers to estimate demand systems, I propose an adaptation so they can be used to estimate a heterogeneous *supply* system.[3] Since I observe the location of each housing unit on Airbnb, variations in the distribution of neighborhoods demographics allow me to estimate the heterogeneity in the cost of home-sharing. Moreover, price variations at the daily level allow me to estimate the heterogeneity in the price coefficient. In addition, to address unobservable costs, I use a measure of short-term demand seasonality as the price instrument. Since the high-frequency daily data results in a large number of market-share equations to match, I employ the MPEC procedure developed in Dubé, Fox and Su (2012) to improve the numerical performance of the estimator.

With the estimated model parameters, I conduct two counterfactual analyses to evaluate the distributional impact of Airbnb on the participants of the housing market. To evaluate the welfare impact through the rent channel in the long-term rental market, I perform a counterfactual analysis where all housing units made available by absentee landlords on Airbnb are returned to the long-term rental market. To evaluate the welfare impact via the host channel in the short-term rental market, I perform a counterfactual analysis where the hosts are no longer allowed to participate in home-sharing.[4]

The key findings are threefold. (i) The net impact of Airbnb aggregated across all renters is a loss of \$178mm per annum (p.a.), because the losses from the rent channel at \$201mm p.a. dominate the gains from the host channel at \$23mm p.a. (ii) While the median renter loses \$125 p.a., more significant losses are suffered by renters who are high-income,

[3]BLP methods only require aggregate data, which is particularly useful in my setting given that part of the ongoing legal feud in New York City concerns the regulator's inability to obtain individual host-level data.

[4]Since these counterfactuals reverse the entry of Airbnb, I interpret the negative of the compensating variation therein as the welfare and distributional implications of Airbnb, as both the long-term and the short-term rental models are static. However, it ignores dynamic considerations such as switching costs. Therefore, the results here should be interpreted as a static approximation.

educated, and white, because they demand housing types that are more desirable in the short-term rental market. (iii) There is a divergence between the median and the tail: The equilibrium rent increase affects all 2.1 million renter-occupied units in New York City, but the host gains accrue heavily to a small fraction of households with particularly low costs of sharing, including low-income families.

Specifically, to compute the welfare loss through the rent channel, I link the detailed Airbnb penetration data across neighborhoods and housing types with the estimated long-term rental demand model. Combined with the assumption that the total supply of physical structures available for housing is fixed, I back out the entire vector of price changes across all housing types because of the supply squeeze due to Airbnb. Across New York City, I estimate that about 0.68% of the housing units were likely to have been reallocated,[5] resulting in an equilibrium price increase of 0.71%. The compensating variation for the median renter earning $47,000 annually is $128 p.a. When aggregated across all renters, it amounts to a total transfer of $200mm p.a. from renters to property owners, or $2.7bn in NPV terms.[6] It also leads to a welfare loss of $1mm p.a. for renters displaced from the city.

The presence of severe housing supply restrictions is the main driver for the elevated and widespread equilibrium price increase in the long-term rental market. When a number of housing units of type h leave the long-term rental market due to Airbnb, three forces act. First, because the supply of type h is inelastic, a permanent reduction leads to a price increase in h. Second, displaced renters may try to substitute to another housing type h' in the city. Since the supply of h' is also inelastic, the price of type h' also increases. Finally, the price increase in h' creates a feedback loop as some return to choose type h, thereby pushing up the price of h even further. It continues until enough renters leave the city. Such spillover effects are the primary reasons for an elevated equilibrium price response even in neighborhoods that are relatively far from the city center and have low levels of direct

[5]This is based on the number of housing units marked as available on Airbnb for over 180 days in 2018.

[6]Using a capitalization rate of 7.5% based on New York City hotel REITs as of 2018, based on data compiled by CBRE Research.

Airbnb activity.

In terms of the distributional impact via the rent channel, I find the most significant welfare losses, when measured in dollar terms, are suffered by renters who are high-income, educated, and white.[7] I find that the median renter in the top income quintile suffers a loss of $167 p.a., whereas the median renter in the bottom income quintile suffers $123 p.a. In terms of education, the median renter with a college degree loses $156 p.a., compared to a loss of $120 p.a. for those without a college degree. Across race and ethnicity, the median white renter loses $152 p.a., whereas the loss is $134 p.a. for African American renters and $113 p.a. for Hispanic renters.

Contrary to what typical political or media narratives purport, welfare differences are primarily driven by the actual geographical patterns of Airbnb usage, where its penetration in New York City tends to be higher in educated, high-income, and predominantly white neighborhoods. Although the impact on higher-income renters is exacerbated by the fact that they also tend to have a greater willingness to pay for housing amenities in general, I show that the role of geography is dominant.

Next, to compute the welfare gains via the host channel, I use the parameters estimated from the supply model to compute the compensating variation if residents no longer had the option to share their homes. The distribution of the welfare gains has a large mass close to zero and a heavy right tail. I find that the host surpluses are irrelevant for the vast majority of the residents: The median resident gains a surplus of only $0.4 p.a., and the 75th-percentile resident makes $5.9 p.a., as the fraction of residents who have actually hosted guests on Airbnb remains low, averaging to only 0.8% of housing units. However, the expected gain to households in the right tail (above the 99th percentile) amounts to over $300 p.a. When aggregated across all renters, the total host surplus produced by supplying rooms to short-term visitors on Airbnb amounts to $23mm per year, or approximately

[7] My focus here is to conduct a positive analysis to estimate the welfare impact in dollar terms. I do not impose any differential social welfare weights across individuals of different levels of education, race, or income. Insofar as a municipal planner wishes to use an alternative set of welfare weights, the analysis here provides the basis for which those weights could be applied.

$300mm in NPV terms.[8]

In terms of the distributional impact via the host channel, larger gains accrue to lower-cost suppliers, which tend to be young and educated households with no children. Some low-income households also benefit more as they are better able to take advantage of peak demand in the short-term rental market. As lower-income households tend to be more price sensitive, their supply is more elastic. Conditional on being in the right tail of the host surplus distribution, households in the bottom income quintile expected a gain of $454 p.a., compared to a gain of $233 p.a. for the top income quintile. As a result, given the pattern of short-term rental activity, Airbnb does not necessarily exacerbate the income inequality *among* renters. However, Airbnb is likely to increase wealth inequality, depending on the status of property ownership.

Taken together, I find that the welfare loss via the rent channel at an NPV of $2.7bn dwarfs the gains from the increased utilization via the host channel at an NPV of $300mm. Since housing supply is inelastic, reallocation by absentee landlords raises rents on *all* units, not just for the particular unit moved to Airbnb. Nonetheless, the cost of home-sharing remains high for most people, evidenced by the fact that only a small fraction of households become hosts. As the losses from the rent channel are widespread and the gains from the host channel are concentrated, the net welfare impact remains negative for the vast majority of renters. With 67% of the housing units being renter-occupied, the median household in the city experiences a loss of $114 p.a. As a result, even a simple voting model favors restricting Airbnb reallocation.

From the perspective of the social planner, since the rent increase results in a mere transfer to absentee landlords, the overall welfare impact of Airbnb remains positive because it also includes the economic gains accrued to all hosts, as well as the surpluses accrued to tourists net of hotel losses. However, the substantial transfer from renters to property owners

[8]Note that the surplus number estimated here is only for resident hosts sharing rooms in their homes. Absentee landlords who have reallocated to Airbnb also obtain surpluses from providing short-term rental services, presumably greater than their foregone rent. Their gains also matter to the social planner, but they are not the focus of my supply estimator.

reflects the regulatory conundrum caused by the severe housing supply restrictions in place. More broadly, this paper shows that taking into account the existing market structure and regulatory constraints is essential to evaluate whether gains from technological innovations could be shared equitably throughout society.

Related Literature The literature on the sharing economy is nascent but rapidly expanding. Theoretical work such as Filippas, Horton and Zeckhauser (2019) explore the impact of the sharing economy on asset ownership in the long-run. Within the empirical literature, many study the design of peer-to-peer platforms (Edelman, Luca and Svirsky, 2017; Jaffe, Coles, Levitt and Popov, 2019). Related to outcomes in the housing market, Horn and Merante (2017) and Barron, Kung and Proserpio (2018) provide evidence of Airbnb leading to increased rent and home prices using panel data following the entry of Airbnb. Koster, Van Ommeren and Volkhausen (2019) take advantage of a panel regression-discontinuity design around city borders and find evidence of depressed home prices following short-term rental restrictions in Los Angeles. Using detailed transaction records from Barcelona, Garcia-López, Jofre-Monseny, Martínez-Mazza and Segú (2020) employ a battery of reduced-form specifications, yielding qualitatively similar results. Using a structural approach, Farronato and Fradkin (2018) estimate the welfare impact of Airbnb on tourists and hotels. This paper is the first to develop a structural model to directly quantify the distributional impact of home-sharing in a city on the participants of its housing market.

Compared to reduced-form analyses, a structural model of an integrated housing market developed in this paper offers several important advantages. First, by imposing market-clearing conditions, the structural model readily characterizes the spillover effect of Airbnb entry in one part of the city on prices of other parts of the city.[9] Second, the model explicitly incorporates the notion that local residents are not only drivers of long-term rental demand, but also a source of short-term rental supply when they become Airbnb hosts themselves.

[9]When regressing price changes on Airbnb intensity, reduced-form analyses may suffer an interference between treatment and control: Neighborhoods in a given geographic region are generally not independent but rather imperfect substitutes for long-term residents.

The benefits accrued to these local hosts is an integral component of the welfare analysis, since the widespread of peer production is a key feature of such platform services. Lastly, by recovering the heterogeneity in both demand and supply preferences, the model provides an estimate of the distributional impact along several key demographic characteristics. The distributional impact of such technological innovation is a key input into public policies, especially when there are equity concerns.

This paper also highlights the fact that supply constraints in the housing market (Saiz, 2010; Gyourko, 2009; Gyourko and Molloy, 2015; Baum-Snow and Han, 2020) have a large effect on how the efficiency gains from the sharing economy are distributed, thereby complementing the existing literature on the economic impact of housing constraints and related housing policies (Ganong and Shoag, 2017; Hsieh and Moretti, 2019; Diamond, McQuade and Qian, 2019; Favilukis, Mabille and Van Nieuwerburgh, 2018). In addition, since the ability to rent to short-term visitors alters the cash-flow-generating ability of the underlying housing asset, this paper is also related to studies that estimate the determinants of housing value, ranging from school quality, crime, nearby foreclosures, and various environmental factors.[10] Notably, beyond the direct price effects of short-term rental demand, Almagro and Domínguez-Iino (2020) complement this work by studying changes in endogenous neighborhood amenities in a dynamic framework, which is particularly relevant when tourists account for a substantial fraction of the local population. Methodologically, this paper treats the housing stock as inherently heterogeneous in multiple dimensions explicitly,[11] which allows for a detailed characterization of the distributional implications.

This paper adds to the growing literature that employs structural models to capture equilibrium effects and consumer welfare due to innovations. Buchak, Matvos, Piskorski

[10]This is an extensive literature that employs a variety of methods, including both choice-based sorting (Ferreyra, 2007; Timmins, 2007; Banzhaf and Walsh, 2008; Bayer, Keohane and Timmins, 2009; Klaiber and Phaneuf, 2010; Tra, 2010; Galiani, Murphy and Pantano, 2015) and reduced-from estimates (Black, 1999; Greenstone and Gallagher, 2008; Campbell, Giglio and Pathak, 2011; Davis, 2011; Autor, Palmer and Pathak, 2014, 2017).

[11]The differentiated nature is characterized in Rosen (1974); Epple (1987), where Wong (2019) provides a comparison between the hedonic marginal willingness-to-pay and those estimated by discrete choice models.

and Seru (2018) show that technology provides an avenue to dampen capital requirements in mortgage lending. Higgins (2019) argues that the adoption of financial technology exhibits two-sided network effects, which benefit wealthy consumers. Rich structural models are also used to empirically assess the impact of existing or proposed market regulations, especially taking into account supply-side responses and consumer heterogeneity. For instance, Robles-Garcia (2018) and Benetton (2018) emphasize the shifting of market power among mortgage originators and brokers. Cuesta and Sepulveda (2019) study the impact of price regulation on consumer credit, whereas Nelson (2019) examines its distributional impact across different types of borrowers.

The remainder of the paper proceeds as follows. Section 1.1.1 starts with a stylized model, highlighting the key innovations of Airbnb. Section 1.2 discusses the background and the data used for the analysis. Section 1.3 presents the main structural model, followed by Section 1.4, describing how the models are estimated and the parameters obtained in . Section 1.5 performs the counterfactual analysis examining the welfare and distributional impact of Airbnb via the rent channel and via the host channel, respectively. Section 1.6 concludes.

1.1.1 A Stylized Model

In this subsection, I present a stylized version of the model highlighting the key innovations that Airbnb brings to the housing market. First, the long-term rental market and the short-term rental market become integrated as absentee landlords can offer their housing units in either market. Second, residents can act as peer suppliers to participate directly in the production of short-term rental services. The former improves the allocative efficiency, whereas the latter increases the utilization of existing homes. Despite its stylized nature, it illustrates the key outcomes of this paper, namely, the welfare loss through the rent channel and the welfare gain through the host channel. It also motivates my use of a model-based approach and clarifies the main assumptions made.

In Figure 1.1, Panel A on the top left started with the long-term rental market, whereas

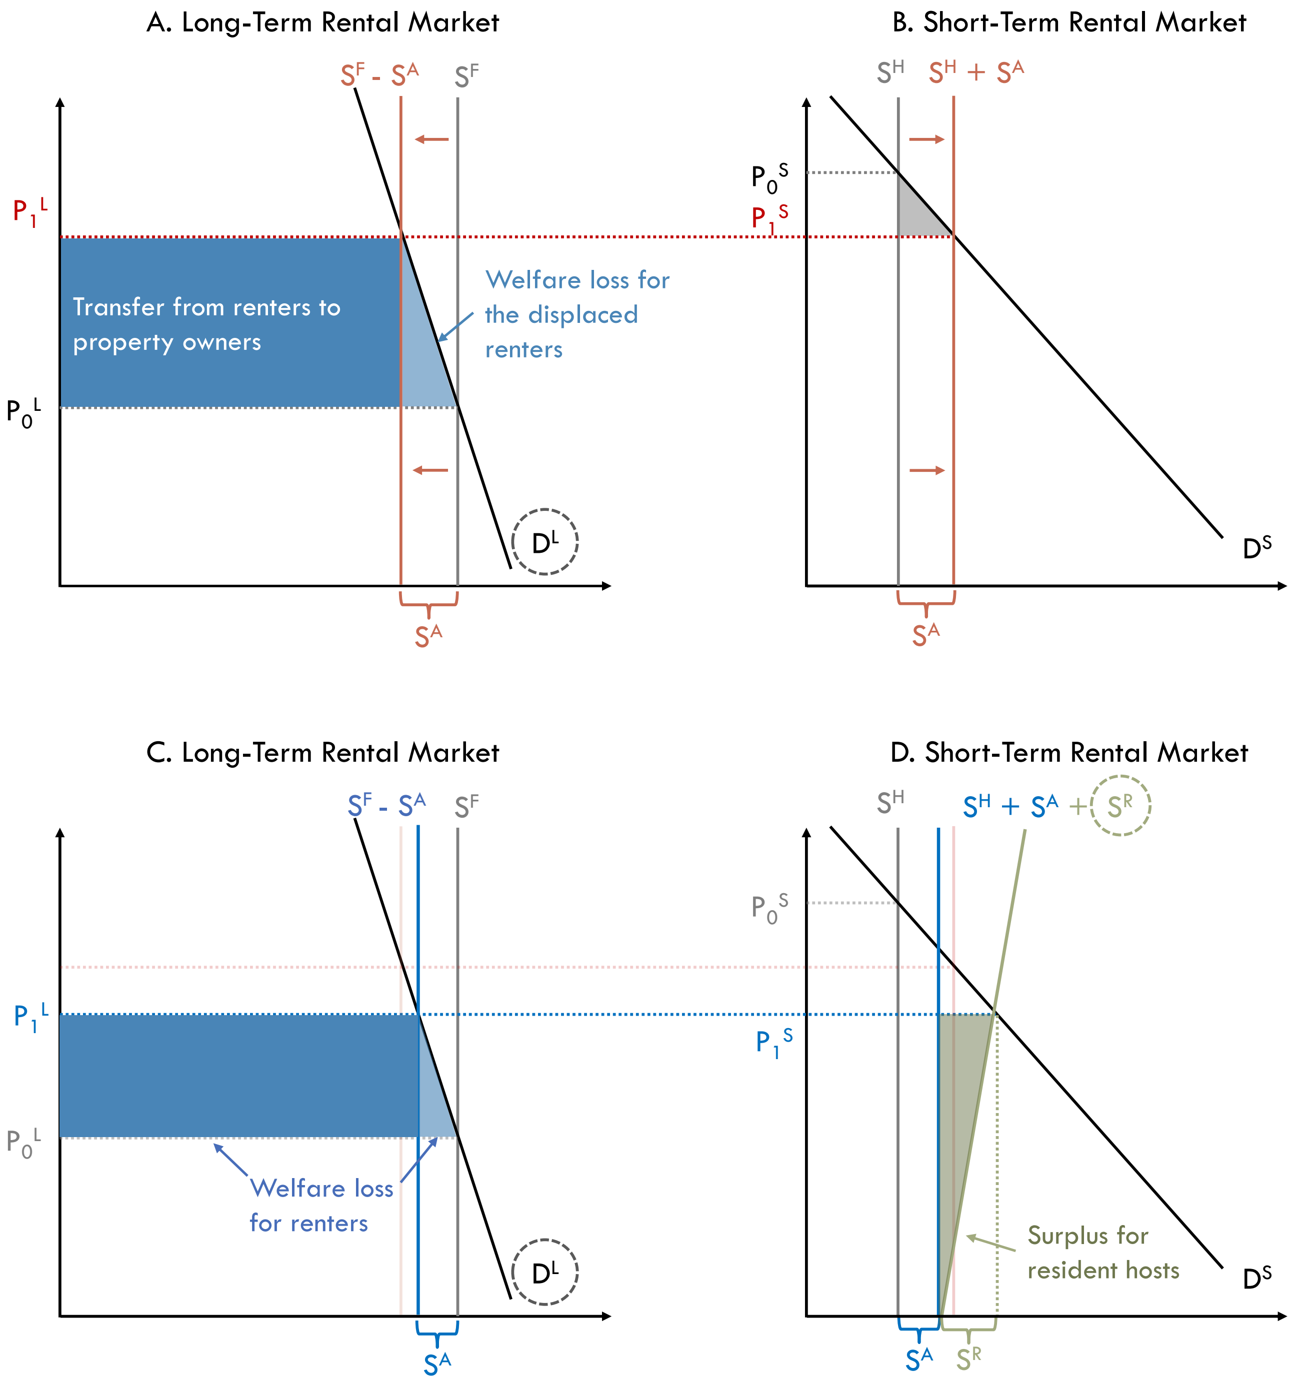

Figure 1.1: *An Illustrative Model of an Integrated Housing Market*

Panels A and B illustrate the market equilibrium when Airbnb allows for a more efficient allocation of housing units *between* the long-term and the short-term rental market. Panels C and D illustrate the market equilibrium when Airbnb also allows for increased utilization of existing homes already occupied by long-term tenants.

Panel B started with the short-term rental market. Before the arrival of Airbnb, these two markets were completely separate.[12] Both markets were in equilibrium with different market-clearing prices at p_0^L and p_0^S, respectively.[13] In other words, physical structures built for residential purposes could only be rented to long-term tenants. It used to be prohibitively costly for a property owner to rent out a residential housing unit on a short-term basis.[14]

As a large-scale home-sharing intermediary, Airbnb has rapidly brought down search costs and reduced asymmetric information between hosts and guests. An absentee landlord is no longer confined to renting in the traditional long-term rental market but gains the option to participate in the newly available short-term rental market.[15] For the purpose of this stylized model alone, assume the cost of operating in either market by absentee landlords is zero. If prices are higher in the short-term rental market $p_0^S > p_0^L$, then absentee landlords will be induced to reallocate toward Airbnb and obtain higher prices. As more and more housing units are reallocated, it reduces the price wedge between the two markets. In equilibrium, a no-arbitrage condition pins down the new market price $p_1 = p_1^L = p_1^S$ that clear both markets, as well as the equilibrium number of reallocated housing units S^A.

Since housing supply has been inelastic in New York City, I assume that the total number of physical structures is fixed in the model. As such, any reduction in the number of housing units available for long-term rental results in a higher equilibrium rent $p_1^L > p_0^L$. The blue shaded regions of Panel A illustrate the welfare impact on the renters: The rectangle above the remaining renters represents the welfare transfer to property owners, whereas the triangle above the displaced renters represents the welfare loss for those who leave the city. Importantly, the detailed Airbnb usage records allow me to tabulate the reallocated quantity

[12] In practice, there exist residential apartments that accept guests from the corporate travel market. However, I abstract that component away, since it is much smaller than the overall housing market in NYC. I also assume that hotel operators do not accept long-term tenants.

[13] To make them comparable, the prices may be thought of as the equivalent daily rental rates in the two markets.

[14] The average length of a stay on Airbnb in New York City is 4.5 days.

[15] In many cities, the legality of Airbnb has been hotly debated. In the case of NYC, although renting out a Class A unit for less than 30 days without the presence of its permanent resident is a violation of its Multiple Dwelling Law, it is also estimated that the law is not effectively enforced (Jia and Wagman, 2018).

S^A from the data directly. As a result, once I estimate the slope of the long-term rental demand D^L, together with the observed S^A, I can compute the welfare impact of Airbnb via the rent channel.[16]

The welfare loss through the rent channel is only one part of the overall effect, since Airbnb also allows residents to increase the utilization of their homes without displacing themselves.[17] Panel D in Figure 1.1 illustrates the additional short-term rental supply that is provided by resident hosts, denoted as S^R. It is upward sloping because more residents will find the hassle of home-sharing worthwhile if the price p_1^S is high. Nonetheless, the no-arbitrage condition for property owners still implies that there exist an equilibrium quantity S^A and an equilibrium price $p_1 = p_1^L = p_1^S$ that clears both markets. In Panel D, the green shaded region indicates the surplus accrued to such resident hosts. Therefore, once I estimate the slope of the short-term supply S^R, I can compute their welfare gains via the host channel, which can then be netted against the welfare losses through the rent channel.

Hence, the stylized model illustrates that the key welfare outcomes (shown as the blue and green shaded regions respectively in Figure 1.1) can be computed by estimating the slope of the long-term rental demand D^L and the slope of the short-term rental supply S^R.

The stylized model not only demonstrates the benefits for an equilibrium model-based approach, it also highlights the need to incorporate heterogeneity into the full structural model. A model-based approach captures the equilibrium effects by ensuring the market clearing conditions are satisfied in both markets before and after the entry of Airbnb. Specifically, since it allows households to re-optimize when faced with a new price vector, a fully-specified long-term rental demand model can estimate the welfare change when the *bundle* of housing attributes available has changed. Moreover, a model-based approach featuring heterogeneous preferences can capture the distributional impact over relevant

[16] Correspondingly, in the short-term rental market, since the price has declined $p_1^S < p_0^S$, the net surplus is designated by the gray triangle, enjoyed by tourists visiting the city, and net of hotel losses.

[17] The potential resident hosts on Airbnb include both renters and owner-occupiers. I do not explicitly model whether such home-sharing violates specific leasing agreements or home-association agreements, but they are implicitly incorporated as the cost to home-sharing.

household characteristics, thereby evaluating the net welfare impact by household income, race, education, and family structure. This is especially relevant for those concerned about income or racial disparities. Lastly, quasi-experiments that allow researchers to identify at least the price effect of Airbnb are available only in limited settings,[18] leaving the welfare and distributional issues in some of the most important housing markets unanswered.

The stylized model also clearly shows the limitations and assumptions used. First, when computing the equilibrium price impact of Airbnb via the rent channel, the total supply of physical structures is assumed to be fixed. Another assumption is the absence of negative externalities of short-term visitors on the neighborhood, potentially due to increased noise and traffic. Also, an implicit assumption here is that the ability of residents to act as peer suppliers do not alter their long-term rental demand, which is an empirical simplification.[19] Finally, the tourist welfare will not be explicitly modeled, as the paper focuses on the residents.

1.2 Background and Data

The combination of mobile technology and the improving designs of the reputation systems on two-sided platforms have greatly facilitated the development of the "sharing economy." Although there is no single official definition over what the sharing economy is, several important features stand out: First, it allows existing asset owners to increase utilization by allowing someone else to use their asset temporarily in exchange for payment. The growth and the development of a large-scale intermediary drastically lowers search costs. Second, the platform companies that facilitate the exchange typically do not own the assets themselves, so the services offered on the platform are fulfilled by peers. Although there

[18]Without natural experiments such as Valentin (2019) or Calder-Wang (2019), correlating home price changes with the growth of Airbnb's is susceptible to the usual endogeneity problem, namely that improvements in unobserved neighborhood quality drive both the short-term rental demand and the long-term rental demand.

[19]Given that the expected gains from hosting turn out to be immaterial for the vast majority of households, this assumption is a simplification. I also discuss its likely impact at the end of the paper. In other places where the expected host gains are large, such as vacation markets, this assumption may be more important.

exists an extensive literature on estimating the production function for firms, it is not immediately clear how it extends to such decentralized peer-production processes. In addition, the development of a well-functioning reputation system alleviates the problem of asymmetric information. Even though the buyer and the seller are typically only engaged in a one-time exchange, they interact with the platform repeatedly, thus allowing the platform to aggregate relevant user information over time.

One of the most prominent businesses in this category is the home-sharing platform Airbnb, which is an online marketplace for arranging short-term lodging, especially home-stays. According to its website, the founders of Airbnb started the company in 2008 by allowing guests to sleep on their air mattress in San Francisco after they noticed that the participants of a conference were struggling to find accommodation as local hotels were sold out.[20] The novelty and the scalability of the business model has attracted over $4.4bn in venture capital funding. As of 2019, Airbnb boasts over 7 million listings across the globe, which is significantly higher than any hotel group.[21]

Figure 1.3 shows the rapid growth of Airbnb across cities in the United States. Between 2014 and 2018, the number of reservations made on Airbnb quadrupled in many cities. The largest metropolitan market in the U.S. measured by days reserved is New York, followed by Los Angeles and Miami. Table 1.1 shows that Airbnb booked 5.8 million days of stay in New York City in 2018, which is about 15% of the total number of hotel stays. The average price of an entire home on Airbnb is $224 per night, and the average price of a private room is $86 per night. In 2018, 74,963 listings had active transactions on Airbnb , representing about 2.2% of all housing units. Moreover, Airbnb is by far the most dominant player among short-term rental platforms, capturing over 90% of the market share in New York City.[22]

The average level of Airbnb activity in the city masks the extensive geographic heterogeneity across neighborhoods and boroughs, as illustrated in Figure 1.5. In Brooklyn, the

[20] `https://news.airbnb.com/about-us/`

[21] The largest hotel company in the world, Marriott International operates approximately 1.1 million rooms.

[22] The second largest player in New York City is HomeAway. Given Airbnb's dominance in NYC, I refer to Airbnb and home-sharing platforms synonymously through the paper.

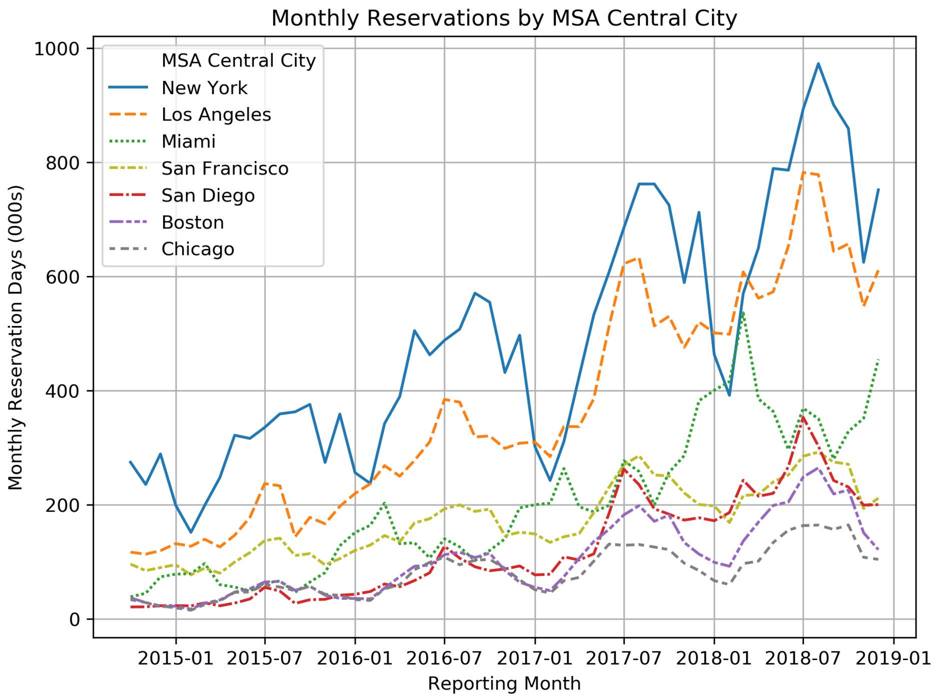

Figure 1.3: *Growth of Airbnb across the U.S.*

The time series plot shows the rapid growth in the number of monthly reservations across select MSAs in the United States. New York is the largest metropolitan market for Airbnb in the United States. Between 2015 and 2018, the number of reservations quadrupled.

Table 1.1: *Breakdown of Airbnb Activity in NYC in 2018*

This table summarizes the transactions on Airbnb in NYC over 2018. On Airbnb, a property is listed as one of three types: entire home/apt, private room, or shared room. Over 95% of the properties are listed as either "entire home/apt" or "private room", which are the focus of this paper. Overall, there are over 74,963 properties on Airbnb that have experienced at least one reservation, accounting for 2.2% of the housing units in the city.

Listing Type	Bedroom(s)	Num Days Reserved (000s)	Num Properties (000s)	Avg Daily Rate
Entire home/apt	All	3,204	38	$224
Entire home/apt	1	1,968	25	$178
Entire home/apt	2	838	9	$257
Entire home/apt	3	289	3	$350
Entire home/apt	4	82	1	$530
Private room	-	2,526	34	$86
Shared room	-	128	2	$59
Total		5,858	75	$156

neighborhood with the highest proportion of housing units active on Airbnb in 2018 is Greenpoint & Williamsburg (9.3%), followed by Bushwick (8.5%). In Manhattan, Chelsea, Clinton & Midtown (6.9%) is at the top, followed by Chinatown & Lower East Side (6.5%). In Queens, the most active neighborhood is Astoria & Long Island City (2.8%). In the Bronx, the most active neighborhood is Concourse, Highbridge & Mount Eden (0.4%), which has much less penetration than the other boroughs.[23]

Just because a housing unit has experienced a booking on Airbnb, it is not obvious what its alternative use would have been if Airbnb had not been invented. Importantly, there are vast variations in terms of a property's calendar availabilities. Figure A.1 shows the distribution of properties by the fraction of its calendar marked available. The distribution is spread out between 0 and 1, but a large mass of 35% are marked as available on Airbnb for over 90% of the calendar days, which indicates they are likely dedicated for short-term

[23] I have excluded Staten Island, as it is much less likely to be in the choice set of a short-term traveler visiting NYC.

rental use only.

Based on calendar availability and the listing type, I approximate the proportion of housing units likely to have been reallocated away from the long-term rental market.[24] Specifically, I designate entire homes available on Airbnb for over 180 days in 2018 as being reallocated by their property owners, whose alternative use would have been long-term rentals. It averages to 0.68% of the rental housing stock, while the distribution also exhibits remarkable heterogeneity across neighborhoods. Figure 1.6 shows the geographic variations of such reallocation follow a similar pattern as the overall level of Airbnb activity in Figure 1.5, but with an even higher concentration in Manhattan. The neighborhood that experiences the highest levels of Airbnb reallocation is Chelsea, Clinton & Midtown (3.4%), followed by Murray Hill, Gramercy & Stuyvesant Town (2.4%), and then by Battery Park City, Greenwich Village & Soho (2.3%). In Brooklyn, Williamsburg & Greenpoint (1.9%) and Bedford-Stuyvesant (1.8%) are affected the most.

Data The primary data source is a full sample of Airbnb listings scraped by a third-party data vendor, AirDNA. AirDNA started scraping the entire website of Airbnb.com comprehensively in late 2014, which is when my data start.[25] For each listing in New York City, the dataset contains detailed information about the property characteristics, including the type of property and the hedonic attributes, such as the number of bedrooms, the number of bathrooms, and other relevant amenities. Broken down by listing type, 51% are entire homes, 45% are private rooms, and less than 5% are shared rooms, as summarized in Table 1.1. Importantly, the dataset also scrapes the latitude and the longitude of the

[24]These are the housing units most likely to be the target of the legislative efforts signed by Mayor de Blasio on August 6, 2018. Int. 981-A requires online short-term rental platforms to report data about transactions. According to Council Member Carlina Rivera, such reported data could be used by the Mayor's Office of Special Enforcement to "pursue more effective oversight and action over the bad actors that exist throughout this largely unmonitored market." `https://www1.nyc.gov/office-of-the-mayor/news/398-18/`

[25]Because hosts need to make their listing information publicly available online to all potential guests, the scraper downloads all the available information. Since the scraper visits each listing multiple times per week, the number of days reserved is then backed out from changes in the calendar availabilities.

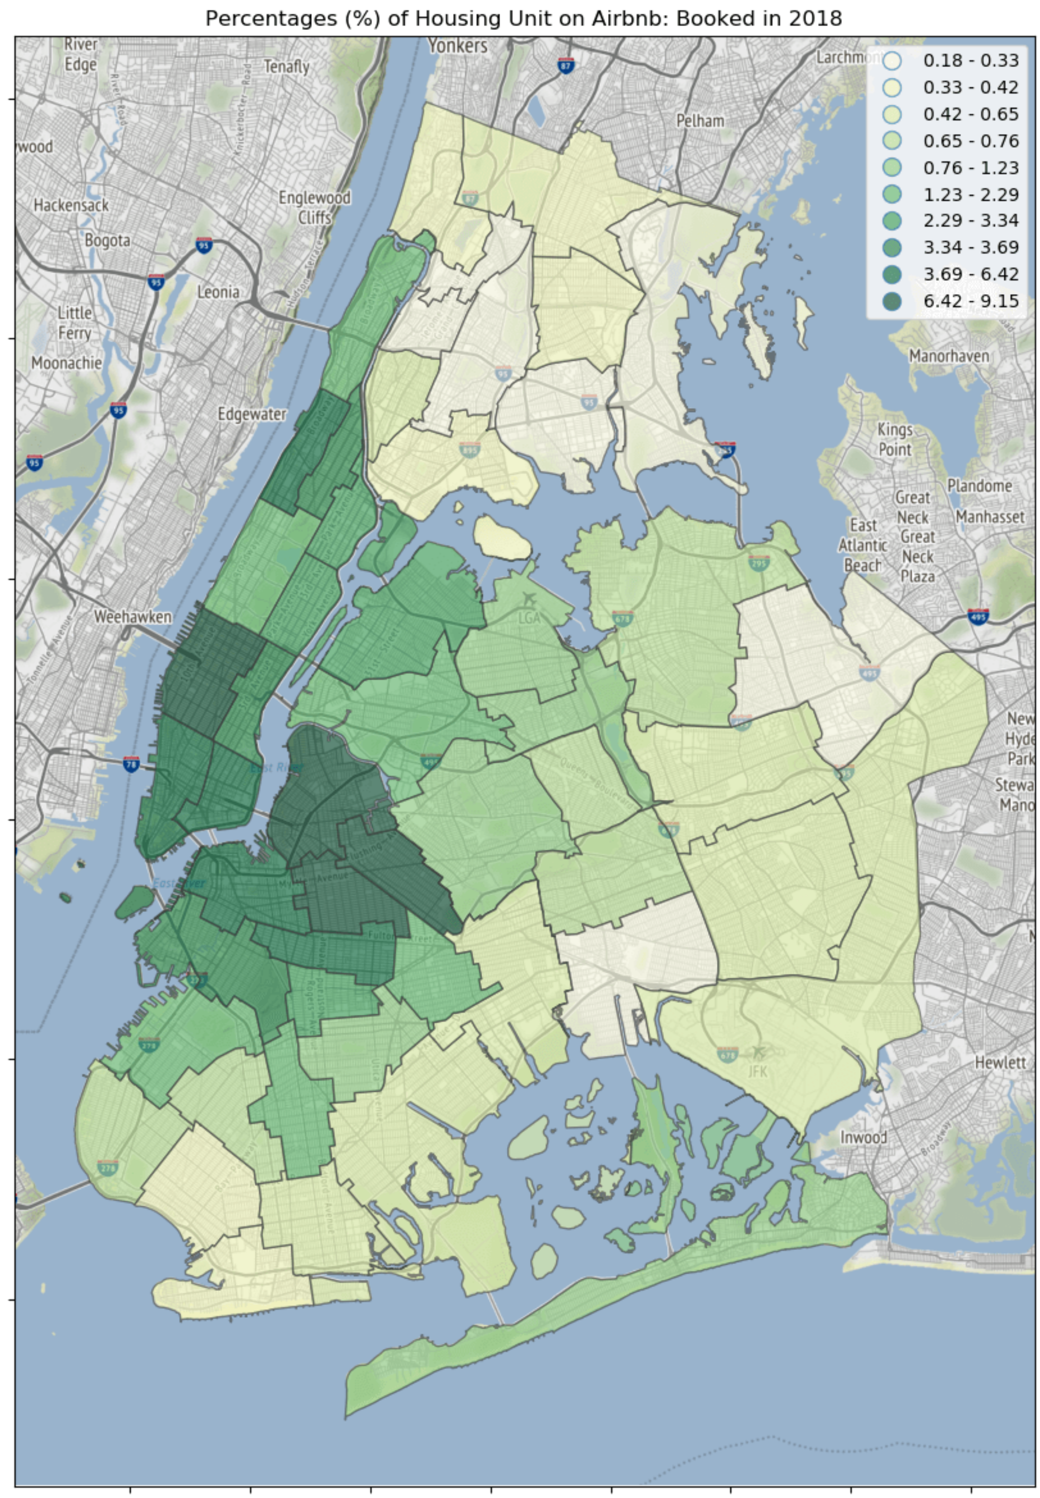

Figure 1.5: *Percentage (%) of Housing Units on Airbnb (Having At Least One Reservation in 2018)*

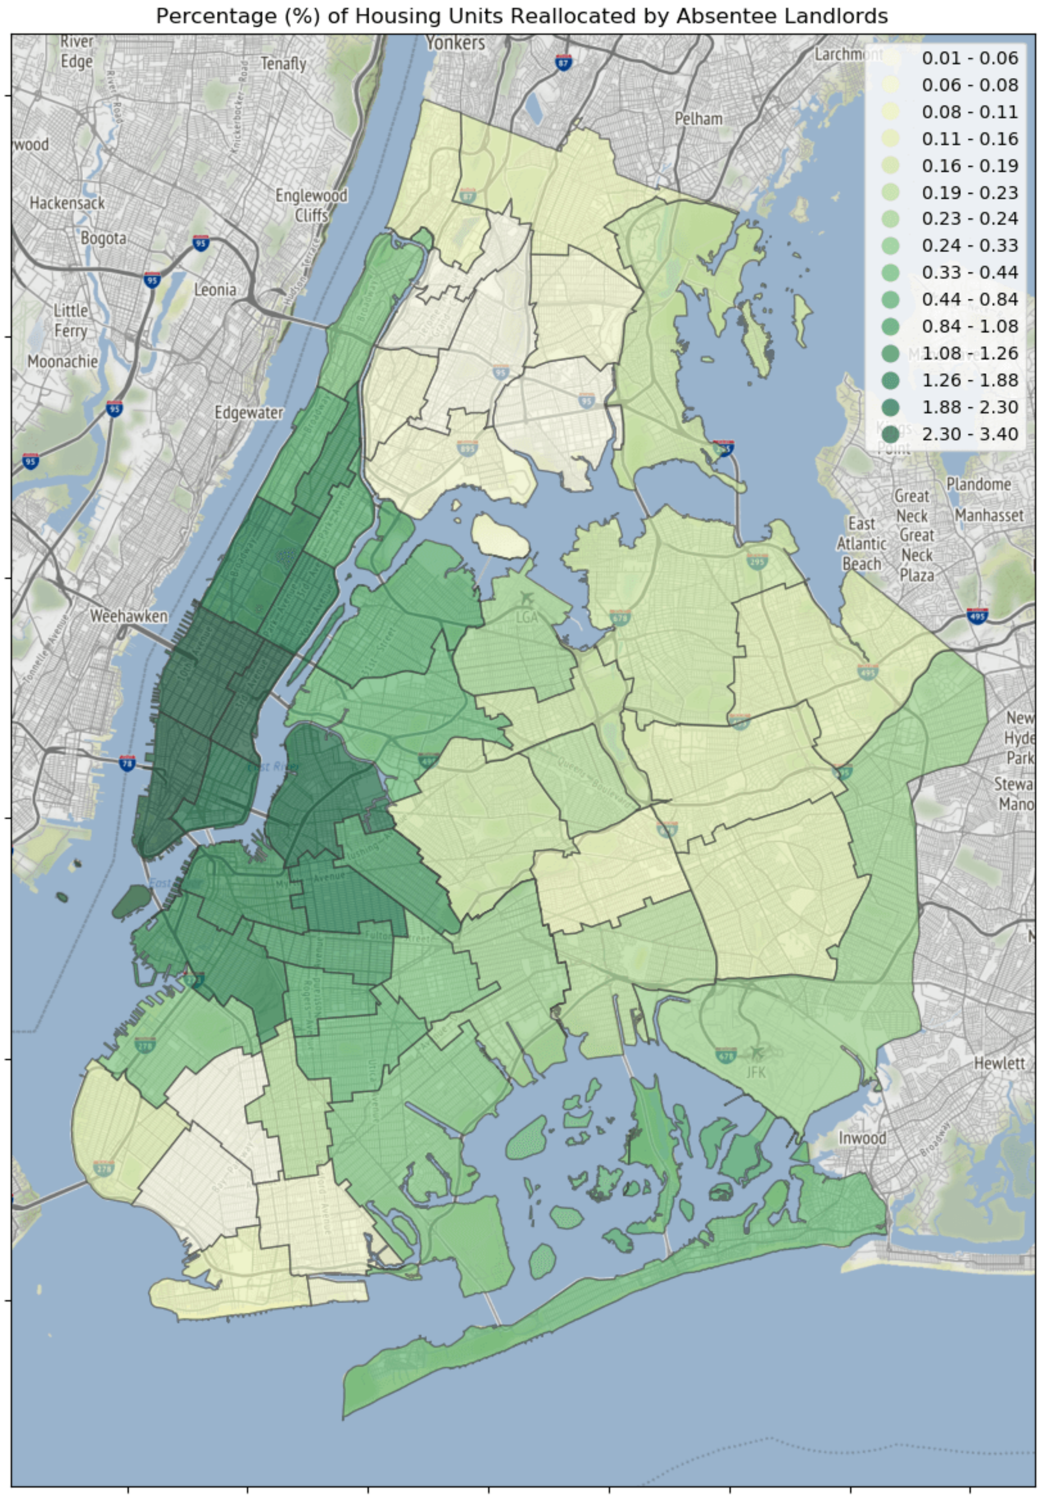

Figure 1.6: *Percentage (%) of Rental Housing Units Reallocated (Available on Airbnb for 180+ Days in 2018)*

property,[26] which allows me to map to its corresponding neighborhood.

Another beneficial feature of the dataset is its high-frequency panel, available at the most detailed daily level. It allows me to capture the seasonal nature of the short-term rental market. For each listing and every day, I observe whether the listing was available on Airbnb, its listed price, as well as whether a reservation occurred on that day. Figure A.2 and Figure A.3 show the time series plot of the total daily quantity and the average transaction price of all private rooms in New York City. There are strong patterns of seasonality. The average transaction price in the sample period for a private room is $73, with an annualized volatility of 42%. The peak demand predictably happens on New Year's Eve each year, with its price averaged at $94. On the other hand, the trough season happens predictably in January and February with depressed quantities and prices. The peak-to-trough ratio of the daily quantity is 3.5x, whereas the peak-to-trough ratio of the average daily price is 1.6x. Overall, the high-frequency dataset allows me to take advantage of the seasonal variations in the short-term rental demand from visitors.

The second dataset is the American Community Survey (ACS) Public Use Microdata Sample. As it contains individual-level data, it is particularly helpful in estimating the housing choice model. I observe the full vector of household demographic characteristics, together with the full vector of housing attributes chosen by the household. The key demographic variables include household income, education, race, ethnicity, age, and family size. The key housing attributes include monthly price,[27] number of bedrooms, building age, and type of building. Moreover, I also observe the location of the home at the neighborhood level.[28] For each neighborhood, I also obtain average neighborhood characteristics, including

[26]Even though the scraper downloads the exact location as shown on the Airbnb website, Airbnb is known to add some noise to ensure host privacy. However, based on Wachsmuth and Weisler (2018), the perturbation is between 0 and 500 feet, which introduces minimal noise in terms of assigning a property to its neighborhood at the Public Use Microdata Area (PUMA) level.

[27]For renters, I observe the monthly rent. For owner-occupiers, I impute an equivalent monthly user-cost based on reported home values, as done by Bayer, Ferreira and Mcmillan (2007).

[28]In the ACS microdata, the location of the home is available at the PUMA level. Fortunately, densely populated New York City contains 55 PUMAs and they could be used as an approximation for neighborhoods. I use the 2017 1-Year Estimates to estimate the long-term rental demand model, which is a 1% sample based on

race, education, and average commute time.

Overall, New York City has 3.14 million occupied housing units, of which renters occupy over 67%. Among the renters, there are substantial demographic disparities across neighborhoods and boroughs. Table 1.2 shows the median renter in Manhattan makes $67,000 a year, while the median renter in Queens makes $52,000, and in Brooklyn $44,000. In contrast, the median renter in the Bronx makes only $31,000. Similar patterns of disparity exist for education. In terms of race and ethnicity, the Bronx and Brooklyn have much higher proportions (greater than 35%) of African American households than Manhattan and Queens. Meanwhile, the Bronx and Queens have higher proportions of Hispanic households than Manhattan and Brooklyn.

Table 1.2: *Income and Demographic Characteristics of NYC Renters*

Based on the American Community Survey 2017 1-Yr estimates, the table summarizes the income and demographic characteristics of New York City renters across all boroughs, excluding Staten Island. Note that there are substantial variations in household income, education, race and ethnicity across the boroughs.

	All	**The Bronx**	**Brooklyn**	**Manhattan**	**Queens**
Annual Household Income ($k)					
0-20%	8	6	7	10	12
20-40%	24	17	22	31	32
40-60%	47	31	44	67	52
60-80%	83	54	78	126	83
80-100%	164	99	156	265	142
Education					
With College	38%	18%	36%	58%	32%
Race / Ethnicity					
White (non-Hispanic)	32%	7%	35%	47%	28%
African American	27%	36%	35%	16%	17%
Hispanic	32%	58%	21%	25%	31%
Asian	11%	2%	8%	12%	22%

The geographic variations of the household demographics allow me to evaluate the

the 2010 U.S. Census. I use other years from 2010 to 2017 for variables involving neighborhood changes. The estimates remain robust if I use the 5-Year Estimates.

distributional impact of Airbnb across these demographic characteristics. The reduced-form correlations in Table 1.3 reveal that there are more Airbnb listings in neighborhoods with more white, educated, and higher income residents. However, a full model is needed to translate this empirical pattern into welfare, which takes into account the equilibrium effects before arriving at the distributional implications.

Lastly, I augment the analysis with data from STR, which tracks supply and demand for hotels. Specifically, I obtain the daily aggregate number of hotel rooms sold in New York City from 2007 to 2018. It is used to construct a seasonality-based demand shifter to estimate the cost of peer supply.

Table 1.3: *Reduced-Form Correlations between Airbnb and Neighborhood Characteristics*

The *dependent variable* is the log of Airbnb share, namely, the share of housing units reallocated away from the long-term rental market. Observations are at the PUMA level. The neighborhood demographics are based on 2017 ACS data, whereas the neighborhood demographic changes are based on the changes from 2010 ACS to 2017 ACS. The correlation table shows that there are more Airbnb reallocation in neighborhoods that are relatively more white, educated, and higher-income. There is also some evidence that there are more Airbnbs in "gentrifying neighborhoods", as measured by improvements in the level of education.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pct White	2.834***							
	(0.780)							
Pct Black		-0.336						
		(0.822)						
Pct Hispanic			-3.438***					
			(0.852)					
Pct Asian				-0.211				
				(1.493)				
Pct College					5.244***			
					(0.784)			
ln(Median Income)						1.803***		
						(0.382)		
Chg in Pct College							0.0974**	
							(0.0475)	
Chg in Pct Black								-0.153***
								(0.0513)
N	52	52	52	52	52	52	52	52

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

1.3 Model Description

1.3.1 A Model of an Integrated Housing Market

In this section, I describe the main components of the model. I specify demand and supply in both the long-term and the short-term rental market, highlighting the key innovations brought about by Airbnb. First, the long-term and the short-term rental market become integrated on the supply side, improving allocative efficiency. Second, the utilization of existing housing units is increased when residents themselves offer space in their homes to host short-term visitors, especially during times when short-term rental demand is high.

A. Long Term Rental Demand I start with a model of residential housing choice using a random utility framework following McFadden (1978) and Bayer, Ferreira and Mcmillan (2007).[29] Each resident faces a discrete choice problem among all housing types in the city, and each housing type is defined by the neighborhood n in which the housing unit is located, as well as the physical characteristics of the unit, which include the number of bedrooms, age of the building, and indicators for the type of the building. As a result, the housing stock in New York City is divided into 1,050 such types. Within each housing type, there are N_h units that are not further differentiated beyond an idiosyncratic component $\epsilon^L_{i,j}$. Hence, the long-term rental utility of household i derived from housing unit j of type h is

$$u^L_{i,j} = \alpha^L_i p^L_h + \boldsymbol{\beta}^L_i \mathbf{X}^L_h + \xi^L_h + \epsilon^L_{i,j} \tag{1.3.1}$$

where the superscript L indicates quantities pertaining to the long-term rental market. p^L_h refers to the price of housing type h, and $\mathbf{X}^L_h$ includes both the physical and the neighborhood characteristics. The neighborhood characteristics include the percentage of the households with a college degree and the percentage of African American, Hispanic, and Asian house-

[29] I use the phrase "long-term rental" to capture the residential housing choices made by long-term residents. It is to be contrasted with the "short-term rental" market where visitors have a demand for housing for typically much shorter periods (e.g. at the daily level). Even though the exposition of the model appears to assume that all households are renters, it can be readily extended to incorporate owner-occupiers in the sense that the model captures the equivalent user-cost of the owner-occupiers.

holds. I also include a measure of its location amenity using the average commuting time.[30] In addition, I allow an unobserved quality component ξ_h^L that could be correlated with price. The price coefficient α_i^L and the coefficients on the housing characteristics β_i^L are determined in a flexible way based on the vector of observable demographics z_i, including household income, race, ethnicity, education, and family size.[31]

Each household i makes an optimal housing choice j by maximizing utility:

$$y_i^L = j \iff u_{i,j}^L > u_{i,-j}^L$$

The set of households who choose housing type h is simply the union of those who choose any housing unit j that is of type h:

$$A_h^L = \bigcup_{j:h(j)=h} \{z_i, \epsilon_{i,\cdot}^L : u_{i,j}^L > u_{i,-j}^L\}$$

The total long-term rental demand for housing type h can be obtained by integrating over all households

$$D_h^L(p_h^L, p_{-h}^L) = \int_{A_h^L} dP(\epsilon^L) dP_D^*(z)$$

where P_D^* represents the empirical distribution of demographic characteristics of all potential city residents.[32]

B. Long-Term Rental Supply The key assumption here is that the total number of housing units of each type in the city is fixed at S_h^F. In a market without the home-sharing platform, the city's housing market would be fully characterized by the market clearing price for each

[30]In the baseline model, I use the average commuting time of workers living in the neighborhood. In a richer model, since I observe the work location at the borough level, the location amenity could be modeled more finely as household-specific.

[31]The model assumes that the unobserved quality is vertical. This assumption is more reasonable when the vector of observable X_h and the vector of demographics z_i are sufficiently rich to capture horizontal sorting.

[32]The market size for living in a city is defined as the relevant metro market, which includes the city itself and the surrounding areas within a commutable distance. In practice, I focus on all the contiguous counties that surround New York City, namely Hudson, Nassau, Westchester, and Bergen.

housing type h in the long-term rental market:

$$\forall h: \quad D_h^L(p_h^L, p_{-h}^L) = S_h^F$$

However, with the home-sharing platform, I will proceed to specify the short-term rental supply before imposing market clearing.

C. Short-Term Rental Supply I categorize the supply of short-term rentals on Airbnb into two different types: those provided by absentee landlords, and those provided by residents.

The first type is provided by absentee landlords who reallocated housing units from the long-term to the short-term rental market. In this case, for every unit reallocated, there is one fewer unit available for long-term tenants.

The second type is provided by the residents directly. In this case, the resident of the housing unit chooses to supply space to the short-term rental market if she finds that the short-term rental income in a given period is greater than the value for her alternative personal use. This second type of short-term supply increases the utilization of housing units already occupied by residents.[33]

C1. Short-Term Rental Supply from Absentee Landlord With Airbnb, an absentee landlord owning a housing unit may now consider short-term rental as an alternative to long-term rental. For each housing type h on day t, the price of the unit in the short-term rental market is $p_{h,t}^A$.

The utility of accepting a short-term rental visitor on day t in housing unit j of type h is:

$$u_{j,t}^A = p_{h,t}^A + \nu_{j,t}^A \tag{1.3.2}$$

where $\nu_{j,t}^A$ represents the cost of operating the short-term rental (including cleaning and

[33]For example, residents may choose to rent out their guest bedroom when they do not have friends or family visiting. In other words, the room would not otherwise be available as a long-term rental unit for another household. Hence, in the baseline model, the supply of short-term rental space offered by residents does not crowd out available space for long-term rental. In the long run, households may adjust their long-term rental demand due to their expected short-term rental income. I consider its likely impact in the last section of the paper.

providing supplies, for instance). However, the decision to reallocate away from the long-term rental is made with the consideration of a much longer time period T (e.g. a year). Hence, an absentee landlord chooses to reallocate if he can make more money in the short-term rental market overall:

$$y_j^A = 1 \iff \frac{1}{T}\sum_{t=0}^{T} \max\{\, p_{h,t}^A + \nu_{j,t}^A,\, 0\,\} > p_h^L + \nu_j^L$$

where ν_j^L is the cost associated with operating property j in the long-term rental market and p_h^L is the long-term rental rate per day.[34] Hence, the total number of housing units of type h that is reallocated by absentee landlords is as follows

$$S_h^A(p_h^L, p_{h,\cdot}^A) = \int_{A_h^A} dP(\nu^A, \nu^L), \quad A_h^A = \bigcup_{j:h(j)=j} \{\nu_{j,\cdot}^A, \nu_j^L : \quad y_j^A = 1\}$$

On a given day t, the total short-term supply is simply the cumulative distribution of those who can operate profitably at the market rate, given that it has already been reallocated to Airbnb:

$$S_{h,t}^A(p_h^L, p_{h,\cdot}^A) = \int_{A_{h,t}^A} dP(\nu^A, \nu^L), \quad A_{h,t}^A = \bigcup_{j:h(j)=j} \{\nu_{j,\cdot}^A, \nu_j^L : \quad y_j^A = 1, \quad u_{j,t}^A > 0\}$$

C2. Short-Term Rental Supply from Residents As the second type of short-term rental supply, city residents can directly host short-term visitors in their homes, which is a more accurate reflection of the spirit of the "sharing economy." In this case, the utility derived from the supply of a short-term rental room by household i in neighborhood n on day t is

$$u_{i,t}^R = \alpha_i^R p_{n,t}^A + \boldsymbol{\beta}_i^R \mathbf{X}_{n,t}^R + \xi_{n,t}^R + \epsilon_{i,t}^R \tag{1.3.3}$$

where $p_{n,t}^A$ denotes the price of an Airbnb private room and $\mathbf{X}_{n,t}^R$ denotes a number of observable shifters, importantly including a constant term representing the negative cost of providing a room. The coefficient β_i^R in front of the constant is modeled as a function

[34] This specification abstracts away from discount rates, uncertainty, and any dynamic considerations such as the option value to convert in the future. In other words, the baseline model is static, but it still allows for short-term rental vacancies due to predictable price variations from short-term rental seasonality.

of household demographics, including age, education, income, and family structure. The price coefficient α_i^R is modeled as a function of household income. The model allows for an unobserved cost at the neighborhood-day level that may be correlated with the prevailing market price $p_{n,t}^A$. A household-specific idiosyncratic taste for home-sharing on a given day is included as $\epsilon_{i,t}^R$.

As an additional housing type, the product here is a private room for short-term use. Therefore, it is only differentiated at the neighborhood and day level.[35] Thus, a given household chooses to supply a room on a given day if the utility of hosting is greater than the outside option, which is normalized to zero. The total short-term rental supply by residents in neighborhood n on day t is the integral of such households in the neighborhood:

$$S_{n,t}^R(p_{n,t}^A) = \int_{A_{n,t}^R} dP(\epsilon^R) dP_{D_n}^*(z), \quad A_{n,t}^R = \{z_i, \epsilon_{i,t}^R : u_{i,t}^R > 0\}$$

where $P_{D_n}^*(\cdot)$ denotes the empirical distribution of the household demographics in neighborhood n.[36]

D. Short-Term Rental Demand The demand for short-term rental is characterized by a discrete choice problem where the choice set includes all short-term rental housing types available on the home sharing platform (including both entire homes of all types h and private rooms in all neighborhoods n) on a given day t, denoted by $D_{h,t}^A(p_{h,t}^A, p_{-h,t}^A)$.[37]

E. Market Equilibrium The long-term rental market is characterized by a sorting equilibrium. Namely, the equilibrium price vector p_h^L for all housing types ensures that the

[35]In other words, a private short-term rental room in a two-bedroom home is not differentiated from that in a three-bedroom home. Moreover, the estimation of the host surplus focuses on the gains derived from such private rooms as opposed to entire homes, as the private rooms are more likely to remain legal under the current regulation in New York City, which requires the permanent resident to be present.

[36]The distribution of the neighborhood demographics D_n is considered constant from day to day and thus does not have a subscript t.

[37]Thus, $D_{h,t}^A$ represents the residual demand for Airbnb after hotels. Because the focus of the paper is not on the welfare of short-term visitors, I do not explicitly estimate their utility. However, one can think of a visitor's utility for a particular short-term rental option being determined by its price and its product characteristics. Then, she makes the optimal choice among them.

demand for each housing type h equals its supply, which is the number of the underlying housing units less what have been reallocated to the short-term rental market:

$$\forall h: \quad D^L_h(p^L_h, p^L_{-h}) = S^F_h - S^A_h(p^L_h, p^A_{h,\cdot}) \tag{1.3.4}$$

The short-term rental market is analogously characterized by the market-clearing of each short-term rental lodging option every day. Namely, the equilibrium price vector $p^A_{h,t}$ for all short-term lodging options h ensures that the short-term rental demand equals the short-term rental supply. For notational simplicity, the short-term lodging options h include all the housing types in the long-term rental market, as well as a private room in neighborhood n.

$$\forall h, t: \quad D^A_{h,t}(p^A_{h,t}, p^A_{-h,t}) = S^A_{h,t}(p^L_h, p^A_{h,\cdot}) + S^R_{h,t}(p^A_{h,t}) \tag{1.3.5}$$

Bayer and Timmins (2005) provide the regularity conditions that guarantee the existence of the sorting equilibrium, namely that the errors ϵ^L and ϵ^A are drawn from a continuous and well-defined distribution function.[38]

1.4 Estimation and Results

In this section, I provide details on how the structural models are estimated. First, I estimate the long-term rental demand model using individual-level data from the cross-section of housing choices. Then, I construct my estimator for the short-term rental supply model, using aggregate data across multiple neighborhoods and time periods. In both cases, I highlight the moment restrictions needed and the identifying assumptions used. Lastly, I discuss the parameters estimated and the corresponding elasticities.

[38] In general, the uniqueness of the sorting equilibrium in the long-term rental market is not guaranteed if a household's long-term rental utility is affected by the choice of other individuals (e.g. preference for neighborhood racial and ethnic composition). Nonetheless, Bayer and Timmins (2005) also show that uniqueness becomes easier to sustain when there is a large number of available choices, which is likely to hold in the residential choice setting.

1.4.1 Estimating the Long-Term Rental Demand

In this section, I describe the two-step procedure for estimating the long-term rental demand using two sets of moment conditions.

In the first step, I estimate the heterogeneous coefficients by taking advantage of the individual-level choice data. I construct moment conditions that match market shares, as well as the covariance between the housing attributes and the demographic characteristics of the households living in those homes. In the second step, I estimate the linear coefficients. To address the concern that there may be an unobservable housing quality correlated with price, I construct a price instrument based on a home's location in the characteristics space while setting the unobserved quality to zero, following Berry, Levinsohn and Pakes (1999) and Bayer and Timmins (2007).

Recall that the utility for a household i considering home j of type h depends on its price p_h^L and its vector of housing attributes X_h^L.[39] Since α_i^L and $\boldsymbol{\beta}_i^L$ are specific to each household i, I parameterize them as the sum of two parts: The first part is common to all households, and the second part is a function of its observable demographic characteristics z_i, where the matrix π^L is fully saturated.

$$u_{i,j}^L = \alpha_i^L p_h^L + (\boldsymbol{\beta}_i^L)^T \mathbf{X}_h^L + \xi_h^L + \epsilon_{i,j}^L$$

$$\begin{bmatrix} \alpha_i^L \\ \boldsymbol{\beta}_i^L \end{bmatrix} = \underbrace{\begin{bmatrix} \alpha^L \\ \boldsymbol{\beta}^L \end{bmatrix}}_{\substack{\text{common to all;} \\ \text{part of the mean utility } \delta_h^L}} + \underbrace{\begin{bmatrix} \pi_{\alpha,1}^L \dots \pi_{\alpha,K}^L \\ \pi_{\boldsymbol{\beta},1}^L \dots \pi_{\boldsymbol{\beta},K}^L \end{bmatrix}}_{\substack{\text{household-specific;} \\ \text{part of heterogeneous utility } \lambda_{i,h}^L}} \begin{bmatrix} z_{i,1} \\ \vdots \\ z_{i,K} \end{bmatrix}$$

Denote $N_h = S_h^F - S_h^A$ as the number of housing units of type h.[40] Assuming the error ϵ^L is i.i.d. type I extreme value, I compute the choice probability analytically:

$$Pr(y_i^L \in h; \delta^L, \pi^L) = \frac{N_h \exp(\delta_h^L + \lambda_{i,h}^L)}{\sum_{h'} N_{h'} \exp(\delta_{h'}^L + \lambda_{i,h}^L)}$$

[39] I index the housing attributes by b, but omit it for simplicity when it does not cause confusion.

[40] Note that one of the housing types is the outside option, whose utility is normalized to zero.

where δ_h^L is the mean utility from housing type h that is common to all households and $\lambda_{i,h}^L$ is the heterogeneous part of the utility specific to household i:

$$\lambda_{i,h}^L = \Big(\sum_k \pi_{\alpha,k}^L z_{i,k}\Big) p_h^L + \Big(\sum_k \pi_{\boldsymbol{\beta},k}^L z_{i,k}\Big)^T \mathbf{X}_h^L$$
$$\delta_h^L = \alpha^L p_h^L + (\boldsymbol{\beta}^L)^T \mathbf{X}_h^L + \xi_h^L$$

Notice that the choice probability is well-defined as long as the parameters δ^L and π^L are provided.

Step 1 Moments In the first step, I construct moments that identify the heterogeneous parameters π^L and the mean utility δ^L.

Since I observe the actual individual-level choices $\mathbb{1}\{y_i^L \in h\}$ from the data, I construct moment conditions that match the market shares, as well as moment conditions that match the covariance between the housing attribute $X_{b,h}^L$ and the average characteristics z_k of households who choose h. For instance, I match the covariance between the number of bedrooms and the average size of the families who choose that type of house. Thus, the set of moment conditions that pin down δ^L and π^L are as follows:

$$\forall h: \quad \mathbb{E}\big[\, Pr(y_i^L \in h;\, \delta^L, \pi^L)\,\big] \quad = \quad s_h^L \tag{1.4.1}$$
$$\forall b, k: \quad \mathrm{Cov}\big(\, \mathbb{E}[z_k | y_i^L \in h;\, \delta^L, \pi^L],\ X_{b,h}^L\,\big) \;=\; \mathrm{Cov}\big(\, \bar{z}_{k,h},\ X_{b,h}^L\,\big) \tag{1.4.2}$$

where the left-hand side denotes the model predictions and the right-hand side is estimated from its empirical counterparts:

$$\hat{s}_h^L = \frac{1}{N}\sum_i \mathbb{1}\{y_i^L \in h\}, \quad \hat{\bar{z}}_{k,h} = \frac{1}{N}\sum_i \mathbb{1}\{y_i^L \in h\} z_{i,k}$$

where $N = \sum_{h'} N_{h'}$. Also, note that the market share moments Eq (1.4.1) take the expectation over all households i. The attribute covariance moments Eq (1.4.2) take the expectation over all housing types h, weighted by the probability that a housing type h is chosen.[41]

[41]The formulation here more closely follows the intuition provided in Berry, Levinsohn and Pakes (2004). Alternatively, Bayer, Ferreira and Mcmillan (2007) formulate an analogous set of moments using a maximum

The identification here relies only on individual rationality, together with the fact that the housing prices and the attributes p_h^L, X_h^L are exogenous from the perspective of a single household making the choice. Since each household is assumed to be infinitesimal, this condition holds.

Step 2 Moments In the second step, I construct moments that identify the linear parameters α^L and β^L as well as the unobserved quality ξ_h^L. Because I allow the market price p_h^L to be correlated with the unobserved quality ξ_h^L, I employ an identification strategy that takes advantage of the shape of the housing characteristics space, following Berry, Levinsohn and Pakes (1999) and Bayer, Ferreira and Mcmillan (2007).

Intuitively, a home that is situated in a crowded part of the housing attribute space has a low equilibrium price, regardless of its own unobserved quality. As such, I construct a price instrument using the characteristics of other homes in the city. Specifically, to characterize the impact of the attribute space on market prices, I compute an alternative vector of equilibrium prices p^{IV} as an instrument for the observed prices p^L by setting the unobserved quality to zero $\xi_h^L = 0$ and resolving the market-clearing conditions across all home types. Thus, the moment conditions for estimating the linear parameters are as follows:[42]

$$\forall h: \quad \mathbb{E}\left[Pr\big(y \in h; \, (p^{IV}, \alpha^L, \beta^L, \xi^L = 0; \, \pi^L, X^{\text{exog}})\big)\right] = s_h^L \tag{1.4.3}$$

$$\forall h: \quad \mathbb{E}\left[Pr\big(y \in h; \, (p^L, \, \alpha^L, \beta^L, \xi^L; \quad \pi^L, X)\big)\right] = s_h^L \tag{1.4.4}$$

$$\mathbb{E}\left[\xi^L p^{IV}\right] = 0 \tag{1.4.5}$$

$$\mathbb{E}\left[\xi^L X_b^{\text{exog}}\right] = 0 \tag{1.4.6}$$

Importantly, the construction of the alternative equilibrium price p^{IV} as an instrument requires a supply-side pricing equation. Unlike a typical Nash pricing equation, the supply-

likelihood estimator, setting its score to zero. In particular, the market share moments correspond to the derivative of the log-likelihood function with respect to the mean utility δ_h^L. The attribute covariance moments correspond to the derivative with respect to each heterogeneous parameter $\pi_{b,k}^L$.

[42]Here, I use the heterogeneous parameters π^L that are estimated from Step 1 in the choice probability.

side pricing equation in the housing context is a market clearing condition with a fixed housing supply, as stated in Eq (1.4.3).[43] Notice that for each guess of the demand parameter α^L, β^L, ξ^L, there exists a corresponding p^{IV}. Hence, the demand parameters and the price instrument are estimated jointly[44] so that all moment conditions listed above are satisfied.

The identification assumption needed is that there exists a subset of housing characteristics $\mathbf{X}^{\text{exog}}$ that is independent of the unobserved quality:

$$\xi_h^L \perp\!\!\!\perp \mathbf{X}^{\text{exog}}$$

Since p^{IV} is constructed by setting the unobserved component ξ^L to zero, then by construction, they are uncorrelated. In the housing context, the challenge is to find a subset of the housing attributes that may be considered reasonably independent of the unobserved ξ^L. Following Bayer, Ferreira and Mcmillan (2007), I use a vector of immutable attributes of the housing stock.[45] For instance, one may not think a two-bedroom home is necessarily of higher or lower unobserved quality when compared to a one-bedroom home. However, I exclude average neighborhood demographics, since it is likely that education and race affect unobserved housing quality.[46] In other words, the price instrument is computed by resolving the pricing equation Eq (1.4.3) with only the exogenous housing attributes.

Step 1 and Step 2 are performed sequentially to estimate the vector of heterogeneous coefficients and the linear coefficients. When the demographics vector z_i is normalized to have mean zero, the ratio of the linear coefficient over the price coefficient $-\beta_b^L/\alpha^L$ represents the willingness-to-pay of the average household for attribute X_b^L.

[43] There are other ways to form instruments based on the shape of the characteristics space, such as those used in Berry, Levinsohn and Pakes (1995). However, using a single price instrument p^{IV} to capture the equilibrium price impact of the attribute space is an approximation of the optimal instrument in the sense of Chamberlain (1987), Reynaert and Verboven (2014), and Berry, Levinsohn and Pakes (1999).

[44] In practice, an iterative procedure is used to find the fixed point.

[45] The attributes included are indicators of the number of bedrooms, type of building, age of building, average commuting time to work centers, and an indicator for being inside the city.

[46] Relatedly, since I do not acquire additional instruments for these endogenous neighborhood attributes, the model only estimates the linear coefficients β^L in front of the exogenous ones. In practice, one may think of the estimated $\hat{\xi}^h$ as the sum of the true unobserved ξ^h and the endogenous component $\beta^L X^{\text{Endo}}$, which I assume to be unchanged in the subsequent counterfactuals.

1.4.2 Estimating the Short-Term Rental Supply

In this section, I describe how I estimate the short-term rental supply. Although BLP methods are widely used by researchers to estimate demand systems (Berry, Levinsohn and Pakes, 1995; Nevo, 2000, 2001), I propose an adaptation so that it can be used to estimate a random-coefficient *supply* system. My adaptation takes advantage of the fact that the location of a housing unit on Airbnb is observed, allowing me to match the "market shares" of home-sharing supply in each neighborhood every day. As such, variations in the distribution of demographic characteristics across neighborhoods and variations due to short-term demand seasonality allow me to estimate the heterogeneity in cost. Lastly, since the high-frequency daily data results in a large number of market-share equations to match, I employ the MPEC procedure developed in Dubé, Fox and Su (2012) to improve the numerical performance of the estimator.

Recall that the utility that a resident host i living in neighborhood n derives from sharing a private room on Airbnb on day t depends on how she values the income from sharing $p^A_{n,t}$ compared to how costly it is to provide such short-term rental services. In order to capture the heterogeneity, I allow α^R_i and $\boldsymbol{\beta}^R_i$ to be household-specific. I parameterize them as the sum of a common component and a component that depends on its demographic characteristics.

$$
\begin{aligned}
u^R_{i,t} &= \alpha^R_i \, p^A_{n,t} + (\boldsymbol{\beta}^R_i)^T \, \mathbf{X}^R_{n,t} + \xi^R_{n,t} + \epsilon^R_{i,t} \\
\begin{bmatrix} \alpha^R_i \\ \boldsymbol{\beta}^R_i \end{bmatrix} &= \underbrace{\begin{bmatrix} \alpha^R \\ \boldsymbol{\beta}^R \end{bmatrix}}_{\text{common to all}} + \underbrace{\begin{bmatrix} \pi^R_{\alpha,1} \dots \pi^R_{\alpha,K} \\ \pi^R_{\boldsymbol{\beta},1} \dots \pi^R_{\boldsymbol{\beta},K} \end{bmatrix}}_{\text{household-specific}} \begin{bmatrix} z_{i,1} \\ \vdots \\ z_{i,K} \end{bmatrix}
\end{aligned}
$$

where $p^A_{n,t}$ denotes the prevailing price of an Airbnb private room in neighborhood n on day t. $\mathbf{X}^R_{n,t} = [1, \ t, \ t^2, Z_{month}, Z_{dow}, Z_{holiday}]^T$ captures features that contribute to the cost of sharing the room. It includes a constant term, a quadratic time trend, and dummies for month fixed effects, day-of-the-week fixed effects, and holiday fixed effects. $\xi^R_{n,t}$ represents an unobserved cost that varies at the neighborhood-day level.[47]

Since the constant term captures the (negative) cost of sharing a room such as the time and the hassle, I parametrize the coefficient β_i^R in front of it as a linear function of household income, age, education, and family structure. Also, I allow the price coefficient α_i^R to be a function of household income, permitting one's price sensitivity to differ by income.[48]

Because each resident is faced with a binary choice between sharing and not sharing, assuming the error is distributed as type I extreme value, I can analytically derive the quantity supplied:

$$s_{i,n,t}^R(\delta_{n,t}^R, \pi^R) = \frac{\exp(\delta_{n,t}^R + \lambda_{i,n,t}^R)}{1 + \exp(\delta_{n,t}^R + \lambda_{i,n,t}^R)}$$

where the mean utility and the heterogeneous utility are defined as follows

$$\lambda_{i,n,t}^R = \Big(\sum_k \pi_{\alpha,k}^R z_{i,k}\Big) p_{n,t}^A + \Big(\sum_k \pi_{\boldsymbol{\beta},k}^R z_{i,k}\Big)^T \mathbf{X}_{n,t}^R$$
$$\delta_{n,t}^R = \alpha^R p_{n,t}^A + (\boldsymbol{\beta}^R)^T \mathbf{X}_{n,t}^R + \xi_{n,t}^R$$

Importantly, note that the market share $S_{n,t}^R = \sum_{i\in n} s_{i,n,t}^R$ is based on the cost of sharing among all residents currently residing in neighborhood n.[49]

As the unobservable cost at the neighborhood-time level $\xi_{n,t}^R$ is allowed to be correlated with the price $p_{n,t}^A$, I instrument the short-term rental price using a measure of tourist demand seasonality. Specifically, I use the number of hotel visits to the entire city of New York on the same day but lagged by seven years. Because Airbnb was essentially irrelevant before 2011,[50] it should be free of reverse causality. Unobserved structural errors in the cost

[47]For example, $\xi_{n,t}^R$ captures unobservable time trends in the cost of sharing, such as the changes in the perceived risk of hosting strangers at home as technology diffuses. It also captures unobserved costs at the neighborhood level, such as fewer house cleaners in a given neighborhood.

[48]Note that in this utility specification, I only consider random-coefficients that can be projected onto observable demographics.

[49]Note that both renters and owner-occupiers are considered as potential suppliers of home-sharing.

[50]My data spans 2014 to 2018. Lagged by seven years, back in 2011, Airbnb had only 300 listings in NYC according to its early employees, and it has since grown over 200-fold. In 2007, Airbnb had not been founded yet. In contrast, the hotels in New York City sold over 25 million nights in 2007.
`https://medium.com/@jgolden/lessons-learned-scaling-airbnb-100x-b862364fb3a7`

of sharing should not affect the total number of New York hotel visits from seven years ago. Moreover, if one worries that the seasonality in the cost of supply remains correlated with lagged hotel demand, I also add a host of calendar-related controls, including month fixed effects, day-of-the-week fixed effects, and holiday fixed effects. What kind of variation does the price instrument leave us with? For example, it captures the impact of foreign holidays. They are persistent over time and affect hotel demand in NYC through increased tourism demand, but they are plausibly uncorrelated with the cost of home-sharing by the city residents.

The moment conditions to match include the market shares in each neighborhood n every day, together with the exogeneity of the linear shifters:

$$\forall n, t: \quad \mathbb{E}_{D_n^*}[\, s_{i,n,t}^R(\delta_{n,t}^R, \pi^R)\,] = s_{n,t}^{R,o} \tag{1.4.7}$$

$$\mathbb{E}[\xi^R Z] = 0 \tag{1.4.8}$$

where $Z = [p^{R,IV}, \mathbf{X^R}]$ includes the lagged hotel visits as the price instruments and $s_{n,t}^{R,o}$ denotes the observed market share.

To numerically estimate the supply system with over 70,000 market-share equations, I cast the problem as a minimization routine over the GMM objective. The Mathematical Programming with Equilibrium Constraints (MPEC) specification was developed in Dubé, Fox and Su (2012).

$$\begin{aligned} \min_{\delta_{n,r}^R, \alpha^R, \beta^R, \pi_\cdot^R, \eta} \quad & \eta^T W \eta \\ s.t. \quad \forall n, t: \quad & S_{n,t}^R(\delta_{n,t}^R, \pi_\cdot^R) = S_{n,t}^{R,o} \\ & \eta = Z'(\delta^R - \alpha^R p^A - \boldsymbol{\beta}^R \mathbf{X}^R) \end{aligned}$$

where W denotes the optimal weighting matrix. The primary advantage of this estimation method is the sparsity structure of the Jacobian and the Hessian. Namely, the mean utility $\delta_{n,t}^R$ only affects the equilibrium in market (n, t) and does not affect other markets.[51] It

[51]See appendix for the derivations of the relevant sparse matrices.

allows the optimizer to perform better numerically, especially when the number of markets is large.

1.4.3 Parameter Estimates

Parameter Estimates from the Long-Term Rental Demand In this section, I describe the key parameters estimated for the long-term rental demand, including both the heterogeneous preference parameters and the linear parameters. Overall, the market-share moments and the attribute covariance moments produce sensible estimates for how housing preferences vary by demographic characteristics. Then, the linear parameters are jointly estimated with the price instrument, providing measures of willingness-to-pay for each household.

Step 1 Results By matching the covariance between housing attributes and household demographics, the estimated heterogeneous parameters reflect sensible choice patterns. Instead of showing the raw parameters, Table 1.4 shows the more interpretable willingness-to-pay in monthly dollars, calculated as $-\pi^L_{b,k}/\alpha^L$, where I take the price coefficient from the second step. The columns in Table 1.4 correspond to the vector of household demographics, whereas the rows correspond to the vector of housing and neighborhood attributes.[52]

Larger households have higher willingness-to-pay for more bedrooms. I also find strong clustering preferences for race and education. African American, Hispanic, and Asian households have significantly higher willingness-to-pay for neighborhoods with higher percentages of their own race and ethnicity. For example, an Asian household is willing to pay \$410 per month to live in a neighborhood that is one standard deviation higher in its percentage Asian.[53] In addition, educated families cluster in neighborhoods with higher

[52]Neighborhood attributes are parsimoniously defined to include all attributes that are the same for all homes within that neighborhood. Specifically, it includes neighborhood demographics such as education and race. It also includes the neighborhood's location amenity as measured by average commute time. Lastly, I include an indicator for the inside option, namely, in New York City.

[53]I do not attempt to explain whether it is an Asian household's preference for its neighbor's ethnicity or for neighborhood businesses that cater to its preferences. I also do not differentiate between whether it is due to housing market preference or discrimination. These parameters simply reflect the underlying choice patterns in the data. As such, I assume whatever preference or discrimination present in the data continues to be present in the counterfactual analysis.

proportions of educated families. The outside option is more valuable for higher income and larger families, whereas racial and ethnic minorities prefer to live in the city, all else being equal. The heterogeneity in front of the price coefficient reflects differential price sensitives, where higher-income and more educated households tend to be less price sensitive.

Step 2 Results Since the market share moments also produce an estimate of the mean utility, Step 2 estimates the linear coefficients, including both the price coefficient and the average willingness-to-pay for each housing attribute.

Table 1.5 column (2) shows that the price coefficient α^L is -2.04, using the instrumented estimator. The F-stat is 15.7. The use of the price instrument has a significant impact, compared to the OLS specification in column (1). To make the coefficients more interpretable, I also transform these coefficients in terms of the willingness-to-pay ($-\beta_b^L/\alpha^L$) of the average household. Table 1.5 column (3) shows that the average household has a higher willingness-to-pay for more bedrooms. It also prefers pre-war structures and fewer units in the building. The average household is willing to pay \$383 for one standard deviation reduction in commuting time.

Since the entire housing stock in the market is divided into 1,050 housing types based on its neighborhood and housing attributes, the demand elasticities estimated vary by housing types. Nonetheless, the average price response to a 1% reduction in the supply of all housing types in NYC is estimated to be 1%, which implies a price elasticity of the aggregate demand of approximately 1.[54]

Parameter Estimates from the Short-Term Rental Supply In this subsection, I describe the key parameters estimated for the supply of short-term rentals by resident hosts.

Table 1.6 column (4) summarizes the coefficients for the main specification. Matching the supply in each neighborhood everyday for four years results in 75,895 market-share equations to match. The linear price coefficient is 0.056, which implies an average supply

[54] Recall that the outside option of the model is to reside in the bordering counties of NYC, including Hudson, Nassau, Westchester, and Bergen.

Table 1.4: *Parameter Estimates for the Long-Term Rental Demand Model (Part I)*

This table provides the heterogeneous coefficients estimated from *Step 1* of the long-term rental demand model. The coefficients on housing attributes are presented as willingness-to-pay in terms of monthly dollars. The coefficient on price is applied to the monthly rent ($k). The omitted categories are studios, those built prior to 1940, and buildings with fewer than 5 units. Neighborhood characteristics are standardized to variance 1. The standard errors for the ratios are computed using the delta method. I highlight the most significant demographic characteristics for each housing and neighborhood attribute. Notice strong demographic clustering along race, ethnicity, and education.

WTP ($)	Ln Income	HH Size	Black	Hispanic	Asian	College
Housing Characteristics						
One Bedroom	75.8	255.1	-40.3	-67.0	-177.1	-86.4
	(31.2)	(78.6)	(55.1)	(54.5)	(75.6)	(48.3)
Two Bedroom	59.1	520.5	98.1	-24.2	-273.4	-212.6
	(28.8)	(156.4)	(63.8)	(54.1)	(101.0)	(76.9)
Three Bedroom	32.4	717.6	143.8	-37.2	-329.1	-214.6
	(28.7)	(214.9)	(80.7)	(64.8)	(121.0)	(82.6)
Four Bedroom	85.0	884.9	-206.5	-328.3	-297.9	-244.9
	(66.7)	(266.0)	(172.3)	(171.9)	(170.2)	(134.0)
Built After 1980	22.4	-35.9	157.9	55.1	42.3	-27.8
	(14.7)	(13.7)	(56.4)	(33.3)	(36.7)	(25.0)
Built 1940-1980	-102.6	6.1	137.0	125.4	-84.8	58.6
	(34.3)	(10.1)	(55.6)	(50.6)	(47.5)	(32.6)
5+ Units	9.7	58.6	-3.8	-118.4	-110.6	-6.7
	(9.5)	(18.6)	(25.9)	(41.3)	(40.1)	(17.6)
Nbhd. Characteristics						
Pct Black (Std)	56.7	-47.9	774.4	330.6	272.1	67.6
	(20.3)	(15.9)	(232.6)	(101.2)	(86.6)	(27.9)
Pct Hispanic (Std)	56.3	-22.8	376.9	469.3	221.5	94.8
	(19.9)	(9.4)	(115.8)	(141.5)	(71.5)	(33.7)
Pct Asian (Std)	47.8	-14.3	98.2	138.0	410.0	-37.7
	(16.9)	(7.1)	(39.0)	(44.9)	(123.9)	(19.2)
Pct College (Std)	145.9	-54.0	185.9	37.0	93.8	260.2
	(45.7)	(18.5)	(68.1)	(32.2)	(44.7)	(81.7)
Inside NYC	-337.8	-421.2	120.0	29.1	299.0	-2.3
	(106.9)	(128.6)	(97.6)	(83.9)	(129.4)	(68.6)
Commuting Time (Std)	38.7	-6.3	127.7	50.7	210.4	8.4
	(19.9)	(11.9)	(56.9)	(43.0)	(76.8)	(33.1)
Monthly Rent	0.33	-0.03	-0.36	-0.23	-0.18	0.21
	0.10	0.02	0.13	0.10	0.10	0.08

Table 1.5: *Parameter Estimates for the Long-Term Rental Demand Model (Part II)*

This table provides the linear coefficients estimated from *Step 2* of the long-term rental demand model. The *dependent variable* is the mean utility of each housing type δ_h^L estimated from *Step 1* using individual-level data. If the long-term rental prices were assumed to be exogenous, column (1) shows that the price coefficient is biased. Column (2) shows the IV regression where I use the alternative equilibrium price vector at $\xi^L = 0$ as the instrument for the observed prices. In Column (3), I show the willingness-to-pay for each of the housing and neighborhood attributes for the average household.

	(1) OLS	(2) Instrumented	(3) ($) WTP Mo.
Monthly Rent ($k)	0.0213	-2.044***	
	(0.0341)	(0.609)	
One-Bedroom	0.425***	0.929***	454.5***
	(0.0447)	(0.188)	(78.2)
Two-Bedroom	0.528***	1.325***	648.2***
	(0.0465)	(0.280)	(93.5)
Three-Bedroom	0.271***	1.392***	681.0***
	(0.0555)	(0.393)	(76.7)
Four-Bedroom	-0.179***	0.904*	442.3*
	(0.0668)	(0.505)	(162)
Built After 1980	-0.114***	0.139	68.2
	(0.0402)	(0.145)	(60.9)
Built 1940 to 80	-0.00917	-0.242**	-118.4**
	(0.0337)	(0.105)	(43.9)
5+ Units	0.00182	-0.209**	-102.3**
	(0.0282)	(0.0974)	(41.2)
Commuting Time (Std)	0.119***	-0.782***	-382.6***
	(0.0215)	(0.279)	(28.2)
Inside NYC	-1.026***	2.536**	1240.7***
	(0.0683)	(1.036)	(147)
N	1050	1050	1050

elasticity of 5.96. The instrument used is the lagged hotel bookings in the city as a measure of tourist demand seasonality. The F-stat is 25.4. In comparison, in columns (1) and (2), the short-term rental prices are assumed to be exogenous, and the price coefficients obtained are biased downward significantly, suggesting the presence of unobserved costs. Besides the quadratic time trend, the main specification also includes an array of calendar-fixed effects (namely, month FE, day-of-the-week FE, and holiday FE). Column (3) shows that these fixed effects have a moderate effect on the price coefficient. The standard errors are clustered at the neighborhood level.

The average cost of home-sharing is high, estimated at \$224 per night, reflecting the fact that the overall share of hosts as a percentage of all residents is still low. In terms of cost heterogeneity, I find that lower-cost suppliers tend to be young and educated households with no children. Dividing the non-linear coefficient π_k^R by the price coefficient α^R, I find that having a college degree is associated with a reduction of \$59 per night in the cost of home-sharing. Having children in the home increases the sharing cost by \$47 per night. Being ten years younger is associated with a reduction of \$18.

Moreover, I find a negative interaction between household income and price, which suggests that lower-income households are more price sensitive. The average supply elasticity would increase to 6.70 for those with a one-standard deviation lower in income.

1.5 Counterfactual Analysis

In this section, I provide estimates of the welfare and distributional impact of Airbnb in two steps. First, I discuss the welfare losses via the rent channel due to housing reallocation by absentee landlords. Second, I discuss the welfare gains via the host channel as residents act as peer suppliers. In the end, I discuss the net effects and the potential implications for the social planner.

Table 1.6: *Parameter Estimates for the Short-Term Rental Supply Model*

This table provides the estimated parameters for the short-term rental supply by resident hosts, using the MPEC procedure. The standard errors are clustered at the neighborhood level. In columns (1) and (2), prices are assumed to be exogenous and not instrumented. In columns (3) and (4), the total hotel bookings in the city (lagged by seven years) are used as the instrument for prices in the short-term rental market. The instrument is strong, with an F-stat of 25.4 in column (4) when controlling for month FE, day of week FE and holiday FE (including Christmas and New Year's Eve) in the corresponding linear specification. Column (5) provides the costs in dollar terms (per diem). Overall, the average cost to host is high, and the low-cost suppliers are those with a college degree, young, and have no children. The average price elasticity is 5.96. Since the interaction between household income and price is negative, low-income households are more elastic, averaging 6.70 for one-standard deviation lower in income.

		(1) Naïve	(2) Naïve	(3) IV	(4) IV	(5) ($) per diem
Linear Coef.	*Non-Linear Coef.*					
Price		0.006	0.007	0.052	0.056	
		(0.002)	(0.001)	(0.002)	(0.002)	
	x ln(income)	-0.018	-0.018	-0.018	-0.011	
		(0.001)	(0.002)	(0.003)	(0.006)	
Cost		15.44	15.51	22.07	21.36	224.3
		(0.10)	(0.09)	(0.12)	(0.11)	(12.7)
	x Has College	-1.17	-2.55	-3.47	-3.27	-58.9
		(0.68)	(0.24)	(0.27)	(0.25)	(4.8)
	x Has Children	2.40	2.58	1.95	2.60	46.7
		(0.42)	(0.36)	(0.53)	(0.44)	(8.1)
	x Age (yr)	0.094	0.093	0.091	0.097	1.8
		(0.005)	(0.005)	(0.006)	(0.006)	(0.1)
	x ln(income)	0.24	-0.14	-0.39	-0.29	-5.1
		(0.09)	(0.13)	(0.26)	(0.48)	(8.7)
Quad. Time		Yes	Yes	Yes	Yes	
Month FE		No	Yes	No	Yes	
Day of Week FE		No	Yes	No	Yes	
Holiday FE		No	Yes	No	Yes	
N		75,895	75,895	75,895	75,895	

1.5.1 The Distributional Impact via the Rent Channel

In this section, I conduct the counterfactual analysis to estimate the impact of Airbnb through the rent channel. Specifically, I compute the vector of counterfactual long-term rental prices p_h^L when all the housing units on Airbnb are "returned" to the long-term rental market.

Overall, I find that housing reallocation by absentee landlords aggregates to a material welfare impact for all renters in the city. Nonetheless, I find that the most significant welfare losses are suffered by renters who are higher-income, more educated, and white.

The patterns of the distributional impact are primarily driven by the geographical patterns of Airbnb penetration but compounded by the clustering preference along demographic lines. Moreover, severe housing supply restrictions result in a general equilibrium price impact of Airbnb being elevated across all NYC neighborhoods.

The Counterfactual Specification

Given the fully estimated long-term rental demand model, I recompute the counterfactual vector of long-term rental prices across all housing types by ensuring that housing demand equals housing supply, when the reallocated units are returned to the supply.

$$\forall h : D_h^L(p_h^{L,\text{ No Airbnb}}, p_{-h}^{L,\text{ No Airbnb}}) = S_h^F \tag{1.5.1}$$

Following McFadden (1978) and Small and Rosen (1981), the compensating variation can then be produced analytically when the errors are logit:

$$CV_i^L = \frac{1}{\alpha_i^L}\Big(\ln \sum_{j\in\mathcal{S}^F\setminus\mathcal{S}^A} \exp\left(V_{i,j}^L\right) - \ln \sum_{j\in\mathcal{S}^F} \exp\left(V_{i,j}^{L,\text{No Airbnb}}\right)\Big) \tag{1.5.2}$$

where $V_{i,j}^L = \alpha_i^L p_h^L + \beta_i^L X_h^L + \xi_h^L$ and $V_{i,j}^{L,\text{ No Airbnb}} = \alpha_i^L p_h^{L,\text{ No Airbnb}} + \beta_i^L X_h^L + \xi_h^L$ represent the non-idiosyncratic component of the utility for household i over home j of type h at the actual and the counterfactual equilibrium prices respectively.[55] $\mathcal{S}^F$ denotes the set of total

[55]Notice that the only difference between $V_{i,j}^L$ and $V_{i,j}^{L,\text{ No Airbnb}}$ is the price, while all other characteristics

physical structures available and $\mathcal{S}^A$ denotes the set of housing units observed to have been reallocated by absentee landlords.[56] Recall that a housing unit is considered to have been reallocated to the short-term rental market if it is marked "available" on Airbnb for over 180 days in 2018.[57] Across New York City, it averages to 0.68% of the rental housing stock but with remarkable heterogeneity across neighborhoods, as illustrated in Figure 1.6. When the counterfactual price vector is recomputed, I find an overall rent increase of 0.71%, also with substantial heterogeneity. Since the overall welfare impact depends not only on the price change in one's own housing type but also the price changes in all other housing types, I present the main results using the compensating variation measure as defined by Equation (1.5.2) and return to a discussion of the counterfactual equilibrium prices afterward.

In the following two sections, I first describe the distributional impact via the rent channel in terms of a variety of demographic characteristics such as household size, race and ethnicity, education, and household income. Then, I provide some characterizations to better understand the distributional patterns found. In particular, I explore (i) the role of geography, (ii) the "spillover" from one neighborhood to other neighborhoods, and (iii) the role of demographic clustering.

X_h^L and ξ_h^L are kept unchanged. Moreover, in computing the counterfactual price equilibrium, the value of the outside option is also kept unchanged.

[56]This formulation makes it explicit that the overall welfare impact of a supply squeeze on renters is comprised of two related components: the impact through an increase in the equilibrium rent and the impact through a reduction in the choice set. In particular, a consistent estimate of the compensating variation takes into account both components.

[57]This is a conservative assumption in ensuring that the unit is no longer available in the long-term rental market. There is also a concern about the selection on quality. Since the long-term rental model captures the unobserved quality for each housing type, insofar as Airbnb reallocation is more concentrated in low-quality housing types, the model fully captures this aspect. However, insofar as Airbnb reallocation might tilt towards lower-quality housing units *within* a given housing type, the model no longer captures it. Indeed, if renters in lower-quality units are more price elastic, then the counterfactual is an upper-bound of the price impact of Airbnb.

Distributional Impact by Demographic Characteristics

Overall, I find the average compensating variation via the rent channel is a loss of \$138 per year, whereas the median loss is \$128 per year and the standard deviation is \$29. To put this in perspective, the median renter's annual household income is about \$47,000. Although the magnitude translates to only 25 basis points of the annual income, the increase in equilibrium rents affects all 2.1 million households who are renters in the city. When aggregated across all renters, it amounts to a direct transfer to property owners of \$200mm per year, or \$2.7bn in NPV terms.

Next, to understand the differential welfare impact along demographic lines, I compute the compensating variation for each household CV_i^L and aggregate them into their respective categories. Importantly, this allows me to take into account the entire correlation matrix across demographic characteristics based on the empirical observations.

For the rest of this subsection, when I refer to "households", I restrict the analysis to all households that are renters. The welfare impact described in this subsection is also restricted to the rent channel via the long-term rental market, to be distinguished from the welfare impact via the host channel in the short-term rental market, which will be discussed in Section 1.5.2.

Household Size I find that Airbnb results in larger welfare losses via the rent channel for smaller households. Figure 1.9a shows that the median welfare loss for households of size one is \$134 p.a. with a standard deviation of \$32. The median welfare loss for households of size four is \$116 p.a. Even though there remains significant variation within each household size category, the most negative welfare losses are still concentrated in small households.

In the Airbnb data, smaller housing units are disproportionately more prevalent in the short-term rental market, compared to its underlying availability. Over 80% of the entire homes listed on Airbnb (with over 180+ days availability) has fewer than two bedrooms in the NYC market. In comparison, 46% of the housing stock available for long-term rental has fewer than two bedrooms, as shown in Figure 1.7. This underlying pattern of Airbnb usage is one of the key drivers for why the welfare impact is more concentrated on smaller households. Nonetheless, despite the large differences in the bedroom count distribution, the differences in the welfare impact across households are much reduced. This is because a reduction in one housing type creates equilibrium price effects on all other housing types as well, which I characterize in greater detail in the next subsection.

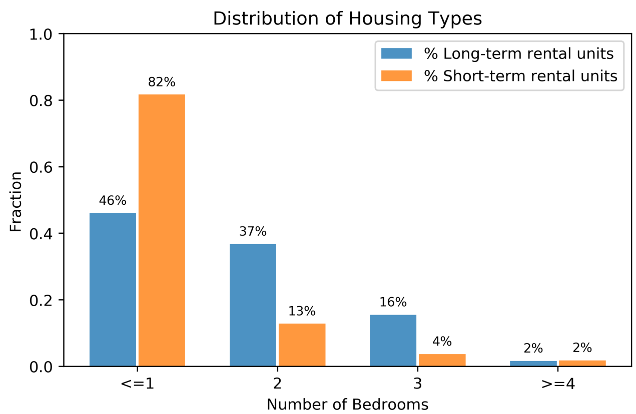

Figure 1.7: *Housing Types Comparison*

A comparison of housing types available on the long-term and the short-term rental market respectively. Notably, smaller housing types are much more prevalent on the short-term rental market than on the long-term rental market.

Race and Ethnicity I find that Airbnb results in larger welfare losses via the rent channel for white renters as compared to African American, Hispanic, or Asian renters. Figure 1.9b shows that the median welfare loss for white households is \$152 p.a. with a standard deviation of \$35 and a significant left tail. The median welfare loss for African American renters is \$134 p.a. The median welfare loss for Hispanic renters is \$113 p.a. The median welfare loss for Asian renters is \$127 p.a. Overall, when measured in dollar terms, the increased rent due to Airbnb reallocation hurts white renters the most, when compared to minority renters.

Figure 1.13 shows that there is significant clustering of housing choices along demographic lines. However, comparing it to the map of Airbnb penetration in Figure 1.6,

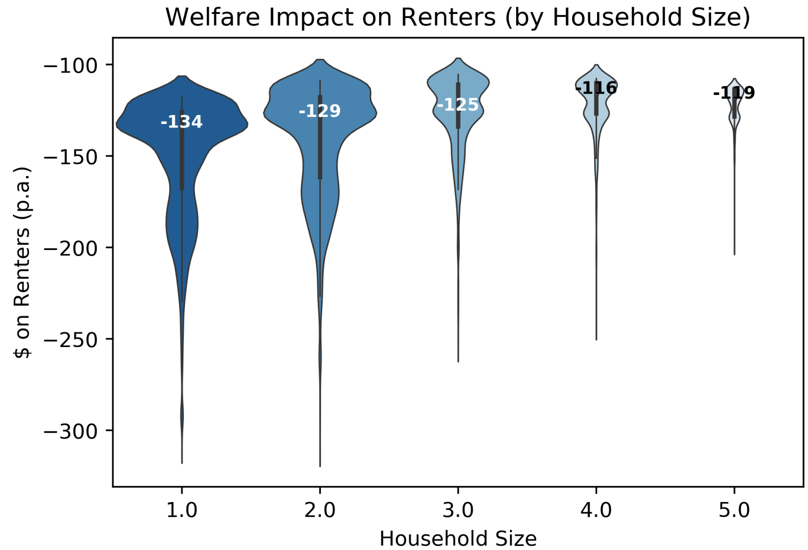

(a) Welfare impact via the rent channel by household size

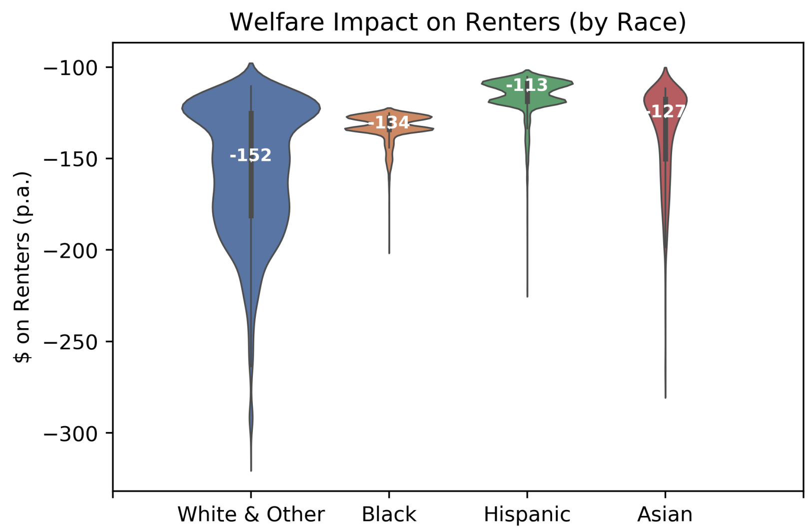

(b) Welfare impact via the rent channel by race and ethnicity

Figure 1.9: *Welfare Impact (Rent Channel) by Household Size and Race*

The welfare loss is suffered more heavily by smaller households. In terms of race and ethnicity, the welfare loss is suffered more heavily by white renters. The labeled numbers indicate the category median. The width of each kernel density plot corresponds to the frequency of the category in the population. The mini-box plot in the center indicates inter-quartile range using the thick black line, and 1.5x inter-quartile range using the thin black line.

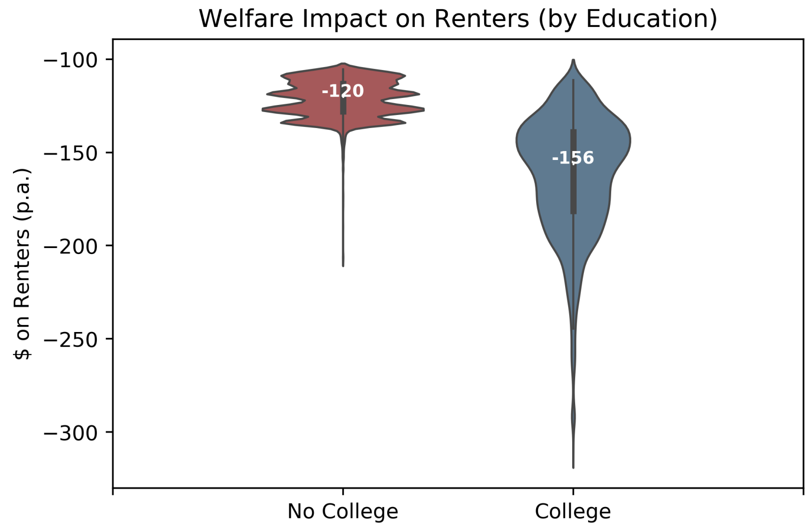

(a) Welfare impact via the rent channel by education

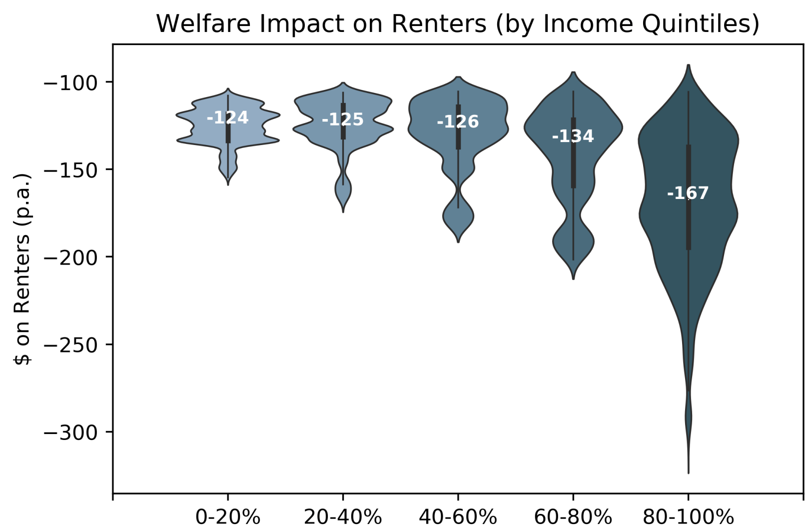

(b) Welfare impact via the rent channel by household income

Figure 1.11: *Welfare Impact (Rent Channel) by Education and Income*

The welfare loss is suffered more heavily by educated and higher-income renters. The labeled numbers indicate the category median. The width of each kernel density plot corresponds to the frequency of the category in the population. The mini-box plot in the center indicates inter-quartile range using the thick black line, and 1.5x inter-quartile range using the thin black line.

one could visually discern that neighborhoods with more Airbnb rentals tend to have greater percentages of white households. Hispanic and Asian households appear to be somewhat less segregated and are generally not concentrated in regions of heightened Airbnb activities. This is also shown in Table 1.3, where I compute the correlation between neighborhood demographics and Airbnb reallocation, quantifying the patterns seen on the maps.[58] Although a municipal planner may choose to place different social welfare weights on different households, this analysis highlights that welfare implications depend heavily on the geographic patterns of Airbnb activity.

Education I find that Airbnb results in larger welfare losses via the rent channel for those with more education. Figure 1.11a shows that the median welfare loss for renters with college degrees[59] is $156 p.a. with a standard deviation of $31. The median welfare loss for households without college degrees is $120 p.a.

Similar to the previous analysis, Figure 1.14 shows the average education attainment across neighborhoods. A comparison with the geographic pattern of Airbnb activity in Figure 1.6 shows that many of the neighborhoods with higher levels of Airbnb penetration also have more educated households. Moreover, given the strong preference for those with college degrees to live in neighborhoods that have more educated households, these renters affected by Airbnb in the city center are more likely to substitute to other educated neighborhoods even if they are further away, thus "spreading" the impact to those who are more similar to them demographically. I discuss this in greater detail in Section 1.5.1.

Household Income I find that Airbnb results in larger welfare losses via the rent channel for those with higher income. Figure 1.11b shows that the median welfare loss for

[58]It is also interesting to note that even though Airbnb activity does not appear particularly correlated with the percentage of African American households in the neighborhood, Airbnb activity tends to be high in "gentrifying neighborhoods". These neighborhoods may be identified in terms of increasing levels of education as well as declining numbers of minority households. Changes in the economic fundamentals are likely drivers for both neighborhood demographic changes and popularity among Airbnb guests. However, insofar as Airbnb reallocation further drives up rents in these areas, it can be viewed as causing the same type of changes as the ongoing gentrification process.

[59]Based on the highest educational attainment of the top earner of the household.

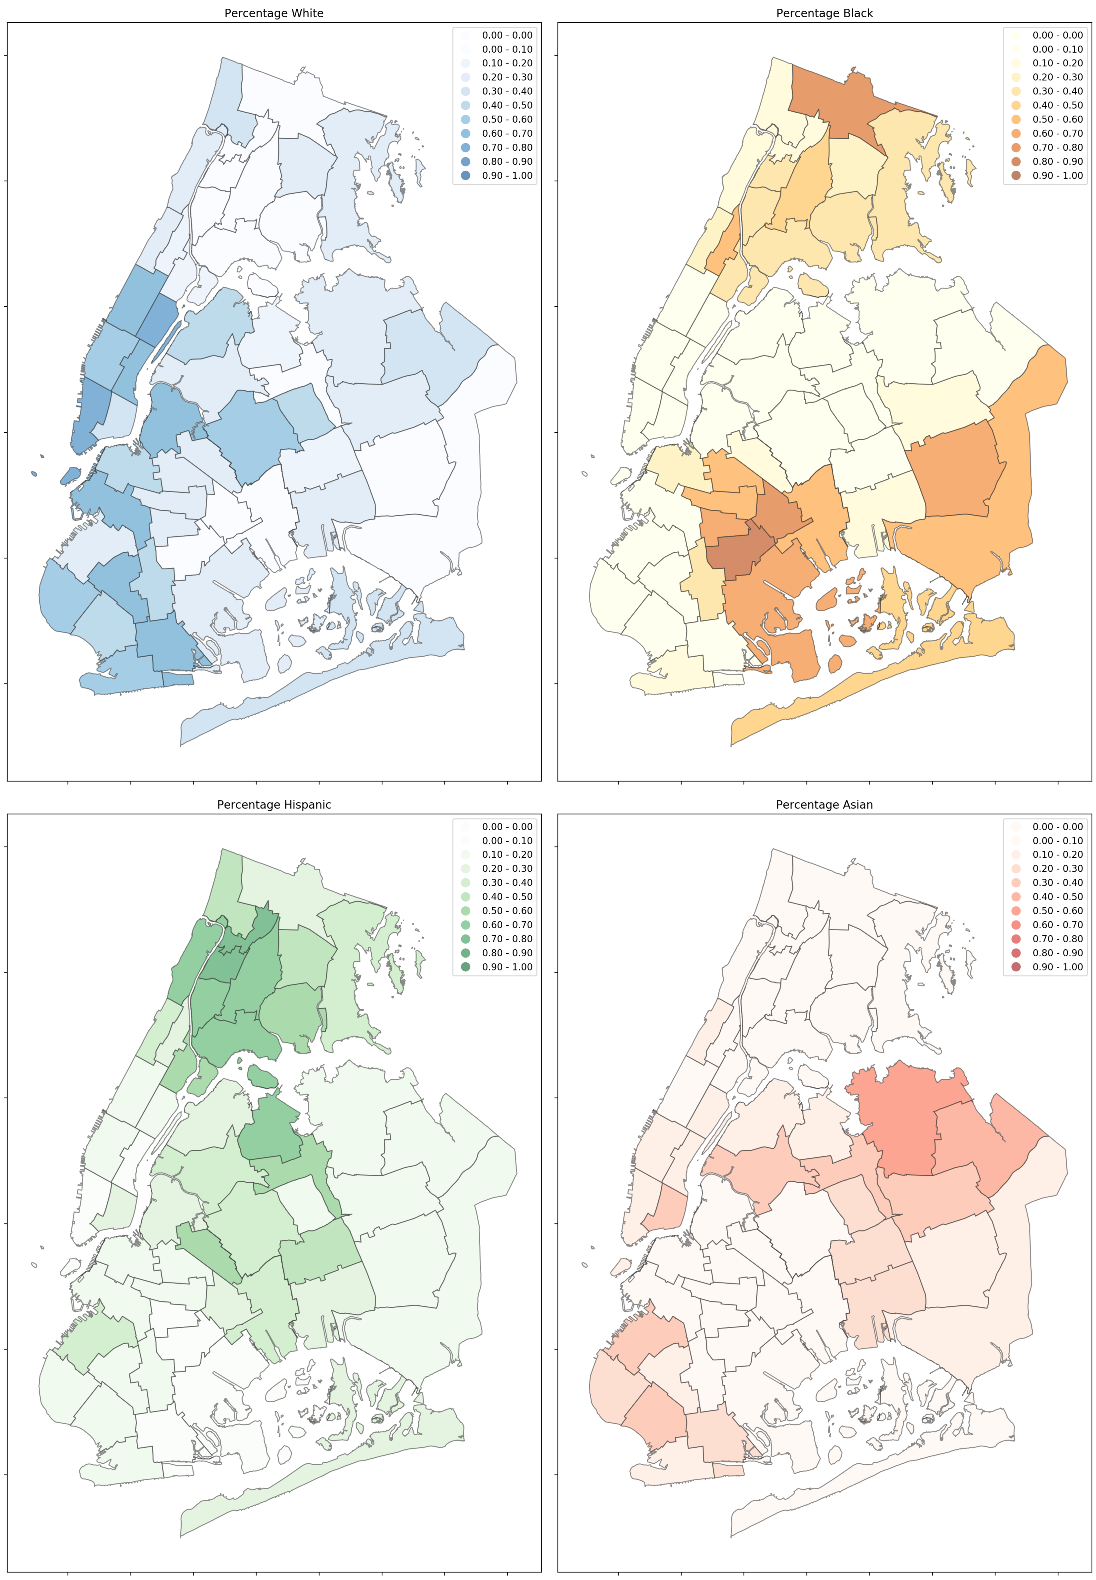

Figure 1.13: *Race and Ethnicity across Neighborhoods*

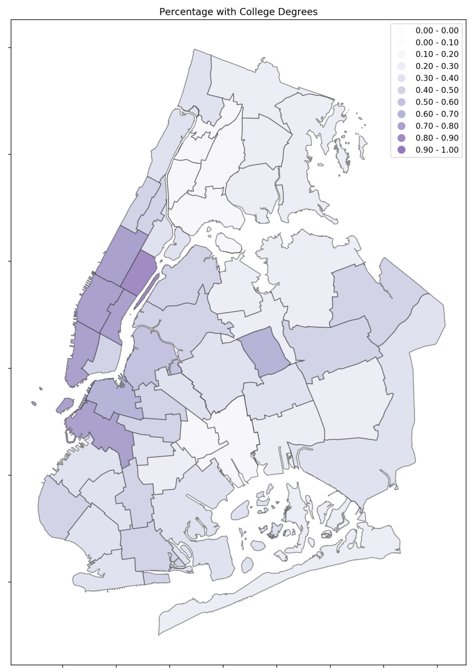

Figure 1.14: *Education Attainment across Neighborhoods*

renters in the top income quintile is \$167 p.a. with a standard deviation of \$38. The median welfare loss for renters in the bottom income quintile is \$123 p.a.

Consistent with the positive correlation between Airbnb activity and median neighborhood income shown in Table 1.3, the counterfactual analysis quantifies the extent to which the welfare losses are greater for higher-income renters. Besides the geography of Airbnb entry patterns, another factor is that the price coefficient α_i^L is smaller in magnitude for higher-income households, indicating that their willingness to pay for housing amenities $-\beta_i^L/\alpha_i^L$ is higher. As a result, when faced with a reduction in the housing supply, the requisite compensating variation for higher-income households tends to be larger, although I show in the next subsection that this factor is secondary to geography.

Decomposing the Distributional Impact

In this subsection, I discuss how the estimated welfare is connected with the data and the model. In particular, I explore (a) the out-sized role that geography plays in determining the distributional impact, (b) the general equilibrium price effect because of housing supply restrictions, and (c) the role of cross-elasticity due to demographic clustering.

A. The Role of Geography Given that I find that the higher-income and more educated renters experience larger welfare losses when housing units are reallocated to Airbnb, how much of the welfare differences across demographic groups is driven by the geographic pattern of Airbnb versus by the fact that these households tend to have a higher willingness-to-pay for housing attributes?

To disentangle these two different channels, I conduct an alternative counterfactual analysis if the housing units on Airbnb were distributed *uniformly* across geographic space and housing types. I find that less than a quarter of the distributional differences are attributable to their higher willingness-to-pay, whereas about three quarters of the differences are attributable to the geographical patterns of Airbnb activity. In the hypothetical counterfactual, if Airbnb had entered uniformly across space, Table 1.7 shows there would be only a $1 difference in terms of the median welfare impact between a household of size one and a household of size five. Similarly, the gap between the median impact on white renters and Hispanic renters would decrease from $39 to $7. The median welfare difference between renters with and without a college degree would decrease from $36 to $9. The gap between the top and bottom income quintiles would decrease from $43 to $9. In other words, since higher-income, and educated renters desire housing and location amenities that are particularly valuable to short-term renters, their housing units are more likely to be reallocated. As a result, they become the demographic group that is affected the most. The effect is further exacerbated as their willingness-to-pay also tends to be higher.

B. The Equilibrium Effects of Supply Restrictions Given that the housing supply in New York City is difficult to expand, the supply restrictions exacerbate the equilibrium price response when the city experiences a supply squeeze in the long-term rental market. I first provide a theoretical derivation to decompose the overall equilibrium effect into (i) the direct effect of a supply reduction on a given housing type, (ii) the spillover effect of a supply reduction on other housing types, and (iii) the additional indirect price impact when

Table 1.7: *Counterfactual using Actual Airbnb Penetration vs. Uniform Airbnb Penetration*

This table compares the counterfactual analysis using the actual Airbnb listing data with an alternative counterfactual analysis using a hypothetical scenario where the penetration of Airbnb is uniform across space and housing types while holding the total supply of Airbnb listings the same. This comparison illustrates that the primary driver of the welfare differences is due to the geography of actual Airbnb listings. The remaining difference in the counterfactual welfare with uniform Airbnb activity captures the higher willingness-to-pay for housing attributes by higher-income and more educated households.

	Actual Penetration		**Uniform Penetration**	
($ p.a.)	Median	Difference	Median	Difference
Household Size				
Household Size = 1	(134)		(142)	
Household Size = 5	(119)		(141)	
Small vs. Large Households		**(15)**		**(1)**
Race / Ethnicity				
White	(152)		(137)	
Black	(134)		(150)	
Hispanic	(113)		(130)	
Asian	(127)		(132)	
White vs. Hispanic		**(39)**		**(7)**
Education				
With College	(156)		(141)	
Without College	(120)		(132)	
With vs. Without College		**(36)**		**(9)**
Household Income				
Highest Quintile	(167)		(144)	
Lowest Quintile	(124)		(135)	
Highest vs. Lowest Income Quintile		**(43)**		**(9)**

the supply of substitute housing types is also restricted. Then, I provide some graphical intuition on the equilibrating process for illustrative purposes. Lastly, I provide numerical estimates of the breakdown between welfare transfer from the existing renters and welfare loss from the displaced renters.

First, consider the market clearing conditions for each type of housing in the long-term rental market:

$$\forall h: \quad D_h^L(p_h^L, p_{-h}^L, s^L) - s_h^L = 0 \tag{1.5.3}$$

where s^L denotes the entire vector of supply of each housing type.[60] Taking the total derivative of the market-clearing condition with respect to s_h and s_h' yields the following relationships:

$$\frac{dp_h^L}{ds_h^L} = \left(\frac{\partial D_h^L}{\partial p_h^L}\right)^{-1} \left(\underbrace{1}_{\text{(i). direct impact of supply reduction}} - \underbrace{\sum_{k \neq h} \frac{\partial D_h^L}{\partial p_k^L} \frac{dp_k^L}{ds_h^L}}_{\text{(iii). indirect impact from price increases of other home types}} - \underbrace{\frac{\partial D_h^L}{\partial s_h^L}}_{\text{reduction in the choice set}} \right) < 0 \tag{1.5.4}$$

$$\frac{dp_h^L}{ds_{h'}^L} = \left(\frac{\partial D_h^L}{\partial p_h^L}\right)^{-1} \left(0 - \underbrace{\sum_{k \neq h} \frac{\partial D_h^L}{\partial p_k^L} \frac{dp_k^L}{ds_{h'}^L}}_{\text{(ii). spillover from price increases of other home types}} - \underbrace{\frac{\partial D_h^L}{\partial s_{h'}^L}}_{\text{reduction in the choice set}} \right) < 0 \tag{1.5.5}$$

Equation (1.5.4) shows that the overall price impact from a supply reduction is a combination of the direct impact from the supply change s_h^L and the indirect impact from price changes of other housing types dp_k^L/ds_h^L. In particular, since the housing supply is *constrained* across all housing types except for the outside option, a supply reduction in a given housing type h will lead to a price increase in other housing types k because of substitution. However, this price increase in p_k^L, in turn, creates additional upward pressure on the original housing type p_h^L.

Conceptually, the housing market with severe supply restrictions is akin to a setting of

[60] Given that the relevant comparative static is with respect to supply, the demand for each housing type is a function of both the vector of home prices and the entire vector of supply because of its effects on the size of the choice set.

imperfect competition with quantity fixing: When a competitor's price increases, the optimal response is to *increase* one's own price when the total quantity is fixed. In general equilibrium, this could lead to a more exacerbated price response than the partial equilibrium effect alone. In the housing context, even though each absentee landlord does not have market power and participates in the market competitively, the overall difficulty in expanding the housing supply acts as the quantity fixing mechanism.

Next, as an illustrative device, Figure 1.15 breaks down the equilibrating process in successive steps of partial equilibria. It shows that the general equilibrium price impact could be greater than the partial equilibrium effects, especially for neighborhoods with little direct Airbnb penetration. To illustration this, I start with the original equilibrium price vector p_h^0 that clears the original long-term rental market with supply s^F. Now, the price vector p_h^1 is allowed to respond to the supply squeeze due to Airbnb s_h^A of its own type in a partial equilibrium manner, namely by assuming all other prices remained unchanged at p_h^0. It is shown in the top left panel of Figure 1.15. However, this partial equilibrium price response is not sufficient to clear the market because the demand substituting into other housing types will push up their prices. Hence, each successive step m generates a new partial equilibrium price vector based on the prices from the previous step $m-1$. These are shown in the top right and the bottom left panel of Figure 1.15.

$$\forall h: \; D_h(p_h^m, p_{-h}^{m-1}) = s_h^F - s_h^A$$

The process continues until the equilibrium is reached for all housing types, as shown in the bottom right panel.[61] Notice that the overall quantity fixing due to supply restrictions results in a general equilibrium price impact that is much greater than the partial equilibrium alone. In neighborhoods with little Airbnb activity, the general equilibrium effects dominate.

Lastly, I quantify the welfare impact because of the changes in the equilibrium prices

[61] It is important to note that there is an outside option that can accommodate the displaced renters.

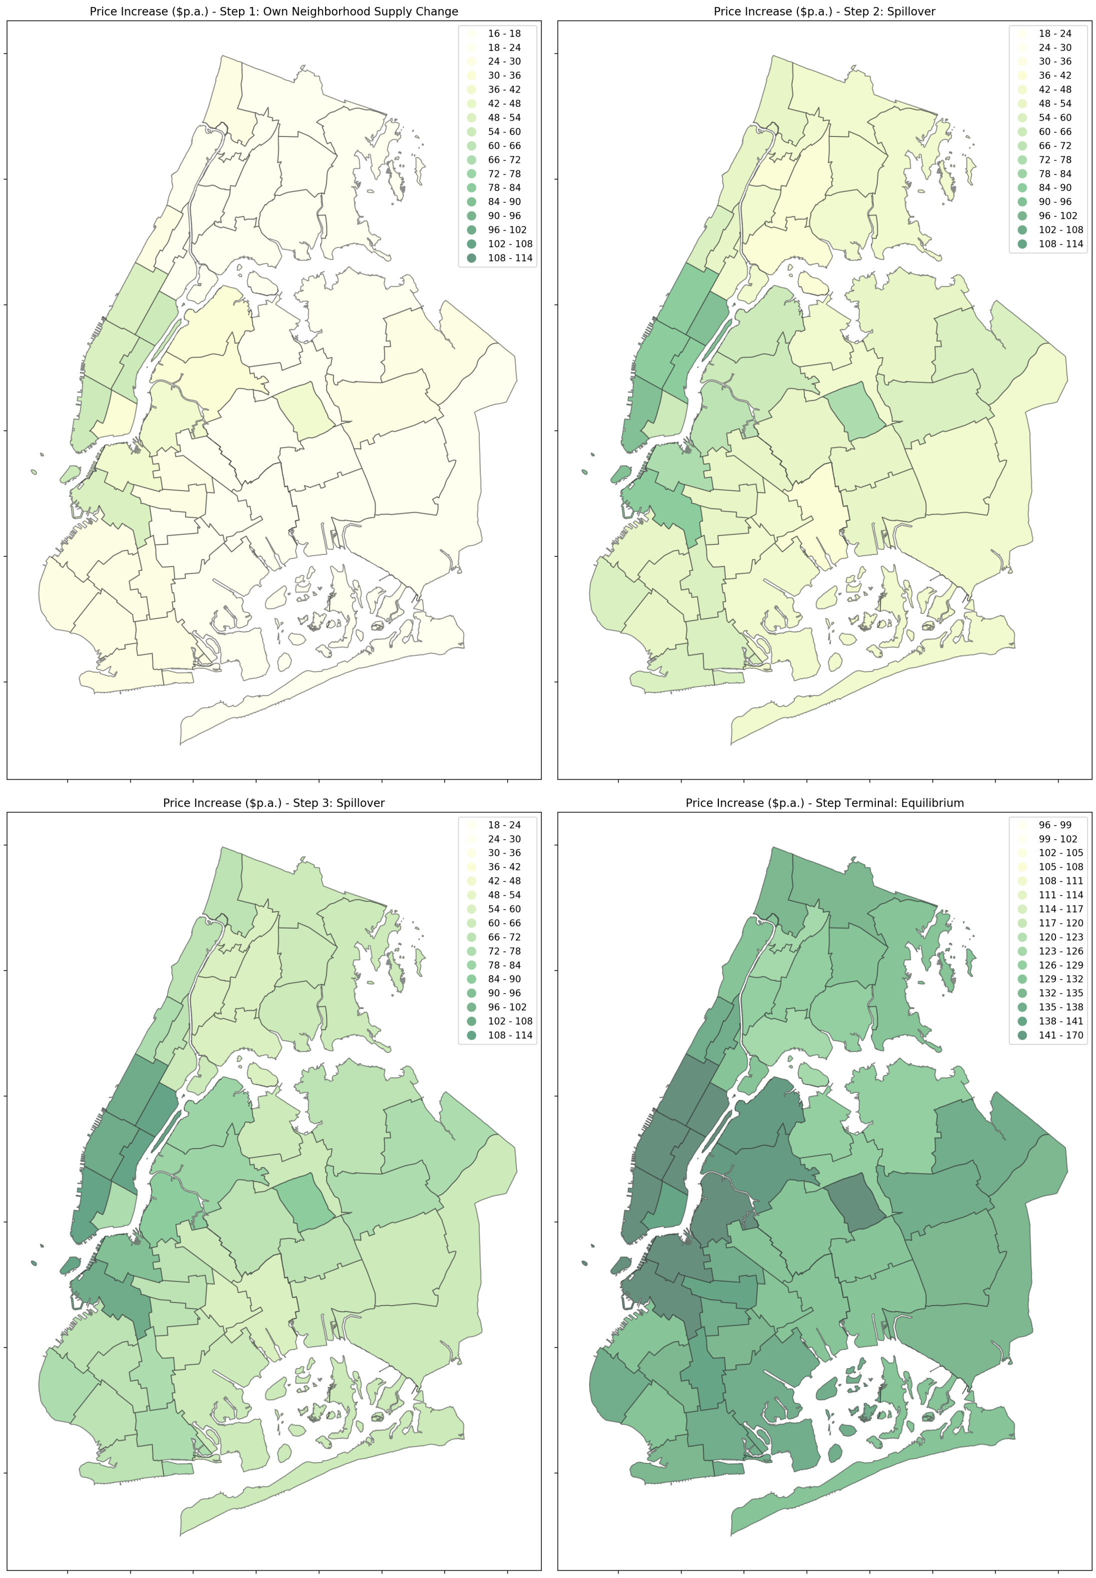

Figure 1.15: *The Equilibrium Effects of Supply Restrictions (Successive Steps of Best Response Function)*

versus the changes in the choice set. Denote $W_i(p, \mathcal{S}) = \left(\ln \sum_{j \in \mathcal{S}} \exp V_{i,j}(p_h)\right)/\alpha_i$. The compensating variation in Eq (1.5.1) can be decomposed into the part driven by the price change and the part driven by the choice set change:

$$CV_i^L = \underbrace{\Big(W_i(p^L, \mathcal{S}^F \setminus \mathcal{S}^A) - W_i(p^L, \mathcal{S}^F)\Big)}_{\substack{\text{Welfare change due to} \\ \text{changes in the choice set}}} + \underbrace{\Big(W_i(p^L, \mathcal{S}^F) - W_i(p^{CF}, \mathcal{S}^F)\Big)}_{\substack{\text{Welfare change due to} \\ \text{changes in the equilibrium prices}}} \tag{1.5.6}$$

I find that the choice set reduction contributed to a welfare impact of \$58mm p.a.[62] and the increase in equilibrium housing prices contributed to a welfare impact of \$201mm p.a. The welfare transfer from renters to property owners is computed by simply summing over the price changes $\sum_h \Delta p_h^L \times (s_h^F - s_h^A)$, which amounts to \$200mm p.a. and is the first-order term (i.e. the dark blue rectangle in Figure 1.1). The welfare loss from displaced renters amounts to \$0.9mm p.a., which is the second-order term (i.e. the light blue triangle in Figure 1.1).

Hence, the vast majority of Airbnb's welfare impact through the long-term rental market amounts to a transfer from the remaining renters to property owners, whereas a small portion is through a reduction in the choice set, and less than 1% is due to the welfare loss of renters displaced outside of the city.

C. The Role of Demographic Clustering In this section, I discuss that one's preference to live closer to other households with similar race and education implies that the distributional impact of Airbnb becomes compounded within demographic lines. In the NYC setting, given the overall geographical patterns of actual Airbnb activity, the clustering preference of white and educated households results in more concentrated welfare losses on white and educated renters.[63] In other words, since there is more Airbnb reallocation in centrally

[62] A reduction in the choice set itself results in a welfare decline simply because there are fewer draws of the idiosyncratic logit error $\epsilon_{i,j}^L$. This makes sense in the housing context, given that the individual-home specific utility is generally important. Even though this welfare impact does not translate into a direct transfer to the property owners, the short-term visitors do benefit from an enlarged choice set of lodging options because of Airbnb. I do not explicitly model the tourists' demand, as the focus of this paper is on the city residents. Still, I provide a back-of-the-envelope analysis at the end based on estimates from Farronato and Fradkin (2018).

[63] On the other hand, the clustering preference of African American and Hispanic households implies that, if more housing units were reallocated to Airbnb in predominantly minority neighborhoods, it would also generate stronger spillover to other minority renters.

located and educated neighborhoods, renters affected in these areas are more likely to substitute to other neighborhoods with high levels of education even if they are located further away, thus "spreading" the price impact more heavily to other educated households.

Since the dominant pattern that emerged from the data is the preference for clustering along demographic lines, including race, ethnicity, and educational attainment, the estimated long-term rental demand model captures such rich substitution patterns. Table 1.8 shows a subset of the cross-elasticities, where a price increase in a given neighborhood leads to higher rates of substitution to other neighborhoods that are closer in the demographic characteristics space.

For example, Forest Hill and Jackson Heights are located in close proximity to each other in Queens, both far from the city center. Neither neighborhoods experienced much Airbnb penetration, with less than 0.5% of the housing stock being affected. However, Forest Hill has a higher proportion of educated households at 63%, whereas Jackson Heights is at 28%. On the other hand, Forest Hill is over 47% white, whereas Jackson Heights is 63% Hispanic. As a result, the model predicts that a price increase in Chelsea and Midtown, which has the highest penetration of Airbnb in the data, generates a much higher rate of substitution toward Forest Hill (0.7%) than Jackson Heights (0.1%), since both Chelsea and Forest Hill have high proportions of white and educated households, as shown in Table 1.8. Consequently, in the counterfactual analysis, the equilibrium price impact due to Airbnb in Forest Hill is 20% higher than in Jackson Heights.

Table 1.8: *Estimated Own- and Cross- Price Semi-Elasticities*

Entry (i, j) corresponds to the percentage (%) of share changes in neighborhood j when the price of housing in neighborhood i increases by $1,000 per month. The table here shows only a subset of the neighborhoods for illustration, whereas the full matrix is 52 x 52. For example, this table shows that when price increases in Chelsea, Clinton & Midtown, households there will substitute away towards similar neighborhoods such as the Upper West Side. However, the substitution towards Forest Hill (an educated and predominantly white neighborhood) is much larger Jackson Heights (a less-educated and predominantly Hispanic neighborhood), even though they are geographically closely located in Queens.

	The Bronx Hunts Point, Longwood & Melrose	*Manhattan* Central Harlem	*Manhattan* Upper West Side & West Side	*Manhattan* Chelsea, Clinton & Midtown	*Queens* Jackson Heights & North Corona	*Queens* Forest Hills & Rego Park
Hunts Point, Longwood & Melrose	(21.9)	0.55	0.09	0.15	0.32	0.20
Central Harlem	0.58	(23.7)	0.20	0.32	0.18	0.30
Upper West Side & West Side	0.10	0.20	(21.0)	1.02	0.10	0.69
Chelsea, Clinton & Midtown	0.12	0.21	0.70	(21.6)	0.12	0.71
Jackson Heights & North Corona	0.62	0.33	0.19	0.32	(18.8)	0.35
Forest Hills & Rego Park	0.15	0.21	0.54	0.81	0.14	(37.5)

To summarize this section, the squeeze on the long-term rental market due to Airbnb results in a moderate yet material welfare transfer from renters to property owners at \$200mm p.a. or \$2.7bn in NPV terms. The median renter making \$47,000 a year loses about \$128 p.a. The general equilibrium price effect is elevated across all renters because housing supply is difficult to expand, which acts as a quantity-fixing mechanism of the market. Across the renters, I find the most significant losses are suffered by higher-income, educated, and white renters, when measured in dollar terms. The distributional differences are driven primarily by the geographical patterns of Airbnb activity and exacerbated by the preference for demographic clustering in housing choices.

1.5.2 The Distributional Impact via the Host Channel

In this section, I estimate the welfare derived from a resident's ability to act as a host on Airbnb. The advent of the sharing economy allows any household to participate directly in the production processes and act as a peer supplier. Overall, I find that the supplier surplus is immaterial for most households. Nonetheless, it does produce a long and heavy tail on the right, which suggests a few households with particularly low costs can benefit tremendously.

The patterns of the distributional impact are driven primarily by a household's cost to share. By observable demographic characteristics, larger surpluses accrue to households that are young, educated, and without children. Moreover, lower-income households have a more elastic supply, which suggests that they benefit more at peak times and locations.

The Counterfactual Specification

To evaluate the welfare gains from direct home-sharing, I perform a counterfactual analysis where the option to host on Airbnb is no longer available to the residents. The compensating variation of household i residing in neighborhood n is computed as follows

$$CV_i^R = \frac{1}{\alpha_i^R} \sum_t \ln\left(1 + \exp(V_{i,t}^R)\right) \tag{1.5.7}$$

where $V_{i,t}^R = \alpha_i^R p_{n,t}^A + \beta_i^R X_{n,t}^R + \xi_{n,t}^R$ and is summed over the course of the year.[64] Note that this specification calculates the *ex-ante* expected welfare from sharing before the actual idiosyncratic component $\epsilon_{i,t}^R$ is realized. Consequently, the distribution of the expected surplus is different from the distribution of the realized surplus, as the latter includes the additional variance due to the variance in the error term.

Overall, the distribution of the supplier surplus is centered close to zero but with a heavy right tail, suggesting that the bulk of the benefits is accrued to a concentrated few. Figure 1.16 shows the surplus distribution of all renters in the city,[65] where the median supplier surplus is only \$0.4 p.a. At the 75th percentile, the supplier surplus remains immaterial at \$5.9 p.a. However, the expected surplus on the very right tail above the 99th percentile amounts to \$307 p.a. When integrated over all renters, the total surplus produced by direct home-sharing amounts to \$23mm a year, or \$300mm in NPV terms.[66]

Supplier Surplus by Demographic Characteristics

One of the key benefits of estimating a random-coefficient supply system is its ability to analyze the supplier surplus by observable demographic characteristics, including income, education, age, and family structure. I discuss how each affects the surplus through one's cost of sharing as well as one's price sensitivity. Since the surplus is immaterial for the majority of the renters, I discuss patterns for both the typical renter and the right tail.

Household Income Even though the home-sharing surplus is immaterial for the majority of the households, higher-income households still expect to have a larger surplus on average. Meanwhile, the lowest-income group enjoys the largest benefits in the tail.

[64] I use 2018 as the base year over which the surplus is computed.

[65] In this section, I focus on the supplier surplus of renters in New York City, to make it directly comparable to the previous section. Nonetheless, the supplier surpluses of all residents, including both renters and owner-occupiers, are estimated and are included in the aggregate measures when appropriate.

[66] In comparison, conditioning on having made a room available on Airbnb in 2018 (totaling 24,100 listings, or 0.8% of the occupied housing units), the median revenue obtained by such resident hosts is \$2,484 in 2018, totaling \$137mm. Note that this represents only a fraction of the total Airbnb revenue in NYC, as the majority of its revenue is earned by housing units operated by absentee landlords.

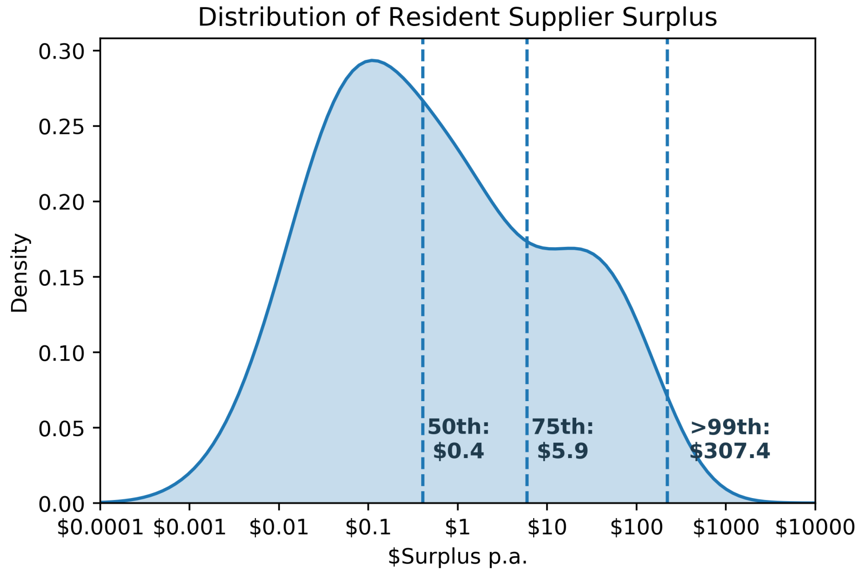

Figure 1.16: *Distribution of Short-Term Rental Surplus*

The kernel density plot of the supplier surplus through Airbnb, namely the welfare impact on renters via the *host* channel. Note that this plot is over the logarithm of the surplus. For the median household, the surplus from being an Airbnb supplier is immaterial at only $0.4 p.a., reflecting the fact that most households do not participate in the sharing platform. The 70th percentile surplus is at $5.9 p.a. The 90th percentile surplus is at $45.6 p.a. However, in the very right tail (>99%), the surplus is substantial at $307 p.a.

Figure 1.18 shows that the median host surplus for renters in the top income quintile is $5.2 p.a., and the median host surplus for renters in the bottom income quintile is $0.1 p.a. Other non-extreme percentiles (e.g. 75th or 90th) of supplier surplus are also generally higher for higher-income households. By contrast, the average surplus above the 99th percentile is significant at $232 p.a. for the top income quintile. But it is even higher at $454 p.a. for the bottom income quintile.

The difference between the median and the tail outcome is a result of two countervailing forces. Intuitively, the decision to share is based on a comparison of the market price $p^A_{n,t}$ and the cost to share $|\beta^R_i|/\alpha^R_i$. On the one hand, higher household income (as well as other correlates of income such as education) results in a lower cost to share with a smaller $|\beta^R_i|$ estimated from the model. On the other hand, higher household income also lowers one's price sensitivity α^R_i, suggesting that the sharing income may not be as valuable. In other

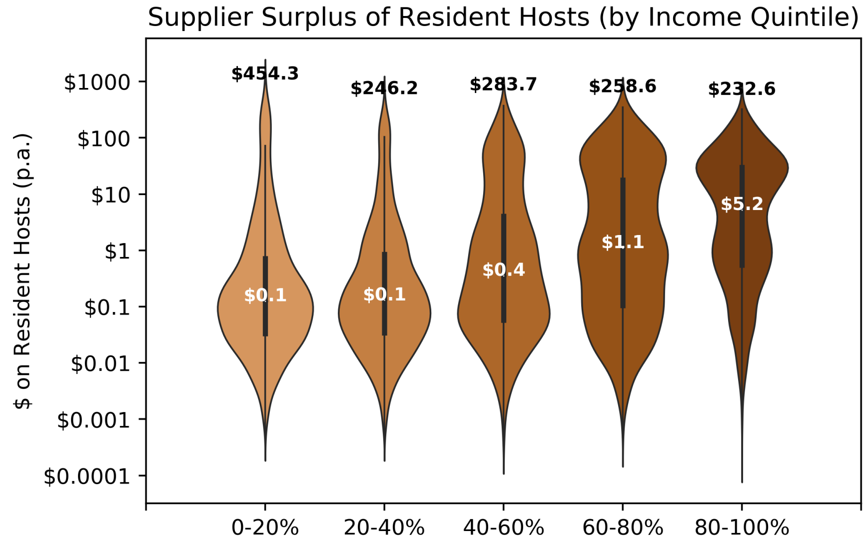

Figure 1.18: *Distribution of Short-Term Rental Surplus by Income Quintiles*

Note that this plot is over the logarithm of the surplus on the vertical axis. The numbers in the middle of the density plot represent the category median, whereas the numbers at the top represent the conditional average above the 99th percentile. The median surplus from home-sharing is immaterial across all income levels, but higher-income groups still have a higher mean. In the tail, the lowest-income quintile accrue larger benefits.

words, lower-income households are more likely to find the hassle of hosting visitors at home worth the money.

Relatedly, as there is significant demand seasonality in the short-term rental market, the sharing income during peak demand times could be particularly valuable for lower-income households

$$\frac{\partial CV_{i,t}^R}{\partial p_{n,t}^A} = \frac{\exp\left(\alpha_i^R p_{n,t}^A + \beta_i^R X_{n,t}^R + \xi_{n,t}^R\right)}{1 + \exp\left(\alpha_i^R p_{n,t}^A + \beta_i^R X_{n,t}^R + \xi_{n,t}^R\right)} \quad \Rightarrow \quad \frac{\partial^2 CV_{i,t}^R}{\partial p_{n,t}^A \partial \alpha_i^R} > 0$$

The cross-partial derivative of the compensating variation with respect to the market price $p_{n,t}^A$ and the price coefficient α_i^R is positive, where α_i^R is decreasing in household income. Given the model parameters, the supply elasticity for the top income quintile is 5.0, whereas the supply elasticity for the bottom income quintile is 7.4.

Overall, the anecdotal narrative that home-sharing could be beneficial for low-income families is borne out only in a limited sense: Conditional on being in the right tail of the surplus distribution, there are more low-income families, and they are more willing to take advantage of the peak demand times. However, for the typical family, regardless of income level, their expected benefit from home-sharing is immaterial.

Other Demographic Characteristics I find that households that are younger and have no children obtain greater supplier surpluses as resident hosts. For the youngest households with no children, the average surplus is expected to be $80 p.a. For households with children, the average surplus is immaterial across all age groups.[67] Since age and family structure only enter the cost term β_i^R in the model, they are not in front of the price coefficient α_i^R. As a result, both the average and the tail of the surplus distribution load heavily on young households with no children.

In terms of education, I find that educated households accrue greater supplier surplus as resident hosts. On average, the expected surplus is $36 p.a. for educated households and immaterial for households without a college degree.[68] Again, both the average and the tail of the surplus distribution load heavily on educated households.

Hence, even a reasonably simple supply system captures important cost heterogeneity of peer suppliers, where larger benefits accrue to those who are young, educated, and without children. In addition, it captures particularly rich dynamics with respect to household income, where a few low-income households could benefit a lot from the home-sharing platform. Nonetheless, for the typical household, the cost to share one's home with visitors is high; thus, the surplus remains immaterial.

---

[67]There are likely richer dynamics with respect to age-related events. However, given that the current model is estimated on aggregate data, it only captures a number of key demographic characteristics.

[68]The model does not necessarily explain why more education is likely associated with a lower cost of sharing, although one might speculate on reasons related to financial sophistication, willingness to adopt new technology, and one's ability to manage online businesses.

1.5.3 The Net Impact on Renters

In this section, I estimate the net welfare impact of Airbnb on renters by combining the welfare losses via the rent channel and the supplier surplus via the host channel described in the two previous sections. The net effect for the median renter is negative as the rent channel dominates, but the right tail is long as a few households benefit heavily from hosting. Nonetheless, on average, the losses remain more significant for educated, higher-income, and white renters.

In addition, the net welfare is divergent between the median and the tail in terms of its spatial distribution, where neighborhoods with the largest median losses also tend to be areas with the largest gains in the tail. I conclude with a discussion of how the social planner's objective may differ from that of the residents.

The Distribution of the Net Welfare Impact

Since the losses from the rent channel are diffused and the gains from the host channel are concentrated, the net welfare impact for the median renter is a loss of $125 p.a, as shown in Figure 1.20. At the 75th percentile, the net welfare impact is -$113 p.a. In fact, the net welfare for over 97% of the renters is negative. Nonetheless, the long right tail in the distribution of host gains results in a similarly long tail in the net welfare impact, averaging to $164 p.a. above the 99th percentile.[69] In the remainder of the section, I decompose the net welfare impact by different demographic characteristics (income, race, education) and by neighborhood geographic location.

Net Welfare by Household Income Overall, higher-income renters experience larger net losses on average and experience smaller net gains in the tail. As such, when restricted to just renters, it does not seem that Airbnb exacerbates income inequality. However, one

[69] As discussed before, the ex-post realized gain for the top percentile outcome is likely more extreme than the ex-ante expected surplus. Hence, the tail statistics should be interpreted accordingly as the ex-ante expectations before the idiosyncratic error terms are realized. I also conduct a robustness analysis with alternative assumptions on the auto-correlation of the error term, where the results are also qualitatively similar.

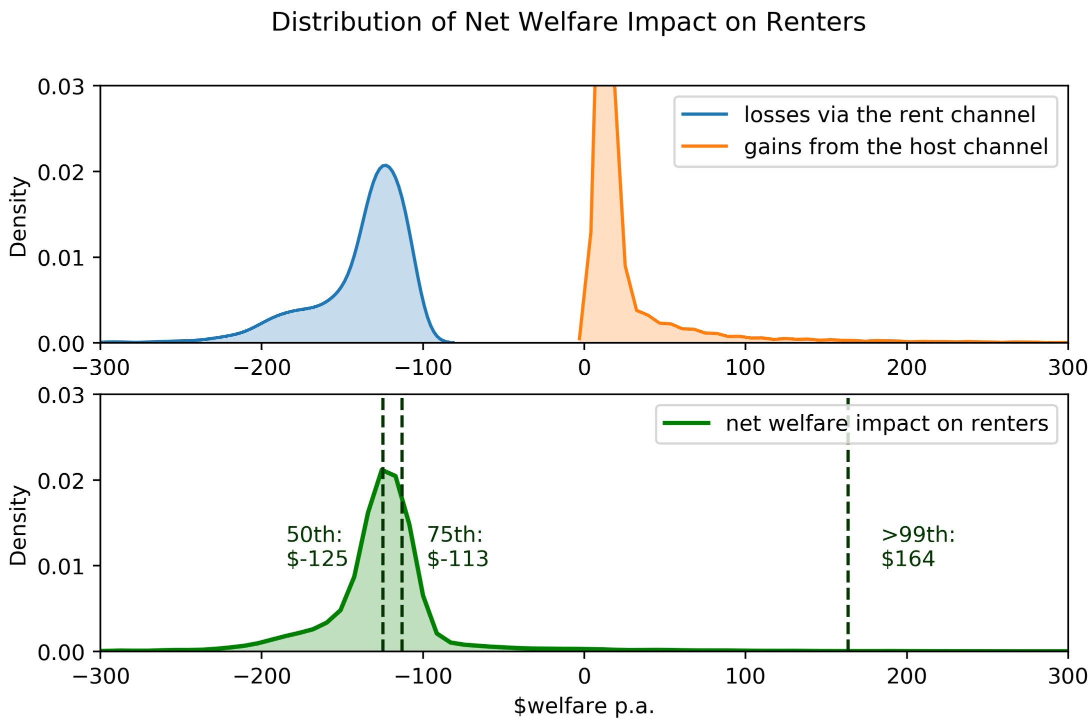

Figure 1.20: *The Net Welfare Impact on Renters*

The net welfare impact of Airbnb on renters combining the loss from the rent channel and the gains from the host channel. The impact on the median renter is -$125 p.a. and the impact on the 75th-percentile renter is -$113 p.a. Note that there is a long tail on the right that the plot does not accommodate.

needs to be more cautious of the overall impact beyond just renters, because unregulated Airbnb will likely worsen the wealth inequality due to the large welfare transfer from renters to property owners.

Table 1.9 shows that renters in the top income quintile lose $146 p.a. on average because of Airbnb, driven primarily by the welfare loss through the rent channel at $169, as the gain from the host channel is only $24. For renters in the bottom income quintile, the average net welfare is -$114, where both the losses from the rent channel and the gains from the host channel are less pronounced when compared with higher-income renters.

In the right tail, the increased price sensitivity of lower-income households results in greater host surpluses, especially if the household lives in areas with heavy short-term rental demand, as discussed in the previous section. In other words, for a small proportion of

Table 1.9: *Net Welfare Impact by Household Income*

This table compares the welfare impact of Airbnb on renters by household income quintiles. Overall, higher-income households experience larger losses on average and experience smaller gains in the tail.

Loss via the Rent Channel	Mean	Median	P25	P75	>P99
0-20%	-125	-124	-134	-118	-115
20-40%	-124	-125	-131	-113	-106
40-60%	-130	-126	-137	-114	-106
60-80%	-142	-134	-159	-122	-105
80-100%	-169	-167	-195	-137	-106

Gain via the Host Channel	Mean	Median	P25	P75	>P99
0-20%	12	0.1	0.0	0.7	454
20-40%	7	0.1	0.0	0.8	246
40-60%	14	0.4	0.1	4.0	284
60-80%	20	1.1	0.1	17.9	259
80-100%	24	5.2	0.5	30.0	233

The Net Welfare Impact	Mean	Median	P25	P75	>P99
0-20%	-114	-122	-131	-114	319
20-40%	-117	-118	-128	-111	101
40-60%	-116	-119	-130	-109	137
60-80%	-122	-125	-139	-112	109
80-100%	-146	-144	-174	-122	73

households, especially low-income ones, when their cost of sharing is low, they can not only make up for the higher rents but also obtain significant surpluses from becoming a host. As higher-income households are less likely to find the hassle of home-sharing worthwhile, their gains in the tail are smaller. Therefore, higher-income renters fare worse both on average and in the tail.

Net Welfare by Education In terms of education, the median net welfare impact is worse for more educated renters. Table 1.10 shows that the median is -\$136 p.a. for renters with a college degree and -\$120 for those without a college degree. As education is associated with a lower cost of home-sharing, I find that educated renters are more likely to be in the right tail of the host surplus distribution than those without a college degree. The divergence of the median and the tail by education is a result of the interaction between the demand and cost parameters in the short-term rental market.

Table 1.10: *Net Welfare Impact by Household Demographics*

This table compares the median welfare impact of Airbnb on renters by household characteristics. Overall, educated and white renters experience greater losses.

	Median Impact \$p.a.			Tail Impact \$p.a.		
	Loss via Rent	Gain via Host	Net Impact	Loss via Rent	Gain via Host	Net Impact
Overall	-128	0.4	**-125**	-109	307	164
Education						
Without College	-120	0.1	**-120**	-105	16	-98
With College	-156	10.8	**-136**	-112	393	253
Race						
Asian	-127	0.7	**-119**	-112	381	245
Black	-134	0.2	**-129**	-125	227	85
Hispanic	-113	0.1	**-111**	-105	232	107
White & Other	-152	1.9	**-130**	-111	326	179

Net Welfare by Race and Ethnicity When evaluating the net welfare impact along race and ethnicity lines, I find that the median white renter loses more than the median African American or Hispanic renter. This difference between the group medians is primarily driven by the differential welfare losses via the rent channel, where the proportion of housing units reallocated to Airbnb is higher in neighborhoods with more white and educated renters. The net welfare impact on the median white renter is -\$152 p.a. In comparison, the net impact on the median African American renter is -\$134 p.a., and for the median Hispanic renter, it is -\$113 p.a.

Even though race and ethnicity itself do not enter into the cost of home-sharing directly, there remains an induced distribution in terms of the host surplus because of the correlation among race, education, and income. As a result, I find more white households in the right tail of the welfare distribution than African American or Hispanic households.[70] White households in the tail make $326 p.a. from hosting, whereas Black and Hispanic households make $227 p.a. and $232 p.a., respectively.

Net Welfare by Geography In addition to decomposing the net welfare impact by demographic characteristics, it is also instructive to decompose the results by neighborhood location, especially since the short-term rental demand varies substantially across the geographic space. Two trends emerge in the spatial distribution of the net welfare: (i) There is a divergence between the experience of the median and the tail renter. (ii) The variation in terms of how much host gains could offset rent increases is driven by the neighborhood's cost of home-sharing, which favors centrally located low-income neighborhoods.

The rightmost panels in Figure 1.22 and Figure 1.23 shows the net welfare impact on the median renter and the right tail respectively. Interestingly, neighborhoods that experience heavier losses for the median renter also tend to experience large net gains in the right tail. For example, the net welfare impact for the median renter in Chelsea, which has high levels of Airbnb penetration, is -$146 p.a., while the tail of its host surplus is above $600 p.a. The irony of the divergence is explained by the overall spatial patterns of short-term rental demand: High short-term demand for the neighborhood drives more reallocation of the housing units away from the long-term rental market, thereby raising rents for everyone in the neighborhood. At the same time, high short-term rental demand also increases the prices that resident hosts can benefit from when they share their homes. However, as low-cost resident hosts are relatively few in number, their large gains affect only the tail of the distribution, not the median.

[70] Although it is not unreasonable to start with a model where race or ethnicity per se do not affect the cost of home-sharing, once controlling for other demographic characteristics, it is possible that African American and other minority hosts might face discrimination from the demand side and thus not be able to obtain as much surplus as peer suppliers. In this case, it would further widen the racial gap found here.

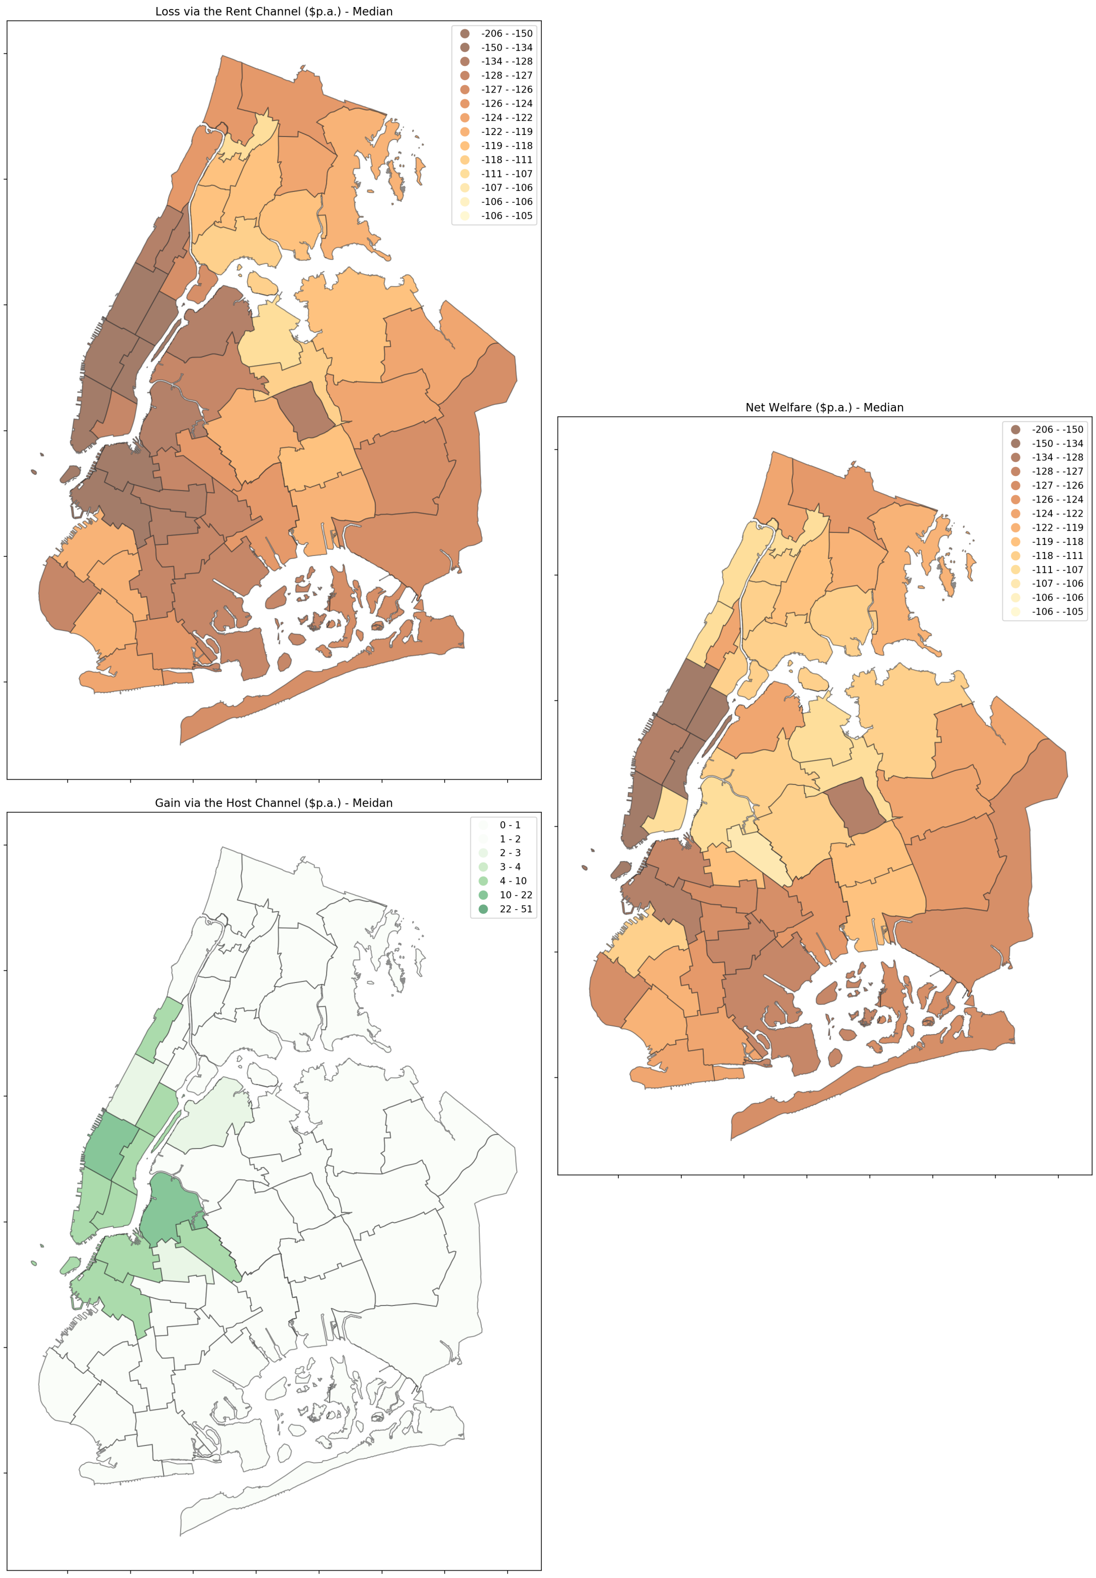

Figure 1.22: *Net Welfare Impact by Neighborhood (Median)*

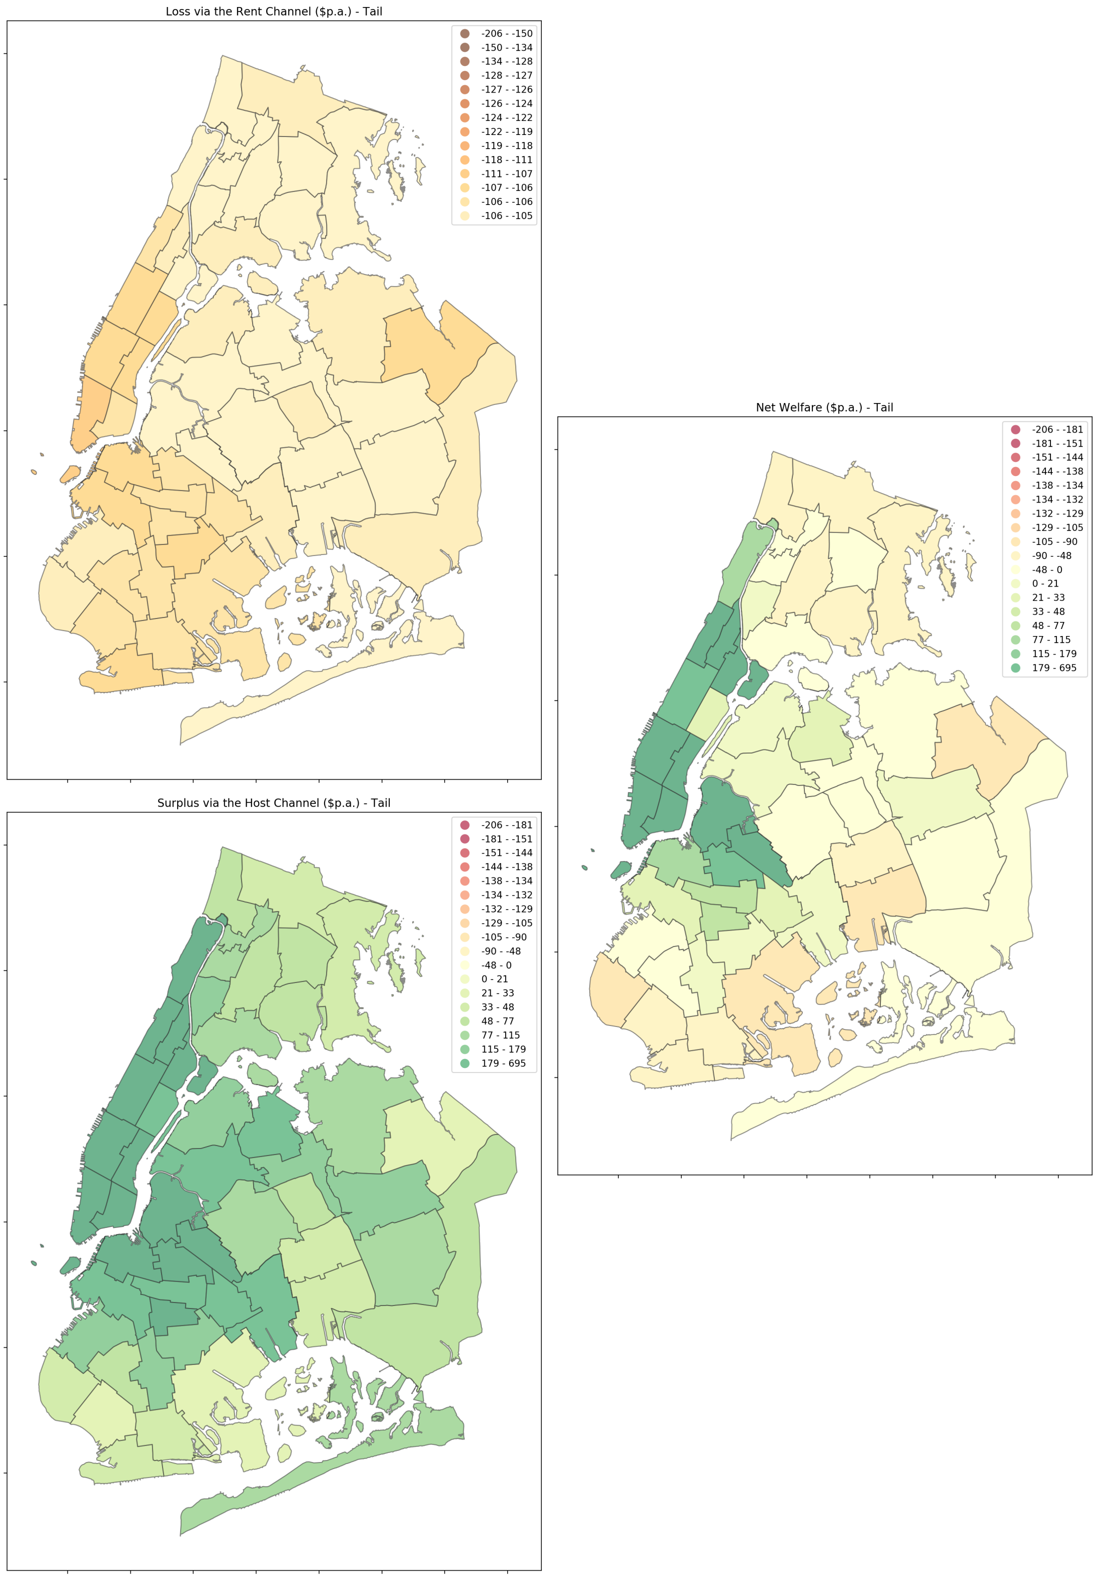

Figure 1.23: *Net Welfare Impact by Neighborhood (Right Tail)*

Although the short-term rental demand drives both the rent increases and the host surpluses, the variation in their difference is affected by the cost of home-sharing, where there are more low-cost hosts in low-income neighborhoods. For instance, the median household income among renters in Bushwick is only 60% of the median household income of the neighboring Park Slope and Williamsburg. As such, residents there are more likely to find home-sharing worthwhile. Despite being a relatively lower-income area, Bushwick is still conveniently located in Brooklyn and attracts short-term visitors. As a result, the surplus from home-sharing is relatively more valuable to its residents.[71]

1.5.4 Implications for the Planner

In this section, I discuss the aggregate welfare impact on all participants of the long-term and short-term rental market by combining model estimates from the previous sections. Despite an increase in the overall efficiency, the crucial trade-off is that the advent of home-sharing tends to benefit those who already own property. In the presence of severe housing market supply restrictions, the welfare transfer from renters to owners becomes substantial. In a city where the majority of the residents are renters, unregulated Airbnb will likely hurt the median person.

The top panel of Table 1.11 summarizes the aggregate impact on renters. On the one hand, the supply model estimates a surplus of \$23mm p.a. from home-sharing. On the other hand, the transfer from renters to absentee landlords amounts to \$200mm p.a., and the welfare loss from displaced renters is about \$1mm p.a. Therefore, the net impact on renters is a loss of \$178mm per year, or \$2.4bn in NPV terms. Although the total surplus from the host channel appears much smaller in magnitude compared to the loss from the rent channel, the relevant comparison for the social planner is different from that of the renter: The welfare gains obtained through the host channel are *true* economic benefits due to the innovation, whereas the welfare loss via the rent channel is not actually lost, but

[71]For example, as measured in terms of a larger residual when regressing the host gains on rent increases at the neighborhood level. Also, it can be visually inspected by comparing the two top panels of Figure 1.22.

primarily a *transfer* to existing property owners.

Table 1.11: *Welfare Impact on All Participants*

This table summarizes the aggregate welfare impact for all relevant participants of the long-term and short-term rental market. The black numbers represent outputs from the model. The gray numbers represent back-of-envelope approximations. In particular, absentee landlords' gains from hosting is based on the difference between their revenue from Airbnb less the forgone rent, assuming the same profitability as hotels. The net welfare impact on tourists and hotels is based on Farronato and Fradkin (2018), which estimates an average consumer surplus of $42 per room night and a decline of 5% in hotel profitability. Hotel usage and profitability are based on data from Smith Travel Research. Although the net welfare impact on renters is negative due to the transfer to property owners, the net impact inclusive of owner-occupiers and landlords is positive.

Welfare Impact on Various Stakeholders ($mm p.a.)			
Renters' loss to absentee landlords	-200		
Renters' loss from displacement	-1		
Renters' gain from hosting	23		
Net Impact on Renters	-178		
Owner occupiers' gain from hosting		5	
Absentee landlords' gain from renters		200	
Absentee landlords' gain from hosting		46	
Net Impact on Housing Market Participants		73	
Net welfare on tourists and hotels			126
Net Impact			198

If I include all participants of the housing market, namely renters, owner-occupiers, and absentee landlords, the aggregate welfare impact then becomes positive, as illustrated in the second panel of Table 1.11. Assuming absentee landlords choose to operate between the long-term and the short-term market rationally, they also accrue surplus from hosting when they operate on Airbnb. A back-of-the-envelope estimation based on their actual Airbnb revenue less the forgone rents results in a surplus of $46mm per year.[72] As owner-occupiers can be viewed as paying rent to themselves, the increase in property value is then

[72] Even though this is not precisely estimated based on an actual cost model of absentee landlords operating on Airbnb, the fact that it is much greater than the welfare loss of the displaced renters ($1mm p.a.) is what makes the net impact on all housing market participants positive.

mostly offset by the increase in the equivalent user-cost of their home.[73] Even though the distributional analysis in the previous section does not suggest that Airbnb exacerbates income inequality *among* renters, the presence of severe housing supply restrictions implies that the reallocation channel opened up by Airbnb will likely increase the wealth inequality between property owners and renters.[74]

Moreover, there is an additional net surplus for tourists. In the third panel of Tabel 1.11, using estimates from Farronato and Fradkin (2018), a back-of-the-envelope approximation produces a substantial gain from the consumer surplus of short-term visitors, net of hotel losses. As such, the total efficiency gain from the social planner's perspective is strictly positive and substantial.

Despite the overall increase in efficiency, the presence of housing supply restrictions results in a net welfare loss for over 97% of the renters. As 67% of the housing units are renter-occupied, the median person in the welfare distribution is a renter. The estimated net welfare on the median resident is a loss of $114 p.a.[75] Consequently, a simple voting model favors regulations to restrict Airbnb, especially restricting the reallocation of housing units by absentee landlords. In practice, it is consistent with the regulatory actions taken in New York. If such restrictions were fully enforced, a large portion of the tourist gains would be erased, whereas the rent transfer would be reversed. However, the analysis from the previous section suggests that enforcing Airbnb restrictions will also likely lead to a bigger reduction of housing costs for higher-income, educated, and white renters.

In other words, even though restricting Airbnb seems to be the most immediate regulatory action that can improve the median voter's outcome, the analysis here suggests that the true underlying challenge remains the inability for the city to expand housing supply. The entry of Airbnb provides a channel for the existing space to be used by the highest value

[73]For owner-occupiers, since they also have the option to leave the city, this suggests they are better off.

[74]In other words, had housing supply been elastic, one would expect a muted price response in the long-term rental market, a bigger reduction in price in the short-term rental market, and the supplier surplus to accrue to those hosts who have low costs of production.

[75]This assumes renters in the city do not own residential properties in the city, which is likely true for the vast majority of the renters.

bidder, but the total quantity restriction raises rents for everyone. Therefore, a reasonable alternative regulatory approach is to allow housing supply to expand more easily.

Model Limitation and Extensions

The structural model in this paper delivers reasonable parameter estimates as well as rich counterfactuals. However, I would like to point out its inherent limitations and the extensions that enrich the existing framework.

Both the long-term rental demand and the short-term rental supply model are static, which ignores transition dynamics. In the long-term rental model, I have assumed away moving costs. As such, the welfare impact found in a frictionless market can only be viewed as an approximation to the actual effects at best. Nonetheless, if moving costs are incorporated, the extent of welfare impact faced by renters in the most affected neighborhood is likely larger, which further exacerbates the losses suffered by higher-income, educated, and white renters. In the short-term rental market, there is an adoption process of the technology, where resident hosts exert a one-time effort to list on Airbnb and incur ongoing expenses to host guests. The static model abstracts away such adoption costs, assuming that the ongoing expenses are dominant.

There are a number of potentially useful extensions. First, the current model focuses on the increased housing costs faced by renters, since the owner-occupiers are always weakly better off. However, it is useful to characterize the gains accrued to the "displaced" owner-occupiers who sell their homes at an increased price and leave the city, benefiting even more than the owner-occupiers who remain in the city.

Second, I have yet to model the joint housing choice problem when households take into account their expected short-term rental surpluses as they make long-term rental choices. Such a joint decision model predicts that households with low costs of sharing (i.e. young, educated families with no children) will move further toward neighborhoods that are popular among tourists. Moreover, the expected short-term rental value will shift out the long-term demand curve, which results in even higher housing prices in equilibrium.

Last, the existing rent control and rent stabilization regulations could be modeled explicitly. Insofar as the current rental control law creates a mismatch in the housing market, the ability to rent part of a home out on Airbnb reverses its distortive effect.[76] In the current model, it is estimated as part of the host gains and captured only in the unobserved cost component at the neighborhood level. Meanwhile, as rent-stabilized units have limited ability to adjust rent, the impact of Airbnb on prices becomes further concentrated on the remaining market-price units, which tend to have higher-income tenants.

1.6 Conclusion

In this paper, I estimate the impact of the sharing economy on the highly contentious New York City housing market.

Such sharing technology operates on two fronts: the reallocation of resources and the increased utilization of resources. In a supply-constrained market, the reallocation of housing to Airbnb leads to an increase in equilibrium rents across all housing units in the city, not just for the specific units removed. It results in a widespread loss for all renters, aggregating to a transfer of $200mm p.a. to property owners. Moreover, the heterogeneity allowed in my structural model shows that more significant losses are shouldered by renters who are higher-income, more educated, and white as they tend to desire housing and neighborhood amenities that are highly desirable to short-term visitors as well. The utilization channel allows residents to provide short-term rental services in their existing homes. The estimated supply model suggests the cost of home-sharing remains high for most people, as evidenced by the fact that less than 1% of the residents become hosts. As a result, the median host gain is immaterial, aggregating to $23mm p.a. across the city. Nonetheless, a small fraction of the households with particularly low-cost of sharing obtains substantial host surpluses, including a few enterprising low-income families taking

[76]In fact, the presence of existing housing-market regulations tends to generate even more regulations, as Airbnb would otherwise provide a channel for regulatory arbitrage. For example, the short-term rental regulation in Los Angeles specifically does not allow units covered by the Rent Stabilization Ordinance to be on Airbnb.

advantage of peak short-term rental demand.

As Mayor de Blasio pushes for stricter enforcement of short-term rental regulations in New York City,[77] it is partially consistent with the goal of alleviating the housing affordability crisis, but it does not address the more fundamental problems created by the housing supply constraints. Banning the reallocation of housing to Airbnb likely reduces the housing costs for all renters, and especially help higher-income and educated renters. However, such ban will be at the expense of significant losses to both existing property owners and potential tourists who never arrive. It also creates additional incentives to evade regulations (Jia and Wagman, 2018). Such near-term regulation against Airbnb masks the underlying challenge, namely the detrimental effects of housing supply restrictions, which afflict not only New York City,[78] but also many other productive cities across the United States.[79] Even without Airbnb, these housing supply restrictions continue to produce significant economic distortions because of the inefficient location choices made by workers and firms (Hsieh and Moretti, 2019).

More broadly, the welfare impact of the sharing economy, as well as many other financial market innovations, can arise from a purely technological aspect as well as a regulatory aspect. The reallocation of housing units due to Airbnb serves as a form of regulatory arbitrage which reduces the price wedge created by the pre-existing allocation of space between the long-term rental market and the short-term rental market. It also responds to regulations that increase the cost of operating in the long-term rental market (e.g., rent control, tenant protection laws) by allowing property owners to generate cash flow in an alternative market. In other contexts, the development of novel financial products often serves as a form of regulatory arbitrage, fueling the growth of shadow banking or shadow

[77]Under the current Multiple Dwelling Law in New York, short-term rentals of Class A properties are not allowed unless its permanent resident is present. If fully enforced, it effectively rules out housing reallocation by absentee landlords, but still helps to protect some of the host gains.

[78]Although physical geography plays a role as Manhattan is a peninsula, the primary source of housing supply restrictions remains regulatory, ranging from restrictive zoning to onerous floor area ratios.

[79]In fact, it is also consistent with the observation that Airbnb tends to face more legal troubles in cities with more housing constraints.

insurance (Buchak, Matvos, Piskorski and Seru, 2018; Koijen and Yogo, 2016).

The utilization aspect of the sharing economy is a more welcoming technological feature, as it allows enterprising residents, including low-income ones, to engage in business activities that would otherwise be costly to start. By aggregating demand and verifying payments, as well as having substantial network effects, these peer-to-peer platforms reduce the barriers to entry for many enterprising individuals. The growth and behavior of a whole class of platform entrepreneurs is an exciting avenue for future research.

Chapter 2

The Value of Information: Why You Should Add the Second Order Conditions

2.1 Introduction

We study the estimation of an optimal choice problem inspired by the empirical example in Pakes, Porter, Ho and Ishii (2015), where agents make investment decisions based on their private information on productivity. In such settings, we worry that some of the instruments used might be weak, which motivates the analysis here.

We first conduct estimation for a continuous choice problem with varying instrument strengths, where we show the effects of explicitly including the second order condition. Then, we analyze the corresponding discrete choice problem using a revealed preference set-up, as in Ciliberto and Tamer (2009), and show how the second order condition can be incorporated there.

Given the inequality nature of both the second order condition and the revealed preferences, we follow Chernozhukov, Hong and Tamer (2007) for estimation. Meanwhile, although the literature on the inference for partial identification proposes various approaches such as Imbens and Manski (2004) and Andrews and Guggenberger (2009), we do not address the inference problem here, because we think that the Monte-Carlo simulation results are effective in conveying our main point. Furthermore, even though we will be working with a specific model in this paper, the key issues raised here can be relevant in a wide variety of settings as long as the estimation relies on optimality conditions combined with instruments, such as Berry, Levinsohn and Pakes (1995).

2.2 The Model

Consider the optimal investment choice of a firm d_i given the investments d_{-i} already made by its competitors in the same market, analogous to the set-up in Pakes, Porter, Ho and Ishii (2015). Suppose the revenue of the firm is as follows:

$$r(d_i, d_{-i}) = A \times \frac{d_i}{d_i + d_{-i}} \tag{2.2.1}$$

where the constant A is known.

The cost of installing d_i units of the investments is quadratic:

$$c_i(d_i) = (\beta_1 + \nu_i)d_i + \beta_2 d_i^2 \tag{2.2.2}$$

$$\mathbb{E}[\nu_i] = 0 \tag{2.2.3}$$

where ν_i represents the firm's independent draw of its idiosyncratic productivity shock that is known to the firm but unobservable to the econometrician. Thus, the firm makes its investment decision based on ν_i and d_{-i} to maximize its profit $\Pi_i(d_i, d_{-i}) = r(d_i, d_{-i}) - c_i(d_i)$.

Lastly, we assume that d_{-i} is drawn from a Poisson distribution whose mean negatively depends on $\nu_i + u_i$, where u_i represents an additional independent cost shock that is only relevant for the competitors.

2.2.1 The Continuous Optimal Choice Problem

Suppose the firm's optimal choice is continuous, that is $d_i \in \mathbb{R}$, we want to obtain an estimate of β_1 and β_2 based on the relevant moment conditions.

Moment Conditions Based on First Order Conditions (FOC)

For ease of notation, denote $c(d_i) = \beta_1 d_i + \beta_2 d_i^2$ and $\Pi(d_i, d_{-i}) = r(d_i, d_{-i}) - c(d_i)$. We form the following moment conditions based on the first order condition of each optimizing firm:

$$\mathbb{E}\left[\Pi'(d_i, d_{-i}) z_i\right] = \mathbb{E}\left[\left(A\frac{d_{-i}}{(d_i + d_{-i})^2} - (\beta_1 + 2\beta_2 d_i)\right) z_i\right] = \mathbb{E}\left[\nu_i z_i\right] = 0 \tag{2.2.4}$$

where z_i is any positive instrument that satisfies $\mathbb{E}[\nu_i z_i] = 0$. We have two valid instruments:

1. $z_i^1 = 1$
2. $z_i^2 = u_i$

Note that d_i and d_{-i} are both endogenous. Here, u_i is a valid instrument for d_i because u_i affects d_i through the number of competitors in the market d_{-i}, but is independent from ν_i. This problem is just identified and we can use the standard IV estimator.

However, in practical settings, one may not observe the cost shock u_i precisely and could suffer weak instrument problems. We model this through scaling u_i by $\pi > 0$ and adding a positive random noise, following Staiger and Stock (1997):

$$z_i^2 = \pi u_i + \epsilon_i$$

$$\epsilon_i \sim \text{Uniform}[0, 1)$$

Moment Conditions Based on Second Order Conditions (SOC)

Given that the firm profit is *maximized*, we also know that $\Pi_i'' \leq 0$. Since the instruments are positive, we can form the following inequality moments based on this second order condition:[1]

$$\mathbb{E}\left[\Pi''(d_i, d_{-i})z_i\right] = \mathbb{E}\left[\left(-A\frac{2d_{-i}}{(d_i+d_{-i})^3} - 2\beta_2\right)z_i\right] \leq 0 \tag{2.2.5}$$

Combining with Eq (2.2.4), we obtain a lower bound $\underline{\beta}_2^j$ and an upper bound $\bar{\beta}_1^j$ from each instrument:

$$\beta_2 \geq \underline{\beta}_2^j := \frac{\mathbb{E}\left[\left(-A\frac{d_{-i}}{(d_i+d_{-i})^3}\right)z_i^j\right]}{\mathbb{E}\left[z_i^j\right]}$$

$$\beta_1 \leq \bar{\beta}_1^j := \frac{\mathbb{E}\left[\left(A\frac{d_{-i}}{(d_i+d_{-i})^2} - 2\underline{\beta}_2 d_i\right)z_i^j\right]}{\mathbb{E}\left[z_i^j\right]}$$

Geometrically, Figure 2.1 shows that each moment equality condition generated by the FOC identifies a line in the space of (β_1, β_2), where their intersection produces the IV estimator. However, the moment inequality condition generated by the SOC further restricts each line to a *ray* starting at $(\bar{\beta}_1^j, \underline{\beta}_2^j)$. If an instrument becomes weak, producing an intersection that is not on the ray, the SOC restriction will become binding.

Simulation Results

We run simulations to illustrate the properties of the estimators.

First, Figure 2.3 shows that as the instrument weakens, the IV estimator becomes increasingly noisy and biased, exhibiting the classical weak instrument problem.

Next, we add the inequality moments generated by the SOC to the equality moments generated by the FOC, where the estimation is conducted following Chernozhukov, Hong and Tamer (2007) using an identity weighting matrix. Figure 2.5 shows that this noticeably "tucks in" one of the tails.

Therefore, even when a problem has enough equality restrictions for identification, incorporating the second order condition could still improve the efficiency of the estimator, especially when some of the instruments are weak.

[1]Although in this example the inequality can be applied at the observation level, practically, any measurement error in Π_i'' would require the expectation operator.

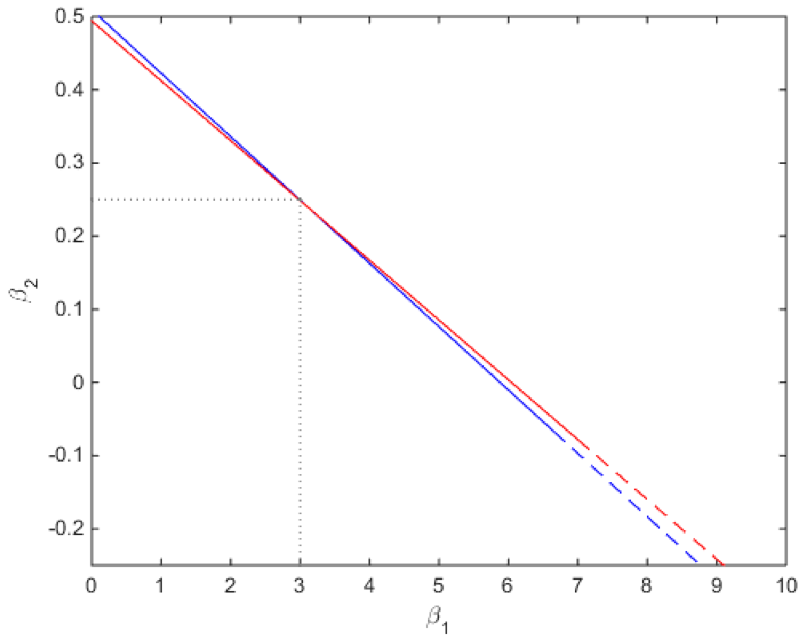

Figure 2.1: *Parameter Identification for the Continuous Problem*

Each FOC produces a line and the intersection produces the IV point estimate, while the dashed part shows the portion ruled out by the SOC. The true parameter value $\beta_1 = 3$ and $\beta_2 = 0.25$.

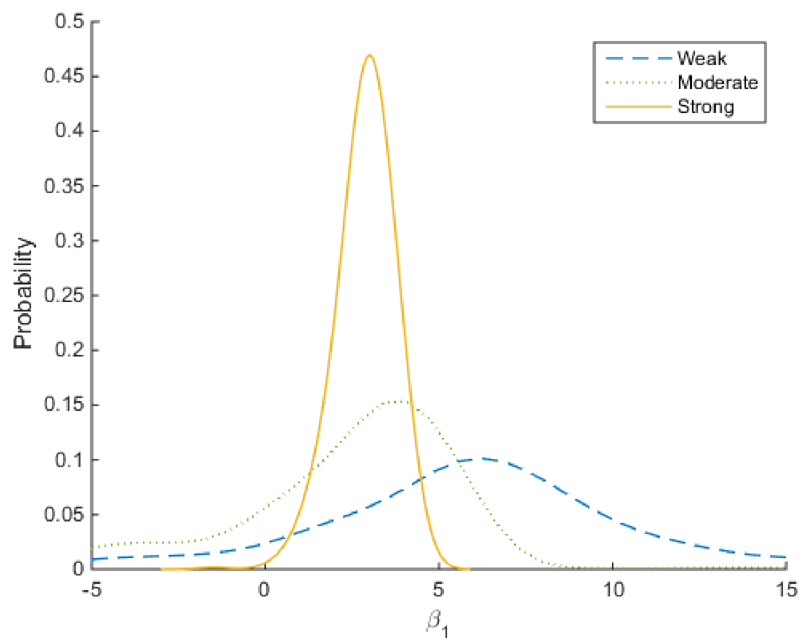

Figure 2.3: *The Effects of Instrument Strengths*

The effects of the instrument strengths. $\pi = 0.02, 0.1, 0.5$ are used for the weak, moderate and strong label respectively.

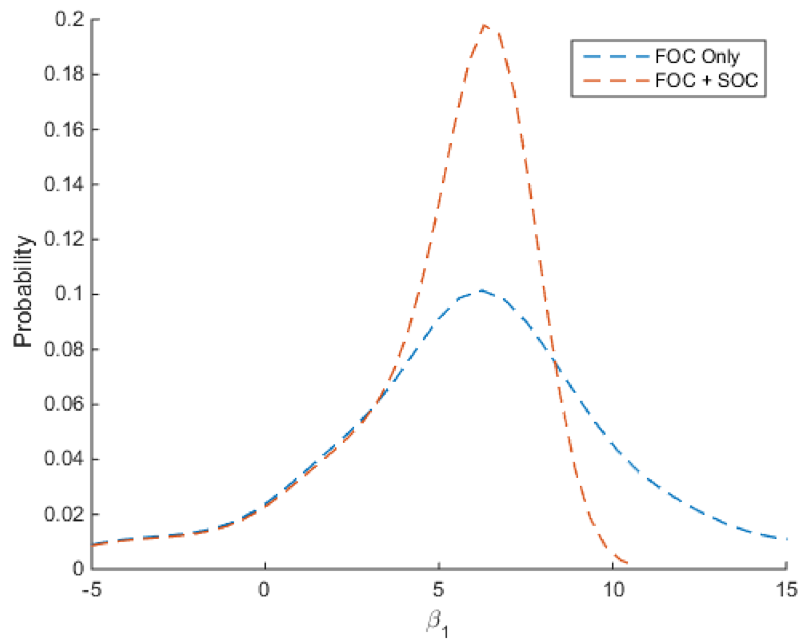

Figure 2.5: *Impact of Moments from the Second Order Condition*

The instrument is weak. Comparing to only using the FOC restrictions, including moments from the second order condition "tucks in" one of the tails.

2.2.2 The Discrete Optimal Choice Problem

In this section, we study the corresponding discrete choice problem, which is analogous to the previous section except that the firm can no longer choose any investment $d_i \in \mathbb{R}$, but only discrete units with a discretization step of S. Specifically, $S = 1$ implies that $d_i \in \mathbb{Z}$. The revenue function, the cost function and the agent's information set remain the same.

Inequality Moment Conditions Based on Optimality

Based on revealed preferences, namely $\Pi_i(d_i, d_{-i}) \geq \Pi_i(d_i - S, d_{-i})$ and $\Pi_i(d_i, d_{-i}) \geq \Pi_i(d_i + S, d_{-i})$, we can construct the following moment inequality restrictions for the same positive instruments:

$$\mathbb{E}\left[\left(\frac{\Pi(d_i, d_{-i}) - \Pi(d_i - S, d_{-i})}{S}\right) z_i\right] \geq \mathbb{E}\left[\nu_i z_i\right] = 0 \tag{2.2.6}$$

$$\mathbb{E}\left[\left(\frac{\Pi(d_i, d_{-i}) - \Pi(d_i + S, d_{-i})}{S}\right) z_i\right] \geq \mathbb{E}\left[-\nu_i z_i\right] = 0 \tag{2.2.7}$$

The estimator will find the bounds of the identified set if feasible, and minimizes the deviations otherwise. Meanwhile, combining (2.2.6) and (2.2.7), we obtain

$$\mathbb{E}\left[\left(\frac{\Pi(d_i+S,d_{-i})-\Pi(d_i,d_{-i})}{S}-\frac{\Pi(d_i,d_{-i})-\Pi(d_i-S,d_{-i})}{S}\right)z_i\right]\leq 0 \tag{2.2.8}$$

which resembles SOC because it computes the difference of the first derivative $\Pi'(d_i, d_{-i})$ estimated above and below d_i.

Simulation Results

We run simulations using the two inequality moment conditions constructed in (2.2.6) and (2.2.7).

To build intuition, we show in Figure 2.7 the identified set using the constant only. The intersection of the two inequalities forms a "wedge", which contains the ray constructed by the FOC and SOC of the corresponding continuous problem up to an approximation term.[2]

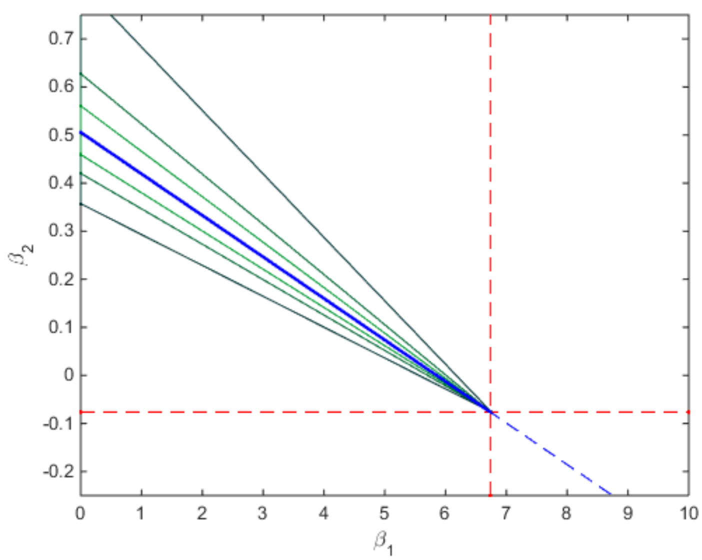

Figure 2.7: *Identification from Moment Inequalities Conditions*

Identification using the constant. The thick blue ray shows the FOC and SOC restrictions of the continuous problem, where the start of the ray is emphasized by the red-dashed lines. The pair of green lines shows the "wedge" identified by the moment inequalities, which becomes "thinner" as S decreases.

---

[2]Note that $\mathbb{E}\left[\left(\frac{\Pi(d_i+S;\hat\beta)}{S}-\frac{\Pi(d_i;\hat\beta)}{S}\right)z_i\right]=\mathbb{E}\left[\left(\Pi'(d_i;\hat\beta)\right)z_i\right]+\frac{1}{2}\mathbb{E}\left[\left(\Pi''(d_i;\hat\beta)\right)z_i\right]S+\mathcal{O}(S^2)\leq 0$, provided $\mathbb{E}\left[\left(\Pi'(d_i;\hat\beta)\right)z_i\right]=0$ and $\mathbb{E}\left[\left(\Pi''(d_i;\hat\beta)\right)z_i\right]\leq 0$ and the third term is not too large.

Then, we show in Figure 2.9 the effects of the instrument strengths. With $S = 1$ fixed, the bounds of the identified set is much less sensitive to the weakening of the instrument, compared to the IV estimator of the corresponding continuous problem. This nice behavior is due to the implicit incorporation of the SOC as shown in Eq (2.2.8).

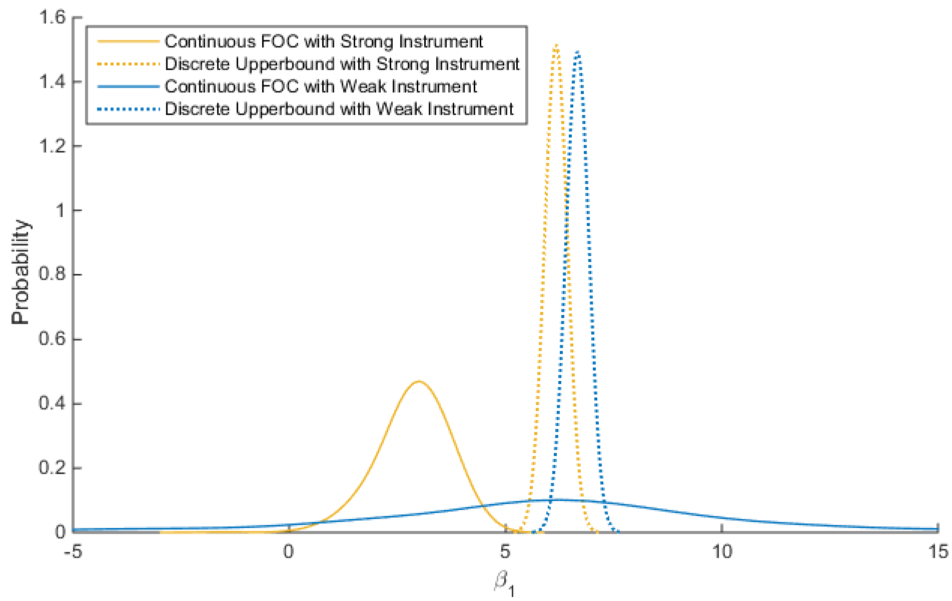

Figure 2.9: *Discrete vs. Continuous Problem*

The upper bound of the identified set of the discrete problem (the dotted lines) is much less sensitive to the weak instrument than the corresponding parameter estimates obtained from the FOCs of the continuous problem (the solid lines).

Next, Figure 2.11 shows as the discretization step size decreases, the bounds estimated from the discrete problem starts to resemble the IV estimator of the continuous problem, increasingly breaching the second order condition. To understand this, take the limit of Eq (2.2.6) and (2.2.7) with $S \to 0$:

$$\mathbb{E}\left[\Pi'_{-}(d_i, d_{-i})z_i\right] \geq 0$$

$$\mathbb{E}\left[\Pi'_{+}(d_i, d_{-i})z_i\right] \leq 0$$

Since Π is differentiable, we recover the first order condition:

$$\mathbb{E}\left[\Pi'(d_i, d_{-i})z_i\right] = 0 \tag{2.2.9}$$

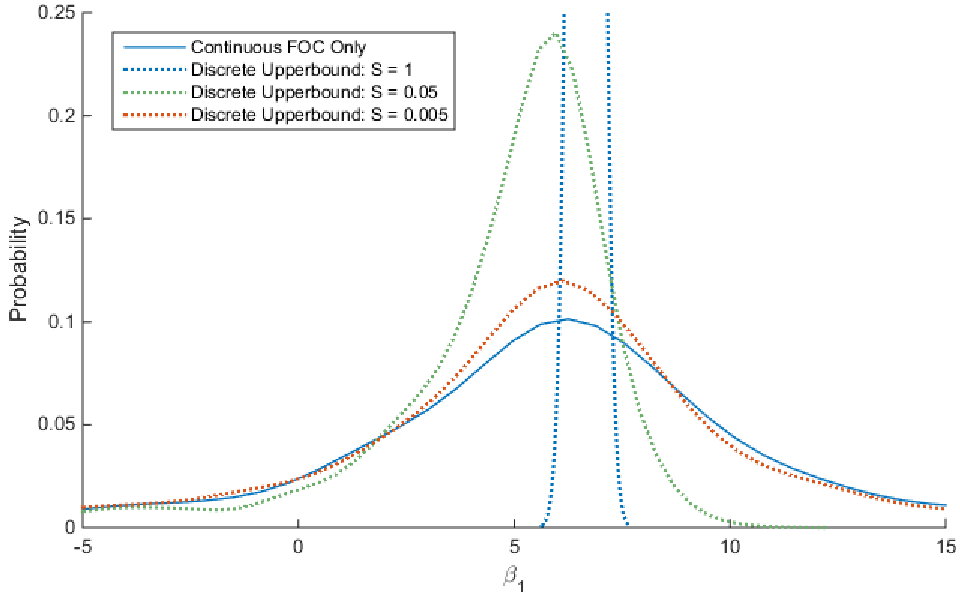

Figure 2.11: *The Distribution of the Upper Bounds*

With large S, the upper-bound of the identified set is further from the true value, but the distribution of the bound itself is narrow. As S decreases, the distribution starts to resemble the IV estimator of the continuous problem.

However, we can rewrite Eq (2.2.8) as

$$\mathbb{E}\left[\left(\Pi''(d_i, d_{-i})S + \mathcal{O}(S^2)\right) z_i\right] \leq 0 \tag{2.2.10}$$

Notice that the strength of the second order condition is scaled by S. As $S \to 0$, the revealed preference set-up converges to that of the FOC only and the SOC loses its effect.

To address this perverse behavior, we suggest explicitly constructing the additional moment for the SOC as in Eq (2.2.8) but scaled by $1/S$. In the limit when $S \to 0$, this becomes the explicit addition of the SOC moments to the continuous problem, shown in Figure 2.13.

Relatedly, by adding moment conditions that look beyond the "immediate neighbor" for $N \times S$ steps away, one also improves the relevance of the SOC by a factor of N. However, one needs to trade off these additional moments with potentially larger confidence sets. Indeed, Pakes, Porter, Ho and Ishii (2015) included larger steps ($d = \pm 2$) and found the

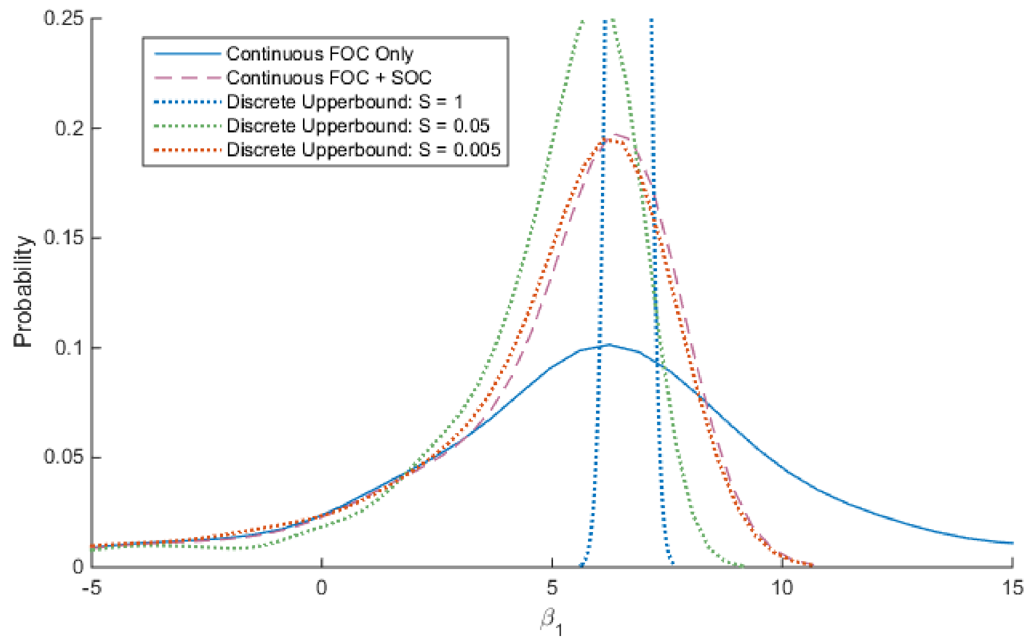

Figure 2.13: *Improvements Introduced by the Additional Moment Condition*

With additional moment inequalities specified by (2.2.8) scaled by $1/S$, as S decreases, the distribution of the upperbounds starts to resemble the FOC + SOC estimator of the continuous problem.

estimate of the identified set unchanged.

2.3 Conclusion

Using a simple optimal choice setting, we showed why it can be useful to include additional moment conditions for both the continuous and the discrete choice problem. Therefore, regardless whether there are already enough moments for identification, we suggest empirical researchers to consider explicitly incorporating moments based on the second order condition, which may be particularly useful when there are weak instruments.

Chapter 3

And the Children Shall Lead: Gender Diversity and Performance in Venture Capital[1]

3.1 Introduction

Homophily-driven biases can be a powerful force that inhibits diversity in organizations. Gender hiring bias has been shown to persist over time in many highly compensated professions. To overcome these barriers, policymakers have often attempted to actively promote diversity in the workplace. Most recently, California passed a law that mandates gender diversity on boards of companies incorporated in the state. Whether enacted by politicians or senior executives, many of the measures that are adopted assume that greater diversity naturally leads to better performance. Others are skeptical that there is a measurable improvement in performance when diversity is mandated. Most of the research on whether or not greater diversity leads to improvement in organizational performance has been hampered by the inability to identify exogenous variation in diversity which is needed for causal inferences. Still, other work has been done in artificial settings outside

[1]Co-authored with my advisor Paul A. Gompers

of a real business context in which true long-run profit motives would be present. Our paper makes two important contributions to the literature on diversity by using a novel experimental design. First, we show that when existing partners in a venture capital firm have a higher proportion of daughters relative to all children, hiring biases against women are reduced. Second, our reduced-form regressions show a strong relationship between the relative number of daughters that senior partners have and deal/fund-level performance. Lastly, we instrument a firm's gender diversity induced by hiring a female investor with our children data, providing suggestive evidence that greater exogenous gender diversity leads to improvement in performance.

Our institutional setting, venture capital, has a number of important attributes that make it an ideal setting to explore the performance implications of diversity. Venture capital firms tend to be small with typically less than a dozen investment professionals. The decision-makers are easy to identify (partners), and performance (fund-level returns and deal-level outcomes) can be precisely estimated. Through unique data, we are able to identify hiring events for senior investment professionals at venture capital firms. Calder-Wang and Gompers (2017) show that only about 8.5% of new hires in the venture capital industry are women. Prior work by Gompers, Mukharlyamov, Weisburst and Xuan (2021) has shown that approximately 75% of venture capital firms have never had a senior investment professional who is a woman. Our experimental design is to gather data on the gender of venture capitalists' children. Our results show that when existing partners have a higher number of daughters relative to the total number of children, hiring biases against women are reduced. When existing partners have more daughters, the probability of hiring a female investor is increased substantially. The relative effect of having a daughter rather than a son for all senior partners at a firm translates into a 4.4% increase in the probability of hiring a woman.[2] Compared to a baseline level of 9.9% of hiring a woman in our sample, the relative effect represents a 45% increase in the probability of hiring a female investor. Additionally, our results for hiring more women only exist for senior partners' children. This makes sense

[2]The standard deviation of the number of daughters for a senior partner is 0.90.

given that senior partners typically retain decision rights over new hires.

Because the gender of one's children is usually thought to be exogenous, the gender diversity induced by having more daughters, controlling for the total number of children, can be used to estimate the impact of gender diversity on firm performance in venture capital. We examine both deal-level outcomes as well as fund-level excess returns. In reduced-form regressions, the gender of partners' children has strong and significant effects on both. In instrumental variable regressions, our results suggest that greater gender diversity has economically and statistically significant effects on deal-level outcomes and fund-level excess returns. Success rates on individual deals improve by 4.7% for a 5% increase in gender diversity (namely increasing the fraction of women hired from a baseline level of approximately 10% to 15%). This represents a 17% increase compared to the baseline success rate of 27.3%. Our results are robust to various measures of the relative ratio of daughters to total children as well as alternative measures of venture capital performance. The relevant exclusion restriction here is that the impact of having daughters affects venture capital performance only through the proportion of female partners hired. We test and rule out a number of alternative explanations, ranging from whether having more daughters alters the gender composition of the entrepreneurs invested, to whether raising daughters measurably improves the productivities at an individual level. Taken together, we think this framework provides suggestive evidence that gender diversity improves venture capital performance, although we acknowledge that there may be other alternative channels through which children's gender can affect investment performance that we cannot rule out.

Related research has explored the gender bias in hiring as well as various treatments that can reduce the gender bias in hiring. In their paper, Goldin and Rouse (2000) find that introducing blind auditions dramatically increased female representation in the major orchestras in the United States. Bohnet, van Geen and Bazerman (2016) find in an experimental setting that joint evaluation of job candidates can reduce gender bias in hiring versus separate candidate assessment. However, besides direct interventions at the hiring stage, subtle debiasing effects related to an increase in exposure have been considered as an alternative, albeit outside of the labor market. In the political arena, Beaman, Chattopadhyay, Duflo, Pande and Topalova (2009) show that when voters were exposed to female chief councilors, the likelihood of a woman winning an unreserved councilor or Pradhan seat in India increased. In the more recent theory and experimental literature, Bordalo, Coffman, Gennaioli and Shleifer (2016) show that stereotypes are developed by overweighting representative members of a group. Under this framework, gender stereotypes could lead to persistently homogeneous organizations if they are small and make infrequent hiring decisions, like our venture capital setting. Thus, the first part of our paper contributes to this literature by providing real-world, empirical evidence of the relevance of gender exposure effect on hiring decisions in the labor market.

Our choice of exogenous variation is motivated by research that has explored the effect of parenting on social preferences. For example, Warner (1991) surveys parents and finds that fathers of daughters tend to show greater support for feminist causes. Similarly, Warner and Steel (1999) show that fathers of daughters have greater support for gender equity than do fathers of sons. More recent works have also demonstrated that decision-making of fathers can be influenced by the gender of their children. Washington (2008) finds that US Congressmen vote more liberally, especially on issues affecting women, if they have more daughters. Cronqvist and Yu (2017) show that CEOs who have more daughters are more likely to adopt socially responsible corporate policies. Glynn and Sen (2015) demonstrate that Federal Court judges with more daughters tend to decide cases on women's issues more liberally and that the effect is largely driven by Republican-appointed judges. Finally, Bennedsen, Nielsden, Perez-Gonzales and Wolfenzon (2007) explore the effects of gender birth order and the fraction of children that are girls on the likelihood that a family firm appoints a non-family CEO. Like our work here, they use the gender of a family CEO's children as an instrument for the appointment of a non-family CEO successor. In instrumental variable regressions, they find that family CEO succession reduces performance relative to a non-family CEO.

Our results on hiring decisions suggest that having daughters has a dramatic debiasing effect on hiring even in an industry in which gender diversity is severely lacking. The demographic patterns and trends surveyed in Calder-Wang and Gompers (2017) highlight the overall lack of gender diversity in venture capital. Women have entered into venture capital at a rate much lower than their entry rates into other highly compensated professional fields such as medicine or law, both of which is approaching equity at the junior levels. The representation of women in advanced degrees in science and technology and MBAs (as a precursor to entry into venture capital) are much higher than the representation of women in the innovation sector. The percentage of venture capital partners who are female has not increased measurably over the past twenty-five years, persistently hovering around 10%.

There is certainly a multitude of factors that might explain the persistent low fraction of women in venture capital. We do not attempt to disentangle the factors here, but we want to highlight the role of homophily, especially in small teams. As surveyed in Mcpherson, Smith-lovin and Cook (2001), the notion that "similarity breeds connection" has robust and profound effects in network structures of every type, including "marriage, friendship, work, advice, support, information transfer, exchange, co-membership, and other types of relationship." Moreover, the typical venture capital firm is small in size, with a median of three partners in our data set. Hiring decisions are made infrequently. Most venture capital firms only make infrequent senior hires, e.g., perhaps once every three to five years. Aggregate new hiring in this industry is driven by the (aggregated) decisions of small teams. From social psychology, small groups are more likely to be homophilous and to have biases aggregated into expressed decision-making (Klocke, 2007). Thus, a slight preference over certain demographic characteristics, like gender, could aggregate into a sustained overall lack of gender diversity at the industry level.

A direct implication of this "birds of a feather" phenomenon is that venture capitalists prefer to hire, invest in, or coinvest with those that are similar to themselves in characteristics such as gender and ethnicity. Cohen, Frazzini and Malloy (2008) show that homophily also works at the school ties level in the investment management arena between buy-side analysts and CEOs. Moreover, Gompers, Mukharlyamov and Xuan (2016) show that coinvestment patterns in venture capital are driven by social similarities, which means venture capitalists who are more similar in gender, ethnicity, school background, and work history are more likely to collaborate. Solal (2019) looks at televised entrepreneurial pitch competitions and finds strong gender matching between investors and entrepreneurs. Similarly, Ewens and Townsend (2020) find gender segregation on AngelList in which male investors show more interest in male-founded companies and female investors show more interest in female-founded companies.

Our next contribution is to use a more credible empirical strategy to estimate the impact of diversity on firm performance in a real business setting. Even though we are by no means the first to use the gender of one's children as a randomization device, the venture capital setting, with our rich person and investment-level data, gives us the unique ability tightly link the family characteristics of the key decision-makers with every hiring decision and investment outcome. To our knowledge, we are the first to map such exogenous variations to actual firm outcomes and use it to deduce the performance effects of diversity.

Sociology-based research has tended to look at ex-post data and measure correlations with performance. Results on gender diversity have been by and large equivocal. Furthermore, the setting does not allow for causal interpretations of results. Still, other papers have looked at experimental settings and assigned members based on gender to various "team-based" projects. These works, however, tell us little about whether or not the kinds of complex problems in business are affected by diversity. Bernile, Bhagwat and Yonker (2018) use the local availability of diverse directors as an instrument and find that greater board diversity leads to lower volatility and better performance. Several recent papers have looked at mandated board diversity. Schwartz-Ziv (2017) looks at mandated board diversity in Israel for firms with any government ownership and finds that boards with equal numbers of men and women are more active, but she does not find a performance effect. Other papers find more mixed results. Our paper differs from papers that look at mandated diversity, because forced diversity could have different results on performance from diversity resulting from debiasing hiring.

Theory also does not help when trying to understand whether firm diversity increases or decreases performance. One conjecture is that the more characteristics a pair of individuals have in common, the better the pair is likely to perform. This better performance can result from easier communication, the ability to better convey tacit information, or the ability to make joint decisions in a timely and productive manner (Mcpherson, Smith-lovin and Cook, 2001; Cohen, Frazzini and Malloy, 2008; Ingram and Zou, 2008).

On the other hand, however, homophily could induce social conformity and groupthink that can lead to inefficient decision-making (Asch, 1951; Janis, 1982; Ishii and Xuan, 2014). Individuals in homophilic relationships often have an enhanced desire for unanimity and ignore, or insufficiently consider, the disadvantages of the favored decision as well as the advice from experts outside the group. Nonetheless, other research has suggested that the presence of salient demographic differences legitimizes divergent perspectives and thus improves decision-making (Phillips, Liljenquist and Neale, 2009; Phillips and Loyd, 2006; Sommers, 2006). Consequently, under an alternative hypothesis, more diverse firms might perform better because decision-making under uncertainty is improved. Therefore, estimating the performance impact of diversity in a non-laboratory setting using a credible strategy is an important step to guide any subsequent attempts to enact sensible diversity-related policies.

The rest of the paper is organized as follows: Section 3.2 discusses our data. Our methodological approach is outlined in section 3.3. Section 3.4 presents a discussion of our results, both the hiring level regressions as well as the performance results. Section 3.5 concludes.

3.2 Data Collection

The core data used in this paper comprise several parts. The first element of our data involves collecting a comprehensive data set of all venture capital partners as well as their demographic and family information. The second element consists of a panel data set of venture capital firm hiring events. The final data entail the deal-level and fund-level performance for each of our venture capital firms.

We start with VentureSource, a database that contains detailed information on venture capital investments. Our data cover the period from 1990 through mid-2016. We start our analysis in 1990 because the data become reasonably comprehensive at that point in time. The unit of observation in the data is venture capital-backed companies. For each portfolio company, we have the identities of the individuals involved with the firm including founders, venture capital investors, angel investors, board members, and early hires. We focus on the venture capitalists on the boards of directors. Venture capitalists who never serve on a board will not be identified in our data. We believe this is reasonable because most venture capitalists serve on the board of directors for companies for which they are the lead investor. Similarly, most venture capitalists highlight their active involvement in their portfolio companies via board representation. In addition to information about the people involved in the company, we also have information on the portfolio company's location and industry. A venture capitalist enters the data in the year they make their first investment for which they sit on the board of directors.

For each individual venture capitalist in the data set, we collect a broad range of biographical information such as gender, ethnicity, education, and prior job experience. We collect this information from a variety of sources, including a leading online resume website, web searches, SEC filings, and news articles. In particular, venture capitalist genders are primarily determined based on first names. In the cases of unisex names, we determine gender by reading news articles and web pages mentioning or containing pictures of the individual. Our overall match rates for gender exceeds 99%. A full detailed summary of the data is presented in Calder-Wang and Gompers (2017).

Our empirical approach is to focus on the effects of children's gender on the hiring choices of venture capital firms and how exogenous changes in gender diversity associated with children's gender affects venture capital investment outcomes. We therefore set out to collect a novel data set on the family information of venture capital partners including the number of children as well as the gender and age of each child which we summarize in Table 3.1. We obtain information from a total of 1,310 individuals from various sources including college and business school directories and reunion books (61.6%), direct email solicitation (34.7%), and Marquis Who's Who database (2.9%). For email solicitation, we sent out over 3,000 emails and obtained 454 responses. If we do not obtain a child's gender explicitly but have the child's name, we assign a best-guess gender based on the first name. Overall, we are able to identify gender for over 98% of venture capital partners' children in our data.

Panel (a) of Table 3.1 provides descriptive statistics for our data on children. Venture capital partners in our data set have on average 2.39 children and 1.14 daughters as of 2016. For 70.5% of the children we obtain their exact ages as well. Panel (a) also summarizes the gender and ethnic breakdown of our sample. Our sample mirrors the industry results in Calder-Wang and Gompers (2020). 9.9% of the venture partners for whom we have children information are female, 87% are white, 4.4% are South Asian, 5.3% are East Asian, and 3% are Hispanic. Panel (c) shows the distribution of boards for the venture partners in our sample. 35.2% have served on two or fewer boards. 43.9% have served on five or more company boards.

In constructing our sample, as long as we have children information on at least one partner from a given firm, we include that firm in our sample. We do not believe that this creates issues for our results because the partners from whom we obtained information are typically more senior and have an important role in making hiring decisions. Similarly, there should be no bias from using all firms for which we have children's gender for at least one partner. Table 3.2 compares the characteristics of the firms in our sample (i.e., for whom we have data on the gender of partners' children) with those for whom we have no data on children. In particular, our sample includes firms that have more partners (6.96 vs. 2.07), were founded earlier (1995 vs. 2003), and are more likely to be US-based (82% vs. 61%). Although our sample differs from those not in our sample, they do hire similar proportions of women (8.1% vs. 9.2% and statistically not different).

We define two measures of deal success. Our most conservative measure of success

Table 3.1: *Children Data Collection*

This table reports the characteristics of the venture capital partners from whom we collect children information.

(a) Characteristics of Venture Capital Partners

	N	Mean	SD	Min	Max
Number of Children	1310	2.389	1.07	0	7
Number of Daughters	1310	1.143	0.90	0	5
Number of Sons	1310	1.237	0.98	0	5
Female	1310	0.099	0.30	0	1
Whites	1310	0.869	0.34	0	1
South Asian	1310	0.044	0.21	0	1
East Asian	1310	0.053	0.22	0	1
Hispanic	1310	0.030	0.17	0	1
African American	1310	0.003	0.06	0	1
Children Age Available	1310	0.705	0.46	0	1

(b) Source of Children's Information

	N	Percent (%)
Email	454	34.7
Harvard Reunion Book	301	23.0
HBS Alumni Directory	299	22.8
Stanford Reunion Book	85	6.5
Princeton Reunion Book	74	5.6
Yale Reunion Book	48	3.7
Marquis	38	2.9
Other	11	0.8
Total	1310	100.0

Other includes Wikipedia, New York Times, Penn Alum Directory, and Qualtrics

(c) Career Deal Count

	N	Percent (%)
1	292	22.3
2	169	12.9
3	161	12.3
4	113	8.6
5 or More	575	43.9
Total	1310	100.0

Table 3.2: *Firm Sample Selection*

(a) This table characterizes the venture capital firms in our sample. Each observation is a venture capital firm.

	N	Mean	SD	SE	25%	50%	75%
VC Firms in Sample							
Average Partner Count	301	6.96	4.77	0.13	3.85	5.92	8.63
VC Founding Year	301	1995.2	7.33	0.16	1989	1997	2000
Firm Deal Count	301	64.5	75.3	0.50	20	38	78
Fraction of US Based Deals	301	0.82	0.30	0.03	0.81	0.97	1
Firm IPO Count	301	8.83	15.0	0.22	1	3	11
Firm IPO Rate	301	0.11	0.11	0.02	0.033	0.080	0.18
Firm Success Rate	301	0.23	0.13	0.02	0.14	0.22	0.30
Total Number of Hires	301	12.8	10.9	0.19	6	9	16
Total Number of Female Hires	301	1.12	1.68	0.07	0	1	2
Average Female Hired Ratio	301	0.081	0.11	0.02	0	0.029	0.14
VC Firms Not in Sample							
Average Partner Count	5757	2.07	1.92	0.02	1	1.33	2.50
VC Founding Year	5748	2003.4	6.88	0.03	1999	2002	2009
Firm Deal Count	5757	5.42	10.3	0.04	1	2	5
Fraction of US Based Deals	5757	0.61	0.46	0.01	0	1	1
Firm IPO Count	5757	0.51	1.61	0.02	0	0	0
Firm IPO Rate	5757	0.092	0.23	0.01	0	0	0
Firm Success Rate	5757	0.16	0.29	0.01	0	0	0.25
Total Number of Hires	5757	2.66	3.10	0.02	1	2	3
Total Number of Female Hires	5757	0.24	0.61	0.01	0	0	0
Average Female Hired Ratio	5757	0.092	0.24	0.01	0	0	0

(b) Sample Representativeness

	Percent (%)	Total N
% VC Firms in Sample	4.97	6058
% Deal in Sample	38.34	50543
% IPO in Sample	47.66	5579

is whether the company in which the venture capitalist invested goes public in an IPO. Because many successful companies are acquired by larger companies for a profit, we define successful deals as those that either go public in an IPO or get acquired for a higher value than the total investment in the company. We obtain acquisition values from Capital IQ when available. If we are unable to identify an acquisition value, we do not consider the investment a success. The IPO and success rates are modestly higher in our sample of venture capital firms. 11% of the deals for our sample firms go public and 23% go public or are acquired for more than the invested capital vs. 9.2% and 16.0% for firms not in our sample, averaging over firms. Economically, we believe that this is a relevant sample because these firms make disproportionately more deals (64.5 vs. 5.42) and hire more people (12.8 vs. 2.6). The empirical results from this group of firms are of great economic importance given they represent a large fraction of all deals (38.3%) done. Additionally, this selection is unlikely affected by the gender breakdown of the children, which is also what we need for the internal validity of the empirical results.

Next, we construct a panel of gender breakdowns for each firm's new hires, which allows us to test whether the gender of an existing partner's children can have an effect on the hiring of women. While we do not directly observe exactly when a particular venture capital partner is hired by a firm, we estimate the "hiring" event as the year before the person first sits on the board of a venture capital-backed company and represented the particular venture capital firm. Moreover, we approximate the "active" period of a partner's career as the year before the first board seat and three years after the last observed board seat.

Table 3.3 summarizes information on the venture capital firms and hiring level information. We have data on 1,645 venture capital partners in 301 venture capital firms who were hired during our sample period by the firms for which we have children information. 9.9% of the hires are female. The general pattern of low female hiring rates is consistent with the results of Calder-Wang and Gompers (2017). Our firms are larger with the average firm employing 12.8 partners over the entire sample period. At the time of the hiring events, the average number of daughters at the firm is 0.98 per partner and the average number of sons is 1.10. The average daughter ratio is 0.48 and approximates the birth rates by gender in the general population.

Panel (b) of Table 3.3 shows that our firms account for 10,987 deals of which 13.4% go public and 27.3% are successful. We match venture capital firms to the Preqin fund database. Preqin is relatively comprehensive on amounts raised but has data on only a fraction of fund returns. We identify fundraising information on 1,263 funds for the firms in our sample. The average fund raised $517.5 million while the median fund raised $230 million. We are able to obtain fund return information for 395 funds. The average fund internal rate of return (IRR) is 14.3% and the median fund IRR is 9.3%. Because investment outcomes and fund returns are highly dependent upon market conditions, we match our funds to median benchmark fund IRRs for funds raised in the same year, the same geographic region, and having the same investment strategy. We compute fund excess IRRs by subtracting the median fund benchmark IRR from the funds' IRR. The average fund excess IRR is 3.9%.

Table 3.4 summarizes the distribution of female hires by firms. We have partners' children gender information for 301 venture capital firms. 58.5% of our firms have never hired a female investor. 22.6% have hired exactly one female investor. 18.9% have hired more than one female investor. Not surprising, the number of women hired is monotonically related to firm size.

Even though venture capital firms are very small in size, we still examine the fraction of females hired as a percentage of all hires at firms of various sizes in Table 3.4. This controls for any correlation between the number of hires and the female hired ratio. The average female hired ratio for firms with fewer than five partners is 10.9%. As firms grow, there is no significant trend in the fraction of total hires that are females. For firms with 15 or more partners, the female hired ratio is 11.1%. The standard deviation of the female hired ratio is also similar across venture capital firm size. This gives us confidence that there is significant heterogeneity of the propensity to hire a female within firm size groups.

Table 3.6 tabulates the hiring rate for females by the time period of the hire. The female

Table 3.3: *Summary Statistics*

(a) Venture Capital Firm Characteristics and Children Metrics: Hiring Level Observations

	N	Mean	SD	25%	50%	75%
Female	1645	0.099	0.30	0	0	0
Partner Count	1645	12.9	9.59	6	10	15
VC Firm Age	1645	13.2	7.42	7	12	18
Avg Daughters	1645	0.98	0.59	0.50	1	1.33
Avg Sons	1645	1.10	0.69	0.60	1	1.50
Avg Daughters (Senior)	1617	0.97	0.62	0.50	1	1.33
Avg Sons (Senior)	1617	1.08	0.71	0.50	1	1.50
Avg Daughters (Junior)	486	0.91	0.77	0	1	1.33
Avg Sons (Junior)	486	1.00	0.80	0	1	1.40
Daughter Ratio	1602	0.48	0.26	0.33	0.50	0.67
Average Daughter Ratio	1601	0.49	0.27	0.33	0.50	0.67
Daughter-Heavy Partner Fraction	1645	-0.069	0.61	-0.50	0	0.33
First Daughter Partner Fraction	1602	0.50	0.23	0.40	0.50	0.60
At Least One Daughter Fraction	1645	0.69	0.33	0.50	0.75	1

(b) Deal Performance: Deal-Level Observations

	N	Percent (%)	SD	25%	50%	75%
IPO	10987	13.4	0.34	0	0	0
Success	10987	27.3	0.45	0	0	1

(c) Fund Performance: Fund-Level Observations

	N	Mean	SD	25%	50%	75%
Excess Return	395	0.039	0.18	-0.039	0.0050	0.070
Net IRR	395	0.14	0.22	0.023	0.093	0.18
Median Fund Benchmark	434	0.10	0.082	0.034	0.100	0.15
Quartile	431	2.30	1.00	1	2	3
Amount Raised (USDmm)	1263	517.5	1192.8	90	230	500

Table 3.4: *Number of Female Hires*

This table breaks down the firm sample by the number of women hired during a firm's entire history.

	N	Percent (%)	Firm Size
Never Hired Women	176	58.5	5.9
Hired One Women	68	22.6	9.1
Hired Two Women	27	9.0	11.4
Hired Three Women	17	5.6	12.7
Greater Than Three	13	4.3	27.9
Total	301	100.0	8.4

Table 3.5: *Female Hired Ratio by Firm Size*

This table breaks down the hiring sample by the size of the firm.

	N	Female Hired (%)	SD	SE
Fewer than 5 Partners	129	10.9	0.31	0.049
5 - 9 Partners	653	9.2	0.29	0.021
10 - 14 Partners	405	9.4	0.29	0.027
More than 15 Partners	458	11.1	0.31	0.026
Total	1645	9.9	0.30	0.013

hired ratio does not vary substantially over time. Before 1994, the female hired ratio was 10.3%. The female hired ratio increased to 12.3% between 2005 to 2009, but declined to 8.1% between 2010 and 2016. These results are consistent with the industry-wide summaries in Calder-Wang and Gompers (2017) that showed no meaningful trend in the hiring of female venture capital investors.

Table 3.6: *Hiring Patterns Over Time*

This table breaks down the hiring sample by the year of the hire.

	N	Female Hired (%)	SD	SE
Before 1994	97	10.3	0.306	0.056
1995-1999	355	8.7	0.283	0.028
2000-2004	508	9.3	0.290	0.024
2005-2009	463	12.3	0.329	0.027
After 2010	222	8.1	0.274	0.035
Total	1645	9.9	0.299	0.013

Table 3.7 tabulates the ratio of deals done by the woman in our sample by industry. Across the 10,937 deals, only 7.0% of the deals are led by women venture capital partners. Healthcare has the highest percentage of female-led deals at 13.5%. The consumer goods industries and consumer services industries have female lead investors serving on the board 9.3% and 7.63% of the time. Information technology has the lowest rate of female-led deals at 4.1%.

Finally, in Table 3.8 we examine the demographics and career statistics for male and female hires in our sample. We include data on all partners who are hired, not just those from firms for which we have information on the gender of partners' children. First, we look at schooling. We tabulate the fraction of hires that have undergraduate degrees from a top ten college. Top ten colleges are defined as the ten most frequent undergraduate institutions, namely Harvard, Stanford, University of Pennsylvania, Princeton, Yale, Dartmouth, University of California, Berkeley, Cornell, MIT and Duke. A slightly higher fraction of male hires (30%) went to a top ten school than women (25%). Next, we look at the fraction of hires with an MBA and the fraction with an MBA from a top five program. Top five MBA

Table 3.7: *Industry Patterns*

This table summarizes the fraction of deals made by women across industries.

	N	Female (%)	SD	SE
Business and Financial Services	1975	5.6	0.219	0.011
Consumer Goods	86	9.3	0.292	0.058
Consumer Services	1335	7.6	0.251	0.014
Energy and Utilities	180	4.6	0.186	0.032
Healthcare	2409	13.5	0.328	0.012
Industrial Goods and Materials	148	10.1	0.291	0.044
Information Technology	4804	4.1	0.188	0.006
Total	10937	7.0	0.244	0.005

programs are defined as the five most frequent business schools, i.e., Harvard, Stanford, University of Pennsylvania, Columbia, and the University of Chicago. Nearly half of all new venture capital partners have an MBA degree. 53% of male hires and 48% of female hires have an MBA degree. 34% of male hires have an MBA from a top-five program while 31% of women have a top five degree. Finally, we look at the fraction of venture capitalists with a graduate degree e.g., masters' degree, PhD, JD, or MD, excluding MBAs. 38% of men and 43% of women who are hired as venture capitalists have a graduate degree other than an MBA.

Table 3.8 also tabulates career statistics for these hires. On average, male venture capital hires do more deals on which they serve on the board (6.37) than their female counterparts (5.10) over the course of their careers.[3] Interestingly, success rates are virtually identical for both men and women. Male venture capital hires have a 22% success rate on their investments and 11% go public. For female venture capital hires, 22% are successfully exited while 13% go public.

[3] As noted earlier, we can only identify a partner's connection to a deal if they are explicitly noted on the board of directors. Our venture capitalists almost certainly have done more deals than this. Amornsiripanitch, Gompers and Xuan (2019) show that venture capitalists get board seats approximately one third of the time.

Table 3.8: *Partner Characteristics by Gender*

	Men	Women	Difference	p-Value
Top 10 Colleges	0.30	0.25	0.048	0.065
MBA	0.53	0.48	0.053	0.063
MBA (Top 5)	0.34	0.31	0.024	0.373
Graduate Degree	0.38	0.43	-0.047	0.090
Success	0.22	0.22	-0.0043	0.794
IPO	0.11	0.13	-0.020	0.115
Deal Count	6.37	5.10	1.27**	0.003
N	3463	333	3130	.

Notes: This table includes information from all partners in the firm sample. Top 10 colleges are defined as the ten most frequent undergraduate institutions, namely Harvard, Stanford, University of Pennsylvania, Princeton, Yale, Dartmouth, UC Berkeley, Cornell, MIT and Duke. Top 5 MBA are defined as the five most frequent business schools, namely Harvard, Stanford, University of Pennsylvania, Columbia, and University of Chicago. Graduate degree includes masters' degree, PhD, JD, and MD.

3.3 Methodology

The work of Gompers, Mukharlyamov and Xuan (2016) and Calder-Wang and Gompers (2017) suggest that homophily is a strong force that affects collaboration and hiring decisions in the venture capital industry. Our empirical approach is to examine whether having daughters debiases venture capital hiring decisions. From the work of Warner and Steel (1999) and Washington (2008), we know that the gender of one's children affects parental behavior in the political arena. Politicians with more daughters are more likely to support feminist policies and women's issues relative to other issues. In this paper, we examine whether the same type of debiasing affects hiring decisions in venture capital. Also, because the gender of one's children is exogenous, we examine how differences in children's gender affects investment performance and, conditional on the validity of our exclusion restriction, whether greater gender diversity affects that performance.

The thought experiment is as follows. A venture capital partner and his/her spouse decide to have a child. Nature randomly assigns the gender of the child. Importantly, our empirical set-up conditions on the total number of children, while estimating the relative effect of having a daughter versus a son, which we refer to as the "daughter effect" in this paper. One can interpret the coefficient on the daughters' variable as the effect of replacing one son with one daughter.

$$Y_{i,t} = \beta_1 \text{ \#Daughters}_{i,t} + \beta_2 \text{ \#Children}_{i,t} + \text{Controls}_{i,t} + \epsilon_{i,t} \tag{3.3.1}$$

Our first analysis looks at the hiring events for our firms. For each hiring event, we run a regression in which $Y_{i,t}$ is the gender of the hire i that occurs at time t. On the right-hand side of Eq (3.3.1), *number of daughters* and *number of children* refer to the average number of daughters or children among the existing partners of the firm. We also divide partners who were present at the time of the hire into senior and junior partners. Senior partners are defined as those with an investment tenure of more than three years.[4] We control for a variety of other venture capital firm characteristics that may influence firm hiring decisions. These include the age of the venture capital firm at the time of the hiring event, the average age of the existing partners, the number of active partners in the firm, and the size of the fund defined as the logarithm of the capital per partner.

In Eq (3.3.1), β_1 identifies the relative effect of having an additional daughter as compared to an additional son. It is important that we condition on the total number of children because we know that people who choose to have more children are more likely to have different beliefs (Washington, 2008). However, once we condition on the total number of children, the gender distribution can be more reliably thought of as a random variable uncorrelated with the error. Additionally, since the total number of children, the number of daughters, and the number of sons are linearly dependent, we cannot differentiate whether the venture capital behavior is related to having a daughter, not having a son, or a combination of both.

The important identifying assumption is that conditioning on the total number of

[4]By this definition, senior partners account for 54% of the partner sample and they account for 73% of the hiring-partner pairs. At the hiring event level, over 95% of the hires are made with at least 1 senior partner present. The senior partners are more active, taking a median of 7 board seats vs 2 board seats for junior partners.

children, the number of daughters is exogenously assigned by nature. This requires that parents are not giving birth using a gender-based stopping rule or practicing any type of direct sex-selection. It is this natural experiment setting that allows us to identify a causal relationship between the relative number of daughters and the female hired ratio as well as its effect on venture capital performance.

We first rule out sex-selection that can skew the sex ratio. Given that direct sex-selection through abortions is uncommon in the US, it is not surprising that we find that male-to-female ratio in our sample of children is not statistically different from the natural male-to-female birth ratio in the overall population. This is true if we condition on the total number of children, or if we examine various subgroups, namely the senior partners, the junior partners, the male partners, and the female partners. Being able to recover the natural sex birth ratio in all subsamples gives us confidence in the integrity of our data. As such, we do not find evidence of sex-selection in our data.

Next, we want to rule out gender-based stopping rules. If parents employ a gender-based stopping rule which stipulates that they keep having children until they have at least one son, then conditioning on the total number of children, those who have more daughters would be more likely to be using such a stopping rule. To provide support for this identifying assumption, we run a number of tests. In particular, we find that having a first-born daughter does not predict the total number of children, consistent with the findings in Washington (2008). We tabulate these results in the appendix (Table C.1). Further, we also do not find statistical evidence of gender-stopping rules by testing whether the gender distribution is different from that of a binomial distribution with the natural sex birth rates conditioning on the total number of children. As such, the gender of the partners' children in our sample is considered truly random, and hence uncorrelated with the error. Our estimation of the form in Eq (3.3.1) can then identify the impact of the children's gender on female hiring.

In alternative specifications, we also consider other measures of children's gender breakdown, including the average ratio of daughters, the proportion of partners who have more daughters than sons, as well as the proportion of partners who have at least one daughter. All results are robust to these alternative specifications for the gender makeup of the existing partners' children. Additionally, we include control variables for firm size (partner count), venture capital firm age, the average existing partners' age, log capital per partner, and year fixed effects.

In addition to examining the effects of children's gender on hiring decisions, we instrument for gender diversity induced by having a hiring a female investor using children's gender to examine the causal effect of diversity on venture capital investment performance. These results are dependent upon the validity of our exclusion restriction which we discuss in detail after presenting our instrumental regression results. The performance effects are examined in two ways. First, we simply look at the reduced-form regression results: We examine a performance regression where deal- or fund-level performance is on the left-hand side and a variety of controls are on the right-hand side, including data on the gender of children for partners who were present when the current partners were hired (more details below on how this is constructed).

Our performance results exploit the exogenous nature of a venture capital partners' children's gender. We use the "number of daughters" relative to the total number of children as an instrument for the "female hired ratio." In this instrumental variable framework, we look at the performance effect of the exogenous component of shocks to gender diversity for a venture capital firm that is associated with the gender of existing partners' children. Our measure of a shock to the firm's gender diversity is the female hired ratio, i.e., looking at the time of a deal, what fraction of the active partners who were hired are female:

$$\text{Female Hired Ratio}_{i,t} = \gamma_1 \text{ \#Daughters}^h_{i,t} + \gamma_2 \text{ \#Children}^h_{i,t} + \text{Controls}^h_{i,t} + \epsilon_{i,t} \tag{3.3.2}$$

$$\text{Performance} = \theta_1 \text{Predicted Female Hired Ratio}_{i,t} \tag{3.3.3}$$

$$+ \theta_2 \text{ \#Children}^h_{i,t} + \text{Controls}^h_{i,t} + \omega_{i,t} \tag{3.3.4}$$

We employ a linear instrumental variable regression framework for estimation. Eq (3.3.2) and (3.3.4) present our instrumental variable set-up. In Eq (3.3.2), the dependent variable is a measure of gender diversity *Female Hired Ratio* for firm i in year t. It is defined as the number of female partners who have been hired at any prior time in the firm's existence who are still active (defined as having done at least one deal in the last three years) divided by the total number of hires who are still active using the same definition. As an instrument, we use the average number of daughters for the partners who were present at the time when an active partner was hired. The purpose of this procedure is to capture the numbers of daughters that are relevant for the hiring of the active partners (who were hired before and sometimes many years before the deal year), rather than the number of daughters at the time of the deal itself. This procedure also makes it consistent with the hiring specification outlined in Eq (3.3.1). We denote such variables by a superscript h. Additionally, there are a number of controls including the average number of children as well as other firm-level characteristics, similarly averaged over the hires, such as the firm size, average partner age, partner count, and log capital per partner.

Eq (3.3.4) represents our second stage regression. The endogenous regression involves regressing gender diversity on the deal- or fund-level outcomes. Here, *Predicted Female Hired Ratio* can be thought of as the fitted value from the first stage of the instrumental variable using the average number of daughters and various controls for deal n in year t. In addition to controlling for firm-level characteristics, we also control for deal-level characteristics, including industry, round, and country. Besides the random assignment of the children gender, for the identification of θ_1, the exclusion restriction required is that the gender of partners' children affects firm performance only through the gender of the hiring decisions made. We will turn to test possible alternative channels through which a partner's children's gender might affect investment performance after our main results.

In this set-up, we can estimate the effect of gender diversity on performance in venture capital. As discussed above, we run a variety of robustness checks throughout the results to ensure that our findings are not sensitive to the measure of the prevalence of daughters or the measure of the gender diversity of the venture capital firms. Our reduced-form results are also robust to a randomization style inference rather than a conventional inference, in which we make simulation draws of randomly assigned gender for the children in our data set.

3.4 Empirical Results

This section presents our empirical findings. We first analyze the causal relationship between the gender of existing partners' children and the hiring of female investment partners. Then, we analyze the reduced-form relationship between the gender of existing partners' children and investment performance. Finally, we use an instrumental variable framework to estimate the impact of the female hires on venture capital firm performance.

3.4.1 Effects on Venture Capital Hiring

In Table 3.9, we show the effect of daughters on the gender of new hires. As discussed earlier, our dependent variable is one if the gender of a new hire is female and zero otherwise. We express data on children by averaging across all the partners present when the individual was hired. We include the average number of daughters that existing partners have as well as the average number of children.[5] We also include a variety of firm-level controls including firm size (number of existing partners), firm age, the average partner age, and the size of the fund measured as log capital per partner. In Column (1), we observe a positive and significant coefficient on the average number of daughters, implying a positive relationship between having more daughters (holding the number of children constant) and the probability that the new hire is female. It is also important to note that holding the number of daughters constant, increasing the average number of children is correlated with a reduction in the probability of hiring a female. Adding additional firm-level controls does not change the magnitude of the effect that daughters have on the hiring decisions, with the coefficient remaining statistically significant at 5%. We also see that the hiring effect is limited entirely to senior partners. The gender of junior partners' children has no effect

[5] As previously discussed, our results are robust to expressing gender ratios in a variety of ways.

on the gender of a hire controlling for senior partners' children. Here, senior partners are those that have an investment tenure of more than three years. We expect that long-standing partners are more likely to have a greater role in hiring new partners. In this specification in Column (4), conditioning on the total number of children, the relative effect of having one more daughter for all existing senior partners increases the probability of hiring a female by 5.1%. Given that, on average, firms have a female hired ratio of 9.9%, this is a substantial increase of 50%.[6]

Figure 3.1 shows the main result from these regressions. In the first panel, we divide firms into those in which the existing partners have more daughters, have an equal number of daughters and sons, and have more sons. Firms with more daughters and an equal number of daughters and sons have a higher percentage of females that are hired (10.9% and 10.1%) than firms with more sons (9.2%). The pattern is even stronger when we look only at the gender of senior partners. For firms in which the existing senior partners have more daughters, the percentage hires that are female is 11.1%. Females represent 10.0% of new hires for firms in which existing senior partners have an equal number of sons and daughters. Finally, for firms in which existing senior partners have more sons, women represent 9.0% of the new hires.

We also run the hiring regressions with several alternative measures of the gender composition of existing partners' children. This is motivated by the concern that the potential effect may not be linear in the number of daughters relative to the total number of children. The dependent variable is a binary indicator of whether a given hire is a woman. We look at the original measure, the average number of daughters at the firm as well as the daughter ratio (defined as the ratio of total number of daughters to children at the firm), the average daughter ratio (defined as the average of the daughter-to-children ratio over active partners), daughter-heavy partner fraction (the fraction of partners with more daughters than sons less those with fewer daughters than sons), first daughter partner

[6]The standard deviation of the number of daughters is 0.9, implying that an increase of 1 daughter is slightly more than a 1.1 standard deviation increase.

Table 3.9: *Hiring Level Regression*

The dependent variable is a binary indicator of whether a given hire is a woman. We use the children metrics for the existing partners the year before the hire. *Avg Daughters* is the average number of daughters of the partners at the firm. *Avg Children* is the average number of children at the firm. Partners are identified in the deals they make when they take a board sear. We define whether a partner is present by the time window in which we observe them making deals. We extend it for two years at the beginning and three years at the end to approximate their active years at the firm. Senior partners are defined as those who make deals for more than three years. To approximate for hiring rather than founding, the sample is restricted to firms that have more than three active partners at the time and have been in existence for more than three years. Standard errors are clustered at venture capital firm and year level.

	(1) Female	(2) Female	(3) Female	(4) Female	(5) Female	(6) Female
Avg Daughters	0.0396**	0.0439**				
	(0.0174)	(0.0185)				
Avg Children	-0.0160	-0.0202*				
	(0.0109)	(0.0118)				
Avg Daughters (Sr)			0.0456***	0.0514***	0.0444***	0.0501***
			(0.0165)	(0.0174)	(0.0162)	(0.0172)
Avg Children (Sr)			-0.0129	-0.0162	-0.0123	-0.0157
			(0.0102)	(0.0112)	(0.0102)	(0.0112)
Avg Daughters (Jr)					0.0185	0.0178
					(0.0252)	(0.0251)
Avg Children (Jr)					-0.00456	-0.00448
					(0.0135)	(0.0137)
VC Firm Age		0.0000974		-0.0000857		-0.000102
		(0.00131)		(0.00134)		(0.00134)
Avg Partner Age		0.00115		0.00109		0.00104
		(0.00115)		(0.00116)		(0.00116)
Partner Count		0.000496		0.000847		0.000802
		(0.000874)		(0.000878)		(0.000895)
Log(Capital)		-0.000678		-0.000621		-0.00142
		(0.00640)		(0.00682)		(0.00685)
Year	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1645	1573	1617	1546	1617	1546

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

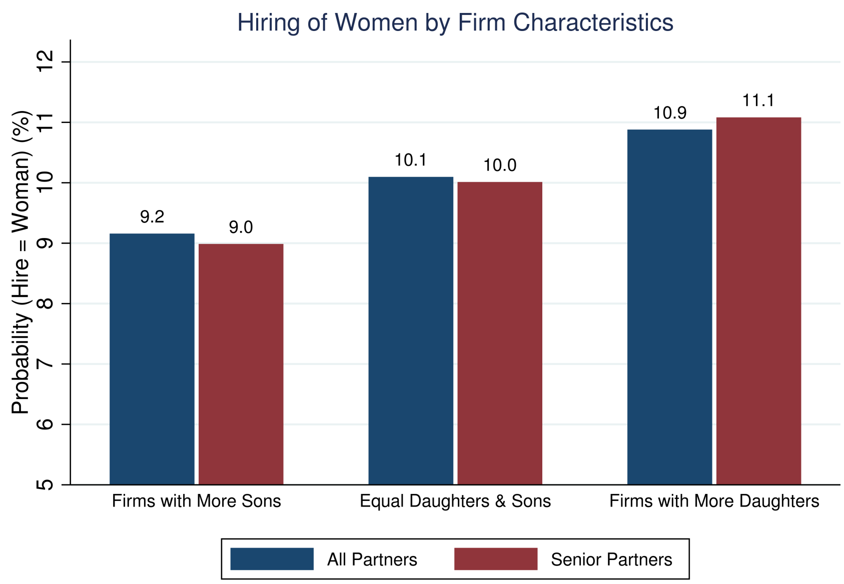

Figure 3.1: *The Probability of Hiring a Woman*

This figure plots the probability of hiring a woman based on the existing partners' children information. Firms are categorized into those with more sons, equal number of daughters and sons, and more daughters.

fraction (the fraction of partners at the firm whose first child is a daughter), and at least one daughter fraction (the fraction of partners who have at least one daughter at the firm). In Table 3.10, with the same controls including holding constant the number of children, we observe that the first five variables are all positive and the first four are statistically significant. The only measure of daughter intensity that is not positive is the fraction of partners that have at least one daughter, but the standard error is large, suggesting this definition of daughter-heaviness is not particularly informative. Results are qualitatively identical if we use data on all partner's children, shown in the appendix (Table C.4).

Table 3.10: *Hiring Level Regression (Alternative Measures of Daughters)*

The dependent variable is a binary indicator of whether a given hire is a woman. *Avg Daughters* is the original measure, the average number of daughters at the firm. *Daughter Ratio* is defined as the ratio of total number of daughters to the number of children at the firm. *Average Daughter Ratio* is the average of the daughter-to-children ratio over active partners. *Daughter-Heavy Partner Fraction* is the fraction of partners with more daughters than sons, less those with fewer daughters than sons. *First Daughter Partner Fraction* is the fraction of partners at the firm whose first child is a daughter. *At Least One Daughter Fraction* is the fraction of partners who have at least one daughter at the firm.

	(1) Female	(2) Female	(3) Female	(4) Female	(5) Female	(6) Female
Avg Daughters (Senior)	0.0514***					
	(0.0174)					
Daughter Ratio (Sr)		0.0657**				
		(0.0323)				
Average Daughter Ratio (Sr)			0.0593*			
			(0.0311)			
Daughter-Heavy Partner Fraction (Sr)				0.0360**		
				(0.0143)		
First Daughter Partner Fraction (Sr)					0.0479	
					(0.0341)	
At Least One Daughter Fraction (Sr)						-0.0117
						(0.0232)
Avg Children (Senior)	-0.0162	0.00988	0.00954	0.0106	0.00799	0.00722
	(0.0112)	(0.0106)	(0.0106)	(0.00985)	(0.0105)	(0.0104)
VC Firm Age	-0.0000857	-0.000209	-0.000219	-0.000116	-0.000242	-0.000265
	(0.00134)	(0.00139)	(0.00139)	(0.00134)	(0.00139)	(0.00135)
Avg Partner Age	0.00109	0.00122	0.00124	0.000899	0.00131	0.00123
	(0.00116)	(0.00119)	(0.00119)	(0.00116)	(0.00120)	(0.00118)
Partner Count	0.000847	0.00120	0.00114	0.000900	0.00102	0.000554
	(0.000878)	(0.000963)	(0.000963)	(0.000877)	(0.000962)	(0.000879)
Log(Capital Per Partner)	-0.000621	-0.000274	-0.000405	-0.00115	-0.000398	-0.00139
	(0.00682)	(0.00713)	(0.00708)	(0.00681)	(0.00709)	(0.00675)
Year	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1546	1484	1484	1546	1485	1546

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Since the source of randomization is the gender of the children, we also conduct statistical inference using a randomization test. Specifically, we randomly assign the gender of the children in the data set of all partners, holding the birth years and the total number of children the same as the original data set. We regress the gender of the hire on the children and firm-level characteristics just as before. Compared to the coefficient distribution produced by 1000 simulations, the true coefficient has a p-value smaller than 5% for the specification with all partners and less than 1% for the specification with senior partners, both shown in the appendix. Taken together, we are confident that when existing partners have relatively more daughters, there is a positive relationship with hiring more female investors.

3.4.2 Effects on Venture Capital Performance

In the prior section, we established a link between having a greater fraction of children who are daughters and hiring more female partners. In this section, we explore the performance implications of these effects. We first look the reduced-form regressions to explore the relationship between children's gender and performance. Clearly, given that the gender of children is randomly assigned, it is exogenous relative to investment performance. We regress the deal- or fund-level performance on children's gender. Since multiple deals or funds can be associated with a given venture capital firm, we make sure the firm identity, the fund identity, and the deal are all appropriately matched for the purpose of our reduced-form regression:

$$Y_{i,n,t} = \alpha_1 \ \#\text{Daughters}^h_{i,t} + \alpha_2 \ \#\text{Children}^h_{i,t} + \text{Controls}_{i,n,t} + \omega_{i,n,t} \tag{3.4.1}$$

At the deal-level, $Y_{i,n,t}$ is a success indicator for a deal n made by firm i in year t, and it is defined as successful if the investment exited via an IPO or high value acquisition. $\#\text{Daughters}^h_{i,t}$ refers to the average number of daughters by partners of firm i who contributed to the hiring of active partners present in year t.[7] Besides the firm-level controls

[7]In the case where a deal is funded by a number of venture capital firms, it will be counted as separate

such as firm age, firm size (partner count), fund size (log capital per partner), and partner age, we also add deal-level controls including the industry, the country, the funding round. Analogously, for the fund-level regressions, $Y_{i,n,t}$ is the net IRR achieved by the fund, while $\#\text{Daughters}^h_{i,t}$ is similarly defined for the fund raising year t.

In Table 3.11, the dependent variable is a binary "success" indicator based on whether the deal has resulted in an IPO or a successful acquisition where the acquisition value is greater than the amount of capital invested. We see a positive and significant coefficient on the number of daughters across all specifications controlling for the number of children. Like the hiring results, we find the effect of children's gender to be larger for senior partners. In the main specification with senior partners' daughter information in Column (4), the point estimate suggests that a relative increase of one daughter on average leads to an increased probability of success by 3.2%. Compared with the overall success rate of 27.3%, this is an economically meaningful magnitude. Therefore, in a reduced-form, we find strong evidence of a relationship between the gender of a venture capitalists' children and performance.

We also find a positive significant coefficient for the firm size. Firms with more partners have greater investment success. Similarly, venture capital age is positively related to success rates. This is consistent with the survival of better performing firms and the persistence in venture capital investment performance (Kaplan and Schoar, 2005; Gompers, Kovner, Lerner and Scharfstein, 2010). Surprisingly, we find that venture capital partner age is negatively related to success controlling for firm age and size.[8] As before, we further support our findings by performing a randomization test by comparing the actual coefficients with the distribution of simulated coefficients obtained using the same specification but with randomly assigned children's gender with details in the .

So far, we have been measuring performance at the deal-level with binary outcomes, but

---

observations.

[8]We also present the reduced-form result if we focus just on IPO in the appendix. IPO alone may not be a good measure of success because IPO rates have generally declined over the past decade and the importance of high value acquisitions have increased. We find moderately statistically significant results for the number of daughters of all existing partners and we find that the t-statistics increases if we focus only on the senior partners' children gender.

Table 3.11: *Daughter Effect on Performance (Deal-Level Reduced-Form)*

This table reports reduced form results of the deal level sample. The dependent variable *Success* equals to 1 if the portfolio company went public or was acquired with acquisition value greater than invested amount. Independent variables are the averages of existing partners' children and firm characteristics when current partners were hired. The sample of deals are restricted to those made after the first fund is raised and before 2014 to allow the investment outcomes to have time to realize. Standard errors are clustered at venture capital firm, year level.

	(1) Success	(2) Success	(3) Success	(4) Success
Avg Daughters	0.0225***	0.0235***		
	(0.00856)	(0.00870)		
Avg Children	-0.0209***	-0.0166***		
	(0.00599)	(0.00615)		
Avg Daughters (Senior)			0.0298***	0.0318***
			(0.00896)	(0.00904)
Avg Children (Senior)			-0.0241***	-0.0230***
			(0.00668)	(0.00665)
VC Firm Age		0.00212**		0.00212**
		(0.00101)		(0.00100)
Avg Partner Age		-0.00146**		-0.00157**
		(0.000683)		(0.000663)
Partner Count		0.00297***		0.00316***
		(0.00112)		(0.00110)
Log(Capital Per Partner)		0.00597		0.00525
		(0.00521)		(0.00521)
Year FE	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes
Observations	10435	10435	10435	10435

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

there may be a meaningful difference between two "successful" exits in terms of the actual rate of return that is achieved. Our deal-level analysis is limited by the lack of comprehensive deal-level return data as well as the fact that we do not have the structure of the deals and share class preferences which affect the ultimate realized IRR for any venture capital investment. Fortunately, we are able to match a meaningful portion of the venture capital funds in our sample to the Preqin Funds database in which we can access the fund-level IRRs. We have return information for 395 of 1263 funds in our sample and perform the same reduced-form regression as before controlling for log fund size. Because IRRs vary by investment focus and year, we use excess IRR, defined as the fund-level IRR minus the median fund return for venture capital funds raised in the same year and geographic region.

Despite the limited sample size, consistent with the findings in the deal-level sample, Table 3.12 shows positive and statistically significant coefficient for the number of daughters, i.e., our reduced-form regression indicates a positive relationship between the fund return and the number of daughters controlling for the total number of children. Like all the previous results, the effect of children's gender is stronger for senior partners. In Column (4), we find that the relative effect of having a daughter over a son is a 4.56% increase in excess return for the fund. In comparison, the average net IRR is 14.0% and the average excess return is 3.9% for the funds in our sample.

Our two main results establish that having a greater number of daughters controlling for the number of children for venture capital partners, especially for senior partners, leads to a significant increase in the proportion of female partners hired. We also saw in the reduced-form regression, that there is a significant improvement in the firm's investment performance. Not only does the statistical significance remains robust across different specifications, but the economic magnitude of the estimated coefficients is meaningful: The relative effect of having a daughter instead of a son increases the female hired ratio by about 5.1%, compared with a base rate of 9.9%. It lifts deal success by about 3.2% relative to an overall success rate of 27.3%.

Table 3.12: *Daughter Effect on Performance (Fund-Level Reduced-Form)*

This table reports reduced-form result in the fund level sample. The dependent variable is the excess return of the fund, defined as the net internal rate of return less the median fund benchmark. The median fund benchmark is defined as the median fund return in each region and year, as provided by Preqin. Independent variables are the averages of existing partners' children and firm characteristics when the current partners were hired. Partners are considered current if they are active as of the fund closing year.

	(1) Excess Return	(2) Excess Return	(3) Excess Return	(4) Excess Return
Avg Daughters	0.0439**	0.0388*		
	(0.0192)	(0.0198)		
Avg Children	-0.0275**	-0.0152		
	(0.0119)	(0.0136)		
Avg Daughters (Senior)			0.0480**	0.0456**
			(0.0198)	(0.0205)
Avg Children (Senior)			-0.0345***	-0.0267**
			(0.0116)	(0.0126)
VC Firm Age		0.000740		0.000722
		(0.00243)		(0.00243)
Avg Partner Age		-0.00309*		-0.00272*
		(0.00164)		(0.00152)
Partner Count		-0.000288		0.0000214
		(0.00116)		(0.00119)
Log(Capital Per Partner)		-0.0160*		-0.0152
		(0.00969)		(0.00980)
Year FE	Yes	Yes	Yes	Yes
Observations	371	371	371	371

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

3.4.3 Instrumental Variable Regression

Having established a strong, positive relationship between having more daughters relative to the total number of children and hiring female investors as well as fund performance, we next explore an instrumental variables specification in which we identify exogenous increases in gender diversity and its effect on investment performance. In particular, we use the average number of daughters and the average number of children of the existing partners as an instrument for the variations in the female hire ratio. For the specification to be a feasible empirical strategy, we need the instrument to be relevant for a firm's gender diversity, and the hiring regression suggests this is likely the case. We also need the gender of these children to be randomly assigned, independent of potential outcomes for the firm, which is also very likely. Finally, the relevant exclusion restriction is that having daughters only affects venture capital investment performance through the proportion of female partners hired.

We are sympathetic to the possibility that the gender of partners' children can affect performance through alternative channels. We discuss some of these alternative channels through which our exclusion restriction could be violated. Additional data are collected to test these channels to the extent possible, and the analyses follow after the instrumental variable results. For example, we do not find evidence that having more daughters increases the percentage of female entrepreneurs within a partner's investment portfolio. We also do not find that general sensitivity affects the role allocated to female investors, i.e., when senior partners have more daughters, female investors are not assigned more board seats nor do they have longer investment tenures. Moreover, we do not find that individuals with more daughters are more successful themselves, the improvement in performance is a broader firm-level improvement. We do find an interaction effect in which the performance of female venture capitalists is enhanced by having senior partners with more daughters. Taken together, we find the exclusion restriction plausible, but we acknowledge that there may be other alternatives that our data are unable to rule out. We provide more details and suggestions for future research at the end of the section.

We employ a linear two-stage least square (2SLS) estimation of our instrumental variable regressions. Table 3.13 presents both ordinary least squares (OLS) and 2SLS estimates for our deal-level performance regressions in which success is our outcome measure. In the OLS regression, we use the actual female hired ratio at the time of the deal, and in the 2SLS, we use the predicted value, *Predicted Female Hired Ratio*, from the first stage regression as our measures of shocks to gender diversity. Our OLS results show that the female hired ratio is not related to deal-level performance. The coefficients are small and negative. By contrast, the instrumental variable results are positive and significant. When we instrument for *Female Hired Ratio* with the average number of daughters for all partners, *Predicted Female Hired Ratio* is positive and significantly related to deal-level success. When we use the gender of senior partners' children as instruments, the results are even stronger. The coefficient of 0.942 in Column (6) implies that if the female hired ratio increases by 5%, the deal success rate would increase by 4.7%. With an overall success rate of 27.3% in our deal-level sample, this represents a 17% increase in the success rates. As we saw with the reduced-form, venture capital firm age and size (partner count) are positively related to performance.

Comparing the OLS with the IV, we believe there could be a number of omitted variables that can cause the OLS estimator to be either biased upward or downward. On the one hand, one might a priori expect higher quality firms to hire more diverse candidates, biasing the OLS upward. However, any "window-dressing" motives in hiring women or minorities by firms can produce a number of negative effects or be correlated with different firm characteristics under which female investors perform poorly, possibly biasing the OLS downward. Additionally, given the cyclical nature of the venture capital business, there could be time-varying omitted variables (for instance unobserved over-optimism during booms) that influence both the hiring of women and the subsequent performance. Given such an array of possible omitted variables in the OLS, we view our daughter-instrumental variable framework, despite all its limitations, as a valuable contribution to understanding the performance impact of diversity.

In Table 3.14, we present results for the first stage regressions corresponding to Columns

(3) through (6) in Table 3.13. The dependent variable is, as discussed in Eq (3.4.1), the *Female Hired Ratio* for deal n, in year t, for firm i. As our hiring regressions demonstrated, the average number of daughters for existing partners, controlling for the average number of children, is positive and statistically significant. Once again, the gender of senior partners' children has a more pronounced effect on hiring in the first stage. The economic significance of the effect is also significant.

Table 3.13: *Deal-Level Instrumental Variable Regression*

This table reports regression result of deal success in deal-level sample using the average number of daughters as the instrument. The dependent variable *Success* equals to 1 if the deal went public or was acquired with acquisition value greater than invested amount. *Female Hired Ratio* is the number of active female partners divided by the total number of active partners. In the instrumental variable regression, the instruments are the average number of existing partners' daughters when the hires (now active partners) were made. The sample of deals are restricted to those made after the first fund is raised and before 2014 to allow the time for realization of investment outcomes. Standard errors are clustered at venture capital firm and year level.

	(1)	(2)	(3)	(4)	(5)	(6)
	Success	Success	Success	Success	Success	Success
Female Hired Ratio	-0.0106	-0.0236	0.823**	0.895**	0.873***	0.942***
	(0.0383)	(0.0380)	(0.352)	(0.387)	(0.299)	(0.315)
Avg Children	-0.0133**	-0.00907	0.00124	0.00678		
	(0.00534)	(0.00553)	(0.00920)	(0.0101)		
Avg Children (Senior)					-0.00320	-0.000771
					(0.00723)	(0.00766)
VC Firm Age		0.00254**		0.000241		0.000190
		(0.00102)		(0.00156)		(0.00147)
Avg Partner Age		-0.00158**		-0.00102		-0.000753
		(0.000683)		(0.000855)		(0.000872)
Partner Count		0.00253**		0.00313**		0.00299**
		(0.00112)		(0.00133)		(0.00131)
Log(Capital Per Partner)		0.00707		0.00498		0.00496
		(0.00519)		(0.00647)		(0.00656)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Instrumented for Female Hired Ratio						
Average # Daughters	N/A	N/A	X	X		
Average # Daughters (Senior Partner)					X	X
First Stage F-stat			17.08	15.79	25.93	25.16
Observations	10435	10435	10435	10435	10435	10435

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table 3.14: *Deal-Level First-Stage Regression*

This table reports the first stage results of the deal level IV regression. The dependent variable *Female Hired Ratio* is the number of active female partners divided by the total number of active partners at the time of the deal. Independent variables are the averages of existing partners' children and firm characteristics when current partners were hired. Standard errors are clustered at venture capital firm and year level.

	(1) Female Hired Ratio	(2) Female Hired Ratio	(3) Female Hired Ratio	(4) Female Hired Ratio
Avg Daughters	0.0274***	0.0262***		
	(0.00663)	(0.00660)		
Avg Children	-0.0269***	-0.0261***		
	(0.00410)	(0.00442)		
Avg Daughters (Senior)			0.0341***	0.0338***
			(0.00670)	(0.00673)
Avg Children (Senior)			-0.0239***	-0.0236***
			(0.00444)	(0.00460)
VC Firm Age		0.00210***		0.00204***
		(0.000762)		(0.000759)
Avg Partner Age		-0.000493		-0.000866*
		(0.000474)		(0.000454)
Partner Count		-0.000181		0.000181
		(0.000656)		(0.000659)
Log(Capital Per Partner)		0.00110		0.000307
		(0.00391)		(0.00393)
Year FE	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes
Observations	10435	10435	10435	10435

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

In Table 3.15 we estimate OLS and instrumental variable regressions for fund-level excess IRR. We again use the *Female Hired Ratio* as the measure of gender diversity. Like the deal-level results, *Female Hired Ratio* is only weakly correlated with excess fund IRR in the OLS regressions. When we run the 2SLS, however, we find that *Predicted Female Hired Ratio* is positively and statistically significantly related to fund excess IRR, although results are somewhat weaker than our deal-level specifications. The lower significance level is driven primarily by the smaller sample size. Deals are collapsed into fund returns, reducing the number of observations by a factor of 10. Similarly, we only have return data on approximately one-fourth of our funds. This means fund return observations are only about 2.5% of the number of deal outcome observations. The economic magnitude of the effect also appears reasonable. A 5% increase in the female hired ratio increases fund excess IRR by between 4.2% and 4.7%.

In Table 3.16, we tabulate the results of the first stage regression for our 2SLS estimation of the impact of gender diversity on performance. *Female Hired Ratio* for these regressions is defined as the ratio of females who were hired at any time in the past who were active in the fund divided by the total number of historical hires who were active in the fund. The results look qualitatively identical to the first stage in the deal-level regression in Table 3.13. The differences arise because the analysis is at the fund-level and we have only 371 fund-level observations for excess IRR.

Table 3.15: *Fund-Level Instrumental Variable Regression*

This table reports regression result of success in the fund level sample. The dependent variable is the excess return of the fund, defined as the net internal rate of return less the median fund benchmark. The median fund benchmark is defined as the median fund return in each region and year, as provided by Preqin. *Female Hired Ratio* is the number of active female partners divided by the total number of active partners. In the instrumental variable regression, the instruments are the average number of existing partners' daughters when the hires (now active partners) were made. Standard errors are clustered at venture capital firm and year level.

	(1) Excess Return	(2) Excess Return	(3) Excess Return	(4) Excess Return	(5) Excess Return	(6) Excess Return
Female Hired Ratio	0.0253	0.0343	0.819*	0.641*	0.942**	0.779*
	(0.0798)	(0.0841)	(0.421)	(0.357)	(0.471)	(0.398)
Avg Children	-0.00574	0.00573	0.00647	0.0157		
	(0.00894)	(0.0107)	(0.0123)	(0.0130)		
Avg Children (Senior)					-0.00496	0.00107
					(0.0109)	(0.0109)
VC Firm Age		0.00116		0.00380		0.00437
		(0.00250)		(0.00251)		(0.00268)
Avg Partner Age		-0.00373**		-0.00463**		-0.00379**
		(0.00167)		(0.00185)		(0.00180)
Partner Count		-0.000804		-0.00251		-0.00316*
		(0.00113)		(0.00166)		(0.00187)
Log(Capital Per Partner)		-0.0153		-0.0105		-0.00701
		(0.00960)		(0.0115)		(0.0122)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Instrumented for Female Hired Ratio						
Average # Daughters	N/A	N/A	X	X		
Average # Daughters (Senior Partner)					X	X
First Stage F-stat			11.67	13.46	9.99	12.15
Observations	371	371	371	371	371	371

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table 3.16: *Fund-Level First-Stage Regression*

This table reports the first stage results of the fund level sample. The dependent variable *Female Hired Ratio* is the number of active female partners divided by the total number of active partners at the time of the deal. Independent variables are the averages of existing partners' children and firm characteristics when current partners were hired. Standard errors are clustered at venture capital firm and year level.

	(1) Female Hired Ratio	(2) Female Hired Ratio	(3) Female Hired Ratio	(4) Female Hired Ratio
Avg Daughters	0.0536***	0.0606***		
	(0.0157)	(0.0165)		
Avg Children	-0.0415***	-0.0481***		
	(0.0105)	(0.0127)		
Avg Daughters (Senior)			0.0510***	0.0585***
			(0.0161)	(0.0168)
Avg Children (Senior)			-0.0314***	-0.0357***
			(0.0108)	(0.0115)
VC Firm Age		-0.00477**		-0.00468**
		(0.00185)		(0.00188)
Avg Partner Age		0.00240		0.00137
		(0.00160)		(0.00145)
Partner Count		0.00346**		0.00408***
		(0.00147)		(0.00157)
Log(Capital Per Partner)		-0.00873		-0.0105
		(0.0112)		(0.0111)
Year FE	Yes	Yes	Yes	Yes
Observations	371	371	371	371

Standard errors in parentheses
* $p < 0.1$, ** $p < .05$, *** $p < .01$

One potential concern is that having more daughters may lead the firm to invest in more companies with female founders. If the average quality of those entrepreneurs is higher than male entrepreneurs because they are overlooked by other firms, then their success rates would also be higher. As such, it would constitute an alternative channel in which the gender of partners' children may affect firm investment performance, but not via the channel of increased gender diversity at the firm itself. To test this, we collected data on the gender of the founders of venture capital-backed companies for the venture capital firms in our sample. The sample consists of 13,000 founders for portfolio companies of the venture capital firms in our sample. On average, portfolio firms in our sample have 2.1 founders, while only 6.1% of them are female. In Table 3.17, we test whether venture capital firms with more daughters invest in more female founders. The dependent variable is the fraction of the portfolio company's founders who are women. We do not find any evidence that having more daughters leads to more investment into female-founded companies at the firm-level. If we look at the firm-level founder ratio, a greater number of daughters relative to the total number of children leads to a slightly lower percentage of female founders. The presence of a female venture capitalist, however, in Column (2) is significantly related to the fraction of female founders. Therefore, if the channel through which children's gender affects performance is through investing in more female entrepreneurs, then it is only through actually hiring a female investor in which that channel operates. Having daughters in and of itself does not increase the fraction of female entrepreneurs. In Column (3), we match companies to the individual venture capitalist who invested in the deal. There is no significant relation between having more daughters and investment in more female founders at the individual level. The results remain robust regardless of whether we measure the daughters of all existing partners or just senior partners. In alternative specifications where the dependent variable is an indicator of the presence of any female founders in an investment, we also do not find it correlated with the number of daughters. Moreover, we also do not find evidence that individual venture partners with more daughters invest more in female founders. Overall, we conclude that the number of daughters does not seem to affect firm performance through the gender of the founders in which they invest.

One may also contend that the effects of having daughters comes not only from the extensive margin of hiring more woman, but also the intensive margin such as assigning existing female employees more responsibilities, mentoring them better, as well as other unobserved channels. We test this hypothesis by directly controlling for the gender of venture capitalists as well as its interaction with the fraction of daughters that existing senior partners have and find that our results remain robust. In Table 3.18, we run a deal-level performance regression, where Column (1) shows that a higher daughter to children ratio leads to better firm performance, consistent with the previous section. Column (2) shows that, on average, deals led by female venture capitalists do not perform significantly differently from their male counterparts. Interestingly, the interaction term in Column (3) suggests that female venture capitalists perform better in firms with more daughters, consistent with this alternative explanation. In some sense, this is not unexpected if we believe that having more daughters has a subtle debiasing effect on how people work with female colleagues. When existing partners have more daughters, not only does the probability that a firm hires a female investor increase, but existing partners are likely to serve as better mentors and, hence, those women perform better. The coefficient in front of the daughter ratio is somewhat smaller and remains positive and significant when controlling for this interaction. Therefore, our results indicate that firm performance improvement is not entirely driven by hiring female investors who perform better than her colleagues. Rather, the entire firm performs better. This may indicate better decision-making through greater diversity or greater overall deal flow. We discuss these specific mechanisms and potential future research that can address this question in our conclusion.

Table 3.17: *Daughter Effect on Entrepreneurs*

The dependent variable *Female Founder Ratio* measures the fraction of a portfolio company's founding team who are women. Independent variables are the averages of existing partners' children and firm characteristics at the time of the investment.

	(1) Female Founder Ratio Female Founder Ratio	(2) Female Founder Ratio	(3) Female Founder Ratio	(4) Female Founder Ratio
Daughter Ratio (Individual)			0.0115 (0.00832)	0.0117 (0.00830)
Number of Children (Individual)			-0.00535 (0.00377)	-0.00516 (0.00375)
Daughter Ratio (Firm)	-0.0150** (0.00660)	-0.0154** (0.00658)		
Number of Children (Firm)	0.000226 (0.00218)	0.000262 (0.00217)		
Female VC		0.0199** (0.00813)		0.0242 (0.0169)
Partner Count	-0.000349 (0.000389)	-0.000362 (0.000388)	-0.000297 (0.000871)	-0.000383 (0.000881)
VC Firm Age	0.0000346 (0.000377)	-0.0000438 (0.000372)	-0.000439 (0.000751)	-0.000536 (0.000732)
Avg Partner Age	0.00000925 (0.000255)	0.0000388 (0.000253)	-0.000104 (0.000559)	-0.0000226 (0.000548)
Log(Capital Per Partner)	-0.00264 (0.00204)	-0.00276 (0.00204)	-0.00630 (0.00449)	-0.00595 (0.00447)
Year FE	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes
Observations	10081	10081	2647	2647

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table 3.18: *Daughter Effect on Performance Controlling for Venture Capitalist Gender*

This table reports a reduced form analysis for the deal level sample, controlling for the gender of the deal-maker. The dependent variable *Success* equals to 1 if the portfolio company went public or was acquired with acquisition value greater than invested amount. The independent variable *Female VC* is a binary indicator for whether the individual partner who made the investment is a woman. *Daughter Ratio* is the ratio of total number of daughters to the number of children at the firm, which is a fraction between 0 and 1. All other dependent variables are the same as the reduced-form regression. Standard errors are clustered at venture capital firm and year level.

	(1) Success	(2) Success	(3) Success
Daughter Ratio (Sr)	0.0654***	0.0657***	0.0549***
	(0.0175)	(0.0176)	(0.0179)
Female VC		-0.0117	-0.0852**
		(0.0171)	(0.0362)
Female VC x Daughter Ratio (Sr)			0.145**
			(0.0633)
Avg Children (Sr)	-0.00850	-0.00852	-0.00842
	(0.00614)	(0.00613)	(0.00612)
Partner Count	0.00350***	0.00350***	0.00365***
	(0.00112)	(0.00112)	(0.00112)
VC Firm Age	0.00211**	0.00216**	0.00203**
	(0.00102)	(0.00102)	(0.00102)
Avg Partner Age	-0.00145**	-0.00146**	-0.00144**
	(0.000696)	(0.000696)	(0.000694)
Log(Capital Per Partner)	0.00486	0.00493	0.00503
	(0.00541)	(0.00540)	(0.00538)
Year FE	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes
Observations	10081	10081	10081

Standard errors in parentheses
* $p < 0.1$, ** $p < .05$, *** $p < .01$

Another alternative channel for the performance impact of children's gender could be that having more daughters in some way directly improves a partner's investment-related skills, e.g., negotiation or communication skills, allowing venture capitalists to better source or close deals. We test this alternative by controlling for the children's gender of the individual venture capital partner who actually made the investment. In Table 3.19 Column (3), we find that venture capital partners with more daughters themselves do not have more successful deals. The coefficient is negative, but statistically insignificant. We caveat this by noting that the sample size is reduced because we only have children information for a relatively smaller set of venture capital partners. We also acknowledge that there are potentially other alternative explanations that we are not able to directly test. Taken together, we think the instrumental variable framework provides us with suggestive evidence that greater gender diversity has a positive impact on venture capital performance.

We also explore the possibility that when senior partners have daughters, they give more authority to female investors. Perhaps the greater authority given to female investors in these firms is the channel through which performance improves. In Table A7 we look at the number of board seats allocated to male and female investors. We do not find that senior partners having daughters increase the number of board seats held by female investors. Similarly, in Table A8, we look at the career tenure of venture capitalists dependent upon the gender of the investor and interacted with the number of daughters and number of children for senior partners. Once again, we do not find that female investors have longer career tenure in firms in which senior partners have relatively more daughters. These results appear to rule out the possibility that improvement in performance is driven by giving female investors a greater role in the firm.

In this section, we discuss a number of additional robustness tests. One concern regarding the sample is that about 34% of the children information is obtained from email solicitations. If the respondents are self-selected in terms of their parental involvement, this could bias the results. When we run the same analysis excluding email respondents while including only those whom we obtain information from public sources, we find that the daughter effects on female hiring and in the reduced-form deal performance regressions remains robust (Table C.3). Similarly, the instrumental variable results continue to hold. Statistical significance is slightly weaker, however, due to the reduced sample size. All of our results continue to hold with similar magnitude and statistical significance. Additionally, we

Table 3.19: *Daughter Effect on Performance Controlling for Individual Venture Capitalist Family Characteristics*

This table reports a reduced-form analysis for the deal-level sample, controlling for the children characteristics of the deal-maker. The dependent variable *Success* equals to 1 if the portfolio company went public or was acquired with acquisition value greater than invested amount. *Daughter Ratio (Individual)* is the ratio of number of daughters to children for a given venture capital partner, which is a fraction between 0 and 1. *Number of Children (Individual)* denotes the total number of children for a given venture capital partner. All other dependent variables are the same as the reduced form regression. Standard errors are clustered at venture capital firm and year level.

	(1) Success	(2) Success	(3) Success	(4) Success
Daughter Ratio (Individual)			-0.0198	-0.0212
			(0.0252)	(0.0251)
Number of Children (Individual)			-0.00352	-0.00439
			(0.00923)	(0.00918)
Daughter Ratio (Firm)	0.0654***	0.0657***		
	(0.0175)	(0.0176)		
Number of Children (Firm)	-0.00850	-0.00852		
	(0.00614)	(0.00613)		
Female VC		-0.0117		-0.112***
		(0.0171)		(0.0280)
Partner Count	0.00350***	0.00350***	0.00230	0.00269
	(0.00112)	(0.00112)	(0.00256)	(0.00254)
VC Firm Age	0.00211**	0.00216**	0.00394**	0.00439**
	(0.00102)	(0.00102)	(0.00196)	(0.00195)
Avg Partner Age	-0.00145**	-0.00146**	0.000542	0.000162
	(0.000696)	(0.000696)	(0.00143)	(0.00143)
Log(Capital Per Partner)	0.00486	0.00493	0.0107	0.00909
	(0.00541)	(0.00540)	(0.0101)	(0.0101)
Year FE	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes
Observations	10081	10081	2647	2647

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

also examine the results for the gender of male and female partners separately. Our results remain for the male partners, while results for female partners are not significant due to the dramatically smaller sample size for female partners. Another concern regarding the sample is that we only obtain information about birth years in 70% of the children. Among them, since over 90% of the children were born before the parent takes his or her first board seat, we assume that for those children we did not have the birth year, they were present throughout the parent's career tenure. As a robustness check, if we simply drop those for whom children's birth years were not available, we find that our main results still hold (Table C.2). Relatedly, we investigate whether the age of children matters for hiring and performance. In Table 3.20, we include both the number of daughters over the age of 12 and the number of daughters under 12 at the time of the hiring. Interestingly, our results show that the number of teenage daughters, rather than the number of pre-teen daughters, matters more for the hiring of female investment partners. This might suggest that older daughters have more of an effect on the attitudes of their fathers. This is consistent with fathers observing potential gender biases that their daughters face as they get older. Finally, we also conduct extensive robustness analyses to variations in the outcome measures. Results are robust to using IPO as the only success measure (Table C.5). Results are also robust, and in fact even stronger, when we restrict the sample to U.S.-only deals and U.S.-focused funds, presumably because the data quality is the highest for them (Table C.6 and Table C.7).

Table 3.20: *Hiring Level Regression (Daughter Age Effects)*

This table compares the impact of older daughters and younger daughters on firm hiring. The sample includes partners whose children age information is available. The dependent variable is a binary indicator of whether a given hire is a woman. Standard errors are clustered at venture capital firm and year level.

	(1) Female	(2) Female	(3) Female	(4) Female
Avg Daughters	0.0478**			
	(0.0199)			
Avg Daughters (>= 12 Years)		0.0468**		0.0622**
		(0.0238)		(0.0253)
Avg Daughters (< 12 Years)			0.0145	0.0367
			(0.0225)	(0.0236)
Avg Children	-0.0219*	-0.0105	-0.00761	-0.0212*
	(0.0122)	(0.0103)	(0.0113)	(0.0122)
VC Firm Age	0.0000906	0.000177	0.000423	0.0000257
	(0.00137)	(0.00137)	(0.00139)	(0.00137)
Avg Partner Age	0.000758	-0.000726	0.00146	-0.000126
	(0.00126)	(0.00140)	(0.00149)	(0.00150)
Partner Count	0.000622	0.000352	0.0000357	0.000684
	(0.000920)	(0.000907)	(0.000899)	(0.000921)
Log(Capital Per Partner)	0.00239	0.00177	0.00213	0.00219
	(0.00684)	(0.00688)	(0.00686)	(0.00686)
Year	Yes	Yes	Yes	Yes
Observations	1428	1428	1428	1428

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

3.5 Conclusion

Persistent gender disparity in a variety of professions has been a focus of both academic research and popular media. Diversity has been lauded as an important cornerstone of modern civil society and our contemporary workplace, but there have been few rigorous studies, to our knowledge, that estimate the causal economic impact of a diverse workforce in a real business setting. In this paper, we address the effects of gender diversity by collecting a unique data set of the gender of venture capitalists' children and taking advantage of a research design, in which this gender is exogenous to the individual partner. Combined with the time series of investment professional hiring and deal/fund-level performance, we establish that a increase in the number of daughters relative to the number of children leads to a significant and economically meaningful increase in the proportion of females hired. In reduced-form regressions, a higher relative fraction of daughters is related to increases in deal success rates and fund-level excess IRRs. Exploiting the exogenous nature of children's gender, when the relative fraction of daughters is used as an instrument for shocks to gender diversity, the results suggest that the exogenously induced increase in firm gender diversity leads to improvement in venture capital performance. These results provide evidence in a real business setting with strong profit motives that performance can be improved with greater gender diversity. As discussed in Gompers and Kovvali (2018), however, it is worth noting that this result does not necessarily imply that implementing a blunt gender quota would bring about the same positive outcomes. Improvement in diversity through having daughters is facilitated by debiasing existing partners. Gompers and Kovvali (2018) discuss how mandatory diversity programs often lead to resentment and reductions in performance. Our result that female venture capitalists perform better in firms in which senior partners have more daughters relative to total children suggests that this is potentially the case in venture capital as well. Having daughters might make senior partners better mentors for female venture capitalists.

Our work highlights the importance of understanding the role that this subtle removal of gender bias has for increasing diversity. The subtle debiasing effects of having daughters that prior research has shown to influence US Congressmen's votes (Washington, 2008) and Federal judicial rulings (Glynn and Sen, 2015) also play a role here in causing fathers to increase the likelihood that they hire a female investor. Our results suggest that diversity achieved through genuine removal of a bias or a change in beliefs could lead to better economic outcomes than mandated gender ratios. Future research efforts should explore other means of achieving similar debiasing.

There are several potential explanations for the mechanism by which a diverse investment team performs better. First, having daughters reduces the bias that one has towards women, which leads to more female hires. Given that the pool of potential female investors is relatively untapped, these female investors could be of higher quality than the counterfactual male hires. The higher quality hires then generate higher returns. Our results on the educational background of male and female hires, however, do not find substantial differences in the background of male and female hires. It is possible, however, that there are unobserved measures of quality that we cannot identify. A second potential explanation is that having a diverse set of backgrounds around the table to make decisions about investments could reduce correlated errors in judgment. Since homophily in hiring in venture capital is strong, most venture capital firms are populated by men of the same ethnicity with similar schooling and work histories. Different perspectives can reduce groupthink and allow venture capital firms to avoid costly investment mistakes. This explanation would be consistent with overall firm improvement. Examining how deal investment memoranda change after firms hire female investors could potentially shed light on whether decision-making improves with diversity. Similarly, collecting information on how investment decisions are made (Gompers, Gornall, Kaplan and Strebulaev, 2020) can also help to establish if diversity operates through the decision-making processes. Third, because so much of venture capital investment success is driven by having access to the best deals, having more diverse backgrounds can attract broader deal flow and, hence, increased average quality of deals. Collecting data on the deal funnel at venture capital firms can be a fruitful way to explore this channel. We believe that future research on these potential mechanisms will be very fruitful for understanding the sources of performance improvement that greater gender diversity engenders.

References

ALMAGRO, M. and DOMÍNGUEZ-IINO, T. (2020). Location Sorting and Endogenous Amenities: Evidence from Amsterdam. *Working Paper*.

ALONSO, W. (1964). Location and Land Use. *Harvard University Press*.

AMORNSIRIPANITCH, N., GOMPERS, P. A. and XUAN, Y. (2019). More than Money: Venture Capitalists on Boards. *Journal of Law, Economics, and Organization*, **35** (3), 513–543.

ANDREWS, D. W. and GUGGENBERGER, P. (2009). Hybrid and Size-Corrected Subsampling Methods. *Econometrica*, **77** (3), 721–762.

ASCH, S. E. (1951). Effects of Group Pressure upon the Modification and Distortion of Judgments. In *Groups, Leadership and Men*, Oxford, England: Carnegie Press, pp. 177–190.

AUTOR, D., PALMER, C. and PATHAK, P. (2017). Gentrification and the Amenity Value of Crime Reductions: Evidence from Rent Deregulation. *NBER Working Paper*.

AUTOR, D. H., PALMER, C. and PATHAK, P. A. (2014). Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge Massachusetts. *Journal of Political Economy*, **122** (3), 661–717.

BANZHAF, H. S. and WALSH, R. P. (2008). Do People Vote with Their Feet? An Empirical Test of Tiebout's Mechanism. *American Economic Review*, **98** (3), 843–863.

BARRON, K., KUNG, E. and PROSERPIO, D. (2018). The Sharing Economy and Housing Affordability : Evidence from Airbnb.

BAUM-SNOW, N. and HAN, L. (2020). The Microgeography of Housing Supply. (2010), 1–59.

BAYER, P., FERREIRA, F. and MCMILLAN, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, **115** (4).

—, KEOHANE, N. and TIMMINS, C. (2009). Migration and Hedonic Valuation: The Case of Air Quality. *Journal of Environmental Economics and Management*, **58** (1), 1–14.

— and TIMMINS, C. (2005). On the Equilibrium Properties of Locational Sorting Models. *Journal of Urban Economics*, **57**, 462–477.

— and — (2007). Estimating Equilibrium Models for Sorting Across Destinations. *The Economic Journal*, **117**, 353–374.

Beaman, L., Chattopadhyay, R., Duflo, E., Pande, R. and Topalova, P. (2009). Powerful Women: Does Exposure Reduce Bias? *Quarterly Journal of Economics*, **124** (4), 1497–1540.

Benetton, M. (2018). Leverage Regulation and Market Structure: An Empirical Model of the UK Mortgage Market. *SSRN Working Paper Series*.

Bennedsen, M., Nielsden, K. M., Perez-Gonzales, F. and Wolfenzon, D. (2007). Inside the Family Firm: The Role of Families in Succesion Decisions and Performance. *The Quarterly Journal of Economics*, **122** (2), 647–691.

Bernile, G., Bhagwat, V. and Yonker, S. (2018). Board Diversity, Firm Risk, and Corporate Policies. *Journal of Financial Economics*, **127** (3), 588–612.

Berry, S., Levinsohn, J. and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, **63** (4), 841–890.

—, — and — (1999). Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy. *American Economic Review*, **89** (3), 400–430.

—, — and — (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy*, **112** (1), 68–105.

Black, S. E. (1999). Do Better Schools Matter? Parental Valuation of Elementary Education. *Quarterly Journal of Economics*, **114** (2), 577–599.

Bohnet, I., van Geen, A. and Bazerman, M. H. (2016). When Performance Trumps Gender Bias: Joint Versus Separate Evaluation. *Management Science*, **5**, 1225–1234.

Bordalo, P., Coffman, K., Gennaioli, N. and Shleifer, A. (2016). Stereotypes. *Quartely Journal of Economics*, **131** (4), 1753–1794.

Buchak, G., Matvos, G., Piskorski, T. and Seru, A. (2018). Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks. *Journal of Financial Economics*, **130** (3), 453–483.

Calder-Wang, S. (2019). No Land Like Disneyland: The Impact of Airbnb Restriction on Home Prices. *Working Paper*.

— and Gompers, P. A. (2017). Diversity in Innovation. *NBER Working Paper*.

Campbell, J., Giglio, S. and Pathak, P. (2011). Forced Sales and House Prices. *American Economic Review*, **53** (9), 1689–1699.

Chamberlain, G. (1987). Asymptotic Efficiency in Estimation with Conditional Moment Restrictions. *Journal of Econometrics*, **34** (3), 305–334.

Chernozhukov, V., Hong, H. and Tamer, E. (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, **75** (5), 1243–1284.

Ciliberto, F. and Tamer, E. (2009). Market Structure and Multiple Equilibria in Airline Markets. *Econometrica*, **77** (6), 1791–1828.

COHEN, L., FRAZZINI, A. and MALLOY, C. (2008). The Small World of Investing: Board Connections and Mutual Fund Returns. *Journal of Political Economy*, **116** (5), 951–979.

CRONQVIST, H. and YU, F. (2017). Shaped by Their Daughters: Executives, Female Socialization, and Corporate Social Responsibility. *Journal of Financial Economics*, **126** (3), 543–562.

CUESTA, J. I. and SEPULVEDA, A. (2019). Price Regulation in Credit Markets: A Trade-Off between Consumer Protection and Credit Access. *Working Paper*.

DAVIS, L. W. (2011). The Effect of Power Plants on Local Housing Values and Rents. *Review of Economics and Statistics*, **93** (4), 1391–1402.

DIAMOND, R., MCQUADE, T. and QIAN, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, **109** (9), 3365–3394.

DUBÉ, J.-P., FOX, J. T. and SU, C.-L. (2012). Improving the Numerical Performance of BLP Static and Dynamic Discrete Choice Random Coefficients Demand Estimation. *Econometrica*, **80** (5), 2231–2267.

EDELMAN, B., LUCA, M. and SVIRSKY, D. (2017). Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*, pp. 1–34.

EPPLE, D. (1987). Hedonic Prices and Implicit Markets : Estimating Demand and Supply Functions for Differentiated Products. *Journal of Political Economy*, **95** (1), 59–80.

EWENS, M. and TOWNSEND, R. R. (2020). Are Early Stage Investors Biased Against Women? *Journal of Financial Economics*, **135** (3), 653–677.

FARRONATO, C. and FRADKIN, A. (2018). The Welfare Effects of Peer Entry in the Accommodation Market: The Case of Airbnb. *NBER Working Paper*.

FAVILUKIS, J. Y., MABILLE, P. and VAN NIEUWERBURGH, S. (2018). Affordable Housing and City Welfare. *SSRN Electronic Journal*.

FERREYRA, M. M. (2007). Estimating the Effects of Private School Vouchers in Multidistrict Economies. *American Economic Review*, **97** (3), 789–817.

FILIPPAS, A., HORTON, J. J. and ZECKHAUSER, R. J. (2019). Owning, Using and Renting: Some Simple Economics of the 'Sharing Economy'. *Working Paper*.

GALIANI, S., MURPHY, A. and PANTANO, J. (2015). Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment. *American Economic Review*, **105** (11), 3385–3415.

GANONG, P. and SHOAG, D. (2017). Why Has Regional Income Convergence in the U. S. Declined? *Journal of Urban Economics*, **102**, 76–90.

GARCIA-LÓPEZ, M. À., JOFRE-MONSENY, J., MARTÍNEZ-MAZZA, R. and SEGÚ, M. (2020). Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona. *Journal of Urban Economics*, **119** (June 2019).

Glynn, A. N. and Sen, M. (2015). Identifying Judicial Empathy: Does Having Daughters Cause Judges to Rule for Women's Issues? *American Journal of Political Science*, **59** (1), 37–54.

Goldin, B. C. and Rouse, C. (2000). Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians. *American Economic Review*, **90** (4), 715–741.

Gompers, P. A., Gornall, W., Kaplan, S. N. and Strebulaev, I. A. (2020). How Do Venture Capitalists Make Decisions? *Journal of Financial Economics*, **135** (1), 169–190.

—, Kovner, A., Lerner, J. and Scharfstein, D. (2010). Performance Persistence in Entrepreneurship. *Journal of Financial Economics*, **96** (1), 18–32.

— and Kovvali, S. (2018). The other diversity dividend. *Harvard Business Review*, **2018** (July-August).

—, Mukharlyamov, V., Weisburst, E. and Xuan, Y. (2021). Gender Gaps in Venture Capital Performance. *Journal of Financial and Quantitative Analysis*, pp. 1–58.

—, — and Xuan, Y. (2016). The Cost of Friendship. *Journal of Financial Economics*, **119** (3), 626–644.

Greenstone, M. and Gallagher, J. (2008). Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program. *Quarterly Journal of Economics*, (August), 951–1003.

Gyourko, J. (2009). Housing Supply. *Annual Review of Economics*, **1**, 295–318.

— and Molloy, R. (2015). Regulation and Housing Supply. *Handbook of Regional and Urban Economics*, **5**, 1289–1337.

Higgins, S. (2019). Financial Technology Adoption. *Working Paper*.

Horn, K. and Merante, M. (2017). Is Home Sharing Driving up Rents? Evidence from Airbnb in Boston. *Journal of Housing Economics*, **38**, 14–24.

Hsieh, C.-T. and Moretti, E. (2019). Housing Constraints and Spatial Misallocation. *American Economic Journal: Macroeconomics*.

Imbens, G. W. and Manski, C. F. (2004). Confidence Intervals for Partially Identified Parameters. *Econometrica*, **72** (6), 1845–1857.

Ingram, P. and Zou, X. (2008). Business Friendships. *Research in Organizational Behavior*, **28**, 167–184.

Ishii, J. and Xuan, Y. (2014). Acquirer-target Social Ties and Merger Outcomes. *Journal of Financial Economics*, **112** (3), 344–363.

Jaffe, S., Coles, P., Levitt, S. and Popov, I. (2019). Quality Externalities on Platforms: The Case of Airbnb. *Working Paper*.

Janis, I. L. (1982). *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*. Houghton Mifflin.

Jia, J. and Wagman, L. (2018). Platform, Anonymity, and Illegal Actors: Evidence of Whac-a-Mole Enforcement From Airbnb. *SSRN Electronic Journal*.

Kaplan, S. N. and Schoar, A. (2005). Private equity performance: Returns, persistence, and capital flows. *Journal of Finance*, **60** (4), 1791–1823.

Klaiber, H. A. and Phaneuf, D. J. (2010). Valuing Open Space in a Residential Sorting Model of the Twin Cities. *Journal of Environmental Economics and Management*, **60** (2), 57–77.

Klocke, U. (2007). How to Improve Decision Making in Small Groups: Effects of Dissent and Training Interventions. *Small Group Research*, **38** (3), 437–468.

Koijen, R. S. J. and Yogo, M. (2016). Shadow Insurance. *Econometrica*, **84** (3), 1265–1287.

Koster, H., Van Ommeren, J. and Volkhausen, N. (2019). Short-term Rentals and the Housing Market: Quasi-experimental Evidence from Airbnb in Los Angeles. *Working Paper*.

McFadden, D. (1978). Modeling the Choice of Residential Location. *Spatial Interaction Theory and Planning Models*, pp. 72–77.

Mcpherson, M., Smith-lovin, L. and Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, **27** (2001), 415–444.

Mills, E. S. (1967). An Aggregative Model of Resource Allocation in a Metropolitan Area. *American Economic Review: Papers & Proceedings*, **57** (2), 197–210.

Muth, R. F. (1969). Cities and Housing: the Spatial Pattern of Urban Residential Land Use. *University of Chicago Press*, (18), 200.

Nelson, S. T. (2019). Private Information and Price Regulation in the US Credit Card Market. *Working Paper*.

Nevo, A. (2000). Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry. *The RAND Journal of Economics*, **31** (3), 395.

— (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*, **69** (2), 307–342.

Pakes, A., Porter, J., Ho, K. and Ishii, J. (2015). Moment Inequalities and Their Application. *Econometrica*, **83** (1), 315–334.

Phillips, K. W., Liljenquist, K. A. and Neale, M. A. (2009). Is the Pain Worth the Gain? The Advantages and Liabilities of Agreeing with Socially Distinct Newcomers. *Personality and Social Psychology Bulletin*, **35** (3), 336–350.

— and Loyd, D. L. (2006). When Surface and Deep-level Diversity Collide: The Effects on Dissenting Group Members. *Organizational Behavior and Human Decision Processes*, **99** (2), 143–160.

Reynaert, M. and Verboven, F. (2014). Improving the Performance of Random Coefficients Demand Models: The Role of Optimal Instruments. *Journal of Econometrics*, **179** (1), 83–98.

Robles-Garcia, C. (2018). Competition and Incentives in Mortgage Markets: The Role of Brokers. *Working Paper*.

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, **82** (1), 34–55.

Saiz, A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics*, **125** (3), 1253–1296.

Schwartz-Ziv, M. (2017). Gender and Board Activeness: The Role of a Critical Mass. *Journal of Financial and Quantitative Analysis*, **52** (2), 751–780.

Small, K. A. and Rosen, H. S. (1981). Applied Welfare Economics with Discrete Choice Models. *Econometrica*, **49** (1), 105–130.

Solal, I. (2019). The Gender of Money: How Gender Structures the Market for Entrepreneurial Capital. *Working Paper*.

Sommers, S. R. (2006). On Racial Diversity and Group Decision Making: Identifying Multiple Effects of Racial Composition on Jury Deliberations. *Journal of Personality and Social Psychology*, **90** (4), 597–612.

Staiger, D. and Stock, J. H. (1997). Instrumental Variables Regression with Weak Instruments. *Econometrica*, **65** (3), 557.

Timmins, C. (2007). If You Cannot Take the Heat, Get out of the Cerrado... Recovering the Equilibrium Amenity Cost of Nonmarginal Climate Change in Brazil. *Journal of Regional Science*, **47** (1), 1–25.

Tra, C. I. (2010). A Discrete Choice Equilibrium Approach to Valuing Large Environmental Changes. *Journal of Public Economics*, **94** (1-2), 183–196.

Valentin, M. (2019). The Effects of Regulating the Housing Short-Term Rental Market: Evidence From New Orleans. *Working Paper*, **16802**, 1–38.

Wachsmuth, D. and Weisler, A. (2018). Airbnb and the Rent Gap: Gentrification through the Sharing Economy. *Environment and Planning A: Economy and Space*, **50** (6), 1147–1170.

Warner, R. L. (1991). Does the Sex of Your Children Matter? Support for Feminism among Women and Men in the United States and Canada. *Journal of Marriage and the Family*, **53** (4), 1051.

— and Steel, B. S. (1999). Child Rearing as a Mechanism for Social Change: The Relationship of Child Gender to Parents' Commitment to Gender Equity Stable. *Gender and Society*, **13**, 503–517.

Washington, E. L. (2008). Female Socialization: How Daughters Affect Their Legislator Fathers' Voting on Women' s Issues. *American Economic Review*, **98** (1), 311–332.

WONG, M. (2019). A Tractable Framework To Relate Marginal Willingness-To-Pay in Hedonic and Discrete Choice Models. *Journal of Housing Economics*.

Appendix A

Appendix to Chapter 1

A.1 Supplementary Figures

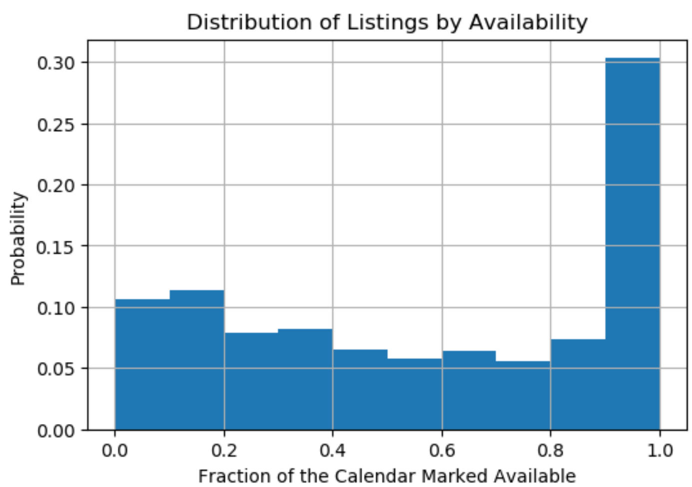

Figure A.1: *Distribution of Listings by Calendar Availability*

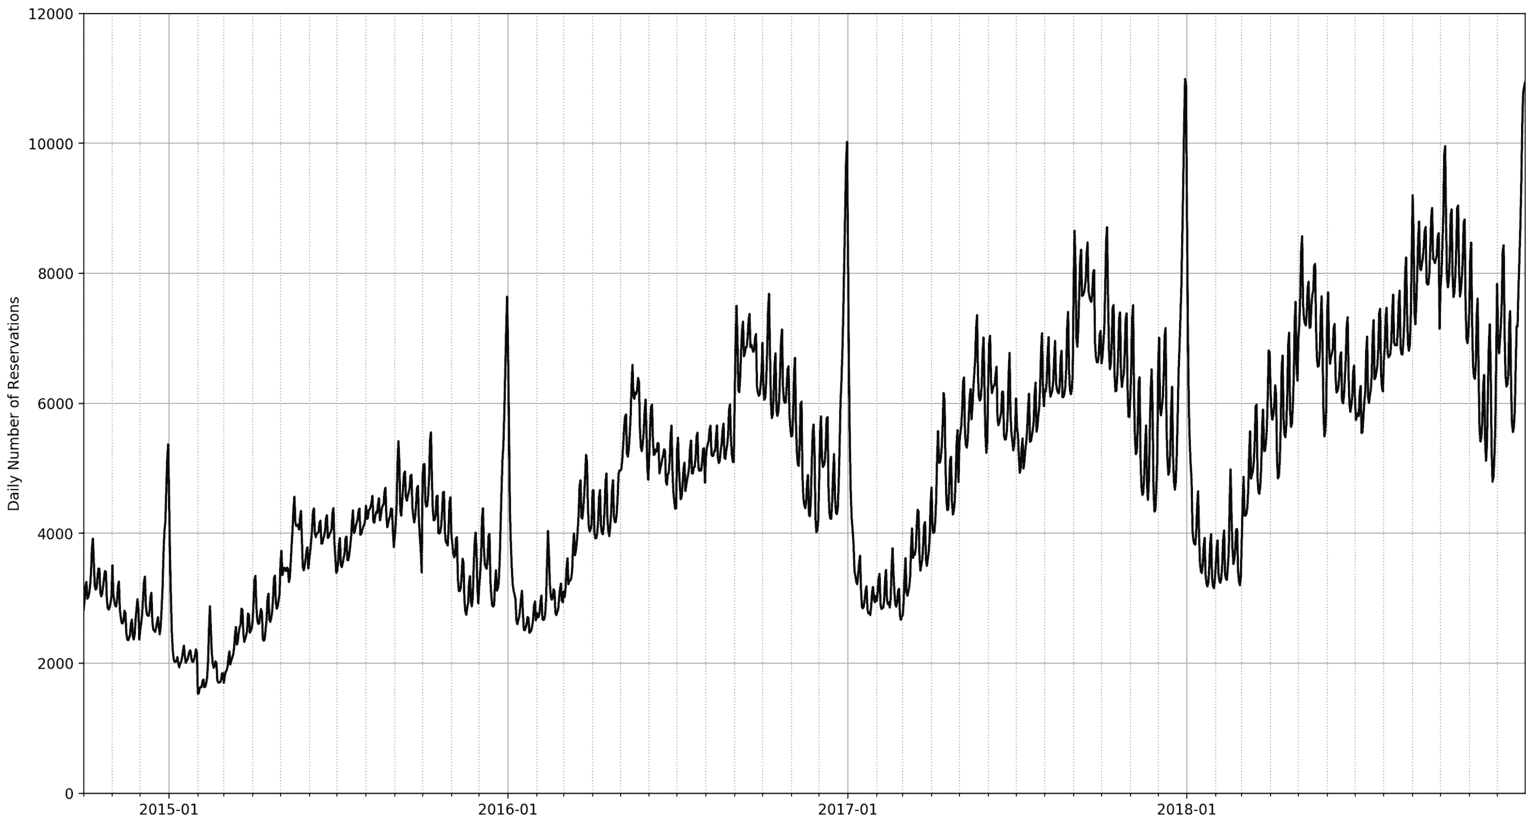

Figure A.2: *The Daily Number of Reservations of Private Rooms Sold on Airbnb in New York City.*

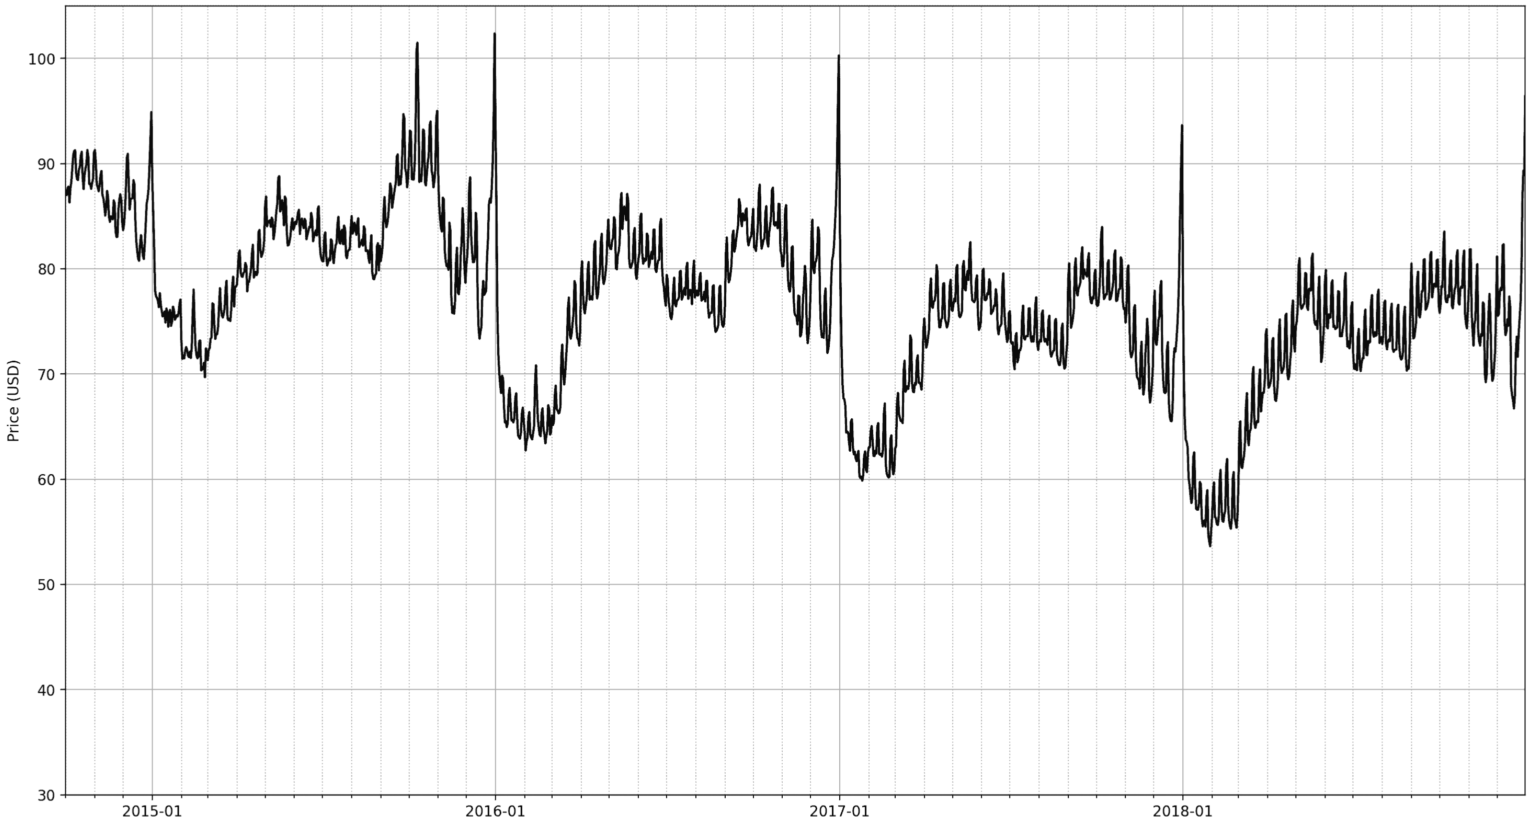

Figure A.3: *The Daily Average Price of Private Rooms sold on Airbnb in New York City.*

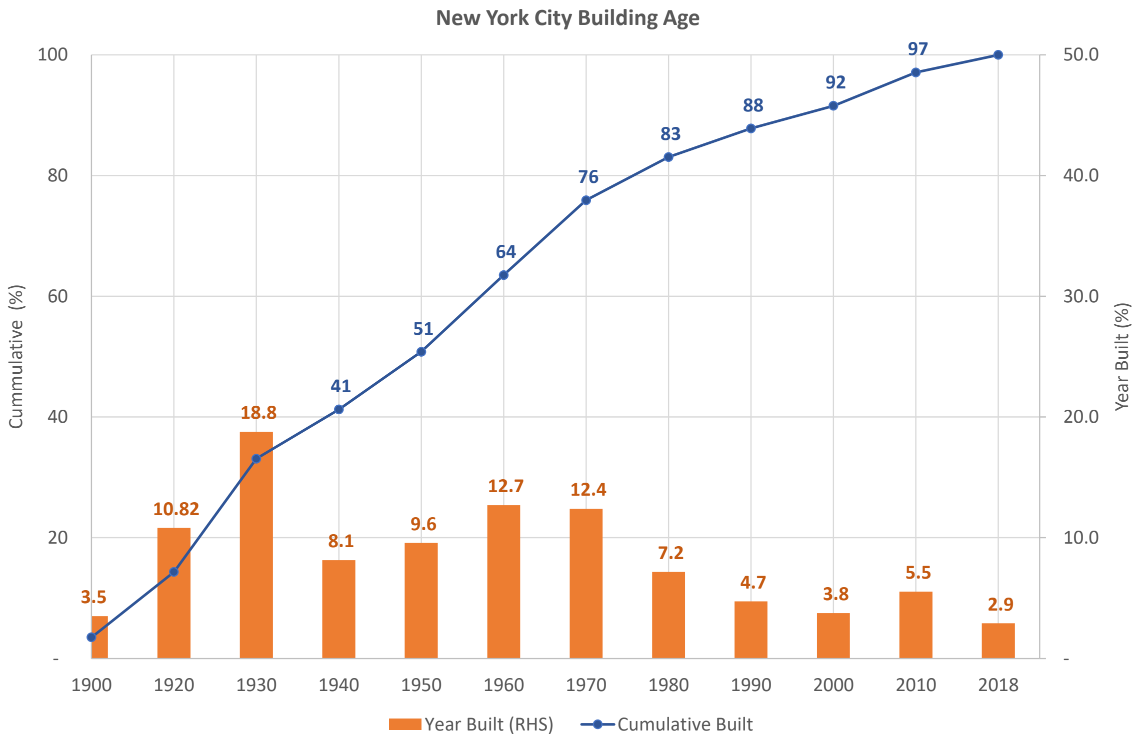

Figure A.4: *Building Age of NYC Housing Units*

Data based on ACS 2017 and New York City Housing and Vacancy Survey (2017). Over 40% of the housing units were built prior to 1940. Housing construction since the 1980s has remained depressed. 88% of the housing were built prior to 1990. Only 2.9% of the housing stock was built post-2010, whereas 3.5% of the units were built prior to 1900.

A.2 More on LTR Demand Estimation

The main idea is that one can use a supply-side pricing equation to produce a price instrument. The key intuition is that the availability of similar housing characteristics in the market has an impact on the equilibrium price of a given house, but is uncorrelated with the unobserved quality. Moreover, if I have two or more housing characteristics, then their relative availability allows this identification strategy to work even with only one cross-section of the market.

In this section, I will first describe the model, followed by the identification strategy, and end with some findings from a simulation study.

A.2.1 Model

Consider the following model with N households indexed by i and M homes indexed by j in a given housing market. Each home is endowed with a vector of physical features $X_{j,k}$ with $k = 0, \ldots, K$, where there are at least two features $K \geq 2$. Let $X_{j,0}$ denote the indicator for the outside option, with its price normalized to zero. Moreover, each home has an unobserved quality component ξ_j.

There is a random coefficient $\nu_{i,k}$ associated with each of the K features, drawn randomly from a normal distribution $\sigma \times \mathcal{N}(0, I)$. Here, I assume σ is known to simplify the exposition.[1]

The utility for household i renting home j is as follows:

$$U_{i,j} = \alpha p_j + \sum_k \beta_{i,k} X_{j,k} + \xi_j + \epsilon_{i,j} \tag{A.2.1}$$

where $\beta_{i,k} = \beta_k + \nu_{i,k}$.

Given that each household maximizes its utility, the choice probability of a household i

[1] In the actual model, it is captured by the coefficients in front of household demographics, which are estimated offline first using individual level choice data.

over home j becomes:

$$P_{i,j}(p, \mathbf{X}; \Theta) = \frac{exp(\delta_j + \lambda_{i,j'})}{\sum_{j'} exp(\delta'_j + \lambda_{i,j'})} \tag{A.2.2}$$

$$\delta_j(p, \mathbf{X}; \Theta) = \alpha p_j + \sum_k \beta_k X_{j,k} + \xi_j \tag{A.2.3}$$

$$\lambda_{i,j}(p, \mathbf{X}) = \sum_k \nu_{i,k} X_{j,k} \tag{A.2.4}$$

where $\Theta = (\alpha, \vec{\beta}_k, \xi)$.

The model is closed by a sorting equilibrium where the price p^e clears the market: The demand for each home equals the observed supply, namely $\forall j : s_j^F = 1$.

$$\forall j : \sum_i P_{i,j}(p^e, \mathbf{X}; \Theta) = 1 \tag{A.2.5}$$

A.2.2 Estimation

The key identification assumption is that the unobserved quality ξ is *independent* of the physical features of the home $\mathbf{X}$.

As a result, the unobserved quality ξ is uncorrelated with the price instrument z, which is constructed as the market clearing price with ξ set to zero.

$$\forall k : \mathbb{E}[\xi X_k] = 0 \tag{A.2.6}$$

$$\mathbb{E}[\xi z] = 0 \tag{A.2.7}$$

$$\forall j : \sum_i P_{i,j}\big(z, \mathbf{X}; (\alpha, \vec{\beta}, \xi = 0)\big) = 1 \tag{A.2.8}$$

$$\forall j : \sum_i P_{i,j}\big(p^e, \mathbf{X}; (\alpha, \vec{\beta}, \xi)\big) = 1 \tag{A.2.9}$$

Remark A.2.1. Why does the instrument work? The key intuition is that the supply-side model, namely that the demand clears the fixed supply of homes, implies that homes with rarer characteristics are going to have higher equilibrium prices, *compared with* homes with housing characteristics that are more common. And this component of variation (the scarcity premium) is uncorrelated with the unobserved quality of the particular home.

Suppose that there are only two relevant housing characteristics in the market: X_1 indicates that a home was built within the last 5 years, which is relatively rare in New York City, and X_2 indicates that a home has two bedrooms (as opposed to just one), which is more common. Both are desirable features. Without loss of generality, assume that the mean utilities for the two features are the same $\beta_1 = \beta_2$.

The key identification assumption is that the price premium observed for homes in brand-new buildings is higher than the price premium for homes with an extra bedroom; this is not because homes with two bedrooms have unobservably lower quality but rather because it is a much more common housing characteristic than being brand-new.

Remark A.2.2. It is essential that there be a distribution of preferences over the physical features of the homes. This is because for the IV strategy to work, the model has to produce equilibrium prices that are higher for rarer features, even if the valuations for it by the average household (namely, the household with $\nu_i = 0$) might be completely identical. In other words, the distribution of preferences over characteristics means that those who care a lot about a rare characteristic will bid up its price in equilibrium. In practice, as it is very likely that households of different income levels will have different price sensitivities, there will be a distribution over the willingness-to-pay $\beta_{i,k}/\alpha_i$ for housing attributes.

Remark A.2.3. It is also essential that there are at least two relevant housing characteristics $K \geq 2$ in the market for the estimation to work with just one cross-section. Otherwise, the concept of the rarity of a characteristic is undefined. Because I can compare the rarity of one housing characteristic with another housing characteristic in just one cross-section, it is also the reason why the instrument works even with just one cross-section of the market.

Remark A.2.4. It is also important that $\mathbf{X}$ is independent of ξ. Because the instrument z is a non-linear function of $\mathbf{X}$ as well as the parameters, the independence assumption ensures that z will be uncorrelated with ξ as a result.

Remark A.2.5. In general, there is a hump-shaped relationship between the mean utility (δ_j) and the price (or the price instrument): Namely, the mean utility for very expensive and very cheap homes tends to be lower than the mean utility of an average home. For expensive homes with rare and sought-after characteristics, prices will be too high compared to what an average household is willing to pay, resulting in low δ. On the other end, homes without these desirable characteristics are not going to have prices that are low enough for the average household to justify its lack of these characteristics. This is an inherent feature of the model, rather than evidence of the presence of ξ.

A.2.3 Simulation Study

In this section, I generate a very simple dataset with only two binary features X_1, X_2 with known parameters, where the first housing characteristic X_1 is much rarer than the second one X_2, even though the mean utilities on them are the same.

I show that the price instrument constructed in Eq (A.2.8) indeed recovers the true parameter, unlike the OLS. Moreover, I will also show that the typical hedonic regression will result in estimated WTP for amenities significantly different from the true model parameters, whereas the proposed IV estimation strategy does recover the WTP for the average household as specified by the model.

Simulation Set-up

More specifically, I have $N = M = 200$. Here is the breakdown of their characteristics with one home being the outside option (indicator by X_0):

(i) 20 homes have $X_1 = 1, X_2 = 1$

(ii) 20 homes have $X_1 = 1, X_2 = 0$

(iii) 120 homes have $X_1 = 0, X_2 = 1$

(iv) Remaining 39 homes have $X_1 = 0, X_2 = 0$

The true parameters of the simulation are set as follows:

$$\alpha = -1, \quad \beta_0 = 3, \quad \beta_1 = 1, \quad \beta_2 = 1 \tag{A.2.10}$$

Moreover, the unobserved quality ξ is drawn from a normal distribution with a standard deviation of 1/200. ν_i are drawn from a normal distribution with a standard deviation of 1 for both characteristics. With the model fully specified as above, I can solve the model by computing the equilibrium price $p^e(\mathbf{X};\alpha,\vec{\beta},\xi)$. Then, the mean utility δ is computed as follows:

$$\delta_j = \alpha p_j^e + \beta_0 X_{j,0} + \beta_1 X_{j,1} + \beta_2 X_{j,2} + \xi_j \tag{A.2.11}$$

Simulation Results

I compute the price instrument $z(\mathbf{X};\alpha,\vec{\beta},\xi=0)$ that clears the market. I checked that the price instrument is indeed uncorrelated with the unobserved quality with $cov(z,\xi)=0$.

I find that the IV recovers the true price coefficient, whereas the OLS produces a biased estimate of them, as shown in A.1.

I also find that the estimation strategy correctly recovers the true value of both amenities ($-\beta_k/\alpha$), whereas a hedonic regression of price on amenities greatly overestimates the value of the rarer amenity X_1 and underestimates the value of the more common amenity X_2, as illustrated in Table A.2.

Note that the result *does not* require $\beta_1 = \beta_2$. I can vary the parameter vector to different values and the estimation strategy will still work, as long as there are at least two characteristics.

Table A.1: *Regression on Mean Utility*

The *Dependent variable* is the mean utility δ_j. The price instrument is constructed using the market clearing prices assuming $\xi = 0$. The F-stat of the first stage is well above 10 (at 306.0). The OLS produces a biased estimate of the price coefficient, whereas the IV recovers it.

	(1) OLS	(2) First Stage	(3) IV
Price	-0.635***		-1.042***
	(0.0363)		(0.0589)
Price Instrument		0.960***	
		(0.0549)	
X1	0.498***	0.0549	1.057***
	(0.0499)	(0.0754)	(0.0810)
X2	0.715***	0.0316	1.033***
	(0.0283)	(0.0428)	(0.0460)
Inside Option	1.907***	0.120	3.125***
	(0.109)	(0.164)	(0.176)
N	200	200	200

Standard errors in parentheses

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

Table A.2: *Estimation of the Willingness to Pay for Amenities*

In this table, I compare the amenity value estimated by the model $(-\beta_k/\alpha)$ with the hedonic regression. In particular, even though the mean utility of the two characteristics X_1 and X_2 are identical, the relative scarcity of X_1 pushes the hedonics to produce a much greater coefficient.

	(1) Hedonic	(2) Sorting
X1	1.374***	1.015***
	(0.00139)	(0.0204)
X2	0.781***	0.991***
	(0.00122)	(0.0120)
Inside Option	2.994***	2.999***
	(0.00110)	(0.000760)
N	200	200

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3 More on STR Supply Estimation

In this section, I provide a number of relevant computational details about the estimation procedure used to estimate the supply system. Following Dubé, Fox and Su (2012), I formulate the GMM objective function as a mathematical program with equilibrium constraints (MPEC). With the problem appropriately defined, I provide the analytical Jacobian and Hessian used to estimate the parameters, where I highlight the sparsity features of the problem that make the computation reasonably efficient.

A.3.1 Problem Formulation

Without loss of generality, let q index a total of Q markets. Let $Z_q = [\, p_q^{IV} \;\; X_q^R \,]^T$ be a vector of exogenous shifters and $X_q = [\, p_q \;\; X_q^R \,]^T$ be the endogenous shifters. Let $Z = [Z_1, \ldots, Z_Q]^T$ and $X = [X_1, \ldots, X_Q]^T$ be the data matrix of size $Q \times (X+1)$. Let D_q denote the empirical distribution of the demographic characteristics in market q.

The requisite moment condition is

$$\mathbb{E}\left[(\delta_q - \theta_1^T X_q)^T Z_q\right] = 0 \tag{A.3.1}$$

To solve for the supply system coefficients, I formulate the problem as follows

$$\min_{(\delta,\theta_1,\theta_2,\eta)} \quad \eta^T W \eta \tag{A.3.2}$$

$$\text{s.t.} \quad \forall q : \; S_q(\delta_q, \theta_2; X_q, D_q) = S_q^o \tag{A.3.3}$$

$$\eta = Z'(\delta - \theta_1^T X) \tag{A.3.4}$$

As such, I denote the Lagrangian of the problem as follows

$$f_q(\delta_q, \theta_2; X_q, D_q) = S_q(\delta_q, \theta_2; X_q, D_q) - S_q^o \tag{A.3.5}$$

$$g(\delta, \theta_1; Z, X) = \eta - Z^T(\delta - \theta_1^T X) \tag{A.3.6}$$

$$G(\eta; W) = \eta^T W \eta \tag{A.3.7}$$

$$\mathcal{L}(\delta, \theta_1, \theta_2, \eta) = G(\eta; W) + \sum_q \lambda_f^q f_q(\delta, \theta_2; X_q) + \lambda_g^T g(\delta, \theta_1; Z, X) \tag{A.3.8}$$

Note that θ_1 denotes the linear coefficients of the model, whereas θ_2 denotes the non-linear coefficients of the model. Further, let $\theta_2 = [\pi_b^k]$ where $b = 1, \ldots, (X+1)$ indexing the product characteristics and $k = 1, \ldots, K$ indexing the demographic characteristics. As such, the heterogeneous component of the home sharing is

$$\lambda_{i,q} = \sum_b \big(\sum_k \pi_k^b z_{i,k}\big) X_q^b \tag{A.3.9}$$

where $z_{i,k} \sim P^*_{D_q}$ is drawn from the empirical distribution of the demographics in market q. With logit error, the market share S_q is thus computed as

$$S_q(\delta_q, \theta_2; X_q, D_q) = \frac{1}{N_q} \sum_q P_{i,q}(\delta_q, \theta_2; X_q, D_q) = \frac{1}{N_q} \sum_q \frac{\exp(\delta_q + \lambda_{i,q})}{1 + \exp(\delta_q + \lambda_{i,q})} \tag{A.3.10}$$

A.3.2 Analytical Derivations

Analytical Derivatives of MPEC The gradient of the objective function is

$$\nabla_{(\delta,\theta_1,\theta_2,\eta)} G(\eta; W) = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 2W\eta \end{bmatrix} \tag{A.3.11}$$

The Jacobian for the constraints is

$$\nabla_{(\delta,\theta_1,\theta_2,\eta)}(f,g) = \begin{bmatrix} \frac{\partial f}{\partial \delta} & 0 & \frac{\partial f}{\partial \theta_2} & 0 \\ \frac{\partial g}{\partial \delta} & \frac{\partial g}{\partial \theta_1} & 0 & I_g \end{bmatrix} \tag{A.3.12}$$

where

$$\frac{\partial f_q}{\partial \delta_q} = \frac{1}{N_q} \sum_i P_{i,q}(1 - P_{i,q}), \quad \frac{\partial f_q}{\partial \delta'_q} = 0, \quad q \neq q' \tag{A.3.13}$$

$$\frac{\partial f_q}{\partial \pi_k^b} = \frac{1}{N_q} \sum_i P_{i,q}(1 - P_{i,q}) z_{i,k} X_q^b \tag{A.3.14}$$

$$\frac{\partial g}{\partial \delta} = -Z^T, \quad \frac{\partial g}{\partial \theta_1} = Z^T X \tag{A.3.15}$$

Note that the upper-left $Q \times Q$ block of $\frac{\partial f}{\partial \delta}$ contains only diagonal terms $\frac{\partial f_q}{\partial \delta_q}$. When the number of markets Q is large, this results in a sparse Jacobian, which is particularly attractive computationally.

Analytical Hessians of MPEC Lagrangian Next, I derive the analytical Hessian of the MPEC Lagrangian.

$$\nabla^2 \mathcal{L}(\delta, \theta_1, \theta_2, \eta) = \begin{bmatrix} \frac{\partial^2 \mathcal{L}}{\partial \delta^2} & 0 & \frac{\partial^2 \mathcal{L}}{\partial \delta \partial \theta_2} & 0 \\ 0 & 0 & 0 & 0 \\ \frac{\partial^2 \mathcal{L}}{\partial \theta_2 \partial \delta} & 0 & \frac{\partial^2 \mathcal{L}}{\partial \theta_2^2} & 0 \\ 0 & 0 & 0 & \frac{\partial^2 \mathcal{L}}{\partial \eta^2} \end{bmatrix} \tag{A.3.16}$$

where

$$\frac{\partial^2 \mathcal{L}}{\partial \delta_q^2} = \lambda_f^q \frac{\partial f_q^2}{\partial \delta_q^2} = \lambda_f^q \frac{1}{N_q} P_{i,q}(1 - P_{i,q})(1 - 2P_{i,q}) \tag{A.3.17}$$

$$\frac{\partial^2 \mathcal{L}}{\partial \delta_q \delta_q'} = 0, \quad q \neq q' \tag{A.3.18}$$

$$\frac{\partial^2 \mathcal{L}}{\partial \delta_q \partial \pi_k^b} = \lambda_f^q \frac{\partial f_q^2}{\partial \delta_q \partial \pi_k^b} = \lambda_f^q \frac{1}{N_q} P_{i,q}(1 - P_{i,q})(1 - 2P_{i,q}) z_{i,k} X_q^b \tag{A.3.19}$$

$$\frac{\partial^2 \mathcal{L}}{\partial \pi_k^b \partial \pi_{k'}^{b'}} = \sum_q \lambda_f^q \frac{\partial f_q^2}{\partial \pi_k^b \partial \pi_{k'}^{b'}} = \sum_q \lambda_f^q \frac{1}{N_q} P_{i,q}(1 - P_{i,q})(1 - 2P_{i,q}) z_{i,k} X_q^b z_{i,k'} X_q^{b'} \tag{A.3.20}$$

$$\frac{\partial^2 \mathcal{L}}{\partial \eta^2} = 2W \tag{A.3.21}$$

Again, notice the upper-left $Q \times Q$ block representing $\frac{\partial^2 \mathcal{L}}{\partial \delta^2}$ has non-zero entries only on the diagonal term. Thus, even with a large Q, the Hessian remains sparse, making it computationally more tractable.

Appendix B

Appendix to Chapter 2

B.1 Simulation Details

$\beta_1 = 3, \beta_2 = 0.25, A = 500$. ν_i and u_i are drawn independently from Uniform$[-2.5, 2.5)$. d_{-i} is drawn from a Poisson distribution with $\lambda = 50 + 50 \times (1 - 0.2(\nu_i + u_i))$. The number of simulation draws $ns = 500$. The sample size for each draw $N = 500$.

Appendix C

Appendix to Chapter 3

C.1 Supplementary Tables and Figures

Table C.1: *Impact of the First Child's Gender*

This table reports result from regressing number of children on the gender of the first child. Each observation is an individual partner. The dependent variable is the number of children a partner has. The independent variable *First Child is Daughter* is a binary indicator on whether the partner's first child is a daughter.

	(1) Number of Children	(2) Number of Children	(3) Number of Children
First Child is Daughter	0.0568	0.0466	0.0421
	(0.0621)	(0.0627)	(0.0624)
Partner Age		0.0159***	0.0148***
		(0.00249)	(0.00249)
Female Partner			-0.372***
			(0.100)
Constant	2.370***	1.481***	1.579***
	(0.0366)	(0.143)	(0.145)
Observations	1310	1235	1235

Standard errors in parentheses
* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.2: *Robustness: Hiring Level Regression when Children Age Available*

The dependent variable is a binary indicator of whether a given hire is a woman. We use the children metrics for the existing partners the year before the hire. The sample is further restricted to hires where existing partners' children age information is available. Standard errors are clustered at venture capital firm and year level.

	(1) Female	(2) Female	(3) Female	(4) Female
Avg Daughters	0.0419**	0.0478**		
	(0.0185)	(0.0199)		
Avg Children	-0.0174	-0.0219*		
	(0.0114)	(0.0122)		
Avg Daughters (Senior)			0.0462**	0.0514**
			(0.0199)	(0.0211)
Avg Children (Senior)			-0.0209*	-0.0249**
			(0.0112)	(0.0123)
VC Firm Age		0.0000906		0.000248
		(0.00137)		(0.00143)
Avg Partner Age		0.000758		0.000847
		(0.00126)		(0.00125)
Partner Count		0.000622		0.000314
		(0.000920)		(0.000975)
Log(Capital Per Partner)		0.00239		0.00281
		(0.00684)		(0.00685)
Year	Yes	Yes	Yes	Yes
Observations	1493	1428	1493	1428

Standard errors in parentheses
* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.3: *Robustness: Hiring Level Regression Excluding Email Respondents*

The dependent variable is a binary indicator of whether a given hire is a woman. We use the children metrics for the existing partners the year before the hire. The sample is further restricted to hires where existing partners' children age information is not solicited via emails. Standard errors are clustered at venture capital firm and year level.

	(1) Female	(2) Female	(3) Female	(4) Female
Avg Daughters	0.0256	0.0280*		
	(0.0168)	(0.0170)		
Avg Children	-0.0111	-0.0128		
	(0.0105)	(0.0113)		
Avg Daughters (Senior)			0.0333**	0.0373**
			(0.0156)	(0.0158)
Avg Children (Senior)			-0.00508	-0.00590
			(0.0101)	(0.0109)
VC Firm Age		0.00130		0.00121
		(0.00138)		(0.00154)
Avg Partner Age		0.000368		0.0000712
		(0.00106)		(0.00111)
Partner Count		0.000274		0.000950
		(0.000966)		(0.00111)
Log(Capital Per Partner)		-0.00161		-0.00342
		(0.00726)		(0.00812)
Year	Yes	Yes	Yes	Yes
Observations	1263	1240	1190	1168

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.4: *Hiring Level Regression All Partners (Alternative Measures of Daughters)*

The dependent variable is a binary indicator of whether a given hire is a woman. *Avg Daughters* is the original measure, the average number of daughters at the firm. *Daughter Ratio* is defined as the ratio of total number of daughters to the number of children at the firm. *Average Daughter Ratio* is the average of the daughter-to-children ratio over active partners. *Daughter-Heavy Partner Fraction* is the fraction of partners with more daughters than sons, less those with fewer daughters than sons. *First Daughter Partner Fraction* is the fraction of partners at the firm whose first child is a daughter. *At Least One Daughter Fraction* is the fraction of partners who have at least on daughter at the firm. Standard errors are clustered at venture capital firm and year level.

	(1) Female	(2) Female	(3) Female	(4) Female	(5) Female	(6) Female
Avg Daughters	0.0439**					
	(0.0185)					
Daughter Ratio		0.0563*				
		(0.0341)				
Average Daughter Ratio			0.0496			
			(0.0326)			
Daughter-Heavy Partner Fraction				0.0358**		
				(0.0151)		
First Daughter Partner Fraction					0.0329	
					(0.0365)	
At Least One Daughter Fraction						-0.0140
						(0.0266)
Avg Children	-0.0202*	-0.000209	-0.000757	0.00236	-0.00262	-0.00132
	(0.0118)	(0.0111)	(0.0111)	(0.0104)	(0.0109)	(0.0113)
VC Firm Age	0.0000974	0.0000790	0.0000398	0.000137	0.0000214	-0.0000290
	(0.00131)	(0.00132)	(0.00131)	(0.00131)	(0.00132)	(0.00134)
Avg Partner Age	0.00115	0.00144	0.00145	0.000937	0.00154	0.00134
	(0.00115)	(0.00118)	(0.00118)	(0.00114)	(0.00118)	(0.00116)
Partner Count	0.000496	0.000618	0.000607	0.000555	0.000477	0.000225
	(0.000874)	(0.000882)	(0.000884)	(0.000871)	(0.000883)	(0.000877)
Log(Capital Per Partner)	-0.000678	-0.0000502	-0.0000546	-0.00124	-0.0000240	-0.000776
	(0.00640)	(0.00654)	(0.00652)	(0.00640)	(0.00652)	(0.00636)
Year	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1573	1533	1532	1573	1533	1573

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

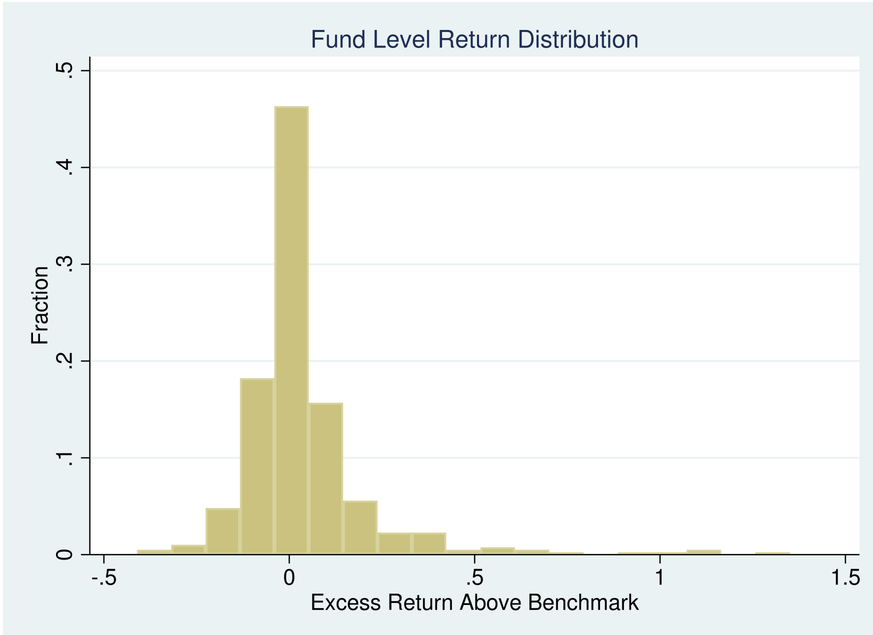

Figure C.1: *Fund Return Distribution*

Table C.5: *Robustness: Deal-Level Instrumental Variable Regression with IPOs Only*

This table reports regression result of deal success in deal-level sample using the average number of daughters as the instrument. The dependent variable*IPO* equals to 1 if the deal went public. *Female Hired Ratio* is the number of active female partners divided by the total number of active partners. In the instrumental variable regression, the instruments are the average number of existing partners' daughters when the hires (now active partners) were made. Standard errors are clustered at venture capital firm and year level.

	(1) IPO	(2) IPO	(3) IPO	(4) IPO	(5) IPO	(6) IPO
Female Hired Ratio	0.0272	0.0186	0.326	0.345	0.374*	0.402*
	(0.0305)	(0.0301)	(0.247)	(0.267)	(0.205)	(0.215)
Avg Children	-0.0160***	-0.0145***	-0.0108*	-0.00892		
	(0.00404)	(0.00407)	(0.00640)	(0.00676)		
Avg Children (Senior)					-0.00731	-0.00628
					(0.00491)	(0.00504)
VC Firm Age		0.00182**		0.00100		0.000869
		(0.000722)		(0.000998)		(0.000920)
Avg Partner Age		-0.000467		-0.000269		-0.000370
		(0.000498)		(0.000543)		(0.000563)
Partner Count		0.00149*		0.00170**		0.00183**
		(0.000786)		(0.000851)		(0.000836)
Log(Capital Per Partner)		0.00698*		0.00624		0.00594
		(0.00380)		(0.00403)		(0.00407)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Instrumented for Female Hired Ratio						
Average # Daughters	N/A	N/A	X	X		
Average # Daughters (Senior Partner)					X	X
First Stage F-stat			17.08	15.79	25.93	25.16
Observations	10435	10435	10435	10435	10435	10435

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.6: *Robustness: Deal-Level Instrumental Variable Regression with U.S. Deals*

This table reports regression result of deal success in deal level sample using the average number of daughters as the instrument. The dependent variable *IPO* equals to 1 if the deal went public. *Female Hired Ratio* is the number of active female partners divided by the total number of active partners. In the instrumental variable regression, the instruments are the average number of existing partners' daughters when the hires (now active partners) were made. The sample of deals are further restricted to portfolio companies located in the United States. Standard errors are clustered at venture capital firm and year level.

	(1)	(2)	(3)	(4)	(5)	(6)
	Success	Success	Success	Success	Success	Success
Female Hired Ratio	0.00153	-0.0151	0.839**	0.931**	0.888***	0.968***
	(0.0407)	(0.0401)	(0.352)	(0.403)	(0.294)	(0.322)
Avg Children	-0.0143**	-0.0108*	0.00157	0.00781		
	(0.00560)	(0.00584)	(0.00984)	(0.0114)		
Avg Children (Senior)					-0.00364	-0.000781
					(0.00769)	(0.00834)
VC Firm Age		0.00283***		-0.000467		-0.000555
		(0.00107)		(0.00194)		(0.00174)
Avg Partner Age		-0.000993		-0.000758		-0.000455
		(0.000731)		(0.000891)		(0.000896)
Partner Count		0.00237**		0.00343**		0.00326**
		(0.00114)		(0.00142)		(0.00136)
Log(Capital Per Partner)		0.00756		0.00243		0.00221
		(0.00537)		(0.00699)		(0.00693)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE, Round FE, Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Instrumented for Female Hired Ratio						
Average # Daughters	N/A	N/A	X	X		
Average # Daughters (Senior Partner)					X	X
First Stage F-stat			17.73	15.59	27.79	25.79
Observations	9346	9346	9346	9346	9346	9346

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.7: *Robustness: Impact on Career Outcomes (Deal Count)*

This table reports the Poisson regression result of deal count in the hiring sample. The dependent variable is deal count, defined as the number of deals a venture capitalist has been a board member on. The sample is further restricted to portfolio companies located in the United States and to venture capitalists with two or more deals. Standard errors are clustered at venture capital firm and year level.

	(1) Deal Count	(2) Deal Count	(3) Deal Count	(4) Deal Count
Deal Count				
Female VC	-0.0105	-0.0629	0.00747	-0.187
	(0.0826)	(0.293)	(0.0848)	(0.291)
Daughter Ratio	-0.0191	-0.0356		
	(0.110)	(0.118)		
Avg Children	-0.00736	-0.00724		
	(0.0410)	(0.0421)		
Female VC x Daughter Ratio		0.157		
		(0.300)		
Female VC x Avg Children		-0.0113		
		(0.108)		
Daughter Ratio (Sr)			-0.00414	-0.0193
			(0.109)	(0.116)
Avg Children (Sr)			-0.0395	-0.0473
			(0.0367)	(0.0390)
Female VC x Daughter Ratio (Sr)				0.106
				(0.307)
Female VC x Avg Children (Sr)				0.0720
				(0.105)
Partner Count	-0.0580***	-0.0582***	-0.0609***	-0.0613***
	(0.0160)	(0.0160)	(0.0157)	(0.0158)
VC Firm Age	0.0273***	0.0272***	0.0280***	0.0281***
	(0.00491)	(0.00491)	(0.00485)	(0.00483)
Avg Partner Age	0.00432	0.00426	0.00448	0.00449
	(0.00379)	(0.00379)	(0.00386)	(0.00385)
Log(Capital Per Partner)	-0.0708***	-0.0712***	-0.0701***	-0.0702***
	(0.0233)	(0.0232)	(0.0230)	(0.0230)
Constant	0.511**	0.529**	0.568**	0.567**
	(0.226)	(0.227)	(0.228)	(0.231)
Year, Industry FE	Yes	Yes	Yes	Yes
Observations	940	940	922	922

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.8: *Robustness: Impact on Career Outcomes (Partner Tenure)*

This table reports the regression result of partner tenure in the hiring sample. The dependent variable is tenure, defined as the number of years a partner has worked as a venture capitalist. The sample is further restricted to portfolio companies located in the United States and to venture capitalists with two or more deals. Standard errors are clustered at venture capital firm and year level.

	(1) Tenure	(2) Tenure	(3) Tenure	(4) Tenure
Female VC	-0.818	2.453	-0.589	2.737
	(0.712)	(2.156)	(0.707)	(2.059)
Daughter Ratio	0.562	0.839		
	(1.102)	(1.124)		
Avg Children	0.486	0.613		
	(0.383)	(0.403)		
Female VC x Daughter Ratio		-1.879		
		(2.594)		
Female VC x Avg Children		-1.173		
		(0.933)		
Daughter Ratio (Sr)			0.629	1.029
			(1.105)	(1.126)
Avg Children (Sr)			0.339	0.445
			(0.378)	(0.401)
Female VC x Daughter Ratio (Sr)				-3.259
				(2.639)
Female VC x Avg Children (Sr)				-0.858
				(0.894)
Partner Count	-0.552***	-0.547***	-0.561***	-0.554***
	(0.189)	(0.188)	(0.192)	(0.190)
VC Firm Age	0.312***	0.312***	0.282***	0.283***
	(0.0508)	(0.0506)	(0.0501)	(0.0499)
Avg Partner Age	0.147***	0.149***	0.184***	0.185***
	(0.0442)	(0.0443)	(0.0432)	(0.0432)
Log(Capital Per Partner)	0.0991	0.105	-0.0112	-0.00310
	(0.262)	(0.263)	(0.266)	(0.267)
Constant	11.24***	10.64***	10.49***	9.994**
	(3.884)	(3.894)	(3.919)	(3.932)
Year, Industry FE	Yes	Yes	Yes	Yes
Observations	940	940	922	922

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

Table C.9: *Robustness: Fund-Level Instrumental Variable Regression with U.S. Funds*

This table reports regression result of success in the fund level sample, restricting to funds targeting the United States only. The dependent variable is the excess return of the fund, defined as the net internal rate of return less the median fund benchmark. The median fund benchmark is defined as the median fund return in each region and year, as provided by Preqin. *Female Hired Ratio* is the number of active female partners divided by the total number of active partners. In the instrumental variable regression, the instruments are the average number of existing partners' daughters when the hires (now active partners) were made. Standard errors are clustered at venture capital firm and year level.

	(1) Excess Return	(2) Excess Return	(3) Excess Return	(4) Excess Return	(5) Excess Return	(6) Excess Return
Female Hired Ratio	0.0527	0.0610	0.916*	0.664*	1.022**	0.827*
	(0.0979)	(0.101)	(0.480)	(0.387)	(0.517)	(0.424)
Avg Children	-0.00384	0.00969	0.00735	0.0185		
	(0.0106)	(0.0123)	(0.0121)	(0.0132)		
Avg Children (Senior)					-0.00332	0.00339
					(0.0117)	(0.0118)
VC Firm Age		0.000660		0.00293		0.00344
		(0.00306)		(0.00303)		(0.00322)
Avg Partner Age		-0.00457**		-0.00560***		-0.00496**
		(0.00188)		(0.00203)		(0.00199)
Partner Count		0.00112		-0.00115		-0.00222
		(0.00181)		(0.00247)		(0.00277)
Log(Capital Per Partner)		-0.0130		-0.0103		-0.00771
		(0.0116)		(0.0145)		(0.0156)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Instrumented for Female Hired Ratio						
Average # Daughters	N/A	N/A	X	X		
Average # Daughters (Senior Partner)					X	X
First Stage F-stat			10.76	12.64	9.66	11.92
Observations	301	301	301	301	301	301

Standard errors in parentheses

* $p < 0.1$, ** $p < .05$, *** $p < .01$

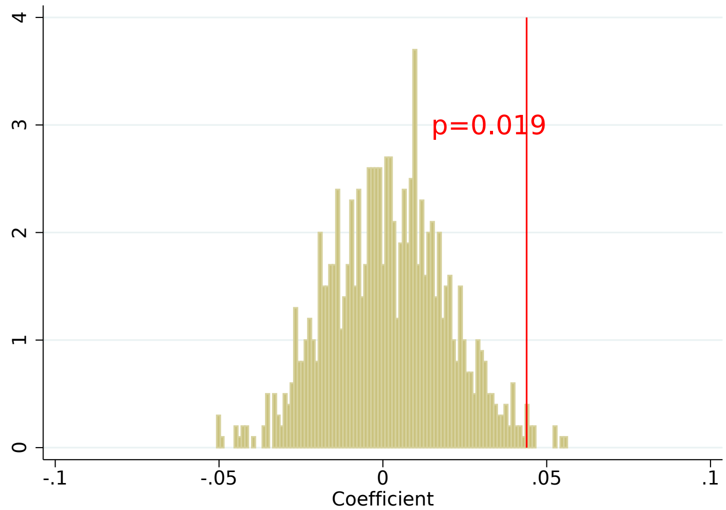

(a) Randomization test on the number of daughters by all existing partners. Specifically, we randomly assign the gender of the children in the data set of partners, holding the birth years and the total number of children same as the original data set. We regress the gender of the hire on the children and firm-level characteristics as before. The chart displays the true coefficient for *Avg Daughters* in the hiring regression, compared to the coefficient produced by the gender permutation simulated 1000 times.

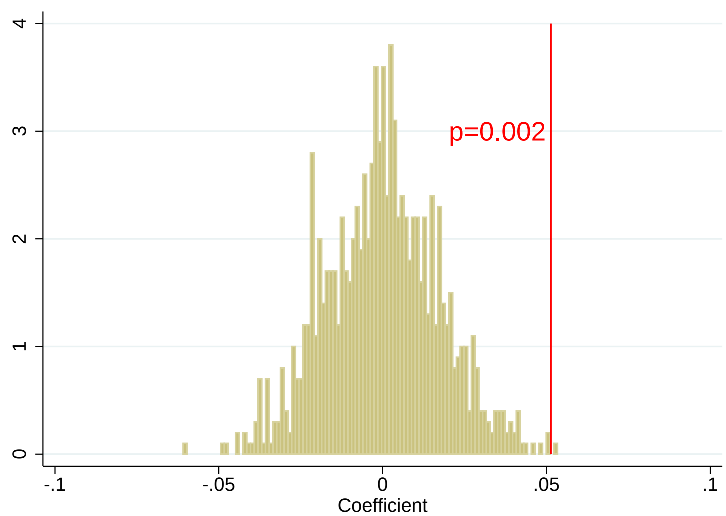

(b) Randomization test on the number of daughters by senior partners.

Figure C.2: *Randomization Inference: Hiring Level Regression*

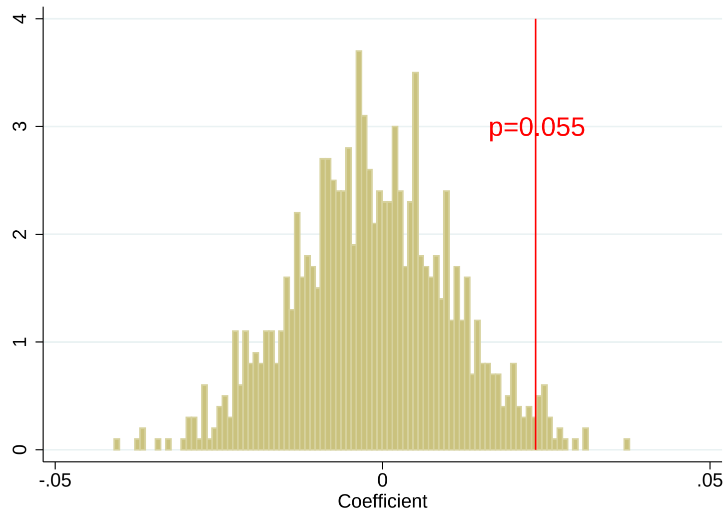

(a) Randomization test on the number of daughters by all existing partners. Specifically, we randomly assign the gender of the children in the data set of partners, holding the birth years and the total number of children same as the original data set. We regress the deal-level success on the children and firm-level characteristics as before. The chart displays the true coefficient for *Avg Daughters* in the deal-level reduced-form performance regression, compared to the coefficient produced by the gender permutation simulated 1000 times.

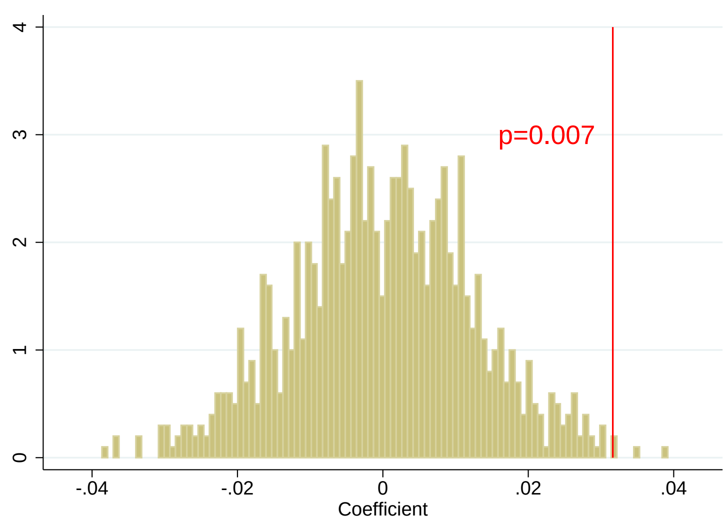

(b) Randomization test on the number of daughters by senior partners.

Figure C.3: *Randomization Inference: Deal-Level Performance Regression*