Essays on Housing Economics and Family Dynamics

by

Tereza Ranošová

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2023

Doctoral Committee:

Professor John Leahy., Co-Chair
Professor Ana Reynoso, Co-Chair
Professor Charles Brown
Professor Toni Whited

Tereza Ranošová
tranosov@umich.edu
ORCID ID 0009-0004-1929-7761

ACKNOWLEDGEMENTS

This dissertation would not be possible without the encouragement, help and support of a great number of people. First and foremost, I am very grateful to the members of my dissertation committee: John Leahy, Ana Reynoso, Charles Brown and Toni Whited, for their encouragement and valuable advice. John Leahy's introduction to macro theory provided a foundation to which I often come back. More importantly, his second-year macro class devoted to recent housing-related research in macroeconomics taught me what makes a good research paper while helping me realize that housing is a topic I would like to study. I am grateful he pushed me to be better and helped me to stay focused on the bigger picture, while allowing me to keep my interests broad. When I learned Ana Reynoso is joining the University of Michigan's Economics Department, my enthusiasm for research was reinvigorated. Her research is on the very forefront of studying the incentives within marriages, a topic that has always fascinated me. She quickly became a mentor to me and guided me through the process of finishing my degree. She continues to be an inspiration to me, both as a researcher and as a person, and I will always be thankful for her help. Charlie Brown is without exaggeration one of the kindest and most dedicated professors I have ever met. I am grateful for his detailed and insightful comments on my work. Finally, I extend my gratitude to Toni Whited. She provided careful and generous methodological support, but more importantly, her encouragement gave me the confidence not to shy away from technical challenges and to gain skills I will surely fall back on in my future research endeavors.

In addition to my committee, I am indebted to a number of current and former faculty and staff of the Department of Economics. Thank you to Martha Bailey, Mel Stevens, John Bound, Pablo Ottonelo, Andrés Blanco, Gabriel Ehrlich, David Lam, Christopher House, Matthew Shapiro, Dmitriy Stolyarov and Michael Mueller-Smith. The administrative office in the Economics Department as well as the staff at the International Center provided excellent support. I want to thank Laura Flak in particular, for making my transition to the United States as smooth as it could be and for always going the extra mile to support her students. My time at Michigan would hardly be the same without being embraced as a member of the RSQE forecasting team. They helped me improve my technical and writing skills, provided funding for my studies and gave me a taste of what economics can do outside of research. Thank you to Gabriel Ehrlich, Daniil Manaenkov, Jacob Burton, Owen Kay, Kyle Henson, Michael McWilliams and Donald Grimes. I am also grateful to the Panel Study of Income Dynamics Help Desk for granting me access to the confidential geolocated data and for being a great support throughout the project. Looking back, there is no way I would have made it this far without the help and support of my teachers and mentors in Prague at the Institute of Economics Studies as well as at CERGE-EI. In particular, I would like to thank Barbara Pertold-Gebická who have guided me through my first ever attempt at research and Filip Matějka who have encouraged me to apply to a PhD program abroad.

I feel so incredibly lucky to have been part of the most wonderful community of graduate students. Many of them provided invaluable feedback to my work during countless discussions and seminars. Many of them have also become my close friends, hopefully for life. In particular, I could not have done this without Stephanie Karol, John Olson, Theodor Kulcyczki and Hanna Onyshchenko, who have stood by me, encouraged me countless times and believed in me probably more than I believed in myself. For being a source of support and companionship through the ups and downs of a graduate degree, both professionally and personally, I would also like to thank Tyler Radler, Florian Trouvain, Hayk Hambardzumyan, Shwetha Raghuraman, Max Huppertz, Hayley Abourezk-Pinkstone, Keshav Garud, Russel and Emilie Morton, Brian Cevallos Fujiy, Nathan and Faith Mather, Caitlin Hegarty, Leticia Juarez, José-Ramón Morán Van Gelderen, Mike Mei, Luis Baldomero-Quintana, Anastasiia Chaikina, Michael Ricks, María Aristizábal-Ramírez, Zsigmond Pálvölgyi and Levi Turner. I am also grateful to my family and friends outside the department or back in Prague. First and foremost thank you to my sister, my father and especially my mother, Jarmila Ranošová, a trailblazer in our family who is an endless source of inspiration on that even from humble beginnings you can make it far, but it is also ok to take a step back an enjoy the ride. Special thank you also to our community of Czechs in Chicago, especially Vladimír Lizák and Fabio Gottlicher.

Without a doubt the person I am most grateful to is my husband Oskar Gottlieb, a brilliant mind, a kind friend and the most supportive life partner. I thank my husband for being my biggest cheerleader. As a researcher in the field of family economics I am constantly reminded that gender equality in career outcomes is definitively not possible without an equitable distribution of responsibilities within the household. Yet, this is still rare. I also thank my husband for being the exception.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation focuses on two of the biggest kinds of long-term decisions in one's life, on housing and on marriage. In all three chapters I focus on previously unexplored aggregate implications of simple differences between the incentives faced by different demographic groups. Chapter I shows that when dealing with long commutes in an urban housing market couples have a technological advantage over singles in being able to specialize. Therefore, sprawled metro areas with long commutes motivate couple formation, while postponing marriage contributes to gentrification. Chapter II explores how a preference of men to match with younger women naturally leads to divorce, and the more so when an unusually large cohort enters the marriage market. Chapter III shows that the reaction of housing markets to monetary policy depends on the age-distribution of the population.

In Chapter I "Commuting and the value of marriage" I study how and why long commutes are costly to different types of households (namely singles versus couples) and what are the implications for metropolitan housing and marriage markets. Over time, as metro-areas sprawled to the suburbs, long commutes became common. In this paper I combine motivating evidence with a structural model to show that even though long commutes are particularly detrimental to married women's labor market outcomes, in terms of welfare it is singles who lose the most. First, I show that the gender gap in commuting among singles is negligible. Second, men in couples (not women) have much longer commutes than single men, and job access alone cannot explain this difference. This together with other observations suggests that commuting features gains from specialization harnessed within couples, allowing men to take better jobs. I embed this feature in a quantitative spatial model with endogenous marriage and location choices that successfully captures the commuting and location patterns by marital status. In a joint housing and marriage market equilibrium, as metro areas sprawl, commuting increases most for men in couples and employment falls most for women in couples, contributing to gender gaps in both outcomes. However, in terms of welfare singles lose more than couples, increasing the value of marriage. Couples are able to partially evade commuting costs through specialization, lower housing costs and redistributing resources within the household.

Chapter II "Did the baby boom cause the US divorce boom?" connects two major

demographic 'booms' that the United States experienced during the second half of the twentieth century, in births after the second world war and in divorces 25 years later. This paper argues that the two booms are linked. As the baby-boom generations were entering marriageable age, men in previous cohorts were faced with exceptionally good remarriage prospects motivating them to rematch. The cohorts who ultimately divorced most were the ones with the biggest increase in remarriage opportunities for men. Using cross-state variation in the size of the baby-boom, I show that marriages in the pre-boom generations were more likely to divorce the bigger the relative supply of young women. This conclusion is robust to instrumenting the size of the baby-boom with WWII mobilization rates. Lastly, I construct a simple dynamic marriage market model which can generate a divorce boom caused by a baby-boom, and can account for between a seventh and a third of the rise in divorces in the 1970s.

In Chapter III "Housing market channel of monetary policy: the role of residents in their 50s" I provide empirical evidence that the age composition of the population matters for the transmission of monetary policy through the housing market. The response of housing demand to changes in interest rates is a key mechanism of monetary policy. I show that the effect of monetary policy shocks, identified through high-frequency event studies, on housing prices depends on the age-structure of the economy in a non-trivial way. Both across U.S. metro areas and across states, local housing prices drop more after monetary policy tightens whenever the share of population between 50 and 65 years of age is higher. If the share of population in a metro area 50-55 years old increases by one percentage point, a one standard deviation monetary policy shock depresses housing prices by an additional 0.413 percent after 3 quarters. A stronger investment motive in the demand for housing by this age group is a possible mechanism. This differential reaction of housing prices is already detectable by the quarter of the shock, and is followed by a differential response in employment starting about four quarters after the shock.

CHAPTER I

Commuting and the Value of Marriage

Abstract

Over time, as metro-areas sprawled to the suburbs, long commutes became common. In this paper I combine motivating evidence with a structural model to show that even though long commutes are particularly detrimental to married women's labor market outcomes, in terms of welfare it is singles who lose the most. First, I show that the gender gap in commuting among singles is negligible. Second, men in couples (not women) have much longer commutes than single men, and job access alone cannot explain this difference. This together with other observations suggests that commuting features gains from specialization harnessed within couples, allowing men to take better jobs. I embed this feature in a quantitative spatial model with endogenous marriage and location choices that successfully captures the commuting and location patterns by marital status. In a joint housing and marriage market equilibrium, as metro areas sprawl, commuting increases most for men in couples and employment falls most for women in couples, contributing to gender gaps in both outcomes. However, in terms of welfare singles lose more than couples, increasing the value of marriage. Couples are able to partially evade commuting costs through specialization, lower housing costs and redistributing resources within the household.

1.1 Introduction

Over the 20th century the geographic footprint of US metropolitan areas grew enormously. Figure 1.1 shows that the share of U.S. population living in the suburbs increased from 7 percent in 1910 to 50 percent in 2000. Figure 1.1 illustrates this point within the Panel Study of Income Dynamics (PSID), the primary dataset used in this paper. The distance from a residence to the city center increased from over 13 miles in 1970 to almost 19 miles by 2010, and so did the distance between residence and an average job in the metro area. In this paper I focus on an overlooked aspect of suburban long commutes: the differential impact on couples and singles, and men and women within couples, operationalized through a joint housing and marriage market equilibrium. I show that while long commutes are most detrimental to married women's labor market outcomes, this does not necessarily mean their welfare is also most affected.

A range of policies can affect commuting.[1] A long policy discussion about the pros and cons of suburban sprawl and long commutes (see Glaeser and Kahn (2004), Ewing and Hamidi (2015), Ehrlich et al. (2018) for reviews) focuses on the trade-off between productivity returns to agglomeration and costs of commuting.[2][3] The differential effect on the behavior and welfare of singles and couples, the nature of commuting costs within households, and the ways in which housing policy effects can be operationalized through a joint housing and marriage market, are overlooked aspects of this debate.

I start by observing that singles and couples differ markedly in their commuting and residential location decisions. I show that men in couples have much longer commutes than all others. Single men and women both single and in couples have similar commutes. I show that the choices of residential location (and thus job access) cannot explain this difference. Rather the commuting margin plays a role in within couple job taking behavior, increasing specialization on commuting as well as time. I then embed this feature in a joint urban spatial housing equilibrium and marriage market equilibrium model, and show that within this framework long potential commutes are most costly to singles, even though observable labor market outcomes of married women are the most affected. As a result, longer potential commutes actually incentivize couple formation.

The conclusion that long commutes decrease the welfare of singles more than that of married women might come as a surprise. First, there is now a robust body of evidence documenting that long commutes are a contributing factor to gender gaps in the labor market, particularly for married women.[4] For example Black et al. (2014) and Farre et al. (2020) provide descriptive and quasi-experimental evidence suggesting that in metropolitan areas

[1]Bento et al. (2005) discusses how variables that can be affected by policy, such as population density restrictions on new development, public transit supply, density of the road network and distribution of jobs, correlate with average commutes across the Unites States. Gyourko and Molloy (2015) reviews the literature on housing regulations that discourage density, and thus encourage urban sprawl.

[2]See Fu and Ross (2013),Yinger (2021), Boehm (2013), Harari (2020) for examples.

[3]Recently, the COVID pandemic also reignited the discussion on benefits of work-from-home options (Delventhal et al., 2022) and the interaction of working from home with time spent in home productions Leukhina and Yu (2022).

[4]Despite considerable convergence over the past century, women participate in the labor market less than men, and when they do, they work shorter hours and earn lower wages (for reviews and overall trends see Blau and Kahn (2013), Blau and Kahn (2007), Blau and Kahn (2017),Petrongolo and Ronchi (2020)).

with long commutes women tend to work less. Several early papers document there are substantial gender gaps in commuting, with men commuting more than women (Madden (1981), White (1986), Turner and Niemeier (1997), Tkocz and Kristemen (1994)). Moreover, a streak of recent papers documents that women have a lower willingness to trade off a long commute for a higher wage, contributing to gender wage (and other labor market outcomes) gaps (Rosenthal and Strange (2012), Gutierrez (2018), Liu and Su (2020), Barbanchon et al. (2020), Caldwell and Danieli (2023), Borghorst et al. (2021), Moreno-Maldonado (2022)). With the exception of Gutierrez (2018), the gender gaps in commuting are left unexplained or interpreted as a difference in preferences.[5] While welfare implications are rarely discussed in this literature, it is often implicit that the increased gender gaps in labor market outcomes are undesirable for women and that policies implying long commutes are worse for women than men. This would be true if gender gaps in commuting were caused by particular distaste for commutes among women. Such a mechanism is, however, not supported by the range of empirical evidence I provide in this paper.

Second, as metro-areas sprawl while jobs are concentrated in the city, suburban neighborhoods lose more access than central ones. At the same time, couples are more likely to live in the suburbs. Thus geographically sprawled areas hurt job access of couples more than that of singles. Plus long distances exacerbate the collocation issue couples are facing. Yet, in this paper I show that even though sprawl does hurt couple's job access more and the labor market outcomes of women in couples more, it is overall singles who loose most in terms of welfare.

To make welfare conclusions about the impact of long commutes, it is necessary to know more about the underlying motivations for residential and job choices, as they are both endogenous. To this end, I collect a range of motivating evidence about commuting and location choices of singles and couples. The most important and novel observation is that the gender gap in commuting arises solely through men dramatically increasing their commuting after they form couples. Specifically, using a geolocated PSID sample, I first show that men substantially increase their commuting after forming a couple, that this is not true for women and that there is essentially no difference in commuting between single men and single women. Thus, gender gaps in commuting cannot be a result of gendered preferences. Second, while it is true that couples are more likely to live in the suburbs than singles, this difference is not big enough to explain the gap in commuting between single men and men in couples. Specifically, I show that the gap in commuting reduces only marginally and remains large

[5]Interestingly, both Barbanchon et al. (2020) and Liu and Su (2020) state that there is heterogeneity in the gender-gap by marital status, with married women being least willing to trade off higher wages for longer commutes.

and statistically significant after controlling for various measures of how much residential location is suburban or how much job access it has. Theoretically, couples could achieve short commutes of wives by systematically prioritizing her job access when choosing where to live. However, I find no evidence of this mechanism. Combining the main sample with the geographic distribution of jobs from the LEHD Origin-Destination Employment Statistics (LODES), and assigning individuals to their respective labor markets based on their most common lifetime industry and earnings segment, I show that in fact, couples locate weakly closer to the kind of jobs the husband typically works in. Alternatively, a wife might commute less than her husband, because even if most of her potential jobs are far away, she searches for a local alternative or drops out of the labor force if no convenient jobs are available. I show evidence consistent with this second mechanism. Within couples, actual commutes of husbands are more correlated with distances to potential jobs (i.e. potential commutes) then are actual commutes of wives. On the other hand, labor market participation and hours of wives are more negatively correlated with potential commutes. Thus, overall when jobs in the husband's labor market are further away from the couple's residence, husbands simply commute more. When wives' jobs are further away, they are more likely to work locally, reduce hours, or not work at all. Lastly, I confirm that couples and singles also choose different residential locations within a metro-area. Couples systematically live further away from the city center, and consequently further away from jobs.

Motivated by this evidence, I construct and estimate a quantitative urban spatial housing market and marriage market equilibrium model.[6] Singles and couples choose a residential neighborhood within a metropolitan area, accept or reject job offers and choose how to allocate their time. I overlay this structure with a simple marriage market equilibrium. The difference between singles and couples is a crucial feature of the model motivated by the empirical evidence that both commuting and residential location differ substantially by relationship status. Nevertheless, modeling this heterogeneity is very rare in quantitative urban economics. Most closely related to this paper is Tscharaktschiew and Hirte (2010), who construct a quantitative spatial equilibrium with couples and singles choosing a location. However, their paper has implications that do not square with the evidence presented here (for example singles flocking to the suburbs, and higher wage earners within couples commuting less). I model the choices of couples explicitly as a collective decision resulting from bargaining between two partners with potentially conflicting interests, as in Browning et al. (2014).[7] This

[6]The spatial equilibrium portion is standard, based on a discrete choice of location as in McFadden (1977), Redding and Rossi-Hansberg (2017), Ahlfeldt et al. (2015) and many others.

[7]Thus, I relate to the literature on gender differences in labor market outcomes within the context of household specialization. See Gronau (1977), Chiappori et al. (2002), Cherchye et al. (2012), Blundell et al. (2016), Bertrand et al. (2015), Bianchi et al. (2000)).

again is a methodological contribution, as modeling household bargaining over residential location is rare in quantitative urban economics. With the exception of Chiappori et al. (2018a), who show that ignoring the bargaining process within couples in urban models results in biased measures of value of time, the urban economics literature typically relies on a 'unitary' representation of the household.[8][9] Lastly, I endogenize the decision to form a couple and the required within-couple distribution of resources. To the best of my knowledge this is the first paper constructing and estimating a quantitative spatial equilibrium model of a metropolitan area with a combined housing and marriage market equilibrium, showing how the effects of a housing policy can be operationalized through a joint equilibrium outcome.[10]

Several possible mechanisms could explain gender gaps in commuting within couples. However, none of the mechanisms common in the literature can also explain the gap between single and coupled men.[11] I propose that the observed patterns can be rationalized if commuting imposes costs on households in a form that rewards specialization – when one spouse takes a local job or stays at home, the other is freed to work far away, accepting better jobs. I propose a simple parametrization of a household-level cost of commuting that features gains from specialization and show that it allows the model to match the observed patterns of commuting and residential location. The cost captures the intuition that households value if someone is close by, to deal with emergencies, accept packages or pick up children from school. However, one person per household is quite enough, and there is no added benefit when two people are working close to home at the same time. I estimate the model with a moment based procedure, targeting moments summarizing the distribution of people and

[8]Taking the potential conflict between wives' and husbands' location priorities seriously is more common in papers studying cross-metropolitan-area mobility. (Costa and Kahn, 2000) suggest that two-career couples locate in a large metro areas to solve the collocation problem whose career to prioritize. Several papers (Compton and Pollak (2007), Gemici (2008), Chauvin (2018), Venator (2020)) show cross-metro mobility is typically associated with labor market improvements for the husband, and losses for the wife.

[9]Even among unitary representations of the household, those models that consider explicit specialization by gender are not quantitative. Black et al. (2014), Abe (2011) present theoretical illustrative models where a fixed cost of commuting increases labor force participation gaps by gender. Madden (1977) and Gutierrez (2018) present theoretical spatial models where commuting returns increasing with hours (though higher wages) can explain why women in couples commute less. None of these confront their quantitative predictions with the data.

[10]Moreno-Maldonado (2022) constructs a quantitative model of choosing location across metro areas and labor supply, where women's labor supply declines in large cities due to higher commuting costs. Fan and Zou (2021) present a pioneering model of location choice across metro-areas, with joint local marriage and labor market clearing.

[11]For example, Gutierrez (2018) studies only couples and shows that theoretically gender differences in commuting can arise because returns to commuting scale with hours, but commuting itself is a fixed cost with respect to hours. As husbands work longer hours than wives, they are more willing to commute. While I incorporate this mechanism in my model, I argue it is not the principal driving force behind the observed commuting differentials, because it does not explain why men in couples commute so much more than single men.

jobs, labor market behavior of couples and singles, and residential location and commuting behavior patterns presented in the empirical section.

Within this framework, I show how longer potential commutes benefit couples and encourage more marriage while simultaneously increasing gender gaps in labor market outcomes. Couples in metro-areas with long commutes become more specialized, with one member (typically the wife) staying home or taking a local job and spending more time in home production. This allows husbands to accept high-value jobs without worrying about commuting. As suburbs are less and less convenient in terms of jobs access, housing rents in the suburbs fall compared to the city. Because singles lack the technological advantage of being able to specialize, they are more incentivized to flock to the city and overpay on housing. While wives lose by not being able to keep jobs they like, within a marriage market equilibrium they are compensated with more leisure. Thus, when the housing and marriage market equilibrium re-clear, marriages end up being more valuable for both men and women and welfare falls most for those who are single.[12]

In the next section I describe the commuting and residential patterns of singles versus couples, as well as the evidence that men and women in couples react differently to long potential commutes. In section three I discuss the model, its structure and estimation. Section four presents the results of a counter-factual simulation, changing the urban landscape towards more sprawl that requires longer commutes. Finally, section five compares aspects of the counter-factual simulation with variation across U.S. metro areas, providing further validation for the model mechanisms.

1.2 Commuting and residential location of couples and singles

1.2.1 Data and Measurement

The primary data source for this section is the geocoded restricted version of the Panel Study of Income Dynamics (PSID), with residential location data available up to the level of a Census tract. With this information I assign each response to a 2010-defined metropolitan statistical area (MSA) and compute an euclidean distance between the (2010 population weighted) centroid of the tract of residence and the centroid of the largest Census place within the MSA (distance to center d_c). In addition, the PSID includes four variables allowing

[12] In this paper, I abstract from divorce for the sake of simplicity. If the ability to specialize on the commuting margin adds to the value of marriage, it should also lower the probability of divorce. On the flip side, if an individual (more often the wife) is working close to home, allowing their partner to accept longer commutes, increasing their commute can be especially costly to the couple. Recent evidence by Hrehova et al. (2021) shows that when a commute is increased by the business relocating, the worker is more likely to divorce later.

Figure 1.1: Suburban sprawl and commuting

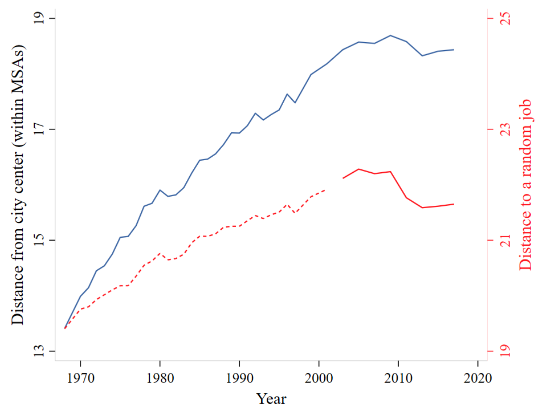

Miles. Sample: PSID 18-50, normalized to white men, 35 years old, in couples. Job distribution: LODES (U.S.CensusBureau, 2021), imputed before 2002.

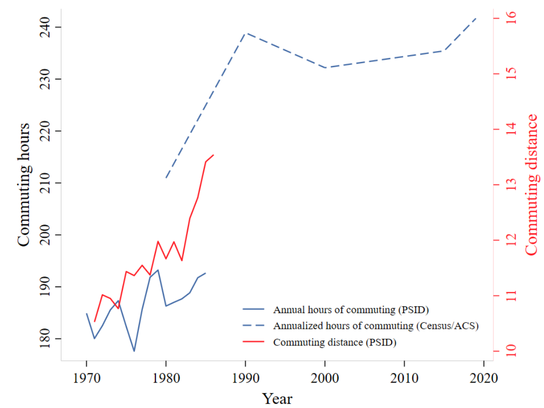

Sample: 18-50, normalized to white men, 35 years old, in couples. Sources: PSID, Census IPUMS 1980-2000, ACS IPUMS 5-year 2010, 2015, 2019.

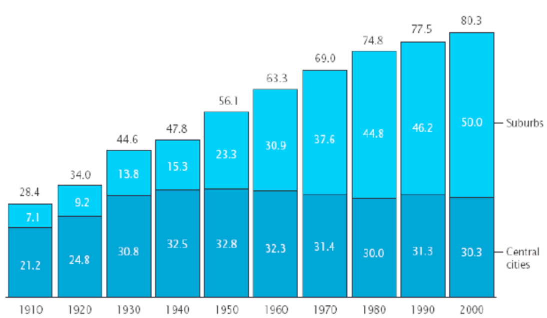

Percent of total population living in a metropolitan area: central cities versus suburbs. Source: (Hobbs and Stoops, 2000)

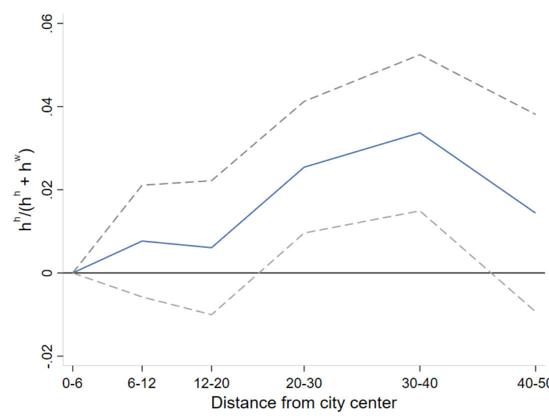

$\frac{H^h}{H^h+H^w}_{t,i} = \sum_{j=2}^{N} \alpha_i^{\text{dist bin } j} + \alpha_i^{ah} + \alpha_i^{aw} + \alpha_i^{educh} + \alpha_i^{educw} + \alpha_t + \alpha_i^{raceh} + \alpha_i^{\#children} + \epsilon_{i,t}$. Source: PSID sample of couples (as defined bellow). The solid line plots the difference between couples living 0-6 miles from the center and the rest of the locations in the share of household market hours performed by husband. The dotted line shows the 95% confidence intervals.

me to study commuting. First, in waves 1971-1986 the PSID includes a typical commuting distance in miles for the head and the wife (with 1971-1974 and 1977 only asking the head of the household). This is the primary commuting variable in the analysis, labeled d. Second, in waves 1970-1981 and 1983-1986 the PSID includes annualized hours of commuting for the head and the wife (with 1973, 1974 and 1977 only asking the head of the household).[13] Third, in 2011-2017 both the head and the wife are asked about typical duration of a one way commute (I annualize this report assuming each person works 5 days a week). Lastly, in 2013-2017 the geocoded restricted version includes the census tract of the current job. After restricting only to people whose job is in the same metro area as their residence (thus avoiding distances that are unreasonable to be an actual daily commute), I construct a 'distance to job' measure by computing the euclidian distance between the centroid of the tract of residence and the tract of work. In almost all waves these variables are only asked of people who worked over the last year. I use the alternative commuting measures to confirm robustness of the main results to alternative definitions and time periods.

To study the labor market behavior I use annual hours of work and labor income. To measure time in home production I use annual hours of housework.[14] To study behavior before and after forming a couple I construct tenure within a couple by assigning the first observed year of cohabitation in the PSID as the year a person stopped being single. For couples that are already observed in the first wave in 1968 I use the year of marriage, whenever available, to represent the year the couple was formed. Throughout this section, single is used to describe people in the PSID who have not been observed in a couple before.[15]

To study the distribution of jobs within metro areas I utilize the publicly available counts of jobs in a census bloc (counting jobs that are part of the unemployment insurance reporting system) per industry (19 categories) and earnings segment (3 categories) provided by the Census Bureau as part of the LEHD Origin-Destination Employment Statistics (LODES) available in 2002-2017 (U.S.CensusBureau, 2021).[16] To extend the sample size I than backfill the job distribution from 2002 to all pre-2002 waves of the PSID. I compute a matrix of

[13]The timing of this variable is somewhat convoluted - combining a typical commute of the current job with the work-schedule of last calendar year. I keep the timing tied to the year of the wave, as the first component is more relevant.

[14]According to the documentation, this variable should only include hours of housework, not childcare. However, Gayle et al. (2015) and others use this measure to be a combination of plain housework and childcare. Namely, Gayle et al. (2015) show that subtracting typical PSID housework hours of singles from PSID housework hours of women in couples results in a measure of childcare that matches well with childcare hours reported in ATUS.

[15]Most importantly, singles do not include divorcees.

[16]2-digit industry categories and 3 earnings segments is the level of differentiation available in the LODES data (U.S.CensusBureau, 2021), which I use to construct distribution of jobs in metro-areas. The segments are separated by monthly earnings at or bellow \$1250, between \$1250 and \$3333, and above \$3333.

distances from each tract to each tract within all metro areas. Then, by weighting distances by the number of jobs I can compute 1) a distance to an average job in an MSA for each census tract for each year and 2) a distance to an average job in each industry-segment combination in an MSA for each census tract for each year (with pre-2002 years using the 2002 distribution). For each individual in the PSID that works at least in one wave when industry classification is available I select their most common industry and most common earnings segment (normalized to 2002 dollars for pre-2002 waves) and label this their 'labor market'. The associated distance to jobs in their labor market is interpreted as the distance to other potential jobs the individual would be a good fit for, a 'potential commute' or 'distance to opportunities' (labeled d_o, measured in miles). Distance to an average job across all labor markets is labeled d_j.

In the analysis bellow, I restrict the sample to individuals 18-50 years old who live in a metro area of at least 250 thousand residents per the 2010 Census. Moreover, for each individual I select their most common metro area over their observed lifetime in the PSID and exclude periods when this individual did not live in this MSA, so that all location changes are within the same area. Lastly, I only use single people who have not been in a couple before and couples for whom this is their first match, as far as it can be determined in the PSID.[17]

1.2.2 Commuting of couples and singles

This section presents a set of descriptive facts about how commuting behavior changes when men and women move from being single to forming a couple.[18] Tables 1.1 and 1.2 show results of regressing commuting outcomes on an indicator of whether an individual is in a couple (this can be a marriage or a cohabitation, to the extent it can be identified within PSID), metro-area, age and time fixed effects and additional demographic controls (with i standing for an individual and t for the wave of PSID). The analysis is done separately for men and women

$$d_{it} = \beta \cdot \text{In couple}_{it} + \alpha_t + \alpha_{age} + \alpha_{msa} + X_i + \epsilon_{it} \tag{1.2.1}$$

Consistently, men in couples have considerably longer commutes and spend more time commuting than single men.[19] However, for women there is very little difference between

[17]Table A.2 presents summary statistics for the sample starting in 1969, the first year geographic information is available, and since 1990, a subsample used large parts of the analysis.

[18]Since commuting is only defined for people with a job, all analysis in this subsection is done using a subsample of working individuals.

[19]I study both commuting time and commuting distance and treat them as providing information about the same behavior. Table A.5 in the appendix shows that this pattern holds in the cross-section using more recent variables in the PSID – typical commuting time (available in waves 2011, 2013, 2015 and 2017) and distance to work (available in waves 2013, 2015 and 2017).

Table 1.1: Commuting differences between singles and individuals in couples

	Commuting distance (miles)									
	Men					Women				
Singles (mean)			8.900					8.495		
In couple	2.708	2.555	2.369	2.388	2.238	.297	-.036	-.023	-.086	-.106
	(.674)	(.638)	(.662)	(.633)	(.660)	(.658)	(.646)	(.632)	(.663)	(.628)
d_o		.207					.117			
		(.062)					(.027)			
d_c			.206					.108		
			(0.032)					(.023)		
d_o *bins**				x					x	
d_c *bins**					x					x
N	24299	23243	24299	23243	24299	13641	13238	13641	13238	13641
N clusters	155	153	155	153	155	144	142	144	142	144

SEs statistics in parentheses. SEs clustered at the MSA level.
All regressions include year, age, MSA fixed effects, education and race dummies and cohort of birth.
The sample includes only individuals that are observed in a couple at some point.
* Includes dummies for d_x in intervals of 0-6, 6-12, 12-20, 20-30, 30-40 and over 40 miles.

couples and singles. The differences are large in scale compared to the baseline. Single men commute about 9 miles on average. Men in couples commute 20 to 30 percent more.[20] This is not a result of selection into being in a couple, as the sample excludes singles whom I never observe forming a couple later on.

A potential explanation for why men in couples commute more than single men is that couples typically move to the suburbs, thus further away from jobs. In section 1.2.4 I show that couples do indeed live further away from the city centers (more often in the suburbs) than singles. However, in tables 1.1 and 1.2 I show that this difference in residential location cannot account for the observed commuting differences. Specifically, I show that the commuting gap between single and married men reduces only marginally and remains large and statistically significant after controlling for various measures of how much their residential location is suburban (adding additional controls to equation 1.2.1). In column 3, I include the distance from residential location to the city center d_c as a control. While living further away from the city correlates with longer commutes, the gap between single and coupled men remains well above 2 miles. Column 2 presents the result when d_o, the

[20]The raw mean of commuting distance in the PSID sample is 10.6 miles with a standard deviation of 11.3 miles.

Table 1.2: Commuting differences between singles and individuals in couples

	Commuting time (annual)									
	Men					Women				
In couple	35.253	36.270	33.893	34.010	31.860	-24.027	-24.022	-23.415	-24.138	-23.680
	(8.528)	(8.291)	(8.735)	(8.112)	(8.469)	(6.742)	(7.147)	(7.187)	(7.288)	(7.393)
d_o		.715					-.135			
		(.515)					(.435)			
d_c			.862					-.241		
			(0.294)					(.337)		
d_o *bins**				x					x	
d_c *bins**					x					x
N	24181	22993	24181	22993	24181	15003	14475	15003	14475	15003
N clusters	154	152	154	152	154	147	144	147	144	147

SEs statistics in parentheses. SEs clustered at the MSA level.

All regressions include year, age, MSA fixed effects, education and race dummies and cohort of birth.

The sample includes only individuals that are observed in a couple at some point.

* Includes dummies for d_x in intervals of 0-6, 6-12, 12-20, 20-30, 30-40 and over 40 miles.

potential commute constructed with LODES data on jobs distributions, is added as a control instead. This accounts more directly for the job access lost when living in the suburbs. While a longer potential commute is associated with a longer actual commute, the gap between single and coupled men is only marginally affected. Columns 4 and 5 repeat the exercise, including instead dummies for several bins of d_o or d_c, showing the results are not driven by the linearity of the specification.[21] Again, for women there is almost no change in commuting before and after forming a couple, with or without controlling for where they live.

Next, I use within-person variation to show how commuting of men and women evolves before entering a couple through spending 5, 10 and more than 15 years in the couple. Figures 1.2 plot coefficients $\beta_{(-10)}$,..., β_{15} from the following regression, where α_i stands for a person fixed effect. β_5 measured the difference in commuting between those who are 5-9 years in a

[21]Similarly, the results are also robust to including polynomials of d_o or d_c instead.

relationship compared to the baseline of between 1 and 5 years before forming a couple.

$$\begin{aligned} d_{it} = {} & \beta_{(-10)} \cdot (\text{More than 5 years before forming a couple})_{it} + \beta_0 \cdot (\text{0-4 after forming a couple})_{it} \\ & + \beta_5 \cdot (\text{In couple for 5-9 years})_{it} + \beta_{10} \cdot (\text{In couple for 10-14 years})_{it} \\ & + \beta_{15} \cdot (\text{In couple for 15 and more years})_{it} \\ & + \alpha_t + \alpha_a + \alpha_g + \alpha_i + \epsilon_{it} \end{aligned} \tag{1.2.2}$$

The results mimic the cross-sectional comparison. For men, commuting distance increases after at least 5 years in a relationship to a level 2-3 miles higher than the commute of single men 5-1 years before they enter a relationship and flattens after.[22] The pattern is analogous for commuting time.[23]

The picture for women, however, is starkly different. In the cross-section, women in couples spent fewer hours a year commuting (likely confounding working fewer days with commuting shorter daily distances, as the time measure is annual). Using person fixed effects, and as such comparing women who worked both before and after forming a couple, I see that there is essentially no effect of forming a couple on commuting.

A second key observation is that this stark gender difference in commuting behavior that emerges within couples is not present among single people. Table 1.3 shows that across a variety of measures of commuting, gender differences are stark in couples, but are negligible among singles.[24] This observations disqualifies differential distaste for commuting by gender as the primary driver of gender gaps in commuting. Since men and women behave similarly as singles, but starkly different within couples, it has to be a dynamic of within household optimization that explains gender differences in commuting.

[22]Notice, though data is limited for this subsample, there is no evidence of a pre-trend, as commuting is actually higher for men 10-6 years before entering a couple than for men 5-1 years before settling with a partner.

[23]Tables A.4 and A.3 in the appendix repeat the analysis in tables 1.2 and 1.1 with person fixed effects. Qualitatively, the patterns are robust to comparing explicitly men and women before and after they form a couple. Men commute more, while women do not change their commutes. Quantitatively, the differences are smaller. This is not surprising – given the limited number of PSID waves that offer commuting information, there is only a limited number of observations that have commuting information both available before and after forming a couple. Those that do are observed in only very fresh couples. The pattern in 1.2 shows that commuting gaps take about 5-10 years to materialize.

[24]Table A.6 in the appendix confirms this pattern in the 2000 Census data.

Figure 1.2: Event studies of commuting with respect to forming a couple

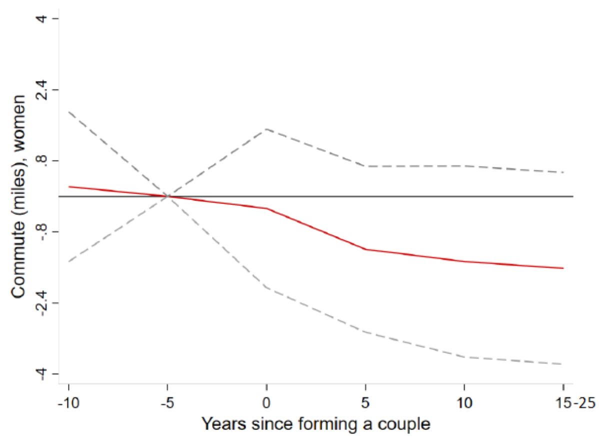

women

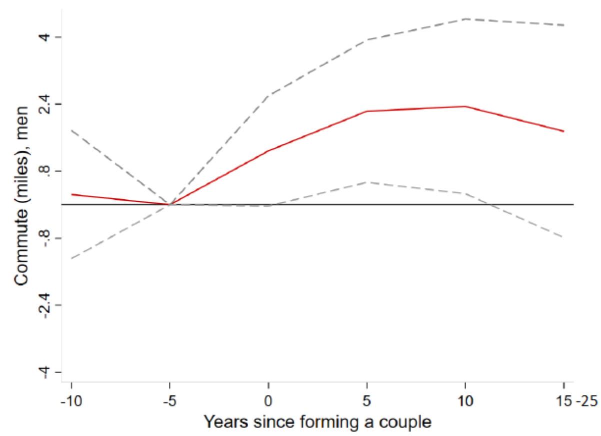

men

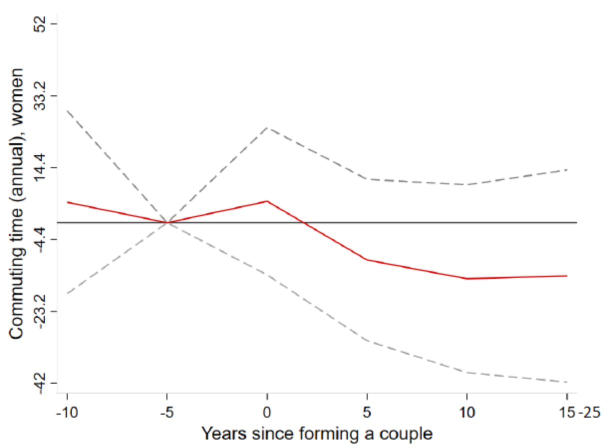

women

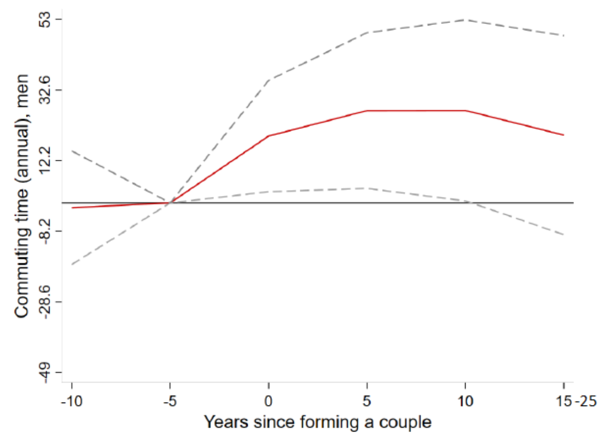

men

Source: PSID. Plotting coefficients $\beta_{(-10)}$, $\beta_{(-5)}$, β_5, β_{10}, β_{15} and the respective 95% confidence intervals from fixed effects regressions of the form 1.2.2 with the category ”in couple for 5 or fewer years” excluded and normalized to 0. Outcomes: commuting distance in miles (one way) and annual hours spent commuting. Notice these regressions include person fixed effects, therefore they are identified from differences in commuting over lifetime as a person moves from being single to living with a partner and from living with a partner for a short versus a longer time, after regressing out age effects. Sample: commuting distance is available in waves 1975-1976, 1978-1986 plus in 1969-1974, 1977 for heads of households only; commuting time is available in waves 1969-1972,1975-1976, 1978-1986 plus in 1973-1074, 1977 for heads of households only.

Table 1.3: Commuting differences between men and women when single and when in couples

	Commuting distance (typical)	Commuting time (annual)	Commuting time (annualized)	Distance to work (tract to tract)
All (mean)	10.646	173.274	183.402	9.039
Man	**.008**	**-4.656**	**15.550**	**-.773**
	(.697)	(9.220)	(8.791)	(.563)
In couple	-.637	-31.045	-7.600	-.252
	(.542)	(5.931)	(6.383)	(.420)
Man in couple	3.867	74.676	24.197	2.630
	(.775)	(10.060)	(10.582)	(.728)
N	25078	26942	9189	4843
N clusters	145	150	165	148
In couple at some point	x	x		
	1970-1986		2011-2017	2013-2017

All regressions include year, age, MSA, education, race, cohort controls. *SEs* statistics in parentheses. SEs clustered at the MSA level.

1.2.3 Potential commutes and labor market attachment within couples

The previous section shows that within couples there is a gender gap in commuting. Two sets of mechanical explanations are possible. First, it could be that couples locate close to the chosen job location of the wife, more than that of the husband. Second, when a couple forms, women with long potential commutes drop out of the labor market or switch to local jobs, while men keep their jobs with long commutes or switch to potentially better jobs even further away.[25] In other words, either couples chose their residential location closer to the wife's jobs, or men and women in couples differ in how they accept jobs given their residential location.

To discriminate between these two proximate causes, I take advantage of the data on the distribution of jobs in a metro-area. I compute the difference between husband and wife in their potential commutes (defined as the distance to potential jobs in their typical industry and earnings segment $d_{o,it}$), by running the following regression $d_{o,it} = \beta \cdot \text{Man}_i + \alpha_{couple} + X_i + \epsilon_{it}$, where α_{couple} stands for couple fixed effects and β measures the within-couple gender gap in job access. If couples on average systematically prioritized job access of wives, their residential

[25]This is consistent with the dramatic drop in labor market attachment of women after forming a couple, as illustrated in figure A.1, compared to men, who actually slightly increase their labor market attachment after forming a couple).Both figures A.1 are based on regressions with person fixed effects, ruling out any possible explanation of the selection of working men and non-working women into coupling.

location would be on average closer to the wife's type of job. Table 1.4 provides evidence against the residential location channel. There is no statistically significant difference in potential commute within couples. If anything, husbands have weakly shorter potential commutes.

Table 1.4: Difference between men and women within couples in potential commutes

	d^o			
Man	-0.039 (.080)	-0.064 (.090)	-0.118 (.074)	-0.109 (.080)
X_i:				
Industry+segment fixed effects		x		x
Education, race, cohort		x		x
age, year		x		x
N	47482	47130	96412	96023
Sample	$\geq$ 1990		$\geq$ 1969	

SEs in parentheses, clustered at the MSA level.
All regressions include couple fixed effects.

$d_{o,it} = \beta \cdot \text{Man}_i + \alpha_{couple} + X_i + \epsilon_{it}$ where β measures the difference within couples between men and women in their distance to an average job in their assigned industry and earnings segment.

The lack of a gender gap in potential commutes suggests that the gender difference in commuting within couples happens because husbands and wives take jobs differently. Next, I provide more direct evidence of this mechanism. Similar to Gutierrez (2018) I study variation within heterosexual couples, thus comparing men and women living in the same location. Consider the following regression (where i stands for an individual, c stands for a couple, a stands for age and t stands for time).

$$\text{comm}_{it} = \beta_d \cdot d_{oit} + \beta_{wd} \cdot d_{o,it}\text{woman}_i + \beta_w \cdot \text{woman}_i + \alpha_c + \alpha_a + \alpha_t + \alpha_{ind,seg} + \epsilon_{it} \quad (1.2.3)$$

The left-hand side variable is one of the measures of commuting available in the PSID. Table 1.5 shows the results for commuting distance, annual hours spent commuting, annualized usual time spent commuting and distance to work (euclidean distance. tract to tract). The first row presents the estimate for β_d, showing that for all measures of commuting being the one whose potential jobs are further away from the place of residence is associated with longer commutes for men. This is reassuring as it validates that the chosen measure of access to potential jobs correlates strongly with actual commutes. The second row (estimates of

Table 1.5: Actual commutes and potential commutes within couples

	Commuting distance (miles)		Commuting time (annual)		Commuting time (annualized)		Distance to work (tract to tract)	
Distance to jobs (d^{opp})	0.614	0.706	4.241	5.348	4.614	6.268	0.306	0.570
	(0.117)	(0.107)	(1.033)	(0.971)	(0.984)	(1.135)	(0.137)	(0.167)
$d^{opp}\cdot$ Woman	-0.098	-0.106	-1.581	-1.478	-1.681	-1.689	-0.128	-0.121
	(0.047)	(0.050)	(0.612)	(0.649)	(0.324)	(0.378)	(0.046)	(0.032)
Woman	-1.453	-0.134	-35.342	-21.611	-5.412	8.89	0.940	1.760
	(0.896)	(0.981)	(10.400)	(11.847)	(7.930)	(8.697)	(0.780)	(0.650)
X_i:								
'Labor market' fes		x		x		x		x
Couple fes	x	x	x	x	x	x	x	x
N	19836		21244		8824		3350	
N clusters	146		145		159		131	

SEs in parentheses. SEs clustered at the MSA level.
All regressions include year, age, MSA fixed effects.

Sample: all waves when a selected commuting variable is available. For commuting distance in miles and annual commuting time this requires using distribution of jobs from (mostly) 2002 backfilled to the 1970s. Annualized typical commuting time and distance to work use the actual distribution of jobs in the respective wave (2011-2017).

β_{wd}) shows the main coefficient of interest. The association between actual commutes and potential commute is not symmetric by gender – it is weaker for women. Whenever a couple lives further away from the wife's job opportunities, her commute increase less than it would for her husband.

Notice that by including couple fixed effects (α_c) I rely on variation in differences between husband and wife for couples where both of them work and they each work in a different kind of job (industry and/or earnings segment). As a byproduct, I am, by definition, only comparing people who live in the same location, with the location of the job determining commuting. This is important as it eliminates potential differences among people in how one's residential location is convenient for job access in general. Moreover, in columns 2, 4, 6 and 8 I include fixed effects for industry and earnings segment interactions. This way I am netting out systematic gender differences in working in generally more or less accessible industries. Overall, there is a strong pattern in couples of women's commutes being less associated with distance to opportunities than men's.

Next I repeat the analysis with labor market behavior on the left hand side. Table 1.6 presents the results. In the first three columns I see that within couples, for men their distance to opportunities is associated with lower hours and a lower probability of employment. This

Table 1.6: Work attachment and potential commutes within couples

	Hours	Working	Hours (positive)	Housework hours	log wage
Distance	-5.462	-0.002	-4.935	-1.064	0.00143
to jobs (d^{opp})	(2.221)	(0.001)	(1.911)	(1.236)	(0.00118)
d^{opp}. Woman	-5.146	-0.002	-2.484	3.285	-0.00118
	2.356)	(0.001)	(1.722)	(0.884)	(0.00062)
Woman	-557.429	-0.0762	-433.625	416.334	-0.109
	(44.939)	(0.014)	(31.969)	(17.319)	(0.020)
X_i:					
'Labor market' fes	x	x	x	x	x
Couple fes	x	x	x	x	x
Both working			x		x
N	59872		49120	58892	47918
N clusters	177		177	177	177

SEs in parentheses. SEs clustered at the MSA level.
All regressions include year, age, MSA fixed effects.

All results are based on waves 1990-2017 to avoid excessive backfilling of the jobs distribution. Table A.10 shows analogous analysis of hours spent working over samples when commuting variables are available, providing a more direct link to table 1.5.

suggests that a long potential commute disincentivizes work, either because commuting takes out of the time endowment or is costly for other reasons, leaving less time for work, or because jobs that are further away from industry centers are less desirable to spend time in. The second row shows that this association is again not gender-neutral: it is stronger for women, the opposite pattern to what I observe in commuting. Column 4 shows that the distance to potential jobs correlates negatively with hours of housework for men, but it is positively associated for women. The last column shows that long potential commutes are weakly associated with higher wages. However, this is less true for wives (though this result is only marginally significant).

Overall, a clear pattern emerges. Men and women in couples do not react symmetrically to poor job access. While husbands go for desired jobs even when they are far away and spent a long time commuting, wives tend to take a more local job, cut their hours or drop out altogether, spending more time on housework, potentially also settling for a lower paying job. This again suggests that couples behave as if husband's commuting was less costly to them than that of wives.[26]

[26] A potential shortcoming of this analysis is that commuting variables are available only in selected waves, and for commuting distance in miles and annual commuting time the jobs distribution has to be imputed from the first available datapoint, typically from 2002. Table A.10 in the appendix repeats the analysis of hours, cutting the sample to only waves when a respective commuting variable is available. While the

1.2.4 Residential location of couples and singles

In this section I show that couples and singles also choose differently when picking a residential location within the metro-area. Mimicking the analysis of commuting, table 1.7 shows the results of regressing distance of the census tract of residence to the center of the MSA in miles (d^c_{it}) on an indicator of whether an individual is in a couple (this can be a marriage or a cohabitation, to the extent it can be identified within PSID), metro-area, age and time fixed effects and additional controls (with i standing for an individual and t for the wave of PSID).[27]

$$d^c_{it} = \beta \cdot \text{In couple}_{it} + \alpha_t + \alpha_{age} + \alpha_{msa} + X_i + \epsilon_{it} \tag{1.2.4}$$

Table 1.7: Distance to the city: couples versus singles

	Distance to center		Distance to center < 10 miles		Distance to jobs		Distance to jobs in own industry/ segment	
In couple	1.830 (.338)	.966 (.197)	-.071 (.0132)	-.047 (.0097)	1.399 (.312)	.369 (.183)	1.754 (.322)	.436 (.192)
X_i:								
Education, race, cohort	x		x		x		x	
Person fes		x		x		x		x
Sample:					$\geq$ 1990		$\geq$ 1990	
N	160549	209337	160549	209337	108662	105099	89873	88970
N clusters	181	181	181	181	183	183	183	183

SEs statistics in parentheses. SEs clustered at the MSA level.
All regressions include year, age, MSA fixed effects.

Columns 1,3,5 and 7 only use people in couples or singles who are later observed in a couple.

In the cross-section, after controlling for age, education and race dummies, couples live on average almost 2 miles further away from the center than singles. The third column shows that for singles the probability of living less than 10 miles from the center is about 7 percent

raw association between distance to opportunities and hours is not robust to using only older waves when commuting information was available, the gender difference is. When using only recent samples that include annualized commuting time and distance to work, the gender difference is not significant. This is likely because the sample is substantially smaller, lacking enough variation within couples in their industry-segment combination. When I extend it moderately to include waves from 2000 onward, the result reemerges and is quantitatively similar to using older waves.

[27] Unlike for commuting, analysis in this section is not excluding people who drop out of the labor force.

Figure 1.3: Event studies of distance to the city with respect to forming a couple

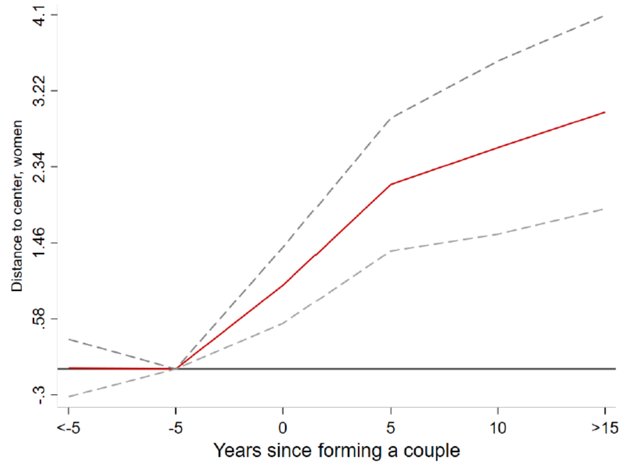

women

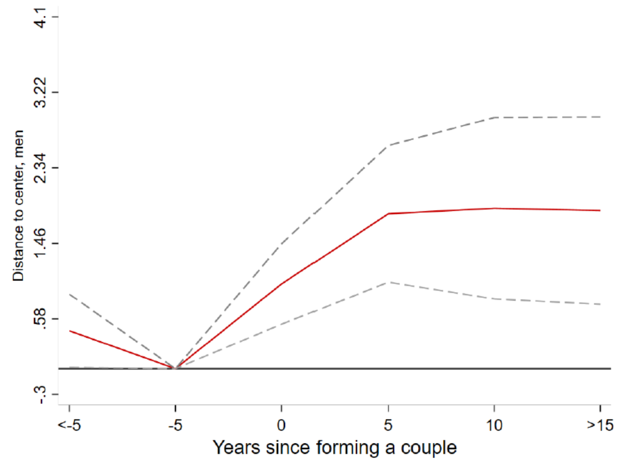

men

Plotting coefficients β_{-10}, β_0, β_5, β_{10}, β_{15} and the respective 95% confidence intervals from fixed effects regressions of the form: $d^c_{it} = \beta_{-10} \cdot$ In couple in more than 5 years$_{it}$ $+ \beta_0 \cdot$ In couple less than 5 years$_{it}$ $+ \beta_5 \cdot$ In couple for 5-10 years$_{it}$ $+ \beta_{10} \cdot$ In couple for 10-15 years$_{it}$ $+ \beta_{15} \cdot$ In couple for more than 15 years$_{it}$ $+ \alpha_t + \alpha_a + \alpha_g + \alpha_i + \epsilon_{it}$ with the category "5-1 year before forming a couple" excluded and normalized to 0.

higher. Columns 2 and 4 show that selection of singles to the city center is robust to looking strictly at the panel variation, after including person fixed effects. Figures 1.3 show that the pattern of moving to the suburbs is comparable for men and women and stabilizes after about 5 years of cohabitation. In metropolitan areas, jobs are typically more concentrated than people. When couples move to the suburbs, they are also moving further away from jobs. This is illustrated in columns 4-8 in table 1.7. People in couples live further away from an average job and further away from an average job in their most typical industry and earnings segment. Figures 1.4 show this pattern with the within-person variation. Unlike with commuting, the move to the suburbs is very similar for men and women.

To summarize, I show the following facts. First, on average men commute more then women. Second, this gap arises purely because men increase their commutes substantially after they form couples. This gender gap is not present among singles and men in couples commute much more than single men, while there is little difference for women. Third, couples are more likely to live in the suburbs and further away from jobs than singles. However, this difference is not the main driver of the commuting gap between single men and men in couples. In fact, the commuting gap is very robust to controlling for aspects of residential location that measure distance to the city or jobs. Fourth, there is no evidence that couples locate further away from the husband's potential jobs. This suggests that gender gaps in commuting within couples are not facilitated by the choice of residential location that

Figure 1.4: Event studies of job access with respect to forming a couple

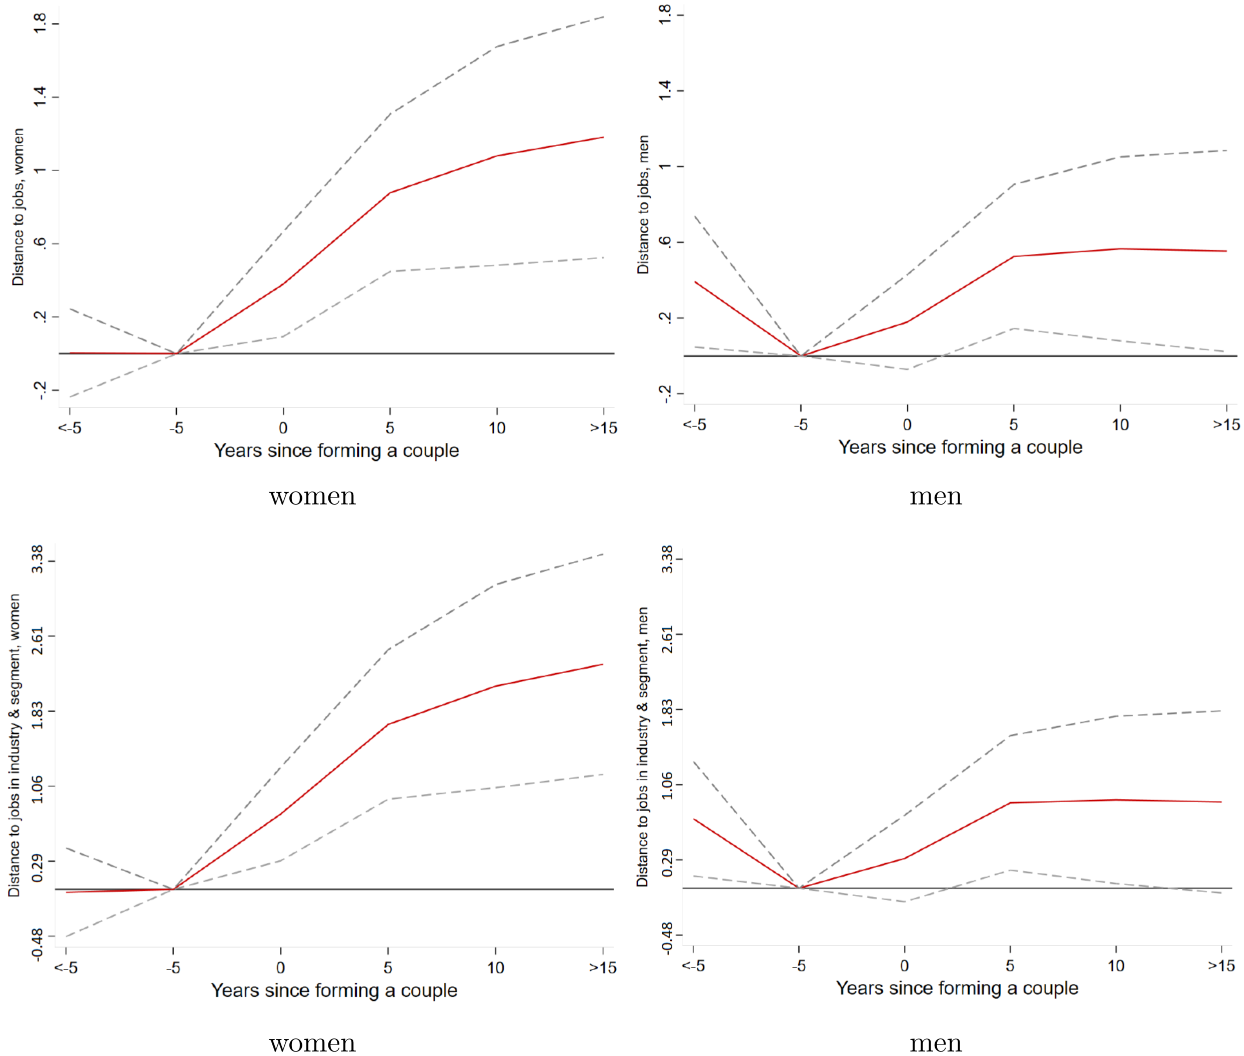

Analogous regressions for figures 1.3. With distance to center d^c replaced with the average distance to jobs in the metropolitan area of residence d_j, and the average distance to such jobs restricted to the individuals most common industry and earnings segment d_o. Only data after 1990 are used, to not backfill job location information by more than 12 years. The results, however, are very similar when a shorter or longer samples are used.

prioritizes job access for women. Fifth, I show direct suggestive evidence that when faced with long potential commutes, couples are willing to accept them for men while women in couples opt for a local job, shorter hours or drop out altogether. Overall, this set of facts suggests that there is something about commuting that motivates couples to specialize on this margin. Furthermore, this gendered specialization markedly changes the behavior of men. Specialization is allowing men in couples to behave as if commuting is less costly for them than for everyone else, including singles. Through specialization, men in couples can accept potentially better jobs with longer commutes.

1.3 Model

In this section I present a structural spatial equilibrium and marriage market equilibrium model of a metro-area capable of replicating the salient features of commuting and location decisions as presented above. Because gender differences in wages and productivity of time in home production do not lend themselves naturally to explain difference in commuting between single men and men in couples, I posit that working close to home matters to households beyond the time lost commuting in a way that rewards specialization. I introduce a convenient functional form that matches commuting patterns in the data. This cost of commuting is capturing the intuition that it is convenient if at least one member of the household works close by, but one is quite enough. As a result, couples have a technological advantage over singles in being able to specialize in commuting, with one of them working close by and freeing the other one to accept jobs far away from their the residence.

In its basic structure the model is a spatial equilibrium of a single metro-area with fixed housing supply per neighborhood, where residential rents are clearing the markets for housing and a bargaining weight clears the marriage market. Crucially, agents in the model are differentiated by gender and relationship status. This is directly motivated by the data, which show that men and women behave very differently depending on whether they are single or in a couple. To capture the transition from singlehood to forming a couple I use a simple overlapping generations structure. The population consists of three overlapping generations, one of singles and two of potentially living in a couple (1 period represents roughly 10-15 years). Moreover, I include a simple marriage market equilibrium, to endogenize the share of the population who is married and the within-couple distribution of resources.

With the focus on the individual decision of households differentiated by relationship status the rest of the model is kept rudimentary. Matching into couples happens at random. The spatial structure is simple. There is a city and suburbs to capture the typical degree of centralization of economic activity and the differences in access to jobs between singles and couples. The suburbs are further differentiated into two locations, one offering more opportunities to men and one to women. This is necessary to capture the potential for disagreement within couples about whose career to prioritize and the degree to which the collocation issue itself can be responsible for long commutes in couples.[28] Thus, the metro-area is composed of three equally-sized neighborhoods organized in a triangle (see figure 1.5).

The labor market is structured as a distribution of offers that individuals can accept or

[28]There is no robust difference between men and women in how much the kind of jobs they typically work in are offered in the city versus the suburbs. As a result, having only two locations would not be appropriate to capture the potential for disagreement between men and women in couples about where to locate within a metro-area.

reject coming from two distinct sectors. The purpose of differentiated jobs is to capture the potential trade-offs between short commutes and better monetary and non-monetary benefits from working. The purpose of two sectors (one with more men and one with more women) is to capture the potential for disagreement couples are facing in whose job access to prioritize. Each individual is exogenously assigned to a sector T and draws offers only from that sector. Each job is a bundle of a location (j), a wage (w) and a utility match shock (ξ) representing non-monetary benefits, where both w and ξ are higher for jobs that are close to other jobs in the same sector and ξ scales with hours (to mimick that monetary benefits scale with hours as well). Job characteristics are exogenous; there is no firm decision or labor market clearing. To allow for the idea that location choices are shaped by job access, households learn the location j of a job offer before they choose where to live. In the data, however, differences in commuting (between men and women in couples and between couples and singles) are not explained away by differential job access, but by differences in how jobs are accepted. Specifically, there is no evidence that within couples short commutes of women are facilitated by couples locating close to the wife's opportunities. Thus, one job offer would not provide enough flexibility in shaping commutes to capture the patterns in the data. Therefore, I allow a share π of the population to always get an additional option to work locally. This flexibility gives households more agency to affect their final commutes.

Figure 1.5: Spatial structure of the model metro-area

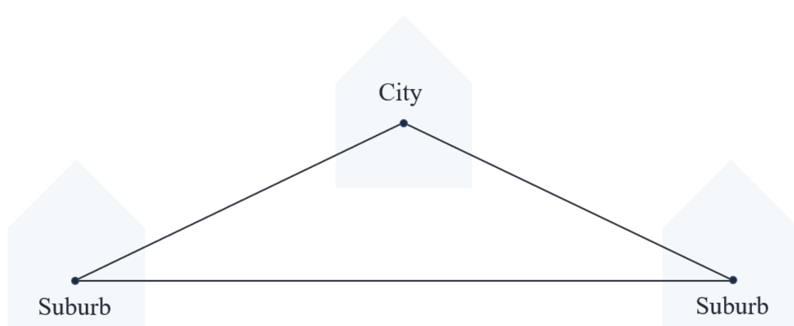

I fit the model to match location, commuting and work patterns in metropolitan areas in the United States, study how these patterns change if potential commutes increase and what the implications are for the welfare of singles and couples.

1.3.1 Household choices and timing

There are two main types of households choosing where to live, work and how to spend their time – singles and couples. Each individual goes trough three life stages. Everybody is single in the fist stage. After the first life stage, a person decides whether to marry and stay married forever or whether to stay perpetually single.

Figure 1.6: Model timing: lifecycle

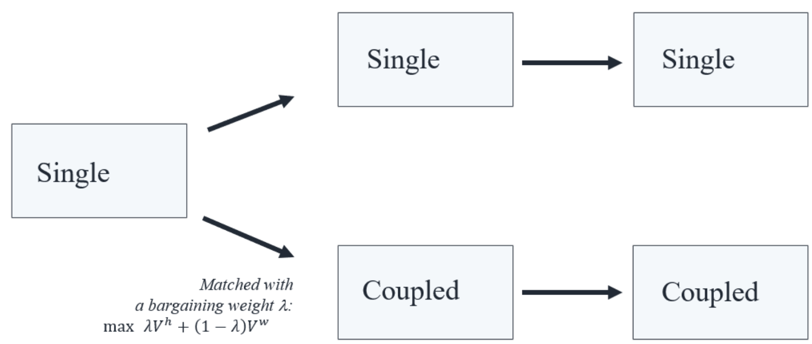

Couples differ from singles on three dimensions. First, couples have a more complex optimization problem. Whereas individuals maximize their own utility, couples maximize a weighted average of the utilities of the husband and the wife, where both live by definition in the same location. Second, couples derive more value from time spent on home production of a public good. This difference is capturing the fact that couples are more likely to bring up children in their households.[29] Third, singles and couples have potentially different preferences over location amenities, with couples appreciating the suburbs more (potentially for better schools and an overall good environment for children).

Within a life stage, decisions are made sequentially. Figure 1.7 presents the timeline. Each period starts with everybody drawing a job and learning its location. With this information in hand (and taking residential rents as given) households choose where to live. Next, a share of households π learns about local jobs (including the match shock) and decides whether to take it. Whoever is left without a job learns the match shock of their original offer and decides whether to take it or drop out of the labor force. After jobs are assigned, all households make decisions on time use, consumption and housing quantity demanded.[30]

[29] See figures A.2 in the appendix.

[30] The sequential nature of these decisions is posited for simplicity. The order of receiving a local offer first or second does not meaningfully affect results.

Figure 1.7: Model timing: within each stage of life

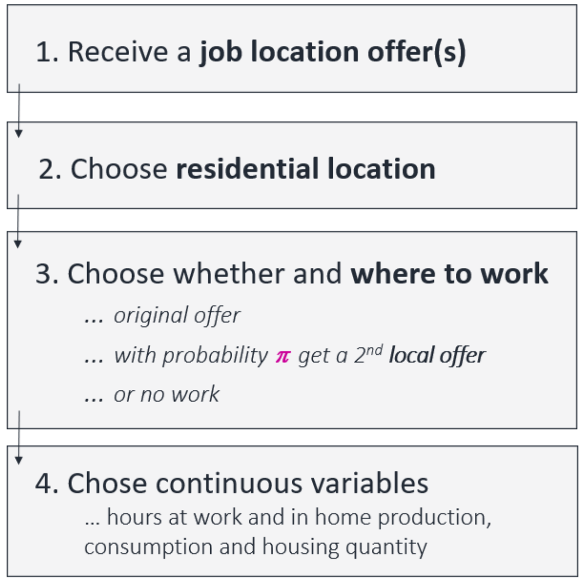

Equation 1.3.1 presents the optimization problem of a single person, after they have settled to a location i with rent $R(i)$ and a job $(j_*, w(j_*, T), \xi(j_*, h, \xi_0, T))$ (a non-monetary benefit ξ has an idiosyncratic random component ξ_0 and both monetary and non-monetary benefits to jobs scale with hours and are higher in industry centers, thus depending on both location and what sector one belongs to). Notice the model abstracts from borrowing and lending, so each period a household spends all their income. As a result, the decision problem is static.

$$
\begin{aligned}
U^s(i, j_*, \xi_0, T) = \max_{c^s, l^s, h^s, H^s, x^s} \; & u_c(c^s) + u_l(l^s) + u_H(H) + a^s(i) + \Pi^s + \xi(j_*, h^s, \xi_0, T) \qquad (1.3.1)\\
\text{s.t. } c &= h^s \cdot w(j_*, T) - R(i)H^s \\
1 &= h^s + \beta \cdot d_{i,j_*} + l^s + x^s \\
h^s &\leq \bar{h} \\
\Pi^s &= -F^s(d_{i,j_*}) + u_x^s(x^s)
\end{aligned}
$$

Utility is derived from c consumption, l leisure, H housing quantity, $a^s(i)$ amenity derived from living in a location i, $\xi(.)$ a transitory non-monetary benefit of a job, and Π^s the value of home production and additional household costs of commuting. Time is constrained to sum up to 1: $1 = h + b \cdot d + x + l$ (where h stands for time spend at work, x for home production and $b \cdot d$ for commuting). Commuting distance is given by d_{i,j_*}, where $d_{(.,.)}$ is the matrix of distances between neighborhoods and $d_{i,0} = 0 \forall i$ (0 is indicating not having a job). The value generated at home depends on time put in as well as on commuting: $\Pi^s = F^s(d) + u_x^s(x)$, where F is decreasing in d and u_x is increasing in x. The effect of commuting on household value generated at home reflects the costs of commuting beyond time use – the option value of being near home in the case of emergencies in the household (such as accepting packages and letting in maintenance personnel). Section 1.3.3 includes a detailed discussion of this parametrization choice. There is an upper bound on work hours, reflecting that the labor market typically does not allow for complete flexibility in the kind of employment contracts offered.

Equation 1.3.2 presents the optimization problem solved by a single person when choosing a job, given their residential location i and a job location offer j. $A_j \in \{0, 1\}$ indicates whether a person accepts job location j.

$$\begin{aligned} V^s(i,j) &= \pi \cdot E_{\xi^i}\Big[\max_{A_i}\ E_{\xi^j}\big[\max_{A_j} v(A_i, A_j|\xi^i, \xi^j)\big]\Big] \\ &\quad + (1-\pi)\cdot E_{\xi^j}\Big[\max_{A_j} v(0, A_j|0, \xi^j)\Big] \\ v(A_i, A_j|\xi^i, \xi^j) &= U^s(i, j_*, \xi_0) \\ \text{where } j_* &= A_i \cdot i + (1-A_i)A_j \cdot j \\ \xi_0 &= A_i \cdot \xi^i + (1-A_i)A_j \cdot \xi^j \end{aligned} \tag{1.3.2}$$

A single person knows the location of their potential job j. A share of the population learns about a local job and decides whether to take it. This would give a short commute which is weighed against the potential of a better match. All who did not take the local job accept the offer in j or drop out of the labor force. j_* is the final optimally chosen job location (either i, j or 0). Figure 1.8 summarizes the job-taking decision. Settling into jobs, each single person decides on hours of leisure, work and home production, on consumption and quantity of housing. At the end of the period job ties are severed.

A couple acts to maximize a weighted sum of the husband's and the wife's utility, with λ representing the bargaining power of the husband. The maximization problem of a couple within each period, given residential location i and jobs $(j_*^g, w(), \xi^g())$ for $g \in \{h, w\}$ is

Figure 1.8: Choice of where and whether to work

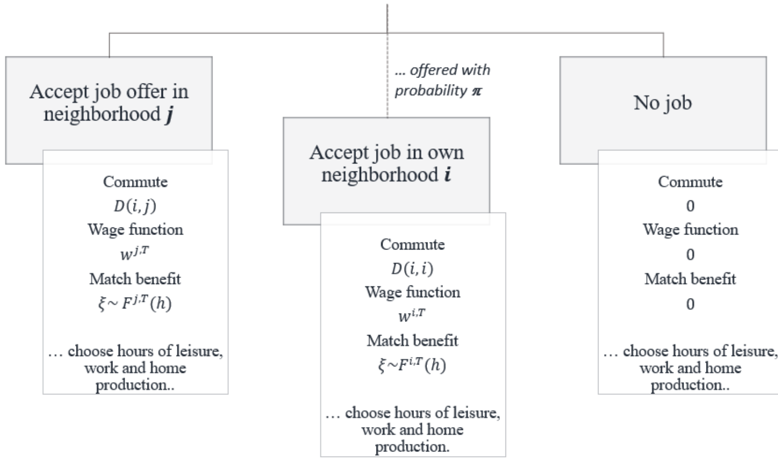

presented in equation 1.3.3.

$$
\begin{aligned}
U(i, j_*^h, j_*^w, \xi_0^h, \xi_0^w, T^h, T^w) = \max_{c^h, l^h, h^h, x^h, c^w, l^w, h^w, x^w, H} \Big[\qquad\qquad (1.3.3) \\
\lambda\big[u_c(c^h) + u_l(l^h) + \xi(j_*^h, h^h, \xi_0^h, T^h)\big] \\
+ (1-\lambda)\big[u_c(c^w) + u_l(l^w) + \xi(j_*^w, h^w, \xi_0^w, T^w)\big] \\
+ u_H(H/2) + a^c(i) + \Pi\Big] \\
\text{s.t. } c^h + c^w = h^h \cdot w^h(j_*^h, T^h, h) + h^w \cdot w^w(j_*^w, T^w, w) - R(i)H \\
1 = h^g + b \cdot d_{i,j_*^g} + l^g + x^g \text{ with } g \in \{h, w\} \\
h^g \leq \bar{h} \\
\Pi = u_x^c\left(P(x^h, x^w | d_{i,j_*^h}, d_{i,j_*^w})\right) - F^c(d^h, d^w)
\end{aligned}
$$

The value generated at home depends on time put in housework as well as on commuting: $\Pi^c = u_x^c\left(P(x^h, x^w | d^h, d^w)\right) - F^c(d^h, d^w)$, where P is an increasing function of x^h and x^w, while F^c is a (weakly) increasing function of d^h and d^w. Moreover, for couples, there is a complementarity between time and commuting: the productivity of x^h versus x^w is decreasing in $d^h - d^w$. The quantification section includes a detailed discussion of the parametrization choices.

The optimization problem of a couple choosing jobs is presented in equation 1.3.4, given residential location i. A couple knows the locations of their potential jobs j^h, j^w. First, a share of households learns about local jobs (located in i) and decide whether the husband, the wife or both should take it. This would give a short commute which is weighed against the potential of a better match. All who did not take the local job accept their initial offer or settle into non-participation in the labor market. j_*^g is the final optimally chosen job location (either i, j^g or 0). Lastly, the couple chooses housing continuous variables: housing quantity, consumption, hours of work and hours of home production for both partners. At the beginning of second period, old job ties are severed, new offers ar presented, news jobs chosen and continuous variables are re-optimized.

$$\begin{aligned}
V^C(i, j^h, j^w) &= \pi \cdot E_{\xi^{i,h},\xi^{i,w}}\Big[\max_{A_i^h, A_i^w} E_{\xi^{j^h},\xi^{j^w}}\big[\max_{A_j^h, A_j^w} v(A_i^h, A_i^w, A_j^h, A_j^w|\xi)\big]\Big] \qquad (1.3.4)\\
&+ (1-\pi) \cdot E_{\xi^{j^h},\xi^{j^w}}\big[\max_{A_j^h, A_j^w} v(0, 0, A_j^h, A_j^w|\xi)\big]\\
&\text{with } \xi = (\xi^{i,h}, \xi^{i,w}, \xi^{j^h}, \xi^{j^w})\\
v(A_i^h, A_i^w, A_j^h, A_j^w|...) &= U(i, j_*^h, j_*^w, \xi_0^h, \xi_0^w)\\
&\text{where } j_*^g = A_i^g \cdot i + (1 - A_i^g)A_j^g \cdot j^g\\
&\qquad \xi_0^g = A_i^g \cdot \xi^{i,g} + (1 - A_i^g)A_j^g \cdot \xi^{j,g}\\
&\text{with } g \in \{h, w\}
\end{aligned}$$

1.3.2 Job offers

In the data, men and women systematically work in different sectors of the economy. Moreover, similar jobs cluster together within geographic areas. This sets up couples for a potential disagreement about whether to locate closer to the husband's or wife's potential jobs. To capture this tension, I classify all offers into two distinct labor market sectors $T \in \{1, 2\}$. Each sector has a hub in the city and in one of the suburbs. Figure 1.9 summarizes the spatial distribution of first job offers. Red and green distinguishes sectors 1 and 2, showing that the first suburb has more offers in sector 1. In every period each individual is assigned to one sector, and only draws job offers from that sector. Sector assignments are random, but more men belong to sector one ($T = 1$).

Jobs coming from locations where the sector is concentrated in come with better benefits, both monetary and non-monetary. Specifically, both w and ξ are decreasing in $d'_{o,T}(j)$, the implied distance in a location j to a random first job offer from a labor market sector T. I use first offers which are exogenous, not actual distribution of jobs which is endogenous, for tractability of the solution. Nonetheless, the spatial distribution of jobs and the distribution of first offers is closely linked.

There is also a fixed gender pay gap for men and women in couples. There is no gender wage-gap for singles.[31] Overall, the wage function is defined as

$$w(j,T,g) = w_a \cdot e^{-w\left(d'_{o,T}(j)\right)+1_{g==h}\cdot 0.5 w_{gap} - 1_{g==w}\cdot 0.5 w_{gap}}$$

where $w\left(d'_{o,T}(j)\right) = w_{\bar{\Xi}} \cdot (b \cdot d'_{o,T}(j) - b \cdot \bar{d}'_{o,T})$. The non-monetary benefit has a deterministic component (mimicking the structure of wages) and a stochastic component (smoothing the participation choice and making sure the solutions are continuous). Overall, the non-monetary benefit is defined as

$$\xi(j,h,\xi_0,T) = \xi_0 + \Xi\left(d'_o(j,T),h\right)$$

, where ξ_0 is stochastic and $\Xi\left(d'_{o,T}(j),h\right) = -h \cdot \bar{\Xi} \cdot (b \cdot d'_{o,T}(j) - b \cdot \bar{d}'_{o,T})$ where h stands for hours of work. Overall, the heterogeneity of jobs captures agglomeration effects, where labor market sector hubs concentrate more productive and more innovative companies that provide their workers with a higher job satisfaction and better benefits.

Figure 1.9: Job offers in location

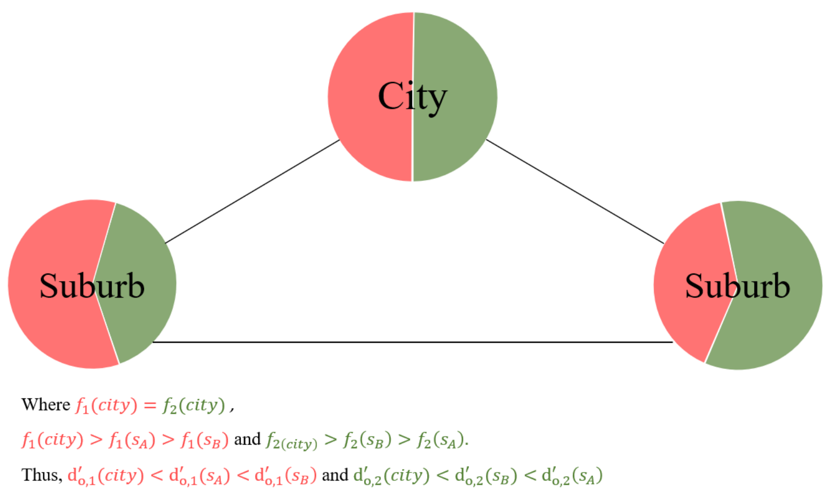

Red and green distinguishes sectors 1 and 2, showing that the first suburb has more offers in sector 1. Specifically, $f_T(j)$ is the share of first job offers in a labor market sector T in the location j. $d_{o,T}(j)$ is the implied distance in a location j to a random first job offer from a labor market sector T, that influences the benefits (monetary and non-monetary) of a job.

[31] Including a smaller gender gap for singles would not change any principal conclusions of the paper. Gender gaps in couples are imposed on the model to properly capture the incentives for whether a wife or a husband, or both, participate in the labor market, as the endogenously arriving gender wage gap from commuting incentives is too small. This paper does not aim to explain why wage gaps are much larger for people in couples.

1.3.3 Value of working close to home that rewards specialization

Men in couples commute much more than single men, and much more than women in general. In this section I propose a functional form for the cost of commuting that rewards specialization on this margin, and allows a model to quantitatively match the differences observed in the data. I posit that working close to home has additional benefits to all households beyond time use. For singles, this is a simple fixed cost of working, that scales with time spent commuting.

$$F(d) = \phi \cdot (b \cdot d) \tag{1.3.5}$$

Working close to home can be beneficial for many reasons, to be around in case of emergencies around the house, to accept deliveries, to walk the dog or to pick up a kid from a local school. Importantly, these kinds of benefits do not scale naturally with both partners being involved. If one member of the household is available by working close by, there is little harm in the other one working far away from home. This suggests that couples have an advantage over singles by sharing a household together, and that specialization of one member of the household working close to home is often the efficient choice. Equation 1.3.6 suggests a functional form that captures this intuition for couples

$$F^c(d^h, d^w) = \phi \cdot \min\Big\{b \cdot d^w, b \cdot d^h\Big\} \tag{1.3.6}$$

$F(d^h, d^w)$ implies that one person working close to home benefits the whole household. It is weakly increasing in the commuting time of husband and wife. I impose the scale of the value of working close to home ϕ to be the same for couples and singles (i.e. $F(d) = F^c(d, d)$). ϕ is identified from the variation in commuting between singles and married men and women.

The principal reason to include this additional cost in the model is that even though there are several potential explanations for gender differences in commuting, none of them have quantitatively important implications for the gap in commuting between singles and couples. Next, I go through the potential explanations for gendered commuting behavior in more detail.

In principle, a gender gap in commuting could be rationalized by assigning men and women different preferences (for example with $b^w > b^h$), as is sometimes implicitly or explicitly assumed in the literature. This, however, is rejected by the data as there is essentially no difference in commuting between single men and single women.

Theoretically, differences in commuting within couples could by caused by differences in bargaining power. If λ is low, husband's interests are not considered. If individual utility is decreasing in commuting time, the couple could prioritize short commutes of wives. This is the primary channel in Chiappori et al. (2018a). However, other features of the data do not support this explanation. If women commute less because couples prioritized their offers when choosing a residential location, we would see couples living systematically closer to the jobs in the wife's sector. In other words, $E(d_o^h - d_o^w)$ would be positive in the data (where d_o^h is the distance within MSA to an average job in the husband's most common industry and earnings segment, i.e. the husband's potential commute). However, the husband's potential commute is on average about the same or weakly shorter than the wife's potential commute (see table 1.4). This is true both when comparing raw means an when controlling for education, age and labor market sector fixed effects. Matching this moment in the data thus disqualifies a bargaining advantage as the primary factor explaining that women in couples have shorter commutes.[32]

Differences in commuting within couples could also come from differential job access to sectors dominated by men versus women (for example, if jobs where men work were concentrated far away from residential areas, it would be less surprising that men commute more). However, the same argument applies – women in couples do not systematically have better job access than men. Moreover, this mechanism would create a commuting gap among single men and women as well.

Standard explanations in the literature for gender differences in the labor market include either hard-wired preferences for women to stay at home, differences in the value of leisure, productivity in home production or differences in compensation in the labor market. All of these are a feature of the model. However, none of these lend themselves naturally to explain the observed gaps in commuting, unless commuting has additional costs to households beyond time use that reward specialization. First, wage gaps actually incentivize a reverse gap in commuting between men and women. Since commuting takes out of the time endowment, if the husband's time is more valuable, the couple is motivated to locate close to the husband's offer or to let him accept a local offer to avoid losing his valuable time. Second, a gender gap in home productivity also does not naturally motivate couples to prefer the wife's commute to be shorter than the husbands. This is because the decline in available time due to commuting $b \cdot d_{i,j}$ is offset by a decline in market hours, leaving time in home production and leisure unaffected.

Both wage gaps and differential productivity in home production do generate a gap in commuting between husbands and wives in the model through selection, motivating women to drop out of the labor force more than men, and with women being more likely to drop out

[32]Moreover, higher bargaining power does not necessarily lead to shorter commutes. Within the household, a long commute if more efficient for the household as a whole can be compensated with shorter hours of work or home production, with consumption or with a better job match.

when their potential commute is long. However, this channel has no potential in matching the gap between single men and husbands, and is not quantitatively sufficient to match the within-couple difference.

Similarly, gender gaps in commuting within couples could come from a technological complementarity between commuting and time in home production. For example, it is conceivable that it is more efficient for the same person to be able to pick up children from school and to take care of them afterwards. This also matches the fact that couples in suburbs are more specialized on housework and that women do more housework if the couples lives further away from her job opportunities. I include this mechanism as a channel that motivates couples to specialize in a gendered manner. Specifically, I extend the typical CES production function for home production to scale the wife's relative advantage up if her relative commute is shorter $\kappa_w(d^h - d^w) = \kappa_w^0 + (d^h - d^w) \cdot \kappa_d$:

$$P(x^h, x^w | d^h, d^w) = (\kappa_w(d^h - d^w) x_w^{1-\eta_x} + (1 - \kappa_w(d^h - d^w)) x_h^{1-\eta_x})^{\frac{1}{1-\eta_x}} \tag{1.3.7}$$

By making $\kappa_w(d^h - d^w)$ increasing in $d^h - d^w$, I facilitate that commuting in couples leads the individual to put in fewer hours in home production, letting the partner compensate for the loss. Since $\kappa_w > 0.5$, $\kappa_d > 0$ motivates the couple to prioritize short commutes of wives and helps to match the sensitivity of home production hours to differences in the distance to opportunities within couples from table 1.6.

Lastly, part of the difference in commuting between husband and wives and between husbands and singles is rationalized by the fact that benefits to commuting in the model scale with hours. Since husbands work longer hours than other groups, their returns to commuting are bigger. This mechanism is a part of the model and thus is accounted for, though would not be enough on its own.

1.3.4 Choosing residential location and marriage

Each period every household chooses a residential location, either the city or one of the two suburbs. This is a standard discrete choice as in McFadden (1977), comparing systematic benefits (access to current job offers, access to other potential jobs, amenity values and costs of housing) with idiosyncratic amenity preferences per location ϵ_i. For singles:

$$\max_{i=c,s_A,s_B} V^s(i|j,T) + \epsilon_i$$

For couples:

$$\max_{i=c,s_A,s_B} V^c(i|j^h, j^w, T^h, T^w) + \epsilon_i$$

With ϵ_i following a Type-1 extreme value distribution, the choice probabilities can be solved in closed form.

The choice of marriage for a man (h) and a woman (w) is done by comparing the expected value from remaining single for 2 periods and the expected value of being married (given the husband's bargaining weight λ). As in Choo and Siow (2006), I assume that in addition to the systematic component of utility in the married or single state each individual receives an idiosyncratic payoff θ^g that is specific to him or her. The expected value from remaining single for 2 periods is defined by plugging optimal choices of time use, spending, job taking and residential location in the period utility functions

$$u^s + \theta_s = 2 \cdot E_{T,\epsilon_i,j,\xi_0^i,\xi_0^j} \left(u_c^s(c) + u_l(l) + u_H(H) + a^s(i) + \Pi^s(x,d) + \xi \right) + \theta_s$$

where the expectation is taken over job-match shocks for the offered and local jobs ξ_i, ξ_j, draw of job offer location j, the idiosyncratic location preferences ϵ_i, and ultimately the labor market sector assignment T.

Similarly, the expected values in marriage for a man and a woman is defined as

$$u^h(\lambda) + \theta_h =$$
$$E_{T^g,\epsilon_i,j_t^g,\xi_{i_t}^g,\xi_{j_t^g}^g} \left(\sum_{t=1}^{2} u_c(c_t^h) + u_l(l_t^h) + u_H(H_t/2) + a^c(i) + \Pi(x_t^h, x_t^w, d_t^h, d_t^w) + \xi_t^h \right) + \Theta + \theta^h$$

$$u^w(\lambda) + \theta_w =$$
$$E_{T^g,\epsilon_i,j_t^g,\xi_{i_t}^g,\xi_{j_t^g}^g} \left(\sum_{t=1}^{2} u_c(c_t^w) + u_l(l_t^w) + u_H(H_t/2) + a^c(i) + \Pi(x_t^h, x_t^w, d_t^h, d_t^w) + \xi_t^w \right) + \Theta + \theta^w$$

Expectations are taken over job-match shocks for the offered and local jobs for two periods and for both partners, draws of job offer locations, the idiosyncratic location preferences ϵ_i, and ultimately the labor market sector assignments T for self and the partner.

In addition to idiosyncratic preferences θ_g (where $g \in \{h, w\}$), I allow for unaccounted-for benefits to marriage (a constant Θ). Note that the Pareto weight does not depend on the realization of uncertainty. This implies full commitment and efficient risk sharing within the household. Moreover, I assume that assignment to a labor market sector T is randomly reshuffled after marriage (preserving the gender composition of each). This way with matching at random there is no differentiation in the incentive to marry by labor market sector T. This is a simplifying assumption that preserves the internal logic of one common marriage market.

A man decides to enter the marriage market if, given λ,

$$u^h(\lambda) + \theta_h > u^s + \theta_s$$

As with residential location, I assume that the idiosyncratic payoffs θ_g and θ_s, observed prior to the marriage decision, follow the Type-1 extreme value distribution with a zero location parameter and the scale parameter σ_m. Thus the proportion of men or women $g \in h, w$ who would like to be married has a closed form and is given by

$$p^g(\lambda) = \frac{e^{\frac{u^g(\lambda) - u^{g,s}}{\sigma_m}}}{1 + e^{\frac{u^g(\lambda) - u^{g,s}}{\sigma_m}}}$$

Denoting M and F as the supply of men and women in the marriage market, an equilibrium bargaining weight λ satisfies

$$M \cdot p^h(\lambda) = F \cdot p^w(\lambda)$$

Assuming a gender-balanced metro-area, with equal number of men and women, this equation boils down to a simple equilibrium condition.

$$u^h(\lambda) - u^{h,s} = u^w(\lambda) - u^{w,s}$$

1.3.5 Equilibrium

There are four overlapping markets: three housing markets and one marriage market. With three discrete locations, there are three prices $\{R(i)\}_{i=c,s_A,s_B}$ to clear three housing markets. The bargaining weight λ is endogenous in the model and serves as a price clearing the marriage market. Supplies of housing $\{H_i\}_{i=c,s_A,s_B}$ are fixed in each residential location, and they sum up to 1 (equal to the total population of the metro-area). Individuals are differentiated by gender $g \in \{h, w\}$ (with an equal number of men and women living in the metro-area) and labor market sector assignment $T \in \{1, 2\}$ (with an exogenous distribution $\{s^g(T)\}_{g\in\{h,w\},T\in\{1,2\}}$). Moreover, the location of first job offers is drawn from an exogenous distribution $\{f_T(j)\}_{T\in\{1,2\},j\in\{c,s_A,s_B\}}$. The matching to couples is random with respect to T. However, marrying is a choice, so the share of people married is endogenous. Thus, the overall distribution of different types of households is endogenous in the model.

Definition 1. *Given fixed supplies of housing units per location $\{H_i\}_{i\in\{c,s_A,s_B\}}$, exogenous distributions of individuals to sectors $\{s^g(T)\}_{g\in\{h,w\},T\in\{1,2\}}$, and of locations of first job offers $\{f_T(j)\}_{T\in\{1,2\},j\in\{c,s_A,s_B\}}$, a housing and marriage market equilibrium is a set of* ***rents*** *per location $\{\boldsymbol{R(i)}\}_{i\in\{c,s_A,s_B\}}$ and a* ***bargaining weight $\boldsymbol{\lambda}$****, such that choices are optimal, the* *choice probabilities to enter the marriage market* $\{\boldsymbol{p^g}\}$ *are equal for men and women*

$$\boldsymbol{p^h} = \boldsymbol{p^w}$$

and the choice probabilities to live in a location $\{\boldsymbol{P^s_{j,T}(i)}\}$ *and* $\{\boldsymbol{P^c(i)_{(j^h,j^w),(T^h,T^w)}}\}$ *and the housing demands* $\{\boldsymbol{H^{s,g}_{i,j,T}}\}$ *and* $\{\boldsymbol{H^c_{i,(j^h,j^w),(T^h,T^w)}}\}$ *are such that the housing markets clear*

$$\begin{aligned}
H_i = & \sum_{g\in h,w}\ \sum_{\text{sector } T}\ \sum_{\text{offer in } j} N^{s,g}_{i,j,T}\cdot \boldsymbol{H^{s,g}_{i,j,T}} \\
& + \sum_{(T^h,T^w)}\sum_{(j^h,j^w)} N^c_{i,(j^h,j^w),(T^h,T^w)}\cdot \boldsymbol{H^c_{i,(j^h,j^w),(T^h,T^w)}} \\
& + \sum_{(T^h,T^w)}\sum_{(j^h,j^w)} N^c_{i,(j^h,j^w),(T^h,T^w)}\cdot \left(\sum_{(j_2^h,j_2^w)} f_{T^h}(j_2^h) f_{T^w}(j_2^w)\cdot \boldsymbol{H^c_{i,(j_2^h,j_2^w),(T^h,T^w)}}\right)
\end{aligned}$$

where

$$\begin{aligned}
N^{s,g}_{i,j,T} &= \left(\frac{1}{3} + (1-\boldsymbol{p^g})\frac{2}{3}\right)\cdot s^g(T)\cdot f_T(j)\cdot \boldsymbol{P^s_{j,T}(i)} \\
N^c_{i,(j^h,j^w),(T^h,T^w)} &= \boldsymbol{p^h}\frac{2}{3}\cdot s^h(T^h)s^w(T^w)\cdot f_{T^h}(j^h)f_{T^w}(j^w)\cdot \boldsymbol{P^c_{(j^h,j^w),(T^h,T^w)}(i)}
\end{aligned}$$

Equilibrium prices $\{R(i)\}_{i=c,s_A,s_B}$ are to be interpreted as residential rents per unit of housing in each neighborhood.[33]

1.3.6 Selecting parameter values

I populate the metro area with a fixed number of individuals equal to the number of housing units, half men and half women. Each location has the same number of housing units. Overall, I fit the model to match moments in the data summarizing the distribution of people and jobs within the urban space, time use of couples and singles, and residential location and commuting behavior patterns presented in section 1.2. These are created using several data sources as described in the empirical section: the geocoded PSID sample described above, the LODES jobs data and the 2000 Census as well as 2006-2010 ACS IPUMS samples (Ruggles et al., 2019).[34] Table A.11 presents the list of targeted data moments $\bar{m}$ used in the estimation routine.

I set preferences over consumption, leisure, housing quantity and home production to

[33] Section A.2.1 in the appendix presents details on how the model is solved.

[34] Section A.3.1 in the appendix describes the details.

have the constant relative risk aversion functional form (with $z \in \{c, l, H, x\}$ and $g \in \{s, c\}$ denoting singles or couples).

$$u_z(z) = \Omega_z^g \cdot \frac{z^{1-\omega_z}}{1-\omega_z}$$

Preferences are imposed to be the same for singles and couples, except for the value of home production and amenities ($\Omega_z^s = \Omega_z^c$ for $z \in \{c, l, H\}$). Since childcare is part of producing value at home and couples are more likely to have young children in their household (see figures A.2), it is easily imaginable that couples put a higher value on home production time (i.e. that $\Omega_x^s < \Omega_x^c$). Time endowment is set to one (but time variables are scaled in utility to be in the same scale as consumption).

Section A.3 in the appendix describes in detail what parameters are identified by what variation in the data. The spatial structure of the metro-area and the distribution of jobs is identified from distances between two random jobs (any and within the same labor market), share of jobs and of people close to city center, and distances to a random job from a place of residence. Moreover, I target average $|d_o^w - d_o^h|$, the average absolute value of the difference between potential commutes of husband and wife. This statistic determines the potential of disagreement within couples – the larger this difference in absolute value, the bigger the challenge for a couple to balance living close to opportunities for both household members. Amenities are identified to match the residential location choices of couples and singles as well as the price gradient between the city and suburbs. The distribution of stochastic match shocks ξ_0 helps to fit labor force participation of men and women in couples. To identify the benefits (monetary and non-monetary) to working close to a sector hub, I include moments from tables 1.5 and 1.6 in the estimation. Hours of work and housework in the PSID, as well as a share of income spent on housing from Bureau U.S.BureauofLaborStatistics (2020), is used to identify preference parameters over continuous variables. Bargaining weight λ is identified using the model's implied derivative of marriage rates of men versus women with respect to variation in the ratio of men and women, matching that to the equivalent variation in the data across metro-areas (mimicking the identification argument in Gayle and Shephard (2019)). Table A.12 presents a complete list of parameters to be estimated.

I estimate the model with a moment based procedure.[35] There are 26 parameters to be estimated and 44 moments used in estimation. A subset of the parameters α^1 is fit directly within the estimation routine to exactly match a moment condition at each iteration, using current guesses of other parameters combined with moments in the data. This partition decreases the number of parameters that are estimated via a global search, decreasing the computational burden in estimating the model. Letting $\alpha = [\alpha^1, \alpha^2]$ denote the $Bx1$

[35] A small subset of the parameters is calibrated outside the estimation routine.

parameter vector, the estimation problem may be formally described as

$$[\alpha^1, \alpha^2] = \arg\min_{\alpha^2} \; [m(\alpha) - \bar{m}]^T W [m(\alpha) - \bar{m}] \tag{1.3.8}$$

$$\text{s.t. } \alpha^1 = f(\alpha^2, \bar{m}) \tag{1.3.9}$$

W is constructed based on the inverse of the variance-covariance matrix of the data[36].

1.3.7 Fit of the model

Table 1.8 highlights that the model matches very well the commuting patterns of couples and singles, for men and women. Specifically, the large difference between the commute of husbands and single men as well as the small difference for women is successfully captured by the model. This is a combination of the couples moving to suburbs (thus further away from jobs in general) and men in couples being more willing to accept long commutes, wherever they live. Men in couples accept longer commutes, because couples are facing a collocation issue, husbands having higher returns to commuting due to their longer hours and especially because of the benefit of working close to home that rewards specialization. Effectively, husbands have a lower cost of commuting compared to single men, because responsibilities around the house are already covered by the wife.[37]

Couples are more likely to live in the suburbs, and suburbs have on average longer commutes. Still this alone cannot account for the difference in commuting between single men and men in couples in the model (as well as in the data). The bar graph in 1.8 illustrates this point, showing that if singles made the residential choices of couples, their commutes would increase by about a mile. This is because distances between neighborhoods, distribution of jobs and distribution of location choices between city and suburbs of couples and singles in the model are constrained to match corresponding moments in the data. Table 1.9 presents these moments and their fit. The model, with only three locations, is capable of capturing the distances between jobs and people and between one random job to another, as well as the distribution of jobs and people between suburbs and city quite well. Couples are less likely to live in the city, so they live on average further away from jobs. There are more jobs in the

[36]For moments from different samples I set the covariance to zero. For moments within the same sample I compute the variance-covariance matrix using influence functions of individual moments, and clustering at the MSA level. Moreover, I increase the weight of the most crucial moments (see details in the appendix).

[37]A.3.3 shows that an equivalent model without the specialization-rewarding cost of commuting fails to match the difference between single men and men in couples.

Table 1.8: Commuting moments data versus model

Moment	Model value	Data value
Average commute of single d^s	8.642	8.667
$d_h^s - d^h$	**-2.582**	**-2.708**
$d_w^s - d^w$	-0.307	-0.297

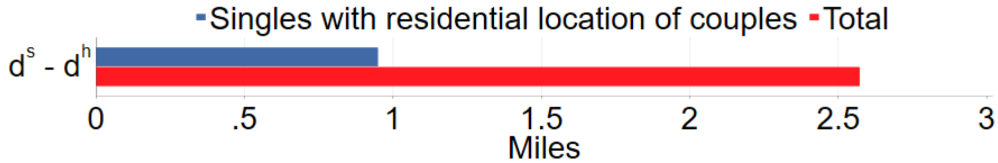

d^s is the average commuting distance of singles in miles. $d_h^s - d^h$ is the difference in commuting distance between single men and men in couples. $d_w^s - d^w$ is the equivalent for women. The blue bar presents the commuting difference between single men and men in couples that is accounted for by their differences in residential location.

city than in the suburbs. There are also slightly more people in the city (living in smaller units). Table A.11 in the appendix shows the fit on all targeted moments.

Table 1.9: Moments describing the spatial structure of a metro area: data versus model.

Moment	Model value	Data value
Distance to an average job for a couple (d_j^h)	21.729	20.277
Distance to an average job in own labor market for a husband (d_o^h)	21.667	20.027
Distance between 2 random jobs	18.811	17.300
Distance between 2 random jobs of the husbands labor market	18.632	16.267
Distance between husbands and wifes actual jobs	12.940	9.740
$\lvert d_o^w - d_o^h \rvert$	1.785	1.862
$P(city \mid couple) - P(city \mid single)$	0.133	0.070
$d_o^s - d_o^h$	-1.694	-1.693
$d_j^s - d_j^h$	-1.640	-1.410
$d_o^w - d_o^h$	0.046	0.028
Share of jobs in city	0.496	0.498
Share of population in city	0.340	0.392

Moments describing the spatial structure of a metro area, as well as commuting and location preferences. In the data, a 'city' is defined as a radius around city center of 10 miles.

Table 1.10 compares a key result from table 1.1 that is not targeted in estimation, a difference in commuting between single men and men in couples after netting out the potential commute differences d_o, between the model and the data. While the pattern is similar, the model in fact attributes a bigger role to the move to the suburbs in the commuting gap between singles and couples. Therefore, if anything, the model likely underestimates how much specialization on commuting within couples allows husbands to accept long commutes, beyond how far away from jobs they move. Together tables 1.10 and 1.8 provide insight into why a move to the suburb alone does not account for how much men commute more after they form a couples. Couples live about 1.5 miles further from a random job in their labor market. So even if jobs were taken entirely randomly, husbands would only commute about 1.5 miles more. However, jobs and locations are not random – people prefer shorter commutes, all else equal. Thus an increase of 1.5 miles in the distance from a random job translates to less than a mile of increase in actual commuting distance. The rest has to be accounted for by a change in behavior towards jobs. To sum up, the difference between couples and singles in the share of living in the suburbs is quite simply too small to explain the commuting differences.

Table 1.10: Difference in commuting between husbands and single men, controlling for d_o: data (as in table 1.1) versus model equivalent.

	Commuting distance (miles) men	
	Data	Model
In couple	**2.555**	**1.513**
	(.638)	
d^o	.207	0.651
	(.062)	
N	23243	

Available in 1975-1976, 1978-1986; plus in 1969-1974, 1977 for heads of households only. Sample of only those in a couple, or those that have never been observed in one, but eventually they will be.

All regressions include year, age, MSA, education, race, cohort controls. *SEs* statistics in parentheses. SEs clustered at the MSA level.

The estimation routine also does not target any differences in labor market outcomes within couples based on whether they live in the city or the suburbs, similar to the behavior presented in figure 1.1. Here I present the fit on these non-targeted moments. In the data I regress a gender gap measure in a couple on binned distance of their residence from city center in a metro area, controlling for dummies for demographic characteristics of the couples. I compare the estimated differences $\alpha^{\text{dist bin } j}$ in gender gaps between those living close to city center and those living further away to the difference between city and suburbs in the model (which in the model presents a distance of over 20 miles). Figure 1.10 visualizes the comparison between the data and the model, for share of market hours ${\frac{h^h}{h^h+h^w}}_{t,i}$, share of housework hours ${\frac{x^h}{x^h+x^w}}_{t,i}$ and difference in commuting distance $(d^h - d^w)_{t,i}$. Overall, the model is successful in capturing how much gender gaps are larger in the suburbs.

Lastly, I check whether the distribution of wages in the data with respect to how close a job is to the city center aligns with that in the model. Table 1.11 presents the results.

Figure 1.10: Gender gaps within couples by distance of residential location from city center: data vs model

Share of household market hours provided by the husband

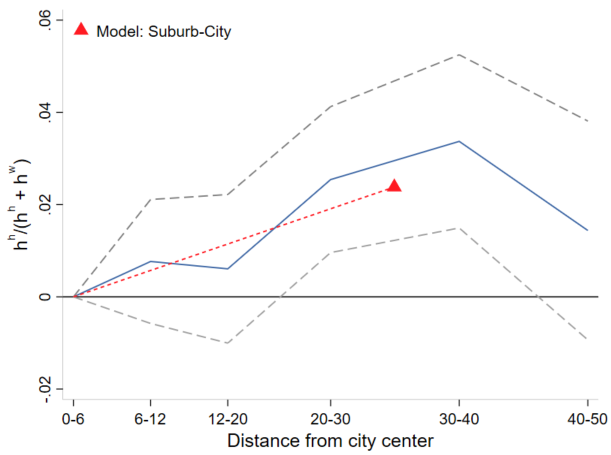

Share of household housework hours provided by the husbands

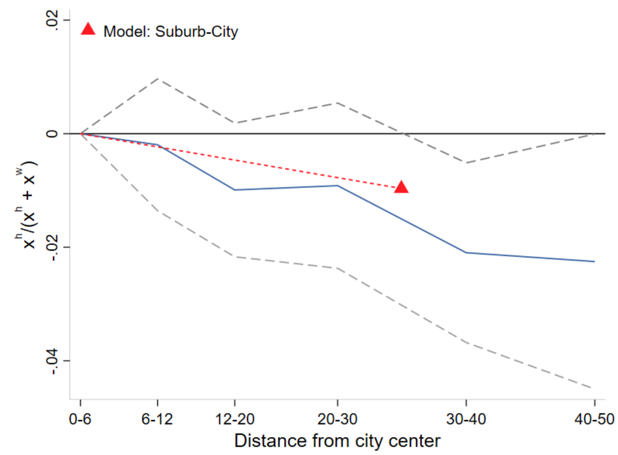

[

.] Difference in commuting distance $(d^h - d^w)_{t,i}$

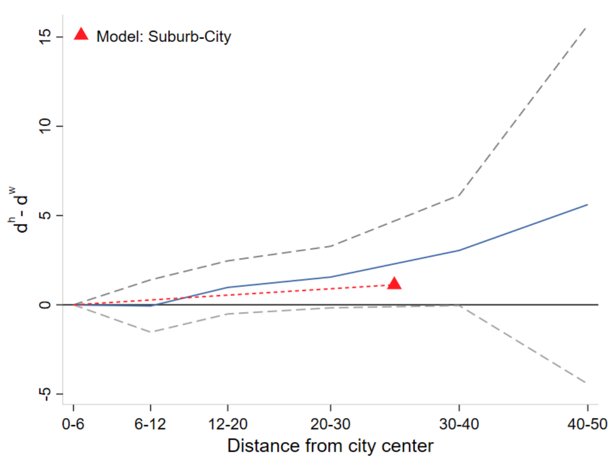

As in figure 1.1, I regress a gender gap measure in a couple on binned distance from city center in a metro area, controlling for dummies for age and education of both spouses, race of the head and number of children.

$$\frac{h^h}{h^h+h^w}_{t,i} = \sum_{j=2}^{N} \alpha_i^{\text{dist bin } j} + \alpha_i^{ah} + \alpha_i^{aw} + \alpha_i^{educh} + \alpha_i^{educw} + \alpha_t + \alpha_i^{raceh} + \alpha_i^{\#children} + \epsilon_{i,t}$$

Table 1.11: Wage gradient by distance of the job to city center: data versus model

	Data	Model
$d_{\text{job to city}}$	-0.00387	-0.00790

The first column is based on the PSID sample 18-50 years old with data available on the actual census tract of the job (waves 2013-2017). The coefficient presented comes from the following regression $\log(wage)_{it} = \beta d_{it}^{\text{job, city center}} + \gamma X_i + \alpha_{age} + \alpha_t + \alpha_{msa} + \epsilon_{it}$ where X_i includes race, gender, in couple, education, industry and education-cross-industry dummies.

1.4 Commuting and the value of marriage

Over the 20th century U.S. metropolitan areas have been sprawling out in space, increasing the necessary commutes one has to accept to work in a desirable job. In this section I mimic this trend by changing the geographic size of the model metro area. Specifically, I re-solve the model with different values of $D = D(1,2) = D(1,3)$ (implying $D(2,3) = D \cdot \frac{\bar{D}(2,3)}{\bar{D}(1,3)}$), keeping all the other parameters the same. This means that the metro area is stretched out in space, without any change in amenities or productivity (materialized as wages or non-monetary benefits) at work. Figure 1.11 summarizes this counter-factual. Lower values of D represent dense metro-areas with a short average distance from suburbs to the center. As such, this counter-factual also mimics a policy intervention that makes suburbs more or less accessible, for example, by (dis-)investing in public transit.

Figure 1.11: Connectivity between suburbs and the city

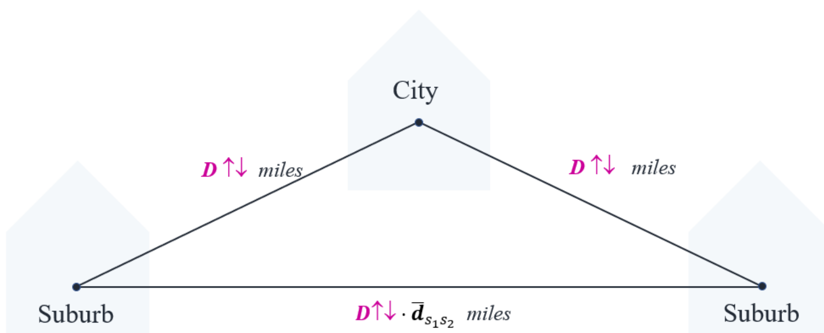

I study the effect on welfare of men and women and singles and couples, asking for whom commuting is ultimately most costly in a combined housing and marriage market equilibrium. All comparisons are between different steady states – alternative scenarios where the metro area would develop differently. I measure welfare for each subgroup (single and married men and women) as period utility averaged over the respective population: $W = \int_i u^*(i)$.[38] Moreover, I show how job access, gender gaps within couples, residential rents and sorting changes.

Overall, increasing D is a negative technology shock to the metro-area. All groups are hurt by it, on average. However, not everybody is affected the same way. Figure 1.12 presents the first set of results – the differential incidence of losing job access. On the horizontal access is the respective distance between suburbs and city D, with the middle point representing the baseline value. The first figure shows that the distance from a residential location to a

[38]Since people decide whether to marry before all heterogeneity is revealed for the period, there is no difference in an average period utility of the original singles versus the new singles.

random job increases more in the suburbs than in the city. This is both because the city is positioned in the center, and because more jobs are offered in the city. Because couples are more likely to live in the suburbs, their job access deteriorates more compared to singles. In other words, the incidence of a policy that makes cities less easily accessible from the suburbs is higher on couples.

Figure 1.12: Job access when metro area grows in space

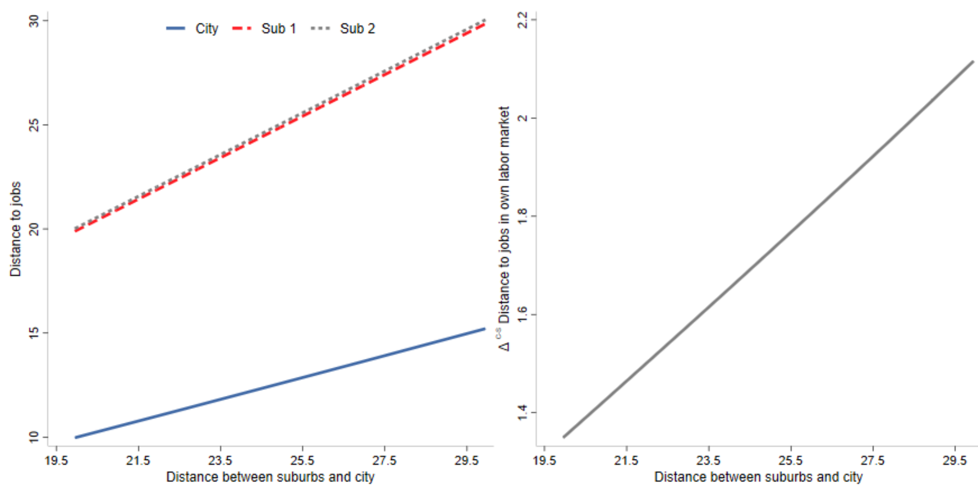

Counter-factual simulations of the model, varying the distances between neighborhoods while keeping the shape of the metro-area fixed. Distance to a random job by location. Distance to a random job: difference between singles and couples

Figure 1.13: Commuting and work gender gaps when metro area grows in space

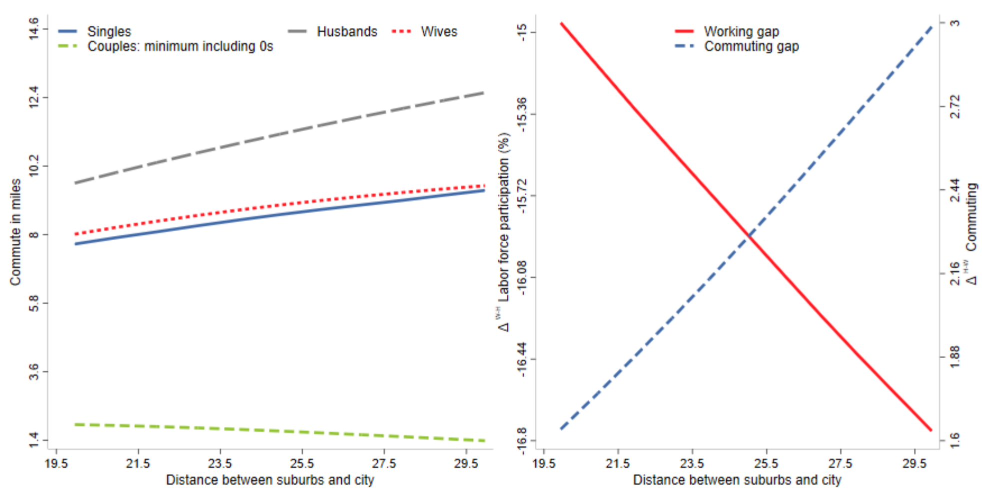

Commuting of all subgroups. Gender gaps within couples.

And yet, they are ultimately less affected in terms of welfare. This is because jobs are not taken randomly with respect to residential location. Households hustle to make their commutes short – by moving close to jobs and by taking jobs close to home. Moreover, couples specialize with one (more often the wife) taking a local job or no job at all while the other accepts long commutes. Because commuting is in part costly on the household level through only the shortest commute within household, couples are rewarded for their specialization. Figure 1.13 shows that husband's commute increases the most. Wives and singles also increase their commuting, but less so. The second figure shows that indeed within couples gender gaps increase, both in labor force participation and in commuting. Overall, women in couples are more sensitive in their labor market outcomes to long commutes than men. Figure 1.13 also shows how gender gaps help couples evade part of the commuting costs. Even though all groups of people commute more on average, within couples at least one is often close by (assigning 0 commute to those who drop out of the labor force).

Figure 1.14: Housing when metro area grows in space

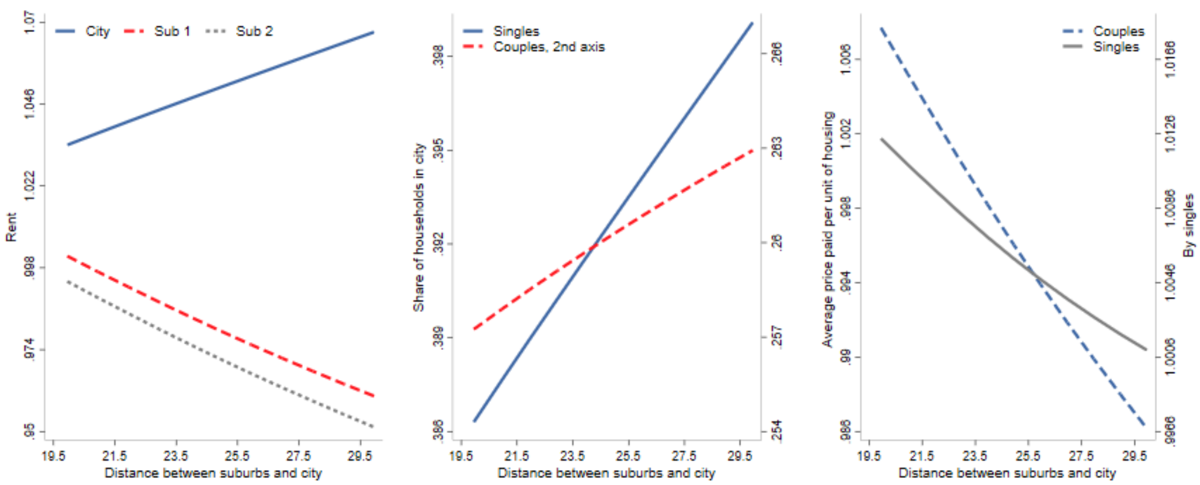

Housing rents, sorting and housing costs of couples and singles.

On top of the endogenous responses in labor market behavior and specialization within households, the housing market is affected by long commutes from suburbs to the city. Figure 1.14 presents the results. Housing rents increase in the city, but fall in the suburbs. All households try to flock to the city, however singles are more motivated than couples, because they cannot evade commuting costs through specialization and they care less about suburban amenities. The spatial segregation of couples and singles within a metro area increases. As a result, singles are now forced to overpay more for housing. Overall, housing prices fall, because households have on average less time to work and thus less income to spend.

So who looses the most from long commutes? Husbands commute the most and their commuting increases more. Wives are most affected in terms of their labor market outcomes.

Figure 1.15: Welfare when metro area grows in space

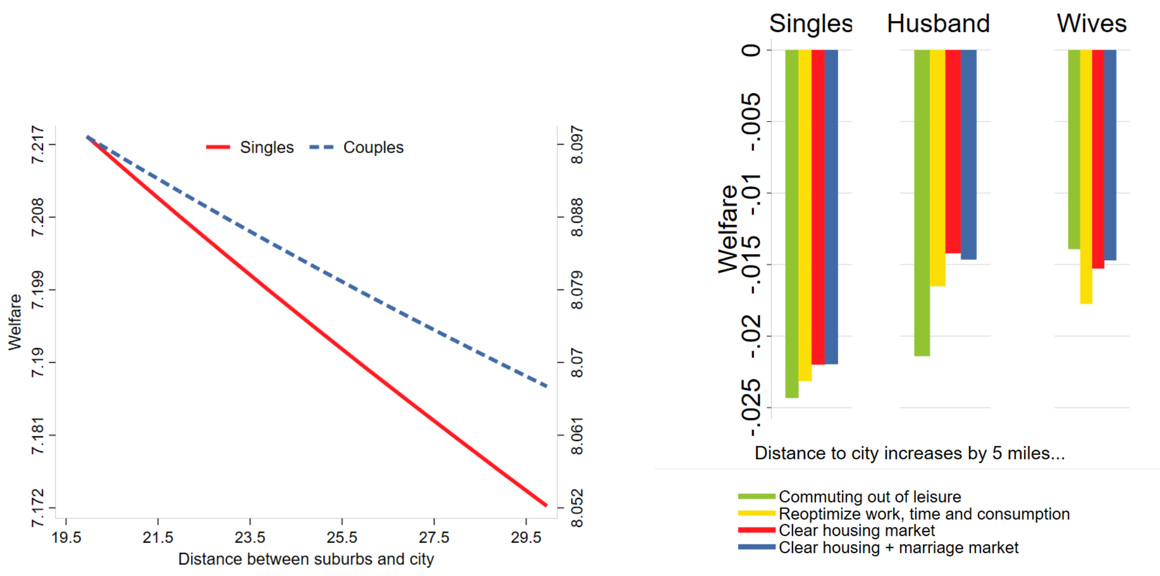

Change in welfare for singles, husbands and wives between the baseline metro area, and a sprawled metro area. The second figure presents a decomposition of the drop in welfare when the suburbs are 5 miles further away from the city, depending on which parts of the model are re-solved. 0.01 drop in welfare represents approximately a 1.5% decline in consumption.

Couples overall lose most job access. However, it is singles who lose the most when commutes are longer. Both husbands and wives change their behavior more than singles, but those large observable changes in behavior are in fact a sign that they have the added flexibility to do something about an inconvenient situation. Overall, couples benefit from evading commuting costs through specialization within the household, as housing prices adjust in the housing equilibrium and distribution of tasks and resources within the couple adjusts in the marriage market equilibrium. Figure 1.15 presents this result. While welfare falls for everybody, it falls less for couples.

The endogenous responses of behavior and prices in the the joint housing and marriage market equilibrium only reinforce this result. Figure 1.15 presents the decline in welfare for singles, husbands and wives when the distance between suburb and cities increases by 5 miles. The green bar shows the effect when longer commutes simply subtract from leisure, without re-solving the housing and marriage market equilibrium. In this case, husbands loose markedly more than wives, precisely because they are the ones who are locked into the longest commutes. The yellow bars show that when households re-optimize but the housing and marriage market does not re-clear, it is the wives who lose more. This is precisely because they change their labor market behavior, dropping out of jobs they liked into worse local jobs or out of the labor force altogether, to diminish the burden of commuting costs on the household. Moreover, the gap between singles and couples widens, as specialization

Figure 1.16: Value of marriage when metro area grows in space

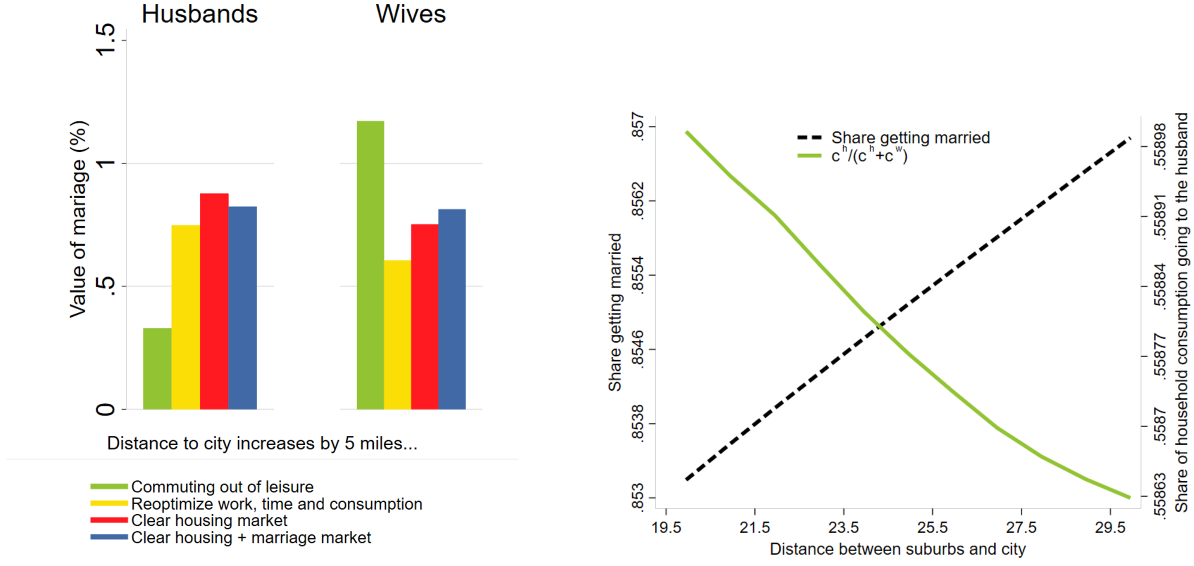

Change in welfare and the value of marriage between the baseline metro area, and a sprawled metro area where the suburbs are 5 miles further away from the city. Value of marriage is defined as the difference in period welfare between a single and a person in a couple $\Delta^{h-s}W$ and $\Delta^{w-s}W$. New marriage market equilibrium: share of people getting married and bargaining position of husbands – determined by the price in the marriage market.

helps couples evade the commuting costs. The red bars shows the effect on welfare when the housing market re-clears. The blue bars show the final result within a full housing and marriage market equilibrium. The new prices help couples further (both husbands and wives), because they can now enjoy cheaper prices in the suburbs.

Figure 1.16 presents this explicitly, showing the effect of longer distances between suburbs and city on the value of marriage. The value of marriage increases for both men and women more after the new equilibrium is achieved. Couples save on housing, husbands keep their long commutes or opt for even better jobs, wives take local jobs or stay at home, but are ultimately compensated with more leisure within the household. This way couples evade part of the commuting costs. Figure 1.16 shows that the share of people marrying increases, while the bargaining weight of husbands (here presented as the share of couple consumption going to the husband) adjusts.

Figure 1.17 shows how the distribution of tasks and resources within couples is reorganized. The first figure shows that when metro-areas have long commutes, wives gain leisure compared to husbands, but loose more on non-monetary benefits from work. Figures 2 and 3 decompose the change in the value of marriage accounted for by different sources of welfare for men and women. For both men and women, marriage becomes more valuable partially through home production and the household value of somebody working close to home. Moreover for husbands, the value of marriage increases most through better jobs that they enjoy more. On the other hand, wives take worse jobs, but their marriage is more valuable for them through more leisure.

Figure 1.17: Welfare effects decomposed into elements of utility

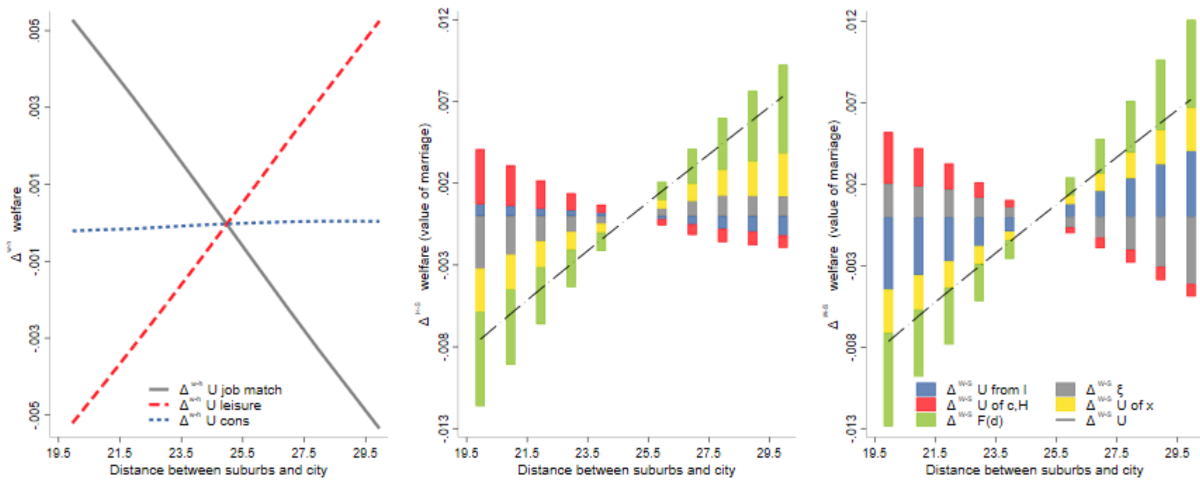

Decomposing the husband-wife welfare gap: effect of consumption, leisure and non-monetary benefits from work. Decomposing the change in the value of marriage: accounted for by leisure, consumption, household value of having somebody work close to home, non-monetary benefits from work and home production.

To summarize, the counter-factual simulation shows that if metro-areas spread out, couples gain compared to singles and marriage becomes more valuable for both men and women. This is facilitated through larger gender gaps in the labor market, more residential sorting and cheaper housing in the suburbs. While longer potential commutes are costly to all, single people have the most to lose.

1.5 Testing model predictions with cross metro-area correlations

In this section I provide further validation for the model, by comparing the counterfactual simulations with variation across U.S. metro areas. First I replicate results by Black et al. (2014) showing that metro areas with longer commutes have larger differences in labor force participation between men and women in couples. Using the 2000 IPUMS Census sample I run the following regression

$$Working_{im} = \beta C_m \cdot (\text{woman}_i) + \gamma X_i + \delta X_m + \epsilon_{i,m}$$

where i stands for an individual, m for a metro are, C_m is the average annualized hours of commuting in a metro m and the sample is restricted to people in couples. β is the coefficient of interest – it shows the differential impact of living in a place of long-commutes on men and women.

Table 1.12: Cross-metro area variation in work and commuting compared to model simulations

	Working	Commute (annualized)
$C_m\cdot$ (woman)	-0.0552 (0.0101)	-0.239 (0.0379)
C	x	x
C· (age, race, educ)	x	x
1-digit industry		x
Sample :	couples	couples

SEs in parentheses, clustered at the MSA level.

All regressions include age, education, region, race dummies and MSA size polynomial.

Data

Source: IPUMS 2000 Census 5% sample. Sample: 18-50 years old, married or cohabiting, MSAs of at least 250k people. "Working" is equal to one if the person worked at least for 1 week in the past year and is scaled up by 100 so that results are interpreted as percentage point changes. Industry dummies are for 1-digit NAICS codes. D_m is the average of annualized commuting hours for all residents of the MSA that do not work from home.

	Working	Commute
	-0.0540	-0.666

Implied by the D counter-factual.

Model simulations

Table 1.12 shows the results. In metro-areas with 16.5 more average hours of commuting per year (roughly corresponding to 1 mile) the gender gap in labor-force participation in couples is higher by almost a whole percentage point. I repeat the exercise with commuting itself on the left hand side (and using a sample of working individuals), where d is a commute of an individual (measured in annualized hours).

$$d_i = \beta_c C_m \cdot (\text{woman}_i) + \gamma X_i + \delta X_m + \epsilon_{i,m}$$

The results show that in metro-areas with longer average commutes the difference between the time spent commuting of wives and husbands increases. When the average commute in a metro-area increases by an hour, husbands' commute increases more than wives' by an average of 0.24 hours. Qualitatively, this is exactly what happens in the model. Quantitatively, counter-factual exercises above imply a somewhat bigger effect on the gender gap in commuting.

Table 1.13: Cross-metro area variation in marriage compared to model simulations

	(Ever married or cohabiting)$\cdot$100			
C	.00465 (.0097)	.0158 (.0048)		
$C_{husbands}$			.0114 (.0067)	.0134 (.0031)
Politics and church-affiliation proxies		x		x
Sample:	$30 \leq \text{age} \leq 50$			

SEs in parentheses, clustered at the MSA level.

All data regressions include age, education, region, race dummies and MSA size polynomial.

Data

	(Ever married)$\cdot$100	
C	0.0110	
$C_{husbands}$		0.00768

Implied by the D counter-factual.

Model simulations.

Source: IPUMS 2000 Census 5% sample. Sample: 30-50 years old, MSAs of at least 250k people. The outcome variable is equal to one if the person is married, divorced, separated, widowed or currently cohabiting. Columns 3 and 4 replace C_m with an average commute in an MSA among married men.

Next I test the model prediction that larger average commutes are actually conducive of couple formation, by making single life disproportionately costly compared to being in a couple and being able to specialize. I focus on the sub-population of 30-50 years of age, corresponding to the population in the model that is either in a couple or perpetually single. Using the 2000 IPUMS Census sample I run the following regression

$$\text{Ever in couple}_{im} = \gamma C_m + \gamma X_i + \delta X_m + \epsilon_{i,m}, \quad \forall i : \text{age}_i \geq 30, \text{age}_i \leq 50$$

Ever in couple$_{im}$ is a dummy variable equal to 1 if the person is married, currently cohabiting with a partner or has ever been married. C_m, again, is the average annualized hours of commuting in a metro are m. Table A.16 shows that, at least when metro-area-level controls X_m include religious participation and proxies for political affiliation, the estimate of γ is positive and statistically significant. Across metro areas those with a longer average commute tend to have fewer people staying perpetually single. This correlation in the data could be caused by a selection effect – metro-areas with more couples have higher average commutes because it is the married men who commute most. Columns 3 and 4 in table A.16 show the result is robust to replacing C_m with the average commute among only married men, avoiding this type of selection.

Table 1.13 again compares the cross-metro area correlations to the equivalent change in marriage rates implied by the increase in distances between neighborhoods in the model simulations. As in the data, simulated metropolitan areas with longer average commutes have a higher share of the population eventually marrying.

1.6 Conclusion

In this paper I show that longer potential commutes make marriage more valuable by making living alone relatively more costly. This is despite the fact that long commutes hurt labor market prospects of married women, and couples lose more job access than singles.

First, using the geolocated PSID I identify patterns in the data that suggest that commuting plays a role in specialization within couples. I show that there is a large and robust difference in commuting between single men and men in couples beyond what can be accounted for by couples moving to the suburbs. This wide margin cannot be easily explained with usual approaches to modeling the costs of commuting or gender gaps in other labor market outcomes within couples. I argue that there is likely an aspect to commuting costs that rewards specialization on this margin within a household.

I propose a simple functional form that captures this intuition and when added to a standard collective labor supply model is capable of matching the large gap in commuting between men in couples and single men, as well as other salient features of the data. I embed this behavior in a quantitative spatial equilibrium model of a metro-area, contributing to the urban economics literature by seriously distinguishing between the incentives of couples and singles in this setting. Moreover, I overlay the spatial equilibrium structure with a simple marriage market clearing, endogenizing both the share of individuals choosing marriage and the distribution of resources between a husband and a wife. I show how increasing potential commutes, through lower connectivity between neighborhoods or suburban sprawl, affects behavior and welfare of singles and couples. While long potential commutes increase the gap in labor market outcomes between married men and women, residential sorting of singles to the city and the rent differential between suburbs and the city, they also make marriages more valuable.

As metro-areas grow out in space while jobs concentrate in central cities, average commutes increase. I argue that there is an aspect of commuting costs that creates a wedge between singles and couples. In section A.4 in the appendix I discuss two additional implications of this result. Recently, the COVID pandemic reinvigorated the discussion about the benefits of allowing employees to work from home. The results in this paper imply that while women in couples are most likely to be motivated to enter the labor force when more work from home options are available, it is singles who would benefit the most in terms of welfare (taking into account only the non-commuting aspect of working from home, without changing any benefits of the job). Second, in this model as in the data singles are more likely to live in the city, both because they value suburban amenities less and because they appreciate short commutes more. In recent decades we have seen a marked decline in the share of population getting married, especially through increasing the age at first marriage. I show that a natural implication of a decline in marriage is gentrification – a steepening of the distance price gradient in metropolitan areas. Both of these observations touch on timely topics in labor and housing economics and would be a fruitful direction of future research.

CHAPTER II

Did the Baby Boom Cause the US Divorce Boom?

Abstract

The United States experienced two major demographic 'booms' during the second half of the twentieth century, in births after the second world war and in divorces 25 years later. This paper argues that the two booms are linked. As the baby-boom generations were entering marriageable age, men in previous cohorts were faced with exceptionally good remarriage prospects motivating them to rematch. The cohorts who ultimately divorced most were the ones with the biggest increase in remarriage opportunities for men. Using cross-state variation in the size of the baby-boom, I show that marriages in the pre-boom generations were more likely to divorce the bigger the relative supply of young women. This conclusion is robust to instrumenting the size of the baby-boom with WWII mobilization rates. Lastly, I construct a simple dynamic marriage market model which can generate a divorce boom caused by a baby-boom, and can account for between a seventh and a third of the rise in divorces in the 1970s.

2.1 Introduction

Between 1960 and 1980 the divorce rate in the United states more than doubled. This was also the time when large post-World-War-II cohorts started entering the marriage market. In this paper I propose a causal link between the two concurrent events. I hypothesize that more marriages started to break up during this time, because the marriage market flooded with younger generations much bigger than the previous ones. As men typically match with younger women, a large young cohort presents a supply shock to the remarriage options of husbands in existing partnerships. At the same time sharp cohort size growth results in a shortage of eligible men, making it more worth while for young women to match with already married men.

Divorce is a large and common disruption in the lives of American families. By age 45 one third of first marriages end in divorce (Stevenson and Wolfers, 2007) and 4 out of 10 children in the U.S. will experience the divorce of their parents (Garcia-Moran, 2018). It represents a large income shock, especially for women and children.[12] In fact, Moffitt (1992) notes that most exits and entrances into welfare are associated with changes in family structure (not with changes in e.g. labor market circumstances) and Chetty et al. (2014) shows that the share of single-parent families is the strongest predictor of geographical variation in intergenerational mobility. Yet why people divorce is not very well understood.

This paper argues that a major reason for divorce is rematching of men to younger women (whom they prefer) which causes marriages among peers to break up. Survey evidence consistently shows that men prefer younger women, while women prefer men of their own age (or older).[3] Many features of existing marriage and divorce behavior are consistent with asymmetric age preferences causing divorce. Marriages are consistently forming more between younger women and older men, a tendency that is stronger for later-age marriages and for remarriages (England and McClintock, 2009). Divorced people rarely stay single. Stevenson and Wolfers (2007) show that by age 45, 69% of those who divorced have remarried. Men remarry more than women and soon after the divorce (Browning et al., 2014).

Showing that divorces are at least partially driven by preferences of men for younger women also sheds light on an important gender asymmetry in the lifecycle of marriage. In particular, if men prefer younger women while women prefer men of the same age, the remarriage prospects of the wife deteriorate faster relative to the husband as a couple ages together. Consistent with this hypothesis, Bruze et al. (2015) finds that as men and women get older, husbands receive a larger share of the marital surplus. This mechanism has consequences for the inherent stability of different marriage matchings, the importance of common property as an insurance mechanism (Lafortune and Low, 2017) and the welfare implications of policies concerning division of property at divorce, joint taxation, spousal

[1]Depew and Price (2018), Ananat and Michaels (2008) and others use the sex of the first child as an instrument for divorce to confirm this link is in fact causal.

[2]See Chiappori et al. (2018b) for a review of the evidence on negative effects of divorce on reported subjective wellbeing beyond the loss of income. See Boca (2003) for a review on the negative effects of divorce on children.

[3]See for example Bozon (1991), Kenrick et al. (1996). Kenrick and Keefe (1992) shows evidence of age preferences from dating newspaper advertisements. Recently Low (2023b) showed that this result holds even when controlling for physical appearance. These preferences are often rationalized with gender differences in fecundity by age (for example inGiolito (2004), Díaz-Giménez and Giolito (2013) and Low (2023a)). Shephard (2019) shows using an estimated a random search steady state matching model with limited commitment that the increase in utility for a man whose wife is around five years younger is equivalent to his private consumption being 50% higher.

income effects on social security and others.

In this paper I provide evidence that the entry of baby-boomers on the marriage market in the 1970s was at least partially responsible for the rise in divorces. First, in section 2.2 I show that on the aggregate the timing of the divorce boom and the incidence among cohorts matches up with this hypothesis. Divorce rates started increasing sharply when the cohorts who should be most affected by the entry of baby-boomers were in their prime divorcable age. Moreover, these cohorts also had the highest chances of ultimately having divorced. Section 2.3 presents cross-state evidence that the pre baby-boom cohorts divorced more by 1980 if born in a state with a bigger baby-boom between 1930 and 1950. I show that this correlation is robust to controlling for differences in a variety of socio-demographic characteristics. Lastly, I confirm that the estimated effect of a steep cohort size increase on divorces of preceding cohorts is in fact stronger if the size of the early baby-boom is instrumented with WWII mobilization rates (a strategy motivated by Doepke et al. (2013)). This provides further evidence that the cross-sectional correlation is causal and not driven by common state specific factors such as a decline in pro-family norms.

In section 2.4 I show that a simple repeated matching marriage market model, where divorce is driven by remarriage, can generate a divorce boom out of a baby-boom entry. Despite its simplicity, this approach is novel in capturing the proposed mechanism. Models of the marriage market currently used in the literature do not provide a useful starting point for this hypothesis, because they are either ignoring divorce and remarriage altogether, or divorce is not attributed to remarriage and rematching is only possible out of singlehood.[4] The within period matching market is based on Choo and Siow (2006).[5] The model matches the timing of the boom and the incidence across cohorts, but misses the size of the boom and the persistence of the divorce bust after 1990. Quantitatively, it explains between a fifth and a third of the divorce boom (depending on the metric chosen).

This paper contributes to the literature studying the effects of cohort size variation on the marriage markets. A long tradition in demography since Groves and Ogburn (1928), and more recently Abramitzky et al. (2011), has recognized that a variation in cohort size combined with a common age gap in marriages causes a 'marriage squeeze', affecting *marriage rates*.

[4]Notable exceptions of models that do allow for rematching 'on-the-job' (building on Mortensen (1988)) are Cornelius (2003) and Burdett et al. (2004), who work with stylized theoretical random search models with non-transferable utility. In spirit, the model closest to the hypothesis in this paper is Shephard (2019), which builds a random search model with rematching with age as a key variable of interest. It recognizes that divorce at older ages can happen, because men realize their preferred (younger) partners are now available. However, this paper still does not allow rematching straight from marriage (that would motivate divorce) and the model is too complex to study dynamics (the paper only studies a steady state).

[5]Choo and Siow (2007) and Choo (2015) also extend Choo and Siow (2006) to a dynamic setting. In many aspects these models are more complex, yet none of them allows divorce motivated by rematching (in fact they assume exogenous divorce).

Bergstrom and Lam (1991), Bergstrom and Lam (1994) and more recently Bhrolcháin (2001) show that the effects on marriage rates are usually largely mitigated by adjustments in the *age gap* itself. Brainerd (2017), and indirectly Bronson and Mazzocco (2018), provide evidence that *men investing* more or less in marriageable capital such as education and experience can also be a margin on which the marriage market adjusts to a cohort size variation. I propose that *divorce and remarriage* (serial monogamy) is an alternative mechanism through which the marriage market absorbs a large increase in cohort size.

This paper also adds to the literature trying to understand the rise in divorces in the US in the 1970s. Most existing explanations relate to the concurrent decrease in the wage-gap/increase in female labor supply (e.g. Ruggles (1997), Weiss and Willis (1997)).[6] McKinnish (2007) shows that sexually integrated workplaces cause divorce (through remarriage) and Greenwood et al. (2016) shows that a technological progress in the household sector can be behind both an increase in divorces and an increase in female labor supply. A large literature discusses whether the boom in divorces is caused by liberalized divorce laws in the early 1970s (see e.g. Friedberg (1998), Wolfers (2006)), though Stevenson and Wolfers (2007) concludes that despite apparent conflict in this literature, liberalized divorce laws had at most a small effect on divorce rates. In a way, this paper suggests that divorce laws might have been liberalized because of a mounting demand for divorces as the baby-boomers were entering marriageable age. Croix and Mariani (2015) hypothesizes that throughout history marriage norms have been changing to accommodate demand for polygyny in otherwise monogamous societies. This theory is consistent with the divorce laws changing in the 1970s.[7].

2.2 Aggregate evidence

The United States experienced two major demographic 'booms' during the second half of the twentieth century. Between 1945 and 1960 the number of births increased sharply (see figure 2.1b). Approximately 25 years later divorces started to rise, peaking around 1980 (see figure 2.1a). This paper argues that the two booms are linked. As the baby-boom generations were entering marriageable age, men in previous cohorts were faced with exceptionally good remarriage prospects. This hypothesis requires that the divorce boom must have a strong cohort component, because the entry of baby-boom generations is expected to affect the cohorts immediately preceding more than others. Indeed, figure 2.2a shows that there is a boom and bust also in the share of men ever-divorced by cohort. Men born between 1930 and

[6]However, it is also natural to suspect reverse causality from divorce to female labor supply, as showed in Johnson and Skinner (1986) and most recently by Goldin and Katz (2016).

[7]Other potential suspects include social attitudes towards marriage (Cherlin, 2004) and generosity of the welfare state (Moffitt, 1997)

Figure 2.1: Divorce boom and Baby-boom.

(a) Divorce boom (b) Baby-boom 25 years earlier

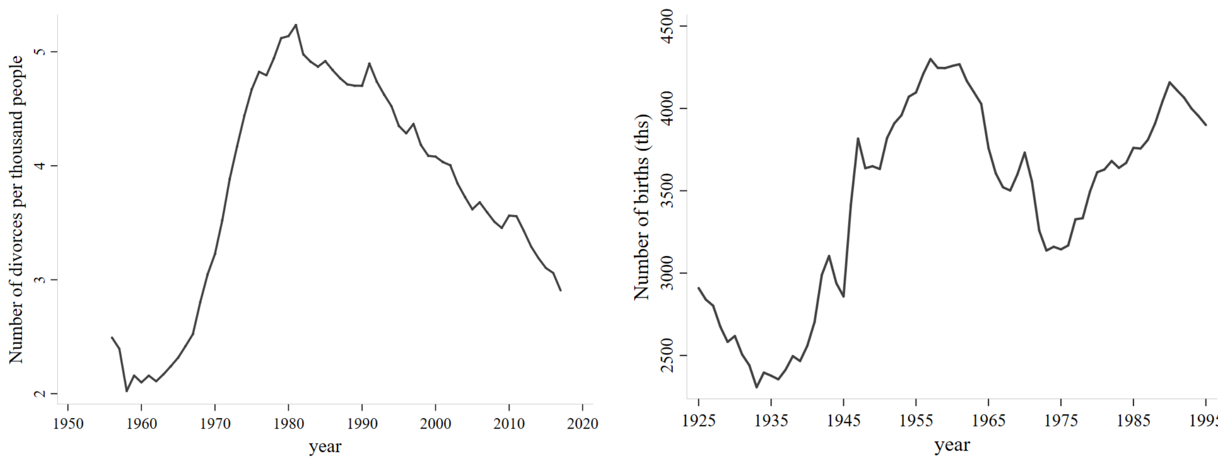

Source: Vital Statistics (divorce rate in panel 2.1a is based on a subsample of states with coverage over the whole sample).

1940 were the ones for whom the chances of divorce increased sharply. The cohorts of men born between 1940 and 1950 had the highest chance of ever experiencing a divorce. By the time the divorce rate started increasing, between the years 1960 and 1980, these men were entering their 30s, which is a common age to get divorced.[8] Figure 2.2a shows that the peak of the divorce boom coincides with the years when cohorts who divorced most were 34 years old. This was also the time when the age distribution in the marriage market relevant for men in their 30s shifted dramatically. Figure 2.2b shows a sharp increase between 1960 and 1980 in the number of people older than 34 compared to the number of people younger than 34 (by up to 6 and 12 years respectively), peaking around 1980. This means that by 1980 the prospects for a 34 year old man to remarry to a younger woman were especially high as there were many more younger women compared to the number of men who would be competing with him. Subtracting the lag of 35 years in figure 2.2b implies that the remarriage prospects were most favorable for men born between 1940 and 1950, exactly the cohorts who divorced the most.[9]

[8] See appendix B.1.2.

[9] The divorce rates are also high for the cohorts 1950-1955, even though the age distribution has already been stabilizing. It is very possible that this persistence in divorce risk is still endogenous to the mechanism discussed here. Early baby-boom cohorts were increasingly entering in marriages with a lower age-gap (as shown in figure B.13 and predicted by Bergstrom and Lam (1991)) and marriages to divorcees, precisely because young women were in over-supply. These two characteristics correlate with a higher risk of divorce. This can explain why divorce rates remained high for the baby-boomers, despite falling remarriage prospects. Section B.4 in the appendix discusses this hypothesis in more detail. It is also possible that after divorce rates increase, the fundamentals of the marriage market change. For example, Chiappori and Weiss (2003)

Divorce boom across cohorts

(a) Divorce rate (number of divorces a year per thousand people) (left), Share men ever divorced (right)

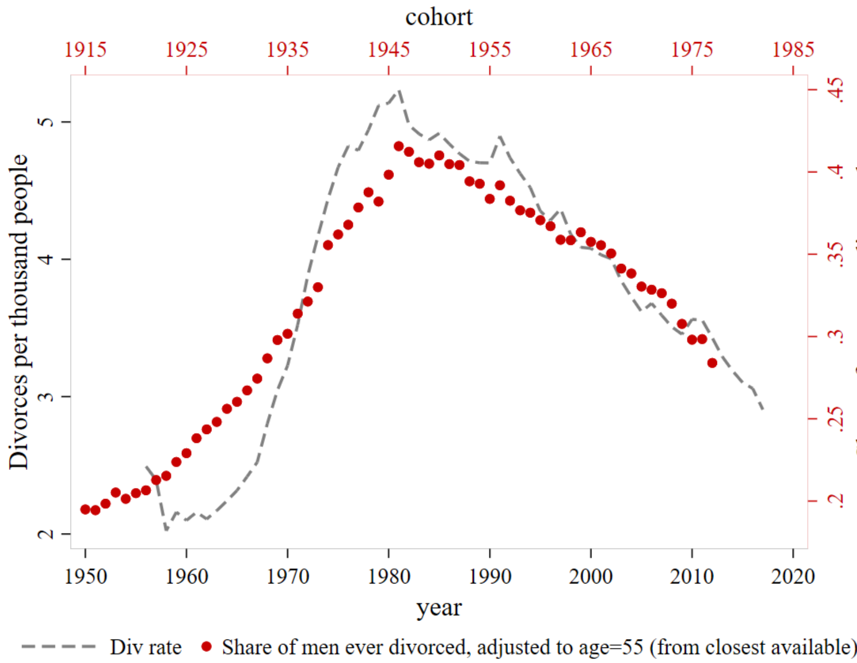

(b) Size of the new entry shock.

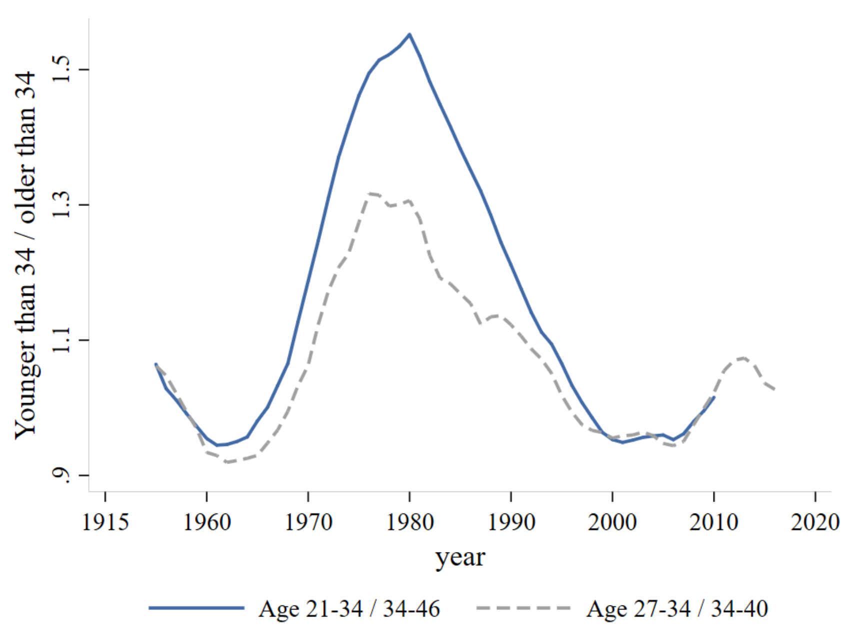

Figure 2.2a: the divorce rate (on the left axis) by year and the share of men in the cohort born 35 years ago who ever went through a divorce (or remarriage), measured at age 50 or closest available and regression adjusted to age 50 (see section in the appendix). Figure 2.2b: change in the supply of younger women to older men, measured by the 0 to 6/12 years younger over 0 to 6/12 years older ratio by age 34 (and so for cohorts born 35 ago). This kind of measure is always high whenever the cohort size is rapidly increasing and is equal to 1 in a stable population. Source: IPUMS Census, Vital Statistics.

Lastly, table 2.1 shows that the age-gap in men's higher order marriage is systematically bigger than in man's first marriages by about 3 years, confirming the pattern that after divorce men typically remarry to younger partners.

Table 2.1: Age gaps in first and higher order marriages

	Mean	sd	mean - sd	mean + sd
First marriages, all	2.92	4.66	-1.74	7.58
First marriages, ages 35-44	2.49	4.12	-1.62	6.61
Higher order marriages, all	6.00	8.00	-1.51	13.89
Higher order marriages, ages 35-44	3.85	6.16	-2.31	10.01

Sample: men. Source: 1960 Census

Overall this shows that the timing of the divorce boom and the incidence across cohorts is consistent with the hypothesis that the baby-boom caused, at least partially, the divorce boom. Section B.1.3 in the appendix shows additional descriptive evidence on the adjustment of the marriage market to the baby-boom. First, the age-gap in marriage was higher and rose more quickly over their lifetime for the men born around 1940 compared to men in other cohorts. These men also stayed single slightly less, almost closing the gender gap in ever getting married. Lastly, compared to men, women born around 1940 who ever went through a divorce remarried less and stayed divorced more.

2.3 Leveraging cross-state variation in the baby-boom

In this section I present evidence that being from a state with a large early baby-boom (large post 1945 cohorts compared to 1935-1944) caused the treated cohorts (of 1935-1944) to divorce more. I focus on this group for several reasons. These were the cohorts who were most likely responsible for the sharp run up of the divorce boom. They were in their prime divorce age in the 1970s[10] and they ultimately ended up having a 10 percentage points higher chance of ever getting divorced than cohorts 10 years older (the sharpest increase compared to other cohort groups: see figure 2.2a). Second, by 1980 they were 35-44 years old, which is old enough to limit the worry of selection into marriage, but young enough for differential

and Chiappori and Weiss (2007) suggest that a small initial increase in divorce rates can be multiplied and turn permanent, because the remarriage market is now more favorable to everybody. This mechanism is requiring that single divorcees are more available to find a match than married individuals, which I am not able to capture in my simple model.

[10]See Section B.1.2 in the appendix.

mortality not to matter. Since remarriages are not observed in the Census for 1990 and 2000, younger cohorts are too young to be studied by 1980 (as many men have not married yet, selection into marriage makes comparison across cohorts difficult to interpret). Third, building on Doepke et al. (2013) I will argue that the relative supply of younger partners available to these cohorts (the size of the early baby-boom of 1945-1954 compared to the pre-baby-boom cohorts of 1935-194) can be instrumented by WWII mobilization rates. This allows me to rule out that the correlation between divorce probabilities and remarriage supply is caused by confounding factors such as persistent cultural norm differences across states towards high fertility and few divorces.

To study the share of a given cohort who divorces by a certain age, I primarily use pooled IPUMS Census data for 1960 and 1980 (Ruggles et al., 2019), focusing on men and women age 35 to 44 (by 1980, the cohorts born 1935-1944 reached this age, cohorts born 1915-1924 who were 35-44 years old in 1960 are used as a control). I restrict the analysis to those born in the United States (as immigrants tend to not fully integrate into local marriage markets).[11] I also exclude those living in group quarters for the same reason.[12] A person is classified as being ever-divorced, if their current marital status is divorced or if their last marriage was not their first.[13] To supplement the main sample, I use IPUMS Census data to construct a state level measure of remarriage opportunities for the cohorts of interest. Specifically, I compute the share of people aged 30-34 in the group 30-44, $\frac{n_{30-34}}{n_{30-44}}_{st}$, in each state and year. Between 1960 and 1980 this measure captures the variation in the steepness of the early baby-boom (of cohorts born 1945-1949). As a robustness check, I repeat the analysis using a more direct measure of this cohort size increase, the share of children 0-4 among children 0-14 30 years ago, $\frac{n_{0-4}}{n_{0-14}}_{s,t-30}$ (using IPUMS Census data from 1930 and 1950). The results are presented in section B.2.1 in the appendix.

The main hypothesis is that people in their late 30s born in states with a steep improvement in remarriage opportunities for men had a higher chance of ever getting divorced. Cross-state differences in fertility can arise for a multitude of reasons, some of which could be persistent and correlate with the likelihood of divorce (violating the exogeneity of the remarriage opportunity measure). For example, cultural norm differences across states towards high

[11]A fact explored in the marriage literature when variation in migration rates can create immigrant specific variation in sex ratios, as in Angrist (2002).

[12]Lastly, people born in Hawaii, Alaska and DC are also excluded from the analysis, as data on WWII mobilization rates are not available for these states.

[13]This information is available in the 1960-1980 Census and the 2008 onwards ACS. Notice it would be unreasonable to use simply the marital status of being divorced, because the main hypothesis is that divorces happened because of a desire to remarry. Unfortunately, this variable is an imperfect proxy, as it also includes people who remarried after their spouse died. This is another reason to use a relatively young age group, for which mortality is still low.

fertility and few divorces would bias the OLS results downwards. To rule out that the results are driven by this kind of omitted variable bias, I instrument the change in the remarriage opportunity measure between 1960 and 1980 with WWII mobilization rates (as used by Acemoglu et al. (2004)). This strategy is based on the evidence in Doepke et al. (2013) suggesting that WWII mobilization was at least partially responsible for the baby-boom. The hypothesized mechanism goes as follow: women (mainly in cohorts 1905-1914, but more broadly in 1905-1924) worked during the war and beyond, this labor supply 'shock' depressed wages of women (being an imperfect substitute for men in the labor market), incentivising the cohorts of 1925-1934 (who were 15-24 years old by 1950) to stay at home and increase fertility (realized between 1945 and 1960).[14][15]

Figure 2.3a displays a cross plot of state level mobilization rates on the x-axis with the change in the remarriage opportunity measure on the y-axis, showing a strong positive correlation (confirming that the fertility variation explained by Doepke et al. (2013) did result in cross-state variation in relative cohort size 30 years later). Second, figure 2.3b shows a positive correlation between the change in the remarriage opportunity measure with the growth in divorce (measured by a change in the share ever-divorced in the relevant age range).

Next, I confirm that this correlation is statistically significant in a regression, robust to controlling for observable differences among states and robust to instrumenting change in the remarriage opportunity measure with mobilization rates during WWII. The main specification (equation 2.3.1) regresses a dummy variable y_{ist} of being ever-divorced (where i stands for an individual, t for a year $\in \{1960, 1980\}$ and s for a state of birth)[16] on state and year fixed effects, the remarriage opportunity measure $\frac{n_{30-34}}{n_{30-44}}_{st}$, and a set of individual level sociodemographic controls. The hypothesis is that people born in states with a bigger increase in remarriage opportunities were more likely to get a divorce ($\beta > 0$).

$$y_{ist} = \alpha_s + \alpha_{1980} + \beta \frac{n_{30-34}}{n_{30-44}}_{st} + \gamma X_{ist} + \epsilon_{ist} \tag{2.3.1}$$

[14]Notice this mechanism should directly affect the cohorts 1915-1924, which serve as a control in the baseline analysis. Specifically the mechanism predicts that these cohorts have a higher labor supply in high mobilization states. Since female labor supply is, if anything, associated with higher chances of divorce, this potential source of invalidity of the exclusion restriction should bias against finding an effect when comparing the treated cohorts to the control. Section B.2 presents a further discussion of this concern.

[15]Section B.1.4 in the appendix provides more detail on the variation in mobilization rates.

[16]I classify people based on their state of birth to avoid spurious correlation in divorcing behavior with internal migration decisions, for marriage or after divorce. Moreover, it is not clear whether the relevant marriage market pool is closer to the state of birth or the state of residence (as individuals can for example temporarily move for a job or to get a degree and then move back 'home' and settle down).

As a consequence, the estimates are in spirit closer to measuring an intent-to-treat effect. When state level remarriage opportunities are assigned based on the state of residence, the results remain qualitatively similar and statistically significant, yet slightly smaller in magnitude and much noisier. This suggests that internal migration decisions mainly add noise to the analysis.

Figure 2.3: Cross-state variation

(a) Mobilization rates affecting the size of the baby-boom.

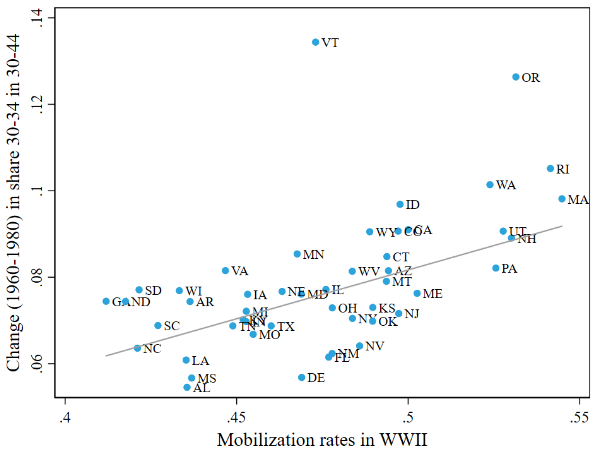

(b) Remarriage opportunities for men affecting divorce.

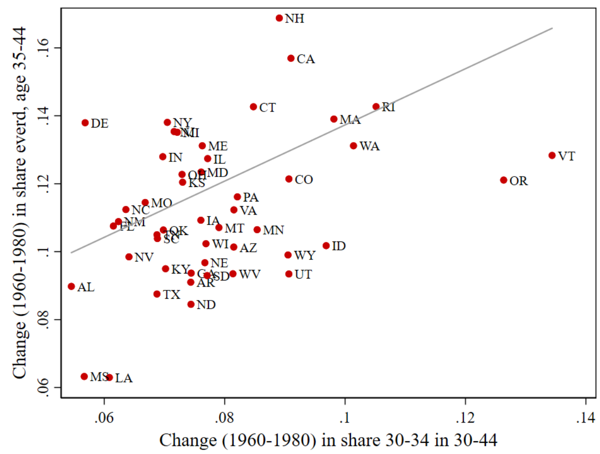

Share of ever-divorced attached by state of birth. Data source: IPUMS Census data from 1960 and 1980, state-level mobilization rates during WWII from Acemoglu et al. (2004). States weighted by population in 1930.

Table 2.2 presents the baseline results. All specifications include dummies for age, sex and race. Columns 1 and 2 present the OLS results, confirming what figure 2.3b has suggested. Individuals born in states with higher remarriage prospects are more likely to go through a divorce (compared to their counterparts 20 years ago). Column 2 adds additional controls, namely dummy variables for levels of education, whether an individual lives in a metro area and farm status. These controls are motivated by baseline cross-state differences (already observed by 1940) between states with high and low mobilization rates (see section B.1.4 for details). If anything, including these controls makes the OLS results slightly bigger in magnitude.

Columns 3 and 4 mirror columns 1 and 2, but instrument $\frac{n_{30-34}}{n_{30-44}}_{st}$ with WWII mobilization rates interacted with a dummy for 1980. The specification of the first stage is directly motivated by Doepke et al. (2013). Results of the first stage regressions are presented in table B.4 in the appendix. Table 2.2 shows that instrumenting the remarriage opportunity measure only strengthens the results. This suggests that perhaps the correlation between remarriage opportunity and divorce chances is mitigated by omitted variables such as persistent trends in pro-family values. Column 5 shows that the result is robust to adding region times year fixed effects (being identified from within region differences in the growth of remarriage options and mobilization rates). Column 6 shows that the main conclusion is robust to only restricting to people who ever got married, so it is not driven by selection into marriage.

Overall, this section presents robust evidence of the mechanism proposed in this paper.

Table 2.2: Increase in the share ever-divorced when remarriage options of men improve

	Ever-divorced					
	OLS		2SLS			
$\frac{n_{30-34}}{n_{30-44}}_{st}$	0.809	0.900	1.485	2.065	1.885	2.011
	(0.238)	(0.296)	(0.387)	(0.457)	(0.429)	(0.444)
X_i:						
sex, race	yes	yes	yes	yes	yes	yes
farm, metro	no	yes	no	yes	yes	yes
educ dummies	no	yes	no	yes	yes	yes
region-year fes	no	no	no	no	yes	no
In ever-married	no	no	no	no	no	yes

N = 2032220 (1902899 in column 5), 48 clusters

SEs statistics in parentheses. SEs clustered at the state level.

All regressions include year, age and state fixed effects.

Pooled cross-section 1960 and 1980 IPUMS Census data, fitting a linear probability model of being ever-divorced and divorced on $\frac{n_{30-34}}{n_{30-44}}_{st}$, a share of 30-34 year olds in 30-44 year olds (aggregated on the state level). Restricting to a sample of ever-married men and women 35-44 years old. All columns include year and state fixed-effects and demographic controls. In columns 3-6 $\frac{n_{30-34}}{n_{30-44}}_{st}$ is instrumented by WWII mobilization rates interacted with an indicator for 1980.

Married men surrounded by an increased opportunity to rematch with a younger spouse were more likely to divorce. On the aggregate $\frac{n_{30-34}}{n_{30-44}}_t$ rose from 0.329 in 1960 to .407 in 1980 a difference of 0.077. Thus, the coefficients presented in table 2.2 suggest the change on the aggregate should imply an increase of between 7 and 15 percentage points in the probability of ever divorcing for 35-44 years old. In the data this measure rose on the aggregate by 12 percentage points. The cross-sectional variation can explain the rise in divorces in this age-group.

2.4 Can a steep cohort size growth cause a boom in rematching?

In this section, I present a simple repeated matching marriage market model calibrated to the pre-divorce-boom marriage market characteristics to illustrate that a baby-boom naturally generates a divorce boom, through the simple channel of men rematching with younger women more. The model is set up with the following strategy in mind: I make assumptions necessary to make the decision about marriage effectively static and happening repeatedly in a competitive matching environment based on Choo and Siow (2006), while keeping initial divorce rates sufficiently low. At the same time, aggregate supply of men and women of different ages is dynamic and exogenous (taken to match the data). Overall, this model can generate a divorce boom (and explain between a third and a seventh of the rise in divorce), despite being very simple and focusing solely on divorce motivated by rematching keeping initial motivations for matching by age fixed.

The setup of the model is as follows. Both men and women are only differentiated by age: $i, j \in \{21, ..., 69\}$. Every period starts with the distribution of existing marriages N_{ij} (number of marriages among men of age i and women of age j) and singles $N_{i\emptyset}$, $N_{\emptyset,j}$. The length of a period is 4 years. Cohorts enter the marriage market in a scattered pattern. Women enter first (a share α_f of women enters at age 21, the rest at age 25). Men enter later (a share α_m of men enters at age 25, the rest at age 29). The main motivation for the staggered entry is to limit initial unnecessary divorces as much as possible. The initial entry in the marriage market is also consistent with the intuition that different people become 'ready for marriage' at different ages. Once a person joins the market, they continue to participate each period, until they die or are widowed. Marriages survive if both partners choose to match with their existing type. Otherwise, a divorce occurs. At the end of a period, the oldest cohort dies. At the start of next period, a new cohort enters the system at age 21.

When a man of type i arrives on the marriage market, he chooses a woman of any age or to stay single, to maximize his lifetime utility from marriage according to the definition in 2.4.1.

$$\begin{aligned}
&\forall i \in 1, ..., I \\
&V(i, \epsilon_p) = \max_{j \in 1,2,...,J,\emptyset} v_{ij} - \tau_{ijt} + \epsilon_{ijp} \\
&\qquad\qquad + \ \beta V(i+1, \epsilon_p)
\end{aligned} \tag{2.4.1}$$

v_{ij} are exogenous systematic gains from a match ij for a man of age i who matches with a woman of age j (or no woman $\emptyset$), respectively. When being born every person p is endowed with a list of age-gap specific preferences $\epsilon_p = (\epsilon_{ijp})_{\{i \in 1,...,I; j \in \emptyset, 1,...,J\}}$ for single-hood and all possible partner ages, given own age. $\tau_t^i = (\tau_{i1t}, ..., \tau_{iJt})$ is a vector of one time endogenous transfers a man of type i would have to pay when matching with a type j. Cohort size variation influences the individual choice through affecting transfers in a standard general equilibrium sense. The optimal choice of a woman is defined equivalently (2.4.2).

$$\begin{aligned} &\forall j \in 1, ..., I \\ &U(j, \epsilon_{p'}) = \max_{i \in 1,2,...,I,\emptyset} u_{ij} + \tau_{ijt} + \epsilon_{ijp} \\ &\qquad\qquad + \ \beta U(j+1, \epsilon_{p'}) \end{aligned} \tag{2.4.2}$$

Since the continuation value does not depend explicitly on the current match, the matching decision as effectively static. This is driven by the assumption that a person enters the market again next period, regardless of what they did this period.

This simple assumption has one exception. To avoid a large surge of 'grey' divorces, I remove widows and widowers from the marriage market. If widows return to the market after their husbands die, they create a sudden increase in the supply of older women, changing the incentives for everybody to match and destabilizing many current marriages. If widowers stay of the market (and preferences for age gaps are stable), existing matches are stable at older ages when new entrants to the market are mostly too young to be relevant rematching partners. However, to preserve the static nature of the matching decision, I assume that this drop-off from the market comes as a surprise to widows and widowers (after their partner dies, they suddenly realize they do not wish to participate in the marriage market anymore). This way the decision to marry a partner of age $I-1$ is not affected by the future market prospects.[17]

Transfers are determined endogenously in an equilibrium defined as follows:

Definition 2. *Given a sequence of cohort sizes* $\{N_t\}$ *a dynamic equilibrium in the marriage market consists of sequences of transfers* $\{(\tau_{i,j,t})_{\{i=1,...,I;j=1,...,J\}}\}$, *numbers of couples of each type* $\{(n_{i,j,t})_{\{i=1,...,I;j=1,...,J\}}\}$ *and numbers of single men* $\{(n_{i,\emptyset,t})_{\{i=1,...,I\}}\}$ *and women* $\{(n_{\emptyset,j,t})_{\{j=1,...,J\}}\}$ *such that*

[17] This assumption is fairly innocuous in steady state with the parametrization bellow. It is equivalent to assuming that widows get a choice between enjoying the period benefit from their previous marriage (perhaps a warm glow from the correct match) or returning to the market. Since transfers are stable in the age gap in steady state, widows would choose to not join the market and rather enjoy the period benefit from their optimal choice. However, in the dynamic marriage market after a cohort-size shock, abandoning the assumption would make the equilibrium challenging to solve for.

1. *Given transfers every period the choice of a partner $j^*(i, \epsilon_p)$ by a man who enters the marriage market solves the maximization problem as defined in 2.4.1.*
   The choice of a partner $i^*(j, \epsilon'_p)$ by a woman who enters the marriage market solves the maximization problem as defined in 2.4.2.

2. *Markets clear, such that*

$$\forall i \in 1, ..., I; j \in 1, ..., I \;\; n_{i,j,t} = \int 1_{j^*(i,\epsilon_p)=j, p \text{ on the market}} \, dp = \int 1_{i^*(j,\epsilon'_p), p' \text{ on the market}} \, dp'$$

$$\forall i \in 1, ..., I; \;\; n_{i,\emptyset,t} + \sum_{j=1}^{J} n_{i,j,t} = m_{i,t}$$

$$\forall j \in 1, ..., J; \;\; n_{\emptyset,j,t} + \sum_{i=1}^{I} n_{i,j,t} = f_{j,t}$$

   where $m_{1,t} = 0$, $f_{1,t} = \lambda_f N_t$, $m_{2,t} = \lambda_m N_{t-1}$,
   $m_{it} = N_{t-i+1} - W^m_{i,t}$, $f_{j,t} = N_{t-j+1} - W^f_{j,t}$ $\forall i > 2, j > 1$.
   with $W^m_{i,t} = \sum_{s=1}^{i-1} n_{i-s,I,t-s}$ *and* $W^f_{j,t} = \sum_{s=1}^{i-1} n_{I,j-s,t-s}$

Lastly, to describe the solution I need to put distributional assumptions on the idiosyncratic preferences each period.

Assumption 1. *Every person when born is endowed with a vector of preferences for being single versus being matched with all possible age-gaps, that are independent within and across individuals, and are distributed identically standard extreme value type I:*

$$(\epsilon_{\emptyset,p}, \epsilon_{a(1)-a(I),p}, ..., \epsilon_{0,p}, ..., \epsilon_{a(I)-a(1),p})$$

Every period these are used to define the relevant vector of preferences: $\epsilon_{ijp} = \epsilon_{a(i)-a(j),p}$ *and* $\epsilon_{i\emptyset,p} = \epsilon_{\emptyset,p}$

Assumption 1 has two notable parts. First, I introduce persistence in preferences. Namely, the idiosyncratic preference for being single is constant for an individual over their life-cycle, and the idiosyncratic preferences defined over age-gaps, instead of a specific age, are constant. This is done to so that people who match with someone because they have a high idiosyncratic error are incentivized to stick with this match. Second I use a convenient distribution that will allow for an approximate closed form solution of period specific matches $n_{i,j,t}$.

Assumption 2. *Entry of men at age $i = 2$ is determined by $F(\epsilon_{8,p}) \leq \alpha_m$. Entry of women at age $i = 1$ is determined by $F(\epsilon_{0,p'}) \leq \alpha_f$.*

Assumption 2 states that for both men and women early entrants on the marriage market are selected not to have a high idiosyncratic preference for the kind of match that will be only possible to realize one period after their entry. This assumption, as well as the staggered timing of entry, is designed to minimize unnecessary divorces as much as possible (as with flexible rematching the biggest challenge in fitting the initial marriage patterns is to match a low initial rate of divorce).

The number of divorces each period is given by

$$D_t = \sum_{i \in \{1,...,I-1\}} \sum_{j \in \{1,...,I-1\}} \lambda \; n_{ij,t-1}(1 - P(j' = j+1|i+1)P(i' = i+1|j+1)) \qquad (2.4.3)$$

The probability that an existing couple who got to try rematching is not divorced is equal to the chances that both the husband and the wife choose to stick with their age-gap $P(j' = j+1|i+1)P(i' = i+1|j+1)$. I determine these by numerically integrating the decisions of a simulated population, given equilibrium transfers τ_{ijt}.

Table 2.3 shows the calibration of the model. Value of singlehood is normalized to 0. I

Table 2.3: Calibrated parameters

ages	$\{21, 25, ..., 69\}$ (I=13)	C	1.25	$\lambda_m = \lambda_f =$	0.75
$u_{\emptyset,j} = v_{i,\emptyset} =$	0	c	0.3791	$u_{ij} =$	C_f
$v_{ij} =$	$C_m - c(\lvert a(i) - a(j) - gap \rvert)^{1+\sigma}$	*(ideal) gap*	4	σ	0.1086

normalize v_{ji} to a constant (same as the peak value of u_{ij}).[18] Systematic gains to marriage for men u_{ij} are assumed to be constant in $a(i) - a(j)$ (maximizing stability of existing matches). Parameters c, gap and σ govern the nature and strength of the preference of men for somewhat younger women. I select them to match as closely as possible the matching patterns by age in 1960.[19]. C is selected to match the share of people ever married in the 1960 Census while keeping share ever-divorced at bay. λ_m and λ_f are selected to minimize the share of people ever-divorced in the steady state. Figure 2.4 plots $\frac{u_{ij}+v_{ij}}{2}$ for selected ages of men and all ages of women.

Table 2.4a summarizes characteristics of the marriage market in the steady state. Overall marriage rates are around 90%, consistent with the data pre-divorce-boom. Women marry

[18] u_{ij} and v_{ji} is not separately identified.

[19] I use the number of matches in the 1960 Census in each cross-age group for men 36-55 and women 31-50 years old, where matching is likely to be settled and everybody is reasonably thought of as participating in the marriage market. I use the result from corollary 1 (taken from Choo and Siow (2006)) and loop though parameters to minimaze the sum of squares between the right and left hand sides, restricting *gap* to be an integer.

Figure 2.4: $\frac{u_{ij}+v_{ij}}{2}$ for selected ages of men and all ages of women.

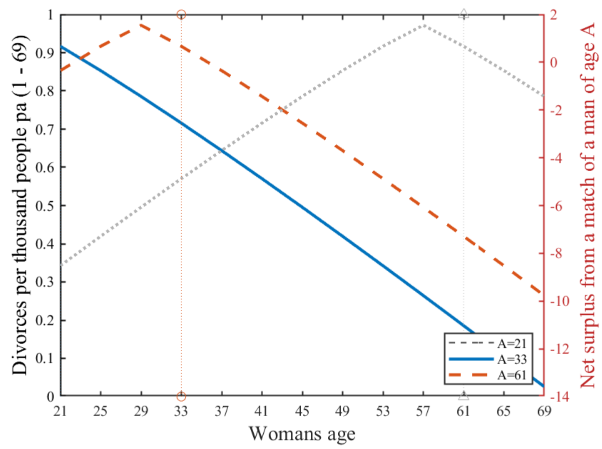

on average for the first time at 23 (very close to 21, the average age at first marriage for women who married around 1950 (Rotz, 2016)). Men marry later (as they do in the data), though the average age is slightly higher than in the data (before 1950 the average was approximately 25 (Rotz, 2016)). The share of men and women who ever divorced, and the age at first divorce, are slightly higher in the model than in the data before the boom while the aggregate divorce rate is slightly lower. This suggests that the model underestimates the share of the population that divorces many times. Overall, the steady state marriage market in the model is a good approximation of the marriage matching by age in 1960.

Starting from a steady state (a constant cohort size), I study the behavior of the model when new cohorts are suddenly bigger. First, I study the reaction to a one time cohort size increase. Figure 2.5 shows the effects on divorcing behavior. Already in the period when the first part of the bigger cohort enters, divorces (and divorce rates) increase. This is purely an effect of rematching of existing couples, as the youngest cohort certainly could not have divorced yet. The spike in divorces continues in the second period, being a combination of the new entrants breaking up couples formed before period 0, and rematching of couples formed in period 0 (who actually experience a higher chance of divorcing then their steady state counterparts). Figure 2.5b shows this explicitly, decomposing the increase in divorces into an effect of a divorce risk and a pure compositional effect of an increase in the number of couples at risk. First, the figure shows the evolution of a sum of probabilities of divorce for couples on the market, $(1-P(j_t = j|i, j_{t-1} = j-1)P(i_t = i|j, i_{t-1} = i-1))$, weighted by the steady state distribution of couples who could have gotten divorced $\{N_{i-1,j-1}\}$. This weighted probability of rematching, summarizing the average risk of divorce for existing couples, increases in year 0 and remains high in year 4. Second, figure 2.5b plots what divorces would have been if in every period the number of divorces was calculated with the steady state probabilities

Table 2.4: Baseline characteristics of the model compared to the pre-baby-boom marriage market.

(a) Baseline (steady state) characteristics of the model

Statistic	Men	Women
Share ever-married	91.7%	91.6%
Share ever-divorced	25.1%	25.1%
Mean age-gap	3.8	
Sd of age-gap	4.8	
Age at 1st marriage	26.2	22.2
Age at 1st divorce	33.2	29.1
Divorces per thousand people	2.2	

(b) Data

Statistic	Men	Women	Source
Share ever-married	93.2%	92.7%	*age 60-69 in 1950 Census*
Share ever-divorced	20.5%	18.4%	*age 60-69 in 1960 Census**
Mean age-gap	3.7		*all married people in 1950 Census*
Sd of age-gap	5.4		*all married people in 1950 Census*
Age at 1st marriage	25.1	22.1	*all ever-married people from 1960 Census**
Age at 1st divorce	34.4	29.2	*women age 40-45, cohort 1920-1934, NSFG*
		~ 34	*women 50-74, SIPP (Goldin and Katz, 2016)*
Divorces per thousand people	2.3		*Vital statistics, average 1956-1960*

* Not available in 1950 Census.

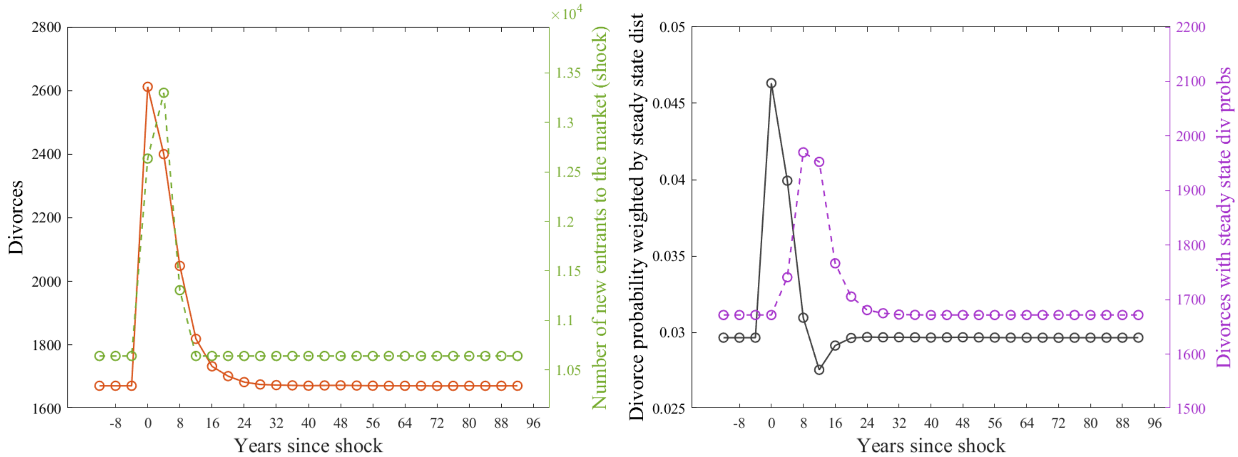

(a) Divorces vs. New entrants

(b) Decomposition of the increase in divorces

Figure 2.5: Impulse response to a one time increase in cohort size by 50%.

$(1 - P(j' = j|i)P(i' = i|j)$. Overall, the increase in divorces caused by a larger number of couples at risk is negligible until 8 years after the initial shock, contributing very little to the overall spike in divorces.

Figures 2.6 and 2.7 show the implications of the model for divorces when the cohort size dynamics are calibrated to match the baby-boom, compared with the observed divorce boom in the data. The solid line in figure 2.6a shows the raw number of births per year in the data aggregated to three years, and replaced by a constant for cohorts before 1928 and after 2009 (to start the model in a steady state)[20]. This represents the exogenous cohort size series as fed into the model.

The model does generate a substantial boom in the divorce rate (calculated as the number of divorces in a period divided by the number of people, including children who have not reached marriageable age yet, and divided by the length of the period in years, to make it comparable with the empirical divorce rate calculated each year). Figure 2.6b shows that the divorce rate in the model rises from 1.8 in 1960 to a maximum of 2.8 in 1988.[21] This shows the model is roughly successful in matching the timing of the divorce boom. Quantitatively, actual divorce rate increased by approximately 150% (3 in absolute terms) in the data and by only 50% (1 in absolute terms) in the model, suggesting this mechanism is only partially responsible for the divorce boom, unless other amplification mechanisms were at play.

Figure 2.7 plots the share of men ever divorced by cohort (both in the model and in

[20]Starting from a level of 2.66 million births per year and finishing at 4 million births per year (which roughly matches the average of a couple of years before 1928 and after 2009)

[21]Notice the divorce rate does not start in steady state, because the upcoming cohort size dynamics affect the denominator.

Figure 2.6: Divorce boom vs baby-boom 25 years ago in the model vs in the data

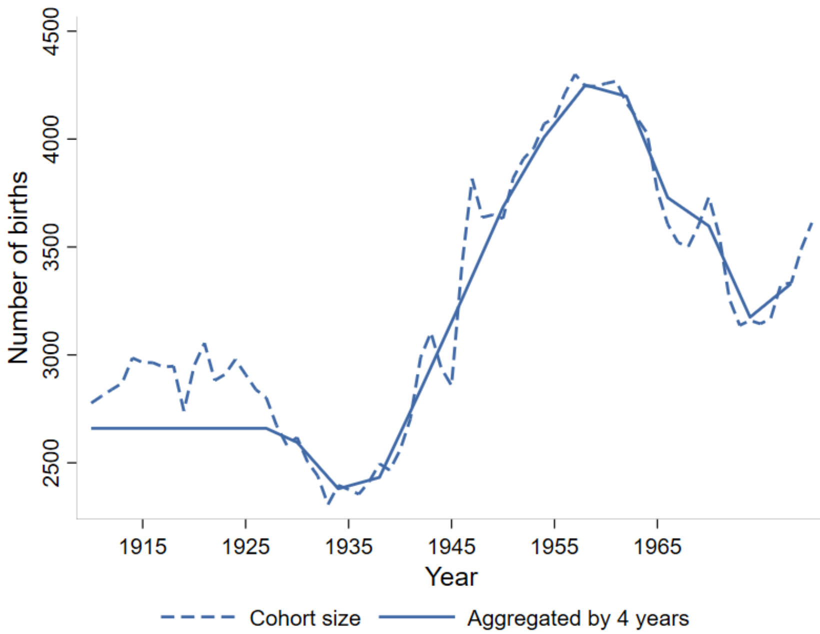

(a) Births (data vs model)

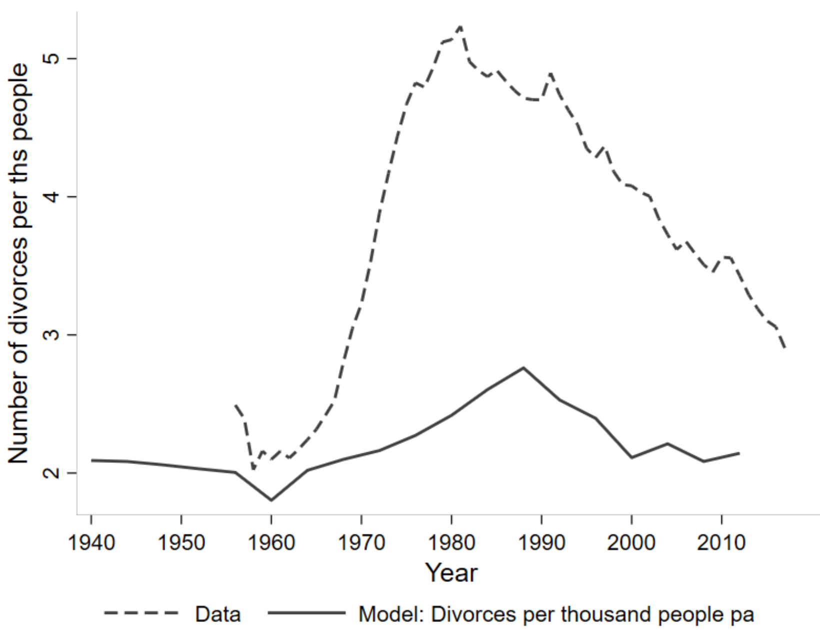

(b) Divorce rate (data vs model)

Figure 2.7: Divorce boom across cohorts in the model vs in the data.

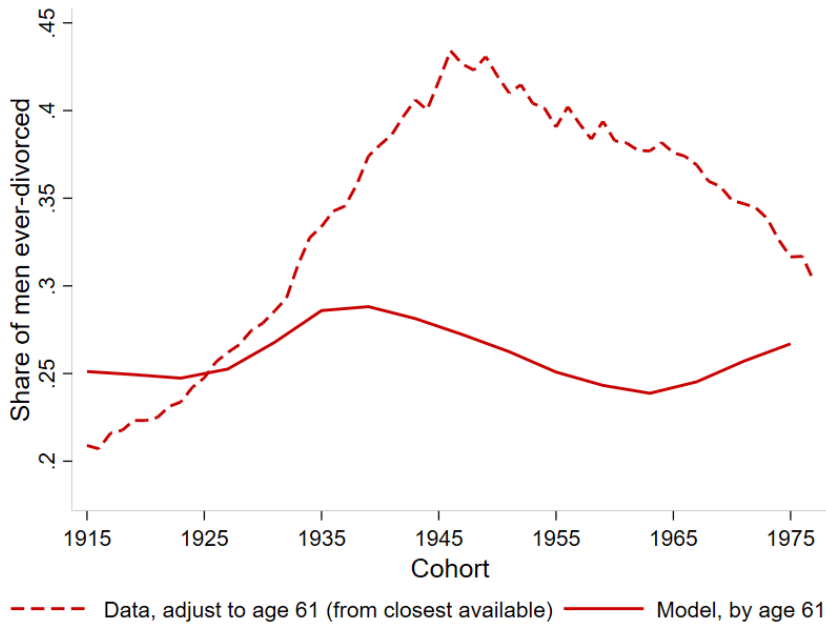

(a) Model

the data). The model does predict that the divorce boom has a strong cohort component. The cohort measure in the model peaks for men born in 1939, which represents the cohorts responsible for the beginning of the divorce boom. Similar to the aggregate divorce rate, the bust in the cohort measure is faster in the model than in the data. Moreover, the share of men ever divorced starts increasing for cohorts who are in their 30s after 2000 (born in late 1960s), which does not happen in the data. Quantitatively, the share of men ever divorced increased by 100% (0.2 in absolute terms) in the data and by only 15% (0.037 in absolute terms) in the model. Overall, the model matches the timing of the start of the divorce boom very well. However, the size of the boom in the data is much larger and the divorce bust after 1980 is less rapid.

Quantitatively, the model explains between a seventh and a third of the divorce boom (depending on the metric chosen). Together, this is consistent with the hypothesis that the proposed mechanism has likely played a substantial role in the run up of the divorce boom in the 1970s. It also suggests that the model is possibly missing a strengthening and persistence generating mechanism. As suggested by Chiappori and Weiss (2003) and Chiappori and Weiss (2007), even a small spark in divorces can increase divorce rates substantially (and potentially permanently). If unmatched divorcées are disproportionately more likely to look for a partner, and if people can influence their access to new matches (for example through an endogenous search effort or divorce to single-hood), then divorce generates divorce. This could also help explain the mismatch between model and data after 1990. The simple model proposed in this paper is not capable of capturing such mechanisms. Other potential explanations for the slow, but more persistent, bust can be a change in the structural reasons for marriage. Starting with the baby-boom cohorts, marriage rates decreased substantially. Lower rates of marriage also imply lower rates of divorce. Similarly, the age-gap in marriage has been trending downwards, suggesting that the nature of age-preferences has also been changing.

2.5 Conclusion

In this article, I propose a causal link between the two demographic 'booms' of the 20th century in the United States. I show that on the aggregate the divorce boom took off right when baby-boomers started to enter the marriage market, creating better remarriage prospects for men in the pre-baby-boom generations. Across cohorts, the divorce boom was more pronounced among men for whom the remarriage opportunities improved the most. Moreover, I provide cross-state evidence that men and women divorced more if born in states with a larger baby-boom, an effect that is robust to controlling for socio-demographic characteristics, and more importantly to instrumenting the size of the baby-boom with WWII mobilization rates. Lastly, I built a simple repeated matching model, which can generate a sizable divorce boom (both in aggregate divorce rates and across cohorts) from a realistically behaved steady state marriage market.

Overall, this paper adds to the literature on large and surprising consequences of cohort size swings. Future research should examine the micro implications of this aggregate dynamic. In particular, if divorce is motivated by remarriage, the threat of divorce is changing over the life-cycle and across cohorts with changes in the underlying age-structure. This can have predictable implications on labor supply, savings, investment in children, and other important decisions that people typically make within marriage. The importance of remarriage considerations for divorce should also motivate development of new structural marriage market models capable of capturing this mechanism, which are now strikingly missing from the literature. Such contributions could also help reconcile the dynamics of divorce after 1990 that this paper fails to explain.

CHAPTER III

Housing Market Channel of Monetary Policy: the Role of Residents in Their Fifties

Abstract

The response of housing demand to changes in interest rates is a key mechanism of monetary policy. This paper shows that the effect of monetary policy shocks, identified through high-frequency event studies, on housing markets depends on the age-structure of the economy in a non-trivial way. Both across U.S. metro areas and across states, local housing prices drop more after monetary policy tightens whenever the share of population between 50 and 65 years of age is higher. If the share of population in a metro area 50-55 years old increases by one percentage point, a one standard deviation monetary policy shock depresses housing prices by an additional 0.413 percent after 3 quarters. A stronger investment motive in the demand for real estate by this age group is a possible mechanism. This differential reaction of housing prices is already detectable by the quarter of the shock, and is followed by a differential response in employment starting about four quarters after the shock.

3.1 Introduction

The volatility of the housing market is key to the understanding of business cycles and the response of the housing market is a key mechanism of monetary policy (Leamer, 2015). In this paper, I study the variation in the effect of national monetary policy on local housing markets. I show that the pass-through from a monetary policy shock to housing prices depends on the age structure of the economy. Specifically, I provide evidence that housing markets with a high share of population in their fifties react most to monetary policy. I test this hypothesis with both cross-state and cross-metro-area variation within the United States, confirm robustness across several dimensions and discuss possible mechanisms.

A large body of work shows that monetary policy tightening puts a downward pressure on housing prices. Jorda et al. (2015) provide robust evidence linking short-term rates, mortgage lending and house prices over 140 years and several countries, using monetary policy shocks identified based on the fact that countries with fixed exchange regimes often see fluctuations in short-term interest rates unrelated to home economic conditions. Paul (2020) provides recent evidence for the U.S. showing that high-frequency surprise jumps in short-term interests rates depress house price growth, and that this effect is time varying and especially low before the 2007-09 financial crisis.

Several papers study the geographic variation in the effects of monetary policy on housing markets. Fuss and Zietz (2016) show that local population growth is a key demand side factor and the percentage of undevelopable land a primary supply side factor that determine how national monetary policy impacts house price inflation rates at the MSA level. Congruently, Aastveit and Anundsen (2022) and Fischer et al. (2021) show that strict local regulatory environments and low housing supply elasticities are associated with larger housing price responses to monetary policy shocks, as quantities cannot adjust to the demand pressure. I propose a new factor that differentiates local markets in how reactive they are to shocks.s

A long tradition in housing economics relates the demographic structure of the economy to housing demand. Mankiw and Weil (1989) argued that the baby boom of 1960s caused a housing boom in the 1980s, as a large cohort entering house-buying age pushed demand up. More recently Takats (2012) uses cross-country variation to show that higher shares of young adults motivate construction. Hiller and Lerbs (2016) come to the same conclusion using variation across German cities. I show that the age-structure also affects the sensitivity of demand to shocks.

Overall, this paper contributes to the growing literature arguing that the economy exhibits time-varying responses to aggregate shocks, which depend on the microeconomic distribution of agents. Several papers study how the economy overall reacts to monetary policy shocks depending on the age structure of the population. The evidence so far is mixed. The most closely related of these is Leahy and Thapar (2022) who use a similar strategy to show that private employment and personal income respond most to a monetary policy shock in states with a high share of the population between 40 and 65 years of age. In section 3.6, I examine in detail how the effect on housing markets and the effect on employment are most likely linked. In contrast with Leahy and Thapar (2022), Berg et al. (2021) show that consumption expenditures for older households are more responsive to monetary policy shocks than for young or middle-aged households, while Sterk and Tenreyro (2018) find that the durable consumption of the young is more responsive than that of the middle aged and old. Wong (2018) also argues that consumption of the young is more sensitive to monetary policy, driven by the behavior of young homeowners. Similarly, Cloyne et al. (2020) shows that people holding a mortgage adjust consumption the most and those are typically younger.

3.2 Data

I measure variation in house price appreciation using the Federal Housing Finance Agency Housing Price Index, on U.S. state and metro-area level. The FHFA (purchase only) index is a weighted, repeat-sales index, measuring average price changes in repeat sales on the same properties. This information is obtained by reviewing repeat mortgage transactions on single-family properties whose mortgages have been purchased or securitized by Fannie Mae or Freddie Mac. The baseline index is available for 100 largest metropolitan areas since 1990. I supplement the data on housing price growth with state-level data on employment in construction, and in real estate, rental and leasing services. State level employment data is available beginning in 1990, provided by the Current Employment Statistics program conducted by the Bureau of Labor Statistics.[1] Lastly, I study the responses of housing permits (all and single-family). State-level data on new private housing units authorized by building permits is collected in the Building Permits Survey and provided by U.S. Census.

To study the responses to monetary policy shocks as dependent by age structure, I collect shares of population of metro areas by age in the 2013-2017 ACS as provided by the NHGIS progam (Manson et al., 2021). Population estimates by age for each year and state were retrieved directly from the U.S. Census website.[2] Shares of population in each age group are multiplied by 100, to be interpreted as percentage points.

I use monetary policy shocks identified using a high-frequency event-study approach, as pioneered by Kuttner (2001). Gurkaynak et al. (2005) study the responses of a range of asset prices around FOMC announcements in a narrow window (10 minutes before and 20 minutes after), so that one can be reasonably certain that any change is caused by the monetary policy change only. Gurkaynak et al. (2005) show that the effects on a variety of asset prices can be well summarized by 2 factors: one that captures variation related to current short-term interest rates (especially current and next month federal funds futures) and a second one that is constrained to only load on future or longer-term interest rates (labeled the 'target' factor and the 'path' factor, respectively). Swanson (2021) extends this analysis to the post great-recession period (capturing 241 FOMC announcements from July 1991 to June 2019), adding a third factor to account for quantitative easing. I use the 'target' factor as identified by Swanson (2021) as the monetary policy shock in this paper, and sum

[1] Construction employment is available for 45 states. Employment in real estate, rental and leasing services is available for 47 states.

[2] https://www2.census.gov/programs-surveys/popest/datasets/ retrieved in January 2021.

over the quarter to aggregate the high-frequency shocks to a quarterly frequency.[3] This shock is designed to reflect variation in a variety[4] of asset prices, yet is primarily identified from movements of short-term interest rates (stripping away the forward guidance aspect of monetary policy). This is important, as a streak of papers (starting with Nakamura and Steinsson (2018)) has shown that FOMC announcements, apart form revealing a change in policy, also have an 'information effect'. A tightening of monetary policy may reflect the fact that the Fed's forecasts about the behavior of the economy in the near future is more positive than the financial markets participants' forecasts. Such information effect thus contaminates this identification strategy. Paul (2020) studies surprise changes in a variety of interest rates around FOMC announcements and shows that while upward jumps in longer-term assets or future short-term rates do correlate with increases in GDP forecasts, this is not true for current short-term interest rates. Specifically, Paul (2020) replicates the methodology of Gurkaynak et al. (2005) and shows that the target factor does not suffer from being contaminated with an information effect.[5] The final sample covers 1991Q3-2019Q2 and is determined by data availability of the monetary policy shock. The monetary policy shock is rescaled to have a standard variation of 1 within the sample. Table C.1 provides summary statistics of the metro-area sample and figure C.1 plots the monetary policy shock.

3.3 Metro-area-level analysis

This section presents the main result of the paper. In metro-areas with a higher share of population in their fifties tightening of monetary policy slows down house price growth more.

To identify differential impact of monetary policy I regress the log-difference of a metro-area specific housing price index $p_{m,t}$ on a set of time and metro-area fixed effects and the share of population in a specific age range interacted with a monetary policy shock $\tilde{\epsilon}_t$.

$$log(p_{m,t+i}) - log(p_{m,t-1}) = \gamma_m + \delta_t + \alpha^{i,a} \cdot \tilde{\epsilon}_t \cdot s^a_{m,2010} + u_{m,t} \tag{3.3.1}$$

I run this regression separately for horizons $i = 0, ...6$ quarters and age ranges $a = 20 -$

[3]As is standard in the literature, I exclude the FOMC announcement on September 17, 2001, which took place before markets opened but after financial markets had been closed for several days following the 9/11 terrorist attacks.

[4]Including federal funds futures contracts of different maturities, Eurodollar futures contracts of different maturities, and the 2-, 5-, and 10-year Treasury yields.

[5]Paul (2020) also discusses that even short-term interest rate surprises around unscheduled meetings sometimes correlate with jumps in GDP forecasts, though at least in a subsample of 1995-2007 this is purely due to the September 2001 meeting, which is excluded in this paper. Moreover, I show in the appendix section that the results are robust to restricting the sample to after 1995 or before 2007, as well as to using the target factor as identified in Paul (2020) from only scheduled meetings.

$25, ..., 70-75, 75+$. Notice that in this specification $s^a_{m,2010}$ (the share of population of area m in an age range a) varies only crosssectionally, while the monetary policy shock $\tilde{\epsilon}_t$ varies over time. Thus while the baseline effect of $\tilde{\epsilon}_t$ is subsumed in the time fixed effects and the baseline effect of $s^a_{m,2010}$ is subsumed in the metro-area fixed effect, a differential response by a varying age-structure can be identified. The left-hand side $log(p_{m,t+i}) - log(p_{m,t-1})$ captures the cumulative change between $t-1$ and a horizon of i quarters, to trace down the effect on an impulse response as in Jorda (2005). The coefficients of interest $\alpha^{i,a}$ captures how much more a 1 standard deviation monetary policy tightening affects the impulse response of the log of housing prices if the share of population in the age range a is 1 percentage point higher. The sample of metro areas is weighted by 2010 population and the coefficients are rescaled by one hundred, so the results can be interpreted in percentage point changes.

Table 3.1 shows the results for horizons of 0-6 quarters and age ranges of 50-55 and 55-60. In areas where the share of population 50-55 years old is 1 percentage point higher, a one standard deviation shock decreases housing prices by a further 0.41 percentage points after 3 quarters and 0.64 percentage points by 6 quarters. Figure 3.1 plots $\alpha^{i,a}$ for a variety of age ranges and a horizon of 3 quarters. Similarly, figure 3.1 plots $\alpha^{i,a}$ for a variety of age ranges and a horizon of 6 quarters. Overall, a clear pattern emerges. Housing markets with a large share of residents in their fifties react substantially more to monetary policy. Metro-areas with a high share of population 35-45 years of age react less than average (though this could be because this share is negatively correlated with the share in their 50s). Other age groups do not show a significant effect.

3.4 State-level analysis

Next, I show that also across states higher share of population in their fifties is associated with a stronger housing market reaction to monetary policy with the following specification:

$$log(y_{s,t+i}) - log(y_{s,t-1}) = \gamma_s + \delta_t + \alpha^{i,a} \cdot \tilde{\epsilon}_t \cdot s^a_{s,t} + \phi^{i,a} \cdot s^a s, t + \beta X_{s,t-1} + u_{s,t} \qquad (3.4.1)$$

$X_{s,t-1}$ includes 4 quarter population growth.[6] First I confirm that the pattern observed across metro areas for house price growth also holds across states. Figures 3.2 plot $\alpha^{i,a}$ for $y_{s,t}$ being a state-level housing price index, a variety of age ranges and a horizon of 3 quarters and 6 quarters respectively. Across states it also holds that higher shares of population between 50 and 60 years old is associated with stronger responses of housing prices to monetary policy. When studying variation across metro-areas it seems that housing markets with a higher

[6]4-quarter growth is chosen, because population estimates are only available annualy.

Figure 3.1: Differential response of metro-area housing prices to monetary policy shocks based on the local age structure

	$\Delta_i \log(p_{m,t+i})$						
	$i=0$	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$	$i=6$
$s^{50-55}_{m,2010} \cdot \tilde{\epsilon}_t$	-.096	-.210	-.292	-.413	-.492	-.572	-.643
	(.032)	(.066)	(.093)	(.107)	(.139)	(.160)	(.180)
	$\Delta \log(p_{m,t+i})$						
	$i=0$	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$	$i=6$
$s^{55-60}_{m,2010} \cdot \tilde{\epsilon}_t$	-.102	-.236	-.324	-.406	-.488	-.542	-.577
	(.027)	(.054)	(079)	(.094)	(.117)	(.138)	(.158)
N				11200			
N clusters				100			
Sample				1991Q3 - 2019Q2			

Standard errors in parentheses, clustered at the MSA level.

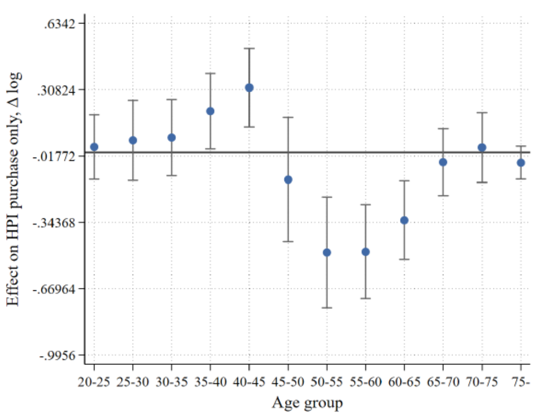

Lag 3

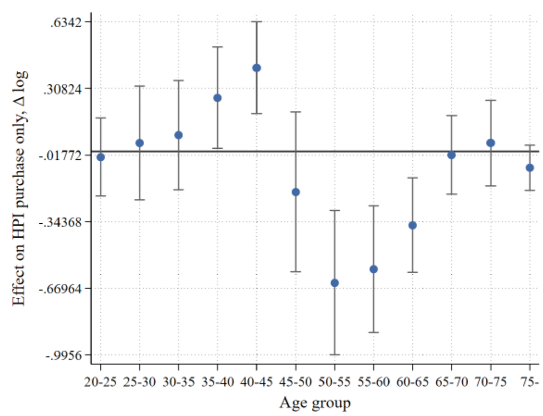

Lag 6

Sample of 100 largest MSAs and metropolitan divisions, weighted by population. Source of price indices: FHFA HPI ©.

Figure 3.2: Differential response of state-level housing prices to monetary policy shocks based on the local age structure

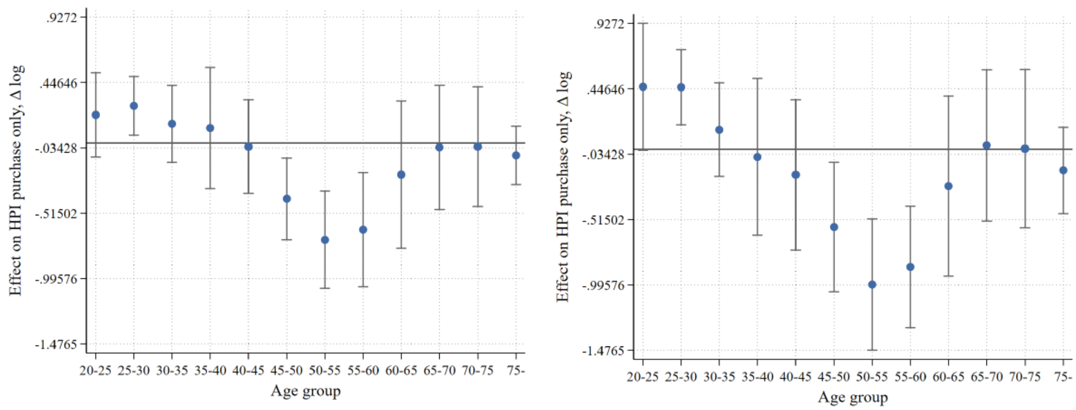

Lag 3 Lag 6

Sample of US states, weighted by population. Source of price indices: FHFA HPI ©. Figures plot the point estimates and 95% confidence intervals of $\alpha^{3,a}$ and $\alpha^{i,6}$ for a variety of age ranges.

share of population between 35-40 responds less to monetary policy shocks. This pattern is not robust to using state-level variation. If anything, it seems that higher shares of population in their twenties is possibly associated with a weaker response to monetary policy.

A stronger reaction in house prices can be caused by a stronger housing demand response or a weaker housing supply response. Thus I next study a set of quantity-related variables associated with the housing market. There is no publicly available data on the number of sales disaggregated by state or metro-area for a long enough time horizon. However, state-level employment by sector is available since 1990. I study two sectors especially associated with activity in the housing market: employment in real estate, rental and leasing services and employment in construction. Figures 3.3 show the results for a horizon of 6 quarters. Real estate, rental and leasing services employment does react more to a monetary policy shock in states with a lot of people in their fifties. However, unlike with prices other age-groups show up as having a significant impact as well. Younger states are associated with weaker monetary policy responses in real estate employment, while older states have a stronger response. The results are similar for construction employment. States with a higher share of population bellow 35 years of age have their construction employment less affected by national monetary policy, while states with a lot of people in their fifties react the most.

Last, I study the behavior of new housing permits to build single-family housing. Figures 3.4 show that housing permits react more to monetary policy in older states and less in younger states. However, unlike with housing prices and sectoral employment the share of population in their fifties does not have a significant effect on the sensitivity of housing

Figure 3.3: Differential response of state-level housing-related employment to monetary policy shocks based on the local age structure

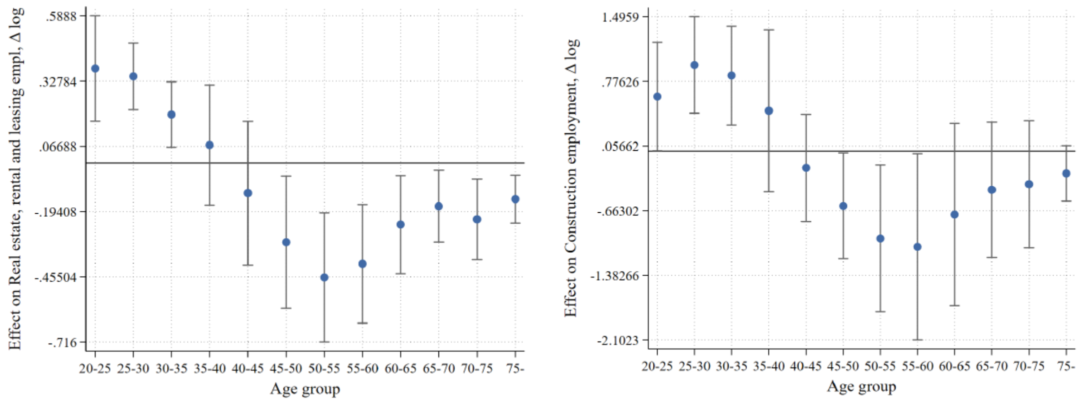

Real estate, rental and leasing emp, lag 6

Construction employment, lag 6

Sample of US states, weighted by population.

Figure 3.4: Differential response of state-level housing starts to monetary policy shocks based on the local age structure

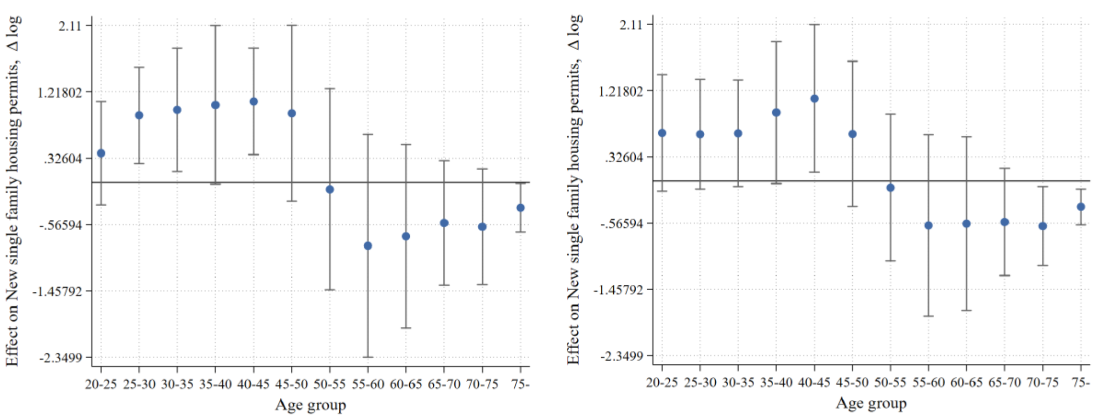

Lag 3, 1-Unit Structures

Lag 6, 1-Unit Structures

Sample of US states, weighted by population.

permits to monetary policy. Still, the pattern is broadly in line with the behavior of prices. Importantly, there is no evidence suggesting that the the supply of housing reacts less to monetary policy in states with a lot of people in their fifties. This implies that the price responses identified above are likely driven by a difference in demand behavior, not supply. Moreover, the inconclusive behavior of permits suggests that the price responses are coming from the market for existing houses and less so for new construction.

3.5 Robustness

Fuss and Zietz (2016), Aastveit and Anundsen (2022) and Fischer et al. (2021) all show that monetary policy affects housing prices more in metro areas that have a low elasticity of housing supply. If the age structure of the population correlates with housing supply elasticity, this effect could confound the mechanism estimated in this paper. Following Aastveit and Anundsen (2022), I use the metro-area housing supply elasticity measure developed by Saiz (2010). Figures C.4 and C.4 repeat the metro-area analysis while adding $\tilde{\epsilon}_t \cdot Saiz\ elasticity_m$ as an additional control, showing that if anything the effects are more pronounced.

Next I test how robust the results are to using different measures of monetary policy shocks. Figure C.6 repeats the metro-area analysis for several different monetary policy shocks identified in the literature. In the first figure, I show that using raw jumps in the current month federal funds futures implied rate, as in Gorodnichenko and Weber (2016) extended with the replication data provided by Paul (2020), gives almost the same results but with less precision. The last figure shows that the overall results are robust to using the 'target factor' identified on data excluding unscheduled FOMC meetings, though the effect of the 50-60 ages share are slightly smaller. Jarocinski and Karadi (2020) also use a high-frequency event study. But instead on relying on the variation in only short-term interest rates, they purge the information effect from jumps in the 3-month rate through a VAR including stock price jumps around FOMC meetings. Again, the second figure shows that housing markets with more 50-60 year olds react mote to this version of a monetary policy shock. Lastly, I show that the results are broadly robust to using alternative monetary policy shocks not reliant on high-frequency event studies as identified in Romer and Romer (2004) and Bu et al. (2021), except that metro areas with more people 65-75 years old exhibit even weaker responses to monetary policy.

Paul (2020) suggests that asset price jumps around unscheduled FOMC meetings before 1995 or after 2007 exhibit more of an information effect. Moreover, only after 1995 did the FOMC began issuing press releases following every meeting. Figure C.5 presents the baseline results starting the sample in 1995. Since the Great Recession starting in 2008 was associated

Table 3.1: Differential response of state-level housing prices: with and without netting the effect on state-level employment

	$\Delta_i \log(p_{s,t+i})$						
	$i=0$	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$	$i=6$
$s^{50-55}_{s,t} \cdot \tilde{\epsilon}_t$	-.188	-.438	-.595	-.710	-.863	-.927	-.992
	(.054)	(.124)	(.177)	(.183)	(.207)	(.230)	(.247)
$s^{50-55}_{s,t} \cdot \tilde{\epsilon}_t$	-.193	-.443	-.419	-.307	-.527	-.541	-.128
	(.052)	(.118)	(.133)	(.132)	(.183)	(.214)	(.304)
$X_{s,t}$:							
$\Delta_i \log(L_{s,t+i})$	x	x	x	x	x	x	x
N	5264						
N clusters	47						
Sample	1991Q3 - 2019Q2						

Standard errors in parentheses, clustered at the state level.

State-level variation. Results from specifications 3.4.1 and 3.5.1. The second line of the table shows the differential effect on house prices when accumulative effects on state-level employment are netted out.

with a major housing market correction, I check that the results are not driven by the post 2008 sample in figure C.5.

Leahy and Thapar (2022) show that employment and personal income in states with a high share between 40 and 65 years of age respond more to monetary policy. A natural question is then whether the differential impact on housing markets is purely a result of differential impact on the whole economy, employment and income, that consequently prompts a shift in demand for housing through an income effect. I rerun the baseline state-level specification with the housing price index controlling for the proximate changes in state level non-farm employment $E_{s,t}$

$$log(p_{s,t+i}) - log(p_{s,t-1}) =$$
$$\gamma_s + \delta_t + \alpha^{i,a} \cdot \tilde{\epsilon}_t \cdot s^a_{s,t} + \phi^{i,a} \cdot s^a s,t + \beta X_{s,t-1} + \beta \cdot (log(E_{s,t+i}) - log(E_{s,t-1})) + u_{s,t} \quad (3.5.1)$$

Table 3.1 presents the baseline state-level results for $a = 50-55$ and horizons of 0-6 quarters contrasted with the results from running the specification 3.5.1. At shorter-term horizons (0-2 quarters) the housing price results are essentially unaffected by controlling for changes in employment. Thus I conclude that the proximate cause of the differential reaction in the housing market is not a differential change in overall employment. At longer-term horizons the independent effect on housing prices diminishes and ultimately becomes subsumed with the variation in employment. This suggests that a stronger reaction in housing prices is ultimately followed by an effect on the whole economy.

3.6 Response of employment with and without the housing market channel

In this section I show that the housing market plays a key role in why employment in states with a large late-middle-aged population reacts most to monetary policy. To study the interaction between the effect on housing markets and the reaction in employment I run the following specification.

$$
\begin{aligned}
&log(L_{s,t+i}) - log(L_{s,t-1}) = \\
&\quad \gamma_s + \delta_t + \alpha^{i,a} \cdot \tilde{\epsilon}_t \cdot s^a_{s,t} + \phi^{i,a} \cdot s^a s, t + \beta X_{s,t-1} + \beta \cdot (log(p_{s,t+i}) - log(p_{s,t-1})) + u_{s,t}
\end{aligned} \tag{3.6.1}
$$

On the left hand side is a cumulative log-change in non-farm employment as in Leahy and Thapar (2022). As a control I add the cumulative log-change in housing prices. Figures 3.5 show the results for a variety of age groups and horizons, with blue markers and solid confidence intervals. In each figure, gray markers represent results without adding $(log(p_{s,t+i}) - log(p_{s,t-1}))$ for comparison. Overall, the effect of age structure on the effectiveness of monetary policy is broadly similar with or without controlling for the reaction of the housing market, confirming this is not the only channel.

However, the exception is precisely in the effect of shares of population between 50 and 60 years of age. At all horizons, controlling for the cumulative log-change in housing prices diminishes the difference between the 50-60 age range and older groups. With or without an effect on housing prices age structure has a significant effect on how much monetary policy affects employment. Yet, when the effect on the housing market is regressed out, the conclusion simplifies. Younger states (with a high share of population bellow 40) react less than older states (with a high share of population above 50). A possible explanation is that the results in Leahy and Thapar (2022) are a combination of two mechanisms. A stronger housing market reaction by people of 50-60 years of age resulting in an overall slowdown in demand and employment, and an unknown mechanism that limits the response of employment whenever the share of population bellow 40 is high.

Interestingly, the results in figures 3.5 are essentially unchanged if only the cumulative

Figure 3.5: Differential response of state-level employment based on local age structure: with and without netting the effect on state-level housing prices

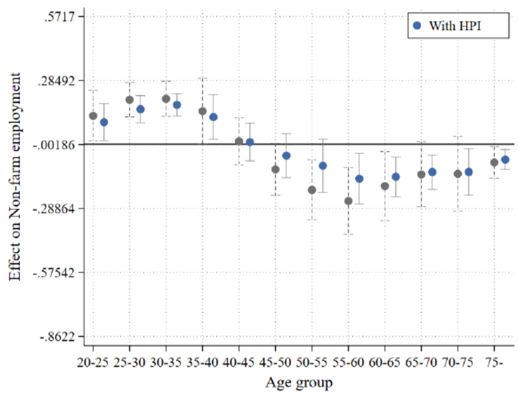

Lag 4

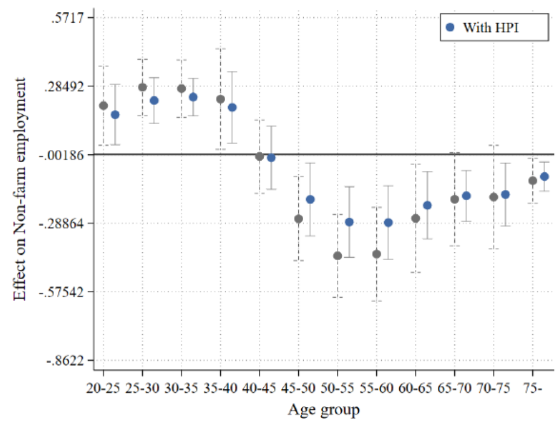

Lag 6

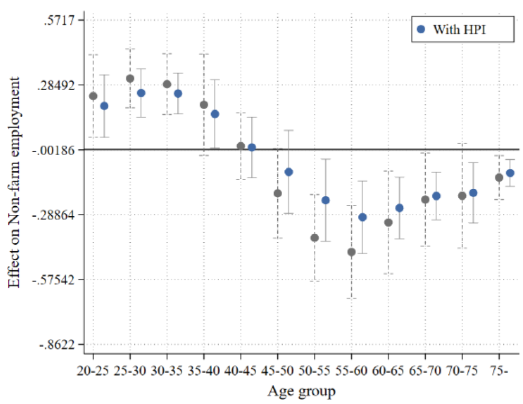

Lag 8

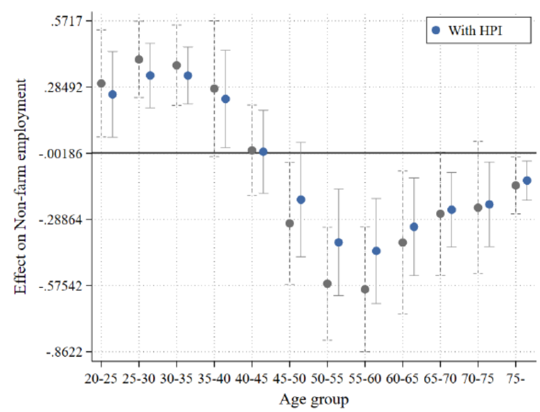

Lag 10

State-level, weighted by population. Source of price indices: FHFA HPI ©. Estimates of equation 3.6.1. Gray dots and dashed confidence interval replicate the results from Leahy and Thapar (2022), extended to match the sample and monetary policy shock used in this paper. Blue dots and solid confidence intervals add the cumulative log-change in housing prices as an additional control.

effect on house prices by the second quarter is used:

$$log(L_{s,t+i}) - log(L_{s,t-1}) = \gamma_s + \delta_t + \alpha^{i,a} \cdot \tilde{\epsilon}_t \cdot s^a_{s,t} + \phi^{i,a} \cdot s^a s, t + \beta X_{s,t-1} + \beta \cdot (log(p_{s,t+2}) - log(p_{s,t-1})) + u_{s,t}$$

This reinforces the conclusion that the housing market mechanism operates sequentially. First housing demand reacts more in locations with a high share of population between 50 and 60 years of age. Next, after several quarters, this change propagates to a stronger reaction in employment.

3.7 Possible mechanisms

The response of housing markets to monetary policy depends pn the age structure of the market in a non-trivial way. It is neither the young places nor the older places that stand out. Prices respond the most in places with a high share of the population in their fifties. It thus begs the question – what makes 50-year-olds special?

One of the salient features that distinguishes this group from others is that they are at the peak of their life-cycle savings behavior. Their incomes are high and their retirement is getting close, motivating them to save more. This also translates in their housing market behavior. They are the most likely to be owners and their housing wealth is at their life-cycle peak. This is documented in the existing literature, for example in Iacoviello and Pavan (2013) (see a reproduced figure from their paper for reference 3.6).

Figure 3.6: Iacoviello and Pavan (2013)

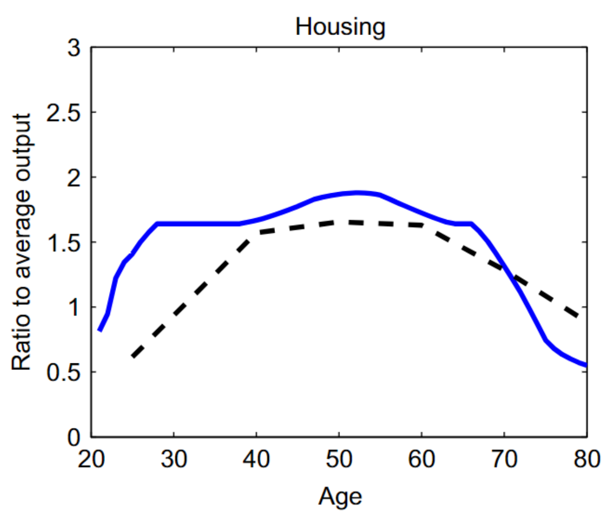

Figure reproduced from Iacoviello and Pavan (2013). Value of housing (the share holding housing times the median holding of housing) to output ratio over the life-cycle. The dashed line represents data from the 1983 Survey of Consumer Finances, while the solid line represents model simulations.

I hypothesize that a possible mechanism explaining the results in this paper stems precisely from this savings motivation. Intuitively, housing serves two purposes to individuals: it provides an immediate consumption value, but it is also a savings vehicle. Since people in their fifties have already accumulated a lot of housing, they do not care as much about the additional unit for immediate consumption. However, they are still very interested in acquiring more to save up. This dichotomy can be illustrated with a simple first order condition for housing stemming from a standard no-risk life-cycle model with consumption, housing (with a price p_t) and a savings/borrowing rate on a safe asset r_t:

$$\frac{MU_{housing}}{p_t \cdot MU_{consumption}} + \left(\frac{p_{t+1}}{p_t}\frac{1}{1+r} - 1\right) = 0 \tag{3.7.1}$$

In this simple setting, there are two aspects of housing demand, the first term represents the immediate value of consuming an additional unit of the housing good, the second term represents the present value of the investment in one unit of housing. If $\frac{MU_{housing}}{MU_{consumption}}$ is lower for people in their fifties, their demand for housing is more sensitive to interest rates. The simplest explanation could come from the life-cycle pattern. Assume the utility from housing is concave, and potentially more concave than utility from consumption. Then though both housing and consumption profiles peak around this age, the peak in housing possibly pushes $\frac{MU_{housing}}{MU_{consumption}}$ down for people in their fifties (giving this term a U-shape). Other factors can be at play. For example, for credit constrained households there is an additional immediate benefit to housing – serving as a collateral. This additional immediate benefit diminishes the response to interest rates just as a high marginal utility from housing consumption does, providing an additional reason for high sensitivity to interest rates of the middle-aged as they are unlikely to be credit-constrained.

The argument so far is based on the incentives of owners. The housing demand for renters pins down the market rent at $\frac{MU_{housing}}{MU_{consumption}} = R_t$ and those who operate in the housing market purely as investors have a demand equation $\frac{R_t}{p_t} + \left(\frac{p_{t+1}}{p_t}\frac{1}{1+r} - 1\right) = 0$, which collapses to equation 3.7.1 except what matters are the preferences of renters. However, if a rental market is present and the behavior of 50-year-olds is driven by their investment motive, they can consider participating in the housing market as downright investors renting out additional properties. Thus, to preserve the argument that a low marginal utility of housing consumption $MU_{housing}$ makes the demand of 50-year-olds for housing more sensitive to interest rates, there has to be a barrier for them as individuals to become investors or to get the same rents as large investors. This is not hard to imagine – with the fixed costs of acquiring an additional property, the agency issues ofr being a landlord, the institutional knowledge of the market that an individual investor is lacking and the often higher property taxes paid on properties that are not a first residence, it is no wonder a 50-year-old saver would opt for investing more in the home they live in (or vacation themselves) rather then branching out into the wild west of owning and renting out additional properties. Overall, the life-cycle profile of housing consumption is a possible mechanism for the results in this paper.

3.8 Conclusion

In this paper I provide evidence that local housing markets are more sensitive to monetary policy if their population is more middle-aged. Specifically, the higher the share of population who are between 50-65 years old, the more housing prices fall (rise) after monetary policy tightens (loosens). This conclusion is true both across U.S. states and metropolitan areas and is robust to alternative sub-samples, alternative monetary policy shocks and to controlling for housing supply elasticity. I also show that employment in housing-related sectors (real estate and construction) falls most after a monetary policy tightening in states with a high share of 50-65 year-olds, while housing permits do not show a clear pattern. This suggests that the differential effect of housing prices is a result of a differential response in demand, not supply.

The effects on employment relate this paper directly to Leahy and Thapar (2022) who show a similar conclusion across U.S. states for overall employment. In this paper I provide evidence that the differential effect on housing prices precedes that on employment but after several quarters they coincide. Moreover, when the effect on housing prices is netted out, the results for employment resemble more a simpler pattern where older places react more strongly to monetary policy than younger places. This suggests the effects on employment in Leahy and Thapar (2022) are a combination of a mechanism where the economy reacts more when its population is older with a mechanism where the housing market (and then consequently the rest of the economy) reacts most when middle-aged. Future research should unpack why employment in older states reacts more to monetary policy. Lastly, I suggest a possible mechanism that could generate the presented results in the housing market. A fruitful area for future research is to look for supportive microeconomic evidence for this mechanism, or to explain the pattern with an alternative one.

APPENDICES

APPENDIX A

Appendix to Chapter 1

A.1 Supplement to empirical observations

In this section I present supplementary empirical evidence to chapter 1. First, table A.1 lists the commuting and location related variables available or possible to construct in the PSID, together with their time span. Table A.2 shows basic summary statistics of the main PSID sample used in the empirical section as well as in the estimation.

Table A.1: Commuting variables available in the PSID

Variable	waves
Distance city center	1969-
Commuting distance	1970-1986 with gaps[1]
Commuting time annual	1970-1986 with gaps[2]
Commuting time usual	2011-2017
Distance to a job	2013-2017
Distance to an average job	1990-2017 (1990-2000 backfilled)
Distance to an average job in ind & seg	1990-2017 (1990-2000 backfilled)

Table A.2: PSID sample summary statistics.

PSID summary statistic	since 1969	since 1990
Age	33	33
Man	49%	50%
2010 Population size of metro-area	4679k	4524k
Having children in the household	62%	58%
Number of children (including zeros)	1.2	1.1
Commuting distance in miles	10.6	
Annual commuting time in hours	173	
Standard deviation of commuting distance in miles	11.3	
Usual annualized commuting time in hours		183
Distance to current job		9
Distance to center in miles	14.7	15.5
'Distance to opportunity'	12.11	12.18
Distance to an average job	12.2	12.3
Share living less than 10 miles from the center	44%	41%
In couple	68%	66%
Tenure in a couple (including negative values for singlehood)	7.8	7.7
Share of men in couples working	98%	97%
Share of women in couples working	76%	82%
Annual hours of work of men in couples (including 0s)	2202	2224
Annual hours of work of women in couples (including 0s)	1201	1415
Annual hours of work of men in couples (when both work)	2249	2293
Annual hours of work of women in couples (when both work)	1578	1724

Age restrictions 18-50. Geographic restriction: in a metro area of at least 250 thousand residents per the 2010 Census. Moreover, for each individual I select their most common metro area over their observed lifetime in the PSID and exclude periods when this individual did not live in this MSA, so that all location changes are within the same area. Lastly, I only use single people who have not been in a couple before and couples for whom this is their first match, as far as it can be determined in PSID. Metro-area assigned as the most frequent metro area within the sample. Row "Tenure in a couple (including negative values for singlehood)" present the sample mean of $Y - Y_{\text{first observed in a couple, or first year of marriage for original sample couples}}$.

Tables A.3 and A.4 show the differences in commuting between singles and couples, measured by directly comparing before and after outcomes (i.e. using person fixed effects). These differences are quantitatively smaller than the main analysis presented in section 1.2, primarily because only a limited number of individuals are observed in both states and the period they spent in being in a couple is very short. This is explicit in figures 1.2, showing event studies with person fixed effects, but pooling the comparison of before and after forming a couple with being in a couple for shorter and longer periods of time. The quantitatively large difference between singles and couples is robust to controlling for person fixed effects and emerges after about 5 years of being in a couple.

Table A.5 shows the differences in commuting by gender and relationship status for two alternative measures of commuting in the psid: annualized hours of commuting (computed

Table A.3: Commuting differences (distance) between singles and individuals in couples – person fixed effects

	Commuting distance (miles)									
	Men					Women				
In couple	1.079	1.095	.964	1.041	.964	-.050	-.100	-.067	.036	-.224
	(.649)	(.616)	(.632)	(.592)	(.620)	(.731)	(.711)	(.719)	(.739)	(.726)
d_o		.469					.322			
		(.071)					(.068)			
d_c			.365					.309		
			(0.047)					(.058)		
d_o *bins**				x					x	
d_c *bins**					x					x
N	24905	23784	24905	23784	24905	16291	15868	16291	15868	16291
N clusters	154	152	154	152	154	146	144	146	144	146

SEs statistics in parentheses. SEs clustered at the MSA level.

All regressions include year, age and person fixed effects.

* Includes dummies for d_x in intervals of 0-6, 6-12, 12-20, 20-30, 30-40 and over 40 miles.

Table A.4: Commuting differences (time) between singles and individuals in couples – person fixed effects

	Commuting time (annual)									
	Men					Women				
In couple	18.781	19.417	16.588	18.802	6.010	5.092	5.782	5.677	5.782	5.134
	(7.9409)	(7.628)	(7.713)	(7.532)	(7.471)	(8.307)	(9.139)	(8.837)	(9.139)	(8.980)
d_o		4.593					2.355			
		(.898)					(.806)			
d_c			3.879					2.309		
			(0.489)					(.623)		
d_o *bins**				x					x	
d_c *bins**					x					x
N	24924	23612	24924	23612	24924	17232	16798	17232	16798	17232
N clusters	154	152	154	152	154	146	143	146	143	146

SEs statistics in parentheses. SEs clustered at the MSA level.

All regressions include year, age and person fixed effects.

* Includes dummies for d_x in intervals of 0-6, 6-12, 12-20, 20-30, 30-40 and over 40 miles.

from an average daily commute report) and a distance from residence to work in miles (census tract to census tract). These are available in more recent waves of the survey, confirming the differences between couples and singles are a persistent pattern.

Table A.5: Alternative measures of commuting

	Commuting time (typical, annualized)		Distance to work (tract to tract)	
In couple	-7.600	-10.443	-.252	-.891
	(6.383)	(6.022)	(.420)	(.341)
Man in couple	24.197	23.536	2.630	2.303
	(10.582)	(10.703)	(.728)	(.703)
Man	15.550	16.464	-.773	-.403
	(8.791)	(8.781)	(.563)	(.562)
X_i:				
Education, race, cohort	x	x	x	x
Distance to center d^c		x		x
N	15215	15208	7922	7922
N clusters	171	170	160	160

SEs statistics in parentheses. SEs clustered at the MSA level.
All regressions include year, age, MSA fixed effects.

$$d_{it} = \beta \cdot \text{In couple}_{it} + \beta_{wc} \cdot \text{Man in couple}_{it} + \beta_w \cdot \text{Man}_i + \alpha_t + \alpha_a + \alpha_{msa} + X_i + \epsilon_{it}$$

Tables A.6 and A.7 show the differences in commuting by gender and relationship status for the commuting variable available in the 2000 Census.

Table A.6: Commuting differences by gender and relationship status in the 2000 Census.

	Commute (annualized)				
Man	-2.705			29.41	-3.946
	(.874)			(1.798)	(0.863)
In couple	-11.68	17.99	-8.157		
	(2.111)	(1.201)	(2.235)		
Man in couple	31.92				
	(2.208)				
Industry 1-digit NAICS dummies.	x	x	x	x	x
Sample:		men	women	couples	singles
N	2286363	1245988	1040375	1565336	721027

SEs statistics in parentheses, clustered at MSA level.
All samples include only people who are married or never married.
All regressions include age, MSA, education, race, cohort controls.

Source: IPUMS 2000 Census 5% sample. Sample: 18-50 years old, married or never married, MSAs of at least 250k people. "In couple" includes married and cohabitation. Industry dummies are for 1-digit NAICS codes.

Table A.7: Commuting differences by gender and relationship status in the 2000 Census.

	Commute (annualized)				
Man	4.085			40.38	3.000
	(0.978)			(1.893)	(0.949)
In couple	-12.35	20.82	-8.561		
	(2.247)	(1.218)	(2.370)		
Man in couple	35.85				
	(2.285)				
Sample:		men	women	couples	singles
N	2286363	1245988	1040375	1565336	721027

SEs statistics in parentheses, clustered at MSA level.
All samples include only people who are married or never married.
All regressions include age, MSA, education, race, cohort controls.

Source: IPUMS 2000 Census 5% sample. Sample: 18-50 years old, married or never married, MSAs of at least 250k people. "In couple" includes married and cohabitation.

Figures A.1 present the well-known result that women are more likely to drop out of the labor force after forming a couple than men. Figures A.2 show the concurrent increase in the number of children in the household.

Figure A.1: Employment with respect to time spend in a couple

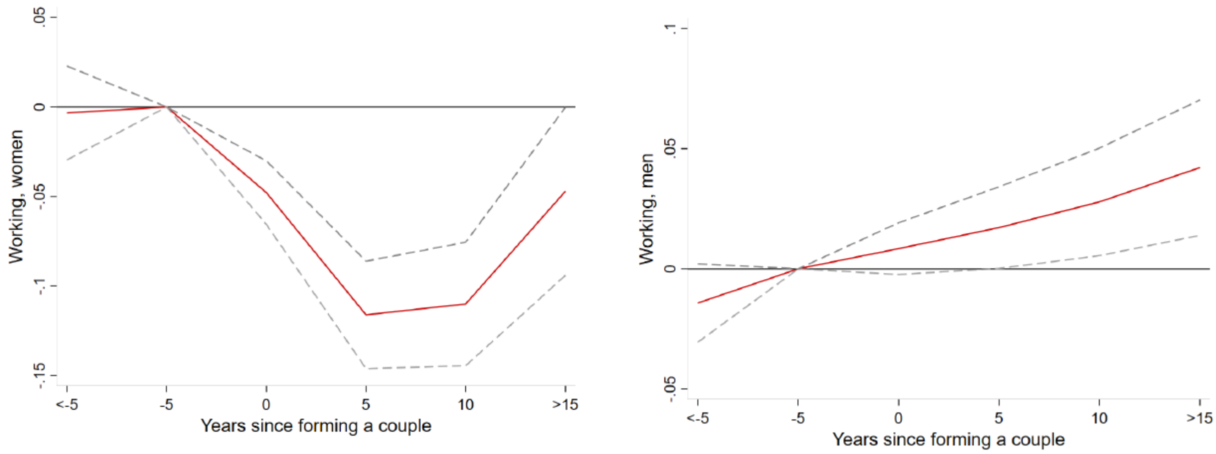

women men

Source: PSID

Figure A.2: Number of children living in the household with respect to number of years in a couple

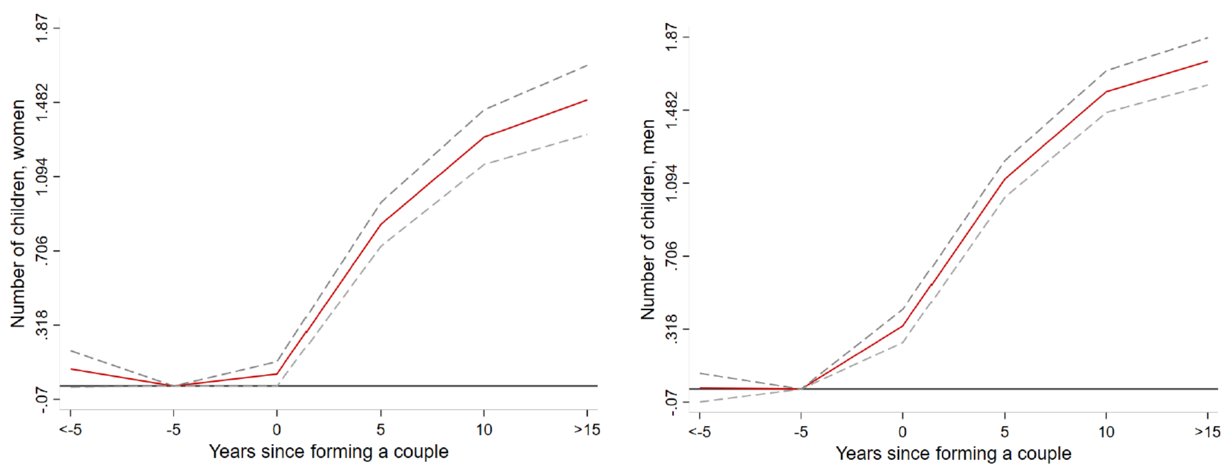

women men

Analogous regressions for figures 1.3. With distance to center d^c replaced with the number of children in the household.

As couples are more likely to have children bellow the age of 18 living with them in the household, I hypothesize that an amenity which is more valued by couples than singles can be the quality of schools. As a proxy, I use the averaged standardized test scores in the public school district of residence administered in 3rd through 8th grade in mathematics and reading, language and arts over 2008-2018 school years, normalized to be comparable nationally and to the middle grade of the data as provided by the Education Opportunity Project (Reardon et al., 2021). The evolution of this proxy measure with respect to tenure in a couple (controlling for year and age dummies) is plotted in A.3, showing that couples do move to better school districts after 5-10 years of being in a couple. It is noteworthy, that this measure of school quality correlates weakly with distance to center. Thus by this measure, suburbs have on average somewhat better schools. However, other measures of school quality (such as financing per pupil or the value-added measures produced by Reardon et al. (2021) fo not correlate systematical with distance to center).

Figure A.3: Average test scores in the school district of residence with respect to time in a couple

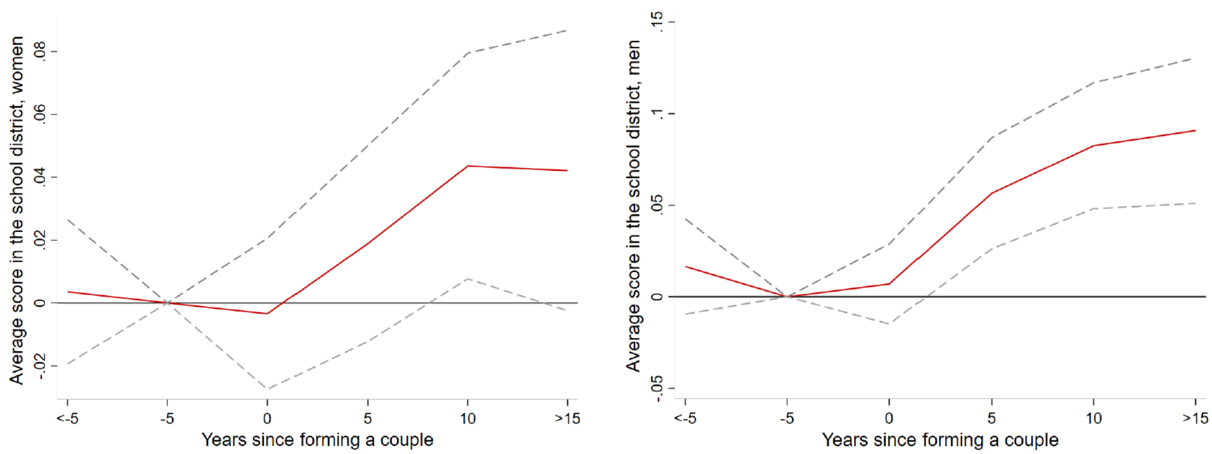

women

men

Analogous regressions for figures 1.3. With distance to center d^c replaced with averaged standardized test scores in the public school district of residence averaged over 2008-2018 (only cross-sectional variation in test scores used).

In my model, I link the specialization on the commuting margin with specialization in time use in the household (as a source of the gendered nature of specialization, not the need for specialization in the first place). For the sake of simplicity, I do not distinguish between childcare and housework in my analysis. From the existing literature, we know that women perform more of both. In this paper I do not aim to uncover the root causes of this gendered specialization; I assume women are more productive in home production as a means to fit the observed patterns. Given the patterns in figures A.2, it would be reasonable to hypothesize that the changes in commuting I observe after men and women form couples are somehow directly related to couples having children in the household, not to the nature of being in a couple alone. Table A.9 shows, however that this is an especially difficult assertion to prove or disprove with this data, because the sample of individuals who ever end up in a couple, have information on commuting and are never observed in a household with a child living in it, is very small. In other words, almost all couples in the sample eventually end up having children, making it difficult to argue that any behavior in these couples is unrelated to future, present or past child-rearing.

Table A.8: Commuting by relationship status, controlling for having children.

	Commuting distance (miles)	
	men	women
In couple	**2.138**	.682
	(.684)	(.653)
children	.969	-1.284
	(.412)	(.394)
N	23243	13238

Available in 1975-1976, 1978-1986; plus in 1969-1974, 1977 for heads of households only. Sample of only those in a couple, or those that have never been observed in one, but eventually they will be.

All regressions include year, age, MSA, education, race, cohort controls.*SEs* statistics in parentheses. SEs clustered at the MSA level.

Commuting differences by relationship status. Separately by gender. As in 1.1, including a dummy for having a child under 18 in the household.

Still, the evidence in tables A.8 and A.9 suggests that specialization in commuting is not purely tied to children, but is more gendered when children are present. Table A.8 shows that controlling for having a child under 18 in the household shrinks the increase in commuting for men after forming a couple marginally (by about half a mile). On the other hand, when controlling for having children women increase commuting after forming a couple by about 0.7 miles (up from the baseline increase of only 0.2 miles), though this difference between single and coupled women is not statistically significant. Correspondingly, table A.8 shows that when a child is present, men commute more while women commute less. Overall, this suggests that child-rearing is related to the gendered nature of specialization in commuting, as there are likely synergies between taking care of children and being close to home. This reasoning motivates the choice of functional form in the model where not commuting and spending time in home production are technologically complementary. However, this evidence also suggests that the sizable increase in commuting after men form couples is not strongly tied to having children. In other words, the benefits to specializing on the commuting margin and thus the increased value of forming couples in long-commute environments come from the nature of sharing a household, not the presence of children. This is supported by the evidence in table A.9. This table shows the increase in commuting for men after they form a couple with samples split by whether the person is ever observed having children in the household or not. While the sample of men who never end up having children is small, the increase in commuting is just as large as for those who end up having children.

Table A.9: Commuting differences, separately by whether a person ever ends up in a couple living with a child or not.

	Commuting distance (miles)			
	men		women	
In couple	**2.392**	**2.691**	.422	.564
	(.687)	(1.248)	(.673)	(1.529)
N	22073	2216	11963	1667
Eventually observed with children	1	0	1	0

Available in 1975-1976, 1978-1986; plus in 1969-1974, 1977 for heads of households only. Sample of only those in a couple, or those that have never been observed in one, but eventually they will be.

All regressions include year, age, MSA, education, race, cohort controls.*SEs* statistics in parentheses. SEs clustered at the MSA level.

Commuting differences by relationship status. Separately by gender. As in 1.1, separately by whether a person ever ends up in a couple living with a child under 18 in their household.

Table A.10 repeats the analysis of hours within couples presented in table 1.6 showing in the second ling that within couples living far away from relevant jobs is associated with women's hours falling more compared to men's. The negative association between job access and hours for men is not robust to different sub-samples, but the gendered nature of the response is. Columns 1 and 2 limit the sample to waves when commuting distance and commuting time respectively are available. Columns 3 and 4 limit the analysis to a sample of waves when the alternative commuting measures, annualized hours and distance to work, are available. These are newer waves and only small samples. Table 5 shows that the result is robuts to using newer waves as long as a sufficient sample is used.

Table A.10: Hours and potential commutes within couples – alternative samples

	Hours				
Distance	5.914	8.209	-8.385	-10.231	-6.593
to jobs (d^{opp})	(3.186)	(1.767)	(4.213)	(5.383)	(2.239)
$d^{opp} \cdot$ Woman	-4.812	-5.685	-1.530	-1.082	-5.274
	(1.981)	(2.123)	(1.926)	(2.159)	(1.698)
Woman	-827.003	-863.17	-552.067	-555.471	-520.638
	(64.583)	(66.743)	(52.937)	(58.368)	(42.665)
Sample:					
Waves when	miles	annual	annualized	distance	
commuting variable is available		hours	annualized	to work	
					year $\geq$ 2000
X_i:					
'Labor market' fes	x	x	x	x	x
Couple fes	x	x	x	x	x
N	33300	35838	11586	8488	27378
N clusters	150	150	160	158	171

SEs statistics in parentheses. SEs clustered at the MSA level.

All regressions include year, age, MSA fixed effects.

Analogous analysis of hours spent working to table 1.6, except over samples when commuting variables are available, providing a more direct link to table 1.5. Using the backfilled d^{opp} from the first year available, typically 2002, in the first two columns.

A.2 Model supplemental information

A.2.1 Model solution

For each location of a job (determining distance and deterministic job characteristics) and rent R, solutions of consumption, time use and housing demand are as follows:

Consumption and housing quantity demand as a function of the budget multiplier and rent for all households can be solved in closed form. For singles:

$$H^s(\mu) = \Big(\frac{\Omega_H}{R\cdot\mu}\Big)^{1/\omega_h}$$

$$c^s(\mu) = \Big(\frac{1}{\mu}\Big)^{1/\omega_c}$$

With income as a function of labor supply $Y(h^s) = h^s \cdot w$.

For couples:

$$H(\mu) = 2^{1-1/\omega_H}\Big(\frac{\Omega_H}{R\cdot\mu}\Big)^{1/\omega_h}$$

$$c^h(\mu) = \Big(\frac{\lambda}{\mu}\Big)^{1/\omega_c}$$

$$c^w(\mu) = \Big(\frac{1-\lambda}{\mu}\Big)^{1/\omega_c}$$

With income as a function of labor supplies $Y(h^h, h^w) = h^h \cdot w^h + h^w \cdot w^w$.

Leisure can be expressed as a function of home production times:

$$l^s(x^s) = \Big(\frac{\Omega_l \mathrm{T}^{1-\omega_l}}{\Omega_x^S (x^s)^{-\omega_x}}\Big)^{1/\omega_l}$$

$$l^h(x^h, x^w) = \Big(\frac{\lambda\Omega_l \mathrm{T}^{1-\omega_l}}{\Omega_x^C (P(x^h, x^w|d^h, d^w))^{(\eta_x-\omega_x)}(x^h)^{-\eta_x}\kappa(d^h-d^w)}\Big)^{1/\omega_l}$$

with wife's leisure derived equivalently. Consequently, labor supply as a function of home production time:

$$h^s(x^s) = 1 - l^g(x^s) - x^s - b\cdot d(i, j^s)$$

$$h^g(x^h, x^w) = 1 - l^g(x^h, x^w) - x^g - b\cdot d(i, j^g)$$

Budget constraint multiplier and home production times (the leftover free variables) solve a system of equations (again for each job location and rent). For singles:

$$Y(h^s(x^s)) = c^s(\mu) + H^s(\mu)$$

$$0 = \frac{\Omega_l \mathrm{T}^{1-\omega_l}}{l^s(x^s)^{\omega_l}} - \mu \frac{\partial Y(h^s(x^s)) + \xi(h^s(x^s))}{\partial h^s}$$

For couples where both work:

$$Y(h^h(x^h, x^w), h^w(x^h, x^w)) = c^h(\mu) + c^w(\mu) + H^c(\mu)$$

$$0 = \frac{\lambda \Omega_l \mathrm{T}^{1-\omega_l}}{l^h(x^h, x^w)^{\omega_l}} - \mu \frac{\partial Y(h^h(x^h, x^w)) + \xi(h^h(x^h, x^w))}{\partial h^h}$$

$$0 = \frac{\lambda \Omega_l \mathrm{T}^{1-\omega_l}}{l^w(x^h, x^w)^{\omega_l}} - \mu \frac{\partial Y(h^w(x^h, x^w)) + \xi(h^w(x^h, x^w))}{\partial h^w}$$

When one in a couple does not work, an equation equalizing the marginal benefit of an hour of work to the marginal utility of leisure is replaced with a leisure equation:

$$0 = -l^h(x^h, x^w) + 1 - x^g - b \cdot d(i, j^g)$$

With solutions for each continuous variable, consumption, housing quantity, labor supply and home production hours for each location of residence, locations of work and labor force participation decision combination, I construct the respective values of discrete choices (participation, job acceptance and location). A single person always works, because their utility approaches infinity at 0 consumption. However, each single person given a choices also decides between taking an offered job in location j (with distance $d_{i,j}$) and a local job (with distance $d_{i,i}$).

For tractability, decisions on what job to take are made sequentially, as stochastic match shocks are revealed one at a time. The match shock ξ has a stochastic component drawn from a uniform distribution. First the local shock is revealed and a decision is made on whether to take it or not, comparing a local known job with the expected value of working in j. Than the originally offered location shock is revealed.[3]. Therefore, each single person chooses a job location by comparing $U^s(i, i, \xi_0^i)$ with $E_{\xi_0^j}(U^s(i, j, \xi_0^j))$. Before match shocks are revealed (when the residential location choice is made) the probability that a single person with an offer in j living in i works in j is given by $P_j' = (1 - \pi) + \pi P_j$, where $P_j = P(U^s(i, i, \xi_0^i) < E_{\xi_0^j}(U^s(i, j, \xi_0^j))$, and is solved for in closed form. A value for a single person with a j offer in hand for each residential location $V^s(i)|t, j$, including conditional expectations of a match shock, can be solved in closed form.

Assuming idiosyncratic preferences for locations ϵ_i have a standard extreme value

[3]This is done to keep matters tractable within couples, avoiding four shocks being realized at once. For singles, the order is without loss of generality, if both shocks come from uniform distributions with the same variance.

distribution, the share of each type of household choosing a location i can be solved in closed form (see McFadden (1977)). Households are differentiated by being a single or a couple, and than by labor market sector assignment and job offers locations of household members. Summing the housing demand over all types of households gives a total demand for housing in each location. This gives a system of three equations and three unknowns which I solve numerically.

A.3 Identification and estimation

In this section I discuss the construction of moments in the data that are used in calibrating and estimating the model and identification of model parameters from these moments. Table A.12 presents a complete list of parameters to be calibrated or estimated. I estimate the model with a moment based procedure. Table A.11 presents the list of data moments $\bar{m}$ used in the estimation routine. A subset of the parameters is calibrated outside the estimation routine. Moreover, a subset of the parameters α^1 is fit within the estimation routine – at each iteration using guesses of other parameters and moments in the data to fit an exact specific moment condition. This partition decreases the number of parameters that have to be searched for numerically, decreasing the computational burden in estimating the model. Letting $\alpha = [\alpha^1, \alpha^2]$ denote the $Bx1$ parameter vector, the estimation problem may be formally described as

$$\begin{aligned} \alpha &= \arg\min_{\alpha^2} \ [m(\alpha) - \bar{m}]^T W [m(\alpha) - \bar{m}] \\ \text{s.t. } \ \alpha &= [\alpha^1, \alpha^2], \alpha^1 = f(\alpha^2, \bar{m}) \end{aligned}$$

W is constructed based on the inverse of the variance-covariance matrix of the data. For moments from different samples I set the covariance to zero. For moments within the same sample I compute the variance-covariance matrix using influence functions of individual moments, and clustering at the MSA level. Moreover, I increase the weight of the most crucial moments (commuting moments, price gradient, job access difference between singles and couples).

A.3.1 Constructing moments

In this section I describe in detail the construction of moments used in estimating the model.

PSID main sample moments Most moments used in estimation and calibration come from a common PSID sample. The publicly available PSID data is linked to confidential identifiers of the census tract of residence. The sample is than restricted to include only people between 18 and 50 years of age, those about whom I can discern whether they have ever been in a couple, those who currently live in a metro-area of at least 250 thousand residents (by 2010) and for whom this is the metro-area they have spent the most number of periods in the PSID sample. Furthermore, I drop observations who have been married (or in a couple as identified in the PSID) before and are now observed as single or in a different couple. This is done so that differences between singles and couples are identified without characterizing divorcees as single, to match the notion of singlehood in the model. All statistics are computed using sample PSID weights (whichever available in each wave).

Average commute of singles d^s is an average commuting distance in miles for those identified as having never been in a couple. $d^s_h - d^h$ and $d^s_w - d^w$ are quantified by running two separate regressions by gender, that control for metro-area, age, education, race and PSID wave dummies. Moreover, in the regressions comparing singles to couples I only use singles that are later (at any point in the future PSID samples) observed in couples. This results in slightly smaller differences between singles and couples, thus choosing a more conservative measure. All hours of work moments ($h^h_{\text{both work}}$-h^s and others) are computed using annual hours of market work. All moments describing hours of home production use data on annual hours of housework as defined in the PSID. Labor force participation is defined as one if an individual worked over the last year at all, and zero otherwise. Differences between two groups are always quantified using a simple regression with the controls as listed above, only using the samples of the two groups being compared.

$P(city|couple) - P(city|single)$ is quantified from a regression of a dummy variable of living in a tract that is less than 10 miles away from the center of the biggest city of a metro-area, using a regression with the controls listed above, and again restricting the sample to exclude single people that are never observed to couple up.

Moments describing a distance to jobs d_j and distance to opportunities d_o were also computed using this sample, except only restricting to waves since 1990, to avoid unnecessary imputation. Construction of these variables are described in the main text. Distance between two random jobs is first computed on the metro-area/year/industry and earnings segment level using the LEHD Origin-Destination Employment Statistics aggregated to a census tract level. They are than matched to individuals in the PSID sample (on metro-area and year, with 2002 being used for PSID waves where no LODES data are available, most common industry and earnings segment). The statistics are than computed on this sample.

Distance between actual jobs of a husband and wife were computed using the job-location census tract identifiers, computing the euclidean distance between the centroids. This information is only available in waves 2013, 2015 and 2017.

PSID moments identified from within-couple variation This set of moments is computed on the sample described above, except that only couples are used and remarried couples are included to increase sample size. All moments in this section are based on within-couple differences, as they are computed using regressions with couple fixed effects. β_b^a is a set of moments mimicking the analysis in tables 1.5 and 1.6, where a denotes the left-hand side variable and b stands for either d or wd, with d marking coefficients on d_o and wd marking coefficients on the interaction term woman $\cdot\, d_o$. a stands for $comm$ (commuting distance in miles), $hours+$ (annual hours of work for those who did any market work last year), $hours$ (annual hours of work including zeros), $work$ (labor force participation), x (annual hours of housework) and $log(w)$ (log of the ratio of annual labor income and annual hours). For all variables except for $comm$ only waves since 1990 were included. The details of this analysis are described in the main text.

$\log(\frac{w^w}{w^h})$ is a measure of gender-wage gap among people in couples computed using within-couples variation. Wage is defined as the ratio of annual labor income and annual hours. The same controls as listed above are included. I also add an interaction between education groups and industries to capture as much as possible the differentiation into different kinds of jobs.

Moments identified in external data $\hat{\lambda}_0$, as described in table A.1, is computed using IPUMS 2000 Census and 2006-2010 ACS (Ruggles et al., 2019). The same sample is used to compute the 'share never married', defined as the ratio between people never married and not cohabiting over all people, in the age-range 30-50. The goal is to use a measure describing a share of population that never ends up married, as of a certain age. This matches the nature of singlehood in periods 2 and 3 in the model.

Next I use NHGIS census-tract data (Logan et al., 2016) from the 2010 Census to compute housing rent gradients. I define $\log(p)$ in the data as the log of the ratio between the median rent in the census tract over the median number of bedrooms in the census tract. I than compute the difference between $\log(p)$ for tracts less than 10 miles away from the center and the rest. Moreover, I compute the $\log(p)$ gradient with the distance to an average job (d_j). This sample is also used to compute the share of overall population living less than 10 miles away from city center. For comparability, I use the 2010 slice of the LEHD Origin-Destination Employment Statistics (LODES) available in 2002-2017 (U.S.CensusBureau, 2021) aggregated to the Census-tract level, to compute the share of jobs located less than 10 miles away from the center of the largest city in the metro-area, restricting to metro areas with at least 250 thousand resident.

I match the industry and earnings segment groups as defined in the LODES data with the measures of industry and labor income from 2006-2010 ACS and 2000 Census IPUMS data, restrict the sample to the age group 18-50. To calibrate the level of gender segregation in the labor market I compute the share of one's own gender in ones own industry and earnings segment group. I take the expenditure share on housing from the 2019 Consumer Expenditures Report (U.S.BureauofLaborStatistics, 2020). Table A.11 shows the list of moments as well as fit. h stands for annual hours of work, x for home production hour, d for commuting distance in miles, w for wage, d_j for distance between residence and a random job, d_o for distance between a residence and a random job in the person's mode industry and earnings segment, p for rent per unit of housing.

Moment	Model value	Data value	Directly used to fit a parameter	Group
Average commute of single d^s	8.6422	8.6669	0	1
$d^s_h - d^h$	-2.5816	-2.7085	0	1
$d^s_w - d^w$	-0.3066	-0.2973	0	1
$h^h_{\text{both work}}$	2143.8515	2206.5421	0	2
$h^w_{\text{both work}} - h^h_{\text{both work}}$	-711.0791	-671.5465	0	2
$h^h_{\text{just husband works}} - h^h_{\text{both work}}$	196.1485	63.1028	0	2
h^s	1893.0719	1873.1261	0	2
$x^w_{\text{both work}}$	1010.8694	973.9142	1	3
$x^w_{\text{just husband works}} - x^w_{\text{both work}}$	691.3780	683.0032	1	3
$x^h_{\text{both work}} - x^w_{\text{both work}}$	-620.9176	-604.5061	1	3
$x^h_{\text{just husband works}} - x^h_{\text{both work}}$	-115.1224	-50.5242	1	3
x^s	492.0368	495.3415	1	3
LFP of wives-husbands	-0.1589	-0.2207	0	2
LFP of husbands	0.9433	0.9738	0	2
$\log(\frac{w^w}{w^h})$	-0.2365	-0.2440	0	2
Expenditure share on housing	0.2032	0.1930	0	NaN
Share of population in city	0.3400	0.3919	0	4
Share of jobs in city	0.4960	0.4978	0	4
Distance to an average job for a couple (d^h_j)	21.7292	20.2769	0	4
Distance between 2 random jobs	18.8111	17.3000	0	4
Dist between 2 jobs of the husbands sector	18.6323	16.2667	0	4
$\lvert d^w_o - d^h_o \rvert$	1.7847	1.8622	0	5
Dist to a random job in own sector for husband (d^h_o)	21.6666	20.0266	0	4
$P(city\|couple) - P(city\|single)$	0.1325	0.0704	0	5
$\log(p)$ distance to jobs gradient	-0.0060	-0.0088	0	5
$\log(p)$ city over suburb	0.0747	0.0728	0	5
$d^s_o - d^h_o$	-1.6941	-1.6931	0	5
$d^s_j - d^h_j$	-1.6400	-1.4099	0	5
$d^w_o - d^h_o$	0.0458	0.0279	0	5
Distance between husbands and wives actual jobs	12.9405	9.7405	0	4
β^{comm}_{wd}	-0.0899	-0.1054	0	6
β^{work}_{wd}	-0.0046	-0.0016	0	6
β^{hours+}_{wd}	-2.4803	-2.4845	0	6
β^{hours}_{wd}	-6.9423	-5.1460	0	6
β^{x}_{wd}	4.3140	3.2846	0	6
β^{comm}_{d}	0.9471	0.7064	0	6

β_d^{work}	-0.0043	-0.0019	0	6
β_d^{hours+}	-20.1448	-4.9347	0	6
β_d^{x}	1.7801	-1.0640	0	6
$\beta_{wd}^{\log(w)}$	-0.0009	-0.0012	0	6
β_d^{hours}	-25.0208	-5.4621	0	6
$\beta_d^{\log(w)}$	-0.0075	0.0015	0	6
Share never married	0.1449	0.1449	1	3
$\hat{\lambda}_0$	0.5494	0.5360	0	3

Table A.11: Moments used in estimation: data versus model. In the data, a city is defined as a radius around city center of 10 miles. In addition, I use a ratio of average commuting time and distance in miles as a scaling factor, I constraint housing prices to be one on average, and I impose that the ratio of men and women in the metro-area is equal to one. Lastly, the distribution of men and women in the two labor market matches that the share of ones own gender in ones labor market in the data is 0.59 percent.

Parameter	Value	Fit directly	SE	t	Groups
Parameter	Value	Fit directly	SE	t	Groups
ϕ	2.652	0	(0.997)	2.660	G1, G4, G5, G6
$D(1,2) = D(1,3)$	24.951	0	(1.726)	14.452	G1, G2, G4, G5, G6
$D(3,2)/D(1,2)$	2.000	0	(0.335)	5.970	G1, G2, G4, G5, G6
$(1\|.)/f(2\|.)$	2.329	0	(0.358)	6.511	G4, G6
$f(2\|1)/f(3\|1)$	1.560	0	(0.179)	8.737	G4, G6
$A(1)$	-0.050	0	(0.048)	-1.047	G4, G6
$A^c(3) = A^c(2)$	0.250	0	(0.106)	2.361	G1, G4, G6
σ_{ϵ_i}	0.453	0	(0.197)	2.302	G1, G4, G6
κ_d	0.488	0	(0.813)	0.600	G1, G4, G5, G6
Ω_l	10.444	1	(9.470)	1.103	G1, G4, G6
ω_c	1.190	0	(0.145)	8.203	G1, G4, G6
ω_l	1.919	0	(0.595)	3.226	G1, G4, G6
Ω_x	0.490	1	(0.045)	10.851	G1, G2, G3, G4, G5, G6
$\bar{\kappa}_w$	0.546	1	(0.025)	21.927	G4, G6
η_x	0.439	1	(0.087)	5.052	G4, G6
ω_x	0.610	1	(0.148)	4.111	G4, G6
Ω_x^s	0.402	1	(0.042)	9.507	G1, G2, G3, G4, G5, G6
w_{gap}	-0.232	0	(0.022)	-10.361	G2, G4
π	0.578	0	(0.064)	9.036	G4, G6
$E(\xi_0)$	0.238	0	(0.058)	4.135	G4, G6
$Var(\xi_0)$	0.395	0	(0.085)	4.659	G4, G6
$\bar{\Xi}$	4.998	0	(2.593)	1.928	G4, G6
w_{Ξ}	7.700	0	(2.549)	3.021	G4, G6
Θ	1.413	1	(0.152)	9.278	G4, G6

λ	0.549	0	
σ_{θ_i}	1.000	1	
$\frac{\Omega_H}{\Omega_c}$	0.197	0	
w_a	24.383	1	
b	0.002	1	
T	5.181	1	
$\bar{h}^U$	0.267	1	
$\frac{\bar{N}^C_{1,1}+\bar{N}^C_{1,2}}{\sum_{u,v}\bar{N}^C_{u,v}}$	0.706	1	

Table A.12: Model parameters. Baseline parameter values, when appropriate standard errors and t-statistics and groups of moments that the parameter is sensitive to (using the measure by Andrews et al. (2017) rescaled by the moments standard deviation and highlighting all groups with at least one moment with sensitivity of at least 10 percent of the maximum). Parameters in the lower part of the table are calibrated outside the estimation routine. Parameters in the upper part with 1 indicated in the third column are fitted directly within the estimation routine to satisfy a particular moment equation.

A.3.2 Identifying parameters

Identification of λ The bargaining weight λ, though technically a price, is treated in practice as a parameter to be estimated (because it is not observed in any form).

While it is in any case identified practically within the model as a market clearing price, solving the marriage market equation given all the other parameters, it is useful to think about external sources of variation for this unobservable number. I build on the identification argument presented in Gayle and Shephard (2019). Given the assumption that allocations within couples are Pareto efficient and λ is constant, equation A.1 presents a useful condition on the value of the bargaining weight (where u^h is the utility in marriage for a man)

$$\frac{\partial u^h(\lambda)}{\partial \lambda} = -\frac{(1-\lambda)}{\lambda}\frac{\partial u^w(\lambda)}{\partial \lambda} \tag{A.1}$$

Given marriage market clearing, $\log(M^m) - \log(M - M^m) = \frac{1}{\sigma_{\theta_i}}(u^h(\lambda) - u^{h,s})$ and $\log(F^m) - \log(F - F^m) = \frac{1}{\sigma_{\theta_i}}(u^w(\lambda) - u^{w,s})$. Assume there is a variable X, that has no impact on the value of the single state and only affects the value in marriage through its influence on the Pareto weight, aka a distribution factor in the sense of Bourguignon et al. (2009). A marginal perturbation in the distribution factor thus gives

$$\frac{\partial(\log(M^m) - \log(M - M^m))}{\partial X} = \frac{1}{\sigma_{\theta_i}}\frac{\partial u^h(\lambda)}{\partial \lambda}\frac{\partial \lambda}{\partial X}$$

$$\frac{\partial(\log(F^m) - \log(F - F^m))}{\partial X} = \frac{1}{\sigma_{\theta_i}} \frac{\partial u^w(\lambda)}{\partial \lambda} \frac{\partial \lambda}{\partial X}$$

Notice the left-hand side is potentially observable. Taking a ratio of these two derivatives thus provides an estimate of the ratio of marginal values of husband versus wife. A typical example of such a distribution factor is a variation in the available supply of men and women M/F.[4]

Thus, I collect $\frac{M}{F}_k$ for a set of metro-areas and years as well as the share of both men and women who are single (s_k^g for $g \in m, f$) and run the following regression

$$\log\left(\frac{1}{s_k^g} - 1\right) = A \cdot \frac{M}{F}_k + B \cdot 1_{g=m} \cdot \frac{M}{F}_k + u_{k,g}$$

Table A.13: Responsiveness of staying single long-term to sex-ratios across US metro areas.

	$\log\left(\frac{1}{s_k^g} - 1\right)$
$\log(\frac{M}{F}_k)$	1.002
	(.317)
$\log(\frac{M}{F}_k)\cdot$ Men	-1.869
	(.162)
$\hat{\lambda}_0$	**0.536**
X_i:	
Polynomials up to 4th order of $\log(\frac{M}{F}_k)$	x
Religious participation by denomination 2000, 2010	x
Vote shares in presidential elections 1996-2012	x
Polynomial of size of MSA, year fes	x
N clusters	166

SEs in parentheses.

$s_k^g = 1 - \frac{\text{married or currently in couple}}{\textit{all}}$ for $g \in h, w$ stands for men or women, in an age range of 25-45. Source of data: 5% IPUMS Census 2000 and 2006-2010 IPUMS ACS, MSAs with at least 250k residents by 2010. $\frac{M}{F}_k = \frac{\text{all men}}{\text{all women}}$, in an age range of 25-45. All controls are also included as interacted with gender

If $\frac{M}{F}_k$ is a distribution factor, $\hat{\lambda} = -\frac{\hat{A}}{\hat{B}}$ could be used as a direct calibration of λ.[5] In this paper, however, $\frac{M}{F}$ affects the relative value of marriage through more than λ. This is because there is a housing market as well as a marriage market. $\frac{M}{F}$ affects the overall share of people being single, thus demand for housing in different locations. Moreover, a change in λ implied by a change in the sex-ratio changes the decisions of couples, impacting their income and thus housing demand.

[4]Gayle and Shephard (2019) use this argument to identify bargaining power from a variation across the population vectors M and F across several marriage markets.

[5]With $c = \frac{\frac{\partial u^h(\lambda)}{\partial \lambda}}{\frac{\partial u^h(\lambda)}{\partial \lambda}} = \frac{A+B}{B}$, $\lambda = \frac{1}{1-c} = -\frac{A}{B}$

Specifically,
$$\frac{\partial u^g(\lambda) - u^{g,s}}{\partial X} = \frac{\partial u^g(\lambda)}{\partial \lambda}\frac{\partial \lambda}{\partial X} + \frac{\partial u^g(\lambda) - y^{g,s}}{\partial p}\frac{\partial p}{\partial X}$$

For the exact identification to be preserved, it would have to hold

$$\frac{\frac{\partial u^h(\lambda)}{\partial \lambda}\frac{\partial \lambda}{\partial X} + \frac{\partial u^h(\lambda) - u^{h,s}}{\partial p}\frac{\partial p}{\partial X}}{\frac{\partial u^w(\lambda)}{\partial \lambda}\frac{\partial \lambda}{\partial X} + \frac{\partial u^w(\lambda) - u^{w,s}}{\partial p}\frac{\partial p}{\partial X}} \sim \frac{\frac{\partial u^h(\lambda)}{\partial \lambda}}{\frac{\partial u^w(\lambda)}{\partial \lambda}} = \frac{\lambda - 1}{\lambda}$$

Thus, I instead collect $\hat{\lambda}_0 = -\frac{\hat{B}}{A}$ as one of the moments that I recreate within the model (using a numerical derivative with respect to $\frac{M}{F}$) resolving the housing and marriage market equilibrium, and collecting the implied changes in the share of single men and women) and use it in estimation. Since $\hat{\lambda}_0$ is now only one of the moment, I also take into account that it is just an imprecise estimate, weighting the associated uncertainty against that of other moments.

Numerically, $\hat{\lambda}$ in the model is very close to the actual λ, suggesting that this identification strategy is still sound even with adding the housing market clearing.[6]

Identification of parameters First, I describe parameters calibrated outside of the estimation procedure. There are two parameters that function as scaling factors. b is a scaling factor between distance in miles and annual hours of time (rescaled to be between 0 and 1). b is calibrated outright to match the ratio between average annual commuting time and average commuting distance in miles in the PSID. Second, I scale the base wage w_a so that prices of housing are around 1[7] and I rescale time inputs in home production and utility so that the average leisure would be close to the average consumption quantity, and both aspects of utility were measured in comparable ranges.[8] Men and women in the model systematically work in different kinds of jobs. Men work more in the labor market that has more jobs in the first suburb. The extent of gender segregation in the labor market is calibrated to match the share of workers of one's own gender in their industry and earnings segment group (as defined in the LODES dataset, the definition of one's labor market in

[6]First, higher bargaining power of husbands allows them to work less. Households have less income on average, housing demand falls, and prices in all neighborhoods fall. This however affects couples and singles about equally. Second, sex ratio different from 0.5 results in a lower overall marriage rate. More singles put pressure on the housing price in the city, favoring marriage over singlehood. This effect, however, is quantitatively minuscule.

[7]Specifically, I utilize moments describing average hours, gender wage-gap in couples, share of couples versus singles, share of couples where both work and share of income spent on housing to have an average demand for housing equal to 1 if price of housing is 1. Since the modeled metro-area has a fixed supply of housing of one unit per person, this ensures equilibrium prices are averaging around 1, whenever the model matches the other moments mentioned.

[8]Specifically, inputs in u_x and u_l are multiplied by T equal to average income per capita.

the data for this paper). The last parameter calibrated outside the estimation routine is an upper bound on annual hours $\bar{h}^U$ equal to 0.2671 rescaled annual hours (equivalent to 9 hours a day, 5 days a week, 52 weeks a year). Having an upper bound on hours allows the model to match the fact that hours of husbands only increase so much when the wife drops out of the labor force. σ_{θ_i} cannot be separately identified and is set to be equal to 1.[9][10]

Out of 9 parameters governing preferences for time, consumption and housing quantity, 6 are fit directly within the estimation routine. η_x, ω_x, Ω_x Ω_x^s and $\bar{\kappa}_w$ (parameters governing home productivity of time and preferences over the resulting public good) as well as Ω_l (the scale of the leisure preference) can be expressed in closed form from first order conditions in the optimization problems of different types of households as functions of allocations of time and other parameters. I replace individual time choices in the first order conditions with their averages from the data and use the current guess of other parameters (ω_l, λ) within the estimation routine. Let L^h, x^h, L^w, x^w be the average leisure and housework hours of husband and wife when both work and $L_0^h, x_0^h, L_0^w, x_0^w$ be the average leisure and housework hours of husband and wife when only husband works.[11] Then combining first order condition of couples where both work versus those where only the husband works, I set

$$\eta_x = \omega_l \cdot \frac{\log(\frac{L_0^w \cdot L^h}{L^w \cdot L_0^h})}{\log(\frac{x^h \cdot x_0^w}{x_0^h \cdot x^w})}$$

$$1 - \kappa_w = \frac{\frac{\lambda}{1-\lambda} \cdot \left(\frac{x^h}{x^w}\right)^{\eta_x}}{\frac{\lambda}{1-\lambda} \cdot \left(\frac{x^h}{x^w}\right)^{\eta_x} + \left(\frac{L^h}{L^w}\right)^{\omega_l}}$$

From there I get the base level of $\bar{\kappa}_w = \kappa_w - b \cdot \kappa_d \cdot (d^h - d^w)$. Using the above, I compute average values of home production time (when both husband and wife works, and when only husband works): $X = (\kappa_w (x^w)^{1-\eta_x} + (1-\kappa_w)(x^h)^{1-\eta_x})^{\frac{1}{1-\eta_x}}$, $X_0 = (\kappa_w (x_0^w)^{1-\eta_x} + (1-\kappa_w)(x_0^h)^{1-\eta_x})^{\frac{1}{1-\eta_x}}$ to get

$$\omega_x = \frac{\omega_l \cdot \log\left(\frac{L_0^h}{L^h}\right) + \eta_x \cdot \log\left(\frac{X_0}{X}\right) - \eta_x \cdot \log\left(\frac{x_0^h}{x^h}\right)}{\log(\frac{X_0}{X})}$$

[9] Importantly, this parameter does not affect any moments used in estimation, except for $\hat{\lambda}_0$, theoretically. However, quantitatively, the effects of σ_{θ_i} on $\hat{\lambda}_0$ are minuscule as well. Thus calibrating this parameter at an arbitrary level does not affect the estimates of the rest of the model.

[10] Alternatively, constraining Θ to 0 would allow identification of σ_{θ_i}.

[11] Using moments in the data as presented in table A.11: $x^w = x^w_{\text{both work}}$, $x^h = x^w + (x^h_{\text{both work}} - x^w_{\text{both work}})$, $x_0^w = x^w + (x^w_{\text{just husband works}} - x^w_{\text{both work}})$, $x_0^h = x^h + (x^h_{\text{just husband works}} - x^h_{\text{both work}})$, $h^h = h^h_{\text{both work}}$, $h^w = h^h + (h^w_{\text{both work}} - h^h_{\text{both work}})$, $h_0^h = h^h + (h^h_{\text{just husband works}} - h^h_{\text{both work}})$, $L^h = 1 - h^h - x^h - b \cdot (-(d^s - d^h) + d^s)$, $L^w = 1 - h^w - x^w - b \cdot (-(d^s - d^w) + d^s)$, $L_0^h = 1 - h_0^h - x_0^h - b \cdot (-(d^s - d^h) + d^s)$, $L_0^w = 1 - x_0^w$.

$$\Omega_x = \frac{\lambda\Omega_l}{1-\kappa_w} X^{(\omega_x-\eta_x)} (x^h)^{\eta_x} \frac{1}{\mathrm{T}\cdot L^h}^{\omega_l}$$

For singles, I allow for a different value of home production derived from time Ω_x^s, calibrated equivalently to Ω_x, and average time allocations of singles from the data. With $L^s = 1 - h^s - b\cdot d^s - x^s$:

$$\Omega_x^s = \Omega_l (x^s)^{\eta_x} \frac{1}{(\mathrm{T}\cdot L^s)^{\omega_l}}$$

The scale of preference for leisure is fit to match average hours performed by singles:

$$\Omega_l = \frac{w_a}{(Y^s-1)^{\omega_c}} \frac{L_s^{\omega_l}}{\mathrm{T}^{1-\omega_l}}$$

In table A.11 the third column indicates the moments that are implicitly directly targeted by this calibration.

Ω_H, ω_c, ω_l are left to be estimated. Ω_H is identified of the share of income spent on housing. A set of moments pertaining to average hours (of husbands and wives, when both work and when only one works) as well as the fact that preferences of individuals in and out of couples are constrained to be the same (except for the value of home production) allow the identification of ω_c and ω_l.

The wage function $w(d_o'(j), g) = w_a \cdot e^{-w_{\Xi}\cdot(b\cdot d_o'(j) - b\cdot \bar{d}_o') + 1_{g==h} w_{gap}/2 - 1_{g==w} w_{gap}/2}$ is decreasing in distance to other jobs in the worker's own labor market, and is lower for women in couples. w_{gap} is estimated to match the observed within-couple gender wage gap in the PSID sample.

Θ is a baseline snifter for the the value of marriage and thus is identified by the share of men and women choosing marriage over perpetual single-hood after the first period (and the fact that the sex-ratio in the metro-area is fixed at 1). It is also a parameter fit directly within the estimation routine.

Next I describe the parameters governing the spatial structure of the modeled metro-area: the distance matrix D and the distribution of job offers f. It is important to identify these parameters separately from commuting behavior, as the parameters of commuting costs are the focus of this paper. The metro-area has 3 locations and is shaped as a triangle. The parameters to be estimated are the distance between suburbs and city $D(1,2) = D(1,3)$ and the distance between the two suburbs (compared to the distance to the city) $D(3,2)/D(1,2)$. There are two labor markets, one offers more jobs in the first suburb, one in the second. The parameters to be identified are the number of jobs (of both types) offered in the city compared to the suburbs $f(1)/f(2)$ and the degree of specialization of each suburb $f(2|1)/f(3|1) = f(3|2)/f(2|3)$. Share of jobs observed in the city (i.e. located less than 10 miles away from the center of the metro-area) identifies the share of job offers in the city. I use a distance between two random jobs to identify the distance between neighborhoods. In addition, I include the distance between two jobs in the same labor market. The difference identifies the degree of concentration of different kinds of jobs in different parts of the metro-area. Increasing $D(3,2)/D(1,2)$ helps to match how much distance to an average job is lower in the suburb than in a city, thus helping to match $d_o^s - d_o^h$. The shape of the metro-area and the distribution of job offers also define the potential for disagreement within couples about whose job offer to locate close to. This is measured by the absolute value of the difference in the distance to opportunities within couples between a husband and a wife $|d_o^w - d_o^h|$, which is also included as a moment in estimation.

Next I describe identification of preferences governing location choices. These include the vector amenity values for singles and couple A^c and A^s, as well as the dispersion of idiosyncratic location preferences σ_{ϵ_i}. Again, it is important to identify these parameters separately from commuting behavior driven by acceptance of different kinds of jobs. First, I constrain $A(2)^s = A(3)^s = 0$ and $A(1)^s = A(2)^c$. Constrains here are necessary. Adding a constant to both A^c and A^s results in exactly the same choices. Similarly, the same differences between couples and singles can be achieved by manipulating any two of the location values. Amenities preferences are identified as residuals – after the value of access to opportunities is taken into account, amenities match the difference in the share of singles versus couples who live in the city, and the price gradient between city and suburbs. Specifically, $A^c(3)$ matches $P(city|couple) - P(city|single)$, $A^c(2) - A^c(3)$ makes $d_o^w - d_o^h$ fit and $A(1)$ matches the price-gradient moments. σ_{ϵ_i} can be identified with the difference between the distance to an average job and the distance to an average job in own labor market – lower dispersion in idiosyncratic preferences matches a higher tendency to sort to the offered job location and potential other good offers.

The value of not commuting is governed by two aspects, the value of time (as governed by the preferences identified above) and the household value of being close to home. ϕ is than identified from the difference in commuting between husbands and single men, and between husbands and wives. The overall level of commuting is identifying π, the share of households who get a local job offer in addition to their initial offered job. Intuitively, there has to be a barrier on how many people are offered a local job wherever they live, so that commutes in the model are large enough to match the data.

Moreover, the model includes a specific interaction between commuting and the productivity of time in home production. Specifically, $\kappa_w(d^h - d^w) = \kappa_w^0 + (d^h - d^w) \cdot \kappa_d$, i.e. wives are more productive at home compared to husbands whenever their commute is short compared to their husbands. Note that individuals who do not work have an implicit commute of 0. κ_d is identified from β_{wd}^x and β_d^x, moments in the data that measure how much housework hours change depending on the difference in the distance to opportunities between the husband and the wife, as described in table 1.6, as well as from the within couple gender difference in commuting.

Each job comes with a non-monetary benefit $\xi = \xi_0 + e^{-\bar{\Xi}\cdot(b\cdot d'_o(j)-b\cdot\bar{d}'_o)\cdot h}$, where ξ_0 is a random variable with a uniform distribution. The parameters to be identified are $\bar{\Xi}$, $E(\xi_0)$ and $Var(\xi_0)$. Moreover, wages are decreasing in the distance to other jobs in your labor-market, through a parameter w_{Ξ}. I include $\beta_{wd}^{\log(w)}$, a coefficient estimate presented in table 1.6, measuring how much within couples a woman's wage is more affected by the couple living far away from other jobs in the wife's labor market than a husband's wage would be. In the model, women take local jobs more often, not taking an advantage of a job that is in a sector hub. Thus when the couple locates far away from the offers in her labor market, her wage does fall more. Similarly, $\bar{\Xi}$ is identified with β_{wd}^{hours+}, a coefficient measuring how much more wife's hours fall when they live far away from opportunities. The higher $\bar{\Xi}$ is, the more jobs outside of sector hubs offer a worse match that scales with hours. Since women in couples are more prone to accept a local job instead of a good job, their hours decrease the more.

To further help estimate the interlink between access to opportunities and labor market behavior I also include the average distance between the actual job of husband and wife (when both work), as well as the other estimates of sensitivity to being far away from opportunities for husband and wives, as shown in tables 1.5 and 1.6. Lastly, I include a share of overall population in the city as a moment to be matched.

In the last column of table A.11 I classify moments into broad groups: commuting, time use and marriage, distribution of jobs and people in space, location choices and sensitivity to opportunities within couples (βs). The last column of table A.12, I compute the sensitivity of each parameter to the moments in the estimation using the measure proposed by Andrews et al. (2017). For each moment and each parameter I compute

$$|Sensitivity| = |-(G'WG)^{-1}G'W|$$

where W is the estimation weighting matrix and G is the numerical derivative of moments with respect to parameters evaluated at the estimated values. Given the scale of the moments is not always comparable, I multiply each element by the standard deviation of the moment (as recommended by Andrews et al. (2017)). For each parameter, I calculate the moment with maximum sensitivity, and consider any moment whose sensitivity is at least 10% of the maximal as being important. As I consider sets of moments, I describe a set as being important if at least one moment from that set is important according to this criterion.

A.3.3 Fit without commuting costs rewarding specialization

In this section I further evaluate the role of the value of working close to home that rewards specialization $F(d^h, d^w)$ in the ability of the model to fit commuting patterns. Specifically, I set $F^g() = 0$ for both singles and couples, and partially re-estimate the model. I keep parameters that fit the structure of the metro area as well as preference parameters fixed, overestimating the share of households receiving a local offer (π), other aspects of jobs ($E(\xi_0)$, $Var(\xi_0)$) as well as amenities ($A(1)$ and $A^c(3) = A^c(2)$) and κ_d – the complementarity between commuting and home production time governing gender-differences in commuting within couples, with resolving the marriage market at a new λ.

Table A.14: Commuting moments fit without a cost of commuting rewarding specialization

Moment	Model value	Data value
Average commute of single d^s	8.670	8.667
$d^s_h - d^h$	**-1.112**	**-2.708**
$d^s_w - d^w$	-0.101	-0.297

Table A.14 shows the fit on the commuting moments, in particular that the increase in commuting for men after they form a couple is not matched as well.

A.4 Other counter-factuals

A.4.1 Work from home

Recently, the COVID pandemic reinvogorated the discussion about the benefits of allowing employees to work from home. In this section, I allow a fixed share of the households to keep whatever job characteristics they were offered, but set their commutes to 0.

Within the joint housing and marriage market equilibrium this generates several endogenous responses. Across all groups there is less commuting. Wives work from home more often than husbands. However, because at baseline husbands commute the most, their overall commuting falls the most, pushing down gender gaps in commuting. Because commuting costs discourage women in couples from the labor force, work from home options reduce gender gaps in working. As commuting costs are partially alleviated across households, everybody is more willing to live in the suburbs. The change in incentives is even stronger for singles than couples. As a result, housing costs in the suburb increase and housing costs paid on average by couples increase, compared to singles.

All benefit from work from home options, here under the assumption that it does not take away from either your monetary or non-monetary benefits from working. However,

Figure A.4: Offering a share of the households a work from home option – commuting and gender gaps.

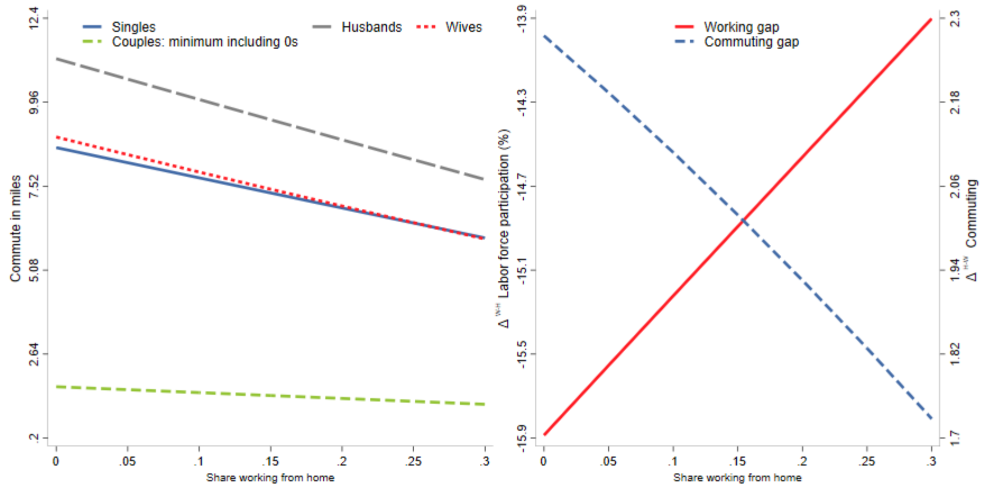

Figure A.5: Offering a share of the households a work from home option - housing costs and sorting.

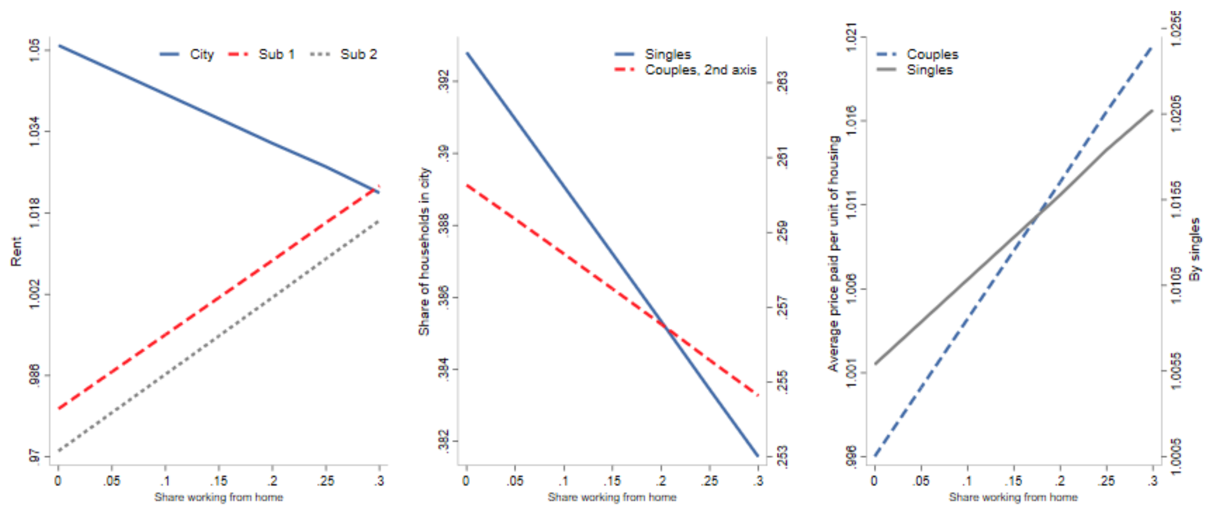

Figure A.6: Offering a share of the households a work from home option - change in welfare for singles and couples, the share of people choosing marriage, bargaining position of husbands.

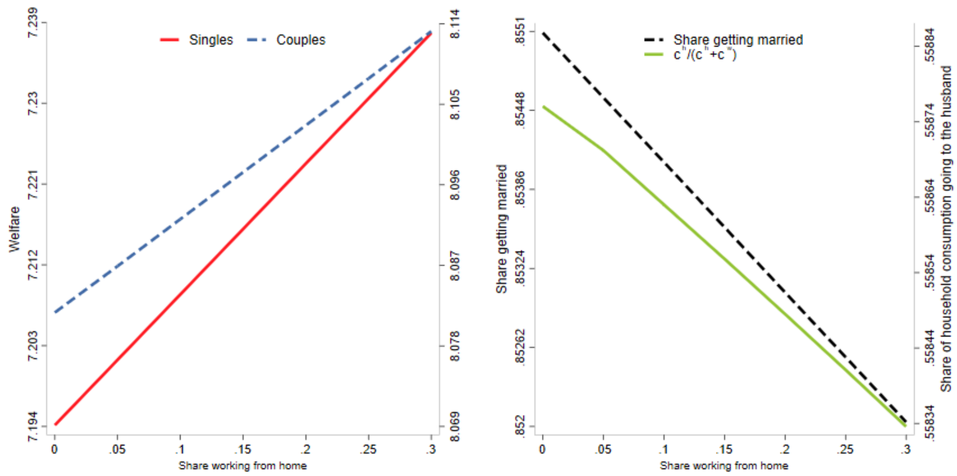

an environment with long commutes is comparatively less costly for couples. Therefore, a widespread work from home option is more appreciated by singles than couples. This is true despite the fact that work from home allows women in couples to come back to the labor force more. As a result, fewer people choose marriage in this metro area.

Figure A.7: Offering a share of the households a work from home option - change in the value of marriage, decomposed into parts of utility.

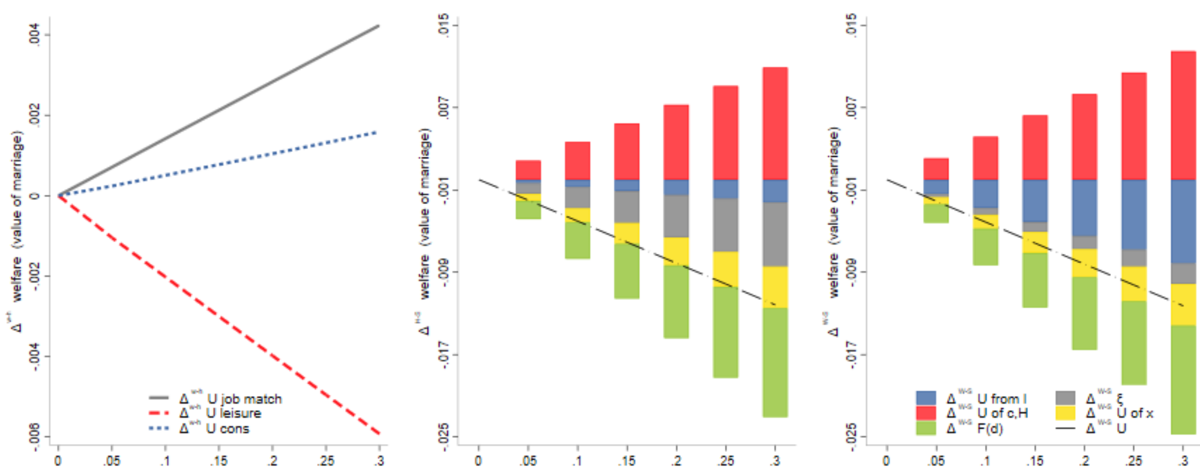

Work from home options present a discrete jump in commuting – a small share of the population can go from a very long commute to a zero commute. However, the rest of the population is still facing the old commuting costs. This heterogeneity implies that some of the women who would like to come back to the labor force still cannot, while some of the women who do come back from the labor force do so even in jobs that are not a good match and bring negative non-monetary benefits. As a result, this policy change is actually not on average appreciated much by wives and results in a renegotiation in the marriage market towards a lower bargaining weight of husbands. If, however, a work from home option was given to everybody but perhaps only for a day a week, women could better sort into jobs and the bargaining weight of husbands within the marriage market equilibrium would increase.

Overall, work from home options make marriages less valuable. This is because the technology of specialization that couples possess is less needed while suburban living is more expensive, so couples do not benefit as much from specialization and lower housing prices within a housing and marriage market valuable. Figure A.7 shows how the decline in the value of marriage for men and women is decomposed into various aspects of utility. Husbands lose some of their privilege of better jobs, because now singles can access them just as easily. Wives lose in terms of leisure compared to single women as they join the labor force, but are somewhat compensated with consumption as they bargain over a higher share of household resources with their husbands.

It is important to know, however, that work from home likely brings complexities not considered by this analysis. For example, especially the non-monetary job benefits (such as socialization and enjoyment) might be both valued more by singles and be diminished by working from home, undermining the idea that work from home benefits singles more than couples. Moreover, Ozimek and Carlson (2023) shows that work from home also prompted more household formation as individuals increased their demand for individual space, pushing up prices even in central cities.

A.4.2 Decline in overall willingness to marry – orthogonal to commuting costs

In recent decades we have seen a marked decline in the share of population getting married, especially through increasing the age at marriage. I show that a natural implication of a decline in marriage is gentrification – a steepening of the distance price gradient. Since couples and singles differ in their location choice preferences, a change in the composition of the metro area results in a shift in demand for city versus suburbs. I model an exogenous decline in marriage as a decline in Θ, the constant gender-neutral shifter to the value of marriage. Individual marries if $u^g(\lambda) + \theta^g > u^{s,g} + \theta^s$, with $\theta^g \sim$ EV1 and

$$u^h(\lambda) = E_{T^g,\epsilon_i,j_t^g,\xi_{i_t}^g,\xi_{j_t^g}^g}\left(\sum_{t=1}^{2} u_c(c_t^h) + u_l(l_t^h) + u_H(H_t/2) + a^c(i) + \Pi(x_t^h, x_t^w, d_t^h, d_t^w) + \xi_t^h\right) + \Theta$$

and

$$u^{s,g} = 2 \cdot E_{T,\epsilon_i,j,\xi_i,\xi_j} \left(u^s_c(c) + u_l(l) + u_H(H) + a^s(i) + \Pi^s(x,d) + \xi \right)$$

Lifetime utility in marriage thus has a constant shifter, Θ that is identified to match the share of people married by a certain age. If Θ declines, the share of population wanting to get married declines (proportionally for men and women). Figure A.8 shows that this is associated with a marked increase in the price of housing in the city, without much change to the price in the suburbs. In other words, an exogenous decline in the willingness to get married causes city centers to become more expensive compared to suburbs.

Figure A.8: Counter-factual simulation: declining value of marriage.

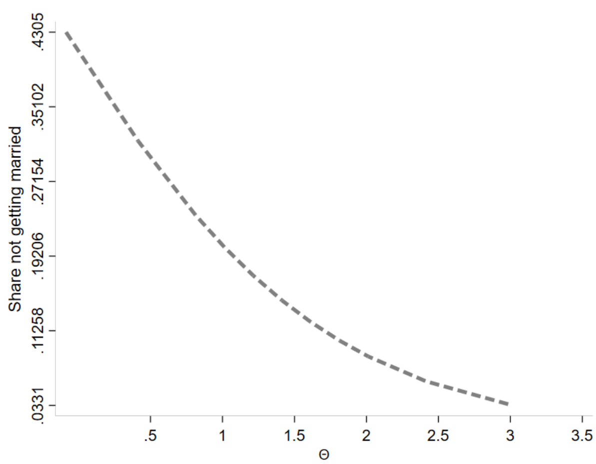

[

.]Share of people not getting married as a function of the exogenous shifter Θ.

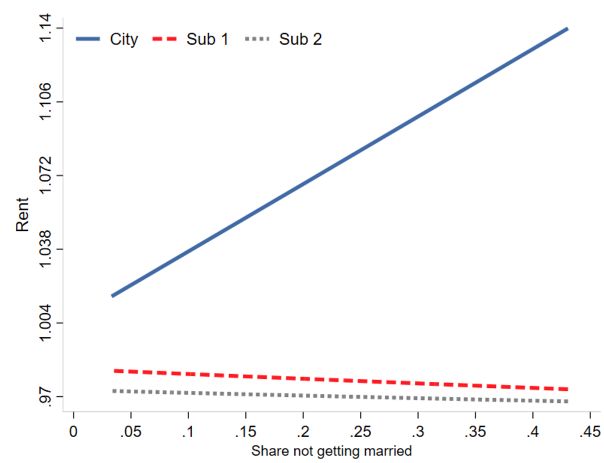

Housing rents per location as a function of the resulting share of people not getting married.

A.4.3 Sprawling into the suburbs

The next counter-factual I study mimics the growth of an existing metro-area by 33% in population and by 33% in housing stock, where this new housing supply is located in a brand new third suburb. Figure A.9 presents this experiment. I add a new neighborhood on the outside of the metro-area that mimics the features of the existing suburban neighborhoods. This new neighborhood is gender-neutral in terms of the jobs offered. Share of first offers coming from the new development is the same as in older suburbs, and so are the amenities offered. Job benefits (monetary and non-monetary) within the whole metro area are scaled to have the same average as the baseline (so there is no change in overall productivity in jobs) and vary based on the distance to other first offers in own sector, according ot the new spatial structure.

Similarly to the main sprawl counter-factual, this form of growth of the metro-area into the suburbs increases the value of marriage. This result is strengthened here, because the metro-area adds a neighborhood that also has couple-targeted amenities (being a suburb). Figure A.10 shows that in the new metro area, there is more marriage while bargaining power adjusts to make men and women enter marriage at the same rates.

Figure A.9: Sprawling into the suburbs – adding a third suburb to grow the metro area (and simultaneously increasing the population by a third).

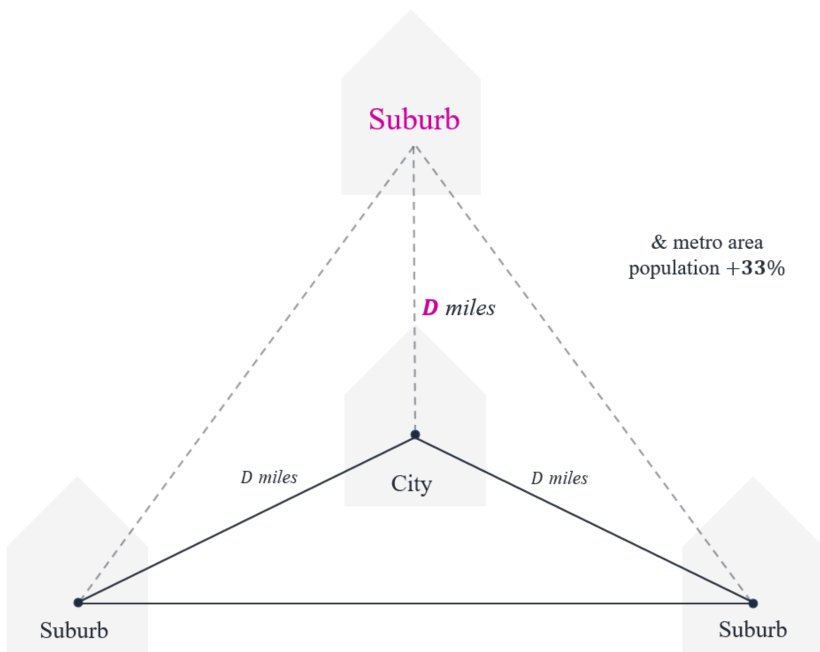

Figure A.10: Sprawling into the suburbs – adding a third suburb to grow the metro area (and simultaneously increasing the population by a third).

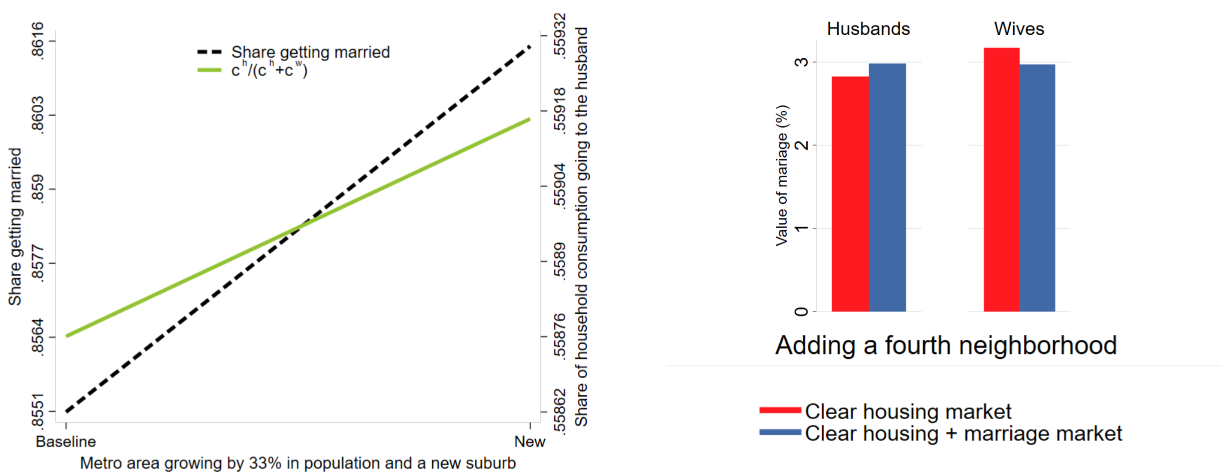

Next I examine how the outcomes of the model change when varying the exact location of the new suburb. Specifically, I change how close or far away the new suburb is from the city. Figure A.11 presents this counter-factual and shows that the further away from city center the new development is located (the more sprawled the new metro area is), the more people get married.

Figure A.11: Sprawling into the suburbs

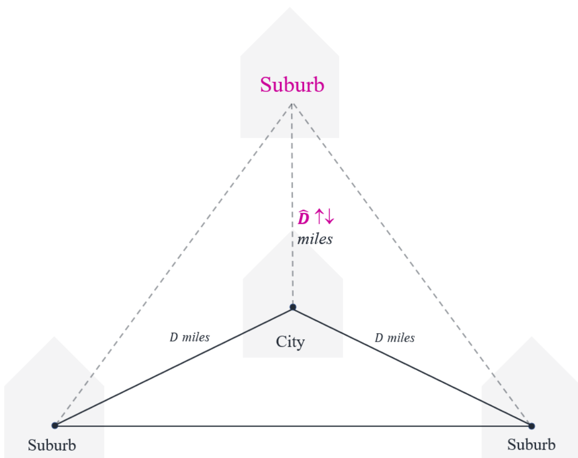

Sprawling into the suburbs – changing how far away from city center the new suburban development is located.

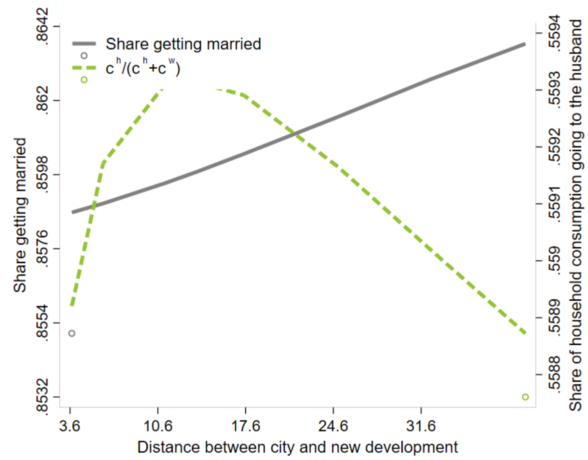

Sprawling into the suburbs – changing how far away from city center the new suburban development is located.

The renegotiation within couples is nonmonotonic with respect to distance. When new development is very close to the city, increasing the distance brings wives more leisure while the job benefits they are losing are small. Therefore, within the marriage market equilibrium, men gain in bargaining weight and are compensated with more consumption. However, at longer distances the gains in leisure for women are not enough to compensate for losing jobs with great benefits. Thus, the bargaining weight starts falling again to result in an overall increase in the value of marriage for women as well as men.

Figure A.12 shows that the same dynamics of specialization play out with four locations, when varying the new suburb's distance (connectivity) to the rest of the metro-area. As in the main analysis, suburban sprawl increases gender gaps within couples in both commuting and employment. This specialization is a key to the mechanism by which marriage becomes more available in sprawled areas.

A.4.4 Sprawling into the suburbs with average amenities

The next counter-factual I study mimics the growth of an existing metro-area by 33% of new housing units and new population and adding a fourth neighborhood with average

Figure A.12: Distance of new suburb: commuting of all subgroups. Gender gaps within couples.

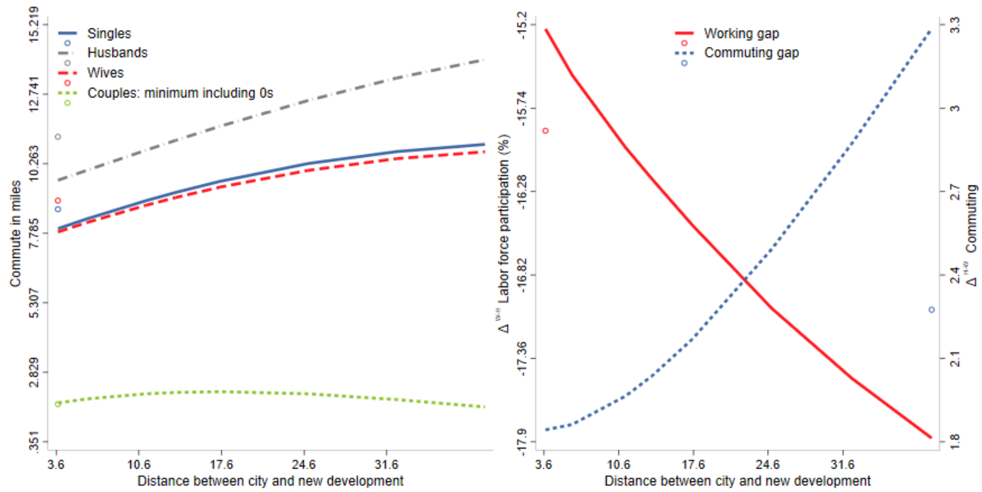

amenities (not specifically catering to couples). Figure A.13 presents this experiment. I add a new neighborhood on the outside of the metro-area.

Figure A.13: Sprawling into the suburbs – adding a new neighborhood with average amenities (and simultaneously increasing the population by a third).

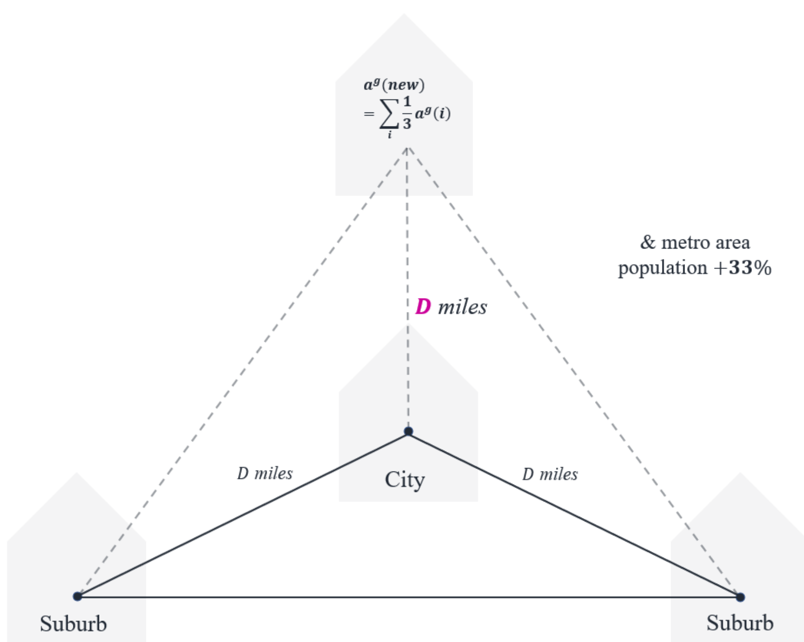

By definition, this new neighborhood is gender-neutral in terms of the jobs offered. Share of first offers coming from the new development is the same as in older suburbs. In terms of amenities and job benefits (monetary and non-monetary), the new neighborhood is an average between the existing neighborhoods, suburbs and city, so the overall distribution of amenities and job qualities is kept constant to the baseline calibration.

Similarly to the main sprawl counter-factual, this form of growth of the metro-area into the suburbs increases the value of marriage. Figure A.14 shows that in the new metro area, there is more marriage while bargaining power adjusts to make men and women enter marriage at the same rates. This conclusion comes about through a combination of a housing and marriage market equilibrium.

Figure A.14: Sprawling into the suburbs – adding a third suburb to grow the metro area (and simultaneously increasing the population by a third).

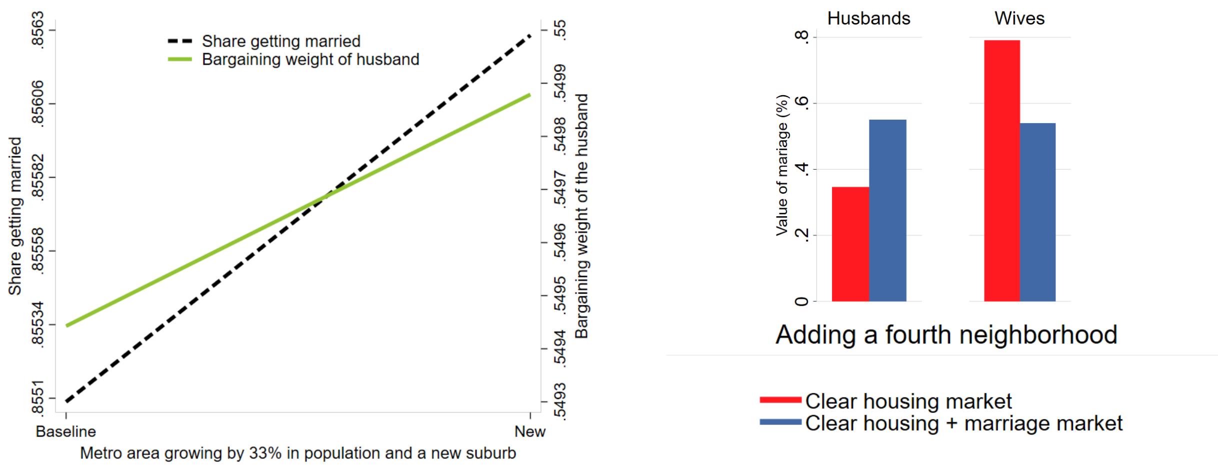

In this version, the value of marriage does increases less than in A.10, because the new neighborhood is suburban in terms of job access, but has amenities that are an average of the existing neighborhoods (not catering specifically to couples). As a result, this version of suburbanization does not actually increase the share of neighborhoods providing suburban amenities that favor couples. Still, the value of marriage increases because couples specialize and pay less for housing, and thus partially avoid the increased commuting costs.

A.5 Testing model predictions with cross metro-area correlations

In this section I show that cross-sectional differences between metro-areas in the US match counter-factual simulations of the model. First I replicate results by Black et al. (2014) showing that metro areas with longer commutes have larger differences in labor force participation between men and women in couples. Using the 2000 IPUMS Census sample I run the following regression

$$Working_{im} = \beta_w C_m \cdot (\text{woman}) + \gamma X_i + \delta X_m + \epsilon_{i,m}$$

where i stands for an individual, m for a metro are and C_m is the average annualized hours of commuting in a metro are m. The sample is restricted to people in couples. β_w is the coefficient of interest – it shows the differential impact of living in a place of long-commutes on men and women. Table A.15, column 6, shows the results. In metro-areas with 16.5 more

Table A.15: The correlation of average commutes and differences in commuting by gender and relationship status across MSAs.

	Commute (annualized)			Working		
$C\cdot$ (in a couple)	0.0667 (0.006)	-0.170 (0.036)		0.0133 (0.0009)	-0.0326 (0.005)	
$C\cdot$ (woman)			-0.239 (0.038)			-0.0552 (0.010)
X_i:						
C	x	x	x	x	x	x
$C\cdot$ *(age, race and educ dummies)*	x	x	x	x	x	x
Sex-couple, age, education, region and race dummies. MSA population.	x	x	x	x	x	x
Industry dummies	x	x	x			
N	1194278	990877	1558750	1776688	1750895	2267949
Sample:	men	women	couples	men	women	couples

SEs statistics in parentheses.

All samples include only people who are married or never married.

Source: IPUMS 2000 Census 5% sample. Sample: 18-50 years old, married or never married, MSAs of at least 250k people. ”Working” is equal to one if the person worked at least for 1 week in the past year and is scaled up by 100 so that results are interpreted as percentage point changes. ”In couple” includes married and cohabitation. Industry dummies are for 1-digit NAICS codes. D_m is the average of annualized commuting hours for all residents of the MSA that do not work from home. Regression in columns 1-2 and 4-5: $d_i = \beta C_m \cdot \text{(in couples)} + \gamma X_i + \epsilon_{i,m}$ for either men or women. Regression in columns 3 and 6: $d_i = \beta C_m \cdot \text{(woman)} + \gamma X_i + \epsilon_{i,m}$ for people in couples.

average hours of commuting per year (roughly corresponding to 1 mile) the gender gap in labor-force participation in couples is higher by almost a whole percentage point.

In column 3, I repeat the same exercise, replacing labor-force-participation with commuting itself on the left hand side (and use a sample of working individuals).

$$d_i = \beta_c C_m \cdot \text{(woman)} + \gamma X_i + \delta X_m + \epsilon_{i,m}$$

The results show that in metro-areas with longer average commutes the difference between the time spent commuting of wives and husbands increases. When the average commute in a metro-area increases by an hour, husbands commute increases by an average of 0.24

hours more than that of wives. Qualitatively, this is exactly what happens in the model. Quantitatively, counter-factual exercises above imply a somewhat bigger effect – an increase of 0.58 hours.

Next I repeat the above analysis, this time focusing on the difference between couples and singles using the following regressions.

$$Working_{im} = \beta_w C_m \cdot \text{(in couples)} + \gamma X_i + \delta X_m + \epsilon_{i,m}$$

$$d_i = \beta_c C_m \cdot \text{(in couples)} + \gamma X_i + \delta X_m + \epsilon_{i,m}$$

If β_w (β_c) is positive, longer average commutes are associated with higher labor force participation (longer commutes) in couples compared to singles. I run this analysis separately for men and women. Columns 1-2 and 4-5 of table A.15 show that both β_w and β_c are estimated to be positive for men and negative for women. In metro-areas with long commutes men in couples work and (conditional on working) commute more than single men. However, women in couples work less and (conditional on working) commute less than single women. Qualitatively, this is exactly what happens in the model. Quantitatively, the model predicts a larger effect on commuting of men while the data suggests a larger effect on commuting of women.

Next I investigate the model prediction that larger average commutes are actually conducive of couple formation, by making single life disproportionately costly compared to being in a couple and being able to specialize. I focus on the subpopulation of 30-50 years of age, responding to the population in the model that is either in a couple or perpetually single. Using the 2000 IPUMS Census samples I run the following regression

$$\text{Ever in couple}_{im} = \gamma C_m + \gamma X_i + \delta X_m + \epsilon_{i,m}, \ \ \forall i : \text{age}_i \geq 30$$

Ever in couple_{im} is a dummy variable equal to 1 if the person has ever been married or is currently cohabiting with a partner. C_m, again, is the average annualized hours of commuting in a metro are m. Table A.16 shows that, at least when metro-area-level controls X_m include religious participation and proxies for political affiliation, the estimate of $\gamma > 0$. Therefore, across metro areas those with a longer average commute tend to have fewer people staying perpetually single. This correlation in the data could be caused by a selection effect – metro-areas with more couples have higher average commutes because it is the married men who commute most. Columns 3 and 4 in table A.16 shows the result is robust to replacing C_m with the average commute among only married men, avoiding this type of selection. It is important to know these results present descriptive and suggestive evidence, not the causal

Table A.16: The correlation of average commutes and the probability of staying perpetually single across MSAs.

	(Ever married or cohabiting)·100			
C	0.00465 (.0097)	0.0158 (.0048)		
$C_{husbands}$			0.0114 (.0067)	0.0134 (.0031)
X_i:				
Age, sex-couple, education, region, race dummies. MSA population polynomial.	x	x	x	x
Presidential elestion results 1996-2008, number of religious congregations and adherens by denomination in 2000 (or 2010 if not available earlier).		x		x
N	2754757	2751511	2754757	2751511
Sample:	$30 \leq$ age ≤ 50			

SEs statistics in parentheses.

Source: IPUMS 2000 Census 5% sample. Sample: 30-50 years old, MSAs of at least 250k people. The outcome variable is equal to one if the person is married, divorced, separated, widowed or currently cohabiting. Columns 3 and 4 replace C_m with an average commute in an MSA among married men.

effect of commuting on marriage rates.

Figures A.15 visualizes the variation in average commuting across metro areas, with and without residualizing with respect to proxies for religious participation and political affiliation.

Figure A.15: Average annualized commuting time, residualized.

Second figure residualizes also with respect to religion and politics proxies. As political affiliation proxies I use county-level shares of votes in presidential elections in 1996-2008 going to the democratic candidate (accessed from Leip (2021)). As religious proxies I use the number of congregations per capita and number of adherents per capita in 2000 (overall and specifically for Evangelical Christian denominations), and number of congregations per capita and number of adherents per capita in 2010 in Black Protestant denominations as provided in Jones et al. (2000) and Grammich et al. (2012).

APPENDIX B

Appendix to Chapter 2

B.1 Appendix: Supplementary graphs and tables

In this section I collect supplementary empirical evidence related to chapter 2.

B.1.1 Age correction

When computing lifetime marriage-related descriptors of cohorts, I face a sampling challenge. The data do not allow me to observe all relevant cohorts at the same age to compare on the aggregate which cohorts ultimately divorced the most. To get around that I make use of the fact that I do observe some of the cohorts multiple times and fit a quadratic age gradient in the probability of being ever-divorced. Specifically, I pool IPUMS 1960-1980 Census and ACS $\geq$ 2008 samples, men of $age \in [40, 80)$, and run the following regression:

$$y_{ict} = \alpha_c + \beta age_{ct} + \beta_2 age_{ct}^2 + \epsilon_{ict}$$

Table B.1 presents the results.

Table B.1: Age gradient

Variable	β	β_2
Ever-divorced (among men)	.0184299	-.0001324

Estimating an age gradient, using IPUMS 1960-1980 Census and $\geq$ 2008 ACS.

Using β and β_2 I adjust the share of men ever divorced ($s_{c,age=a}$) to age 50.

$$s_{c,50} = s_{c,age=a} + \beta(50 - a) + \beta_2(50^2 - a^2)$$

The identifying variation comes from cohorts (1890-1976) who are observed at least twice. Cohorts 1929-1939 are observed both in 1980 and in later ACS. Figure B.1 is equivalent to figure 2.2a, but without the age correction.

Figure B.1: Share ever divorced

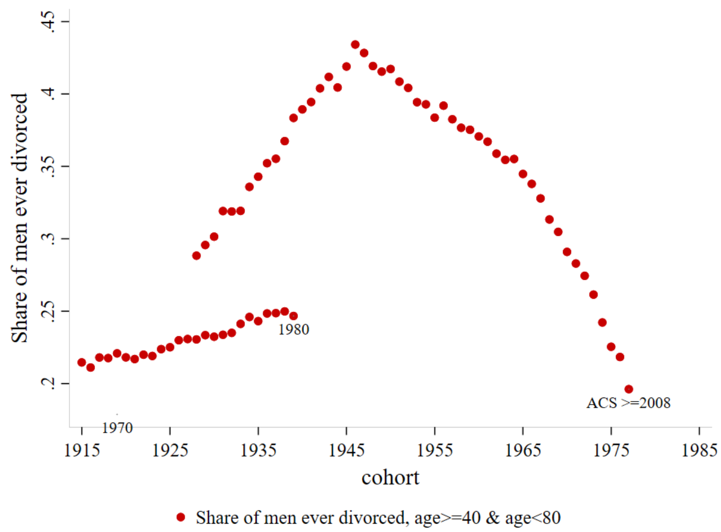

2.2a without an age correction.

B.1.2 What is a common age to first divorce?

People commonly divorce in their 30s. Figure B.2 shows evidence of this from the NSFG. Similarly, Goldin and Katz (2016) show a mean age at divorce for women around 34, and decreasing to 30 with later cohorts, for women observed 50-74 in the SIPP.

Figure B.2: Age at divorce

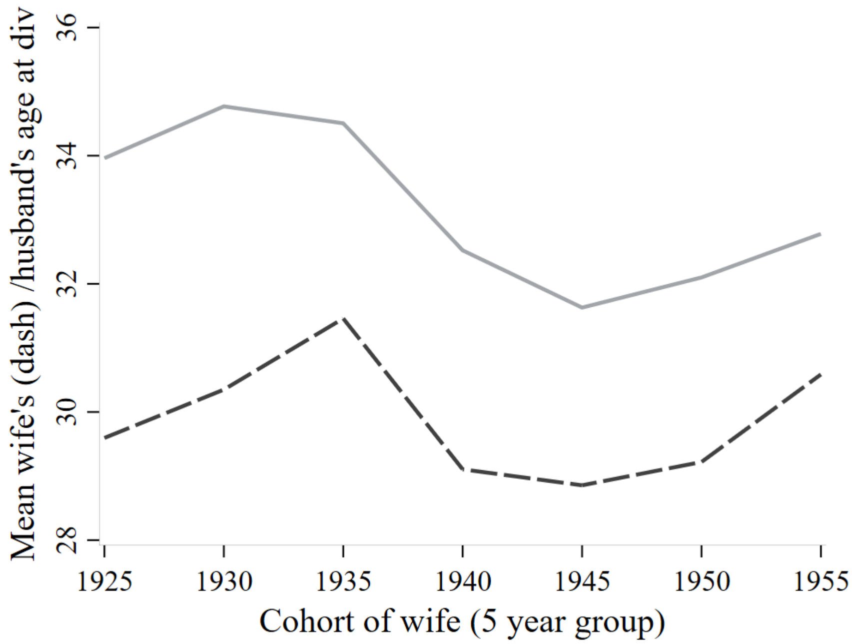

Source: NSFG, women 40-45, ever divorced, age at wife's first divorce.

B.1.3 Supplementary cohort evidence

Especially compared to the downward trend in the age-gap (husband minus wife), men in the pre-boom (and also in the 1930-35 cohorts) lived in marriages with bigger age-gaps on average (see figure B.3). This is consistent with the idea that men in these generations (re)matched with younger (second) wifes.[1]

Among the 'treated' cohorts, men almost caught up with women in whether they ever got married by midage, especially compared to a more downward trend in marriage among men compared to women (see figure B.4). This is consistent with the assertion that this cohort of men faced an exceptionally large pool of eligible women and a low competition from older men.

Compared to other cohorts, women in the treated cohorts were especially likely to stay divorced (not remarry) compared to men (see figure B.5).[2] This is consistent with the hypothesis that the primary cause for divorce was an increase in the remarriage options of the husband.

[1]Analysis by the type of marriage (first versus higher order) shows that this is a composition effect of more remarriage, rather than even first marriages happening more often with larger age-gaps. Since the age-gap in husband's second marriage tends to be higher than in their first marriage, an increase in the share of higher order marriages implies an increase in the age-gap.

[2]Note - this is actually true for the early-boomers as well.

Figure B.3: Age-gap (husband minus wife) for men by age for selected cohorts/ by cohort for selected ages.

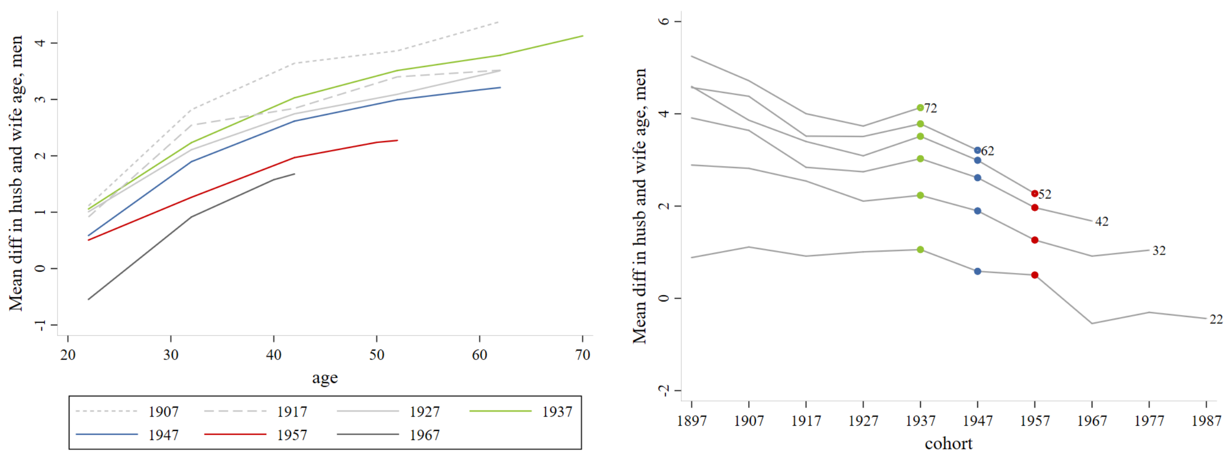

Figure B.4: Gender difference in the share of ever-married.

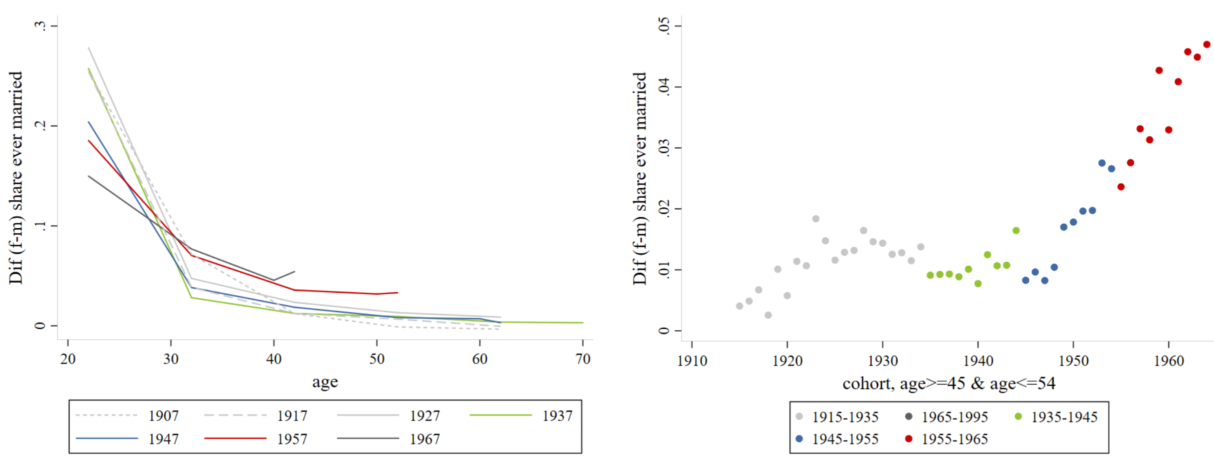

Figure B.5: Gender difference in the share of ever-married/ ever-divorced who live as divorced by cohort for selected ages.

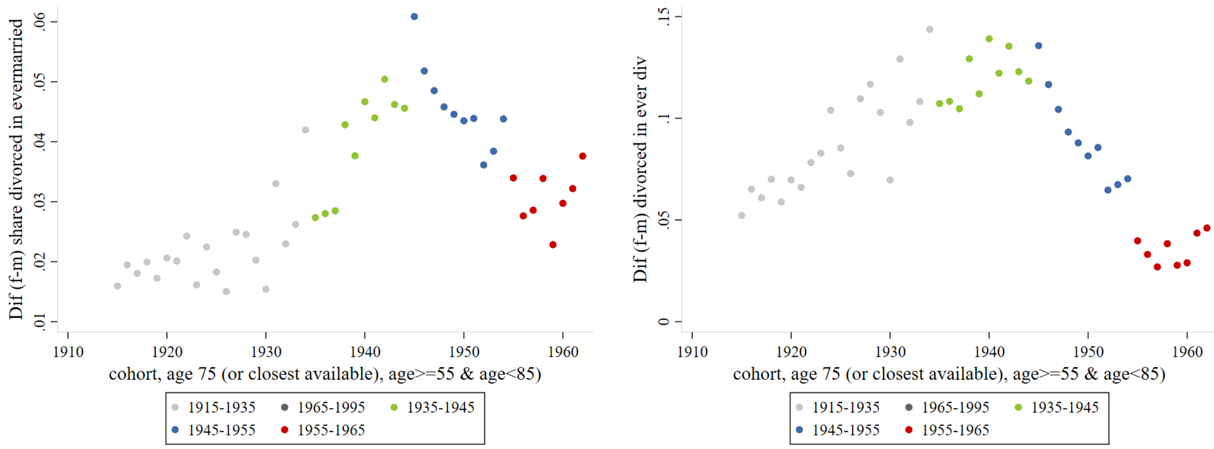

Especially among the 'treated' cohorts, women stayed divorced (single) more than men (even though both divorced a lot).

B.1.4 Supplementary analysis of variation in mobilization rates

Mobilization rate is defined as the fraction of registered men between the ages of 18 and 44 who were drafted or enlisted for WWII, as generously provided by Acemoglu et al. (2004).[3] Figure B.6 shows the geographic distribution of mobilization rate. Table B.2 shows descriptive characteristics of men and women aged 35-44 in high and low mobilization rate states.

Figure B.6: Mobilization rate in WWII (Acemoglu et al., 2004).

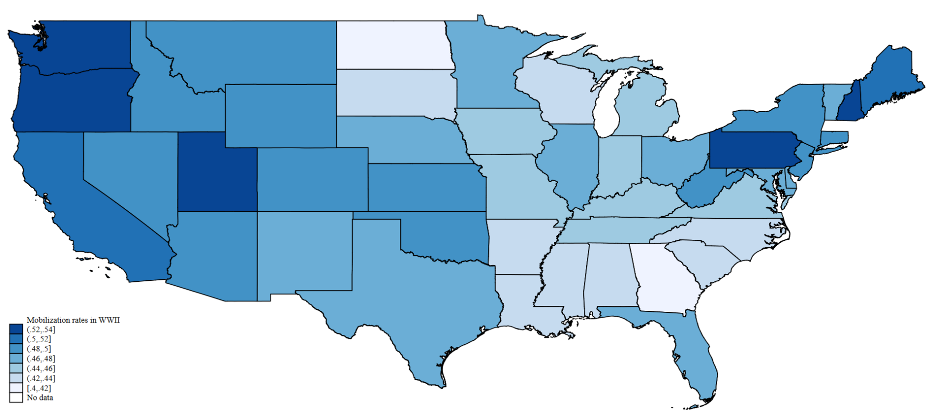

There are baseline differences (even before the WWII) between these states in race,

[3] Original source: published tables of the Selective Service System (1956). Since essentially all men in the relevant age range were registered, mobilization rate is effectively the fraction of men in this age range who have served (Acemoglu et al., 2004).

Table B.2: Descriptive statistics of high and low mobilization states respectively, men and women 35-44 years old.

Cohorts	1940 (1895-1904)		1960 (1915-1924)		1980 (1935-1944)	
	$> p50$	$\leq p50$	$> p50$	$\leq p50$	$> p50$	$\leq p50$
White	0.97	0.83	0.96	0.85	0.94	0.83
Graduated high-shool	0.31	0.24	0.59	0.47	0.83	0.75
Graduated college	0.07	0.05	0.10	0.08	0.22	0.17
Farm	0.11	0.28	0.05	0.12	0.02	0.04
Urban	0.65	0.50	.	.	.	.
In Metro area	0.69	0.44	0.68	0.54	0.74	0.64
Ever-married	0.88	0.90	0.93	0.94	0.93	0.94
Ever-divorced	.	.	0.14	0.17	0.26	0.27

education and rate of urbanization, as already discussed in Acemoglu et al. (2004). At least race and farm are explicitly linked to the mobilization rules (as whites were preferred for mobilization, while mobilization of farmers was discouraged to sustain food production).

Figure B.7 shows differences between high and low mobilization states in share of men ever divorced by age 40 by cohort. Importantly, there are baseline differences between these states: high mobilization states start from a lower level. The figure also shows that starting from a lower level high mobilization states had a bigger boom in divorces among 1925-1955 cohorts, overtaking low mobilization states (consistent with the hypothesis in this paper). This graph looks very similar when splitting by high and low remarriage options for 35 year old men in 1980, $\Delta_{1960-1980}\frac{n_{30-34}}{n_{30-44}}$, or when plotting share ever divorced in ever-married.

The baseline difference in divorce rates strongly motivates investigating what characteristics might account for it and then control for these differences among states in the analysis. Luckily, the differences in the baseline can be largely accounted for by differences in race and education. Using the IPUMS 1960 Census data and ages 35-80 (cohorts 1890-1924, i.e. cohorts that should not be affected by the baby-boom yet), I fit the following regression:

$$\begin{aligned} y_{ist} = \alpha_s + \alpha_{1980} + \alpha_{age} + 0.11\ \textit{Black}\ - 0.01\ \textit{Hispanic} \\ - 0.01\ \textit{Graduated Highschool}\ - 0.04\ \textit{Graduated college}\ + \epsilon_{ist} \end{aligned}$$

I then use the coefficients on race and education to residualize the (age adjusted) share of men ever-divorced in all cohorts 1890-1980. Figure B.8 shows that accounting for these basic demographic differences diminishes the baseline difference in divorce rates to 0, making the states with high and low mobilization rates much more comparable. Reassuringly, controlling

Figure B.7: Share of ever-divorced, adjusted to age 40/50, top and bottom quartiles in terms of WWII mobilization rates. Mobilization rates assigned by state of birth.

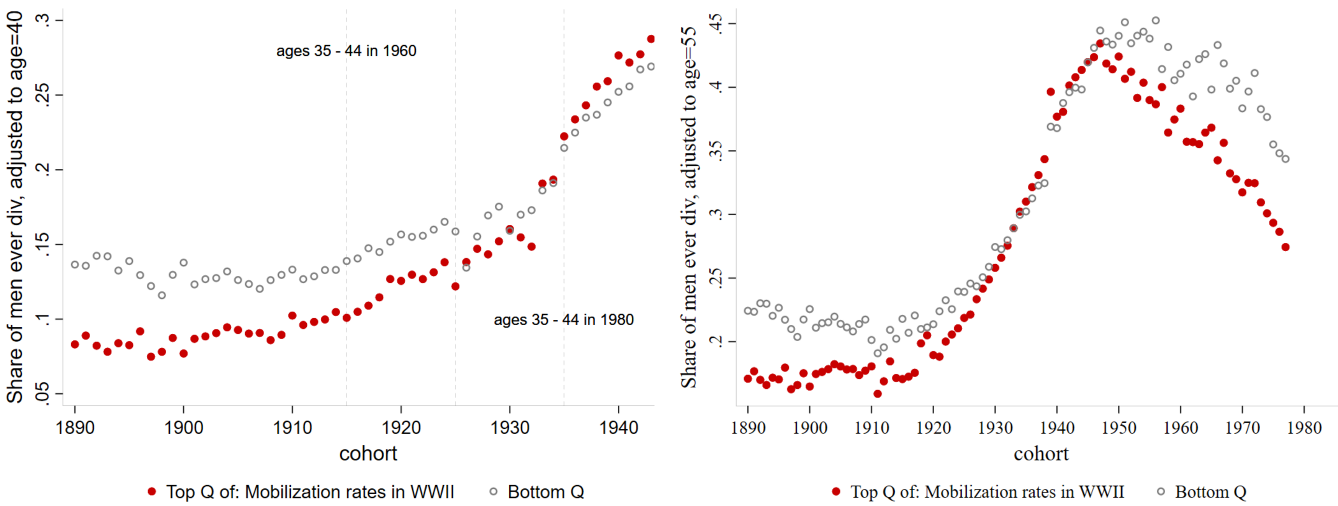

for these baseline demographic differences does not at all undo the conclusion that 40-year-old men born in states with a higher WW II mobilization rate had a bigger increase in the chance of getting divorced by 1980.

Figure B.8: Share of ever-divorced, adjusted to age 40/50 and for baseline differences in race and education, top and bottom quartiles in terms of WWII mobilization rates. Mobilization rates assigned by state of birth.

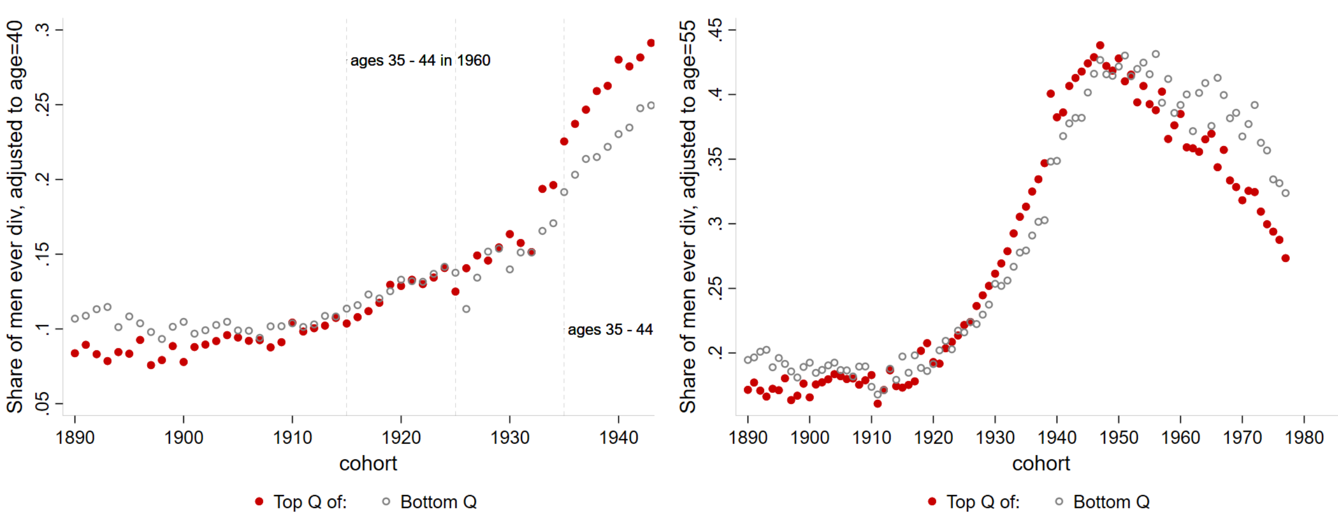

Motivated by the baseline differences presented in table B.2, I control for dummies for race, education and farm status (as detailed as available in IPUMS) in all of the regressions in the main analysis.

B.2 Appendix: Does the effect of mobilization rates come from women's labor force participation or fertility?

The variation in mobilization rates has been shown to increase labor force participation of women during (and a bit after WWII) and decrease labor force participation and increase fertility of young women after the war by Acemoglu et al. (2004) and Doepke et al. (2013) respectively. The story goes as follows. As men were called-of to war, women (mainly in cohorts 1905-1914, but more broadly in 1905-1924) worked more during the war and beyond, gaining experience and quickly gaining a strong position in the labor market. This labor supply 'shock', however, depressed wages of women in general (with women being an imperfect substitute for men in the labor market), especially for potential inexperienced incomers from later cohorts. This incentivized the cohorts of 1925-1934 (who were 15-24 years old by 1950) to stay at home and increase fertility (realized between 1945 and 1960), instead of competing in the overcrowded labor market.

This previous research motivates the use of mobilization rates as an instrument for the size of the baby-boom in this paper. This is mainly because the mechanism of the paper affected cohorts that were generally too old to be relevant to contribute meaningfully to the divorce boom in the 1970s and 1980s (supporting an exclusion restriction assumption necessary for an instrument). Still, a natural question is whether these previously documented mechanisms could not somehow affect divorces directly, not just through shifting the future age distribution in the marriage market.

Namely, one might suspect that an increase in labor force participation of women might make couples more likely to divorce, as women gain work experience and their value outside of marriage increases (as suggested by Ruggles (1997), Weiss and Willis (1997)), thus implying a correlation between the divorce boom and mobilization rates in WWII. This alternative hypothesis is however not consistent with the incidence of WWII effects on labor force participation vs the divorce increase across cohorts. Figure B.9 visualizes that it was mainly cohorts of women born before 1925 who were shown to have higher work attachment in high mobilization states by Acemoglu et al. (2004). These cohorts are not at all the ones responsible for the divorce boom. Moreover, in the baseline cross-sectional analysis these cohorts in fact serve as a control group to the treated cohorts of 1935-1944 (the two groups are indicated by dashed lines in the figure). As a result, if higher labor force participation by women causes divorce, the baseline results presented in table 2.2 should represent a lower bound (with the control being too high). Though some research suggests persistence in women's attachment to labor markets across generations, this is not what happened in the aggregate. In fact, later cohorts of women were less attached to the labor market, which is precisely what is the hypothesized reason for them to instead have more children.

Figure B.9: Effects of mobilization rates across cohorts on female labor force participation and fertility, among cohorts used as treated and control in the baseline analysis.

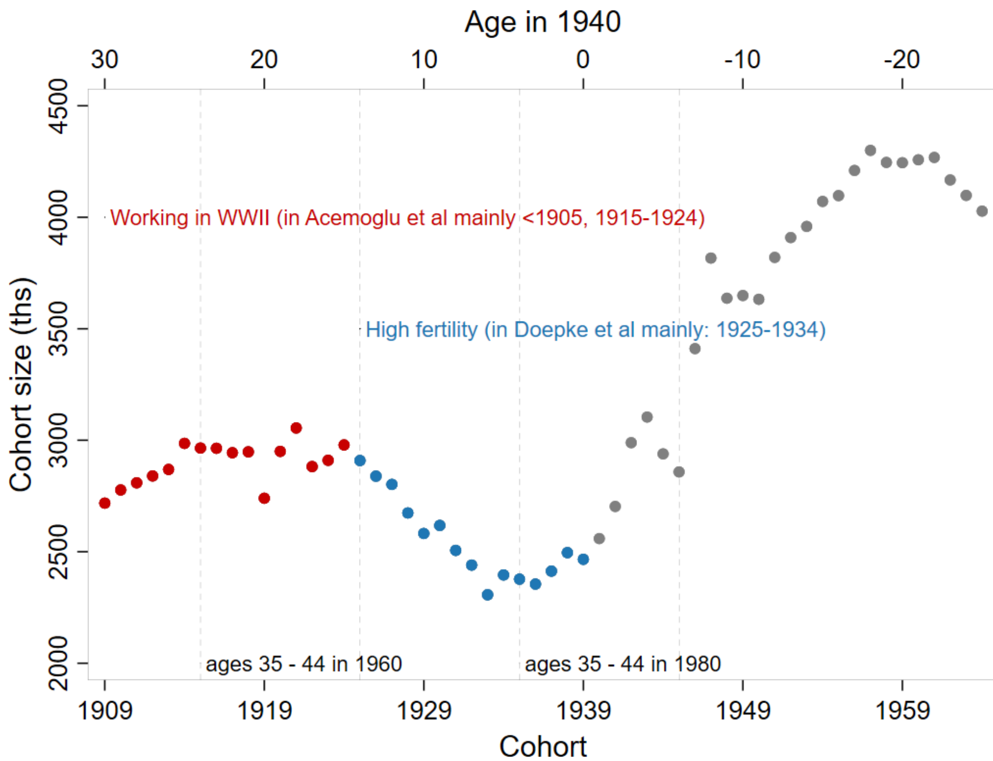

As a further check, I also rerun the analysis presented in table 2.2 for women only and include dummies for labor force participation and (or) employment and confirm that the results are essentially unaffected. Thus it seems that labor force participation of women did not play a major role in the divorce boom, whether it had anything to do with WW II mobilization rates decades earlier or not.

Figure B.9 also shows that the cohorts having a higher fertility compared to previous cohorts in high mobilization states overlap partially with the treated cohorts who started the divorce boom. In other words, the cohorts of 1935-1939 were both mothers to the very late baby-boomers and were themselves affected by the early baby-boomers entering the marriage market. A competing hypothesis could therefore be that couples in high mobilization states started divorcing more because women in these couples had more children and less work attachment, making them more vulnerable to divorce. To test this mechanism, I repeat the baseline analysis, but compare the treated cohorts of 1935-1944 to the cohorts 1925-1934, whose fertility was especially high. Specifically, I repeat the analysis outlined in section 2.3

with a pooled sample of 1970 and 1980 IPUMS Census data, i.e. $t \in \{1970, 1980\}$ in

$$y_{ist} = \alpha_s + \alpha_{1980} + \alpha_{age} + \beta \frac{n_{30-34}}{n_{30-44}}_{st} + \gamma X_{ist} + \epsilon_{ist}$$

If the effect of mobilization rates on divorce is coming from a greater incidence of vulnerable stay-at-home women, comparing two cohorts of high fertility should lead null results. Table B.3 shows that the results are in fact qualitatively the same as in table 2.2. Quantitatively, these results are smaller. This is likely caused by the fact that the cohorts of 1925-1934 and 1935-1944 are close enough so that the entry of baby-boomers affected both of them to a certain degree, yet the effect was bigger on cohorts 1935-1944.

Lastly, I confirm that including age at marriage/ number of children ever born/ number of young children in the household dummies as controls does not diminish the baseline results.

Table B.3: Results with an alternative comparison group

	Ever-divorced					
	OLS		2SLS			
$\frac{n_{30-34}}{n_{30-44}}_{st}$	0.466	0.477	0.884	0.951	1.133	0.916
	(0.156)	(0.163)	(0.231)	(0.290)	(0.360)	(0.297)
X_i:						
farm, metro	no	yes	no	yes	yes	yes
educ dummies	no	yes	no	yes	yes	yes
region-year fes	no	no	no	no	yes	no
In ever-married	no	no	no	no	no	yes
F-stat in 1st stage			56.6	56.6	40.3	56.1

N = 1222220 (1143381 in column 5), 48 clusters

SEs in parentheses. SEs clustered at the state level.

All regressions include year and state fixed effects, and dummies for age, sex, race.

Same setup as table 2.2 except using pooled 1970 and 1980 IPUMS Census data, i.e. measuring growth in divorces between cohorts 1925-1934 and 1935-1944.

B.2.1 Supplementary tables for cross-sectional analysis

First stage This section presents evidence on the strength of the first stage in my IV strategy. Following Doepke et al. (2013), I show that states with higher mobilization rates during WWII have a sharper cohort size growth after WWII, materializing in a sharper increase in the remarriage opportunities in the 1970s for men in their 30s. Specifically, I regress the remarriage opportunity measure, $\frac{n_{30-34}}{n_{30-44}}_{st}$ or $\frac{n_{0-4}}{n_{0-14}}_{s,t-30}$ (notice t still stands for a year 1960 or 1980) on the WWII mobilization rates interacted with a dummy for year 1980, including state and year fixed effects as well as individual level controls used in the main analysis presented in table 2.2.[4] This specification is a panel fixed-effects version of a regression predicting the change in the remarriage opportunity measure with WWII mobilization rates. Table B.4 shows that this strategy has a strong first stage (with F statistics on mobilization rates above 15), with perhaps the exception of the specification regressing $\frac{n_{0-4}}{n_{0-14}}_{s,t-30}$ and including region times year fixed effects, where the F statistic falls slightly short of the rule of thumb of 10 (as suggested by Staiger and Stock (1997) to avoid weak instruments issues as brought to light by Bound et al. (1995)).

Table B.4: First stage results, as relevant for the 2SLS results in table 2.2 and B.5.

$$u_{st} = \alpha_s + \alpha_{1980} + \alpha_{age} + \delta mob \cdot 1_{t=1980} + \gamma X_{ist} + \epsilon_{ist}$$

u_{st}	$\frac{n_{30-34}}{n_{30-44}}_{st}$			$\frac{n_{0-4}}{n_{0-14}}_{s,t-30}$		
δ	0.214	0.214	0.202	0.287	0.283	0.211
F statistic	25.60	25.70	31.13	15.36	15.05	8.76
X_i:						
sex, race	yes	yes	yes	yes	yes	yes
farm, metro	no	yes	yes	no	yes	yes
educ dummies	no	yes	yes	no	yes	yes
region-year fes	no	no	yes	no	no	yes

$N = 2032220$, 48 clusters

SEs clustered at the state level.

All regressions include year, age and state fixed effects.

Proxy cohort size by 1980 with cohort size among children by 1950 To confirm that the baseline result is not driven by selective migration in adulthood I reconstruct the

[4] These controls take out variation in the remarriage opportunity measure correlated with state level compositional changes in sociodemographic variables

measure $\frac{n_{30-34}}{n_{30-44}}_{st}$ in 1950 and 1930 Census as $\frac{n_{0-4}}{n_{0-14}}_{s,t-30}$ and run the following equivalent regression:

$$y_{ist} = \alpha_s + \alpha_{1980} + \beta \frac{n_{0-4}}{n_{0-14}}_{s,t-30} + \gamma X_{ist} + \epsilon_{ist}$$

Table B.5 shows the results.

Table B.5: Baseline with baby-boom size measured at ages 0-14.

	Ever-divorced					
	OLS		2SLS			
$\frac{n_{0-4}}{n_{0-14}}_{s,t-30}$	0.666	0.832	1.109	1.564	1.836	1.085
	(0.112)	(0.130)	(0.242)	(0.301)	(0.561)	(0.242)
X_i:						
sex, race	yes	yes	yes	yes	yes	yes
farm, metro	no	yes	no	yes	yes	yes
educ dummies	no	yes	no	yes	yes	yes
region-year fes	no	no	no	no	yes	no
In ever-married	no	no	no	no	no	yes

N = 2032220 (1902899 in column 5), 48 clusters

SEs statistics in parentheses. SEs clustered at the state level.

All regressions include year, age and state fixed effects.

Same as table 2.2 except $\frac{n_{30-34}}{n_{30-44}}_{st}$ is replaced with $\frac{n_{0-4}}{n_{0-14}}_{t-30}$ (constructed from 1930 and 1950 IPUMS Census data).

The results are broadly similar to table 2.2, though slightly smaller in magnitude.

B.3 Appendix: What is the role of age at marriage?

Between 1960 and 1980 the Census asked both about age at first marriage and the order of current marriage, making it possible to analyze divorce probabilities by age at marriage at selected ages for some of the relevant cohorts. We know from Rotz (2016) (and others) that low age at marriage is strongly predictive of divorce. This is also visible in figure B.10, showing that the share of men who ever divorced is markedly higher, for various cohorts and ages, among those who got married by the age of 20.

Figure B.10: Share of men who ever divorced, among those who ever-married/ among those who married by age 20.

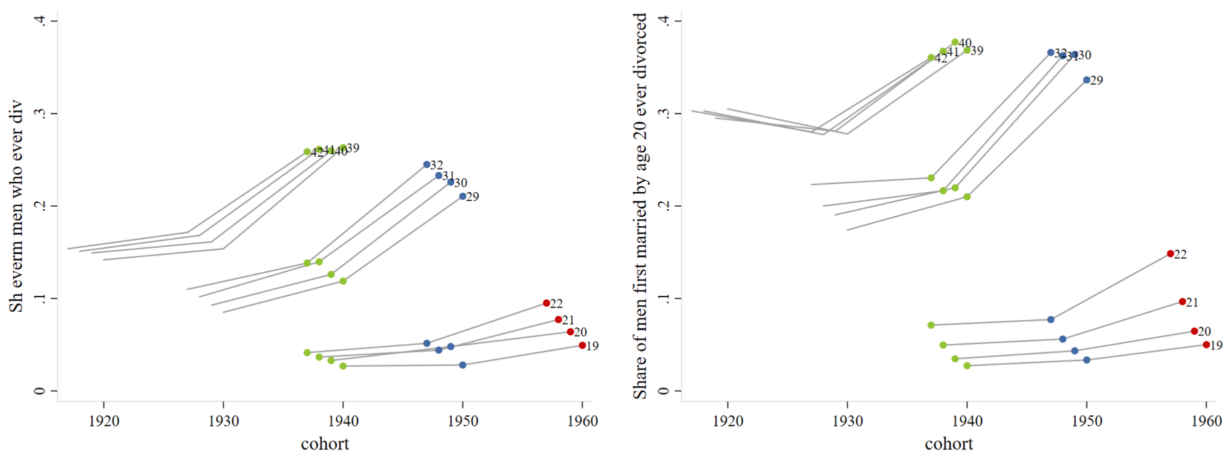

Figure B.11: Share of men and women ever-married in selected by age plotted for selected available ages.

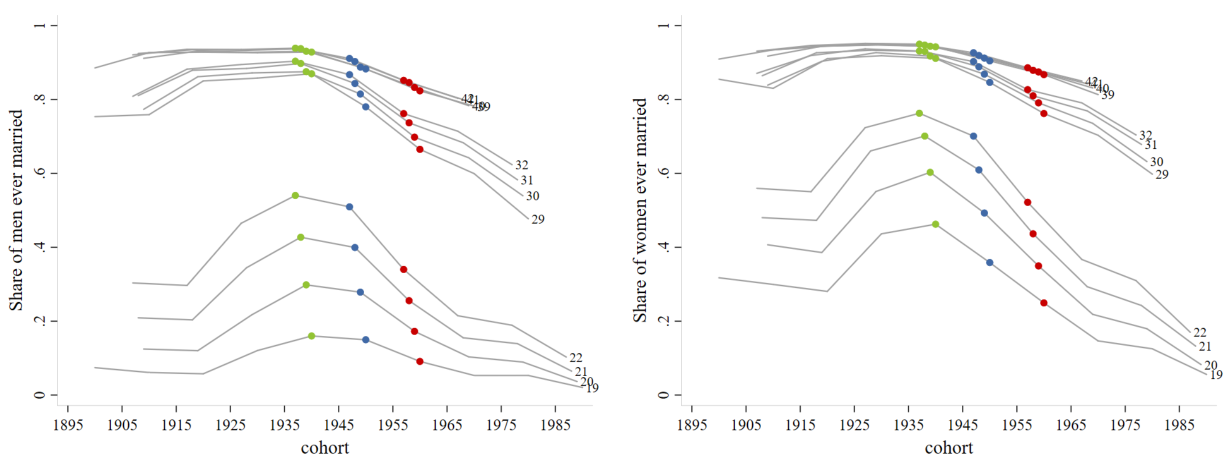

Data source: all IPUMS Census and ACS.

Figure B.12: Share of men/ women ever married by age 19 (using the reported age at marriage) by cohort.

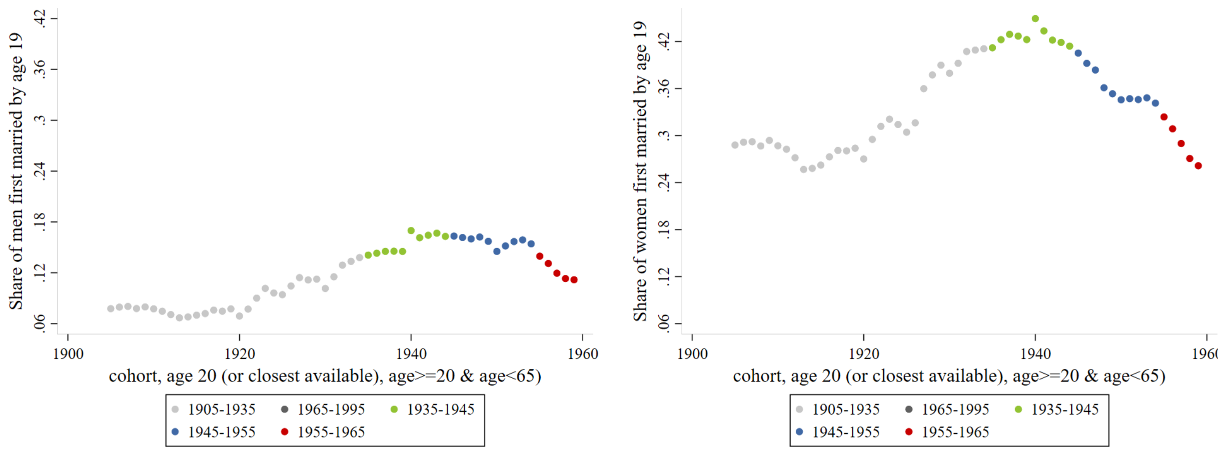

Data source: IPUMS Census, 1960-1980.

At the same time, figures B.11 and B.12 document that women (and men to a smaller extent) married especially young among cohorts 1935-1945 (the ones responsible for the run up in the divorce boom), while men married especially young also in cohorts 1945-1955 (the early baby-boomers) whose divorce rates were still quite elevated. This could suggest that the divorce boom might also have been a result of low age at marriage. However, including age at marriage as a control in my cross-sectional regressions does not diminish the results that pre-boom generations divorced more when faced with a large baby-boom cohort of entrants (see table B.6).

Still, I suspect that low age at marriage among the pre-baby-boom (and somewhat extending to the early baby-boom) generations was endogenous to the cohort size dynamics. The pre-baby-boom cohorts (especially cohorts born during 1930s) were very small (there was a baby-bust during the great depression). This creates a shortage of young women in the marriage market, possibly pushing men not to postpone proposing to lock in a match. If this mechanism is persistent (as young women are matching at earlier ages, their peers also perceive a shortage of partners), it could also explain the low age at marriage extending to the cohorts born after 1935. However, including this mechanism is beyond the scope of this paper.

Table B.6: Baseline, controlling for age at marriage.

	Ever-divorced				
	OLS	2SLS			
$\frac{n_{30-34}}{n_{30-44}}_{st}$	0.569	0.664	1.145	1.647	1.366
	(0.197)	(0.252)	(0.332)	(0.376)	(0.418)
X_i:					
sex, race	yes	yes	yes	yes	yes
farm, metro	no	yes	no	yes	yes
educ dummies	no	yes	no	yes	yes
region-year fes	no	no	no	no	yes

N = 2032220 (1902899 in column 5), 48 clusters

SEs statistics in parentheses. SEs clustered at the state level.

All regressions include dummies for age at marriage, bottom-coded at 15

All regressions include year, age and state fixed effects.

Same setup as table 2.2 except including age at marriage dummies (bottom coded at 15).

B.4 Appendix: Why did early baby-boom cohorts divorce even more than the 1935-1944 cohorts?

In figure 2.2a it is clear that while the treated cohorts of 1935-1945 were the ones form whom divorce rates increased sharply, the cohorts of 1945-1955 (the early baby-boomers) had on average even higher rates of experiencing divorce in their lifetime, even though the remarriage options for men in these cohorts was already waning somewhat. My hypothesis is that a combination of two factors, marrying to divorcees and establishing more marriages with a small age gap, caused the 1945-1954 cohort to divorce even more. Both of these are endogenous responses to the cohort size movement.

As the 1945-1954 cohorts are much bigger than preceding cohorts, women naturally face a shortage of marriage partners (with a standard age-gap in marriage). Bergstrom and Lam (1991) and others show that a way the marriage market adjusts is by increasing the number of partnerships among peers. These kinds of marriages, however, are in general more at risk of divorce due to rematching. This is because a small age-gap implies a bigger pool of women that are younger than the existing wife. In this paper, I also argue that a second adjustment comes in the form of more matches with divorcees. In general, higher order marriages have also been shown to be less stable on average. Together these imply that large baby-boom cohorts entered more in inherently less stable marriages.

Figures B.13 and B.14 confirm that especially among those that married young, the age-gaps in their initial marriages among the early baby-boomers were markedly smaller especially among women (marrying peers more often). Thus different baby-boom women clearly employed different strategies to the apparent shortage of usual partners. Most of them married peers or younger man (entering into marriages with smaller age-gaps) while some of them married divorcees (increasing later-life age-gaps among the treated cohorts (1935-1945) of men).

Figure B.13: Age-gap in existing marriages by age.

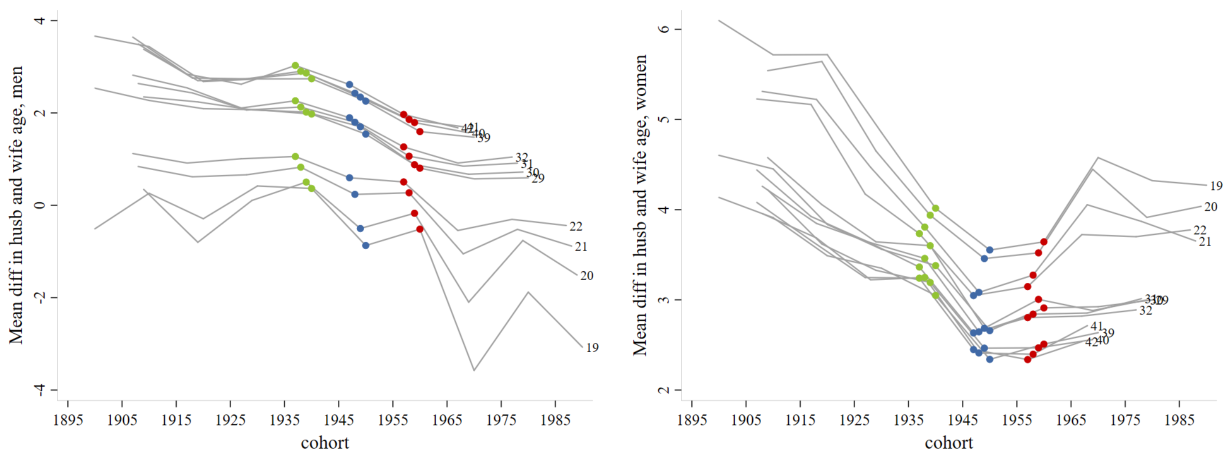

Figure B.14: Age-gap in existing marriages by age, for a group of men/ women who were first married by age 20.

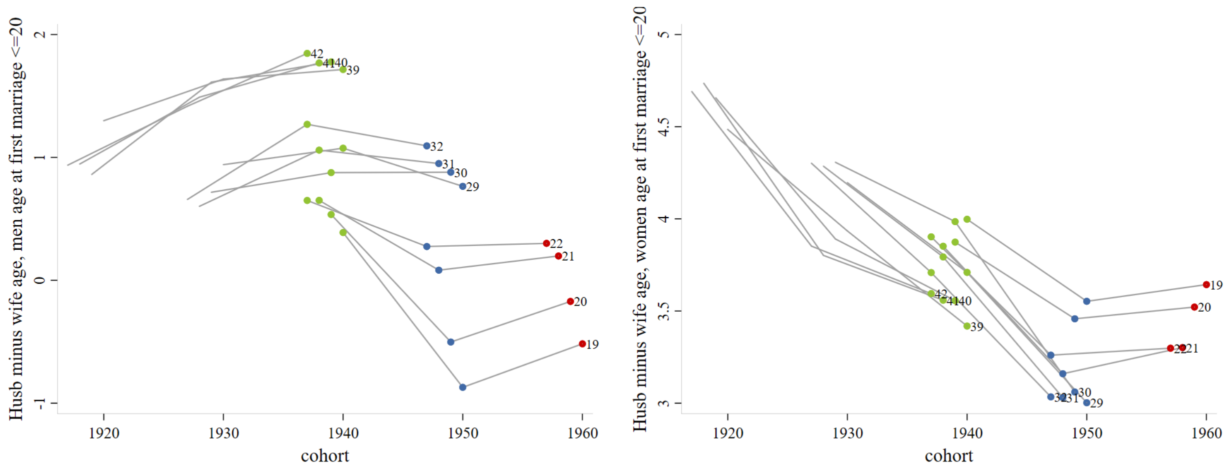

This is a consistently defined group across sample. The sharp increase in the age gap can only happen for two reasons. First, the marriages with small age-gaps broke. Second, some of the marriages at young age broke and rematched with a much younger partner. Both are consistent with the mechanisms studied in this paper. This shows that the increase in age-gap with age is not caused only by selection into first marriage.

Table B.7 and figure B.15 confirms that controlling for age at first marriage[5], an age-gap that is lower (or much higher) than standard is correlated with a higher chance of divorce (using marital histories of cohorts 1920-1970, from the public use versions of the Growth of American Families, National Fertility Surveys and National Surveys of Family Growth as compiled by the Integrated Fertility Survey Serie, (Smock et al., 2015).

Table B.7: Ever-divorced based on age at first marriage

		Ever-divorced
		Age at 1st marriage < 20
Age at first marriage	-0.0327***	-0.0623***
	(0.0013)	(0.0060)
[1em] Age-gap in 1st mar >= 10	0.165***	0.231***
	(0.0189)	(0.0309)
[1em] Age-gap in 1st mar $[5, 10)$	0.0354***	-0.00225
	(0.0132)	(0.0205)
[1em] Age-gap in 1st mar $[1, 5)$	0	0
Age-gap in 1st mar < 1	0.0545***	0.0558**
	(0.0122)	(0.0022)
Observations	51164	23710

SEs statistics in parentheses. SEs clustered at the state level.

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

IFSS (1955 and 1960 GAF, 1965 and 1970 NFS,

1973 1976, 1982, 1988, 1995 and 2002 NSFG waves)

cohorts $1920 - 1970$, age 20-45, ever marrried women, X: cohort and age dummies

Source of data: GAF, NSF and NSFG refer to public use versions of the Growth of American Families, National Fertility Surveys and National Surveys of Family Growth as compiled by the Integrated Fertility Survey Serie, (Smock et al., 2015)

Overall, this suggests that 1980 was special for two reasons: first 'green' treated cohorts of 1935-1945 divorced as blue early baby-boom 1945-1955 cohorts of women searched for suitable partners among already married men. Second, 'blue' marriages of early baby-boomers that happened at a young age, and among peers or with divorcees were selectively more prone to divorce themselves (with a subsequent remarriage or just because value of singlehood of the men went up as the still large 1955-1964 cohorts were entering the marriageable age).

[5]Without controlling for age at first marriage a small gap is associated with delayed marriage, so the positive effect on divorce is not as apparent.

Figure B.15: Ever-divorced by age-gap at first marriage

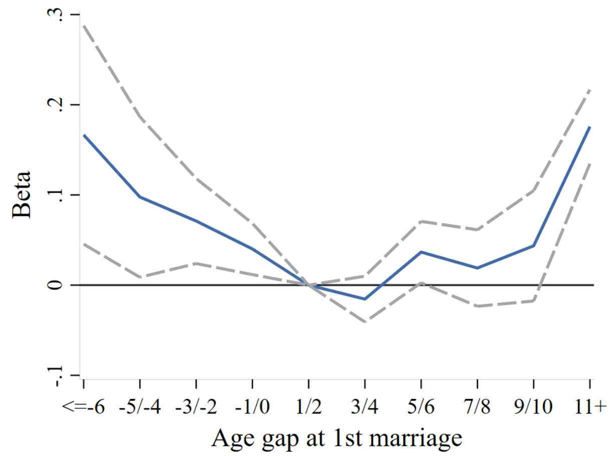

Regressing a dummy of ever being divorced on dummies for age-gap ar first marriage (and age at marriage, age and cohort dummies as controls), plotting coefficients on age-gap dummies (1-2 year age-gap being the base). Source of data: GAF, NSF and NSFG refer to public use versions of the Growth of American Families, National Fertility Surveys and National Surveys of Family Growth as compiled by the Integrated Fertility Survey Series, (Smock et al., 2015)

B.5 Appendix: Model solution

Given the assumptions on idiosyncratic preferences, the within period matching problem is almost fully equivalent to the static model in Choo and Siow (2006), who show that given supplies of men and women of each type, the matches solve a simple system of equations.

Corollary 1. *Given m_{it} and f_{jt}, every period the matching outcomes $n_{ijt}, n_{i\emptyset t}, n_{\emptyset jt}$ satisfy*

$$e^{\frac{\alpha_{ij}-\alpha_{i\emptyset}+\gamma_{ij}-\gamma_{\emptyset j}}{2}} = \frac{n_{ijt}}{\sqrt{n_{i\emptyset t} n_{\emptyset jt}}}$$

and

$$ln\left(\frac{n_{ijt}}{n_{i\emptyset t}}\right) = \alpha_{ij} - \alpha_{i\emptyset} - \tau_{ijt},\ ln\left(\frac{n_{ijt}}{n_{\emptyset jt}}\right) = \gamma_{ij} - \gamma_{j\emptyset} + \tau_{ijt}$$

The strict assumptions of Choo and Siow (2006) are violated in a minor detail, because of the exit of widows and widowers from the marriage market. As a consequence, the distribution of errors in the market is not exactly iid extreme value type I, because widows and widowers are not a random sample of the population. However, numerically the transfers given by $\tau_{ijt} = -ln\left(\frac{n_{ijt}}{n_{i\emptyset t}}\right) + \alpha_{ij}$ with the solution for $\{n_{i,j,t}\}$ as specified in 1 clear the simulated market almost exactly.

APPENDIX C

Appendix to Chapter 3

C.1 Appendix: Descriptive statistics of the sample

Table C.1 presents the summary statistics of the baseline sample. Figure C.1 plots the monetary policy shocks used in the baseline analysis. Figure C.2 shows the shifting age composition in the U.S. over the last century. Figure C.3 shows that across U.S. metropolitan areas, the share of population ages 50-59 correlates negatively with the share of population 35-49 years old. This can explain why in the baseline results metropolitan areas with a high share of 35-49 years old results in lower sensitivity of housing prices.

Figure C.1: Descriptive statistics of the sample

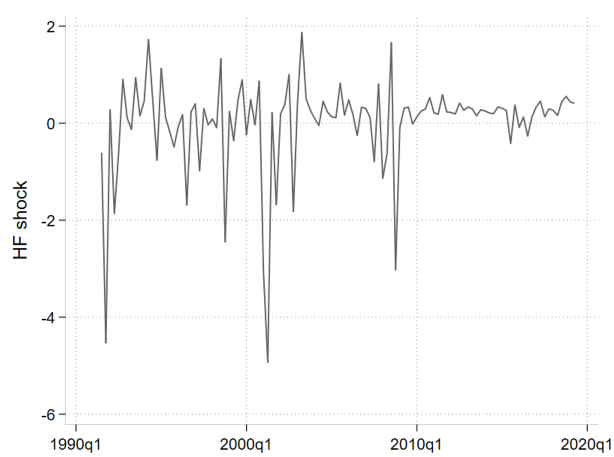

Table C.1: Monetary policy shock from Swanson (2021)

	mean	sd	count
HPI % Δ	0.9	2.1	11200
Share of 50-55 in msa	7.0	0.5	100
Share of 55-60 in msa	6.7	0.6	100
Population in msa	2029769.1	1972484.8	100
HF shock	0.0	1.0	112

Figure C.2: Shifting age composition

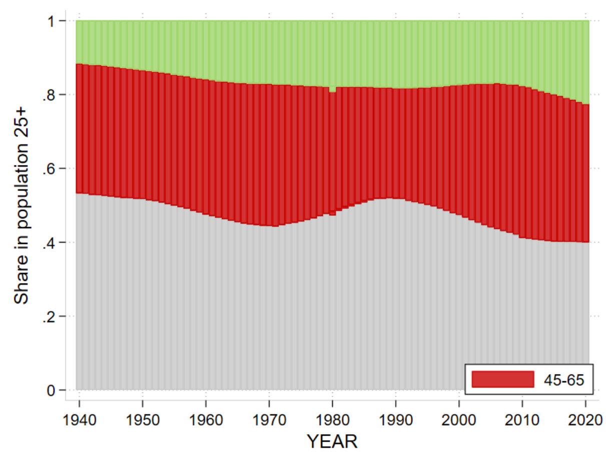

Source: U.S. Census.

Figure C.3: Share in fifties correlates with share 35-45

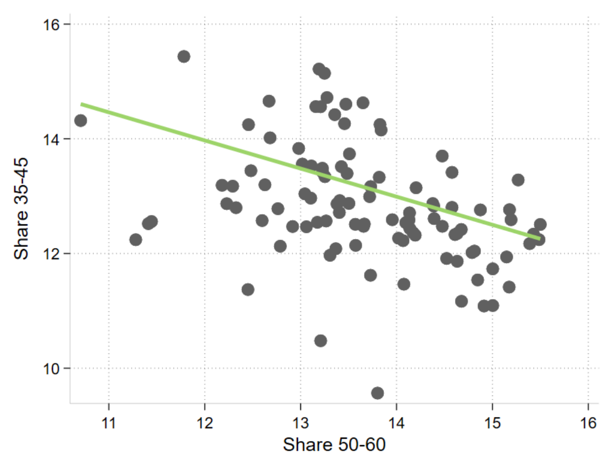

Sample of U.S. metro areas.

C.2 Appendix: Robustness checks

This section presents several robustness checks, as described in the text. Figures C.4 show the results of controlling for local housing supply elasticity in how much monetary policy shocks affect housing prices. Figures C.5 show the results are robust to alternative time subsamples. Figures C.6 show the robustness to using alternative measures of monetary policy shocks available in the literature.

Figure C.4: Differential response of metro-area housing prices: controlling for local housing supply elasticity

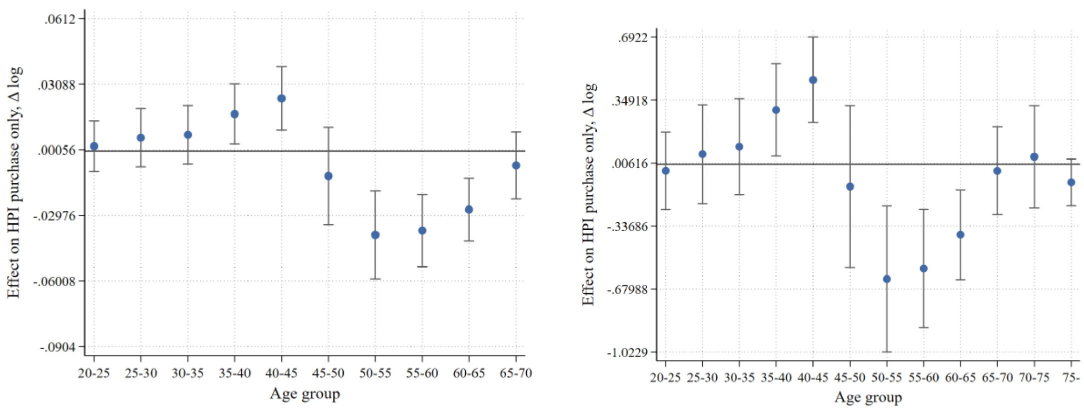

Lag 3 Lag 6

Adding $\alpha^{el} \cdot \tilde{\epsilon}_t \cdot \text{Saiz elasticity}_s$ as an additional control. Metro-level, weighted by population. Lag 3. Source of price indices: FHFA HPI ©.

Figure C.5: Differential response of metro-area housing prices: early and late sub-samples

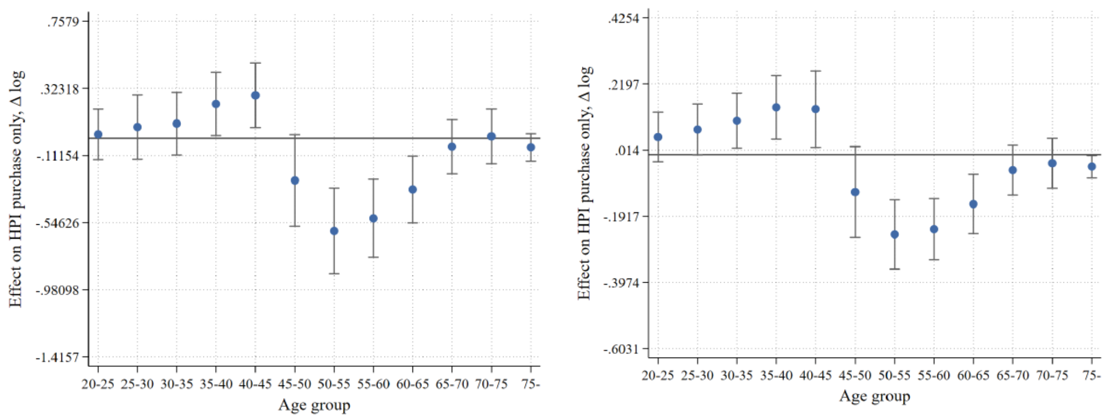

$t \geq 1995$ $t \leq 2007$

Metro-level, weighted by population. Lag 3. Source of price indices: FHFA HPI ©.

Figure C.6: Differential response of metro-area housing prices: alternative monetary policy shock measures

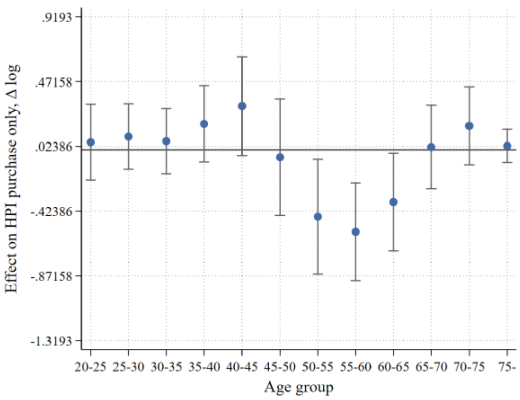

Gorodnichenko and Weber (2016) extended.

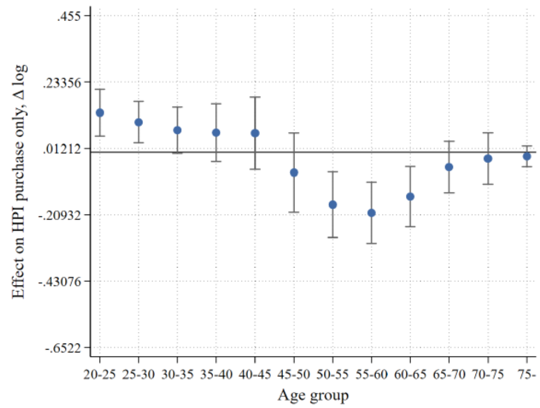

Jarocinski and Karadi (2020)

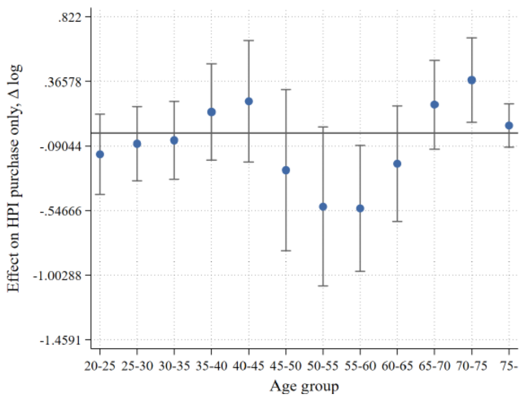

Romer and Romer (2004)

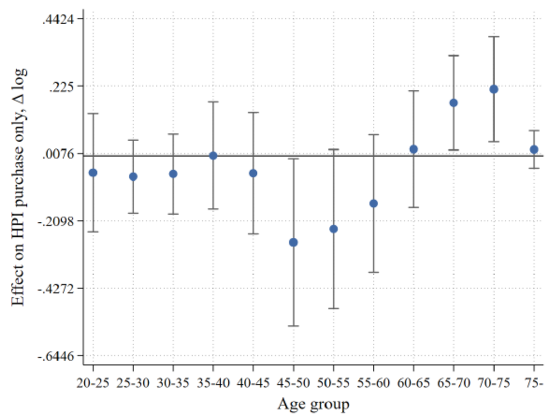

Bu et al. (2021)

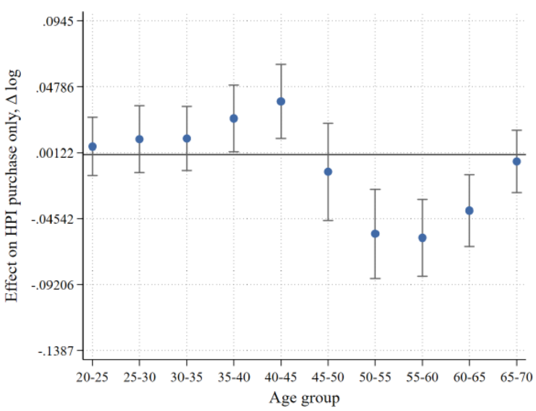

Swanson (2021), reweighted.

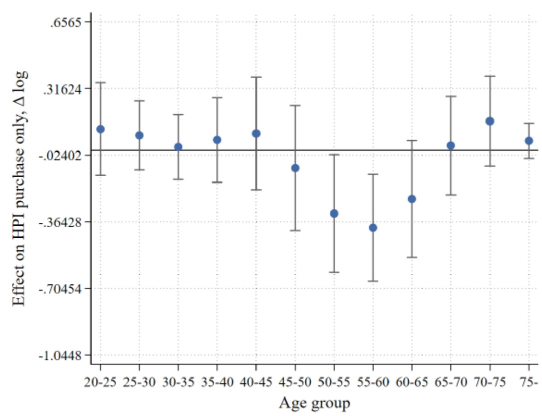

Gurkaynak et al. (2005) as in Paul (2020)

Metro-level, weighted by population. Lag 3. Source of price indices: FHFA HPI ©. (a) Raw 30-minutes window jumps in the price of federal funds rate current month futures around FOMC meetings from Gurkaynak et al. (2005), Gorodnichenko and Weber (2016) and Paul (2020) (Sample: 1991Q1 - 2017Q3). (b) MP shocks constructed in Jarocinski and Karadi (2020) (1991Q1 - 2016Q4). (c) Romer and Romer (2004) shock series, extended by Johannes Wieland and Max Breitenlechner (1991Q1 - 2012Q4). (d) MP shocks constructed by Bu et al. (2021) (1994Q1 - 2019Q3). (e) Baseline MP shock, from Swanson (2021), aggregated using a weighted moving average as in Ottonello and Winberry (2020) (Sample: 1991Q4 - 2019Q2) (f) Target factor from Paul (2020) extracted according to the Gurkaynak et al. (2005) methodology, but excluding all unscheduled FOMC meetings (Sample: 1991Q1 - 2017Q3)

Lastly, I examine the variation in the share of 50-60 year olds across U.S. metropolitan areas. In figure C.7 I visualize the variation on a map. Clearly, a considerable portion of the variation is across-regions. This prompts a suspicion that the measured differential response of housing prices to monetary policy shocks does not arise from the differences in age-structure itself, but from other broad differences across regions. In table C.2 I test how well the results hold up when cross-regional variation is netted out, zooming in at the role of 50-60 year olds and also specifically 55-60 year olds and a lag of 3 quarters. Controlling for a Census region fixed effect dummy one at a time affects the results only marginally. Therefore, I can conclude that the results are not driven by a specific region being very different. However, when netting out all cross-region differences in the response to a shock reduces the baseline result to almost zero, suggesting cross-regional differences in some form are necessary to draw the conclusions of the paper. Upon further investigation, one of the reasons for this is that the within-region variation in the share of 50-60 year olds correlates closely with the within-region variation in the share of an older population in general. Figure C.8 illustrates this point visually. The second map in C.7 also shows that when within-region differences as well as the variation in the share of the population older than 50 is netted out, there is still plenty of variation left. The last column of table C.2 shows that when I net out the differential response in older vs younger areas (including a share of population 50 plus interacted with the shock), the baseline results are robust to including Census region fixed effects interacted with the shock and is thus not driven purely by cross-regional differences.

Table C.2: Differential response of metro-area housing prices: the role of regional and within-regional differences

	$\Delta \log(p_{m,t+i})$, $i=3$					
$s_m^{50-60} \cdot \tilde{\epsilon}_t$	-.387	-.245	-.409	-.204	-.058	-.147
	(.108)	(.105)	(.112)	(.092)	(.094)	(.094)

	$\Delta \log(p_{m,t+i})$, $i=3$					
$s_m^{55-60} \cdot \tilde{\epsilon}_t$	-.378	-.274	-.405	-.240	-.087	-.430
	(.092)	(.094)	(.106)	(.077)	(.071)	(.140)
region 1 $\cdot \delta_t$	x				x	x
region 2 $\cdot \delta_t$		x			x	x
region 3 $\cdot \delta_t$			x		x	x
region 4 $\cdot \delta_t$				x		
$s_m^{50-} \cdot \tilde{\epsilon}_t$						x
N	11200					
N clusters	100					
Sample period	1991Q3 - 2019Q2					

Standard errors in parentheses, clustered at the MSA level.

Sample of 100 largest MSAs and metropolitan divisions, weighted by population. Source of price indices: FHFA HPI ©. Column 1 reproduces the main result from 3.1 with age groups 50-60 and 55-60 and lag of 3 quarters. Columns 2-4 add an interaction of a particular Census region's dummy with the shock. Column 5 adds all of regions 1-3 (adding all regions would result in colinearity). Column 6 adds the region dummies interacted with a shock as well as a share of population 50+ interacted with the shock.

Figure C.7: Map of metropolitan areas by share of 50-60 year olds

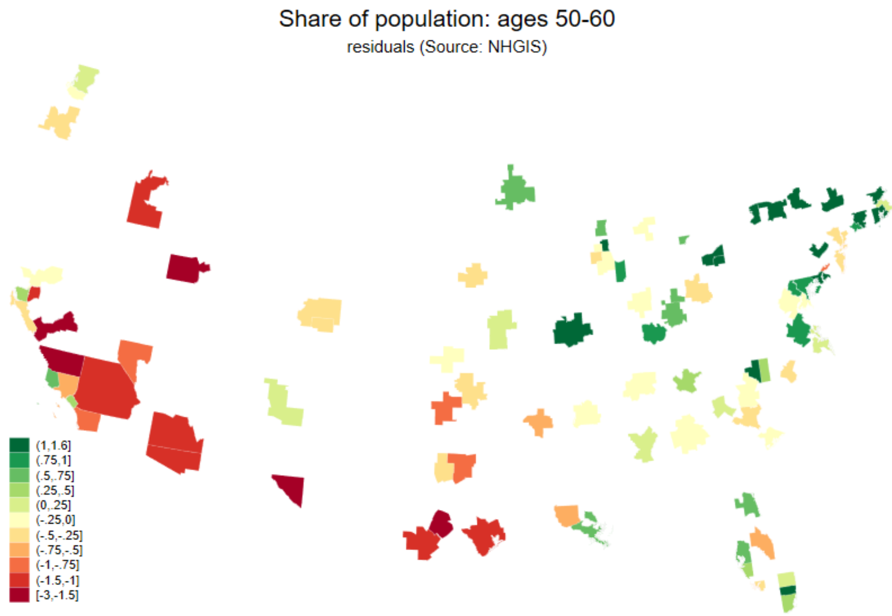

Demeaned share 50-60

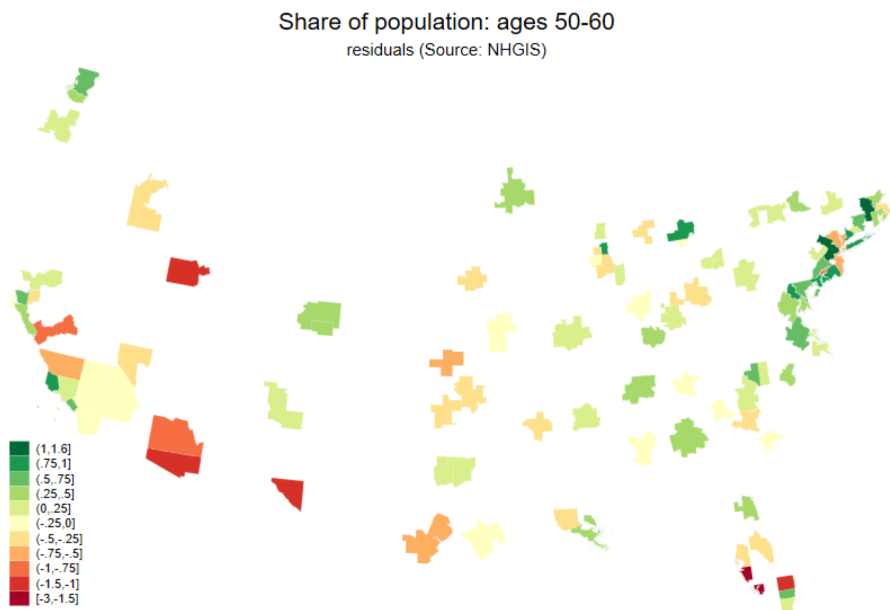

Demeaned share 50-60, within regions & variation in share 50- netted out

Figure C.8: Map of metropolitan areas by share of 50-60 year olds: cont.

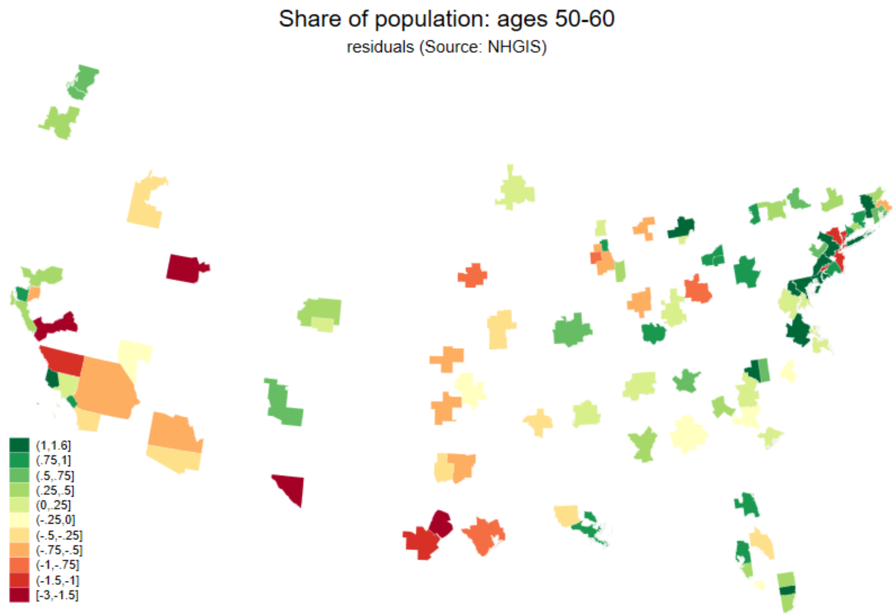

Demeaned share 50-60, within regions

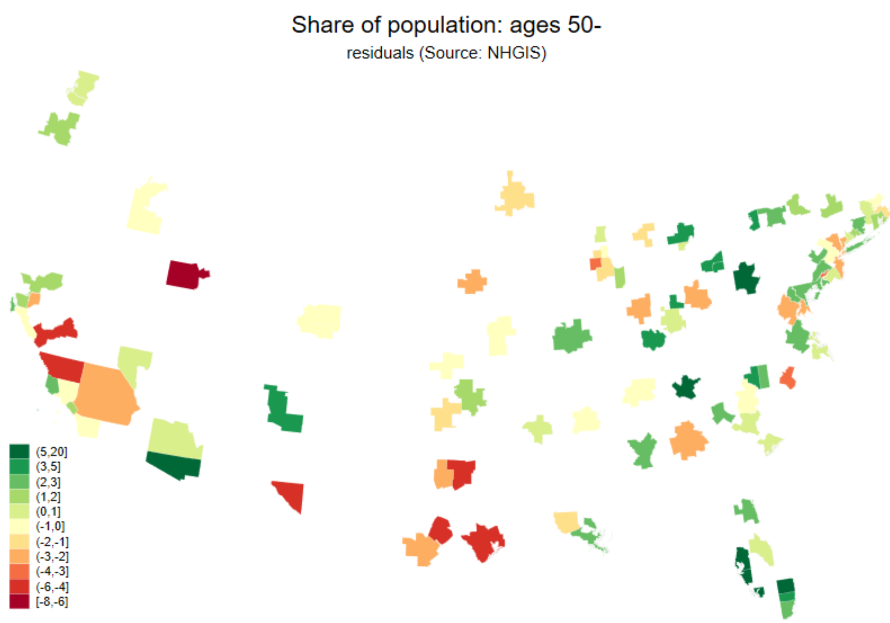

Demeaned share 50-, within regions

Metro-level, weighted by population. The first map shows the geographic distribution of the share 50-60 years old across US metro areas. The second map shows the geographic distribution of the share 50-60 years old across US metro areas with regional differences netted out. The third map shows the geographic distribution of the share 50-60 years old across US metro areas with regional differences netted out. The fourth map shows the geographic distribution of the share 50+ years old across US metro areas with regional differences netted out.

BIBLIOGRAPHY

BIBLIOGRAPHY

Aastveit, K. A. and A. K. Anundsen (2022): "Asymmetric Effects of Monetary Policy in Regional Housing Markets," *American Economic Journal: Macroeconomics*, 14, 499–529.

Abe, Y. (2011): "Family labor supply, commuting time, and residential decisions: The case of the Tokyo Metropolitan Area," *Journal of Housing Economics*, 20, 49–63.

Abramitzky, R., A. Delavande, and L. Vasconcelos (2011): "Marrying Up: The Role of Sex Ratio in Assortative Matching," *American Economic Journal: Applied Economics*, 3, 124–157.

Acemoglu, D., D. H. Autor, and D. Lyle (2004): "Women, war and wages: The effect of female labor supply on the wage structure at midcentury," *Journal of Political Economy*, 112, 497–551.

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015): "The Economics of Density: Evidence From the Berlin Wall," *Econometrica*, 83, 2127–2189.

Ananat, E. O. and G. Michaels (2008): "The effect of marital breakup on the income distribution of women with children," *Journal of Human Resources*, 43, 611–629.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 132, 1553–1592.

Angrist, J. (2002): "How Do Sex Ratios Affect Marriage and Labor Markets? Evidence from America's Second Generation," *The Quarterly Journal of Economics*, 117, 997–1038.

Barbanchon, T. L., R. Rathelot, and A. Roulet (2020): "Gender Differences in Job Search: Trading off Commute against Wage," *The Quarterly Journal of Economics*, 136, 381–426.

Bento, A. M., M. L. Cropper, A. M. Mobarak, and K. Vinha (2005): "The Effects of Urban Spatial Structure on Travel Demand in the United States," *The Review of Economics and Statistics*, 87, 466–478.

Berg, K. A., C. C. Curtis, S. Lugauer, and N. C. Mark (2021): "Demographics and Monetary Policy Shocks," *Journal of Money, Credit and Banking*, 53, 1229–1266.

Bergstrom, T. and D. Lam (1991): "The Two-Sex Problem and the Marriage Squeeze in an Equilibrium Model of Marriage Markets," *Center for Research on Economic and Social Theory CREST Working Paper*, 17.

——— (1994): "The Effects of Cohort Size on Marriage Markets in Twentieth Century Sweden," *The Family, the Market, and the State in Ageing Societies*.

Bertrand, M., E. Kamenica, and J. Pan (2015): "Gender Identity and Relative Income within Households," *The Quarterly Journal of Economics*, 130, 571–614.

Bhrolcháin, M. N. (2001): "Flexibility in the Marriage Market," *Population*, 13, 9–47.

Bianchi, S. M., M. A. Milkie, L. C. Sayer, and J. P. Robinson (2000): "Is Anyone Doing the Housework? Trends in the Gender Division of Household Labor," *Social Forces*, 79, 191–228.

Black, D. A., N. Kolesnikova, and L. J. Taylor (2014): "Why do so few women work in New York (and so many in Minneapolis): Labor supply of married women across US cities," *Journal of Urban Economics*, 79, 59–71.

Blau, B. F. D. and L. M. Kahn (2013): "Female Labor Supply : Why Is the United States Falling Behind ?" *The American Economic Review,Papers and Proceedings*, 103.

Blau, F. D. and L. M. Kahn (2007): "Changes in the Labor Supply Behavior of Married Women: 1980-2000," *Journal of Labor Economics*, 25, 393–438.

——— (2017): "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55, 789–865.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435.

Boca, D. D. (2003): "Mothers, fathers and children after divorce: The role of institutions," *Journal of Population Economics*, 16, 399–422.

Boehm, M. J. (2013): "Concentration versus Re-Matching? Evidence about the Locational Effects of Commuting Costs," *CEP Discussion Paper*.

Borghorst, M., I. Mulalic, and J. van Ommeren (2021): "Commuting, children and the gender wage gap," *Tinbergen Institute Discussion Paper*.

Bound, J., D. A. Jaeger, and R. M. Baker (1995): "Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak," *Journal of the American Statistical Association*, 90, 443–450.

Bourguignon, F., M. Browning, and P. A. Chiappori (2009): "Efficient Intra-Household Allocations and Distribution Factors: Implications and Identification," *The Review of Economic Studies*, 76, 503–528.

Bozon, M. (1991): "Women and the Age Gap Between Spouses: An Accepted Domination?" *Population: An English Selection*, 3, 113–148.

Brainerd, E. (2017): "Uncounted Costs of World War II: The Effect of Changing Sex Ratios on Marriage and Fertility of Russian Women," *Review of Economics and Statistics*.

Bronson, M. A. and M. Mazzocco (2018): "Cohort Size and the Marriage Market: Explaining Nearly a Century of Changes in U.S. Marriage Rates," *Working Paper*, 1–55.

Browning, M., P. Chiappori, and Y. Weiss (2014): *Economics of the Family*, Cambridge University Press.

Bruze, G., M. Svarer, and Y. Weiss (2015): "The dynamics of marriage and divorce," *Journal of Labor Economics*, 33, 123–170.

Bu, C., J. Rogers, and W. Wu (2021): "A Unified Measure of Fed Monetary Policy Shocks," *Journal of Monetary Economics*, 118.

Burdett, K., R. Imai, and R. Wright (2004): "Unstable relationships," *Frontiers of Macroeconomics 1*.

Caldwell, S. and O. Danieli (2023): "Outside Options in the Labor Market," *Review of Economic Studies. Accepted*.

Chauvin, J. P. (2018): "Gender-Segmented Labor Markets and the Effects of Local Demand Shocks," *Inter-American Development Bank Discussion Paper*, IDP-DP-605.

Cherchye, L., B. D. Rock, and F. Vermeulen (2012): "Married with Children: A Collective Labor Supply Model with Detailed Time Use and Intrahousehold Expenditure Information," *American Economic Review*, 102, 3377–3405.

Cherlin, A. J. (2004): "The deinstitutionalization of American marriage," *Journal of Marriage and Family*, 66, 848–861.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014): "Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States," *The Quarterly Journal of Economics*, 129, 1553–1623.

Chiappori, P., A. de Palma, and N. Picard (2018a): "Couple Residential Location and Spouses Workplaces," *Working Paper*.

Chiappori, P., B. Fortin, and G. Lacroix (2002): "Marriage Market, Divorce Legislation, and Household Labor Supply," *Journal of Political Economy*, 110, 37–72.

Chiappori, P.-A., N. Radchenko, and B. Salanie (2018b): "Divorce and the duality of marital payoff," *Review of Economics of the Household*, 16, 833–858.

Chiappori, P. A. and Y. Weiss (2003): "Marriage Contracts and Divorce: an Equilibrium Analysis." *Unpublished manuscript. Department of Economics, University of Chicago*, 1–35.

——— (2007): "Divorce, remarriage, and child support," *Journal of Labor Economics*, 25, 37–74.

Choo, E. (2015): "Dynamic Marriage Matching: An Empirical Framework," *Econometrica*, 83, 1373–1423.

Choo, E. and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114, 175–201.

——— (2007): "Lifecycle marriage matching: Theory and evidence," *Working Paper*.

Cloyne, J., C. Ferreira, and P. Surico (2020): "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *The Review of Economic Studies*, 87, 102–129.

Compton, J. and R. A. Pollak (2007): "Why Are Power Couples Increasingly Concentrated in Large Metropolitan Areas ?" *Journal of Labor Economics*.

Cornelius, T. J. (2003): "A search model of marriage and divorce," *Review of Economic Dynamics*, 6, 135–155.

Costa, D. L. and M. E. Kahn (2000): "Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990," *The Quarterly Journal of Economics*, 115, 1287–1315.

Croix, D. D. L. and F. Mariani (2015): "From Polygyny to Serial Monogamy: A Unified Theory of Marriage Institutions," *The Review of Economic Studies*, 82, 565–607.

Delventhal, M. J., E. Kwon, and A. Parkhomenko (2022): "JUE Insight: How do cities change when we work from home?" *Journal of Urban Economics*, 127.

Depew, B. and J. Price (2018): "Marriage and the economic status of women with children," *Review of Economics of the Household*, 16, 1049–1061.

Doepke, M., M. Hazan, and Y. D. Maoz (2013): "The baby boom and world war II: A macroeconomic analysis," *Review of Economic Studies*, 82, 1031–1073.

Díaz-Giménez, J. and E. Giolito (2013): "Accounting for the timing of first marriage," *International Economic Review*, 54, 135–158.

Ehrlich, M. V., C. A. Hilber, and O. Schoni (2018): "Institutional settings and urban sprawl: Evidence from Europe," *Journal of Housing Economics*, 42, 4–18.

England, P. and E. A. McClintock (2009): "The gendered double standard of aging in US marriage markets," *Population and Development Review*, 35, 797–816.

Ewing, R. and S. Hamidi (2015): "Compactness versus Sprawl: A Review of Recent Evidence from the United States," *https://doi.org/10.1177/0885412215595439*, 30, 413–432.

Fan, J. and B. Zou (2021): "The Dual Local Markets: Family, Jobs, and the Spatial Distribution of Skills," *Working paper*.

Farre, L., J. Jofre-Monseny, and J. Torrecillas (2020): “Commuting Time and the Gender Gap in Labor Market Participation,” *IZA Discussion Paper*.

Fischer, M. M., F. Huber, M. Pfarrhofer, and P. Staufer-Steinnocher (2021): “The Dynamic Impact of Monetary Policy on Regional Housing Prices in the United States,” *Real Estate Economics*, 49, 1039–1068.

Friedberg, L. (1998): “Did Unilateral Divorce Raise Divorce Rates? Evidence from Panel Data,” *American Economic Review*, 88.

Fu, S. and S. L. Ross (2013): “Wage premia in employment clusters: How important is worker heterogeneity?” *Journal of Labor Economics*, 31, 271–304.

Fuss, R. and J. Zietz (2016): “The economic drivers of differences in house price inflation rates across MSAs,” *Journal of Housing Economics*, 31, 35–53.

Garcia-Moran, E. M. (2018): “Differential fecundity and child custody,” *Journal of Economic Dynamics and Control*, 90, 156–170.

Gayle, G.-L., L. Golan, and M. A. Soytas (2015): “What Accounts for the Racial Gap in Time Allocation and Intergenerational Transmission of Human Capital?” *Federal Reserve Bank of St. Louis Working Papers*.

Gayle, G.-L. and A. Shephard (2019): “Optimal Taxation, Marriage, Home Production, and Family Labor Supply,” *Econometrica*, 87, 291–326.

Gemici, A. (2008): “Family Migration and Labor Market Outcomes,” *Manuscript, New York University*.

Giolito, E. (2004): “A Search Model of Marriage with Differential Fecundity,” *Working paper*.

Glaeser, E. L. and M. E. Kahn (2004): “Sprawl and Urban Growth,” *Handbook of Regional and Urban Economics*, 4, 2481–2527.

Goldin, C. and L. F. Katz (2016): “Changes in Marriage and Divorce as Drivers of Employment and Retirement of Older Women,” *NBER Working paper*, 113–156.

Gorodnichenko, Y. and M. Weber (2016): “Are Sticky Prices Costly? Evidence from the Stock Market,” *American Economic Review*, 106, 165–99.

Grammich, C., K. Hadaway, R. Houseal, D. E. Jones, A. Krindatch, R. Stanley, and R. H. Taylor (2012): “2010 U.S. Religion Census: Religious Congregations and Membership Study,” *Association of Statisticians of American Religious Bodies*.

Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2016): “Technology and the changing family: a unified model of marriage, divorce, educational attainment and married female labor-force participation,” *American Economic Journal: Macroeconomics*, 8, 1–41.

GRONAU, R. (1977): "Leisure , Home Production , and Work - the Theory of the Allocation of Time Revisited," *Journal of Political Economy*, 85, 1099–1123.

GROVES, E. R. AND W. F. OGBURN (1928): *Marriage and Family Relations*.

GURKAYNAK, R. S., B. P. SACK, AND E. T. SWANSON (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *SSRN Electronic Journal*.

GUTIERREZ, F. (2018): "Commuting Patterns, the Spatial Distribution of Jobs and the Gender Pay Gap in the U.S." *SSRN Electronic Journal*, 1–37.

GYOURKO, J. AND R. MOLLOY (2015): "Regulation and Housing Supply," .

HARARI, M. (2020): "Cities in Bad Shape: Urban Geometry in India," *American Economic Review*, 110, 2377–2421.

HILLER, N. AND O. W. LERBS (2016): "Aging and urban house prices," *Regional Science and Urban Economics*, 60, 276–291.

HOBBS, F. AND N. STOOPS (2000): "Demographic Trends in the 20th Century," *U.S. Census Bureau, Census 2000 Special Reports*, CENSR-4.

HREHOVA, K., E. SANDOW, AND U. LINDGREN (2021): "Firm Relocations, Commuting and Relationship Stability," *SSRN Electronic Journal*.

IACOVIELLO, M. AND M. PAVAN (2013): "Housing and debt over the life cycle and over the business cycle," *Journal of Monetary Economics*, 60, 221–238.

JAROCINSKI, M. AND P. KARADI (2020): "Deconstructing Monetary Policy Surprises - The Role of Information Shocks," *American Economic Journal: Macroeconomics*, 12, 1–43.

JOHNSON, W. R. AND J. SKINNER (1986): "Labor Supply and Marital Separation," *American Economic Review*, 76, 455–69.

JONES, D. E., C. GRAMMICH, J. E. HORSCH, R. HOUSEAL, M. LYNN, J. MARCUM, K. M. SANCHAGRIN, D. SHERRI, AND R. H. TAYLOR (2000): "2000 U.S. Religion Census: Religious Congregations and Membership Study." *Association of Statisticians of American Religious Bodies*.

JORDA, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

JORDA, O., M. SCHULARICK, AND A. M. TAYLOR (2015): "Betting the house," *Journal of International Economics*, 96, S2–S18.

KENRICK, D. T. AND R. C. KEEFE (1992): "Age preferences in mates reflect sex differences in human reproductive strategies," *Behavioral and Brain Sciences*, 15, 75–91.

Kenrick, D. T., R. C. Keefe, C. Gabrielidis, and J. S. Cornelius (1996): “Adolescents’ Age Preferences for Dating Partners: Support for an Evolutionary Model of Life-History Strategies,” *Child Development*, 67, 1499–1511.

Kuttner, K. N. (2001): “Monetary policy surprises and interest rates: Evidence from the Fed funds futures market,” *Journal of Monetary Economics*, 47, 523–544.

Lafortune, J. and C. Low (2017): “Tying the Double-Knot: The Role of Assets in Marriage Commitment,” *American Economic Review: Papers & Proceedings*, 107, 163–167.

Leahy, J. V. and A. Thapar (2022): “Age Structure and the Impact of Monetary Policy,” *American Economic Journal: Macroeconomics*.

Leamer, E. (2015): “Housing Really Is the Business Cycle: What Survives the Lessons of 2008-09?” *Journal of Money, Credit and Banking*, 47, 43–50.

Leip, D. (2021): “Dave Leip’s Atlas of U.S. Presidential Elections,” .

Leukhina, O. and Z. Yu (2022): “Home Production and Leisure during the COVID-19 Recession,” *B.E. Journal of Macroeconomics*, 22, 269–306.

Liu, S. and Y. Su (2020): “The Geography of Jobs and the Gender Wage Gap,” *Federal Reserve Bank of Dallas, Working Papers*, 2020.

Logan, J. R., B. J. Stults, and Z. Xu (2016): “Validating Population Estimates for Harmonized Census Tract Data, 2000-2010,” *Annals of the American Association of Geographers*, 106, 1013–1029.

Low, C. (2023a): “The Human Capital-Reproductive Capital Tradeoff in Marriage Market Matching,” *Journal of Political Economy*.

——— (2023b): “Pricing the Biological Clock: The Marriage Market Costs of Aging to Women,” *Journal of Labor Economics*.

Madden, J. F. (1977): “Spatial Theory of Sex Discrimination,” *Journal of Regional Science*, 17, 369–380.

——— (1981): “Why Women Work Closer to Home,” *Urban Studies*, 18, 181–194.

Mankiw, N. and D. N. Weil (1989): “The baby boom, the baby bust, and the housing market,” *Regional Science and Urban Economics*, 19, 235–258.

Manson, S., J. Schroeder, D. V. Riper, T. Kugler, and S. Ruggles (2021): “National Historical Geographic Information System: Version 16.0 [dataset],” *Minneapolis, MN: IPUMS*.

McFadden, D. (1977): “Modelling the Choice of Residential Location,” *Spatial Interaction Theory and Planning Models*.

McKinnish, T. G. (2007): "Sexually Integrated Workplaces and Divorce," *Journal of Human Resources*, XLII, 331–352.

Moffitt, R. (1992): "Incentive effects of the US welfare system: A review," *Journal of Economic Literature*, 30, 1–61.

——— (1997): "The Effect of Welfare on Marriage and Fertility: What Do We Know and What Do We Need to Know?" *Institute for Research on Poverty Discussion Papers, University of Wisconsin Institute for Research on Poverty.*

Moreno-Maldonado, A. (2022): "Mums and the City: Household labour supply and location choice," *Working paper.*

Mortensen, D. T. (1988): "Matching: Finding a Partner for Life or Otherwise," *American Journal of Sociology*, 94, S215–S240.

Nakamura, E. and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 133, 1283–1330.

Ottonello, P. and T. Winberry (2020): "Financial heterogeneity and the investment channel of monetary policy," *Econometrica*, 88, 2473–2502.

Ozimek, A. and E. Carlson (2023): "Remote Work and Household Formation," *Working Paper.*

Paul, P. (2020): "The Time-Varying Effect of Monetary Policy on Asset Prices," *The Review of Economics and Statistics*, 102, 690–704.

Petrongolo, B. and M. Ronchi (2020): "Gender gaps and the structure of local labor markets," *Labour Economics*, 64, 101819.

Reardon, S. F., A. D. Ho, B. R. Shear, E. M. Fahle, D. Kalogrides, H. Jang, and B. Chavez (2021): "Stanford Education Data Archive (Version 4.1)." .

Redding, S. J. and E. Rossi-Hansberg (2017): "Quantitative Spatial Economics," *Annual Review of Economics*, 9, 21–58.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084.

Rosenthal, S. S. and W. C. Strange (2012): "Female entrepreneurship, agglomeration, and a new spatial mismatch," *Review of Economics and Statistics*, 94, 764–788.

Rotz, D. (2016): "Why have divorce rates fallen? The role of women's age at marriage," *Journal of Human Resources*, 51, 961–1002.

Ruggles, S. (1997): "The rise of divorce and separation in the United States, 1880-1990," *Demography*, 34, 455–466.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2019): "IPUMS USA: Version 9.0 [dataset]," *Minneapolis, MN: IPUMS.*

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

Shephard, A. (2019): "Marriage Market Dynamics, Gender, and the Age Gap," *SSRN Electronic Journal*.

Smock, P., P. Granda, and L. Hoelter (2015): "Integrated Fertility Survey Series, Release 7, 1955-2002 [United States]." *Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2015-06-18*.

Staiger, D. and J. H. Stock (1997): "Instrumental Variables Regression with Weak Instrument," *The Econometric Society*, 65, 557–586.

Sterk, V. and S. Tenreyro (2018): "The transmission of monetary policy through redistributions and durable purchases," *Journal of Monetary Economics*, 99, 124–137.

Stevenson, B. and J. Wolfers (2007): "Marriage and Divorce: Changes and their Driving Forces," *Journal of Economic Perspectives*, 21, 27–52.

Swanson, E. T. (2021): "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets," *Journal of Monetary Economics*, 118, 32–53.

Takats, E. (2012): "Aging and house prices," *Journal of Housing Economics*, 21, 131–141.

Tkocz, Z. and G. Kristemen (1994): "Commuting Distances and Gender: A Spatial Urban Model," *Geographical Analysis*, 26, 1–14.

Tscharaktschiew, S. and G. Hirte (2010): "How does the household structure shape the urban economy?" *Regional Science and Urban Economics*, 40, 498–516.

Turner, T. and D. Niemeier (1997): "Travel to work and household responsibility: New evidence," *Transportation*, 24, 397–419.

U.S.BureauofLaborStatistics (2020): "Consumer Expenditures Report 2019," .

U.S.CensusBureau (2021): "LEHD Origin-Destination Employment Statistics Data (2002-2018) [computer file]. Washington, DC: U.S. Census Bureau, Longitudinal-Employer Household Dynamics Program [distributor]," .

Venator, J. (2020): "Dual Earner Migration Patterns: The Role of Locational Compatibility within Households," *Working paper*.

Weiss, Y. and R. J. Willis (1997): "Match quality, new information, and marital dissolution," *Journal of Labor Economics*, 15.

White, M. J. (1986): "Sex Differences in Urban Commuting Patterns," *American Economic Review*, 76, 368–372.

Wolfers, J. (2006): "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results," *American Economic Review*, 96, 1802–1820.

Wong, A. (2018): "Population Aging and the Transmission of Monetary Policy to Consumption," *Working Paper*, 1–78.

Yinger, J. (2021): "The price of access to jobs: Bid-function envelopes for commuting costs," *Journal of Housing Economics*, 51, 101742.