

UNPACKING MOVING: A QUANTITATIVE SPATIAL EQUILIBRIUM MODEL WITH WEALTH*

Elisa Giannone

CREI

Nuno Paixão

Bank of Canada

Xinle Pang

SUNY Buffalo

Qi Li

Penn State

December 2025

Abstract

This paper examines how household wealth—both liquid and illiquid—contributes to muted migration responses to local economic downturns. We introduce a novel *housing-wealth lock-in* mechanism: temporary declines in house prices following local shocks erode homeowners’ net worth, reducing their willingness to relocate. We formalize this channel in a quantitative dynamic spatial equilibrium model with uninsurable income risk, incomplete markets, and wealth accumulation in liquid assets and owner-occupied housing. The model generates precautionary migration motives stronger for households with limited wealth and financial insurance, producing heterogeneous migration elasticities even with homogeneous moving costs. We discipline the model using granular Canadian data and exploit oil-producing provinces’ exposure to the 2014–2018 global oil price collapse, which generated sharp, persistent declines in local incomes and house prices. Absent the house price decline, out-migration would have been nearly three times larger. While this wealth effect generates a lock-in analogous to a rent-based lock-in, distinguishing the two is crucial, as they have opposite welfare implications: rent declines lower housing expenditures for renters, mitigating losses by 52 percent, whereas house price declines reduce homeowners’ net worth, constrain relocation, and amplify welfare costs by 45 percent.

JEL codes: G51, R12, R13, R2, R31, R52

Keywords: dynamic spatial models, wealth, incomplete markets, housing

*We want to thank Adam Su for excellent research assistance. We also thank Jason Allen, Lorenzo Caliendo, Katya Kartashova, James MacGee, Alexander Michaelides (discussant), Joan Monras, David Nagy, Giovanni Gallipoli (discussant), Fernando Parro, Giacomo Ponzetto, Esteban Rossi-Hansberg, Stijn Van Nieuwerburgh, Walter Steingress, Micheal Waugh and many seminar participants at the Columbia University, Dallas FED International Economics Conference, EUI, LSE, Minneapolis Fed, NYU, Penn State University, Princeton University, Warwick University, Virtual Macro and Trade Seminar Series, Philadelphia FED, University of Michigan, CMU/UPitts, Yale University (Trade), Fordham University, San Francisco FED, Royal Holloway University, Bank of Canada, CURE, CREI, Dartmouth, NYU-Stern, Yale University (Macro), Richmond FED, NYU-Abu Dabi, Pizzanomics (EIEF), SED, FRB, New York FED, PEJ, University of Toulouse, Science Po University, Bocconi University, University of Washington, Universitat de Barcelona, CMSG, UEA - North American Meetings, Oxford, University of Montreal, St. Louis Fed, OSUS and Cambridge for thoughtful and generous comments. Elisa Giannone acknowledges hospitality from the Minneapolis Fed where part of this work was completed. The views expressed in this paper are those of the authors and do not reflect those of the Bank of Canada.

1 Introduction

Economic disparities across regions remain large and, in many countries, have widened over time. Local economic shocks—such as recessions or financial crises—have heterogeneous and persistent effects across locations, creating spatial wedges in employment and income. Migration could be an adjustment margin: by relocating, households could escape depressed local conditions and attenuate the consequences of adverse shocks. Yet internal migration has declined in developed economies, and a growing empirical literature documents remarkably muted mobility responses even to large negative local shocks.¹

A prominent explanation emphasizes rent-based lock-in: in rent-only frameworks, negative local shocks lower rents, mechanically reducing housing costs and weakening incentives to relocate (e.g., Glaeser and Gyourko, 2005). In many economies, however, homeownership is widespread, so housing-market adjustment operates importantly through changes in asset values rather than solely through the flow cost of housing. This shifts attention to household balance sheets. Standard macro models highlight that wealth and financial access are central for smoothing shocks: wealthier households can draw down savings or borrow, while financially constrained households experience larger consumption drops (e.g., Kaplan, Violante and Weidner, 2014). These considerations naturally extend to spatial adjustment and motivate the paper’s central question: *how does wealth—both liquid and illiquid—shape migration behavior and the response to local economic downturns?*

To this end, we develop a quantitative dynamic spatial equilibrium model with moving costs, income risk, incomplete markets, and wealth accumulation in both liquid assets and owner-occupied housing. This dynamic environment links spatial adjustment to wealth composition. On one hand, *liquid wealth* provides insurance that is partially substitutable for migration: households with larger buffers can smooth shocks without relocating. On the other hand, *homeownership* reduces mobility both through transaction costs and through a *housing-wealth lock-in* effect: when a local downturn temporarily depresses house prices, forward-looking homeowners face an intertemporal trade-off between selling their property at depressed prices (and forgoing future capital gains) and remaining in place despite weaker labor-market conditions. Consequently, out-migration declines with house prices even when transaction costs are absent.

¹These patterns are particularly striking in the context of recent crises. During the U.S. Great Financial Crisis (GFC), regions experiencing large housing and employment contractions exhibited limited out-migration despite severe and persistent income losses (e.g., Dao, Furceri and Loungani, 2017; Kothari, Saporta-Eksten and Yu, 2013; Monras, 2018; Yagan, 2019). Similarly, during the European sovereign debt crisis, deeply affected regions in Southern Europe—such as Spain, Italy, and Greece—saw sharp declines in employment and wages, yet relatively modest migration responses, especially among homeowners (e.g., Mitze, 2019).

Importantly, while both rent-based lock-in and *housing-wealth lock-in* reduce mobility, their welfare implications are opposite: declines in rents mitigate renters’ losses by lowering housing expenditures, whereas house price declines amplify homeowners’ losses by eroding households’ wealth and constraining reallocation. Overall, the paper shows that housing-market dynamics and household balance sheets are central to understanding the migration and welfare effects of local shocks. As governments are increasingly called upon to implement policies that mitigate such shocks, policy responses should distinguish transmission channels to avoid unintended winners and losers.

We begin by documenting new evidence on how mobility varies with wealth determinants in Canada using a rich monthly panel of individuals from *TransUnion Canada*, a credit-bureau dataset covering nearly every person with a credit report and tracking detailed liabilities, credit characteristics, and residential locations over time. We show that migration propensities decline monotonically with borrowing capacity and are lower for homeowners, conditional on age and balance-sheet characteristics.² These patterns motivate a framework in which location choices and wealth accumulation are jointly determined.

We develop a quantitative model with overlapping generations of risk-averse households facing idiosyncratic income risk and incomplete markets. Households can accumulate illiquid wealth by purchasing housing units and save or borrow in a one-period liquid asset subject to a collateral constraint. Every period, households choose where to live, whether to rent or own, non-housing and housing consumption, and savings, taking into account monetary and utility moving costs. Locations differ in productivity, labor-market risk, amenities, and housing-market characteristics. Wages, rents, and house prices are endogenously determined. Combining a spatial equilibrium structure with multiple assets that differ in their degree of liquidity is essential for quantifying how income risk, liquid wealth, and housing wealth determine migration as a self-insurance mechanism and for tracing the aggregate and distributional effects of local shocks.

However, introducing these features comes at the cost of high dimensionality, which poses substantial computational challenges. We address this by combining solution methods from the quantitative spatial literature with those from the macroeconomic literature on heterogeneous agents. Because the joint determination of location choice and consumption–savings decisions

²Conditional on location and time, individuals with a credit score above 800 (super-prime borrowers) are 1.06 percentage points less likely to move to a different city in a given year than those with a credit score below 640 (very poor borrowers). Homeowners are 0.58 percentage points less likely to move than renters. In Appendix B, using direct measures of liquid and illiquid wealth, we show a similar pattern in the U.S., controlling for education, income, industry, occupation, and family composition.

under incomplete markets and non-myopic preferences generates highly nonlinear policy functions, central to our core results, we cannot rely on dynamic hat-algebra for counterfactual analysis, as in spatial models with capital accumulation such as [Kleinman, Liu and Redding \(2023\)](#) and [Dvorkin \(2023\)](#). Instead, we implement a global solution method that explicitly tracks the wealth distribution within and across locations, following the dynamic heterogeneous-agent literature.

We discipline the model to Canada’s 27 largest Census Metropolitan Areas (CMAs). The parameterization requires solving both a stationary equilibrium and the full transition dynamics after a negative local shock. We match cross-sectional moments from 2011–2013 and the observed out-migration from oil-producing provinces following the large and persistent 2014–2018 collapse in global oil prices—Canada’s largest economic downturn in the past three decades outside the COVID period. We use this episode to discipline the model parameterization, particularly the dispersion of idiosyncratic local taste shocks that governs migration elasticities, and to characterize the local downturns studied in the final section. Migration elasticities, the income process, and moving costs are central. We discipline the income process using multiple administrative sources to keep it as data-driven as possible. The baseline parameterization implies an average intercity moving cost of about USD 306,000 (2010 dollars), slightly below the previous literature estimates (e.g., [Kennan and Walker, 2011](#)).

The joint determination of location and consumption–saving decisions lies at the core of the model’s mechanism. Under uninsurable income risk, households can smooth shocks either through financial markets or by relocating; the model therefore generates both precautionary saving and precautionary migration. Wealth and migration act as imperfect substitute insurance channels, delivering endogenous heterogeneity in migration rates across age, homeownership status, and wealth despite homogeneous moving costs.

To quantify these mechanisms, we trace mobility responses to idiosyncratic, transitory-but-persistent negative income shocks and conduct a decomposition that sequentially shuts down income risk, borrowing constraints, and housing transaction costs.³ Eliminating income risk or borrowing constraints reduces the migration response by approximately 73 percent and 33 percent, respectively, relative to the baseline. These effects are concentrated among financially constrained households: in the baseline economy, the migration response of households in the bottom wealth quartile is roughly three times that of those in the top quartile, whereas these

³Because agents are atomistic, these perturbations generate no general-equilibrium effects and prices remain fixed at stationary-equilibrium levels, isolating how individual policy functions depend on each mechanism. Aggregating individual responses then yields a partial-equilibrium migration response to a local productivity shock.

gradients largely vanish once income risk or borrowing constraints are removed. In contrast, housing transaction costs play a minor role for idiosyncratic shocks: eliminating them increases migration response to the shock by only 0.05 percentage points.

In the final part of the paper, we investigate how both liquid and illiquid wealth affect migration during local downturns and the channels through which these effects operate. Leveraging cross-city variation in oil exposure and the 2014–2018 global oil price collapse, we estimate that the downturn generated substantial declines in wages and house prices but only modest out-migration. We then impose an oil-city productivity decline and solve the full transition dynamics, allowing wages, rents, and house prices to adjust endogenously. The model shows that, absent the decline in house prices, homeowners’ migration response would be three times higher, highlighting a housing-wealth channel that suppresses mobility during downturns. At the same time, our framework captures the rent-based lock-in of [Glaeser and Gyourko \(2005\)](#). Distinguishing these channels is crucial because they have opposite welfare consequences: falling rents mitigate renters’ losses, while falling house prices amplify homeowners’ losses. In oil cities, aggregate welfare falls by 1.52 percent in the baseline, compared with 1.30 percent under fixed house prices. Housing price declines account for roughly 45 percent of homeowners’ welfare loss, while declining rents attenuate renters’ welfare loss by 52 percent. Finally, although transaction costs are negligible for idiosyncratic shocks, they constrain mobility during downturns: absent transaction costs, homeowners’ migration response would be three times larger on impact, though transaction costs contribute less to total welfare losses than the house price channel.

Taken together, our results show that housing-market adjustments and household balance sheets are central to understanding the migration and welfare consequences of local shocks. Models that abstract from homeownership and the illiquidity of housing wealth may draw misleading conclusions about spatial adjustment and risk, and may therefore misguide policy in environments where interventions can create distinct winners and losers across the tenure and wealth distributions.

Related Literature. This paper relates to several strands of the quantitative spatial and macro literature. A large literature on location choice emphasizes spatial heterogeneity and moving costs in static or partial-equilibrium settings (see, [Redding and Rossi-Hansberg, 2017](#)). A growing set of dynamic spatial models incorporates forward-looking migration but abstracts from consumption-saving under incomplete markets (e.g., [Caliendo, Dvorkin and Parro, 2019](#); [Desmet and Rossi-Hansberg, 2014](#); [Desmet, Nagy and Rossi-Hansberg, 2018](#); [Eckert and Kleineberg, 2019](#); [Giannone, 2017](#); [Lyon and Waugh, 2018](#); [Oswald, 2019](#)). Recent dynamic general-equilibrium frameworks with capital accumulation allow forward-looking decisions and transition dynamics (e.g., [Cai](#)

et al., 2022; Kleinman, Liu and Redding, 2023), but the key decisions are typically separated across agent types: workers move but are hand-to-mouth, while asset owners accumulate wealth but are immobile. By contrast, we study migration as an insurance margin in an environment where a single agent jointly chooses location and wealth accumulation under incomplete markets.

Bilal and Rossi-Hansberg (2021) formalize how consumption–saving decisions shape location choice via the *location-as-an-asset* mechanism. We build on this framework and introduce an *asset-in-a-location* in the form of homeownership: locations differ not only in contemporaneous housing costs but also in their capacity to create and preserve wealth. We embed two assets with different degrees of liquidity in a dynamic spatial equilibrium model with costly migration, endogenous factor prices, rich spatial heterogeneity, and life-cycle dynamics. This framework links location choice with balance-sheet composition, generating an additional source of sorting across space. It quantitatively matches the heterogeneity in migration rates across demographic groups in the data, even under homogeneous moving costs. Finally, it introduces novel channels through which local shocks are transmitted, with welfare implications that can be opposite to those in frameworks without homeownership.

Glaeser and Gyourko (2005) develop a theory of urban decline with durable housing to explain why negative shocks lead to larger house price declines relative to population. Unlike their rent-only economy, our model allows households to rent or own, endogenizing wealth that fluctuates with house prices. This has two key implications. First, it introduces a new source of lock-in: while a decline in rents reduces renters’ incentives to move, our framework adds the *housing-wealth lock-in* effect, as declines in house prices erode homeowners’ wealth and moving requires forgoing future capital gains. Second, the aggregate welfare implications differ: house price declines amplify losses for homeowners—and therefore aggregate welfare—whereas in rent-only frameworks, rent declines mitigate aggregate welfare losses.

In contemporaneous work, Greaney (2020) develops a related continuous-time framework, focusing on uneven regional growth and wealth inequality rather than migration responses to local shocks.

Naturally, our work also relates to the literature analyzing migration rates and moving costs (e.g., Diamond, McQuade and Qian, 2019; Kennan and Walker, 2011; Molloy, Smith and Wozniak, 2014). By incorporating income risk, incomplete markets, homeownership, and consumption–saving decisions, the benefits of moving become state-dependent. Consequently, the framework generates heterogeneous migration rates even if moving costs are assumed identical across demographic groups, showing that variation in mobility does not necessarily stem from differences in preferences or costs. Moreover, ignoring the substitutability between migration

and financial insurance can lead to overstated moving costs for older, wealthier homeowners.

We also contribute to the macro literature on heterogeneous agents and incomplete markets and idiosyncratic risk (e.g., [Aiyagari, 1994](#); [Huggett, 1993](#); [Imrohoroglu, 1989](#)), including two-asset environments (e.g., [Heathcote, Storesletten and Violante, 2009](#); [Kaplan and Violante, 2014](#); [Kaplan, Mitman and Violante, 2020b](#)). In this literature, financial wealth and access to credit are typically the sole source of self-insurance. Our framework adds migration as an additional (costly) insurance margin, generating *precautionary saving* and *precautionary migration* simultaneously, which allows us to quantify how financial access shapes mobility responses to shocks. We also contribute to the literature on the effects of house price fluctuations. While prior work focuses on how declines in house prices generate negative wealth effects that impact consumption and leverage (e.g., [Berger et al., 2018](#)), we show their implications for migration and its interaction with consumption and wealth accumulation.

Finally, we relate to the literature on *lock-in* effects. Empirical evidence following the 2008 financial crisis shows that negative home equity substantially reduces mobility (e.g., [Andersson and Mayock, 2014](#); [Ferreira, Gyourko and Tracy, 2010](#); [Schulhofer-Wohl, 2011](#)). We complement this research by providing a spatial quantitative framework that formalizes this mechanism. Importantly, the framework suggests that negative equity is not required to deter migration: even transitory declines in housing wealth can reduce mobility, with implications for policymakers concerned about the broader consequences of housing market downturns.

The rest of the paper is organized as follows. Section 2 describes the Canadian *TransUnion* data and documents empirical regularities on migration patterns. Section 3 develops the theoretical framework and discusses the solution method. Section 4 describes the parameterization strategy, model fit, and the quantitative importance of the model’s mechanisms. Section 5 examines how housing frictions and wealth shape the aggregate and distributional effects of local downturns. Section 6 concludes.

2 Empirical Evidence

This section presents evidence that moving propensity decreases with an individual’s ability to access financial markets after controlling for other individual characteristics, particularly age and homeownership, two critical determinants of an individual’s wealth that previous literature has shown to be relevant for individual moving decisions (e.g., [Molloy, Smith and Wozniak, 2014](#)). In and by itself, this evidence does not constitute a new contribution as patterns of migration by these groups or similar dimensions are known to the literature in different countries. Yet, it

reinforces existing evidence on the interaction of migration with wealth and credit, validates our data, and acts as an important input for the quantitative exercise in the subsequent sections.

2.1 Data Description

Our data source is *TransUnion Canada*, one of the two credit reporting agencies in Canada. It collects the individual credit history of about 35 million individuals, covering nearly every person in Canada with a credit report. It is a monthly longitudinal panel of individuals available since 2009⁴ that includes information on borrowers' characteristics such as age, credit score and liabilities. Specifically, we observe credit limits, balances, payments, and delinquency status for different credit accounts such as mortgages, auto loans, credit cards, and lines of credit. Although homeownership status is not directly observed, we infer that an individual is a homeowner if she has a mortgage account with a positive outstanding balance or if a fully amortized mortgage is associated with the current individual's residence. Crucial to our analysis, the data tracks an individual's residence over time, particularly the Forward Sortation Area (FSA) that corresponds to the first three digits of the individual's postal code.⁵ Figure A.1 of Appendix A plots migration rates across different geographic units within Canada and corroborates that whenever the geographic units are comparable, *TransUnion* tracks well the migration rates across Canada. Moreover, *TransUnion* allows us to compute migration rates across Canadian cities at a higher frequency and for different demographic groups, which is impossible using official migration statistics. The green dashed line presents the migration rate between census metropolitan areas (CMA)^{6,7}. Between 2011 and 2019, on average, 1.54 percent of the Canadian population aged between 25 and 85 moved between CMAs per year.

⁴Our analysis starts in 2011 given the limited data coverage before that year.

⁵We restrict our sample to individuals between 25 and 85 years old. Individuals below 25 years old are underrepresented in our data due to the lack of credit history. Anecdotal evidence suggests that young individuals keep the addresses of their relatives as their official residence during school years. Avoiding miscalculating migration rates among post-secondary educated individuals is another reason not to include this demographic group in our sample. We exclude individuals above 85 years old due to the possibility of unreported deaths and to prevent capturing people's movements for nursing homes or similar facilities. The results are not sensitive to this restriction.

⁶A city is defined as a census metropolitan area (CMA) or a census agglomeration (CA) that is formed by one or more adjacent municipalities centered on a population core. A CMA must have a population of at least 100,000, of which 50,000 or more must live in the core. A CA must have a core population of at least 10,000. To be included in the CMA or CA, other adjacent municipalities must integrate with the core, measured by commuting flows derived from previous census place of work data.

⁷Migration rate across CMAs is defined by the number of people reported to change the address to a different CMA divided by the total number of individuals living in a CMA in the previous year.

2.2 Migration Patterns by Demographic Groups

We now document how migration rates vary for different demographic groups. We are particularly interested in analyzing how moving decisions are impacted by financial constraints, notably, individual’s ability to access financial markets. We use two measures of financial constraints, credit score and credit usage. Financial institutions widely use credit scores to determine an individual’s creditworthiness and for loan underwriting and pricing. On average, borrowers with higher credit scores tend to have easier access to credit and more favorable loan terms (Beer and Li, 2018).⁸

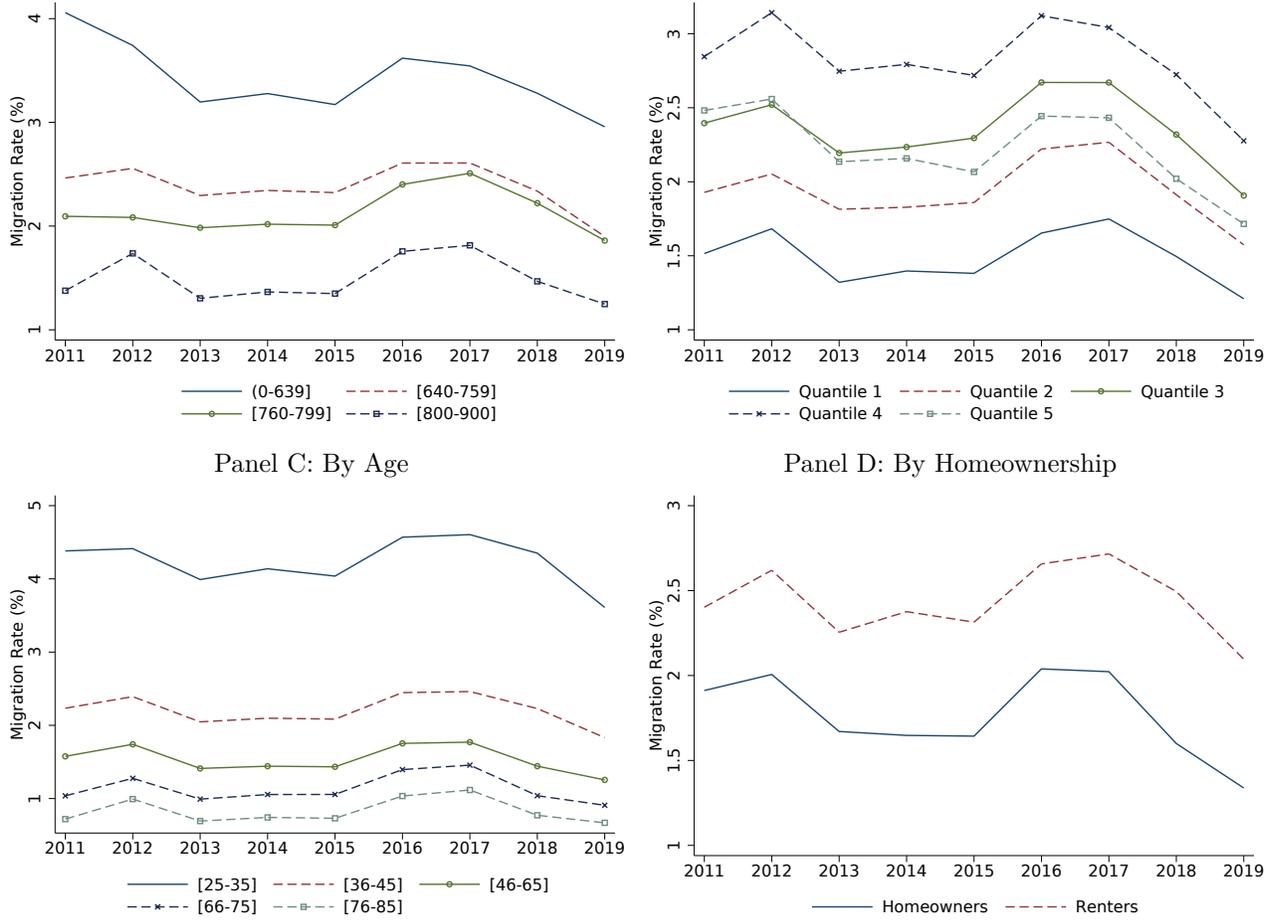
Figure 1 plots annual migration rates across Canadian cities by individual characteristics. In Panel A, we partition the sample into 4 groups of credit score, 0-639, 640-759, 760-799, and 800-900, commonly designated, respectively, as very poor, near prime and prime, prime plus and super prime. In Panel B, we split the sample into quintiles of credit usage. Panel C and D, migration rates are computed by age groups (25-35, 36-45, 46-55, 66-75, 76-85) and homeownership status (renter vs homeowner), respectively.

Panel A and Panel B present evidence of differential migration rates across measures of credit access. According to both measures, more constrained individuals tend to move more frequently than less constrained individuals. Specifically, Panel A shows that migration rate decreases monotonically with credit score and Panel B shows that individuals with higher credit usage rates (more constraint) also move more on average. Panel C shows a monotonically decreasing relationship between age and migration flows for individuals between 25 and 85 years of age. Specifically, individuals between 25 and 35 move, on average, roughly twice as much as people between 36 and 45 and more than four times that of individuals above 65 years old. Panel D shows the difference in migration rates between homeowners and renters. Renters, on average, are 25 percent more likely to move than homeowners. These two last results are consistent with findings for migration flows across US states as in Molloy, Smith and Wozniak (2014), which reinforces the validity of our data in the study of migration both at the aggregate level and by demographic groups.

Regression Framework To account for the correlations between mobility and individual demographics, we formally assess how moving decisions vary with observable characteristics

⁸Credit usage is the total outstanding non-mortgage debt balance divided by the credit limit. We consider any open credit account of credit cards, installments, auto-loans and lines of credit. We abstract from mortgage debt to capture sources of credit that individuals can easily adjust, potentially in response to unexpected life events. We view high credit usage as a proxy for higher financial constraints as it is harder for individuals to increase their debt in the short run if their outstanding debt is already close to the limit.

Figure 1: Migration Rates across Canadian Cities by Demographic Groups
 Panel A: By Credit Score
 Panel B: By Credit Usage



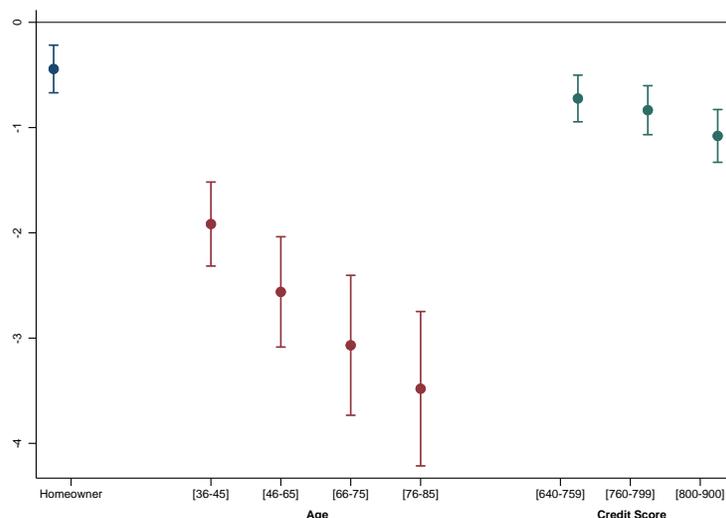
Note: Figure 1 plots the migration rates between Canadian cities (CAs) between 2001 and 2019 by credit score (Panel A), credit usage (Panel B), age (Panel C) and homeownership status (Panel D). The migration rate is defined by the number of people moving across cities divided by the total population in the same set of cities in the year before. *Data Source:* TransUnion.

using the following linear probability specification:

$$\mathbb{1}[Move_{i,z,z',t}] = \beta_0 + \beta_1 X_{i,t-1} + \beta_2 W_{i,t-1} + \delta_{z,t} + \epsilon_{i,z,t} \quad (1)$$

where $\mathbb{1}[Move_{i,z,z',t}]$ is an indicator variable taking the value of 100 if individual i in city z at time t moves to a different city $z' \neq z$, and 0 otherwise ($z' = z$); thus, the coefficients are in units of percentage points. $X_{i,t-1}$ are individual characteristics such as age, homeownership and credit score. We also control for other time-varying characteristics $W_{i,t-1}$ such as credit usage, home equity and delinquencies. Our preferred empirical specification includes city-by-quarter fixed effects to control for changes in local economic conditions. It also controls for trends in migration patterns and for city characteristics such as amenities, long-run productivity levels, and quality of life, among others. In other words, this specification allows comparing individuals

Figure 2: Determinants of Migration Decisions



Note: Figure 2 reports the point estimates of the linear probability model 1. The vertical bands represent 99 percent confidence intervals for the point estimates in each quarter. The first coefficient (in blue) reports the point estimate for the homeowner indicator variable. The second set of estimates under the umbrella “Age” (in red) reports the estimates for age group indicator variables using the 25-35 years old group as the baseline. The third set of estimates under the umbrella “Credit Score” (in green) reports the estimates by credit score group indicator variables relative to the credit score group [0-640]. *Data source:* TransUnion.

within a city in a given period. We cluster standard errors at the city level.

Panel A of Table A.1 in Appendix A reports the results of estimating the specification of equation (1) for individuals that live in CAs for different combinations of controls and fixed effects. Results hold across specifications. Figure 2 reports the estimates of our preferred specification (column 8 of Panel A of Table A.1) with the vertical bands representing 99 percent confidence intervals for the point estimates. The first coefficient in blue reports the estimate for homeowners. After controlling for other individual characteristics, such as home-equity and delinquency status, we find that within a given city and a quarter, homeowners are, on average, 0.58 percentage points less likely to move than renters. The coefficients in red under the “Age” group report the estimates of the relative moving propensity of the different age groups relative to the youth group (individuals with ages between 25 and 35 years old). The likelihood of moving decreases monotonically with age. Individuals between 36 and 45 years old are 1.92 percentage points less likely to move than those between 25 and 35. This difference increases monotonically up to 3.5 percentage points in the group of 76-85 years old.

The last set of coefficients relates to the “Credit Score” group. The probability of moving for individuals with credit scores between 640 and 759 is 0.69 percentage points smaller than the moving probability of those with score below 640, and this difference monotonically increases as credit score goes up. Individuals at the top of the credit score distribution are 1.06 percentage

points less likely to move than those at the bottom of the distribution. Overall, we conclude that after controlling for age, and homeownership, among other individual characteristics, financial constraints are a relevant factor for moving decisions. More financially constrained individuals (lower credit scores) are more likely to move than those less financially constrained in the same city in a given quarter. These results also highlight that age and homeownership alone are not good proxies for the role of credit access in migration decisions. These results contribute to the literature studying population flows in developed countries that mostly use standard surveys or census data that lack information on individuals’ finances and credit scores. In this sense, our findings complement the results in [Bilal and Rossi-Hansberg \(2021\)](#) that analyzes migration flows in France.⁹

We produce a battery of robustness checks in Appendix A. Panel B of Table A.1 replicates the analysis in Panel A of the same table but restricts the sample to CMAs, the 35 largest Canadian cities. Panels A and B of Table A.2 use credit usage (excluding mortgage debt) as a proxy for financial constraints rather than credit score, with the latter panel restricting the sample to CMAs. Table A.2 uses credit usage including mortgage debt instead. It shows that individuals with higher credit usage (more financially constrained) are more likely to move than those with lower credit usage. Due to data limitations, we calibrate the model developed in the next Section to the Canadian CMAs, so this exercise makes the data and the model comparison closer. We validate our findings also using data for the US from the National Longitudinal Survey of Youth in Appendix B. We confirm the same findings for the US, even accounting for education, income and other demographic variables.

3 A Quantitative Heterogeneous Agents Spatial Equilibrium Model with Wealth

Motivated by the empirical evidence above, we develop a quantitative spatial equilibrium model with life cycle dynamics, uninsurable income risk and wealth accumulation through illiquid housing and a liquid asset. After observing an idiosyncratic location preference shock, households jointly choose their consumption-saving profile and where to live, subject to both *monetary* and *utility* moving costs. This framework provides realistic and strong insights into the role

⁹They find that households at the bottom quintile of the financial income distribution tend to move more than their counterparts at the top of the distribution. However, our data allows us to directly identify measures of an individual’s access to credit, a crucial element that determines an individual’s ability to smooth shocks, particularly for low-wealth individuals. Moreover, we have access to credit scores that are a slow-adjusting object, relaxing some of the identification concerns related to the fact that moving and savings are simultaneously chosen.

of income risk, wealth and financial frictions on moving decisions. It highlights *precautionary moving* as an alternative insurance channel against income shocks, one that is particularly relevant for households at the bottom of the wealth distribution.

3.1 Space

The economy is defined by L locations indexed by $l = \{1, 2, \dots, L\}$. Locations differ in four dimensions: productivity, amenities, housing supply elasticities and labor market risk. Location subscripts are omitted unless necessary.

3.2 Household Environment

Demographics Time is discrete. The economy is populated by a measure-one continuum of finitely-lived households. Age is indexed by $q = \{1, 2, \dots, Q\}$. Households live for at most Q periods, but face mortality risk with age-dependent survival probabilities $\{\lambda_q\}$. Households work in the initial \bar{Q} periods and retire after that. Newborns are distributed across locations according to $G(l)$.

Preferences Households value non-durable consumption c , housing services s and location-age specific amenities A_q^l . The location-age specific instantaneous utility function u_q^l is given by:

$$u_q^l(c, s, A_q^l) = \frac{e_q[(1 - \alpha)c^{1-\gamma} + \alpha s^{1-\gamma}]^{\frac{1-\sigma}{1-\gamma}} - 1}{1 - \sigma} + A_q^l, \quad (2)$$

where α measures the relative taste for housing services, $\frac{1}{\gamma}$ measures the elasticity of substitution between housing services and non-durable consumption, and $\frac{1}{\sigma}$ measures the intertemporal elasticity of substitution. Non-durable consumption is the numeraire good. The exogenous equivalence scale $\{e_q\}$ captures deterministic changes in household size and composition over the life cycle. Households discount future periods at rate $\beta > 0$ and leave bequests to future generations when they die. These are captured by a warm-glow bequest motive a la [De Nardi \(2004\)](#):

$$\varphi(a) = \bar{\varphi} \frac{(a + \underline{a})^{1-\sigma} - 1}{1 - \sigma} \quad (3)$$

where $\bar{\varphi}$ captures the intensity of the bequest motive and \underline{a} determines the curvature of the bequest function and hence the extent to which bequests are a luxury good.

Endowments Households receive an income endowment $y_{i,q}^l$ given by

$$\log y_{i,q}^l = \begin{cases} \log w^l + \chi_q + \epsilon_i & \text{if } q \leq \bar{Q}, \epsilon_i \neq \epsilon^1, \\ \log v w^l + \chi_q^l + \epsilon_i & \text{if } q \leq \bar{Q}, \epsilon_i = \epsilon^1, \\ \log w_{ret} + \chi_q^l & \text{if } q > \bar{Q} \end{cases}$$

The income process for working-age households ($q \leq \bar{Q}$) has three components. The first is a location-specific term related to the local wage w^l that is endogenously determined and reflects local productivity. The remaining two components capture individual labor productivity: a deterministic age profile, χ_q^l , which generates the standard hump-shaped pattern of average earnings over the life cycle, and an idiosyncratic component, ϵ_i , which is stochastic and persistent. Following [Guerrieri and Lorenzoni \(2017\)](#), the idiosyncratic productivity component takes values in $\{\epsilon^1, \dots, \epsilon^E\}$, where ϵ^1 represents the unemployment state and is set to zero. Unemployed agents receive a fraction $v < 1$ of the equilibrium wage in their current location, common across cities, consistent with the unemployment insurance system in Canada.

The idiosyncratic component ϵ_i evolves according to a first-order Markov chain, Π^l . Transitions across employment states (ϵ^e , $e > 1$) are common across locations, but transitions into unemployment (ϵ^1) vary by location, making Π^l location-specific. Initial employment status is drawn from the stationary distribution π^l . Households do not observe their idiosyncratic income draw in a prospective location before migrating. Instead, movers draw from the stationary distribution of income shocks in the destination, while stayers draw from the Markov process in their current city. This creates a selection effect: households in high-income states tend to remain, while those in low-income states are more likely to migrate. This mechanism mirrors the ‘match-quality’ logic in [Kennan and Walker \(2011\)](#) and is embedded in the model’s migration elasticities.

Upon retirement, households receive a retirement income w_{ret} , common across locations, and the deterministic age profile component. Households are born with an endowment of wealth that is drawn from a location-specific exogenous distribution, which is correlated with initial income.

Housing Housing services can be obtained by renting ($d = 0$) or owning ($d = 1$). Households have a preference for homeownership: owning a house of size h provides $s = \omega h$ units of effective housing services with $\omega \geq 1$, while renting yields $s = h$. Owner-occupied and rental housing sizes belong to two location-specific finite sets, $\mathcal{H}^{H,l}$ and $\mathcal{H}^{R,l}$, respectively. Rental units are slightly smaller than owner-occupied houses. There is an implicit minimum location-dependent house size requirement capturing the notion that households must consume at least a basic level

of shelter, as is often mandated by city regulations.

A household in location l pays $R^l h$ per period to rent a house of size h and $p^l h$ to purchase a house of the same size. Ownership carries a per-period maintenance cost of $\delta p^l h$, which fully offsets physical depreciation and a property tax of $\psi_h p^l h$. Buying a house involves a proportional transaction cost $F p^l h$. Renters can adjust housing costlessly. When homeowners move, they must sell their property and decide subsequently whether to rent or own in the new location.

Liquid Asset and Wealth Agents can borrow or save through a one-period financial asset b in the international financial market. Positive savings have a fixed exogenous return r . Borrowing is allowed at a fixed exogenous cost $r + \iota$, with $\iota \geq 0$. For simplicity, we define $r^b = r\mathbb{1}[b \geq 0] + (r + \iota)\mathbb{1}[b < 0]$.

Renters face a limit to unsecured borrowing of $\underline{b} \leq 0$. Homeowners can use their housing as collateral but borrowing cannot exceed $\underline{b} + \xi p^l h$. The borrowing constraint is summarized by:

$$b' \geq \underline{b} + \mathbb{1}[d = 1]\xi p^l h, \quad \underline{b} \leq 0 \quad (4)$$

Wealth a is the sum of household's financial assets b and, for homeowners, the market value of their housing $p^l h$:

$$a = b(1 + r) + \mathbb{1}[d = 1]p^l h$$

Location Choice Each period, households draw a vector of L independent Type 1 Extreme Value location preference shocks with a mean zero scale parameter ν and decide where to live. If households decide to move, they incur in a *monetary* moving cost F_m and an *utility* moving cost which depends on the distance between the origin and destination locations. Specifically, *utility* cost of moving from location l to k , $\tau^{l,k}$, is given by:

$$\tau^{l,k} = \tau_0 + \tau_1 D_{l,k} \quad (5)$$

where $D_{l,k}$ is distance between locations. We depart from most of the literature by assuming *homogeneous* moving costs, *i.e.*, moving costs do not depend on households' characteristics such as age or homeownership status. However, as shown later, in the presence of income risk and borrowing constraints, the benefits of moving depend on individual states, endogenously generating heterogeneous migration patterns consistent with the data.

Government The government revenues come from a progressive labor income tax schedule captured by the function $\Psi(\cdot)$ and a proportional property tax ψ_h levied on house values. Revenues also include proceeds from the sale of land permits which are described in Section 3.4.

Expenditures consist of unemployment benefit, pensions for retirees and a public good that is not valued by households.

Timing At the beginning of the period, idiosyncratic location preference shocks are realized and location choices are made. If moving occurs, moving costs are paid, and homeowners sell their property. The idiosyncratic income shock is then observed, after which households choose whether to rent or own and determine consumption, housing services, and financial savings subject to the borrowing constraint. Homeowners pay transaction costs when adjusting housing or when moving, as well as maintenance costs and property taxes. At the end of the period, the death shock is realized and deceased households leave bequests.

3.3 Households' Decisions

State Variables Let X_t be the vector which includes wages w_t^l , housing prices p_t^l , rental prices R_t^l and previous housing stock H_{t-1}^l for each location l .^{10,11} The household's individual state variables is denoted by x_t and includes individual wealth at the beginning of the period a_t , the idiosyncratic productivity shock ϵ_t , age q and the variable \bar{h}_t that incorporates the housing tenure status (d_{t-1}), housing size (h_{t-1}) and location in the previous period. \bar{h}_t equals housing consumption in the previous period, h_{t-1} , if the household was a homeowner that did not move ($d_{t-1} = 1 \wedge l_t = l_{t-1}$) and zero if previously a renter ($d_{t-1} = 0$) or a homeowner that moved ($d_{t-1} = 1 \wedge l_t \neq l_{t-1}$). In a compact way, \bar{h}_t is given by:

$$\bar{h}_t = h_{t-1} \mathbb{1}[d_{t-1} = 1 \wedge l_{t-1} = l_t] \quad (6)$$

Location choice At the beginning of each period, after drawing an idiosyncratic preference shock for each location $\tilde{\epsilon}^l$, but before observing their productivity ϵ , households form beliefs about their productivity shock and aggregate variables (wages, housing prices and rental prices) across different locations and choose the location with the highest value:

$$\tilde{V}_t^l(a_{t-1}^l, \epsilon_{t-1}, q_t, \bar{h}_{t-1}, l_{t-1}) = \max_{k \in \{1, \dots, L\}} \mathbb{E}_t \left\{ V_t^k(a_t^k, \epsilon_t, q_t, \bar{h}_t^k) - \tau^{l,k} + \nu \tilde{\epsilon}_t^{z_i, k} \middle| \tilde{\epsilon}_t^{z_i, k} \right\} \quad (7)$$

Households that move from l to k must pay the *utility* moving costs $\tau^{l,k}$ and the *monetary* moving cost F_m that will impact a_t^k as explicitly below.

¹⁰Alternatively, we could define the aggregate state variable as the distribution of households over age groups, housing tenure, wealth and locations.

¹¹For compactness, value functions are indexed by subscript t to reflect potential changes in the aggregate state.

Tenure choice After location choices are made, the individual productivity in the chosen location is realized and households decide between being a renter ($d = 0$) and a homeowner ($d = 1$) by solving:

$$V_t^l(a_t, \epsilon_t, q_t, \bar{h}_t) = \max \left\{ V_t^{R,l}(a_t, \epsilon_t, q_t, \bar{h}_t), V_t^{H,l}(a_t, \epsilon_t, q, \bar{h}_t) \right\}, \quad (8)$$

where V_t denotes the value function at the beginning of the period in location l , $V_t^{R,l}$ the value of renting and $V_t^{H,l}$ the value of owning. The decision of renting versus owning is based on the comparison of the respective value functions.

Renters' Problem At time t , households of age q in location l with wealth a and income shock ϵ and decide to rent choose how much to consume of non-durable good c , rental units h and liquid savings b that solves:

$$\begin{aligned} V_t^{R,l}(a_t, \epsilon_t, q_t, \bar{h}_t) = & \max_{c_t, h_t, b_t} u_q(c_t, s_t, A_q^l) + (1 - \lambda_q) \varphi(a_{t+1}^l) \\ & + \lambda_q \beta \tilde{\mathbb{E}}_t \left\{ \max_{k \in \{1, \dots, L\}} V_{t+1}^k(a_{t+1}^k, \epsilon_{t+1}, q_t + 1, \bar{h}_{t+1}^k) - \tau^{l,k} + \nu \tilde{\epsilon}_{t+1}^k \right\} \quad (9) \\ \text{s.t. } & c_t + R_t^l h_t + b_t = y_t^{\epsilon, l} + a_t - \Psi(y_t^{\epsilon, l}) \\ & b_t \geq \underline{b} \\ & a_{t+1}^k = (1 + r^b) b_t - F_m \mathbb{1}[l \neq k] \\ & s_t = h_t \in \mathcal{H}^{R,l}, \quad \bar{h}_{t+1}^k = 0 \end{aligned}$$

$y^{\epsilon, l}$ is the pre-tax income consistent with the realization of productivity shock ϵ for working-age workers. As specified in equation (3.2), income depends on the location and age. $\Psi(y^{\epsilon, l})$ is total taxes on labor income. The renter must pay R_t^l per rental unit and savings can be negative but subject to the borrowing constraint. The continuation value has two components. With probability $1 - \lambda_q$ the household dies and leaves bequests. With probability λ_q the household survives and chooses the new location. The last term corresponds then to the expected value of the optimal location choice. $\tilde{\mathbb{E}}_t$ corresponds to expectation over idiosyncratic location preference shocks, idiosyncratic productivity shocks and aggregate state across all locations.¹²

Next period's wealth state a depends on households moving decisions. If renter decides to remain in the same location, no moving costs occur and the next period wealth consists of $(1 + r^b)b_t$. If renters move, monetary moving costs occur and next period's wealth consists of

¹²The expected value of the optimal location choice corresponds to the expectation of value function (7) over the idiosyncratic location preference shocks $\tilde{\epsilon}^l$. Therefore, $\tilde{\mathbb{E}}_t$ differs from \mathbb{E}_t as $\tilde{\mathbb{E}}_t$ includes expectations over the idiosyncratic location preference shocks, while \mathbb{E}_t does not.

$(1 + r^b)b_t - F_m \cdot \bar{h}_{t+1}$ is zero, regardless if renters moves or not.

Homeowners' Problem At time t , a household of age q in location l with wealth a and productivity shock ϵ and decides to own their housing chooses how much to consume of non-durable good c , owned housing units h and liquid savings b that solves:

$$\begin{aligned}
V_t^{H,l}(a_t, \epsilon_t, q_t, \bar{h}_t) &= \max_{c_t, h_t, b_t} u_q(c_t, s_t, A_q^l) + (1 - \lambda_q)\varphi(a_{t+1}^l) \\
&\quad + \lambda_q \beta \tilde{\mathbb{E}}_t \left\{ \max_{k \in \{1, \dots, L\}} V_{t+1}^k(a_{t+1}^k, \epsilon_{t+1}, q_t + 1, \bar{h}_{t+1}^k) - \tau^{l,k} + \nu \tilde{\epsilon}_{t+1}^k \right\} \quad (10) \\
\text{s.t. } c_t + b_t + p_t^l h_t (1 + F \mathbb{1}[h_t \neq \bar{h}_t]) &= y_t^{\epsilon, l} + a_t - \Psi(y_t^{\epsilon, l}) \\
b_t &\geq \underline{b} - \xi p_t^l h_t \\
a_{t+1}^k &= (1 + r^b)b_t - F_m \mathbb{1}[l \neq k] + p_{t+1}^l h_t (1 - \delta_h - \psi_h^l) \\
s_t &= \omega h_t, \quad h_t \in \mathcal{H}^{H,l}, \quad \bar{h}_{t+1}^k = h_t \mathbb{1}[k = l]
\end{aligned}$$

Homeowners' problem differs from renters' problem in three main dimensions. First, homeowners can partially finance their house purchases subject to the collateral constraint $\xi p_t^l h_t$. For simplicity, mortgages are negative short-term safe assets as in [Favilukis, Mabile and Van Nieuwerburgh \(2023\)](#).¹³ Second, homeowners are subject to property taxes ψ_h^l and maintenance costs δ_h per unit of housing value that are paid at time $t + 1$ as in [Favilukis, Mabile and Van Nieuwerburgh \(2023\)](#) and [Kaplan, Mitman and Violante \(2020b\)](#).¹⁴ Third, houses are illiquid assets as households face transaction costs F when buying a new house ($h_t \neq \bar{h}_t$).

As in the renter's problem, next period's wealth also depends on the moving decisions but besides savings and potential moving costs, next period wealth also includes the value of current property at $t + 1$ net of property taxes and maintenance costs.¹⁵ \bar{h}_{t+1} equals housing size h_t if

¹³Consistent with Canadian tax code, we assume that mortgage interest are not tax deductible.

¹⁴As shown in equation (17) below, the physical depreciation is offset by residential investment undertaken by the construction sector. We allow property taxes to vary across locations to match the heterogeneous rental-price ratios across Canadian cities.

¹⁵As in [Kaplan, Mitman and Violante \(2020a\)](#), the timing assumption ensures that a household can sell and buy a new home within the period. Note also that the transaction cost F is not paid by homeowners that remain in the same location with the same housing units ($h_t = \bar{h}_t = h_{t-1}$). These two assumptions, together with the inclusion of the individual state variable \bar{h}_t allow us to write the homeowner problem in a compact manner that implicitly assumes that homeowners trade houses every period even if they remain in the same house. However, if homeowners choose to remain in the same location and consume the same amount of housing services, which corresponds to remain in the same house, then the terms in the budget constraint related to housing (after substituting a_{t+1} condition) collapses to $p_t^l(h_t - h_{t-1} + F \mathbb{1}[\bar{h}_t \neq h_{t-1}]) = 0$. Therefore, if a household does not move and does not adjust its own-housing consumption, it does not obtain any capital gains/losses in case house price changes (i.e., its liquid wealth does not vary with house price movements), even if the household problem is written as homeowners would sell and buy the property every period. Only when homeowners decide to sell the property and move to other location ($\bar{h}_t = 0$) or remain in the same location but consumer a different housing

homeowner remains in the same location and zero if moves.

We can express the law of motion for wealth for both renters and homeowners in the following compact way:

$$a_{t+1}^k = (1 + r^b)b_t - F_m \mathbb{1}[l \neq k] + p_{t+1}^l h_t (1 - \delta_h - \psi_h^l) d_t \quad (11)$$

Migration Given that the idiosyncratic location preference shocks are *i.i.d.* over time and distributed Type-I Extreme Value with zero mean a scale parameter ν , the continuation value in case of survival in equations (9) and (10) can be rewritten as:

$$\lambda_q \nu \log \left(\sum_{k=1}^L \exp \left(\beta \mathbb{E}_t V_{t+1}^k(a_{t+1}^k, \epsilon_{t+1}, q_t + 1, \bar{h}_{t+1}^k) - \beta \tau^{l,k} \right)^{\frac{1}{\nu}} \right). \quad (12)$$

As shown by [McFadden \(1973\)](#), this assumption also implies a closed-form analytical expression for the share of movers across locations. $\mu_t^{l,k}$ denotes the share of households with a given individual state that choose to move from location l to location k and it is given by:

$$\mu_t^{l,k}(a_{t+1}^k, \epsilon_t, q_t, \bar{h}_{t+1}^k) = \frac{\exp \left(\beta \mathbb{E}_t V_{t+1}^k(a_{t+1}^k, \epsilon_{t+1}, q_t + 1, \bar{h}_{t+1}^k) - \beta \tau^{l,k} \right)^{\frac{1}{\nu}}}{\sum_{k=1}^L \exp \left(\beta \mathbb{E}_t V_{t+1}^k(a_{t+1}^k, \epsilon_{t+1}, q_t + 1, \bar{h}_{t+1}^k) - \beta \tau^{l,k} \right)^{\frac{1}{\nu}}} \quad (13)$$

where V_{t+1}^k , a_{t+1}^k and \bar{h}_{t+1}^k are defined in (8), (11) and (6), respectively. a_{t+1}^k , d_t and \bar{h}_{t+1} are optimal savings, housing tenure and housing consumption choices derived from agents' optimization problems (9) and (10).

3.4 Production

There are two production sectors in each location: a tradable good sector which produces non-durable consumption and a construction sector which produces new houses. Productivities are location-specific and labor, supplied inelastically, is perfectly mobile across sectors within the location.

Final Good Sector Each location produces a uniform final good that can be traded across locations in a perfect competition setting. Location-specific productivity has two components: (i) an exogenous component denoted by z^l and (ii) an endogenous agglomeration force governed by parameter ζ that depends on the city's population \bar{N}^l . A representative final good producer amount ($\bar{h}_t = h_{t-1}$) any potential gains or losses are realized.

in location l operates the following technology:

$$Y^l = z^l (N_c^l)^\eta (\bar{N}^l)^\zeta$$

where N_c^l is the effective employment¹⁶ in the final good sector in location l .¹⁷ Because firms in the final-good sector are atomistic, they treat the agglomeration externality $(\bar{N}^l)^\zeta$ as exogenous. This implies that the equilibrium city-level wage in location l is then given by

$$w^l = \eta z^l (N_c^l)^{\eta-1} (\bar{N}^l)^\zeta \quad (14)$$

Construction Sector As in [Kaplan, Mitman and Violante \(2020b\)](#), there is a foreign-owned competitive construction sector that operates in each location with the following production technology:

$$I^l = (z^l N_h^l)^{k^l} (\bar{L}^l)^{1-k^l}$$

where N_h^l is the effective labor employed in the construction sector and \bar{L}^l is the amount of newly available buildable land.¹⁸ The developer solves a static profit maximization problem in a competitive market:

$$\max_{N_h^l} p^l I^l - w^l N_h^l \quad s.t. \quad I^l = (z^l N_h^l)^{k^l} (\bar{L}^l)^{1-k^l} \quad (15)$$

The solution to the developer problem after choosing optimal level of labor allows us to rewrite the investment level I^l as a function of wage w^l and housing price p^l as:

$$I^l = \left(\frac{\kappa^l p^l z^l}{w^l} \right)^{\frac{\kappa^l}{1-\kappa^l}} \bar{L}^l \quad (16)$$

where w^l is given by equation (14) due to free labor mobility across sectors within locations. The housing supply elasticity is given by $\frac{\kappa^l}{1-\kappa^l}$. The overall housing stock in location l evolves

¹⁶As formally defined in Appendix C, city size is defined by the number of households in a given location given by $\bar{N}^l = \sum_{j=1}^L \int_x \mu^{j,l}(x) \bar{N}^j \lambda^j(x) dx + N_0^l$, where N_0^l denote newborns in location l and $\lambda^j(x)$ is the distribution of individual states in location j in the stationary equilibrium. Supply of effective employment reflects both the age composition of workers and their idiosyncratic productivities ϵ , and is formulated in Appendix C.

¹⁷For simplicity, profits are fully taxed by the government.

¹⁸Government issues and sells new permits equivalent to \bar{L}^l units of land to developers in a competitive market as assumed in [Kaplan, Mitman and Violante \(2020b\)](#) and [Favilukis, Mabile and Van Nieuwerburgh \(2023\)](#). This implies that all rents from land ownership accrue to the government and the construction sector makes no profits in equilibrium.

according to

$$H_t^l = (1 - \delta)H_{t-1}^l + I_t^l. \quad (17)$$

Rental Sector Following [Kaplan, Mitman and Violante \(2020b\)](#), we assume that risk-neutral foreign investors can arbitrage between the owned-housing market and the rental market, which connects housing prices and rents in the following way:¹⁹

$$R_t^l = p_t^l - \left(1 - \delta_h - \psi_h^l\right) \frac{\mathbb{E}_t p_{t+1}^l}{1 + r} \quad (18)$$

This non-arbitrage condition assumes that rental sector incurs depreciation and property taxes and discount future profits at the risk-free rate r .

3.5 Equilibrium

Given the set of parameters and the exogenous interest rate r , a competitive equilibrium is a location-specific price vector $\{w_t^l, p_t^l, R_t^l\}_{l=1}^L$ and allocations, namely, housing stock and population (labor supply) consistent with the households and firms optimization and that clear the markets in each location. A stationary equilibrium is one in which all equilibrium objects are time-invariant. A formal equilibrium definition is provided in [Appendix C](#).

3.6 Solving the Model

Rich individual and spatial heterogeneity, together with dynamic consumption–saving under income risk and incomplete markets, generates a very large state space. Credit and housing frictions make the household problem non-convex, and solving counterfactuals requires computing full transition dynamics. These features make the model computationally demanding. Nonetheless, we show that state-of-the-art tools from macro and spatial economics can be combined in a parsimonious way to solve a dynamic spatial equilibrium model with heterogeneous agents, uninsurable income risk, and wealth accumulation in liquid and illiquid assets.

Numerical Computation of Stationary Equilibrium and Transition Dynamics Spatial models with forward-looking migration often rely on dynamic hat-algebra ([Caliendo, Dvorkin and Parro, 2019](#)), its extensions to capital accumulation ([Kleinman, Liu and Redding, 2023](#)), or

¹⁹As presented in [Kaplan, Mitman and Violante \(2020b\)](#), this ‘Jorgensonian’ user-cost formula can be derived from the optimization problem of a competitive rental market that can frictionlessly buy and sell housing units and rents them to households.

log-linear value functions (Dvorkin, 2023). These approaches are not applicable here: borrowing constraints and non-log-linear value functions generate highly non-linear policy rules, preventing closed-form aggregation and ruling out hat-algebra or spectral methods. Instead, we combine a Type-1 Extreme Value structure for location preferences—which yields closed-form migration shares—with global solution methods standard in heterogeneous-agent macroeconomics. The finite life horizon eliminates the need for costly value-function iteration: the terminal value is known in closed form, and value and policy functions are obtained by backward induction. Appendix D provides a full description of the algorithm used to compute both the stationary equilibrium and the transition dynamics.

4 Taking the Model to the Data

In this section, we report how we take the model to the data and that it successfully matches key untargeted moments of the migration and wealth data. We highlight that our model generates heterogeneous moving rates quantitatively consistent with the data, even with homogeneous moving costs. We finally explore the model’s main mechanisms, focusing on the forces that drive migration decisions. Through a decomposition exercise, we show that income risk and financial constraints significantly increase migration propensity, especially for low-wealth individuals. We parameterize the model to match key cross-sectional features of the Canadian economy between 2011 and 2013, as well as the observed out-migration from oil-producing regions following the large and persistent decline in global oil prices between 2014 and 2018. Accordingly, our parameterization strategy requires solving not only for a stationary equilibrium but also for the transition dynamics triggered by a negative local shock. We take the quantitative model to the largest 27 largest CMAs in Canada, which constitute our definition of “cities” throughout the paper.²⁰

Migration elasticity, moving costs, and location-specific amenities are key parameters governing heterogeneous migration decisions. Several studies have estimated these variables in spatial equilibrium models with dynamic location choice but no wealth. In our framework, however, forward-looking agents jointly decide on location and savings, making existing estimates potentially unsuitable.

Nevertheless, estimating migration elasticity in our framework faces challenges similar to those in the existing literature, since location choice, wages, and house prices are jointly

²⁰In 2016, there were 35 cities in Canada with more than 100,000 inhabitants (CMAs), but data limitations constrain our analysis to 27 of them.

determined and may respond to unobserved changes in amenities. Given the heterogeneity and dynamic structure of our model, applying the most recent estimation strategies from the literature is not straightforward.²¹ First, our model features rich heterogeneity, and migration flows at this level of disaggregation are unobserved. Second, linear aggregation to a level with observed migration flows is not possible due to the nonlinearity of equation (13). Third, because agents jointly choose consumption and savings, flow utility is not simply a function of wages and rents, as is typically assumed, but depends on income and wealth. Estimating this equation would require solving for the full policy function for each parameter vector. Instead, we exploit the sharp decline in global oil prices between 2014 and 2018 as an exogenous source of income variation orthogonal to unobserved changes in amenities, and that exogenous variation drove up migration flows. In our procedure, we jointly find the dispersion of idiosyncratic local taste shocks that governs migration elasticity, moving costs, and local amenities, on top of other parameters, such that we match population distribution by age across cities and migration heterogeneity across demographic groups prior to the shock as well as the migration response in cities highly exposed to the oil industry and exogenous changes in oil prices.

4.1 Data Sources

We draw on several data sources described below, with most targeted moments coming from five core datasets and city-level exposure to the global oil price collapse.

TransUnion Canada As in Section 2, we use TransUnion data to compute migration flows across Canadian cities for different demographic groups. To align with the model, we compute migration moments for the subset of cities included in the model and restrict the sample to individuals aged 25 to 85. Migration moments in the initial stationary equilibrium match the average rates between 2011 and 2013.

Labour Force Survey (LFS). We use confidential LFS microdata,²² a monthly survey of individuals aged 15 and older. It provides information on employment status, hours worked, wages, occupation, industry, age, and place of residence. We use the 2011–2013 LFS waves to compute moments related to employment status, wages, and income by city and age, restricting the sample to individuals aged 25 to 64.

²¹For example, [Cai et al. \(2022\)](#) and [Ahlfeldt et al. \(2015\)](#), among others, jointly estimate moving costs, amenities, and migration elasticities using the gravity equation (13) at the skill-location level through pseudo-maximum-likelihood (PML) estimation with instrumental variables. This approach is not feasible in our framework.

²²Detailed information on the public LFS can be found at [StatCan](#).

Survey of Financial Security (SFS). The SFS provides a comprehensive snapshot of the net worth of Canadian households, including detailed information on assets, debts, income, and employment. Location identifiers are not available. We use the 2012 SFS wave²³ to compute moments of the wealth and home-equity distributions. To match the model, we limit the sample to households with heads aged 25 to 85 and exclude households in the top 10 percent of the wealth distribution.

Statistics Canada (StatCan). StatCan provides annual city-level information on wages, income, employment, unemployment, and population by age that we use for 2011–2013. We use these data to construct earnings and unemployment rates across Canadian CMAs.

Teranet. We use the proprietary Teranet dataset, which provides quarterly house price indexes at the FSA level. We construct city-level house price indexes by aggregating FSA data using 2011 population weights.

4.2 Global Oil Crisis and Local Shocks in the Data

Canada is the fifth-largest oil producer in the world, with highly concentrated production in Alberta and Saskatchewan, provinces that include the metropolitan areas of Calgary, Edmonton, Saskatoon and Regina. The economic conditions in these cities are therefore highly sensitive to large fluctuations in global oil prices, which experienced a pronounced collapse between 2014 and 2016 (Figure H.1). As documented in Appendix H, income in oil-exposed locations declined sharply following this episode, in stark contrast with the relatively stable income dynamics observed in the rest of the country.

We exploit this important episode to guide the parameterization of our model and, in Section 5, examine how individuals respond to local downturns and to uncover the mechanisms behind these adjustments. A necessary first step is to identify local aggregate shocks at the city level. To do so, we leverage two key institutional features. First, despite being a major producer, Canada accounts for only a small share of global oil output, making international oil price movements plausibly exogenous to the Canadian economy. Second, the oil sector is highly geographically concentrated, generating substantial cross-sectional variation in exposure to the price collapse. Following Kilian and Zhou (2022), we construct a Bartik-style measure of regional exposure by interacting each city’s pre-determined oil employment share with the time series of real oil

²³We restrict our analysis to 2012 because the previous wave (2005) contains limited variable detail, and the next wave (2016) takes place after the oil shock. Further details are available at [StatCan](#).

prices:

$$\text{Oil Exposure}_{l,t} = \alpha_l \log(\text{OilPrice}_t) \quad (19)$$

where α_l is the share of employment in the oil sector (NAICS 2111) in city l in 2011 and $\log(\text{OilPrice})_t$ is the logarithm of the real oil price at quarter t . We then estimate the impact of this shock on annual income changes using the following panel regression with city and year fixed effects:

$$\Delta inc_{l,t} = \beta \text{Oil Exposure}_{l,t} + \phi_l + \phi_t + \varepsilon_{l,t}, \quad (20)$$

The estimated coefficients, reported in Appendix H, are economically large and statistically significant, indicating that the oil price collapse had substantial effects on local labor markets.

Our estimates imply that the predicted change in income induced by the oil price collapse, $\widehat{\Delta inc}$, accounts for approximately 85 percent of the observed decline in average income across the four CMAs of the oil-producing provinces between 2015 and 2017. The predicted change in income induced by the large decline in oil prices, $\widehat{\Delta inc}$, is a critical object for the rest of the paper. First, we estimate the migration response to this plausibly exogenous local shock to discipline key parameters of the model, most notably the dispersion of idiosyncratic local taste shocks ν . Second, we use $\widehat{\Delta inc}$ to back out the local productivity shocks \hat{Z}_j that we analyze in Section 5.

4.3 Parameters' Choice

We combine several approaches to parameterize the model. As is standard in the literature, a subset of parameters is taken directly from prior studies. Several parameters, such as housing supply elasticities and certain components of the income process, are obtained externally from detailed microdata, without solving the model's equilibrium. The remaining parameters are chosen internally to minimize the distance between a set of model-implied equilibrium moments and their empirical counterparts.

We target several cross-sectional features of the Canadian economy between 2011 and 2013, as well as the migration response to the large decline in oil prices. This requires solving both the stationary equilibrium and the transition dynamics in order to discipline the parameter values.

More precisely, we begin with an initial guess of the structural parameters and a tentative sequence of productivity shocks. Using the procedure described in Section 3.6, we solve the stationary equilibrium and the transition dynamics implied by this guess of \hat{Z} shocks. We then update the parameters based on the gap between the model-generated moments and their empirical targets. In parallel, the sequence of local shocks is updated so that the model-implied

decline in income in oil-exposed cities matches the predicted change in income induced by the large drop in oil prices, $\widehat{\Delta inc}$, as defined in Section 4.2. Appendix F.1 provides a formal description of the parameterization procedure.

4.3.1 Externally Set Parameters

Demographics The model period corresponds to two years of life. Households enter the model at age 25, retire at age 65 ($\bar{Q} = 25$), and die with certainty at age 85 ($Q = 30$). Age-specific mortality risk is captured by the probability of survival, $1 - \lambda_q$, obtained from StatCan.

Preferences We set $1/\gamma$, the elasticity of substitution between non-durable consumption and housing in equation (2), to 1.25, following the estimates in Piazzesi, Schneider and Tuzel (2007). σ is set to 2, implying an elasticity of intertemporal substitution of 0.5. Consumption equivalence scales, e_q , are taken from Auclert, Dobbie and Goldsmith-Pinkham (2019).

Endowments The age- and location-specific income component χ_q^l is taken directly from the data. Because income by age and location is not observed, we use the LFS to compute average wages for employed workers in each city and 5-year age bin, and we set χ_q^l to match the resulting relative wage differences. For retirees, the age component is common across locations and is determined by the relative income of retirees to working-age individuals at the national level.²⁴ The stochastic component of earnings for employed individuals ($\epsilon \neq \epsilon^1$) is modeled as an AR(1) process in logs with annual persistence of 0.91 and the standard deviation of innovations of 0.21, following Berger et al. (2018).

Transitions between employment and unemployment states are city-specific. The steady-state unemployment rate in each city is set to match the average unemployment rate between 2011 and 2013 in StatCan data. Using confidential LFS microdata from 2008 to 2013, we construct income terciles within each age group and city, and compute the monthly transition rates from each income tercile into unemployment and from unemployment into each income tercile. These monthly transition rates are then adjusted to produce a two-year transition matrix consistent with the model’s time period.²⁵

Wealth Distribution We set the annual risk-free interest rate to $r = 1.5\%$, the borrowing wedge ι to 1 percent. Initial bequests mimic the empirical distribution of wealth at the age of 25 years old across cities from SFS 2012.

²⁴The χ_q^l profiles for working-age individuals are rescaled to match both the average wage in each city and the nationwide ratio of aggregate income between working-age and retired households.

²⁵We normalize earnings so that median annual household earnings (64,000 CAD in the 2012 SCF) equal one in the model.

Housing The proportional annual maintenance cost that fully offsets depreciation, δ_h , and the proportional housing transaction costs, F , are set to 1.5 percent and 7 percent, respectively, following Kaplan, Mitman and Violante (2020b). We match the house price to rent ratio in the model to the one observed in the data, which allows us to back a city-specific property tax by inverting equation (18). The collateral constraint ξ equals 0.8 to match the standard 80 percent down-payment in Canada.

Housing Supply Elasticities We estimate city-level housing price elasticities following Guren et al. (2021). Their approach exploits systematic differences in cities’ responses to regional house price cycles. The main advantage of this methodology is that it does not require data unavailable for Canadian cities such as land availability, geographic characteristics and housing regulation in Saiz (2010). Instead, it relies mainly on long series of house prices at a high frequency that we obtain from Teranet. Appendix E.1 reports the full description of this methodology applied to Canadian cities. Figure E.2 plots housing elasticities by city.

Government Labor income is taxed following the functional form in Heathcote, Storesletten and Violante (2009), *i.e.*, $\mathcal{T}(y) = \psi_y^0 y^{1-\psi_y^1}$. ψ_y^0 and ψ_y^1 are chosen to match the federal and provincial effective tax rates across the income distribution. We obtain $\psi_y^0 = 0.92$ and $\psi_y^1 = 0.13$, which implies a mean effective tax rate of 3.7 percent and 15 percent at the 25th and 50th percentile of the income distribution in the model against 3.1 percent and 13.5 percent in the data.²⁶

Elasticity of Labor Demand Following the literature, we set the elasticity of labor demand η to 0.75, which sits right within the range of values used for this parameter.

4.3.2 Internally Set Parameters

Preferences The relative taste for housing services α set to 0.35 to match the median ratio of home equity to total net worth of 0.74, calculated using SFS 2012. The additional utility from owner-occupied housing relative to rental housing ω^l is assumed to be city-specific and chosen to match the homeownership rate in each city, consistent with the national average homeownership rate of 61 percent between 2011 and 2013 (StatCan). Figure F.2c in Appendix F shows that ω^l ranges from 1 to 1.3, with an average of 1.22. Homeownership rates by city are reported in Figure F.5 in Appendix F.

²⁶StatCan.

Endowments To parameterize unemployment benefits, we follow [Birinci, Park and See \(2023\)](#), who document an average earnings loss of 46 percent two years after a job separation caused by a mass layoff in Canada between 2008 and 2010.²⁷ To match this income decline following an unemployment shock, we set $v = 0.31$, which implies that an unemployed individual receives 31 percent of the average income of an employed worker of similar age in the same city.²⁸

Retirement income w_{ret} is set to 4.32 to match the relative income of retirees relative to working age groups. Figure F.4 in Appendix F shows that both in the data and in the model, retirees receive roughly 80 percent of the median income and 47 percent of the peak that occurs between 45 and 50 years old. To match the migration rate of households aged 25 to 35, the standard deviation of the stochastic component of earnings for employed individuals in this age group is set to 0.35.

Wealth Distribution We use 2012 SFS to obtain wealth and home-equity distributions and use several of moments to pin down different parameters. An annualized β of 0.974 generates a median wealth-to-income ratio of 3.39 that matches the data counterpart. The two parameters of the bequest function, $\bar{\varphi}$ and \underline{a} , are chosen to match the 20th and 30th percentile of the wealth to income distribution. These two moments are in the data 0.73 and 1.35, respectively. Setting $\bar{\varphi} = 930$ and $\underline{a} = 22$ implies these two moments in the model of 0.73 and 1.35 against 0.47 and 1.21 in the data. The limit on unsecured credit \underline{b} is to -0.6 to match a 6.6 percent share of households with negative networth in SFS 2012 (7.7 percent in the model).

Housing House grids are assumed to be city-specific. Specifically, the owner-occupied house size set, \mathcal{H}^H , has three elements and the rental housing size set, \mathcal{H}^R , has two elements with the following structure: $\mathcal{H}^{H,l} = [o_1^l \bar{h}^l, \bar{h}^l, o_2^l \bar{h}^l]$ $\mathcal{H}^{R,l} = [o_3^l \bar{h}^l, o_1^l \bar{h}^l]$. As in [Kaplan, Mitman and Violante \(2020b\)](#) we assume that the largest rental unit coincides with the smallest unit of owner-occupied house grid. $\{\bar{h}^l, o_1^l, o_2^l, o_3^l\}$ are chosen to match the average house sale price over median income by city computed using house price data from Teranet and average income from

²⁷The estimated earnings loss after two years includes both individuals who remained unemployed throughout the period and those who returned to employment. Thus, it captures both income loss during the unemployment spell and potential scarring effects once re-employed.

²⁸For tractability, we assume that v is common across cities and that all unemployed agents receive the same benefit regardless of their previous income. Although this is a simplification of the “Employment Insurance” (EI) system in Canada, the model-implied unemployment benefit is broadly consistent with observed replacement rates. In particular, individuals with lower pre-unemployment earnings face higher replacement rates than those with higher earnings. As of 2024, individuals who worked full-time in the year prior to job loss and earned below 30,000 CAD can receive up to 45 percent of previous annual earnings (ranging between 40–50 percent in our sample of cities). For individuals earning between 40,000 CAD and the EI maximum insurable amount of approximately 70,000 CAD, unemployment benefits amount to roughly 42 percent of previous earnings. Since weekly benefits are capped at 695 CAD, the replacement rate declines monotonically with previous income above this threshold.

StatCan, the 30th and 50th percentiles of the distribution of homeowners’ property value over total wealth from 2012 SFS and the average size ratio of owned houses to rental houses from StatCan. The average across cities of $\{o_1^l, o_2^l, o_3^l\}$ is 0.53, 0.75, and 1.38, respectively. The 30th and 50th percentiles of the distribution of homeowners’ property value over total wealth are 0.74 and 0.43 in the data against 0.74 and 0.52 in the model. The average size ratio of owned houses to rental houses is 1.54 both in the data in the model. Panel B of Figure F.5 in Appendix F shows a 97 percent correlation between the house price value over median income between the model and the data.

When solving the stationary equilibrium, we impose that the ratio between average house prices and median income in each location matches the one observed in the data. We then back out the housing stock in each city in the stationary equilibrium consistent with the house price to income ratio observed in the data. Therefore, by inverting equation (17), we obtain the equilibrium measure of construction permits, \bar{L}^l , which are shown in Figure F.2d.

Migration Elasticity The migration elasticity depends on $\frac{1}{\nu}$. To calibrate ν , we target the out-migration response in oil cities driven by the significant decline in oil prices between 2014 and 2016. To obtain the respective empirical response, we regress the annual changes in out-migration rates of individuals aged 25–54 $\Delta\text{out-mig}_{l,t}$ on the predicted income change implied by the decline in oil prices. In particular, we run the following specification:

$$\Delta\text{out-mig}_{l,t} = \delta_0 + \delta_1 \widehat{\Delta inc}_{l,t} + \phi_l + \phi_t + \epsilon_{l,t}, \quad (21)$$

where $\widehat{\Delta inc}$, the predicted change in income implied by the declined oil prices, is obtained from specification (29) as explained in Section 4.2.²⁹ Out-migration rates are computed using TransUnion data. ϕ_l and ϕ_t are city and year fixed-effects. Results are reported in Appendix H.

We then obtain the implied change in out-migration rates in 2016 and 2017, $\widehat{\Delta\text{out-mig}}_{l,16-17}$ for cities in the oil provinces and use it as a target in our parameterization strategy. In particular, we initialize the economy in the stationary equilibrium and impose an unanticipated temporary decline in productivity shock that is consistent with $\widehat{\Delta inc}$ (details below). We then solve the transition dynamics implied by the shock and compare the average out-migration response in Oil cities implied by the model one period after the shock (2 years) against the empirical counterpart, the average of $\widehat{\Delta\text{out-mig}}_{l,16-17}$. This is one of the multiple targets considered in the joint parameterization. As expected, the moment is particularly informative about the migration elasticity, which is found to be 1. The model implied change in out-migration of oil

²⁹ $\widehat{\Delta inc}_{l,t}$ is set to zero for cities outside Alberta and Saskatchewan.

cities is 1.04 against 1.06 percentage points in the data. In Section 5 we explore in more detail the impact of the oil shock on migration as well as other economic variables.

Table F.1 in Appendix F.2 demonstrates the robustness of our parameterization. It reports the implied migration response to the oil shock and the associated loss functions for different values of ν in the range 0.85–1.15, after re-calibrating all other parameters to minimize the distance between equilibrium moments and their data counterparts. We find that both the deviation from the observed out-migration response and the overall loss across all moments are minimized at $\nu = 1$.

Moving Costs Our framework includes both monetary moving costs, F_m , and utility moving costs defined in equation (5). We assume that both types of moving costs do not depend on the individual characteristics, and that utility moving costs depend on the straight line distance, $D_{i,i'}$, between any two pairs of cities.³⁰

Because Canada is a bilingual country, moving between English-speaking and French-speaking regions may involve additional frictions. We also account for the unusually strong transportation links to and from Toronto and Vancouver. To capture these features, we allow the parameter $\tau_{i,i'}^1$ to vary by origin–destination pair. Specifically, we set $\tau_{i,i'}^1 = \bar{\tau}^1$ for moving between two city pairs that are either English (except Toronto and Vancouver) or French-speaking cities. Moves between English- and French-speaking cities are associated with $\tau_{i,i'}^1 = \bar{\tau}_F^1$. Moves from any city (except those in Quebec) to Toronto or Vancouver are associated with $\tau_{i,i'}^1 = \bar{\tau}_E^1$. The parameters $\bar{\tau}^1$, $\bar{\tau}_F^1$ and $\bar{\tau}_E^1$ are chosen to match three moments: the aggregate correlation between distance and out-migration rates, the share of in-migration to Quebec cities (Montréal and Québec City), and the share of in-migration to Toronto and Vancouver. We discipline the parameter τ_0 by matching the average annual out-migration rate of 1.47 percent. Monetary moving cost, F_m , match the average migration rate among households in the lowest wealth quartile. The implied parameter values are $\tau_0 = 6.955$, with distance-related parameters $\bar{\tau}^1 = 0.0065$, $\bar{\tau}_F^1 = 3.5\bar{\tau}^1$ and $\bar{\tau}_E^1 = 0.02\bar{\tau}^1$. Monetary moving costs equal $F_m = 0.23$, corresponding to 7,355 CAD (in 2012 units).

Converting utility-based moving costs into dollar equivalents is non-trivial in our framework because the utility function is nonlinear. To perform the conversion, we compute, for each individual, the change in consumption required to deliver the same lifetime utility in an environment with no moving costs.³¹ Appendix G provides the full procedure. We find that

³⁰Distance is measured as the straight line linking the geographic centers of each pair of cities. We normalize the distance between Guelph and Cambridge-Kitchener-Waterloo (C-K-W), the two closest CMAs in our sample, to one.

³¹Because utility depends on both non-durable consumption and housing services, we rewrite lifetime utility

the moving costs between Canadian cities are approximately 288,000 USD (in 2012 units), as reported in the first row of Table G.1. Kennan and Walker (2011) were the first to estimate within-country moving costs using a forward-looking model of location choice with income risk but no financial markets. They estimate an average moving cost of 312,000 USD (2010 units) for a representative adult male in the United States. Our baseline model implies a moving cost of 306,000 USD (2010 units) for Canadian males (row 3 of Table G.1), slightly below their estimate. Given differences in context and the fact that our model does not include the zero-wealth scenario considered by Kennan and Walker (2011), such cross-study comparisons should be interpreted with caution.

Amenities Amenities are allowed to vary across cities and 5-year age bins and chosen to match the population distribution across cities by age bins. To match the model specification, we restrict population data from Census 2011 to individuals with age between 25 and 85 years old. The total population in the economy is normalized to one in the model. Our distribution of amenities is reported in Figure F.2b in Appendix F.

Productivity and Agglomeration Aggregate productivity depends on two components: city-level exogenous productivity, z^l , and endogenous agglomeration forces governed by ζ . We discipline these objects using two targets: the average wage of the working-age population (ages 25–65) in each location between 2010 and 2013 using the LFS, and the empirical correlation between city productivity and in-migration.

Pre-shock values of z^l cannot be obtained by simply inverting the wage equation (14) in the stationary equilibrium. First, wages depend on the strength of agglomeration forces ζ . Second, N_c^l , the effective labor input in final-good production, reflects both the age composition of workers and their idiosyncratic productivities ϵ , which are not directly observable. To account for potential selection of workers with different idiosyncratic productivities across cities, we jointly choose $\{z^l\}$, ζ , and the remaining parameters internally.

The resulting city-specific exogenous productivity components are reported in Panel A of Figure F.2a in Appendix F. We normalize productivity levels so that the median annual household earnings in 2012 (64,000 CAD) corresponds to one in the model. The agglomeration parameter is set to $\zeta = 0.13$ to match the empirical correlation between city productivity and in-migration of 76 percent (versus 80 percent in the model).

Productivity Shocks Induced by the Oil Crisis As described earlier, the economy is

in terms of an “adjusted-consumption” measure $\omega_{i,t}$ that replicates the same flow utility (net of amenities) as in the stationary equilibrium.

Table 1: Moments for Internal Calibration

Moment	Data	Model	Stationary Distribution	Transition Dynamics
Average out-migration (%)	1.43	1.47	✓	
Correlation (distance, out-migration)	-0.24	-0.23	✓	
Correlation (productivity, in-migration)	0.76	0.80	✓	
Migration rate of the youth	2.93	2.90	✓	
Migration rate of the lowest wealth quartile	2.20	2.22	✓	
Share in-migration to Quebec cities	11%	13%	✓	
Share in-migration to Toronto and Vancouver	35%	22%	✓	
Average retirees income / median income (%)	82	82	✓	
Earning loss after job separation (%)	46	46	✓	
Share with negative assets (%)	6.63	7.70	✓	
20th percentile net worth / income	0.47	0.74	✓	
50th percentile net worth / income	3.39	3.39	✓	
30th percentile house value / net worth	0.47	0.52	✓	
50th percentile house value / net worth	0.74	0.74	✓	
Avg. owned / rented house size	1.50	1.54	✓	
Avg. house price / median income by city		F.5	✓	
Homeownership share by city		F.2c	✓	
Average wages by city		F.3	✓	
Population distribution by age by city		F.2b	✓	
Out-migration response to the Oil Shock (%)	1.05	1.04		✓
Income response to the Oil Shock (%)	-9.62	-9.65		✓

Note: Table 1 reports the targeted moments in the model parameterization. Data sources are described in the main text. The internally chosen parameters are reported in Table F.2 of Appendix F.

initially in a stationary equilibrium without aggregate uncertainty, calibrated to match several cross-sectional moments of the Canadian economy between 2011 and 2013. We then introduce an unanticipated, temporary shock to the exogenous productivity component, \hat{Z}^l , in the four cities in our framework that have high exposure to the oil sector.

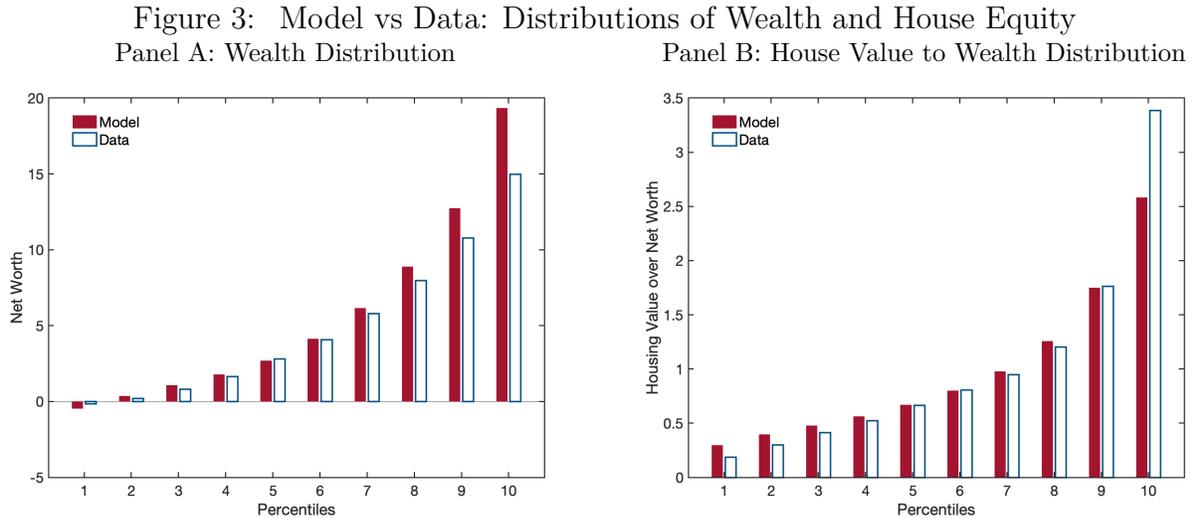
To discipline these shocks, we solve the full transition dynamics and choose \hat{Z}^l such that the model-implied decline in income in those cities matches the observed change in income between 2015 and 2017 triggered by the exogenous collapse in oil prices (the measure $\widehat{\Delta inc}$ defined in Section 4.2).³²

³²Additional details on the magnitude of these shocks are provided in Section 5, where we analyze their

4.4 Model Matching Data

This section documents a set of stationary-equilibrium predictions from the parameterized model. We analyze the distributions of wealth, income, and population. These are central moments that underpin heterogeneous migration behavior and play a fundamental role in the local downturn counterfactual analysis.

Wealth and House Value Distributions Panel A of Figure 3 plots the wealth distribution in the data and model. We explicitly target the second and fifth deciles, but the model is able to reproduce closely the entire wealth distribution in the data below the top decile. Migration decisions for households on the top of the wealth distribution are quite insensitive to individual conditions, the model’s limited fit at the very top is not a major concern for our analysis. Panel B of Figure 3 reports the ratio between housing value and wealth for homeowners. Overall, the model matches closely the data but underestimates slightly the house value in terms of wealth at the top of the distribution. In Figure F.6 of Appendix F, we report the wealth distribution separately by homeownership status.



Note: Figure 3 plots the wealth and the house value to wealth ratio distribution. In Panel A, wealth is normalized by income. Panel B plots the distribution of household house value to wealth ratio. *Data Source: SFS 2010.*

Migration Figure 4 plots the migration rates across demographic groups in the model and data.³³ Migration rates by demographic groups are not targeted in the calibration exercise,

implications for migration responses and other economic outcomes.

³³The age structure in the model does not precisely match that in the data, with the model overestimating the older population. Moreover, the model generates a stronger correlation between age and wealth than observed in the data. To enable a more accurate comparison of migration rates between the model and the data, we construct migration rates by wealth using an age-adjusted approach. Specifically, we compute migration rates

except for the average migration rate, the migration rate for households with ages below 35 years old, the migration rate of the lowest wealth quartile and in-migration to Quebec cities and Toronto and Vancouver. Despite the homogeneous moving costs, we find that the model delivers heterogeneous migration patterns across demographic characteristics consistent with the data. Panel A plots the migration rates in the model and data by homeownership status. As in the data, the model generates higher migration rates for renters than homeowners. The model delivers a yearly migration rate for renters of 2.37 percent and approximately 1 percent for homeowners. In the data, the yearly migration rate is 1.8 percent and 1.23 percent, for renters and homeowners, respectively. In panel B, we observe that the migration rates by age group for in the model replicate very closely the ones in the data. Panel C reports the migration rates by wealth quartiles. In the data we have no household-level wealth information so we are not able to replicate the same migration patterns on the data. Instead, we plot the migration rates by credit score bins.

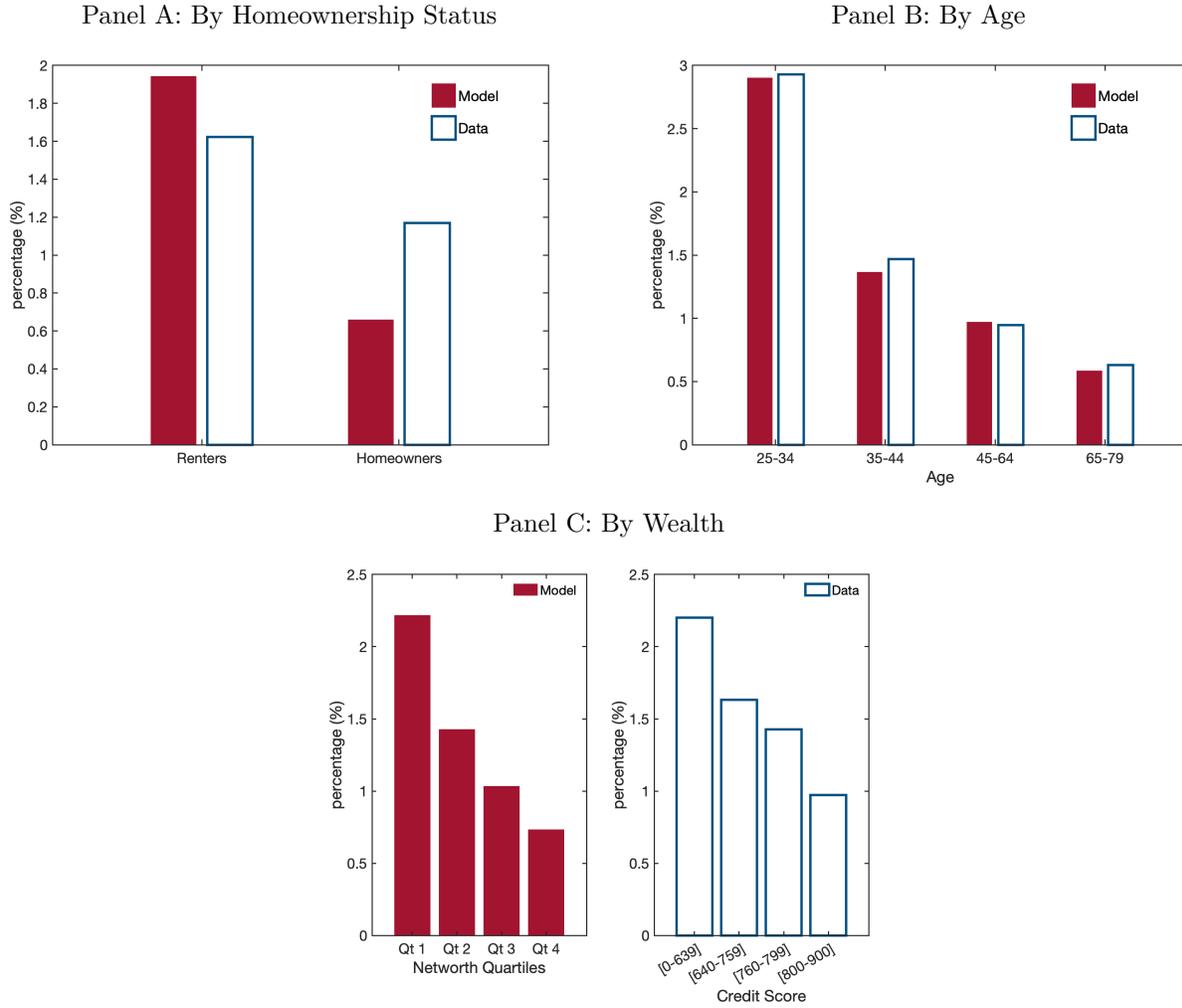
Households at the bottom of the wealth distribution are financially constrained as they are closer to their borrowing limit and have less capacity to adjust their borrowing. Households with lower credit scores are less likely to obtain credit, therefore, more financial constrained. The underlying assumption is that wealth and credit score are highly correlated. We find both in the model and in the data that migration rates decrease with wealth (credit score). In the model, households at the bottom of the wealth distribution are three times more likely to move than those at the top of the distribution. Specifically, on average, the annual moving rate is 2.62 percent and 0.56 percent for households at the first and fourth quartiles of the wealth distribution, respectively. In the data, we observe that households with a credit score below 640 move at an average annual rate of 2.2 percent, while those with a credit score above 800 move at 1.1 percent per year.

Overall, the model performs satisfactorily in matching the heterogeneity in the migration rates observed in the data. This results suggests that heterogeneity in moving costs is less important to match the different migration rates across demographic groups than commonly thought. As it will become clear in Section 4.5, income risk, endogenous wealth and financial constraints generate substantial dispersion in the moving benefits across groups. Consequently, despite having common moving costs, the model generates heterogeneous migration rates consistent with the data.

We now analyze where households move to. Figure 5 plots the share of movers by destination

across wealth quartiles within each age group and then take a weighted average across age groups, both in the data and in the model. An analogous approach is applied when computing migration rates by homeownership status.

Figure 4: Model vs Data: Migration Rates across demographic groups

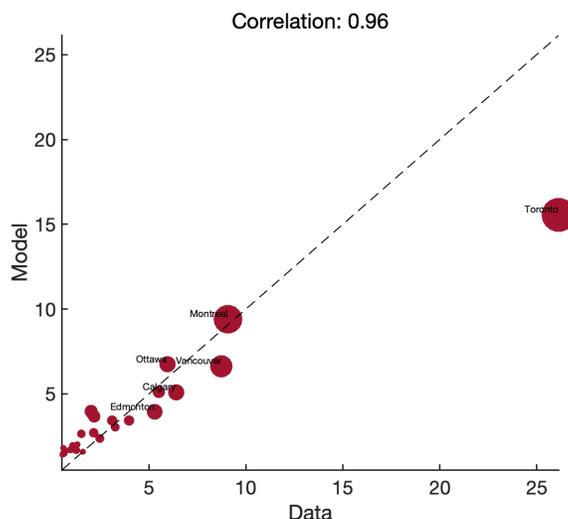


Note: Figure 4 plots the annual migration rates from the model (red bars) and from the data (hollow blue bars) by demographic groups. Panel A plots it by homeownership status, panel B by age and panel C by wealth on the left (model outcomes) and by credit score on the right (data outcomes).

city. Both in the data and in the model, more than 50 percent of the movers choose to move to the five biggest Canadian cities. In the data, there is a disproportionately high fraction of movers to Toronto (26 percent), while Montreal, the second city receiving the highest number of movers, only absorbs 9 percent of the migrants. The model is not able to match this discontinuity observed in the data, partially explained by the lack of heterogeneity in location preference shock. Nevertheless, the model captures the main trends in location choices. Table 2, which reports the correlation between the share of inflow migrants and the characteristics of the destination cities.

There is a very strong correlation between in-migration and the size of the destination city (0.94 and 0.97 in the data and model, respectively). In-migration is also strongly correlated with

Figure 5: Model vs. Data: In-Migration by city



Note: Figure 5 plots the in-migration rates in the model and in the data by destination. Cities are ordered in ascending order by population. *Data source:* TransUnion.

the exogenous component of productivity (0.76 and 0.80 in the data and model, respectively), house prices (0.28 and 0.28 in the data and model, respectively), and amenities (0.27 and 0.29 in the data and model, respectively). Similarly, the correlation for income measures are equally close in data and model: average labor income is 0.28 in the model and 0.25 in the data and average income is 0.29 in the model and 0.22 in the data. Once validated, our calibration strategy, in the next Section, we unpack the different economic forces that drive households' moving choices.

In Appendix F, we present more evidence of the model's ability to match the spatial heterogeneity observed in the data. In Figure F.5, we plot the distribution of median income and house prices across cities. Table F.3 shows the correlation between city characteristics such as house prices, wages, average income, population, productivity, and amenities, both in the data and in the model. Overall, we find strong positive correlations between income, productivity, and house prices in the data and in the model. However, the model underestimates the positive correlation between house prices and population, but matches very well the relationship between population and productivity. In terms of amenities, there is a positive correlation between amenities and house prices, population and productivities both in the data and model. The model, however, overestimates the negative correlation between amenities and income measures.

Table 2: Model vs. Data: Correlation Between Migration and City Characteristics

Characteristics	Correlation	
	Data	Model
Average Labor Income	0.28	0.25
Average Income	0.29	0.22
Productivity	0.76	0.80
House Prices Index	0.28	0.28
Population	0.94	0.97
Amenities	0.27	0.29

Note: Table 2 reports the correlations between the share of in-migrants to a city and a range of city characteristics in the data and in the model. The characteristics include housing prices, wages, average household income, population size, productivity, and local amenities. These correlations highlight the extent to which the model captures the relationships between migrant inflows and the structure of local economic fundamentals. *Data sources:* TransUnion credit panel, StatCan and LFS.

4.5 Model Mechanisms

We study why migration responses differ across demographic groups (Figure 4) even though utility and monetary moving costs are identical across individuals. Two model ingredients are central: market incompleteness and illiquid housing wealth.

What drives migration? In standard spatial equilibrium models, migration is primarily pinned down by moving costs, idiosyncratic preference shocks, and location fundamentals (e.g., productivity and housing costs). With similar moving costs, a finite horizon reduces older agents' expected gains from relocating, helping rationalize the decline in mobility with age. However, preference shocks and labor opportunities alone cannot generate the pronounced heterogeneity in migration propensities across the wealth distribution conditional on age. We therefore focus on three channels that are comparatively underexplored in this context: housing transaction costs, income risk, and financial frictions.

To quantify their role, we conduct a partial-equilibrium shock experiment. Starting from the stationary equilibrium, we trace each individual's moving response to an idiosyncratic, transitory-but-persistent negative income shock.³⁴ We compute the difference in migration

³⁴For any individual with positive mass in the stationary-equilibrium distribution, if they remain in their stationary-equilibrium location, their income is assumed to be 10 percent below its pre-shock level for three periods and then gradually returns to its original level over the subsequent three periods.

propensities with and without the individual-specific shock and average across individuals. Because agents are atomistic, this perturbation generates no general equilibrium effects; prices remain fixed at stationary-equilibrium levels. We repeat the exercise under four environments: (i) the *Baseline* economy; (ii) *No Housing Transaction Costs*, setting $\delta_h = 0$; (iii) *No Income Risk*, imposing $\sigma = \sigma_y = 0$; and (iv) *No Borrowing Limit*, setting $\underline{b} \rightarrow -\infty$.³⁵

Figure 6 summarizes the induced changes in migration propensities: Panel A reports the average response across the four environments, and Panel B reports the average response by wealth.³⁶ Across environments, the negative income shock increases mobility on average. In the *Baseline* economy, the probability of moving rises by 0.6 percentage points, a 40 percent increase relative to the 1.47 percent stationary-equilibrium moving rate. The response is strongly decreasing in wealth: 1.25 percentage points at the bottom of the wealth distribution versus 0.1 percentage points in the top wealth quartile, indicating that wealth is a key determinant of migration decisions.

Housing transaction costs operate in the expected direction. Setting $\delta_h = 0$ raises the average moving response relative to the *Baseline*, but the amplification is quantitatively modest, even among homeowners (see Figure F.7 in Appendix F). By contrast, shutting down either component of market incompleteness—uninsurable income risk or borrowing constraints—substantially attenuates migration responses.³⁷ In the *No Income Risk* and *No Borrowing Limit* economies, the moving response is 0.17 and 0.45 percentage points, respectively—approximately 73 and 33 percent smaller than in the *Baseline* economy.³⁸ Importantly, these gaps widen monotonically toward the bottom of the wealth distribution: relative to *Baseline*, the bottom-of-wealth response is 80 percent lower without income risk and 36 percent lower without borrowing constraints. Taken together, market incompleteness—via income risk and financial frictions—is a key driver of migration responses, especially for low-wealth households.

Why do low-wealth households move more? With incomplete markets and risk-averse households, precautionary behavior limits welfare losses from adverse income realizations. In standard macro models without a spatial margin, this manifests as precautionary savings: households accumulate liquid assets to self-insure against income risk and to avoid binding

³⁵In all cases other than the *Baseline*, we solve for the households’ problems holding prices and all fundamental parameters fixed at their baseline values, except for the parameter modified in each counterfactual, and compare the behavior of the individual after the income shock with their behavior without shock.

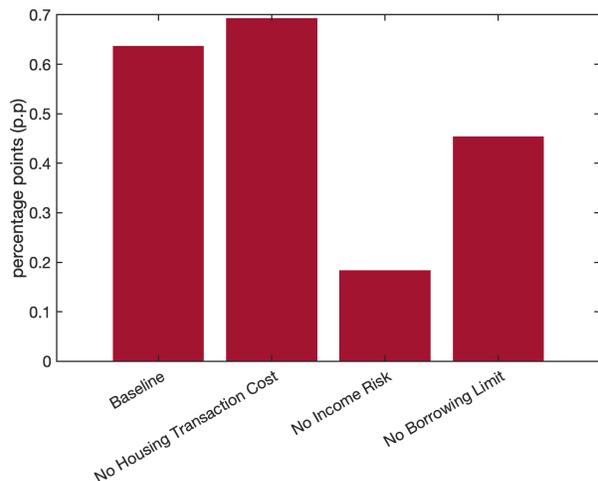
³⁶Migration responses are measured in the first period following the shock. To aggregate individual migration responses, we use the stationary-equilibrium distribution of the *Baseline* economy across all four environments.

³⁷Even when the borrowing constraint is removed, markets remain incomplete because households do not have access to a full set of state-contingent securities spanning all future states.

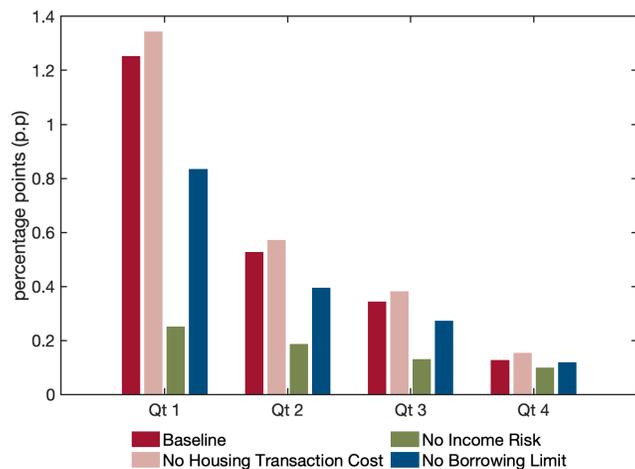
³⁸Appendix F also reports the corresponding results by homeownership status and age.

Figure 6: Decomposition: Change in Migration Propensities

Panel A: Change in Aggregate Migration Propensity



Panel B: Change in Migration Propensity by Wealth



Note: Panel A plots the steady-state-to-shock change in aggregate migration propensity; Panel B plots the analogous change by wealth quartile. Results are for the average response to a negative local productivity shock under four specifications: baseline, No Housing Transaction Costs, No Income Risk, and No Borrowing Constraint.

borrowing constraints. Our framework adds migration as a second, costly margin of adjustment. In this setting, liquid buffers and migration are imperfect substitutes for smoothing utility following income shocks.³⁹

Following a negative shock, households compare the value of staying to the value of relocating. The value of staying depends on wealth and access to financial markets. Relocating allows households to re-optimize spatially—potentially accessing higher-productivity locations or lower housing costs—but requires paying monetary and utility moving costs. Because wealth governs the effectiveness of financial self-insurance, it also governs the attractiveness of relocation. High-wealth households can absorb shocks by drawing down savings or borrowing, making moving relatively unattractive. Low-wealth households face limited financial smoothing; staying often entails large consumption adjustments and substantial utility losses, so relocation becomes a valuable form of insurance via improved future income prospects or reduced housing expenditures. Hence, even with identical moving costs, the value of moving relative to staying declines with wealth, rationalizing stronger migration responses among low-wealth households.

This mechanism is reflected in Figure 6: migration responses for low-wealth individuals decline sharply when income risk is removed or when borrowing constraints are relaxed. In both

³⁹In the standard Bewley–Huggett–Aiyagari class of models (income risk, incomplete markets, concave utility, no migration), households smooth consumption via borrowing/lending and precautionary asset accumulation. Our model retains this mechanism and adds migration as an additional smoothing device; savings/borrowing and migration are therefore imperfect substitutes.

cases, low-wealth households gain effective financial smoothing, raising the value of staying and compressing migration responses across the wealth distribution; the wealth–mobility gradient weakens. These findings indicate that uninsurable income risk and financial frictions are central to generating the negative relationship between migration and wealth observed in the data, and they help explain why frameworks that abstract from either precautionary moving or precautionary savings typically require heterogeneity in moving costs across groups to match observed mobility patterns.

Finally, because households are forward-looking, these forces operate *ex ante* as well as *ex post*. Under income risk and incomplete markets, households sort into locations with higher insurance value, implying that the model jointly delivers *Precautionary Savings* and *Precautionary Moving*. The latter is particularly important at the bottom of the wealth distribution, where migration is a key adjustment margin in response to adverse income shocks.

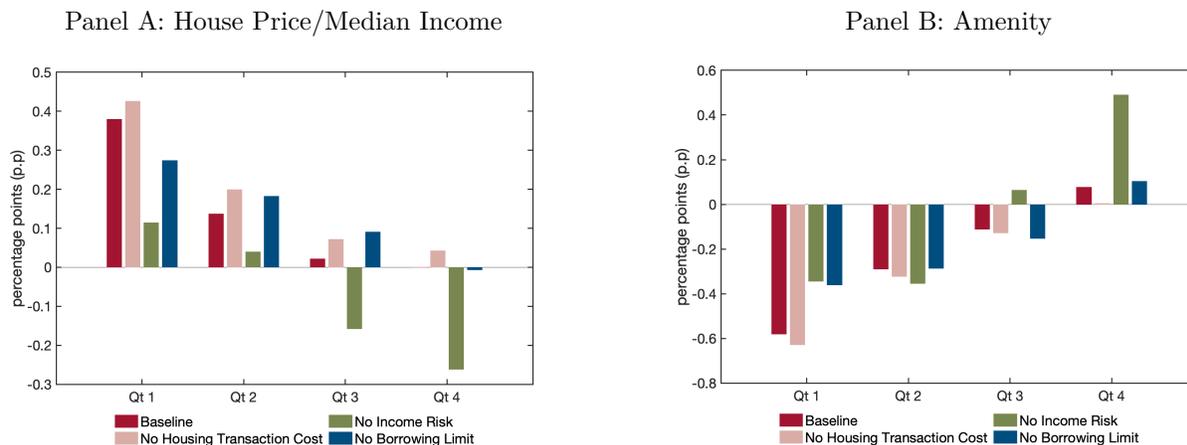
Where do households go? The precautionary moving motive shapes not only migration rates but also destination choice. Following adverse income shocks, low-wealth households disproportionately relocate toward lower housing-cost (lower house price) cities, even when doing so entails lower wages or amenities. This pattern reflects the interaction of housing indivisibilities, borrowing constraints, and uninsurable income risk.⁴⁰ For financially constrained households already at the minimum, housing expenditures cannot adjust further, so shocks are absorbed primarily through sharp reductions in non-housing consumption, generating large utility losses.

Lower house prices ensure against this constraint by reducing the minimum housing outlay required to obtain a given level of housing services. This frees resources for non-durable consumption and precautionary savings, attenuating both contemporaneous and intertemporal utility losses. Consequently, in the baseline economy, low-wealth households choose destinations that trade off wages and housing costs to maximize the insurance value of location, rationalizing the disproportionate reallocation toward cheaper housing markets after adverse shocks.

Figure 7 quantifies this channel by plotting (across the wealth distribution) the shock-induced change in the fraction of movers relocating to lower median house price income ratio destinations. In the baseline economy, low-wealth movers are substantially more likely to select cities with a lower house price income ratio than wealthier movers. When income risk is shut down or borrowing constraints are relaxed, the housing-based insurance channel weakens: financial smoothing (or the absence of risk) makes minimum housing constraints less relevant and lowers

⁴⁰As discussed earlier, households face a lower bound on housing consumption due to a minimum house size. As described in Section 3.2, housing services are supplied in a finite set of discrete sizes.

Figure 7: Decomposition: Destination Choice by Wealth



Note: Figure 7 reports, by wealth quartile, the difference in the share of households that move to locations with a median higher house price relative to median income compared to the rest of the distribution, comparing different model specifications in response to a local negative productivity shock in Panel A. Panel B reports the difference in the share of households that go to locations with higher amenities compared to the rest of the distribution, comparing different model specifications in response to a local negative productivity shock.

the insurance value of low housing costs, reducing the incentive for low-wealth households to move to cheaper destinations. Quantitatively, for the lowest wealth quartile, the share of movers relocating to lower house price income ratio cities falls by approximately 0.3 percentage points (75 percent) in the *No Income Risk* economy and by 0.1 percentage points (25 percent) in the *No Borrowing Limit* economy. If homeownership frictions were not to be present, these results would go the other way round increasing the share of individuals moving to higher median house price income ratio cities.

What are low-wealth households giving up in order to access cheap locations? Amenities as reported in Panel B of Figure 7. Low-wealth individuals after the shock are 0.6 percentage points less likely to move to locations with higher amenities than in the steady state. If income risk or borrowing constraints were not present, this probability would decrease by 50 percent, making high amenities locations more appealing. If housing frictions were not present, low-wealth individuals would be even more less likely to move to high amenity locations.

Overall, households jointly trade off housing costs and labor market conditions when choosing where to move, but spatial insurance operates primarily through lower housing prices. This also implies that the “location as an asset” mechanism of [Bilal and Rossi-Hansberg \(2021\)](#) is quantitatively significant.⁴¹ Consistent with this interpretation, when financial markets provide

⁴¹In [Bilal and Rossi-Hansberg \(2021\)](#), constrained individuals downgrade their location after a negative front-loaded income shock. Because the correlation between income and rents is 1, moving to places with lower housing costs necessarily implies lower-income locations. In our model, households choose along both margins: the correlation between wages and house prices is 68% in the baseline economy (64% in the data). The imperfect correlation partially mutes the downgrading effect, but we still predict that low-wealth individuals facing income

better smoothing or income risk is absent, reliance on cheaper housing markets as insurance weakens and destination choices become less dependent on wealth.

5 The Effect of A Local Downturn

Thus far, we have shown that migration insures households against idiosyncratic income shocks, and that this insurance is most valuable when income risk is high and financial markets are incomplete. It follows that large aggregate shocks should spur migration. Yet, although out-migration rises in the hardest-hit regions, the aggregate response is often viewed as modest relative to the size of the shocks.

This section investigates the mechanisms that attenuate migration responses to aggregate shocks. We focus on mechanisms that arise in a spatial equilibrium model with both *liquid* and *illiquid* wealth held in owner-occupied housing. In particular, declines in house prices generate a *housing-wealth lock-in* effect. Lock-in can also occur without housing wealth, or with only liquid assets, but its underlying forces differ for homeowners. We demonstrate that incorporating homeownership can overturn the aggregate welfare implications of lock-in, even in the absence of housing transaction costs, benefiting some households while harming others.

5.1 Global Oil Crisis

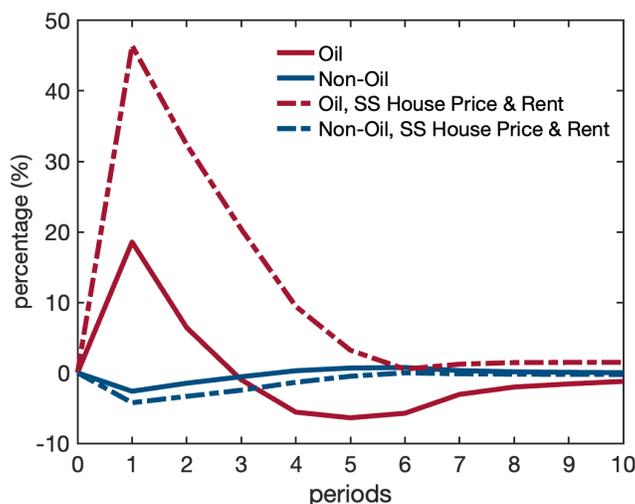
We exploit the 2014–2018 global oil price collapse as a natural experiment capturing the large downturn in oil-producing regions of Canada. As described in Section 4 and Appendix H, global oil prices declined by approximately 62 percent between 2014 and 2016. Leveraging cross-city variation in exposure to the oil sector, we estimate that this shock led to an average decline of 9.62 percent in income among the working-age population in the four cities in the Oil Provinces represented in our model, with effects that persisted over several periods.

To analyze this episode within the model, we impose an unexpected decline in the exogenous component of city-level productivity in oil-producing cities. The magnitude of this productivity shock is chosen to match the observed income decline in the four Oil Province cities, as detailed in Section 4.3.2. The shock is temporary but persistent, with local productivities reverting back to the steady-state level after four periods⁴² (Panel A of Figure J.1 in Appendix J). The economy is initially in its stationary equilibrium at $t = 0$. At $t = 1$, the unexpected shock is

risk and borrowing constraints are more likely to move to locations with lower house prices and income than they would in an environment with perfect consumption smoothing absent “location as an asset.”

⁴²We assume that the exogenous productivity components of cities in Non-Oil Provinces remain fixed at their pre-shock .

Figure 8: Migration Response to the Local Downturn



Note: Figure 8 reports the migration response to the oil-induced local downturn. In solid lines we report the values for the out-migration shock response in the baseline model in red for oil cities and blue for non-oil cities. The dashed lines report the respective values for the case in which house prices and rents are fixed at their pre-shock levels.

realized. Households in all locations have rational expectations and perfectly anticipate the paths of exogenous productivity across cities.

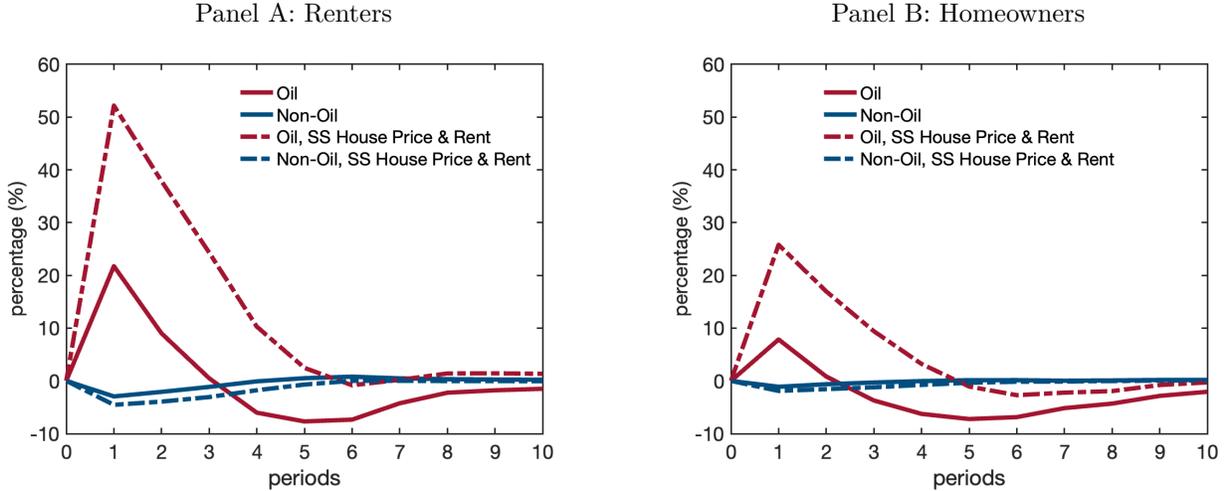
Migration Response. Consistent with the data, out-migration from Oil Provinces rises following the shock.⁴³ As shown in Figure 8, out-migration from Oil Provinces (solid red line) increases by 45 percent in the first period after the shock and gradually returns to its pre-shock level after approximately ten periods. In contrast, out-migration from non-oil cities (solid blue line) declines slightly over several periods.

This aggregate response conceals sharp heterogeneity across groups. As shown in Figure J.4 and J.5, out-migration increases most for younger individuals, renters, and low-wealth households. Tenure differences are particularly pronounced, even after controlling for age.⁴⁴ In the first period, migration increases by 22 and 8 percent above their stationary-equilibrium levels for renters and homeowners, respectively. In the stationary equilibrium, 25–35-year-olds migrate at roughly twice the rate of 36–45-year-olds, yet the shock raises out-migration by 30 percent for the latter, implying a larger proportional response. Households aged 65 and older instead migrate less than in the pre-shock period. Wealth gradients are equally stark: the

⁴³Given the model’s parameterization strategy, aggregate out-migration of individuals aged 25-54 from Oil Provinces in the first period is, by construction, aligned with the shock-induced out-migration estimated in Section 4.

⁴⁴Decompositions by homeownership status and wealth quartile control for age by computing, within each age group, the average migration response separately for renters and homeowners (or for each wealth quartile).

Figure 9: Migration Response to the Local Downturn by Homeownership



Note: Figure 9 reports the migration response to the oil-induced local downturn separately for renters (Panel A) and homeowners (Panel B). In solid lines, we report the values for the out-migration shock response in the baseline model in red for oil cities and blue for non-oil cities. The dashed lines report the respective values for the case in which house prices and rents are fixed at their pre-shock levels.

bottom quartile’s migration rate rises by 27 percent relative to the first period, versus only 11 percent relative to the second quartile. The top quartile exhibits virtually no response.

House Price Dynamics. A key distinction relative to the idiosyncratic shocks analyzed in Section 4.5 is that aggregate shocks generate declines in both house prices and rents. As shown in Panel D of Figure J.1, house prices and rents in oil-producing cities exhibit a highly persistent decline following the shock, remaining below their pre-shock levels even ten periods later.

A similar pattern is observed in the data. Using the same identification strategy employed to estimate the effects of oil price shocks on migration, we estimate equation (21) with house price growth as the dependent variable. Table H.2 in Appendix H shows that the income decline induced by the oil price shock leads to a significant reduction in house prices. In particular, the oil price collapse is associated with a 6.8 percent decline in average house prices in oil-producing cities between 2015 and 2017. The model implies a comparable 4.5 percent decline in these locations.

As discussed above, housing costs play a central role in migration and location choices. At the same time, equilibrium house prices are jointly determined by the spatial distribution of households and their demand for housing. This interaction becomes particularly salient following large local shocks that simultaneously affect income and housing markets. A key feature of our framework is the explicit incorporation of both liquid and illiquid wealth, in the form of homeownership, into a model of location choice. This structure allows us to study how endogenous house price dynamics interact with household balance sheets and influence migration

responses to local shocks, and in particular, whether such dynamics attenuate migration flows following adverse aggregate shocks.

5.2 Housing-Wealth Lock-in Effects

Households' wealth positions play a central role in shaping their responses to local shocks. As shown above, households at the bottom of the wealth distribution are substantially more likely to migrate following a local adverse shock than households at the top of the wealth distribution (Figure 6). In our framework, household wealth comprises both *liquid* and *illiquid* components, with the latter represented by owner-occupied housing. This distinction becomes particularly important when local shocks are accompanied by declines in house prices. As documented in the macroeconomic literature, declines in house prices generate negative wealth effects that depress consumption (Berger et al., 2018; Kaplan, Mitman and Violante, 2020b). Their implications for mobility, however, are less well understood.

To isolate the role of the house price decline induced by local shocks in shaping migration responses, we expose the economy to the same post-shock paths of wages as in the baseline, while holding house prices and rents in all locations fixed at their stationary-equilibrium levels. Appendix J also reports the counterfactual where only house prices are held fixed. Figure 9 shows that, absent the decline in house prices, homeowners would exhibit a stronger migration response to the shock. Quantitatively, the migration response of homeowners would be three times relative to the baseline with endogenous house price dynamics.

The mechanism operates through a housing wealth channel. The temporary decline in house prices induced by the local shock reduces homeowners' net worth, lowering the incentive to sell and relocate in search of better labor market opportunities. Forward-looking homeowners anticipate that, because the shock is temporary, house prices will eventually recover. As a result, they face a trade-off between selling their property at a depressed price—thereby foregoing future capital gains—and remaining in place despite weaker local labor market conditions. This intertemporal trade-off generates a *housing-wealth lock-in* that attenuates migration responses to local shocks. To our knowledge, this source of lock-in has not been explicitly identified in the spatial literature. Appendix J shows that this *housing-wealth lock-in* effect is particularly strong for older and wealthier households, while its impact on renters is negligible. The decline in house prices reduces out-migration among retirees, who do not experience an income decline following the shock. At the same time, lower house prices attract retirees to oil-producing cities, as they take advantage of reduced housing costs during downsizing.

Our framework also features a second, distinct lock-in mechanism operating through rental

markets, qualitatively highlighted in Glaeser and Gyourko (2005). The decline in local rents following the shock partially offsets the negative income shock faced by renters, reducing the gains from relocating. Lower rents act as a form of local insurance, attenuating migration incentives even in the absence of housing wealth. Appendix J documents that this rent-based lock-in effect primarily affects the migration responses of young renters and low-wealth households, as expected.

Overall, we find that declines in both house prices and rents substantially dampen migration responses to local shocks, together reducing migration by approximately 30 percent. While both mechanisms generate lock-in effects, they operate through distinct channels and have opposite welfare implications across households, as documented below.

Welfare Effects We measure the welfare impact of an unexpected shock as the permanent percentage change in consumption-equivalent units of the stationary equilibrium value function. As shown in Appendix I, the welfare effects for agents belonging to group g can be written as:

$$\Delta W_g = \frac{\sum_{i \in g_t} V_{i,t+1}(X^S)^{\frac{1}{1-\sigma}}}{\sum_{i \in g_t} V_{i,t+1}(X^E)^{\frac{1}{1-\sigma}}} - 1 \quad (22)$$

where $V_i(X^E)$ denotes agent i 's value function evaluated at the stationary-equilibrium aggregate state, X^E , and $V_i(X^S)$ denotes the value function evaluated at time $t + 1$ along the aggregate state path, X^S , induced by the unexpected shock that occurs at time $t + 1$. Our objective is to quantify the welfare impact of the shock on a subset of households residing in a given city before the shock. Because migration is endogenous, the composition of households within a city may change following the shock. To ensure we compare the same set of households before and after the shock, we define the group g_t as the set of households residing in a given city immediately before the shock and compute welfare for this fixed group at time $t + 1$, regardless of their mobility decisions between t and $t + 1$.

Table 3 reports the welfare effects of the oil crisis across demographic groups. Columns (1) and (2) present results for oil-producing cities and for Canada as a whole, respectively. The local downturn induced by the oil price collapse leads to an average decline in welfare of 1.52 percent in oil-producing cities and 0.23 percent at the national level. Within oil-producing cities, welfare losses are concentrated among working-age households. Welfare declines by 1.58 percent for individuals aged 25–34 and by 1.69 percent for those aged 35–44, while older cohorts experience smaller losses. Although retirees do not face direct income losses, their welfare nonetheless declines, albeit by less than that of any other age group.

Welfare losses also differ markedly by homeownership status. Homeowners experience substantially larger welfare declines than renters, with average losses of 1.76 percent and 1.17 percent, respectively. Across the wealth distribution, welfare losses are largest for households in the middle of the distribution, amounting to 1.61 and 1.78 percent for the second and third net-worth quartiles, respectively. By contrast, households in the lowest wealth quartile experience the smallest welfare losses. At the national level, the aggregate welfare impact is more muted but exhibits a broadly similar pattern across demographic groups.

These findings may appear surprising at first glance. Despite having more limited access to financial insurance and fewer resources to smooth shocks, younger, renter, and low-wealth households do not experience the largest welfare losses. Instead, welfare losses are concentrated among older, working-age homeowners with relatively high wealth. Overall, the results indicate that local downturns of this type disproportionately affect households with greater exposure to housing markets, highlighting the importance of housing-related mechanisms. Columns (3) and (4) report welfare outcomes under the constant house prices and rents (CHP) scenario.⁴⁵

Aggregate welfare in oil-producing cities would have been 14.5 percent higher if house prices and rents had remained fixed. In other words, the endogenous decline in housing prices and rents amplifies the negative welfare impact of the productivity shock in the baseline economy. This aggregate effect, however, masks substantial heterogeneity across households. The decline in house prices is detrimental for homeowners but beneficial for renters in affected locations. Quantitatively, housing price dynamics account for approximately 45 percent of the homeowners' welfare loss observed in the baseline economy. For renters, the decline in rents attenuates welfare losses by 52 percent. Absent the fall in housing costs, renters' welfare losses would have increased from 1.17 to 1.78 percent, on average. For homeowners, fixing house prices would have reduced welfare losses from 1.76 percent to 0.97 percent.

These opposing effects are also reflected across age and wealth groups. Younger and low-wealth households—who are predominantly renters—would have been worse off had housing costs not adjusted downward. In contrast, welfare losses associated with declining house prices are strictly increasing in age and net worth, two dimensions that are positively correlated with housing equity. For households aged 45–64, preventing the decline in house prices would have mitigated welfare losses by more than 70 percent. Similarly, for households in the top wealth quartile (after controlling for age composition), fixed house prices would have reduced welfare losses by 63 percent.

⁴⁵For ease of interpretation, we report the percentage difference relative to the baseline economy, computed as $100 \times (W^{CHP} - W^{baseline})/|W^{baseline}|$. Positive values indicate smaller welfare losses under CHP, whereas negative values indicate larger losses relative to the baseline.

Table 3: The Welfare Effect of a Local Downturn

Demographics	Baseline		Constant House Prices & Rents	
	Oil Region	Canada	Oil Region	Canada
All	-1.52	-0.23	14.5	8.7
Homeowners	-1.76	-0.27	44.9	40.7
Renters	-1.17	-0.18	-52.1	-44.4
Age 25–34	-1.58	-0.27	-22.2	-22.2
Age 35–44	-1.69	-0.23	37.9	34.8
Age 45–64	-1.26	-0.16	71.4	68.8
Age > 64	-0.93	-0.10	98.9	90.0
Net worth – Q1	-1.31	-0.21	-38.2	-28.6
Net worth – Q2	-1.61	-0.23	10.6	4.3
Net worth – Q3	-1.78	-0.25	36.5	28.0
Net worth – Q4	-1.39	-0.24	63.3	58.3

Note: Table 3 reports the welfare changes of the decline in productivity generated by the oil crisis. The first two columns report the welfare changes in the baseline economy, with the first column for the oil regions and the second for Canada overall. The third and fourth columns report the percentage change in welfare in the constant house price scenario relative to the baseline, computed as $100 \times (W^{\text{CHP}} - W^{\text{baseline}}) / |W^{\text{baseline}}|$. Positive numbers, therefore, indicate smaller welfare losses under constant house prices, while negative numbers indicate larger losses. We estimate welfare changes using the expression in equation (22).

Taken together, declines in house prices and rents generate a lock-in effect that dampens migration responses. Still, the welfare implications differ sharply depending on the source of the lock-in. In renter-only frameworks, such as Glaeser and Gyourko (2005), declining rents generate a welfare-enhancing lock-in effect by lowering housing costs and reducing the need to relocate to higher-income locations while avoiding moving costs. However, these frameworks abstract from homeownership and therefore miss an additional, quantitatively important channel. In our model, declines in house prices also generate a negative *housing-wealth lock-in* effect that discourages mobility and is welfare-reducing. On aggregate, this housing-wealth channel dominates, reversing the positive welfare interpretation of lock-in effects emphasized in the existing literature.

5.3 Other Sources of Mitigation of Migration Response

While *Housing-wealth lock-in* is the novel highlighted mechanism, our framework also features several additional forces that attenuate migration responses to local shocks, especially for homeowners. One such force is the housing transaction cost. Because our model explicitly

incorporates housing as an illiquid asset within a spatial equilibrium, it is particularly well-suited to assess the quantitative relevance of this channel. Finally, we examine how the persistence of the shock affects migration responses at the end of this section.

No Housing Transaction Costs. In the baseline calibration, housing transaction costs amount to 7 percent of the house value.⁴⁶ While this is a sizable cost, Section 4.5 shows that housing transaction costs play only a limited role in response to idiosyncratic shocks. Yet, the relevance of Housing Transaction Costs may nevertheless differ in the context of a local aggregate shock with general equilibrium responses. On the one hand, transaction costs amplify the effective wealth loss associated with selling a house when prices are depressed. On the other hand, because house prices fall following the shock, the nominal value of transaction costs also declines. Figure 10 reports the aggregate migration response and the migration response of homeowners in a counterfactual scenario in which housing transaction costs are eliminated in oil-producing cities for the duration of the shock.⁴⁷ Relative to the baseline, the migration response of homeowners is substantially larger on impact and more persistent when transaction costs are removed. By contrast, the migration response of renters is essentially unchanged.

These results indicate that housing transaction costs are a key determinant of the muted migration response among homeowners following local shocks. In contrast, the welfare gains of eliminating housing transaction costs accrue to young households and renters as reported in Table J.1. Despite experiencing lower current income, forward-looking agents anticipate that house prices will recover after the shock. The combination of temporarily low house prices and zero transaction costs makes the transition from renting to homeownership in affected locations more attractive. Expected future house price appreciation generates a positive wealth effect that is already reflected in welfare at the time the shock occurs. Current homeowners also benefit from lower transaction costs, but only if they adjust their housing consumption or relocate. Naturally, such adjustments occur only when they are welfare improving; however, high moving costs limit the magnitude of the resulting welfare gains for homeowners.

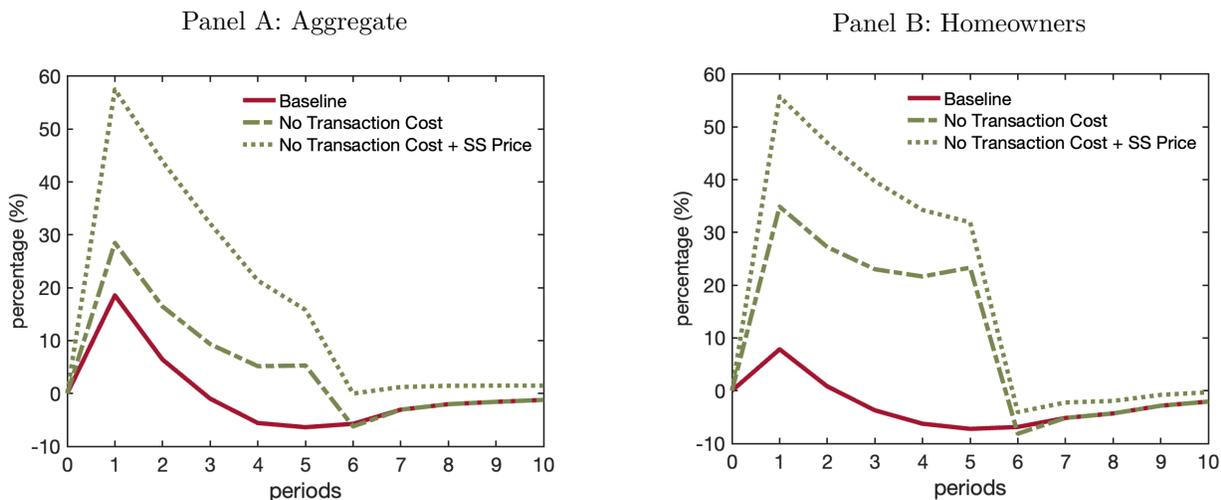
Interestingly, there is a modest interaction between housing transaction costs and the housing-wealth lock-in effect. The magnitude of the lock-in effect is slightly larger in economies without transaction costs, although the difference is not quantitatively large.⁴⁸

⁴⁶On average, this corresponds to approximately 60 percent of median annual income in Canada. This is a sizable cost and could plausibly deter reallocation among homeowners.

⁴⁷In the baseline economy, exogenous productivity in oil-producing cities remains below its pre-shock level for six periods. In the *No Transaction Costs* counterfactual, we set the housing transaction cost parameter $\delta_h = 0$ over the same six periods in oil-producing cities, while imposing the same paths for wages, house prices, and rents as in the baseline economy.

⁴⁸Panel B of Figure 10 shows that, on impact, holding house prices constant and eliminating transaction

Figure 10: No Housing Transaction Costs: Migration Response



Note: Figure 10 reports the migration response to the oil-induced local downturn separately for renters (Panel A) and renters (Panel B) in oil cities. In solid lines we report the values for the out-migration shock response in the baseline model in red for oil cities and in blue in the counterfactual scenario in which housing transaction costs are paid if homeowners move outside oil cities. The dashed line reports the respective values for the case in which house prices and rents are fixed to what they were before the shock.

Permanent Shock. The temporary nature of the shock is crucial for understanding the muted migration response in the baseline scenario. Figure J.9 shows the migration response to a permanent decline in productivity (Figure J.8). Relative to the temporary shock, a permanent shock generates an out-migration response on impact that is approximately 2 times larger and remains about 18 percent higher in the new long-run equilibrium compared with the pre-shock steady state. On impact, these differences are driven primarily by young, renter, and low-wealth households. For wealthy homeowners, two forces are at play. On the one hand, because households do not expect house prices to recover following permanent shocks, the incentive to remain in place to avoid realizing housing wealth losses is reduced. On the other hand, permanent shocks generate much larger declines in house prices (about 15 percent), implying greater housing wealth losses upon migration. Quantitatively, relative to the temporary shock, we find that homeowners' migration responses increase slightly on impact when the house price decline is large and permanent. For completeness, Appendix J.2 provides the full set of results for different demographic groups.

costs raises homeowners' migration rates by approximately 21 percentage points relative to the case with no transaction costs but declining house prices. When house prices and rents are held constant but transaction costs remain positive, the corresponding gap relative to the baseline is about 18 percentage points, as shown in Panel B of Figure 9. Equivalently, removing housing transaction costs has a stronger effect on migration when house prices are fixed than when they adjust endogenously.

6 Concluding Remarks

This paper develops a dynamic life-cycle incomplete-markets model with homeownership embedded in a spatial equilibrium framework. By jointly modeling migration and wealth accumulation in both liquid and illiquid assets, the framework provides a unified theory of how households adjust to local shocks and identifies a new channel that exacerbates the limited migration response to local downturns: the *housing-wealth lock-in* effect. Quantitatively, the model replicates the muted migration response to the oil-induced downturn in Canada and shows that falling house prices generate a powerful *housing-wealth lock-in* effect that substantially dampens relocation, particularly for homeowners. More importantly, declines in house prices and rents amplify welfare losses for homeowners while attenuating losses for renters, while keeping them both geographically locked in. Consequently, policies that lower rents may reduce migration while mitigating renter welfare losses, whereas stabilizing house prices can facilitate homeowner reallocation and raise overall welfare.

Beyond these findings, the model provides a flexible platform to study how local shocks propagate across space. Cross-regional differences in industrial composition imply heterogeneous exposure to technological change, trade shocks, and the transition toward a greener economy. The framework can quantify the resulting reallocation of economic activity, disentangle adjustment through mobility versus balance-sheet responses, and assess welfare consequences—an aspect particularly relevant for decarbonization policies with highly uneven spatial incidence.

A second extension concerns the interaction between migration, wealth accumulation, and the housing affordability crisis. Sustained house price growth, higher rent burdens, and the increasing presence of institutional buyers reshape the distribution of housing wealth and tighten liquidity constraints for younger and low-wealth households, potentially amplifying the *housing-wealth lock-in* effects highlighted here. Incorporating investor demand into the housing sector would allow us to quantify how investor-driven appreciation affects mobility, local adjustment to shocks, and long-run inequality, as well as how these forces reshape the geography of opportunity.

Several avenues for future research follow naturally. More broadly, the framework provides a foundation for a quantitative agenda on the forces shaping the geography of opportunity under structural change. Incorporating demographic dynamics would enable the analysis of how aging and rural depopulation interact with local and aggregate policies, pension reforms, and regionally targeted transfers.

References

- Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf.** 2015. “The Economics of Density: Evidence From the Berlin Wall.” *Econometrica*, 83(6): 2127–2189.
- Aiyagari, S Rao.** 1994. “Uninsured idiosyncratic risk and aggregate saving.” *The Quarterly Journal of Economics*, 109(3): 659–684.
- Andersson, Fredrik, and Tom Mayock.** 2014. “How does home equity affect mobility?” *Journal of Urban Economics*, 84: 23–39.
- Auclert, Adrien, Will S Dobbie, and Paul Goldsmith-Pinkham.** 2019. “Macroeconomic effects of debt relief: Consumer bankruptcy protections in the great recession.” National Bureau of Economic Research.
- Beer, Rachael, Felicia Ionescu, and Geng Li.** 2018. “Are Income and Credit Scores Highly Correlated?” *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System.
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra.** 2018. “House prices and consumer spending.” *The Review of Economic Studies*, 85(3): 1502–1542.
- Bilal, Adrien, and Esteban Rossi-Hansberg.** 2021. “Location as an Asset.” *Econometrica*, 89: 2459–2495.
- Birinci, Serdar, Youngmin Park, and Kurt See.** 2023. “The Heterogeneous Impacts of Job Displacement: Evidence from Canadian Job Separation Records.” , (23-55).
- Cai, Sheng, Lorenzo Caliendo, Fernando Parro, and Wei Xiang.** 2022. “Mechanics of Spatial Growth.” National Bureau of Economic Research.
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro.** 2019. “Trade and labor market dynamics: General equilibrium analysis of the china trade shock.” *Econometrica*, 87(3): 741–835.
- Dao, Mai, Davide Furceri, and Prakash Loungani.** 2017. “Regional Labor Market Adjustment in the United States: Trend and Cycle.” *The Review of Economics and Statistics*, 99(2): 243–257.
- De Nardi, Mariacristina.** 2004. “Wealth inequality and intergenerational links.” *The Review of Economic Studies*, 71(3): 743–768.

- Desmet, Klaus, and Esteban Rossi-Hansberg.** 2014. “Spatial Development.” *American Economic Review*, 104(4): 1211–43.
- Desmet, Klaus, Dávid Krisztián Nagy, and Esteban Rossi-Hansberg.** 2018. “The Geography of Development.” *Journal of Political Economy*, 126(3): 903–983.
- Diamond, Rebecca, Tim McQuade, and Franklin Qian.** 2019. “The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco.” *American Economic Review*, 109(9): 3365–94.
- Dvorkin, M.** 2023. “Heterogeneous Agents Dynamic Spatial General Equilibrium.” *Federal Reserve Bank of St. Louis Working Paper*, 005.
- Eckert, Fabian, and Tatjana Kleineberg.** 2019. “Can we save the American dream? A dynamic general equilibrium analysis of the effects of school financing on local opportunities.” *Unpublished Manuscript, Yale University*.
- Favilukis, Jack, Pierre Mabilie, and Stijn Van Nieuwerburgh.** 2023. “Affordable housing and city welfare.” *The Review of Economic Studies*, 90(1): 293–330.
- Ferreira, Fernando, Joseph Gyourko, and Joseph Tracy.** 2010. “Housing busts and household mobility.” *Journal of Urban Economics*, 68(1): 34–45.
- Giannone, Elisa.** 2017. “Skilled-biased technical change and regional convergence.” In *Technical Report*. University of Chicago Working Paper.
- Glaeser, Edward L., and Joseph Gyourko.** 2005. “Urban Decline and Durable Housing.” *Journal of Political Economy*, 113(2): 345–375.
- Greaney, Brian.** 2020. “The distributional effects of uneven regional growth.” *University of Washington, mimeograph*. [6].
- Guerrieri, Veronica, and Guido Lorenzoni.** 2017. “Credit Crises, Precautionary Savings, and the Liquidity Trap.” *The Quarterly Journal of Economics*, 132(3): 1427–1467.
- Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson.** 2021. “Housing wealth effects: The long view.” *The Review of Economic Studies*, 88(2): 669–707.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante.** 2009. “Quantitative macroeconomics with heterogeneous households.” *Annual Review of Economics*, 1(1): 319–354.

- Huggett, Mark.** 1993. “The risk-free rate in heterogeneous-agent incomplete-insurance economies.” *Journal of Economic Dynamics and Control*, 17(5-6): 953–969.
- Imrohoroglu, Ayse.** 1989. “Cost of business cycles with indivisibilities and liquidity constraints.” *Journal of Political Economy*, 97(6): 1364–1383.
- Kaplan, Greg, and Giovanni L Violante.** 2014. “A model of the consumption response to fiscal stimulus payments.” *Econometrica*, 82(4): 1199–1239.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner.** 2014. “The Wealthy Hand-to-Mouth.” *Brookings Papers on Economic Activity*, 45(1 (Spring)): 77–153.
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante.** 2020a. “Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data.” *Journal of Public Economics*, 189: 104176.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante.** 2020b. “The Housing Boom and Bust: Model Meets Evidence.” *Journal of Political Economy*, 128(9): 3285–3345.
- Kennan, John, and James R Walker.** 2011. “The effect of expected income on individual migration decisions.” *Econometrica*, 79(1): 211–251.
- Kilian, Lutz, and Xiaoqing Zhou.** 2022. “The propagation of regional shocks in housing markets: Evidence from oil price shocks in Canada.” *Journal of money, credit and banking*, 54(4): 953–987.
- Kleinman, Benny, Ernest Liu, and Stephen J Redding.** 2023. “Dynamic spatial general equilibrium.” *Econometrica*, 91(2): 385–424.
- Kothari, Siddharth, Itay Saporta-Eksten, and Edison Yu.** 2013. “The (un) importance of geographical mobility in the Great Recession.” *Review of Economic Dynamics*, 16(3): 553–563.
- Lyon, Spencer, and Michael E Waugh.** 2018. “Quantifying the losses from international trade.” *Unpublished manuscript, NYU Stern*.
- McFadden, D.** 1973. “Conditional Logit Analysis of Qualitative Choice Behaviour.” In *Frontiers in Econometrics*, ed. P. Zarembka, 105–142. New York, NY, USA: Academic Press New York.
- Mitze, Timo.** 2019. “The migration response to local labour market shocks: Evidence from EU regions during the global economic crisis.” *Oxford Bulletin of Economics and Statistics*, 81(2): 271–298.

- Molloy, Raven, Christopher L. Smith, and Abigail K. Wozniak.** 2014. “Declining Migration within the U.S.: The Role of the Labor Market.” *National Bureau of Economic Research*, 20065.
- Monras, Joan.** 2018. “Economic shocks and internal migration.”
- Oswald, Florian.** 2019. “The effect of homeownership on the option value of regional migration.” *Quantitative Economics*, 10(4): 1453–1493.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel.** 2007. “Housing, consumption and asset pricing.” *Journal of Financial Economics*, 83(3): 531–569.
- Redding, Stephen, and Esteban Rossi-Hansberg.** 2017. “Quantitative Spatial Economics.” *Annual Review of Economics*, 9(1): 21–58.
- Saiz, Albert.** 2010. “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics*, 125(3): 1253–1296.
- Schulhofer-Wohl, Sam.** 2011. “Negative Equity Does Not Reduce Homeowners’ Mobility.” *Quarterly Journal of Economics*, 126(2): 1117–1160.
- Yagan, Danny.** 2019. “Employment Hysteresis from the Great Recession.” *Journal of Political Economy*, 109(3): 1117–1153.

Appendix

A	Additional Analysis on the Empirical Evidence	A4
A.1	Migration Patterns in Canada: Census vs TransUnion	A4
A.1	Heterogeneous Migration Responses (Credit Score)	A5
A.2	Heterogeneous Migration Responses (Credit Usage - excludes Mortgage Debt)	A6
A.3	Heterogeneous Migration Responses (Credit Usage - includes Mortgage Debt)	A7
B	Migration and Wealth in the US	B8
B.1	Data	B8
B.2	Empirical Strategy	B8
B.3	Results	B9
B.1	The Determinants of Migration in the US	B10
B.2	The Determinants of Migration in the US with Demographic Controls	B11
B.4	Robustness	B11
B.1	The Determinants of Migration Propensity – Total Wealth	B13
B.2	The Determinants of Migration Propensity – Liquid Wealth	B14
C	Equilibrium	C15
D	Solution Algorithm in Detail	D16
D.1	Stationary Equilibrium	D16
D.2	Transition Path	D17
E	Housing Supply Elasticities Estimation	E19
E.1	Housing Supply Elasticities Estimation	E19
E.1.1	Empirical strategy	E20
E.1	House prices in Vancouver, Winnipeg and the West region	E21
E.1.2	Data	E22
E.1.3	Results	E23
E.2	Housing supply elasticities for Canadian census metropolitan areas	E23
F	Additional Analysis on Model Parameterization	F24
F.1	More on Parametrization Procedure	F24
F.2	Parameterization of the Migration Elasticity and Sensitivity Analysis	F24
F.1	Estimation of the Migration Elasticity	F25
F.1	Sensitivity of Migration Elasticity by Demographic Groups	F25
F.3	Extra Figures and Tables	F27
F.2	Parameter Values	F27

F.2	TFP, Amenity, Utility of homeownership, and Average Income	F28
F.3	Model vs Data: Population and Wages across cities	F29
F.4	Income in the model and the data	F29
F.5	Model vs Data: Homeownership Rates and House Prices across Cities	F30
F.6	Wealth Distribution by Homeownership Status	F30
F.3	Correlations between City Characteristics	F31
F.7	Migration rates by Homeownership and Age; Population Distribution and In- Migration	F32
G	Moving Costs	G33
G.1	Derivation of Moving Costs	G33
G.2	The Size of Moving Costs	G34
G.1	Decomposition of Moving Costs	G35
G.3	Moving Cost: Monetary vs Utility	G36
G.1	Monetary vs Non-monetary Moving Costs	G37
H	Identifying a Local Downturn in the Data	H39
H.1	Background on the 2014–2016 Oil Price Collapse	H39
H.1	Oil Price Evolution	H39
H.2	Oil Price Collapse as a Source of Local Downturn	H40
H.2	Impact of Oil Price Collapse on Income	H41
H.1	Impact of Oil Crisis on Income	H42
H.2	Impact of Oil Crisis on Outflows and House Prices	H42
I	Derivation of Welfare Measure	I43
J	Additional Results on the Impact of a Local Shock	J45
J.1	Local Productivity Shocks, Wages, Income and House Prices	J45
J.2	Out-Migration Response to the Local Downturn	J46
J.3	Out-Migration Response to the Local Downturn by Housing Tenure	J46
J.4	Out-Migration Response to the Local Downturn by Age Group	J47
J.5	Out-Migration Response to the Local Downturn by Net Worth Group	J48
J.1	Lowering the Housing Transaction Costs	J49
J.6	Out-Migration Response to the Local Downturn by Age Group - Lowering the Housing Transaction Costs	J49
J.7	Out-Migration Response to the Local Downturn by Net Worth Group - Lowering the Housing Transaction Costs	J50
J.1	The Welfare Effect of a Local Downturn – Lowering the Housing Transaction Costs	J51

J.2 A Permanent Version of the Local Downturn J52

J.8 Local Permanent Productivity Shocks, Wages, Income and House Prices J52

J.9 Out-Migration Response to a Permanent Local Downturn J53

J.10 Migration Response to a Permanent Local Downturn J53

J.11 Out-Migration Response to a Permanent Local Downturn by Age Group J54

J.12 Out-Migration Response to a Permanent Local Downturn by Net Worth Group . J55

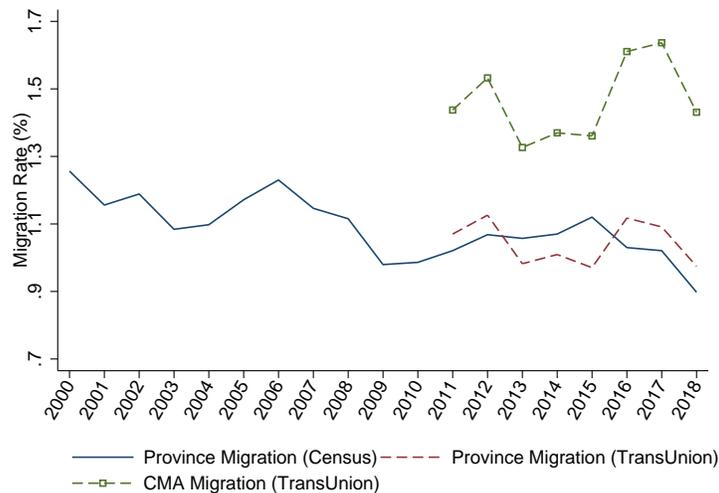
J.2 The Welfare Effect of a Permanent Local Downturn J56

A Additional Analysis on the Empirical Evidence

Definitions

- *Migrants*: All the individuals in our dataset that report living in a different location than the one in the previous period. Locations in our main specifications are CAs and CMAs in the robustness checks.
- *Homeowners*: All the individuals identified in TransUnion with an active mortgage with positive outstanding or a home-equity line of credit above 50,000 CAD or had a fully-amortized mortgage associated with the current address.
- *Credit Usage*: We proxy for borrowing tightness using a credit-utilization measure constructed from TransUnion account-level data. In our baseline specification, credit usage is the fraction of *non-mortgage* credit limit that is currently drawn, computed by aggregating outstanding balances and available limits across open credit cards, installment and auto loans, and lines of credit including home-equity lines of credit (HELOCs). Because mortgage balances are not easily adjustable in Canada (cash-out refinancing is uncommon and refinancing typically requires re-qualification), we exclude mortgages from this utilization measure.

Figure A.1: Migration Patterns in Canada: Census vs TransUnion



Note: Figure A.1 plots the yearly inter-provincial migration rate using Census data (solid blue line) between 2000 and 2018 and using TransUnion data (red dashed line) between 2011 and 2018. The green dashed-squared line plots the yearly migration rates among Canadian CMAs using TransUnion data between 2011 and 2018. *Data Sources*: StatCan and TransUnion.

Table A.1: Heterogeneous Migration Responses (Credit Score)

Panel A: Migration across CAs								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Move=100							
Homeowner	-0.765*** (0.108)			-0.577*** (0.140)	-0.764*** (0.108)			-0.576*** (0.140)
Age [36-45]		-2.057*** (0.221)		-1.925*** (0.203)		-2.056*** (0.221)		-1.924*** (0.203)
Age [46-65]		-2.780*** (0.294)		-2.589*** (0.270)		-2.781*** (0.294)		-2.589*** (0.270)
Age [66-75]		-3.285*** (0.359)		-3.098*** (0.341)		-3.284*** (0.358)		-3.097*** (0.341)
Age [76-85]		-3.621*** (0.380)		-3.504*** (0.375)		-3.620*** (0.380)		-3.503*** (0.375)
Credit Score [640-759]			-1.000*** (0.161)	-0.691*** (0.107)			-1.002*** (0.160)	-0.693*** (0.107)
Credit Score [760-799]			-1.256*** (0.187)	-0.798*** (0.113)			-1.259*** (0.187)	-0.800*** (0.113)
Credit Score [800-900]			-1.873*** (0.231)	-1.058*** (0.124)			-1.874*** (0.231)	-1.059*** (0.125)
Observations	150143346	150143346	150143346	150143346	150143346	150143346	150143346	150143346
Adjusted R^2	0.006	0.011	0.007	0.012	0.006	0.012	0.007	0.012
City Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
Year Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
City \times Year Fixed-Effects	No	No	No	No	Yes	Yes	Yes	Yes

Panel B: Migration across CMAs								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Move=100							
Homeowner	-0.596*** (0.085)			-0.457*** (0.107)	-0.595*** (0.085)			-0.456*** (0.107)
Age [36-45]		-1.486*** (0.171)		-1.389*** (0.159)		-1.485*** (0.171)		-1.389*** (0.159)
Age [46-65]		-2.039*** (0.214)		-1.916*** (0.199)		-2.039*** (0.214)		-1.916*** (0.199)
Age [66-75]		-2.350*** (0.254)		-2.260*** (0.246)		-2.350*** (0.254)		-2.259*** (0.246)
Age [76-85]		-2.541*** (0.275)		-2.514*** (0.275)		-2.540*** (0.275)		-2.513*** (0.275)
Credit Score [640-759]			-0.396*** (0.062)	-0.192*** (0.041)			-0.397*** (0.062)	-0.193*** (0.041)
Credit Score [760-799]			-0.523*** (0.071)	-0.213*** (0.055)			-0.525*** (0.071)	-0.215*** (0.055)
Credit Score [800-900]			-0.997*** (0.112)	-0.405*** (0.074)			-0.998*** (0.112)	-0.407*** (0.074)
Observations	124591109	124591109	124591109	124591109	124591109	124591109	124591109	124591109
Adjusted R^2	0.003	0.008	0.003	0.008	0.003	0.008	0.003	0.008
City Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
Year Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
City \times Year Fixed-Effects	No	No	No	No	Yes	Yes	Yes	Yes

Note: Table H.2 reports the OLS estimates of equation 1 for 2011–2019 period using Credit Score as a proxy for financial access. The dependent variable is a dummy variable that equals 100 in case of moving and zero otherwise. Regressions include home-equity and delinquency status as controls. The sample is restricted to individuals in CAs in Panel A and CMAs in Panel B. Standard errors are presented in parentheses and clustered at the city level. The ***, **, and * represent statistical significance at the 0.001, 0.01, and 0.05 levels, respectively. *Data Source:* TransUnion.

Table A.2: Heterogeneous Migration Responses (Credit Usage - excludes Mortgage Debt)

Panel A: Migration across CAs								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Move=100							
Homeowner	-0.765*** (0.108)			-0.881*** (0.140)	-0.764*** (0.108)			-0.880*** (0.140)
Age [36-45]		-2.057*** (0.221)		-1.943*** (0.211)		-2.056*** (0.221)		-1.941*** (0.211)
Age [46-65]		-2.780*** (0.294)		-2.591*** (0.281)		-2.781*** (0.294)		-2.591*** (0.281)
Age [66-75]		-3.285*** (0.359)		-3.063*** (0.355)		-3.284*** (0.358)		-3.061*** (0.354)
Age [76-85]		-3.621*** (0.380)		-3.372*** (0.381)		-3.620*** (0.380)		-3.370*** (0.381)
Credit Use - Qt2			0.402*** (0.032)	0.411*** (0.031)			0.402*** (0.032)	0.411*** (0.031)
Credit Use - Qt3			0.698*** (0.071)	0.573*** (0.057)			0.698*** (0.071)	0.573*** (0.057)
Credit Use - Qt4			1.043*** (0.116)	0.770*** (0.093)			1.044*** (0.116)	0.770*** (0.093)
Credit Use - Qt5			1.466*** (0.148)	0.930*** (0.108)			1.468*** (0.148)	0.932*** (0.108)
Observations	150143346	150143346	134704719	134704719	150143346	150143346	134704719	134704719
Adjusted R^2	0.006	0.011	0.006	0.012	0.006	0.012	0.006	0.012
City Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
Year Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
City \times Year Fixed-Effects	No	No	No	No	Yes	Yes	Yes	Yes

Panel B: Migration across CMAs								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Move=100							
Homeowner	-0.596*** (0.085)			-0.636*** (0.096)	-0.595*** (0.085)			-0.636*** (0.096)
Age [36-45]		-1.486*** (0.171)		-1.416*** (0.165)		-1.485*** (0.171)		-1.415*** (0.165)
Age [46-65]		-2.039*** (0.214)		-1.944*** (0.207)		-2.039*** (0.214)		-1.944*** (0.207)
Age [66-75]		-2.350*** (0.254)		-2.272*** (0.255)		-2.350*** (0.254)		-2.271*** (0.255)
Age [76-85]		-2.541*** (0.275)		-2.479*** (0.280)		-2.540*** (0.275)		-2.478*** (0.280)
Credit Use - Qt2			0.253*** (0.030)	0.277*** (0.027)			0.253*** (0.030)	0.277*** (0.027)
Credit Use - Qt3			0.426*** (0.070)	0.349*** (0.061)			0.426*** (0.070)	0.349*** (0.061)
Credit Use - Qt4			0.609*** (0.118)	0.428*** (0.105)			0.609*** (0.118)	0.428*** (0.105)
Credit Use - Qt5			0.828*** (0.129)	0.458*** (0.113)			0.829*** (0.129)	0.459*** (0.113)
Observations	124591109	124591109	111876964	111876964	124591109	124591109	111876964	111876964
Adjusted R^2	0.003	0.008	0.003	0.008	0.003	0.008	0.003	0.008
City Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
Year Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
City \times Year Fixed-Effects	No	No	No	No	Yes	Yes	Yes	Yes

Note: Table A.2 reports the OLS estimates of equation 1 for 2011–2019 period using Credit (excludes Mortgage Debt) Usage as a proxy for financial access. The dependent variable is a dummy variable that equals 100 in case of moving and zero otherwise. Regressions include home-equity and delinquency status as controls. The sample is restricted to individuals in CAs in Panel A and CMAs in Panel B. Standard errors are presented in parentheses and clustered at the city level. The ***, **, and * represent statistical significance at the 0.001, 0.01, and 0.05 levels, respectively. *Data Source:* TransUnion.

Table A.3: Heterogeneous Migration Responses (Credit Usage - includes Mortgage Debt)

Panel A: Migration across CAs								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Move=100							
Homeowner	-0.765*** (0.108)			-1.047*** (0.132)	-0.764*** (0.108)			-1.046*** (0.132)
Age [36-45]		-2.057*** (0.221)		-1.929*** (0.207)		-2.056*** (0.221)		-1.928*** (0.207)
Age [46-65]		-2.780*** (0.294)		-2.576*** (0.277)		-2.781*** (0.294)		-2.576*** (0.277)
Age [66-75]		-3.285*** (0.359)		-3.057*** (0.352)		-3.284*** (0.358)		-3.055*** (0.351)
Age [76-85]		-3.621*** (0.380)		-3.376*** (0.378)		-3.620*** (0.380)		-3.375*** (0.378)
Credit Use - Qt2			0.617*** (0.050)	0.486*** (0.037)			0.618*** (0.050)	0.487*** (0.037)
Credit Use - Qt3			0.736*** (0.074)	0.650*** (0.061)			0.736*** (0.074)	0.650*** (0.061)
Credit Use - Qt4			0.831*** (0.094)	0.791*** (0.086)			0.832*** (0.094)	0.792*** (0.086)
Credit Use - Qt5			1.147*** (0.128)	0.789*** (0.102)			1.148*** (0.129)	0.790*** (0.102)
Observations	150143346	150143346	137263522	137263522	150143346	150143346	137263522	137263522
Adjusted R^2	0.006	0.011	0.005	0.012	0.006	0.012	0.006	0.012
City Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
Year Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
City \times Year Fixed-Effects	No	No	No	No	Yes	Yes	Yes	Yes

Panel B: Migration across CMAs								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Move=100							
Homeowner	-0.596*** (0.085)			-0.679*** (0.080)	-0.595*** (0.085)			-0.679*** (0.080)
Age [36-45]		-1.486*** (0.171)		-1.409*** (0.163)		-1.485*** (0.171)		-1.408*** (0.163)
Age [46-65]		-2.039*** (0.214)		-1.941*** (0.206)		-2.039*** (0.214)		-1.941*** (0.206)
Age [66-75]		-2.350*** (0.254)		-2.276*** (0.255)		-2.350*** (0.254)		-2.275*** (0.254)
Age [76-85]		-2.541*** (0.275)		-2.487*** (0.279)		-2.540*** (0.275)		-2.486*** (0.279)
Credit Use - Qt2			0.408*** (0.048)	0.331*** (0.037)			0.408*** (0.048)	0.331*** (0.037)
Credit Use - Qt3			0.448*** (0.074)	0.383*** (0.068)			0.448*** (0.074)	0.383*** (0.068)
Credit Use - Qt4			0.463*** (0.096)	0.404*** (0.101)			0.464*** (0.096)	0.404*** (0.101)
Credit Use - Qt5			0.639*** (0.113)	0.341*** (0.115)			0.639*** (0.113)	0.341*** (0.115)
Observations	124591109	124591109	113953122	113953122	124591109	124591109	113953122	113953122
Adjusted R^2	0.003	0.008	0.003	0.008	0.003	0.008	0.003	0.008
City Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
Year Fixed-Effects	Yes	Yes	Yes	Yes	No	No	No	No
City \times Year Fixed-Effects	No	No	No	No	Yes	Yes	Yes	Yes

Note: Table A.3 reports the OLS estimates of equation 1 for 2011–2019 period using Credit (includes Mortgage Debt) Usage as a proxy for financial access. The dependent variable is a dummy variable that equals 100 in case of moving and zero otherwise. Regressions include home-equity and delinquency status as controls. The sample is restricted to individuals in CAs in Panel A and CMAs in Panel B. Standard errors are presented in parentheses and clustered at the city level. The ***, **, and * represent statistical significance at the 0.001, 0.01, and 0.05 levels, respectively. *Data Source:* TransUnion.

B Migration and Wealth in the US

This section revisits the patterns documented in Section 2 using U.S. longitudinal data. Our goal is twofold: (i) to assess the external validity of our findings by examining them in a different institutional and temporal context, thereby evaluating the robustness and generalizability of the results, and (ii) to exploit richer demographic information than is available in the Canadian TransUnion data. For this purpose, we rely on the National Longitudinal Survey of Youth 1997 (NLSY97), which includes a wide array of personal and household characteristics—such as education, marital status, fertility, income, sectors and wealth (liquid and illiquid)—allowing us to account for confounding factors that cannot be controlled for in our baseline analysis. Incorporating these additional variables enables us to more precisely isolate the relevant mechanisms and better understand the dynamics at play.

These advantages, however, come with some limitations. In particular, the NLSY97 does not report respondents’ exact location, and the individuals in the sample are all of similar age, which restricts our ability to analyze age-specific heterogeneity or geographic exposure with the same level of detail as in the Canadian data.

B.1 Data

The NLSY97 is a nationally representative longitudinal survey of 8,984 men and women born between 1980 and 1984, followed annually since 1997 for over 20 years. It provides rich information on labor market outcomes, including employment histories, job transitions, and income, as well as detailed measures of educational attainment and schooling experiences. The survey also includes extensive data on wealth, covering both liquid and illiquid assets, including homeownership, which facilitates the study of wealth accumulation. In addition, respondents report whether they moved in the past year and their macro region of residence. This combination of income, education, wealth, and migration measures makes the NLSY97 a valuable dataset for studying how socioeconomic conditions shape individual mobility.

B.2 Empirical Strategy

We follow the same empirical strategy as in Section 2 to assess how moving decisions vary with individual characteristics using the NLSY97. Specifically, we estimate the following linear probability model:

$$\mathbb{1}[Move_{i,z,z',t}] = \beta_0 + \beta_1 X_{i,t-1} + \beta_2 W_{i,t-1} + \delta_t + \epsilon_{i,z,t} \quad (23)$$

where $\mathbb{1}[Move_{i,z,z',t}]$ is an indicator variable taking the value 100 if individual i in city z at time t moves to a different city $z' \neq z$, and 0 otherwise, meaning that the coefficients are in units of

percentage points. $X_{i,t-1}$ are individual characteristics such as homeownership and wealth. We also include year fixed-effects to control for changes in economic conditions and region fixed-effects.¹ Since this is a panel and we are already including time fixed effects, we do not include a coefficient for age as most individuals have similar ages.

$W_{i,t-1}$ includes a rich set of time-varying demographic characteristics that are potentially important determinants—and confounders—of wealth and mobility, such as income, education, occupation, industry, marital status, and the number of children in the household.²

B.3 Results

Figure B.1 and column (1) of Table B.1 report the estimated coefficients on homeownership and wealth deciles from specification (23) without additional controls. Despite the differences discussed above, this specification is the closest analogue to the baseline results using the TransUnion Canada data. The main difference is that, in place of credit scores, we use *total* wealth and that we use regional fixed effects for only broad macro regions rather than cities.

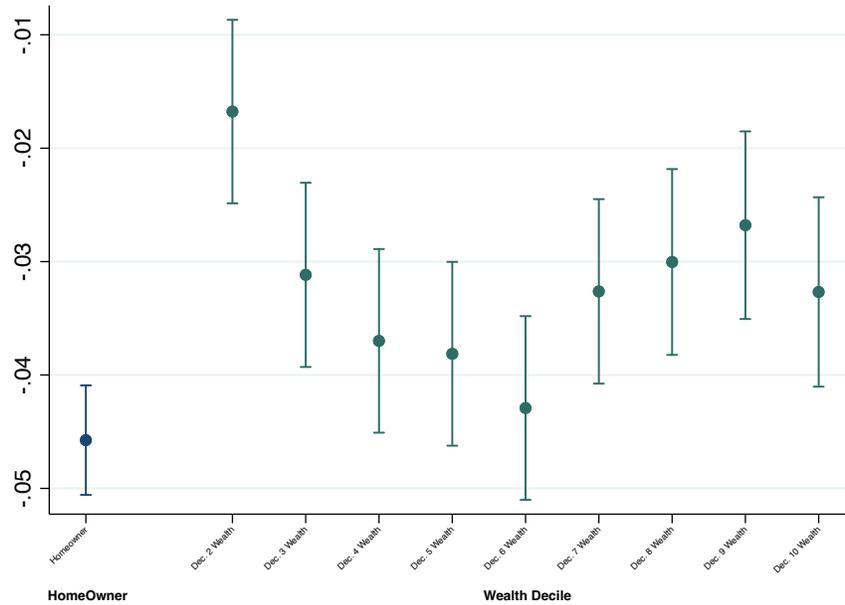
The estimates show that homeownership is negatively correlated with the likelihood of moving: homeowners are substantially less mobile than renters. Furthermore, there is a negative relationship between total *wealth* and migration. Individuals in higher wealth deciles exhibit a lower propensity to move than those in lower wealth deciles. Overall, these patterns confirm the findings from the Canadian TransUnion data on the relationship between migration and wealth determinants. Table B.2 reports the same specification as Table B.1, but isolating liquid wealth only. Our results are robust to different measures of wealth.

However, wealth may simply proxy for underlying individual characteristics that are more directly relevant for migration decisions. To verify that wealth itself plays an important role, we next control for the key demographic factors available in our dataset. Yet, in figure B.1 there could be other confounding factors that might be leading the results in that direction. NLSY data allows us to have substantially more demographic controls than TransUnion Canada. We exploit those controls and, in particular, we highlight one of the main cofounders, which is education. Figure B.2 presents the results of our analysis of migration determinants in the US using the same model outlined in Equation (23), with the inclusion of demographic controls for industry, occupation, income, and number of children, marital status, and, in particular,

¹Unlike in our specification using the TransUnion data, we cannot include city fixed effects because the NLSY97 does not report the exact city of residence. We can only include macro US regions fixed effects. For the same reason, standard errors cannot be clustered at the city level.

²Education is measured in completed years of schooling and entered in bins with cutoffs at 8, 12, 16, and 18 years, corresponding roughly to less than 8th grade, high school graduate, some college (12+4), and college degree, respectively. Industry is defined at the 4-digit NAICS level and occupation at the 4-digit occupation code level. Children denotes the number of children, and married indicates marital status. Region refers to the four macro US regions.

Figure B.1: The Determinants of Migration in the US

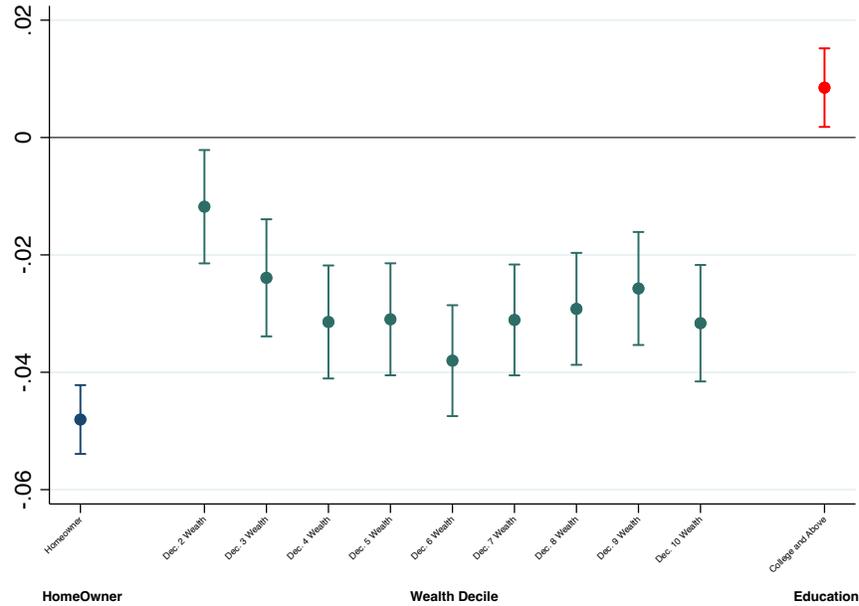


Note: Figure B.1 reports the point estimates of the linear probability model 23. The vertical bands represent 99 percent confidence intervals for the point estimates in each year. The first coefficient (in blue) reports the point estimate for the homeowner indicator variable. The second set of estimates under the umbrella “Wealth Decile” (in green) reports the estimates by wealth decile group indicator variables relative to the bottom wealth decile. *Data Source: NLSY97.*

education. The figure displays the estimated coefficients and corresponding 99 percent confidence intervals for key variables of interest: homeownership, wealth decile, and education. The results indicate that, compared to figure B.1, the relationships are similar and the coefficients change only slightly in magnitude but not at all in sign and significance. Homeownership is negatively correlated with the probability of moving, even after controlling for demographic factors. This is more than those in the bottom wealth decile, which serves to relocate compared to those who do not. Furthermore, our findings reinforce a negative relationship between wealth and migration, controlling for demographics. Individuals in higher wealth deciles exhibit a lower propensity to move than those in the bottom wealth decile, which serves as the baseline.

Figure B.1 also highlights education, one of the main confounders and important driver of wealth accumulation. Education appears to be positively correlated with migration, as individuals with a college degree or higher are more likely to move than those without.

Figure B.2: The Determinants of Migration in the US with Demographic Controls



Note: Figure B.1 reports the point estimates of the linear probability model 23 in a specification where all controls for industry, occupation, and marital status are included. The vertical bands represent 99 percent confidence intervals for the point estimates in each year. The first coefficient (in blue) reports the point estimate for the homeowner indicator variable. The second set of estimates under the umbrella “Age” (in red) reports estimates for age group indicator variables, with the 18-25 years old group as the baseline. The third set of estimates under the umbrella “Wealth Decile” (in green) reports the estimates by wealth decile group indicator variables relative to the bottom wealth decile. *Data Source: NLSY97.*

B.4 Robustness

In this section, we run robustness exercises on the results above: a) We split wealth into total vs liquid (Table B.1 vs. Table B.2); b) we run several specifications controlling for alternative cofounders additionally in column (2)-(6) of Tables B.1 and B.2. We find that independently of whether we account for total wealth or we isolate liquid wealth the results are robust to each other in all the specifications. Note again that the wealth definition used in this context combines liquid and illiquid wealth. Yet, we can run the same analysis for each component separately to corroborate the results further. Tables B.1 and B.2 report the results of this analysis as well as the coefficients of other controls such as income, number of children, and marital status. These characteristics could help explain the heterogeneity in the wealth results we find as pointed out in Molloy, Smith and Wozniak (2014). All specifications include year and macro-region fixed effects. Column (1) in each Table includes only wealth-decile indicators and a homeowner dummy (with the bottom wealth decile and renters as the omitted categories). Columns (2)-(6) then progressively add controls. In column (2) we augment the baseline specification with an indicator for having a college degree or more (“College and Above”). This variable is positively

associated with moving, but the coefficients on the wealth deciles and on homeownership are virtually unchanged. In column (3) we further add a full set of earnings–decile dummies so that wealth effects are estimated conditional on current income. Again, the coefficients on the wealth dummies and the homeowner indicator remain strongly negative and very similar in magnitude to those in columns (1) and (2). Columns (4) and (5) add detailed 4–digit occupation and 4–digit industry fixed effects, absorbing differences in mobility associated with sectoral and occupational sorting. The point estimates for wealth deciles and homeownership in these richer specifications are almost identical to those in column (3), indicating that the negative relationship between wealth and migration, as well as the lower mobility of homeowners, is not driven by differences in occupational or industry composition. Finally, column (6) adds two additional demographic controls: indicators for the presence of children in the household and for being married. Consistent with prior work, children are associated with a significantly lower probability of moving, whereas marital status is essentially unrelated to mobility; yet the coefficients on the wealth deciles and on the homeowner dummy change only slightly.

Comparing across columns (2)–(6) in both Tables, the pattern is very stable: individuals in higher (total or liquid) wealth deciles exhibit a systematically lower propensity to move relative to those in the bottom decile, and homeowners are about 4–6 percentage points less likely to move than renters. These magnitudes are robust to controlling for education, income, fertility, marital status, and detailed occupation and industry fixed effects, confirming that the negative association between wealth and mobility, and between homeownership and mobility, is not an artifact of omitted observable characteristics. The results confirm the figures and our prior knowledge that migration rates do not vary significantly with income and marital status; yet they are negatively correlated with the number of children. Overall, all these regressions are robust to the inclusion of year, industry, and occupation fixed effects.

Table B.1: The Determinants of Migration Propensity – Total Wealth

	(1)	(2)	(3)	(4)	(5)	(6)
Dec. 2 Wealth	-0.017*** (0.004)	-0.013*** (0.004)	-0.013*** (0.004)	-0.011** (0.005)	-0.011** (0.005)	-0.012** (0.005)
Dec. 3 Wealth	-0.031*** (0.004)	-0.026*** (0.004)	-0.026*** (0.004)	-0.024*** (0.005)	-0.024*** (0.005)	-0.024*** (0.005)
Dec. 4 Wealth	-0.037*** (0.004)	-0.033*** (0.004)	-0.032*** (0.004)	-0.032*** (0.005)	-0.032*** (0.005)	-0.031*** (0.005)
Dec. 5 Wealth	-0.038*** (0.004)	-0.034*** (0.004)	-0.034*** (0.004)	-0.032*** (0.005)	-0.032*** (0.005)	-0.031*** (0.005)
Dec. 6 Wealth	-0.043*** (0.004)	-0.039*** (0.004)	-0.039*** (0.004)	-0.039*** (0.005)	-0.039*** (0.005)	-0.038*** (0.005)
Dec. 7 Wealth	-0.033*** (0.004)	-0.029*** (0.004)	-0.029*** (0.004)	-0.032*** (0.005)	-0.032*** (0.005)	-0.031*** (0.005)
Dec. 8 Wealth	-0.030*** (0.004)	-0.027*** (0.004)	-0.027*** (0.004)	-0.029*** (0.005)	-0.029*** (0.005)	-0.029*** (0.005)
Dec. 9 Wealth	-0.027*** (0.004)	-0.025*** (0.004)	-0.024*** (0.004)	-0.026*** (0.005)	-0.026*** (0.005)	-0.026*** (0.005)
Dec. 10 Wealth	-0.033*** (0.004)	-0.031*** (0.004)	-0.031*** (0.004)	-0.031*** (0.005)	-0.031*** (0.005)	-0.032*** (0.005)
Homeowner	-0.046*** (0.002)	-0.047*** (0.002)	-0.047*** (0.002)	-0.053*** (0.003)	-0.053*** (0.003)	-0.048*** (0.003)
College and Above		0.017*** (0.003)	0.017*** (0.003)	0.012*** (0.003)	0.012*** (0.003)	0.008** (0.003)
Nb. Children						-0.012*** (0.001)
Married						0.001 (0.003)
Observations	81419	81419	81419	66062	66062	66062
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes	Yes	Yes
Earn. Deciles	No	No	Yes	Yes	Yes	Yes
Occ. FE	No	No	No	Yes	Yes	Yes
Ind. FE	No	No	No	Yes	Yes	Yes

Notes: The dependent variable is an indicator for moving to a different city between t and $t + 1$, multiplied by 100; coefficients are therefore in percentage points. “Dec. k Wealth” ($k = 2, \dots, 10$) denotes dummy variables for being in the k th wealth decile, with the first decile as the omitted category. “Homeowner” is an indicator equal to one if the individual owns a home. “College and Above” equals one for individuals with at least some college. *Nb. Children* is an indicator equal to one if there are children in the household before migration; *Married* equals one if the individual is married before migration. “Occ. FE” refers to 4-digit occupation fixed effects, and “Ind. FE” to 4-digit NAICS industry fixed effects. “Region” denotes four macro US regions. “Earn. Deciles” indicates whether dummies for earnings deciles are included as controls. In this table, wealth deciles are constructed for both total (illiquid + liquid) net wealth and for liquid wealth, as described in the text. Robust standard errors clustered at the city level are reported in parentheses; ***, **, and * indicate statistical significance at the 1, 5, and 10 percent levels, respectively.

Table B.2: The Determinants of Migration Propensity – Liquid Wealth

	(1)	(2)	(3)	(4)	(5)	(6)
Dec. 2 Liq. Wealth	-0.015*** (0.004)	-0.012*** (0.004)	-0.012*** (0.004)	-0.015*** (0.005)	-0.015*** (0.005)	-0.015*** (0.005)
Dec. 3 Liq. Wealth	-0.032*** (0.004)	-0.028*** (0.004)	-0.027*** (0.004)	-0.028*** (0.005)	-0.028*** (0.005)	-0.028*** (0.005)
Dec. 4 Liq. Wealth	-0.036*** (0.004)	-0.032*** (0.004)	-0.032*** (0.004)	-0.034*** (0.005)	-0.034*** (0.005)	-0.033*** (0.005)
Dec. 5 Liq. Wealth	-0.040*** (0.004)	-0.035*** (0.004)	-0.035*** (0.004)	-0.037*** (0.005)	-0.037*** (0.005)	-0.035*** (0.005)
Dec. 6 Liq. Wealth	-0.038*** (0.004)	-0.034*** (0.004)	-0.034*** (0.004)	-0.037*** (0.005)	-0.037*** (0.005)	-0.036*** (0.005)
Dec. 7 Liq. Wealth	-0.041*** (0.004)	-0.038*** (0.004)	-0.038*** (0.004)	-0.043*** (0.005)	-0.043*** (0.005)	-0.042*** (0.005)
Dec. 8 Liq. Wealth	-0.034*** (0.004)	-0.031*** (0.004)	-0.030*** (0.004)	-0.034*** (0.005)	-0.034*** (0.005)	-0.033*** (0.005)
Dec. 9 Liq. Wealth	-0.029*** (0.004)	-0.027*** (0.004)	-0.026*** (0.004)	-0.031*** (0.005)	-0.031*** (0.005)	-0.030*** (0.005)
Dec. 10 Liq. Wealth	-0.032*** (0.004)	-0.031*** (0.004)	-0.030*** (0.004)	-0.033*** (0.005)	-0.033*** (0.005)	-0.033*** (0.005)
Homeowner	-0.049*** (0.002)	-0.050*** (0.002)	-0.050*** (0.002)	-0.056*** (0.003)	-0.056*** (0.003)	-0.052*** (0.003)
College and Above		0.017*** (0.003)	0.016*** (0.003)	0.011*** (0.003)	0.011*** (0.003)	0.008** (0.003)
Nb. Children						-0.012*** (0.001)
Married						0.001 (0.003)
Observations	81419	81419	81419	66062	66062	66062
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes	Yes	Yes
Earn. Deciles	No	No	Yes	Yes	Yes	Yes
Occ. FE	No	No	No	Yes	Yes	Yes
Ind. FE	No	No	No	Yes	Yes	Yes

Notes: The dependent variable is the annual probability (in percentage points) that an individual moves to a different city between survey waves. “Dec. k Liq. Wealth” are dummy variables for liquid net-wealth decile $k = 2, \dots, 10$; the omitted category is the first liquid-wealth decile. Liquid wealth is defined as non-housing (financial) wealth net of unsecured debt. “Homeowner” is an indicator for owning the main residence before migration. “College and Above” equals one for individuals with at least some college (based on education bins defined in the text). “Nb. Children” is an indicator for having children in the household; “Married” is an indicator for being married. “Year FE” are calendar-year fixed effects; “Region” corresponds to four macro-regions; “Occ. FE” and “Ind. FE” are 4-digit occupation and 4-digit NAICS industry fixed effects, respectively. “Earn. Deciles” indicates whether controls for deciles of pre-migration earnings are included. Robust standard errors clustered at the appropriate geographic level are in parentheses. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

C Equilibrium

Given a vector of individual states $\mathbf{x}_t = (a_t, \epsilon_t, q, \bar{h}_t)$, a competitive equilibrium of the economy consists of endogenous price vectors $\{w_t^l, p_t^l, R_t^l\}_{l=1}^L$, decision rules $\{c_t^l(\mathbf{x}), h_t^l(\mathbf{x}), b_t^l(\mathbf{x}), d_t^l(\mathbf{x}), \mu_t^l(\mathbf{x})\}_{l=1}^L$, distribution of individual states $\{\lambda_t^l(\mathbf{x})\}_{l=1}^L$, and aggregate allocations for population, labor in the construction/final good sector, housing stock, housing investment and government expenditures $\{\bar{N}_t^l, N_{c,t}^l, N_{h,t}^l, H_{t-1}^l, I_{h,t}^l, G_t^l\}_{l=1}^L$ such that:

1. Given endogenous price vectors, the policy functions, $\{c_t^l(\mathbf{x}), h_t^l(\mathbf{x}), b_t^l(\mathbf{x}), d_t^l(\mathbf{x}), \mu_t^l(\mathbf{x})\}_{l=1}^L$, solve the household's problems (8)-(10). μ_t^l denote a matrix of moving probabilities $\{\mu_t^{l,k}(\mathbf{x})\}_{k=1}^L$ defined in equation (13).
2. Given endogenous price vectors, firms in the construction sector maximize profits with associated labor demand and housing investment functions $\{N_{h,t}^l, I_{h,t}^l\}_{l=1}^L$; firms in the final good sector maximizes profits with associated labor demand $\{N_{c,t}^l\}_{l=1}^L$.
3. Wages $\{w_t^l\}_{l=1}^L$ clears the labor market in all locations and satisfy equation (14), where total labor demand equates total supply of effective labor from the working-age population

$$N_{c,t}^l + N_{h,t}^l = \int_{\{x:\epsilon^e, e \geq 2; q \leq \bar{Q}\}} \exp(\epsilon\chi(\mathbf{x})) \bar{N}_t^l \lambda_t^l(\mathbf{x}) d\mathbf{x},$$

where population in each location \bar{N}_t^l is endogenously determined and consistent with the optimal individual moving decisions of survival households satisfying $\bar{N}_t^l = \sum_{j=1}^L \int_{\mathbf{x}} \mu_{t-1}^{j,l}(\mathbf{x}) \bar{N}_{t-1}^j \lambda_{t-1}^j(\mathbf{x}) d\mathbf{x} + \bar{N}_{0,t}^l$, where $N_{0,t}^l$ denote newborns in location l at time t . The population is constant and newborns are distributed across space in proportion to the mass of households of age between 25 ($q = 1$) and 35 ($q = 5$).

4. The rental markets clear at prices $\{R_t^l\}_{l=1}^L$ given by equation (18), and the equilibrium quantity of rental units in each location satisfies $H_t^{R,l} = \int_{\mathbf{x}} h_t(\mathbf{x}) \mathbb{1}[d_t^l(\mathbf{x}) = 0] \bar{N}_t^l \lambda_t^l(\mathbf{x}) d\mathbf{x}$.
5. The equilibrium house price p_t^l clears the owning housing market: $(1 - \delta_h^l) H_{t+1}^l + I_{h,t}^l - H_t^{R,l} = \int_{\mathbf{x}} h_t(\mathbf{x}) \mathbb{1}[d_t^l(\mathbf{x}) = 1] \bar{N}_t^l \lambda_t^l(\mathbf{x}) d\mathbf{x}$, where housing stock evolves according to equation (17).
6. The government budget constraint holds and the expenditures G_t are determined residually as $G_t + \sum_{l=1}^L \int_{\{x:\epsilon^1; q \leq \bar{Q}\}} y_t^l(\mathbf{x}) \bar{N}_t^l \lambda_t^l(\mathbf{x}) d\mathbf{x} + \sum_{l=1}^L \int_{\{x:q > \bar{Q}\}} y_t^l(\mathbf{x}) \bar{N}_t^l \lambda_t^l(\mathbf{x}) d\mathbf{x} = \sum_{l=1}^L \int_{\mathbf{x}} \mathcal{T}(y_t^l(\mathbf{x})) \bar{N}_t^l \lambda_t^l(\mathbf{x}) d\mathbf{x} + \sum_{l=1}^L \left\{ \psi_l p_t^l [H_{t-1}^l - H_{t-1}^{R,l}] + [p_{h,t}^l I_{h,t}^l - w_t^l N_{h,t}^l] \right\}$, where expenditures, unemployment income and pension payments are financed by income taxes, property taxes and revenues from selling new licenses to developers.
7. The aggregate state evolves according to rational expectations; the distribution of individual states evolves consistently with individual optimization.

D Solution Algorithm in Detail

The household value and policy functions are solved via backward induction starting with the final period of life. We set a discrete grid space for wealth and housing. The discrete grid for wealth is uneven with higher concentration near the borrowing limit. Following [Kaplan, Mitman and Violante \(2020\)](#), we consider two grids space for wealth. A coarser grid over which we solved for the value and policy functions and a finer grid (by a factor of three) under which we define homeowners and renters distributions. To update such distributions we interpolate the value functions and associated policy functions. Conditional on moving and housing consumption decisions, we eliminate consumption via the budget constraint and back up liquid assets through the next period wealth equation defined in renters and homeowners problem. We verify ex-post that the upper bound of the wealth grid is not binding. We discretize the AR(1) process for the idiosyncratic component of income endowments using Rouwenhorst’s method ([Rouwenhorst, 1995](#)). Taking the city-specific unemployment rate directly from the data, we build the city-specific first-order Markov chain. The distribution of new-born (age one in the model) across space is read from the data by matching the distribution of individuals with age between 22 and 25 years old across locations. They start their lives as renters and their distribution over assets is read directly from the SFS 2010 that it is assumed to be the same across locations but taking into consideration the correlation with income.

D.1 Stationary Equilibrium

To compute the stationary equilibrium, we follow the following steps.

Step 1. We start by guessing a vector of wages across locations, \mathbf{w}^0 , homeowner and renters distributions, N^H and N^R , respectively, over locations, assets, age and income shock. Given these guesses, we obtain the median income by city. We then back up the house price vector across space \mathbf{p}^0 that matches the house price index to median income ratio obtained directly from the data. Using equation 18, we obtain the rental price vector across space \mathbf{R}^0 .

Step 2. Given price vectors we solve for value functions and policy functions using backward induction. For each asset grid point, we obtain the value function for the last age group (Q) in closed-form defined by the bequest function (3). Using standard grid-search methods, we solve for the wealth, housing and tenure choice policy functions for the age groups $q < Q$ taking as given the value functions across locations for the age $q + 1$ group. For age $q < Q$, we compute expectation income shocks, mortality and location preference shocks as defined in equation (12). Given the value functions, the migration probabilities can be constructed using equation (13).

Step 3. Given the distribution of age one group, we solve forward from age one to age Q to obtain the updated distributions of homeowners and renters across space and individual state.

The distributions are computed following the transition of endogenous states given by housing, savings, homeownership status and location policy functions and exogenous states, age and income and mortality shocks. Transition of endogenous states are computed by interpolating value functions to determine the optimal discrete choice, and then interpolate the associated moving probabilities and policy functions.

Step 4. Given the updated distributions N^H and N^H , we update wages, \mathbf{w}^1 , using labor market clearing condition taking into account the exogenous local unemployment rate. House prices and rental rates are updated as defined in step 1.

Step 5. We repeat steps 2–4 until wages in all the locations converge. Given the equilibrium house price vector, we solve for housing demand in each location that by definition equates to housing supply. By inverting equation 17, we back up the local housing permits consistent with the stationary equilibrium.

D.2 Transition Path

We now present the procedure to compute transitional paths for unanticipated shocks. We assume that the shock is not anticipated in the stationary equilibrium but once it occurs, the full shock path is known by all forward-looking agents. We assume rational expectations.

To compute the transitional path after a given shock, we apply the procedure above to compute the pre-shock stationary equilibrium ($t = 0$) and the new stationary equilibrium consistent with the shock. In the new stationary equilibrium, we don't impose that house prices to median income matches the data. Instead we guess a house price vector and update such guess using local housing market clear conditions taking as given the housing-permits backed up from the pre-shock stationary equilibrium.

The economy starts with the population distribution in the pre-shock stationary equilibrium, and the shock occurs in period 1. We assume that the economy reaches the new stationary equilibrium before period T .

Step 1. We guess wage and house prices paths, $\{\mathbf{w}_t^0\}_{t=1}^T$ and $\{\mathbf{p}_t^0\}_{t=1}^T$, respectively. At period 0 and period T , wages and house prices are equal to those in the pre- and post-shock stationary equilibria, respectively. We use equation (18) to obtain the path of rental prices.

Step 2. Given guessed paths of wages, house prices and rents, we solve backward the value functions and policy functions along the path starting in period $T - 1$ since in period T , value functions and policy functions are known, given by those in the new stationary equilibrium. Migration probabilities are constructed using equation (13).

Step 3. Given the population distribution in the pre-shock stationary equilibrium, we can compute the population distribution by iterating forward from $t = 0$ to $t = T$ following the procedure defined in point 3 of Section [D.1](#) .

Step 4. Given the path of population distribution, we update wage and house prices path guesses, $\{\mathbf{w}_t^1\}_{t=1}^T$ and $\{\mathbf{p}_t^1\}_{t=1}^T$, using the labor market and housing markets clear conditions.

Step 5. We repeat procedures 2–4 until converge in wages and houses prices is obtained in all locations.

Step 6. We check whether wages and prices in period T reach the corresponding levels in the after-shock stationary equilibrium. If not, we increase T .

E Housing Supply Elasticities Estimation

We estimate the first housing supply elasticities for the largest cities in Canada (census agglomerations, CA³) following the approach developed by [Guren et al. \(2021\)](#). This note presents these estimates and the procedure used. This novel approach exploits that house prices in some cities are systematically more sensitive to regional cycles than in other cities. This approach differs from the one used by [Saiz \(2010\)](#) in estimating housing supply elasticities for most metropolitan areas in the United States. He does this by exploiting city-specific building regulations and land unavailability, specifically the land within a 50-kilometer radius of the city center unsuitable for construction due to geographic constraints such as steep slopes or bodies of water. His estimates are widely used in the economic literature in model calibrations and as an instrument for the change in house prices during the boom and bust cycle of the 2000s.

The approach developed by [Guren et al. \(2021\)](#) has two main advantages over the one of [Saiz \(2010\)](#). First, the Saiz measure correlates with other city characteristics such as productivity and growth in demand ([Davidoff, 2016](#)). This raises the concern that higher house price volatility in some cities is not driven by inelastic housing supply, as estimated by Saiz, but by differences in other characteristics such as different industrial composition and different exposure to secular trends, for example, an increase in housing demand in coastal areas with inelastic supply. To address this shortcoming, [Guren et al. \(2021\)](#) employs a panel specification that allows them to control for city-specific trends, different sensitivity to regional business cycles, and changes in the city's population and industry structure. Second, by exploiting the systematic historical sensitivity of local house prices to regional house price cycles, this new approach allows us to estimate housing supply elasticities without resorting to geographical and regulation data across Canadian cities, data that are not currently available for most cities in Canada.

E.1 Housing Supply Elasticities Estimation

We estimate the first the housing supply elasticities for the largest cities in Canada (census agglomerations, CA⁴) following the approach developed by [Guren et al. \(2021\)](#). This note presents these estimates and the procedure used. This novel approach exploits that house prices in some cities are systematically more sensitive to regional cycles than in other cities. This approach differs from the one used by [Saiz \(2010\)](#) in estimating housing supply elasticities for most metropolitan areas in the United States. He does this by exploiting city-specific building regulations and land unavailability, specifically the land within a 50-kilometer radius of the city center unsuitable for construction due to geographic constraints such as steep slopes or bodies of water. His estimates are widely used in the economic literature in model calibrations and as

³<https://www150.statcan.gc.ca/n1/pub/92-195-x/2011001/geo/cma-rmr/cma-rmr-eng.htm>

⁴<https://www150.statcan.gc.ca/n1/pub/92-195-x/2011001/geo/cma-rmr/cma-rmr-eng.htm>

an instrument for the change in house prices during the boom and bust cycle of the 2000s.

The approach developed by [Guren et al. \(2021\)](#) has two main advantages over the one of [Saiz \(2010\)](#). First, the Saiz measure correlates with other city characteristics such as productivity and growth in demand ([Davidoff, 2016](#)). This raises the concern that higher house price volatility in some cities is not driven by inelastic housing supply, as estimated by Saiz, but by differences in other characteristics such as different industrial composition and different exposure to secular trends, for example, an increase in housing demand in coastal areas with inelastic supply. To address this shortcoming, [Guren et al. \(2021\)](#) employs a panel specification that allows them to control for city-specific trends, different sensitivity to regional business cycles, and changes in the city’s population and industry structure. Second, by exploiting the systematic historical sensitivity of local house prices to regional house price cycles, this new approach allows us to estimate housing supply elasticities without resorting to geographical and regulation data across Canadian cities, data that are not currently available for most cities in Canada.

[Guren et al. \(2021\)](#) estimates housing supply elasticities by exploiting systematic differences in cities’ responses to regional house price cycles. [Sinai \(2012\)](#) documents that house prices in some US cities are systematically more sensitive to regional cycles than those in other cities, which is also true for Canada. Let’s consider Vancouver and Winnipeg. [Figure E.2](#) plots the annual log change of real house prices in the West,⁵ Vancouver and Winnipeg from 1992 to 2020. The West region experienced several regional boom-bust cycles throughout the sample period. Vancouver and Winnipeg also experienced several cycles that tended to correlate with the regional ones. However, house prices in Vancouver tended to increase more than those in Winnipeg during periods of regional house price booms. They also decreased by more when regional prices were contracting.

This systematic difference in the sensitivity of house prices in different cities to the regional house price cycles is crucial for the identification strategy described in the next Section.

E.1.1 Empirical strategy

A simple approach to estimating the sensitivity of house prices in different cities to regional house price movements, γ_i , consists of running the regression:

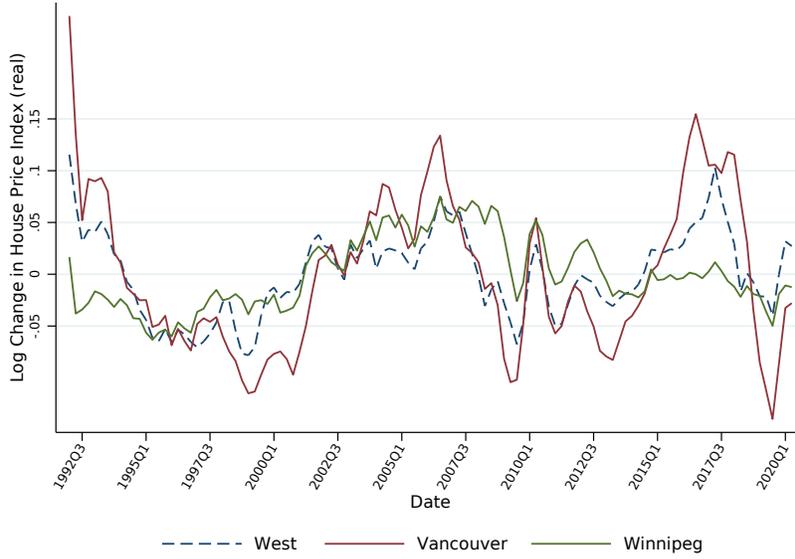
$$\Delta p_{i,r,t} = \phi_i + \chi_{r,t} + \gamma_i \Delta P_{r,t} + \epsilon_{i,r,t} \quad (24)$$

where $\Delta p_{i,r,t}$ denotes the log annual change of real house prices of city i in region r , and $\Delta P_{r,t}$ stands for the log annual change in regional house prices.⁶ This specification includes city fixed

⁵The West region includes all provinces West of Ontario.

⁶Throughout this note, I follow the same notation simplification as in [Guren et al. \(2021\)](#), where $\gamma_i \Delta P_{r,t}$ is used to denote $\sum_i \gamma_i \Delta P_{i,r,t} I_i$, where I_i is an indicator for city i .

Figure E.1: House prices in Vancouver, Winnipeg and the West region



Note: Figure E.1 reports the evolution of house prices in several Canadian regions. All time series correspond to the annual log change in the House Price Index. All series are demeaned relative to the city or region average. The West region includes all provinces west of Ontario. *Data Source:* Teranet.

effects, ϕ_i , to control for unobserved city heterogeneity, and region-time fixed effects, $\chi_{r,t}$, to control for trends at the regional level. Cities with higher $\hat{\gamma}_i$, the estimate of γ_i , are cities that systematically respond to regional shocks with higher fluctuations in higher prices and, therefore, cities with more inelastic housing supply. Therefore, $\hat{\gamma}_i$ denotes the proxy for the inverse of the housing supply elasticity.

This simple approach, however, assumes that local house prices respond differently to regional house price shocks only because of differences in the housing supply elasticity. This assumption seems too restrictive because differences in the structure of the local economy may cause different responses. Applying the example in Guren et al. (2021) to the Canadian context, we suppose that Vancouver has an industrial structure tilted toward highly cyclical durable goods relative to that of Winnipeg. A positive aggregate demand shock would consequently lead to higher increases in employment and house prices in Vancouver than in Winnipeg. Therefore, γ_i would be estimated to be higher in Vancouver than in Winnipeg purely due to reverse causality. Then, variation in $\hat{\gamma}_i$ would reflect not only differences in housing supply elasticities across cities but also potentially other confounding factors.

To address these concerns, we apply a refined version of equation 24 similar to the one proposed by Guren et al. (2021):

$$\Delta p_{i,r,t} = \phi_i + \gamma_i \Delta P_{i,r,t} + \delta_i \Delta y_{i,r,t} + \mu_i \Delta Y_{r,t} + \Gamma X_{i,r,t} + \epsilon_{i,r,t} \quad (25)$$

This version augments equation 24 with local and regional changes in per capita retail, construction and manufacturing employment with city-specific coefficients. The vectors with these changes in employment at the city and regional levels are $\Delta y_{i,r,t}$ and $\Delta Y_{r,t}$, respectively. This specification controls for the different impact across cities of different demand shocks reflected in these industries.⁷ It also includes another set of controls, $X_{i,r,t}$, specifically two-digit industry code shares multiplied by time dummies. This structure allows for non-parametrically controlling for all variation that is correlated with industry structure in the cross-section. I also depart from Guren et al. (2021) by controlling for population growth at the city and regional levels and for real mortgage rates.

Overall, this refined approach implies that $\hat{\gamma}_i$ is estimated using local house price variation that is independent of local and regional changes in employment and all other controls included in $X_{i,r,t}$. It is therefore not subject to the bias resulting from the reverse causality explained before. The key identifying assumption is that, conditional on controls, there are no other aggregate factors sensitive to house prices as captured by γ_i that are correlated with regional house prices in the time series and that differentially impact employment per capita in the same city. However, this approach does not require exogenous variation in regional house prices. Common factors can drive regional house prices, regional economic activity and even local prices and activity.

E.1.2 Data

We estimate the elasticities using the House Price Index developed by Teranet at the forward sortation area (FSA) level. City and regional house prices are built by aggregating them using the 2011 FSA populations as weights.⁸ House prices are converted into a real index by using the gross domestic product (GDP) deflator from StatCan. Following the methodology in Guren et al. (2021), we consider quarterly house prices and calculate annual changes of the log of the House Price Index in real terms. Annual two-digit industry employment codes and population at the CA level are obtained from StatCan. Real monthly mortgage rates are from the Bank of Canada.

Guren et al. (2021) consider four regions when estimating the elasticities for US cities. Given the significant differences in size between Canada and the United States, we consider three regions in Canada: east, west and northern territories.

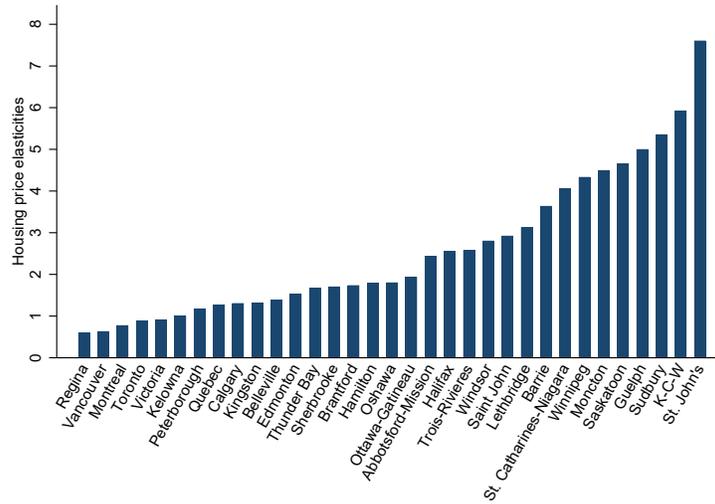
⁷Guren et al. (2021) control for local and regional changes in retail employment only. The correlation between the baseline estimates and the γ_i estimates using this less strict specification is 97 percent. If instead of controlling for changes in employment per capita in these three industries separately, I control only for aggregate changes in per capita employment, the correlation drops to 95 percent. More importantly, if I don't control for any changes in the employment growth across different industries at the city level, the correlation drops to 23 percent, which reflects the importance of controlling for changes in industry composition.

⁸Very similar results are obtained if total dwellings or total occupied dwellings are used as weights.

E.1.3 Results

Figure E.2 plots the estimated housing supply elasticities for Canadian census metropolitan areas (CMAs), specifically, the inverse of $\hat{\gamma}_i$ estimated from equation (25).⁹ The median housing supply elasticity is 2.2 among all CAs and 1.94 if I restrict the sample to CMAs. These estimates imply that a 1 percent increase in house prices in the median Canadian city is associated with an increase in housing supply of 2.2 percent. Alternatively, we can think that, all else equal, a 1 percent increase in housing demand leads to an increase in house prices in the median city of 0.45 percent ($1/2.2$).

Figure E.2: Housing supply elasticities for Canadian census metropolitan areas



Note: Figure E.2 plots the estimated housing supply elasticity, $1/\hat{\gamma}_i$, estimated from equation (25). London and Saguenay are excluded for visualization purposes. The elasticities for these two CMAs are 19.6 and 21.6, respectively. K-C-W stands for Kitchener-Cambridge-Waterloo.

Figure E.2 also shows significant heterogeneity across cities. Going back to the previous example, elasticities in Vancouver and Winnipeg are 0.63 and 4.34, respectively. Assuming that both cities face a 1 percent increase in housing demand, house prices are predicted to increase 1.57 percent in Vancouver and 0.23 percent in Winnipeg. For comparison, the median housing supply elasticity among US metropolitan areas estimated in Saiz (2010) is 2.26, very close to the median elasticity in Canada. Saiz also estimates elasticities of 0.63 and 0.72 in New York and San Francisco, respectively. These values compare closely with Vancouver and Toronto, where estimated elasticities in this note are 0.64 and 0.89, respectively. However, the distribution of elasticities in Canada is more right-skewed than in the United States. A larger share of cities in Canada have very elastic housing supplies.

⁹The procedure estimates housing supply elasticities for 151 CAs, but for clarity the Figure is restricted to the CMAs. For better visualization, London and Saguenay are also excluded from the figure. The elasticities for these two CMAs are 19.6 and 21.6, respectively.

F Additional Analysis on Model Parameterization

This section reports additional details and results related to the model’s parameterization.

F.1 More on Parametrization Procedure

Our parameterization strategy follows a moment-matching (minimum distance) framework. Let φ be a vector of parameters taken from the literature or obtained externally from detailed microdata, and let $\theta \in \mathbb{R}^p$ denote a vector of structural model parameters. The goal is to find θ and \hat{Z} that minimize the distance between a set of equilibrium model moments $m = h(\varphi, \theta, \hat{Z}) \in \mathbb{R}^m$ and their data counterparts $m^d \in \mathbb{R}^m$. Let \mathcal{W} be a $m \times m$ symmetric matrix. A “moment matching” or “minimum distance” estimator of θ and \hat{Z} is given by

$$\min_{\theta, \hat{Z}} \left(h(\varphi, \theta, \hat{Z}) - m^d \right)^\top \mathcal{W} \left(h(\varphi, \theta, \hat{Z}) - m^d \right) \quad (26)$$

In practice, we start with an initial guess for the parameters (θ^0, \hat{Z}^0) , solve the stationary equilibrium and the transition path implied by the guess of \hat{Z} , compute the corresponding equilibrium moments $m^0 = h(\varphi, \theta^0, \hat{Z}^0) \in \mathbb{R}^m$, and evaluate the loss function defined in (26). If the resulting loss is not sufficiently small, we update our guesses iteratively until the loss falls below a pre-specified threshold.

F.2 Parameterization of the Migration Elasticity and Sensitivity Analysis

Appendix F.2 examines the sensitivity of the model’s main outcomes to alternative values of ν , a key parameter governing the strength of spatial adjustment to shocks. Specifically, ν determines how strongly individuals respond to differences in expected utility across locations, thereby shaping both the cross-sectional population allocation and the dynamic response to local shocks. All else equal, a lower ν implies more elastic mobility, leading to faster household reallocation and weaker persistence of local shocks, whereas a higher ν amplifies the role of local frictions and generates more persistent spatial disparities. Its calibration is therefore critical for all counterfactuals, particularly those assessing how local labor market shocks propagate when mobility is limited.

Table F.1 is a by-product of the parameterization procedure. It reports the model-implied out-migration response to the oil shock and the associated loss function for values of ν ranging from 0.85 to 1.15. For each value of ν , all other parameters are chosen to minimize the distance

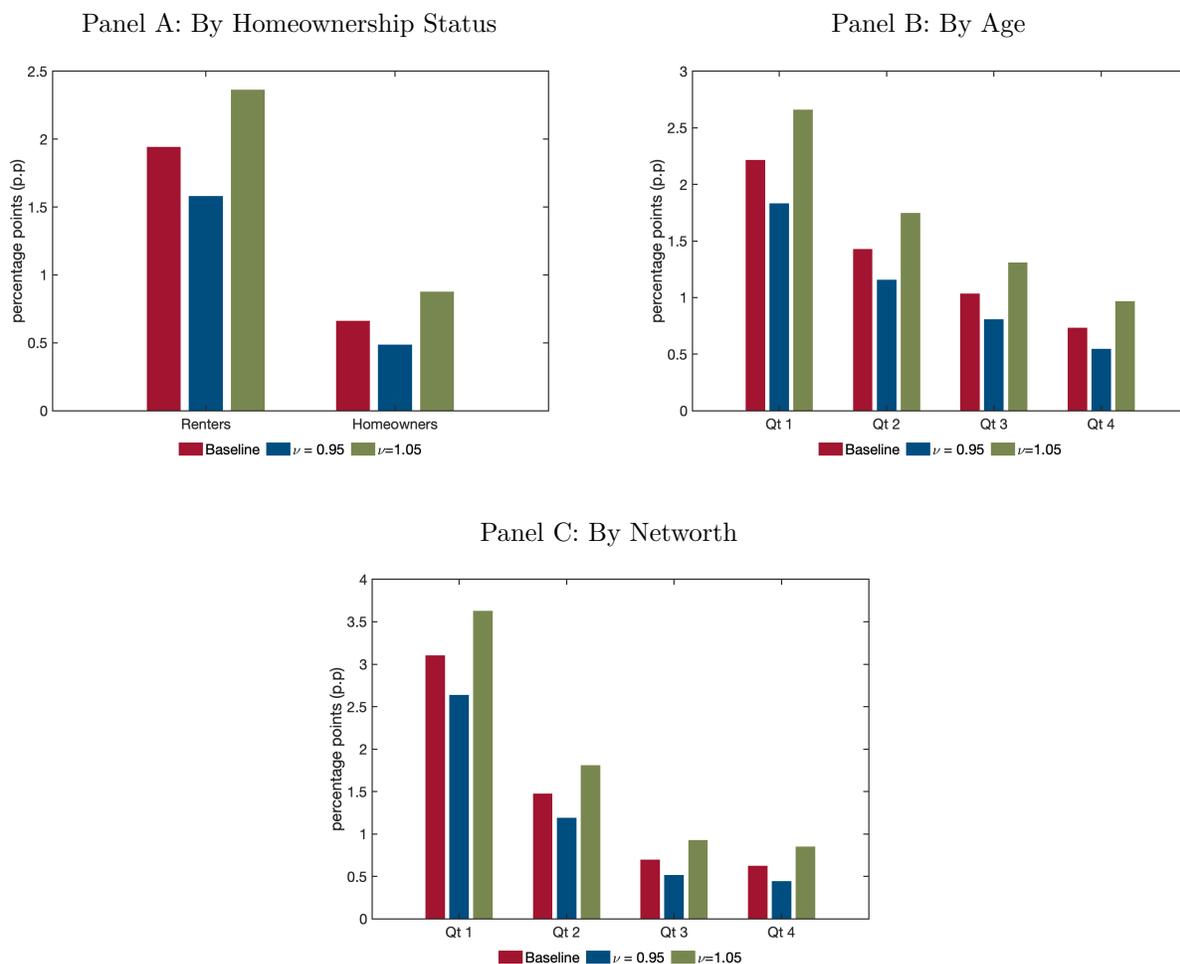
Table F.1: Estimation of the Migration Elasticity

ν	0.85	0.9	0.95	0.99	1	1.01	1.05	1.1	1.15
Out-migration response to the Oil Shock (%)	0.54	0.74	0.82	0.95	1.04	0.96	0.94	0.89	0.87
Loss	0.00432	0.00346	0.00345	0.00337	0.0033	0.00332	0.00366	0.00372	0.00365

Note: Let $h \in \mathbb{R}^m$ be the vector of model moments and m^d be the data counterpart. Loss functions is assumed to be the L2-norm defined as $\frac{1}{m} \sqrt{\sum_{i=1}^m (\frac{h_i - m_i^d}{m_i^d})^2}$.

between targeted moments in the data and their model counterparts. The loss function is minimized at $\nu = 1$, which motivates our choice of $\nu = 1$ in the model parameterization.

Figure F.1: Sensitivity of Migration Elasticity by Demographic Groups



Note: Figure F.1 reports the sensitivity of out-migration propensity when varying ν . In all panels, the solid red bars correspond to the baseline specification in which ν is estimated to be 1. The dark blue bars show the case equal to ν equal to 0.95, while the green bars show the case ν equal to 1.05. Panel A plots the out-migration propensity by homeownership status, Panel B by age group, and Panel C by network quartile.

Sensitivity analysis. To assess how ν affects migration elasticities, we conduct a sensitivity analysis by varying ν around its estimated value of 1, while holding all other parameters fixed at their baseline values. For each alternative value of ν , we re-solve households' value functions and migration decisions under the baseline stationary-equilibrium prices. We then aggregate individual out-migration propensities using the baseline stationary-equilibrium distribution.

Figure F.1 shows that higher values of ν increase overall out-migration propensities relatively uniformly across the wealth and age distributions and across homeownership status. This pattern reflects the fact that, conditional on wages, cost of living, and amenities, a higher ν places greater weight on idiosyncratic moving motives that are unrelated to local economic conditions.

F.3 Extra Figures and Tables

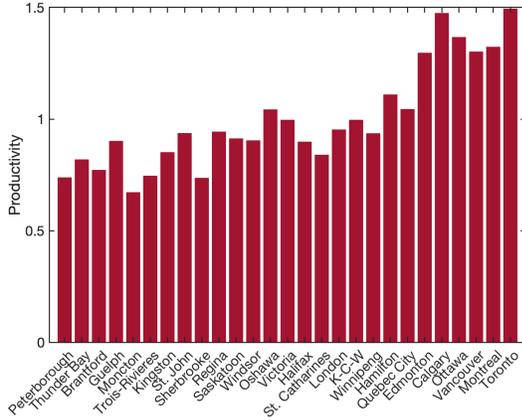
Table F.2: Parameter Values

Parameter	Interpretation	Internal	Value
Space			
L	Number of Locations	N	27
Demographics			
Q, Q	Length of Life, Working Years	N	60, 35
λ_q	Survival probability	N	StatCan
Preferences			
α	Housing consumption share	Y	0.35
$1/\gamma$	Elasticity of substitution (nondurables vs. housing)	N	1.25
β	Discount factor	Y	0.974
σ	Risk aversion	N	2
$\{\omega^l\}$	Additional utility from owning	Y	Figure F.2c
$\{e_q\}$	Equivalence scale	N	Auclert et al. (2021)
$\bar{\varphi}, a$	Bequest	N	930, 22
$\{A_q^l\}$	Amenities	Y	Figure F.2b
Endowments			
ρ_ϵ	Autocorrelation of earnings	N	0.91
σ_ϵ	S.D. of earnings shocks (25–34, 35–64)	N	0.35, 0.21
$\{\chi_q^l\}$	Life-cycle income profile	N	LFS 2012/StatCan
Π_l	Employment–unemployment transition matrix	N	LFS 2008–2013/StatCan
v	Unemployment replacement rate	N	0.31
w^{ret}	Retirement income	Y	4.32
Migration			
ν	Scale of Type 1 E.V. location shocks	Y	1
τ_0	Utility moving cost intercept	Y	6.96
$\bar{\tau}_1, \bar{\tau}_1^F, \bar{\tau}_1^E$	Distance component of utility moving costs (same-language, Eng–Fr, to Toronto/Vancouver)	Y	0.0065; 0.0228; 0.00013
F_m	Monetary moving cost	Y	0.23
Technology			
η	Labor elasticity	N	0.75
ζ	Agglomeration elasticity	Y	0.13
$\{z^l\}$	Exogenous local productivity	Y	Figure F.2a
Housing			
δ_h	Maintenance cost / depreciation	N	0.015
κ^l	Local housing supply elasticities	E	Figure E.2
F	Housing transaction costs	N	0.07
H^R, H^H	Housing size grid (rental, owner)	Y	City-specific, see Section 4
\bar{L}^l	Local land permits (buildable land)	Y	Figure F.2d
ψ_h^l	Property tax rate on housing	Y	Implied by price–rent ratio (eq. (18))
Financial Instruments			
r	Interest rate	N	0.015
ι	Borrowing wedge	N	0.01
\underline{b}	Unsecured borrowing limit	Y	–0.6
ξ	Collateral constraint	N	0.8
ψ_0, ψ_1	Income tax	N	0.92, 0.87

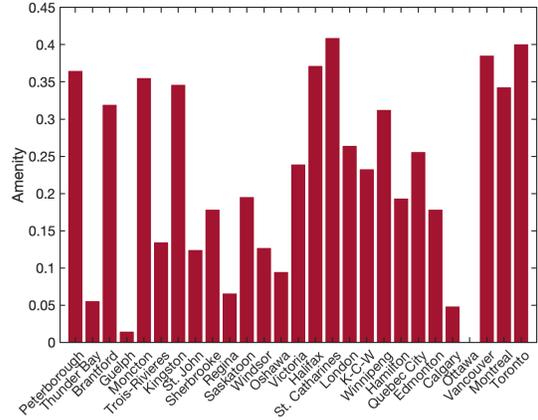
Note: Table F.2 reports the parameter values. The third column, *Internal*, states whether the parameter was internally parameterized (Y), externally taken from the literature or micro data (N). The model period is two years, but all the parameters shown in this Table are annualized. A unit of the final good in the model corresponds to 64,000 CAD, the 2012 median annual household earnings used to normalize incomes in the data.

Figure F.2: TFP, Amenity, Utility of homeownership, and Average Income

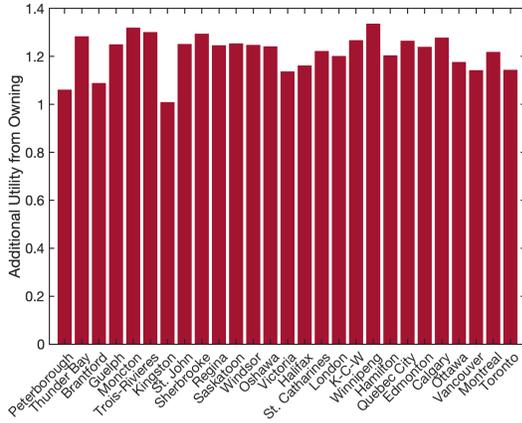
(a) Panel A: Productivity



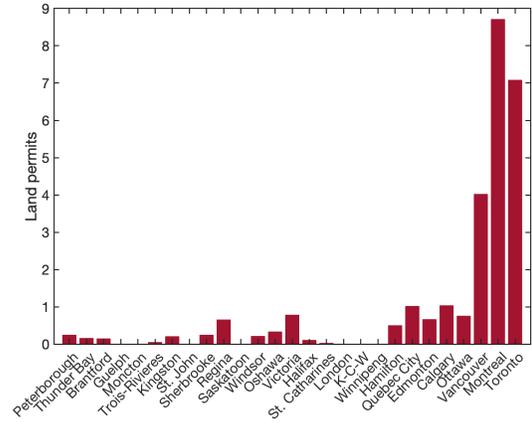
(b) Panel B: Amenity Index



(c) Panel C: Additional Utility of Homeownership

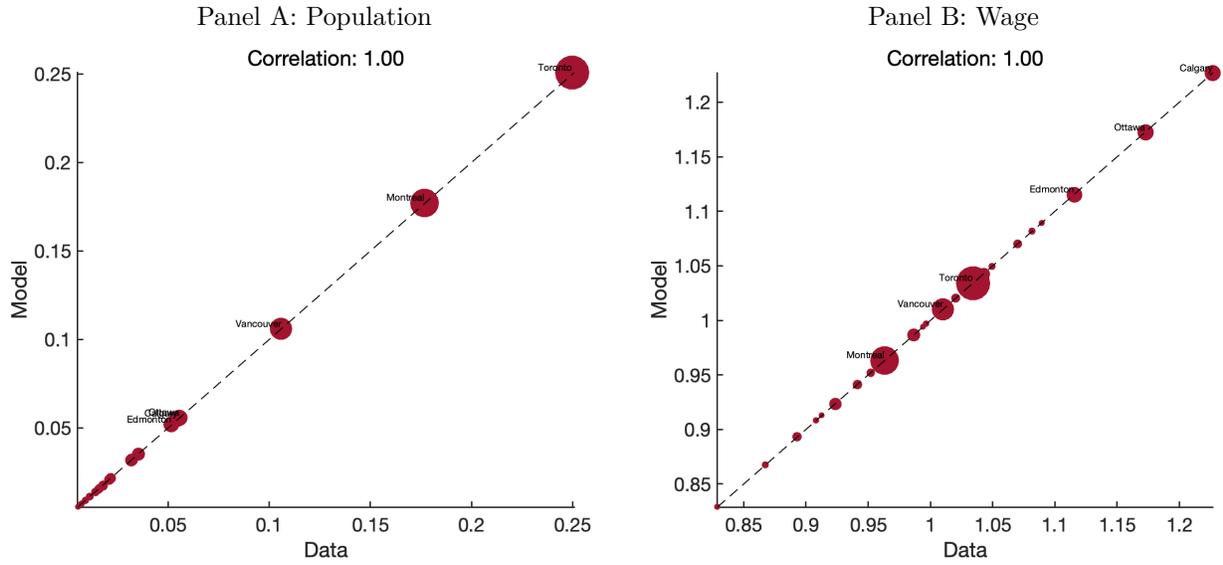


(d) Panel D: Land Permits



Note: Figure F.2 shows several internally calibrated parameters by city. Panel A shows the city-level measures of exogenous productivity, Panel B the average amenity index (aggregated across age groups, normalized to 0 in Ottawa), and Panel C the additional utility of homeownership. Panel D shows the land permits by city. Cities are ordered increasingly by size.

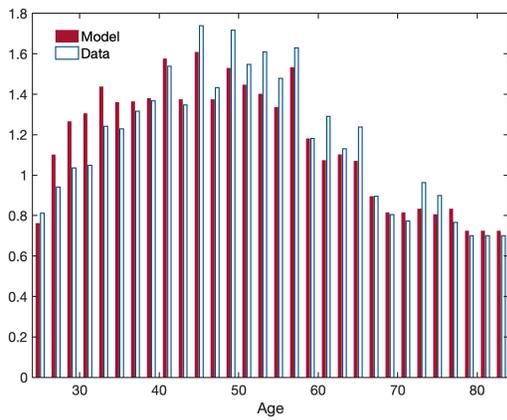
Figure F.3: Model vs Data: Population and Wages across cities



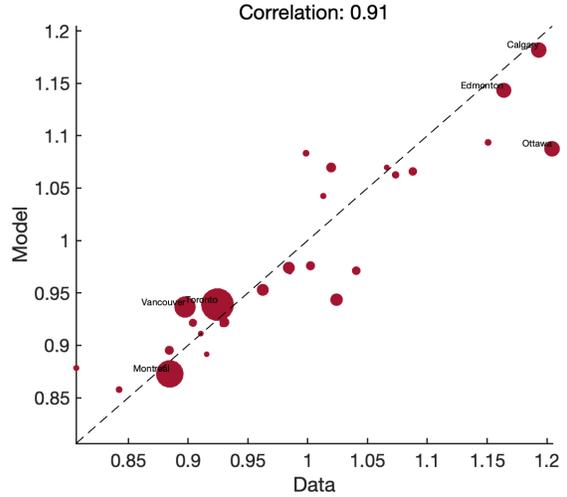
Note: Figure F.3 plots the population (Panel A) and wage (Panel B) by CMA both in the data and in the model. In both panels cities are ordered increasingly by population size. *Data Source:* StatCan.

Figure F.4: Income in the model and the data

Panel A: Average income by age groups

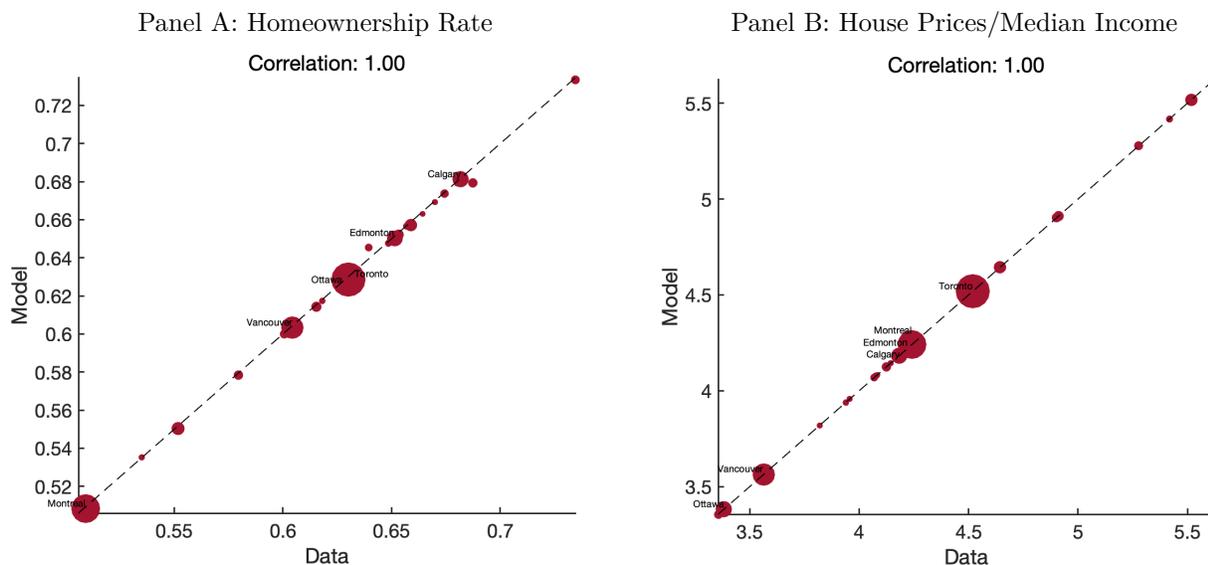


Panel B: Median income by city



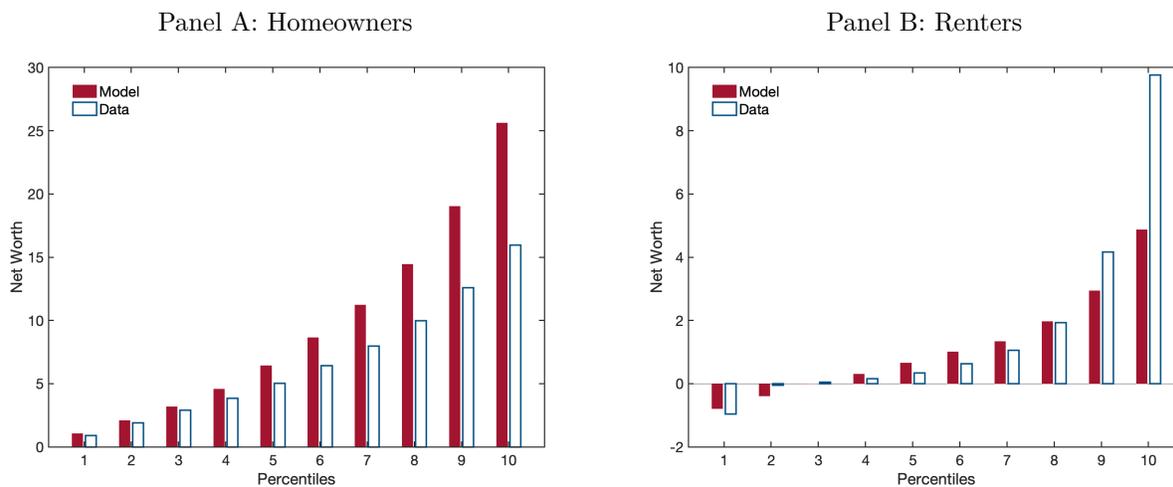
Note: Figure F.4 plots average income by age groups in the model and in the data in Panel A. Panel B plots median income by city in the model and in the data, with cities ordered by population.

Figure F.5: Model vs Data: Homeownership Rates and House Prices across Cities



Note: Figure F.5 plots the homeownership rates (Panel A) and house prices over median income (Panel B) by CMA both in the data and in the model. In both panels cities are ordered increasingly by size. *Data Source:* StatCan.

Figure F.6: Wealth Distribution by Homeownership Status



Note: Figure F.6 plots the wealth to income ratio and the house value to income ratio by CMA both in the data and in the model for homeowners (Panel A) and renters (Panel B). *Data Source:* StatCan.

Table F.3: Correlations between City Characteristics

	House Prices	Wages	Avg. Income	Population	TFP	Amenities
Model						
House Prices	1	0.24	0.23	0.15	0.24	-0.03
Wages	0.24	1	0.96	0.38	0.86	-0.50
Av Income	0.23	0.96	1	0.21	0.73	-0.60
Population	0.15	0.38	0.21	1	0.76	0.33
TFP	0.24	0.86	0.73	0.76	1	-0.09
Amenities	-0.03	-0.50	-0.60	0.33	-0.09	1
Data						
House Prices	1	0.31	0.28	0.36	0.45	-0.06
Wages	0.31	1	0.92	0.19	0.72	-0.62
Av Income	0.28	0.92	1	0.16	0.66	-0.46
Population	0.36	0.19	0.16	1	0.76	0.34
TFP	0.45	0.72	0.66	0.76	1	-0.09
Amenities	-0.06	-0.62	-0.46	0.34	-0.09	1

Note: Table F.3 reports correlations among city-level characteristics in the model (top panel) and in the data (bottom panel). *Data sources:* TransUnion, StatCan, and Teranet.

Figure F.7: Migration rates by Homeownership and Age; Population Distribution and In-Migration



Note: Figure F.7 reports the migration rates by homeownership (Panel A) and age (Panel B) for several model specifications. In panel C, we report the in-migration share. Each bar corresponds to alternative model specifications: baseline economy (red bar), economy with no homeownership (light red bar), economy with no income risk (green bar) and economy with no borrowing limit (blue bar).

G Moving Costs

This section documents the role of moving costs in the model and explains how they are quantified. We begin by deriving the procedure used to convert the estimated utility moving costs into monetary units, taking into account the non-linear preferences that govern households' location choices. We then analyze the magnitude of these costs in the stationary equilibrium and study how they vary across alternative model specifications that remove key mechanisms such as homeownership, income risk, and borrowing constraints. Finally, we analyze how changing monetary vs utility costs shape the moving propensity differently between different wealth groups. Together, these two subsections clarify both the methodology behind our moving cost estimates and the economic channels that shape them. The results will be central for interpreting households' mobility responses in the subsequent local downturn counterfactuals.

G.1 Derivation of Moving Costs

In this section, we show how we convert utility moving costs into monetary units. Given the flow utility depends on non-durable consumption and housing services, we re-write households life-time utility in terms of an “adjusted-consumption” measure, $\omega_{i,t}$, that delivers the same flow-utility (net of amenities) a household would obtain under the stationary equilibrium allocations. For an individual i with housing status d in location l at time t that optimally chooses c_{it} and s_{it} in the stationary equilibrium, the “adjusted-consumption” measure, $\omega_{i,t}$, is given by:

$$u^d(c_{it}, s_{it}, A_q^l) = \tilde{u}(\omega_{i,t}) + A_q^l = \frac{\omega_{i,t}^{1-\sigma}}{1-\sigma} + A_q^l$$

The re-write value functions in terms of $\omega_{i,t}$ given by:

$$V_{i,t}^{d,l} = \tilde{u}(\omega_{i,t}) + A_q^l + (1 - \lambda_i)\varphi(a'_{i,t}) + \lambda_i\nu \log \left(\sum_{k=1}^L \exp \left(\beta \mathbb{E}_t V_{i,t+1}^k - \beta \tau^{l,k} \right)^{\frac{1}{\nu}} \right)$$

deliver the same life-time utility for household i in location l with housing status d than the value functions defined in equations (9) and (10).

We now present our procedure to convert the estimated utility moving costs into a dollar equivalent. Contrary to [Kennan and Walker \(2011\)](#), the utility function is not linear, so the conversion is not direct. Instead, we solve for δ_i , the change in “adjusted-consumption” measure $\omega_{i,t}$ required to achieve the same individual location choices and allocations, and respectively life-time utility, in the absence of moving costs. A household that survives and moves from location l to location j must be indifferent between paying utility cost $\tau^{l,k}$ or facing a cut in

“adjusted-consumption” of $\delta_i^{l,k}$:

$$\tilde{u}(\omega_{i,t}) + \mathbb{E}_t V_{i,t+1}^k - \tau^{l,k} + A_q^l + \nu \tilde{e}_{t+1}^{i,k} = \tilde{u}(\omega_{i,t} - \delta_i^{l,k}) + A_q^l + \mathbb{E}_t V_{i,t+1}^k + \nu \tilde{e}_{t+1}^{i,k}$$

which can be re-written in the stationary equilibrium as

$$\delta_i^{l,k} \frac{\tilde{u}(\omega_i) - \tilde{u}(\omega_i - \delta_i^{l,k})}{\delta_i^{l,k}} = \tau^{l,k}$$

The left-hand side can be approximated by marginal utility evaluated at the “adjusted-consumption” measure ω_i consistent with the stationary equilibrium. Therefore, the utility moving costs evaluated at consumption-equivalent units, $\delta_i^{l,k}$ is given by

$$\delta_i^{l,k} = \frac{\tau^{l,k}}{\tilde{u}'(\omega_i)}$$

As in [Kennan and Walker \(2011\)](#), we compute the moving costs for the average mover, $\bar{\tau}$, given by:

$$\bar{\tau} = \frac{\sum_{k \neq l}^L \tilde{\mu}_i^{l,k} \tau^{l,k}}{\tilde{u}'(\bar{\omega})} + F \quad (27)$$

where $\tilde{u}'(\bar{\omega})$ is the marginal utility of the average mover in the stationary equilibrium, $\tau^{l,k}$ the utility moving costs from l to k , F the monetary moving cost and $\tilde{\mu}_i^{l,k} = \frac{\mu_i^{l,k}}{\sum_{i \neq k}^L \mu_i^{l,k}}$ is the probability of moving from l to k , conditional on moving, for the average mover with $\mu_i^{l,k}$ is defined in equation (13).

G.2 The Size of Moving Costs

Moving costs are often pointed out as the main driver of moving decisions. As previously discussed, different features of the model impact households moving propensities, which in turn lead to different model-implied moving costs estimates. We show that not accounting simultaneously for the precautionary moving and precautionary savings motives leads to higher model-implied moving costs estimates.

Methodology Given that the utility function is not linear, converting the utility moving costs into a dollar equivalent is not direct. Our procedure consists in solving for an individual specific change in consumption required to achieve the same individual location choices and allocations, and respectively life-time utility, in the absence of moving costs.¹⁰ The moving cost

¹⁰Given the flow utility depends on non-durable consumption and housing services, we re-write households

in consumption-equivalent units of an average mover is given by:

$$\bar{\tau} = \frac{\sum_{k \neq l}^L \tilde{\mu}^{l,k} \tau^{l,k}}{\tilde{u}'(\bar{\omega})} + F$$

where $\tilde{u}'(\bar{\omega})$ is the marginal utility of the average mover in the stationary equilibrium, $\tau^{l,k}$ the utility moving costs from l to k , F the monetary moving cost and $\tilde{\mu}_i^{l,k} = \frac{\mu_i^{l,k}}{\sum_{i \neq k}^L \mu_i^{l,k}}$ is the probability of moving from l to k , conditional on moving, for the average mover with $\mu_i^{l,k}$ is defined in equation (13).

Table G.1: Decomposition of Moving Costs

	Baseline	No House	No Income Risk	No Borrow Const
τ_0	6.96	7.08	5.94	6.85
τ_1	0.01	0.01	0.01	0.01
F_m	0.23	0.23	0.23	0.23
CAD 2012				
Moving Costs	287,864	292,865	248,845	283,726
Moving Costs - Males	342,239	348,184	295,850	337,319
USD 2010				
Moving Costs	248,993	253,319	215,244	245,414
Moving Costs - Males	306,967	312,300	265,360	302,555

Note: Table G.1 reports the values of moving costs in monetary terms for the baseline economy, for an economy without homeowners (*No House*), without income risk (*No Income Risk*) and no borrowing constraints (*No Borrow Const*) The top part of the Table reports the corresponding value of the moving costs, τ_0 , τ_1 and F_m . The second part of the Table reports the values of moving costs in monetary terms. “CAD 2012” is Canadian dollars in 2012 units. “USD 2010” is US dollars in 2010 units.

Results Table G.1 reports the model implied parameters for each case in utility terms before being converted into dollar amounts. We find that the utility costs represent a larger fraction of the overall moving costs. We find that the moving costs between Canadian cities are approximately 288,000 CAD (in 2012 units) in the baseline model, as reported in the first row of Table G.1 in Appendix G.3. A model with no homeownership that matches aggregate moving

life-time utility in terms of an “adjusted-consumption” measure, $\omega_{i,t}$, that delivers the same flow-utility (net of amenities) a household would obtain under the stationary equilibrium allocations.

moments implies moving costs that are approximately 1-4 percent higher than in the baseline economy. The results change substantially when we remove income risk. Moving costs drop to approximately 249,000 CAD (in 2012 units). As mentioned before, households have a lower propensity to migrate in the absence of income risk. Thus, lower moving costs are needed to induce households to move at the rate observed in the data. This result is quantitatively large since it corresponds to a drop of 14 percent in moving costs relative to the baseline estimates. Similarly, the last column of Table G.1 reports the results of moving costs for the model without borrowing constraint. In this case, the insurance channel of migration becomes less relevant and households are less likely to move. Under this specification, moving costs are estimated at 284,000 CAD (in 2012 units), 1.4 percent lower than in the baseline case.

The model implied moving costs depend on the strength of the precautionary moving motive. Higher income risk and stronger financial frictions increase the incentive to move. Thus, these channels imply higher moving costs to match the observed ones in the data. Therefore, models where households can only smooth income risk through moving but not through financial markets might imply higher moving costs. Overall, the different estimates across specifications show that not accounting for the ability to smooth shocks simultaneously by either moving or through financial market increases moving costs estimates.

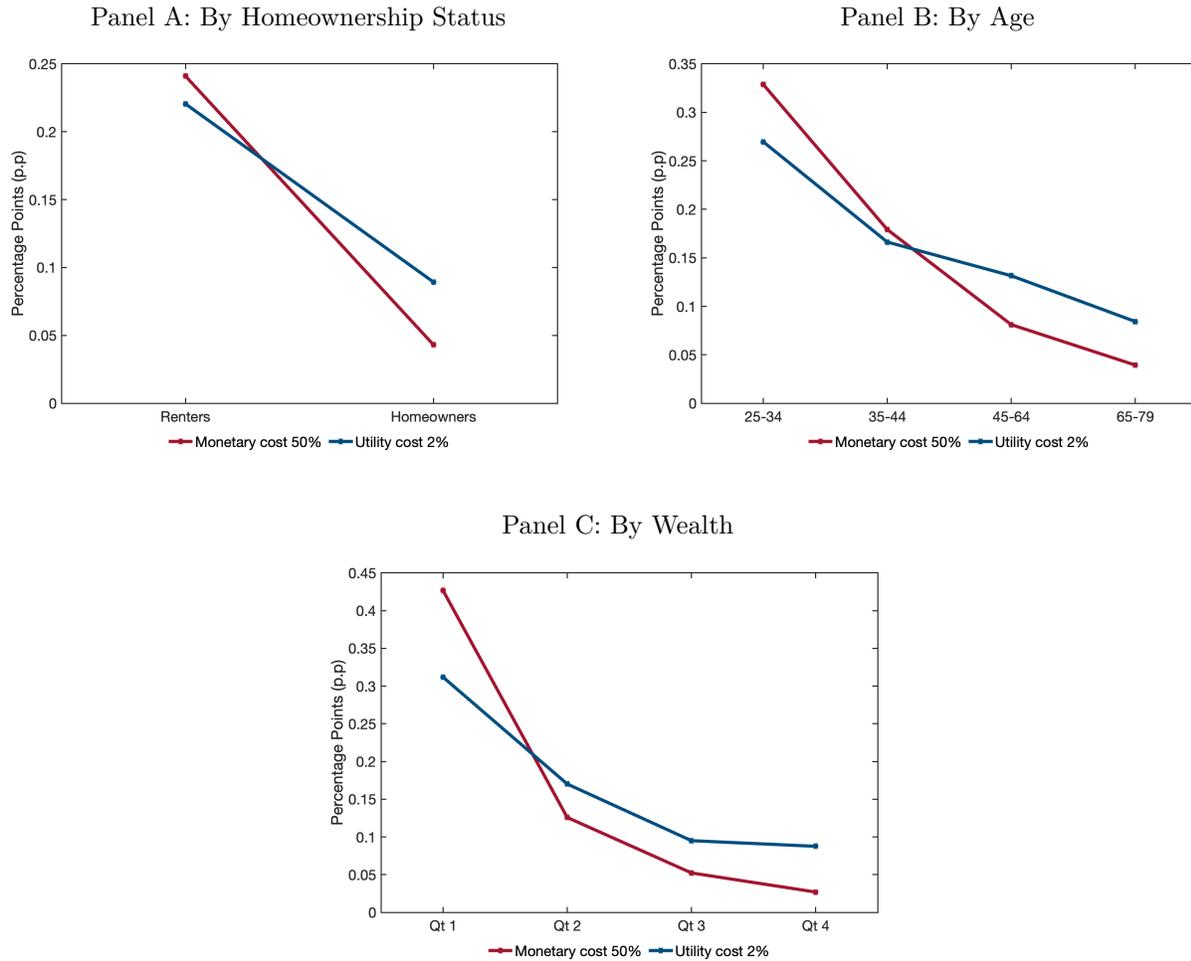
G.3 Moving Cost: Monetary vs Utility

In Appendix G.3, we examine the role of moving costs in shaping relocation decisions by decomposing the effects of monetary and non-monetary components. The model incorporates both types of costs, and each is calibrated to match observed migration rates. Monetary moving costs enter directly into the household budget constraint and therefore interact with borrowing limits and liquid wealth. In contrast, non-monetary (utility) moving costs affect the value of relocating but do not interact with financial frictions. As a result, the two components influence different groups in systematically different ways.

This distinction generates heterogeneous mobility responses across households. Monetary moving costs disproportionately affect low-wealth agents, for whom even moderate costs may bind borrowing constraints or compress available consumption. Reducing these monetary costs, therefore, relaxes liquidity constraints and yields a larger increase in mobility for poorer households. Utility moving costs, by contrast, enter symmetrically across the wealth distribution and thus have a more uniform effect on mobility, since they neither tighten nor relax financial constraints.

To illustrate these mechanisms, we consider counterfactual scenarios that increase aggregate moving rates by a fixed amount, from 1.47 percent to 1.72 percent a year. We implement this decline in two separate ways: (i) by proportionally reducing monetary moving costs; and (ii) by

Figure G.1: Monetary vs Non-monetary Moving Costs



Note: Figure G.1 decomposes the effect of moving costs into monetary and non-monetary components. To compare their effects, we increase aggregate migration rates by the same amount, p.p., either by lowering monetary costs or by lowering utility costs. Because monetary costs are smaller in level, achieving the same aggregate decline requires a 50 percent reduction in monetary costs, compared with only a 2 percent reduction in utility costs. Panel A reports mobility responses by homeownership status, Panel B by age group, and Panel C by net-worth quartile.

reducing utility moving costs. Because monetary costs are calibrated to be smaller in level than utility costs, achieving an equivalent aggregate reduction in mobility requires a substantially larger proportional change. In particular, increasing migration from 1.47 percent to 1.72 percent a year requires a 50 percent reduction in monetary costs but only a 2 percent decrease in utility costs.

The results, reported in Figure G.1, show that the main qualitative insights are preserved regardless of whether we reduce monetary or utility-based moving costs. Lower migration frictions generate a stronger migration response among renters, young workers, and individuals

with low wealth. These findings confirm that our main conclusions do not hinge on the specific way moving costs are modeled, but instead reflect the interaction between homeownership, income risk, and mobility frictions.

At the same time, Figure G.1 also highlights a subtle difference that has been largely overlooked in the migration literature. When we reduce monetary moving costs, mobility rises almost exclusively among low-wealth and liquidity-constrained households, for whom even modest reductions relax binding financial frictions. By contrast, reducing utility moving costs generates the opposite ordering: the largest mobility increases occur among households in the third and fourth wealth quartiles, whose relocation incentives are shaped primarily by non-pecuniary considerations rather than by budget constraints. This reversal in the ranking of mobility responses, depending on which component of moving costs is perturbed, is both novel and economically meaningful. It shows that the composition of moving costs activates entirely different margins of adjustment: monetary costs operate through liquidity constraints, while utility costs operate through the valuation of location amenities and switching frictions. Similar patterns arise across age groups and by homeownership status, reinforcing the idea that the interaction between financial frictions and the structure of moving costs is central to understanding who moves and why. More broadly, these results suggest that models treating moving costs as a single undifferentiated object may miss important sources of heterogeneity in migration behavior and potentially misattribute the mechanisms through which local shocks reshape population flows.

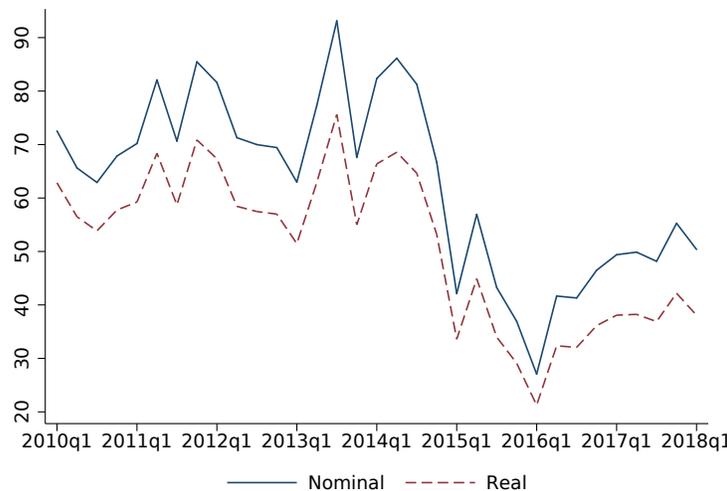
H Identifying a Local Downturn in the Data

This Appendix details how we exploit the collapse in global oil prices to identify local aggregate negative shocks in Canada. We begin with background information on the 2014–2016 oil price collapse and the structure of the Canadian oil sector, and then describe the empirical strategy to obtain measures of negative local shocks for different Canadian cities. We conclude by documenting how the resulting shocks affect key economic outcomes relevant to migration and local housing markets.

H.1 Background on the 2014–2016 Oil Price Collapse

Canada is the fifth-largest oil producer in the world and most of its crude oil production is exported. Because Canada accounts for a relatively small share of the international oil market, global oil price fluctuations are widely regarded as plausibly exogenous to domestic macroeconomic conditions. Production is also highly geographically concentrated, with approximately 90 percent of the production originating in two provinces: Alberta and Saskatchewan.¹¹ These features create substantial cross-sectional variation in exposure to oil price movements across Canadian cities.

Figure H.1: Oil Price Evolution



Note: Figure H.1 reports the evolution of oil prices by quarter between 2010 and 2018, both in nominal and real terms.

The period between 2014 and 2016 witnessed one of the sharpest oil price collapses in modern history. Figure H.1 plots the evolution of the price of the Western Canadian Select crude oil

¹¹Alberta, Saskatchewan, and Newfoundland and Labrador account for 95 percent of Canadian oil production. Since most production in Newfoundland and Labrador occurs offshore, we exclude it from the analysis.

(WCS), measured in Canadian dollars. After remaining relatively stable between 2010 and 2014, WCS prices fell by approximately 62 percent between 2014 and 2016 in both nominal and real terms.¹² Although prices partially recovered thereafter, they remained nearly 40 percent below their pre-2014 average in 2018.

H.2 Oil Price Collapse as a Source of Local Downturn

Oil provinces experienced substantially different economic trajectories after 2014 relative to the rest of Canada. As shown in Panel A of Figure H.2, average income in Alberta and Saskatchewan declined by roughly 9 percent after the shock and remained below pre-2014 levels through at least 2019. In contrast, income continued to grow, albeit modestly, in non-oil-producing provinces. Importantly, prior to 2014, income in oil-producing provinces had been rising at a pace comparable to or faster than the national average, reinforcing the view that the downturn was driven by the oil shock rather than by a pre-existing long-run trend.

Although oil extraction sites themselves are located primarily outside major census metropolitan areas, the oil sector constitutes a substantial share of economic activity in several CMAs. In Calgary, Edmonton, Saskatoon, and Regina, oil-related industries represent approximately 2.38 percent of total employment, against 2.78 percent in the overall province level.¹³ Consequently, these cities also experienced pronounced declines in income after the price shock, with higher magnitude and persistency than in the rest of the provinces (Panel B of Figure H.2).¹⁴

Empirical Strategy and Results To identify city-level local aggregate shocks, we exploit the plausibly exogenous decline in global oil prices with cross-sectional differences in local employment exposure to the oil industry. Following Kilian and Zhou (2022), we build a city-specific *Oil Exposure* taking into account the employment share in the oil industry. Specifically, we build a measure of regional exposure to the oil shock that resembles the standard *Bartik* shocks as follows:

$$\text{Oil Exposure}_{l,t} = \alpha_{l,t-1} \log(\text{OilPrice}_t) \quad (28)$$

where $\alpha_{l,t-1}$ is the share of employment in the oil sector (NAICS 2111) in city l in 2011 and $\log(\text{OilPrice}_t)$ is the logarithm of the real oil price at quarter t .¹⁵ This measure captures how

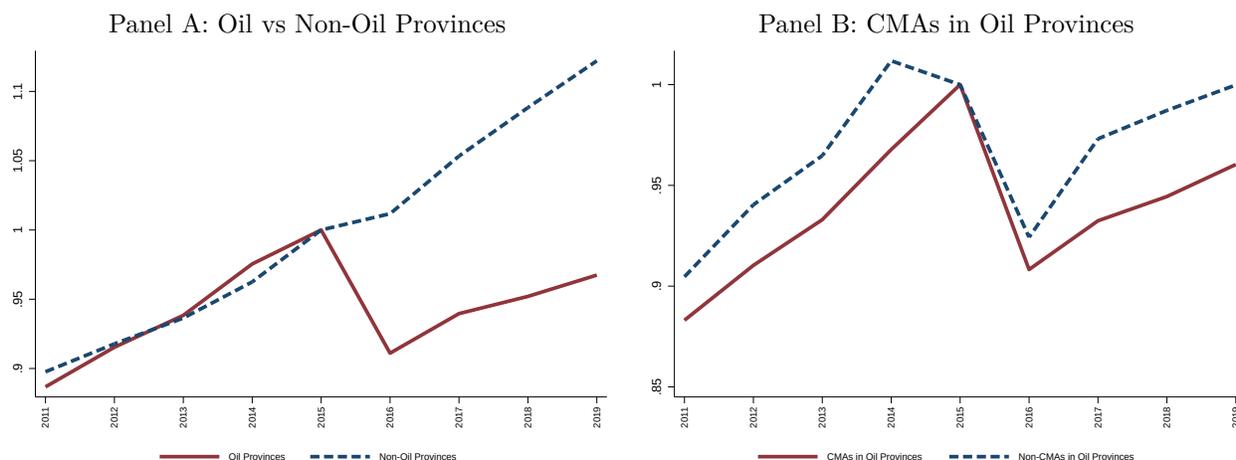
¹²Real prices are computed by deflating WCS by the national wide Canadian consumer price index (CPI). WCS is the oil price measured also used by Kilian and Zhou (2022).

¹³The share of employment in oil industry is 0.48 percent overall in the country.

¹⁴Most of the facilities are located close to oil sands that are a natural mixture of sand, water and bitumen (oil that is too heavy or thick to flow on its own). The oil sands are found in main three regions within the provinces of Alberta and Saskatchewan: Athabasca, Cold Lake and Peace River.

¹⁵There exists pronounced heterogeneity in the employment share of the oil sector even within the oil-producing provinces. For instance, in Wood Buffalo, Cold Lake, and Calgary, the oil sector accounts for 27.06, 8.32, and 4.70 percent of total employment, respectively, whereas in Lethbridge or Regina the corresponding share is below

Figure H.2: Impact of Oil Price Collapse on Income



Note: Figure H.2 shows the evolution of average income for individuals aged 35–54. Panel A reports average income for Oil Provinces (Alberta and Saskatchewan) and Non-Oil provinces. Panel B splits the sample between CMAs located in Oil Provinces and all other locations within those provinces. Average income in each location is normalized to one in 2015. *Data Source:* StatCan.

the global price collapse differentially affected local labor markets depending on their reliance on the oil industry. We then adapt the Bartik-style panel regression model proposed by Kilian and Zhou (2022) and regress changes in annual average income in location l on Oil Exposure $_{l,t}$ controlling for location and year fixed effects:

$$\Delta inc_{l,t} = \beta \text{Oil Exposure}_{l,t} + \phi_l + \phi_t + \varepsilon_{l,t}, \quad (29)$$

Kilian and Zhou (2022) show that under empirically plausible conditions the derivative of the dependent variable with respect to the oil price shock can be given a causal interpretation. We therefore use the predicted change in income, $\widehat{\Delta inc}$, as our measure of the local aggregate shock.

Table H.1 reports the estimates of equation (29). A one-percent decline in oil prices generates a 2.3-percentage-point larger reduction in average income in locations with high exposure to the oil sector (90th percentile) relative to those with low exposure (10th percentile). This differential corresponds to approximately 80 percent and 220 percent of the income decline experienced by the average and median locations, respectively.

Local shocks and economic outcomes To analyze how local economic shocks impact economic outcomes, we follow a similar specification to equation (29), where we regress the

1 percent. More precisely, in 2011 the median and mean values of the oil-sector employment share were 1.62 and 2.90 percent, respectively. The distribution exhibits considerable dispersion, with the 5(10)th and 95(90)th percentiles at 0.10(0.13) and 5.35(4.9) percent.

I Derivation of Welfare Measure

In this section, we describe how we construct the welfare measure defined in equation (22). We follow the standard approach of expressing welfare in consumption-equivalent units. Specifically, we convert lifetime utility into a more interpretable metric by computing a constant level of aggregate-consumption, ω_i , that delivers the same lifetime utility as the stationary-equilibrium allocation.

Let's denote by V_i the life-time utility of household i . The constant consumption ω_i is the solution to

$$V_i = \sum_{t=q}^Q \beta^t \tilde{u}(\omega_i) = \sum_{t=q}^Q \beta^t \frac{\omega_i^{1-\sigma}}{1-\sigma} \quad (31)$$

where the sum is defined over the remaining life years of a household currently q years of age.

As pointed out in Boar and Midrigan (2022), ω_i corresponds also to a measure of welfare adjusted for risk and intertemporal substitution that allows for interpersonal comparisons.

We next compute the welfare impact of an unanticipated aggregate shock. Let X^E denote the constant aggregate state associated with the stationary equilibrium, and let $V_i(X^E)$ be agent i 's value function evaluated at that equilibrium. Let $V_{i,t}(X^S)$ denote the value function along a shock path. We assume that an unanticipated shock occurs at time $t + 1$ and persists until period T , after which the economy returns to the stationary equilibrium path. The aggregate state sequence is therefore given by $X^S = \{X^E, X_1^S, X_2^S, \dots, X_T^S, X^E, \dots, X^E\}$.

We consider a utilitarian social planner who assigns equal weight to all individuals. The welfare of a fixed group of households g_t , with cardinality G , is defined as

$$W_g = \sum_{i \in g_t} \omega_i,$$

where ω_i is the consumption-equivalent measure implied by equation (31), which can be expressed as

$$\omega_i \propto V_i^{\frac{1}{1-\sigma}}.$$

Because the shock is unanticipated and occurs at time $t + 1$, we measure its welfare impact at that date. The welfare change for group g_t is given by

$$\Delta W_g = \frac{\sum_{i \in g_t} \omega_{i,t+1}(X^S)}{\sum_{i \in g_t} \omega_i(X^E)} - 1 = \frac{\sum_{i \in g_t} V_{i,t+1}(X^S)^{\frac{1}{1-\sigma}}}{\sum_{i \in g_t} V_i(X^E)^{\frac{1}{1-\sigma}}} - 1.$$

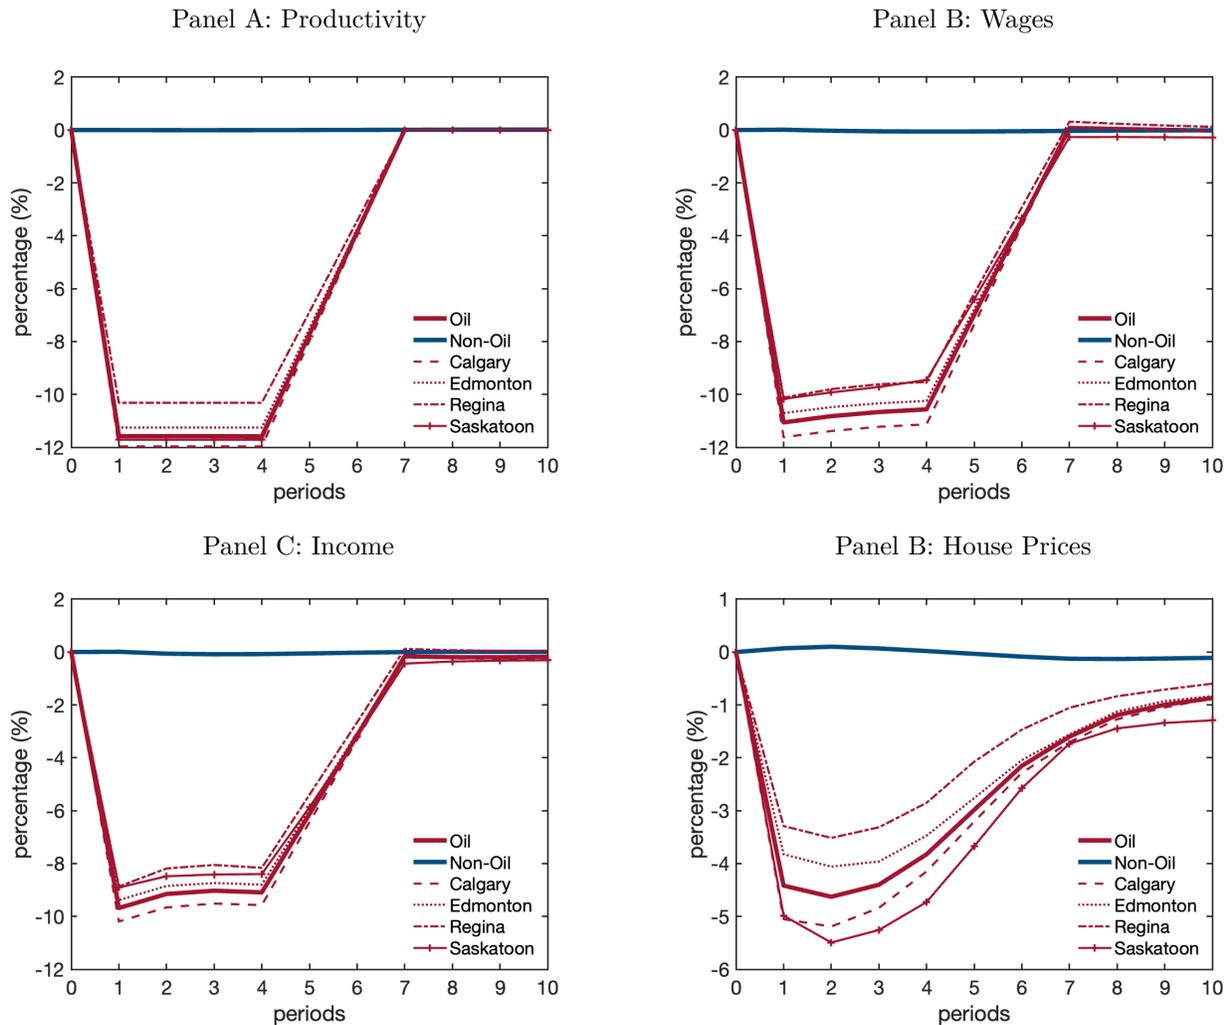
Importantly, the group g_t is defined as the set of households residing in a given city prior to the realization of the shock. Because migration is endogenous, the composition of households living in that city may change after the shock. To isolate the welfare impact of the shock on

the original residents of a location, we compute welfare for this fixed group of households and track their outcomes at time $t + 11$, regardless of whether they choose to remain in the city or migrate elsewhere.

J Additional Results on the Impact of a Local Shock

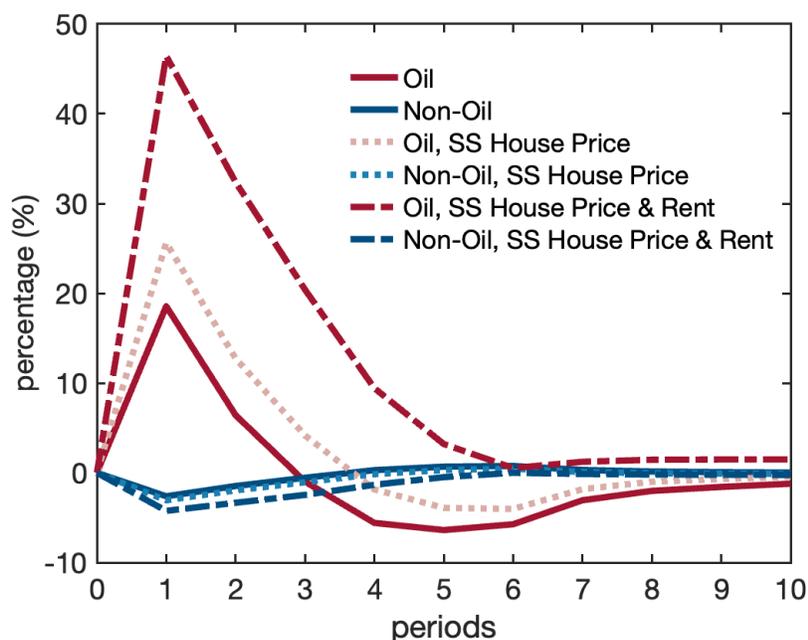
This Appendix reports extra results on the impact of the oil shock in the model. It also details the calculation of the welfare measure then used to assess the welfare impact of the shock under several scenarios in Section I. It then reports the results of the oil shock in the scenario in which housing transaction costs were lowered in Section J.1. It concludes with showing the results of the counterfactual oil shock in which the oil shock were set to be permanent in Section J.2.

Figure J.1: Local Productivity Shocks, Wages, Income and House Prices



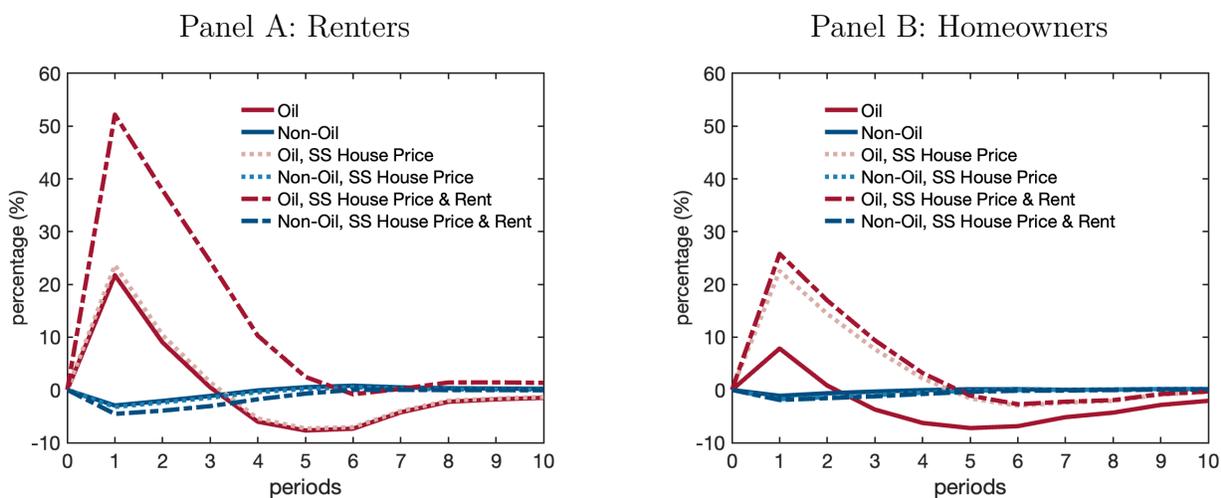
Note: Figure J.1 reports the evolution of aggregate variables in response to the oil shock. Panel A reports the paths of the exogenous productivity component. Section 4.3 describes how these measures are backed out from the data. Panels B, C, and D plot the endogenous paths of local wages, average income, and house prices induced by the shocks shown in Panel A. All series are expressed as percentage deviations from the stationary pre-shock equilibrium.

Figure J.2: Out-Migration Response to the Local Downturn



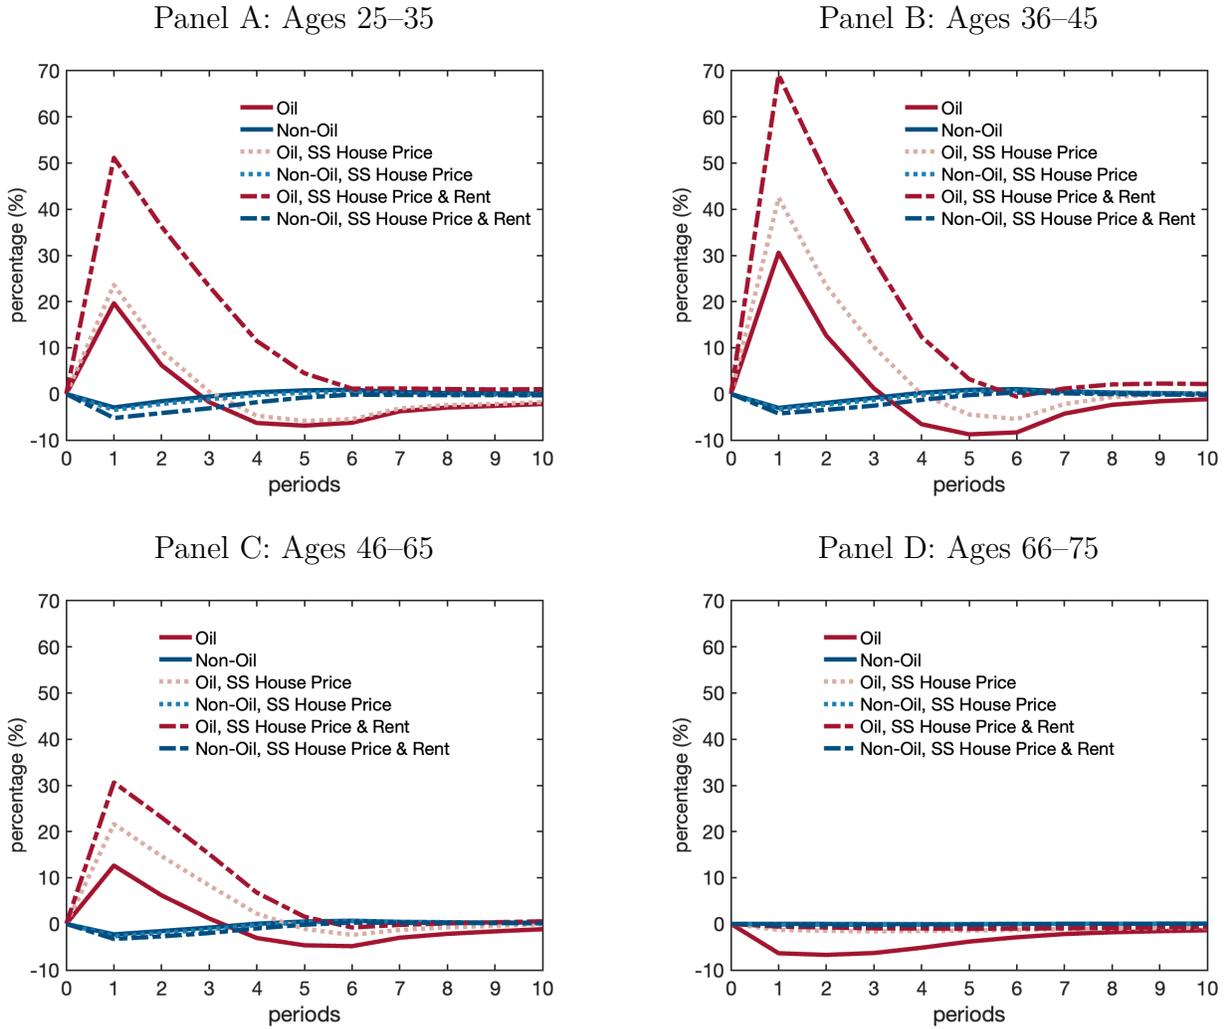
Note: Figure J.2 plots the out-migration response to the local downturn. Solid red (blue) lines correspond to oil-exposed (non-oil) cities in the baseline economy. Dotted lines impose the “SS House Price” counterfactual and dashed lines impose the “SS House Price & Rent” counterfactual. Out-migration responses are reported as percent deviations from the stationary-equilibrium out-migration rate.

Figure J.3: Out-Migration Response to the Local Downturn by Housing Tenure



Note: Figure J.3 plots the out-migration response to the local downturn by housing tenure: renters (Panel A) and homeowners (Panel B). Within each panel, solid red (blue) lines correspond to oil-exposed (non-oil) cities in the baseline economy. Dotted lines impose the “SS House Price” counterfactual and dashed lines impose the “SS House Price & Rent” counterfactual. Out-migration responses are reported as percent deviations from the stationary-equilibrium out-migration rate.

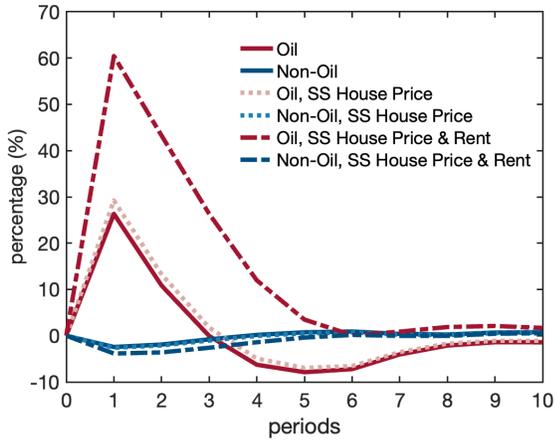
Figure J.4: Out-Migration Response to the Local Downturn by Age Group



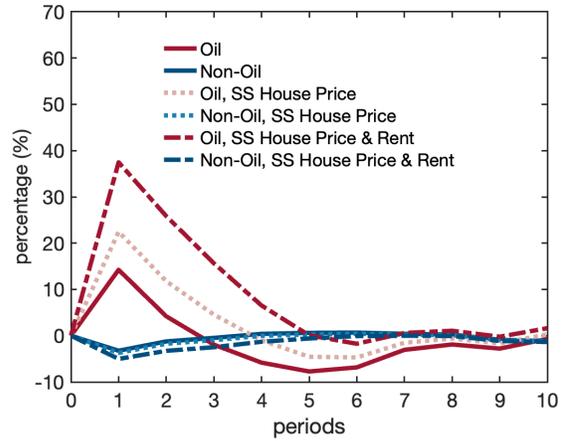
Note: Figure J.4 plots the out-migration response to the local downturn by age group: Panel A ages 25–35, Panel B ages 36–45, Panel C ages 46–65, and Panel D ages 66–75. Within each panel, solid red (blue) lines correspond to oil-exposed (non-oil) cities in the baseline economy. Dotted lines impose the “SS House Price” counterfactual and dashed lines impose the “SS House Price & Rent” counterfactual. Out-migration responses are reported as percent deviations from the stationary-equilibrium out-migration rate.

Figure J.5: Out-Migration Response to the Local Downturn by Net Worth Group

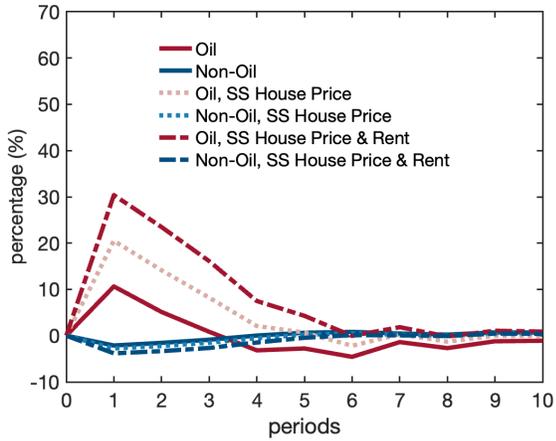
Panel A: Lowest Net Worth Quartile



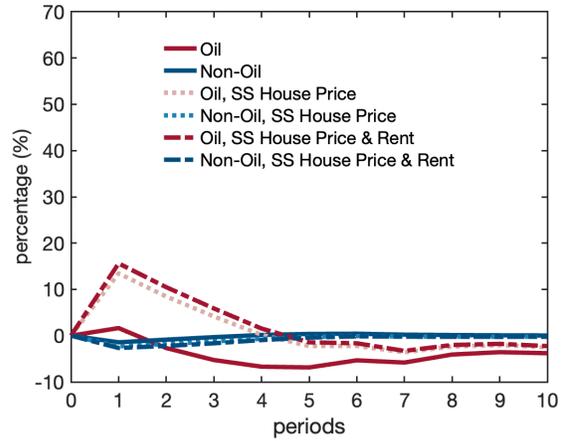
Panel B: Second Net Worth Quartile



Panel C: Third Net Worth Quartile



Panel D: Highest Net Worth Quartile

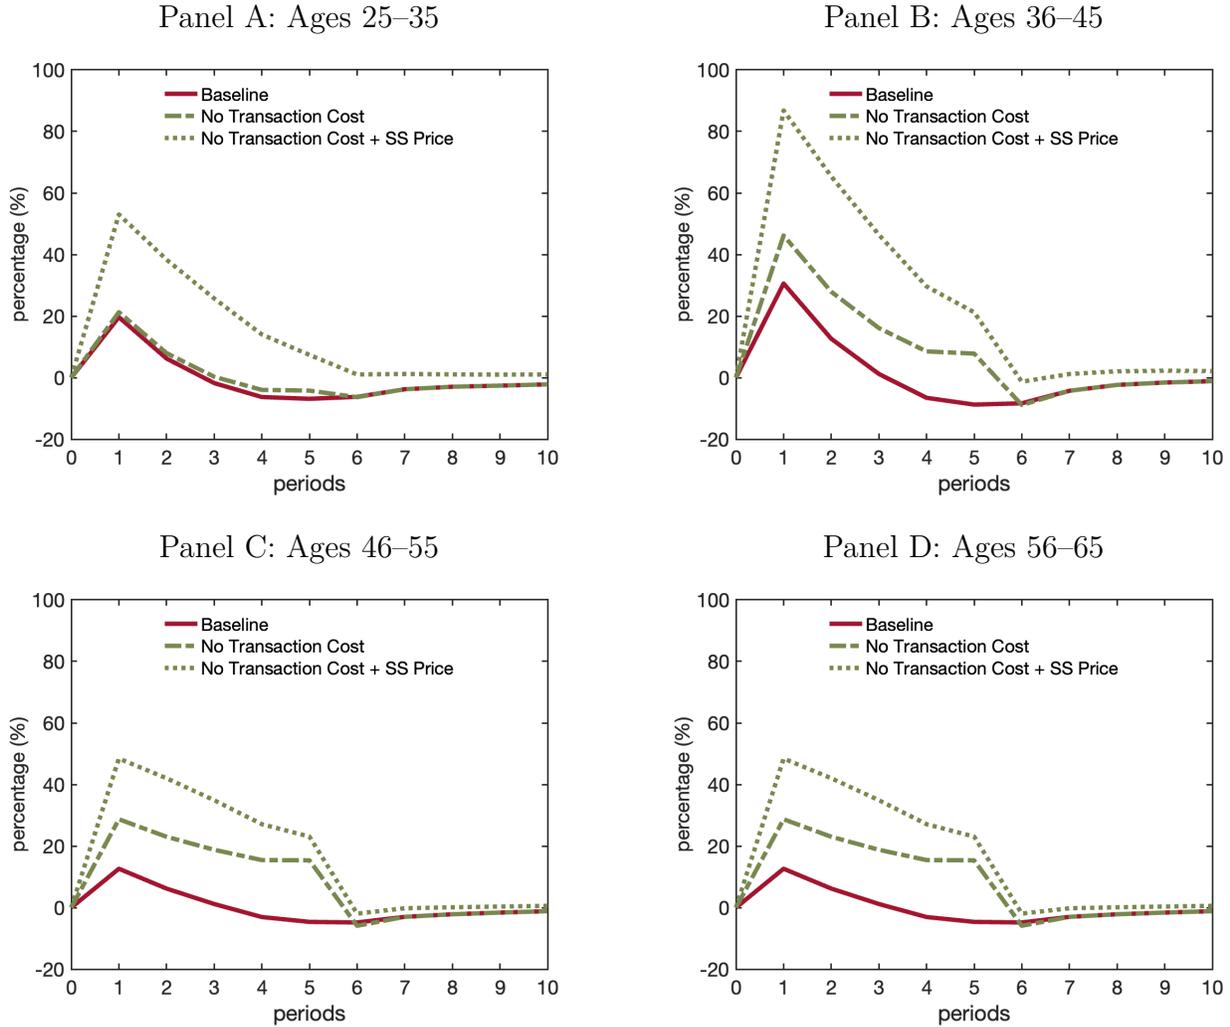


Note: Figure J.5 plots the out-migration response to the local downturn by net-worth quartile: Panel A (lowest), Panel B (second), Panel C (third), and Panel D (highest). Within each panel, solid red (blue) lines correspond to oil-exposed (non-oil) cities in the baseline economy. Dotted lines impose the “SS House Price” counterfactual and dashed lines impose the “SS House Price & Rent” counterfactual. Out-migration responses are reported as percent deviations from the stationary-equilibrium out-migration rate.

J.1 Lowering the Housing Transaction Costs

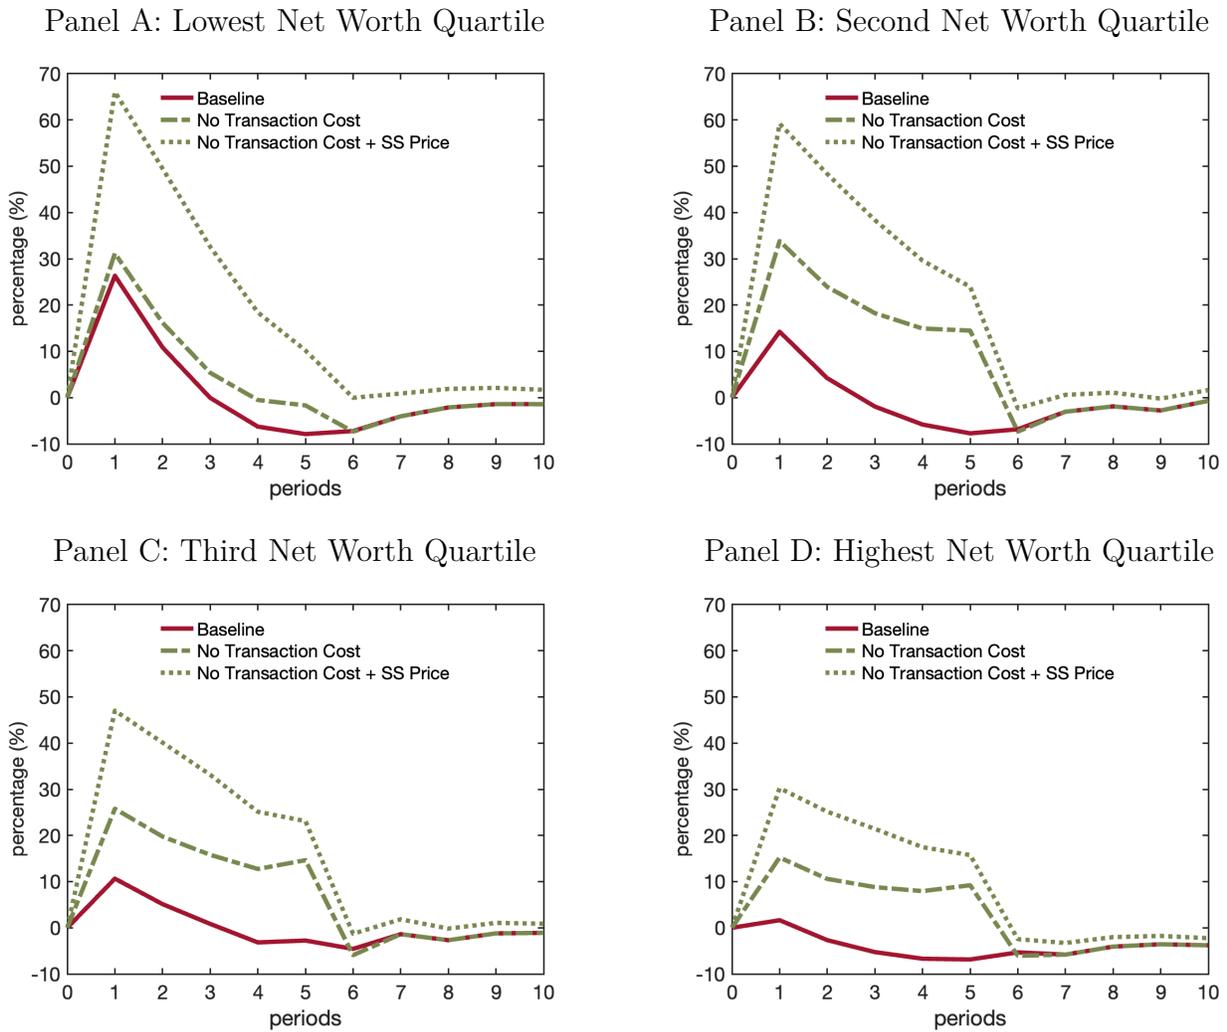
This section reports additional results on the impact of a local downturn on out-migration and welfare in the absence of housing transaction costs.

Figure J.6: Out-Migration Response to the Local Downturn by Age Group - Lowering the Housing Transaction Costs



Note: Figure J.6 plots the out-migration response of oil-exposed cities under the same counterfactuals as Figure 9, by age group: Panel A ages 25–35, Panel B ages 36–45, Panel C ages 46–55, and Panel D ages 56–65. Within each panel, the solid red line reports the baseline economy, the dashed red line imposes “SS House Price & Rent,” and the dotted green line imposes “Lower Transaction Cost.” Out-migration responses are reported as percent deviations from the stationary-equilibrium out-migration rate.

Figure J.7: Out-Migration Response to the Local Downturn by Net Worth Group - Lowering the Housing Transaction Costs



Note: Figure J.7 plots the out-migration response of oil-exposed cities under the same counterfactuals as Figure 9, by net-worth quartile: Panel A (lowest), Panel B (second), Panel C (third), and Panel D (highest). Within each panel, the solid red line reports the baseline economy, the dashed red line imposes “SS House Price & Rent,” and the dotted green line imposes “Lower Transaction Cost.” Out-migration responses are reported as percent deviations from the stationary-equilibrium out-migration rate.

Table J.1: The Welfare Effect of a Local Downturn – Lowering the Housing Transaction Costs

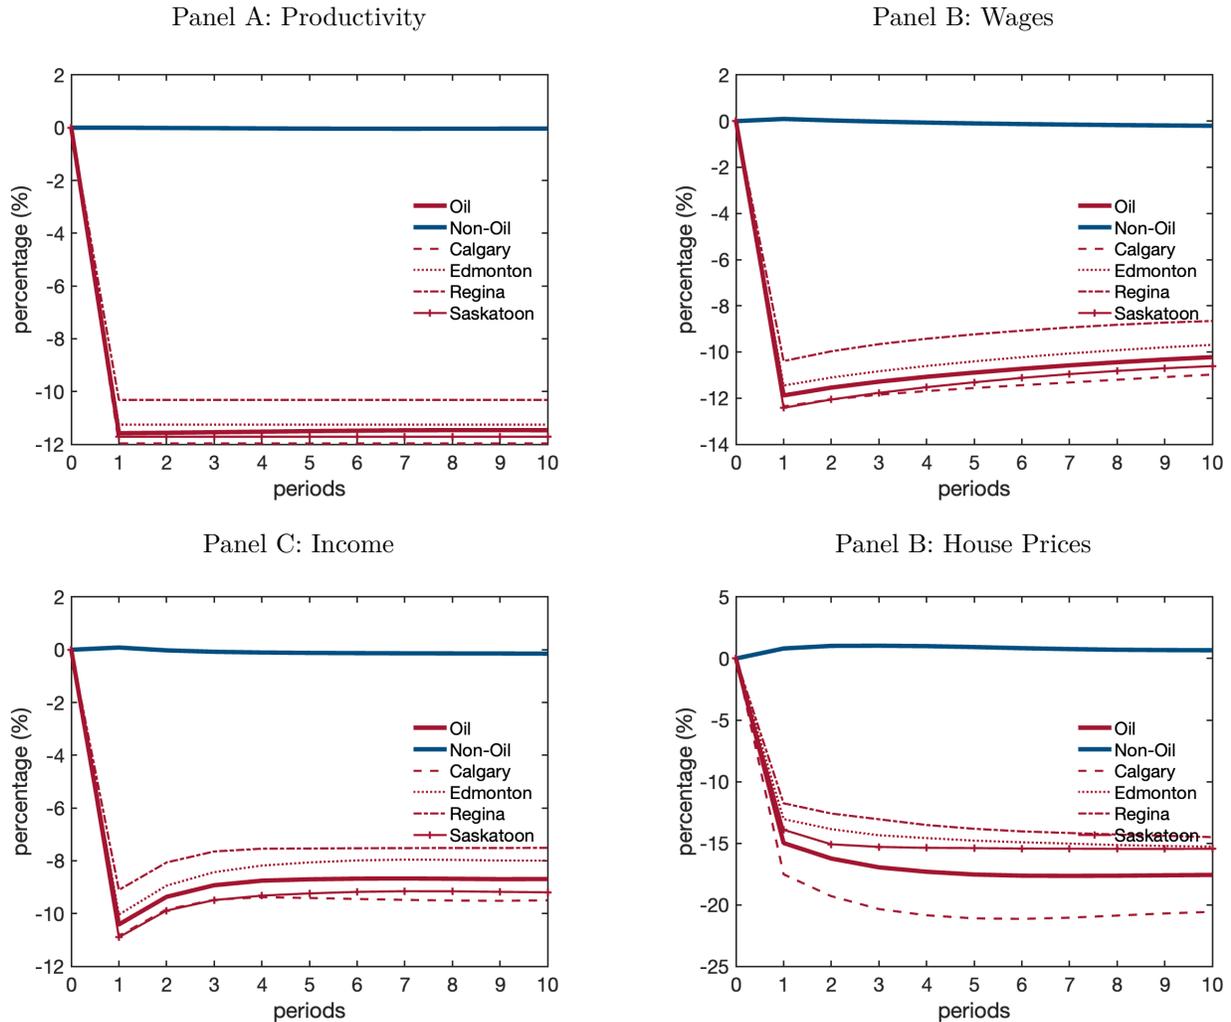
Demographics	Baseline (level,% of lifetime C)		Lower HH Costs ($\Delta\%$ vs baseline)		Lower HH Costs, Const. HP ($\Delta\%$ vs baseline)	
	Oil Region	Canada	Oil Region	Canada	Oil Region	Canada
All	-1.52	-0.23	10.53	8.70	24.34	21.74
Homeowners	-1.76	-0.27	9.09	7.41	54.55	51.85
Renters	-1.17	-0.18	12.82	11.11	-41.03	-33.33
Age 25–34	-1.58	-0.27	14.56	11.11	-8.23	-7.41
Age 35–44	-1.69	-0.23	7.10	4.35	44.97	39.13
Age 45–64	-1.26	-0.16	4.76	6.25	76.19	68.75
Age > 64	-0.93	-0.10	1.08	0.00	100.00	90.00
Net worth – Q1	-1.31	-0.21	10.69	9.52	-28.24	-23.81
Net worth – Q2	-1.61	-0.23	11.18	8.70	22.36	17.39
Net worth – Q3	-1.78	-0.25	12.36	12.00	49.44	44.00
Net worth – Q4	-1.39	-0.24	3.60	4.17	67.63	62.50

Note: Table J.1 reports the welfare effects of the decline in productivity generated by the oil crisis. Columns 1–2 show consumption-equivalent welfare losses (in percent of lifetime consumption) in the baseline economy for households in the oil regions and in Canada as a whole. Columns 3–4 and 5–6 report, respectively, the percentage change in the magnitude of these welfare losses when (i) housing transaction costs are paid by the government, and (ii) housing transaction costs are paid and house prices are held constant. Percentage changes are computed as $100 \times (\Delta W^{\text{alt}} - \Delta W^{\text{baseline}}) / |\Delta W^{\text{baseline}}|$, so positive values indicate smaller (less negative) welfare losses than in the baseline, while negative values indicate larger losses. Welfare changes are computed using equation (22).

J.2 A Permanent Version of the Local Downturn

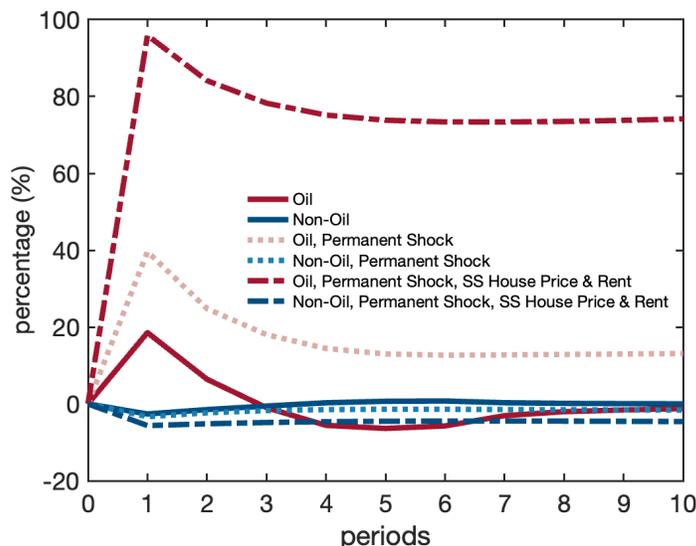
This Appendix presents the out-migration response and welfare implications in the case where the local downturn is permanent rather than transitory.

Figure J.8: Local Permanent Productivity Shocks, Wages, Income and House Prices



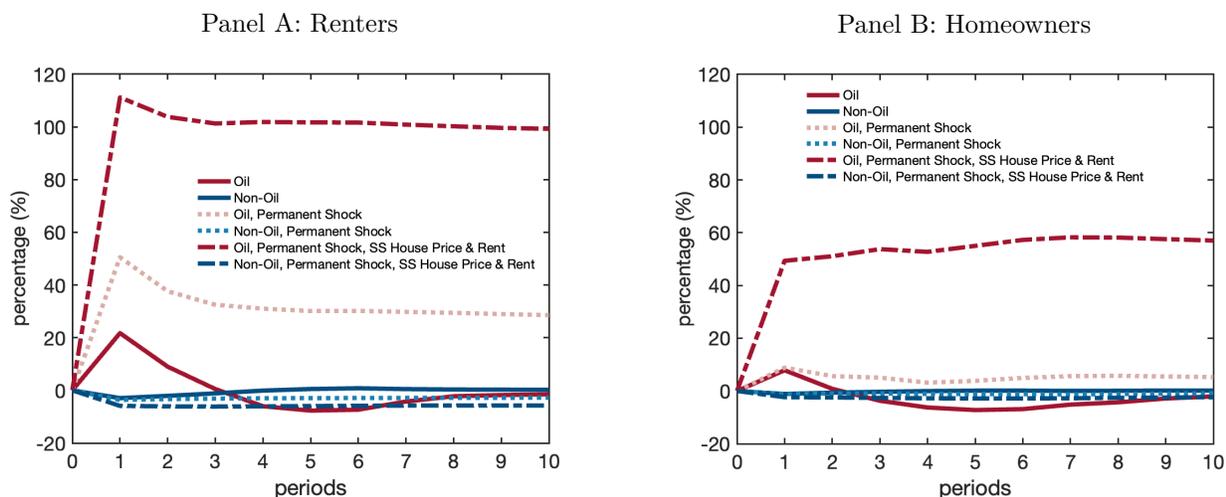
Note: Figure J.8 reports the evolution of aggregate variables in response to the permanent oil shock. Panel A reports the paths of the exogenous productivity component. Panels B, C, and D plot the endogenous paths of local wages, average income, and house prices induced by the shocks shown in Panel A. All series are expressed as percentage deviations from the stationary pre-shock equilibrium.

Figure J.9: Out-Migration Response to a Permanent Local Downturn



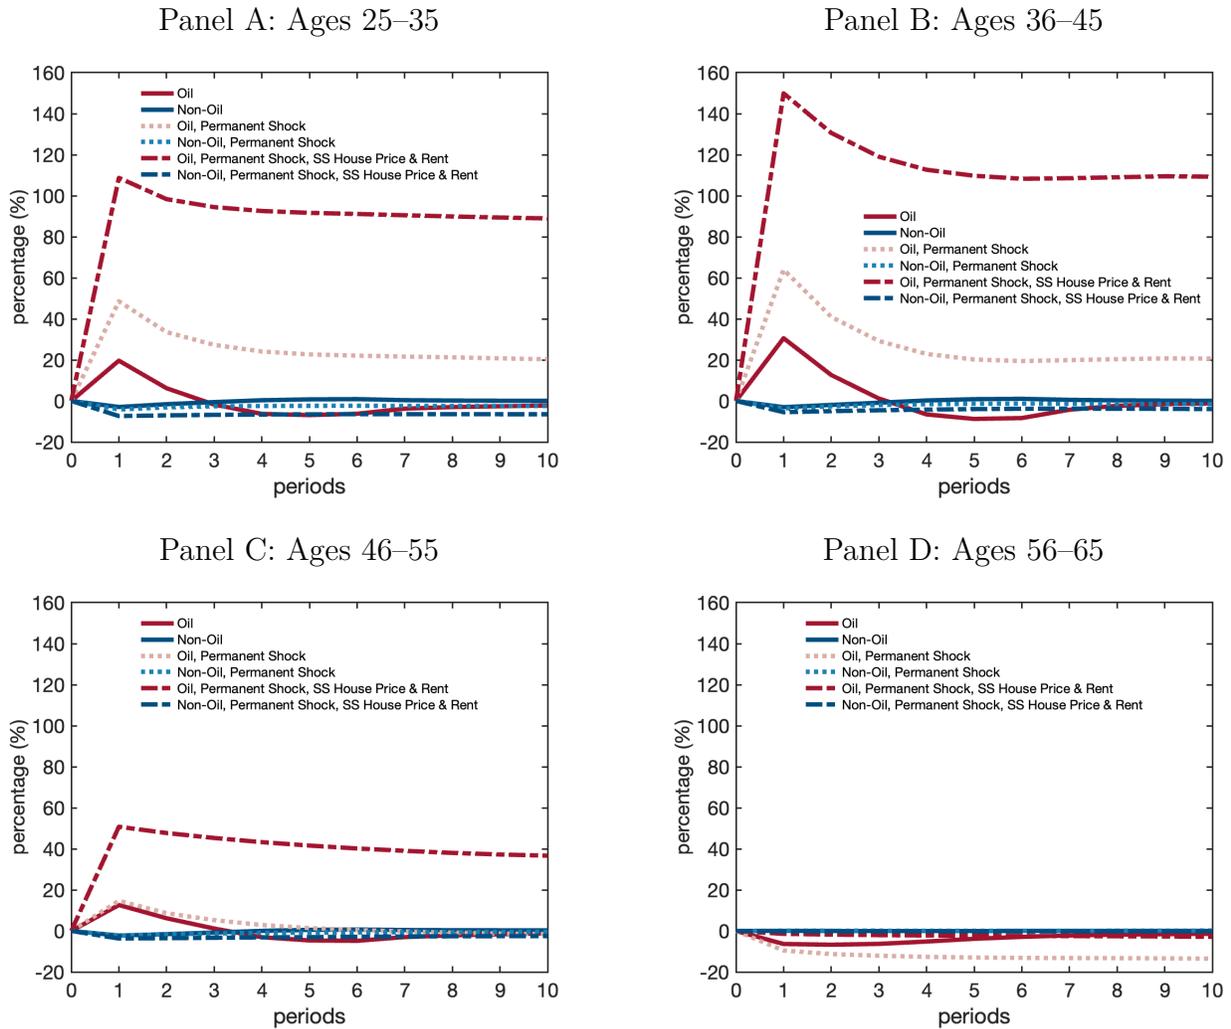
Note: Figure J.9 shows the evolution of out-migration response to the oil shock in oil cities vs non-oil cities if the shock had been permanent. Different from Figure 8, this figure also reports a counterfactual version in which the distribution of individuals is kept fixed to the pre-shock period (in pink dashed for oil regions and in light blue dashed for non-oil regions). The results are reported as percentage-point deviations from the stationary pre-shock equilibrium.

Figure J.10: Migration Response to a Permanent Local Downturn



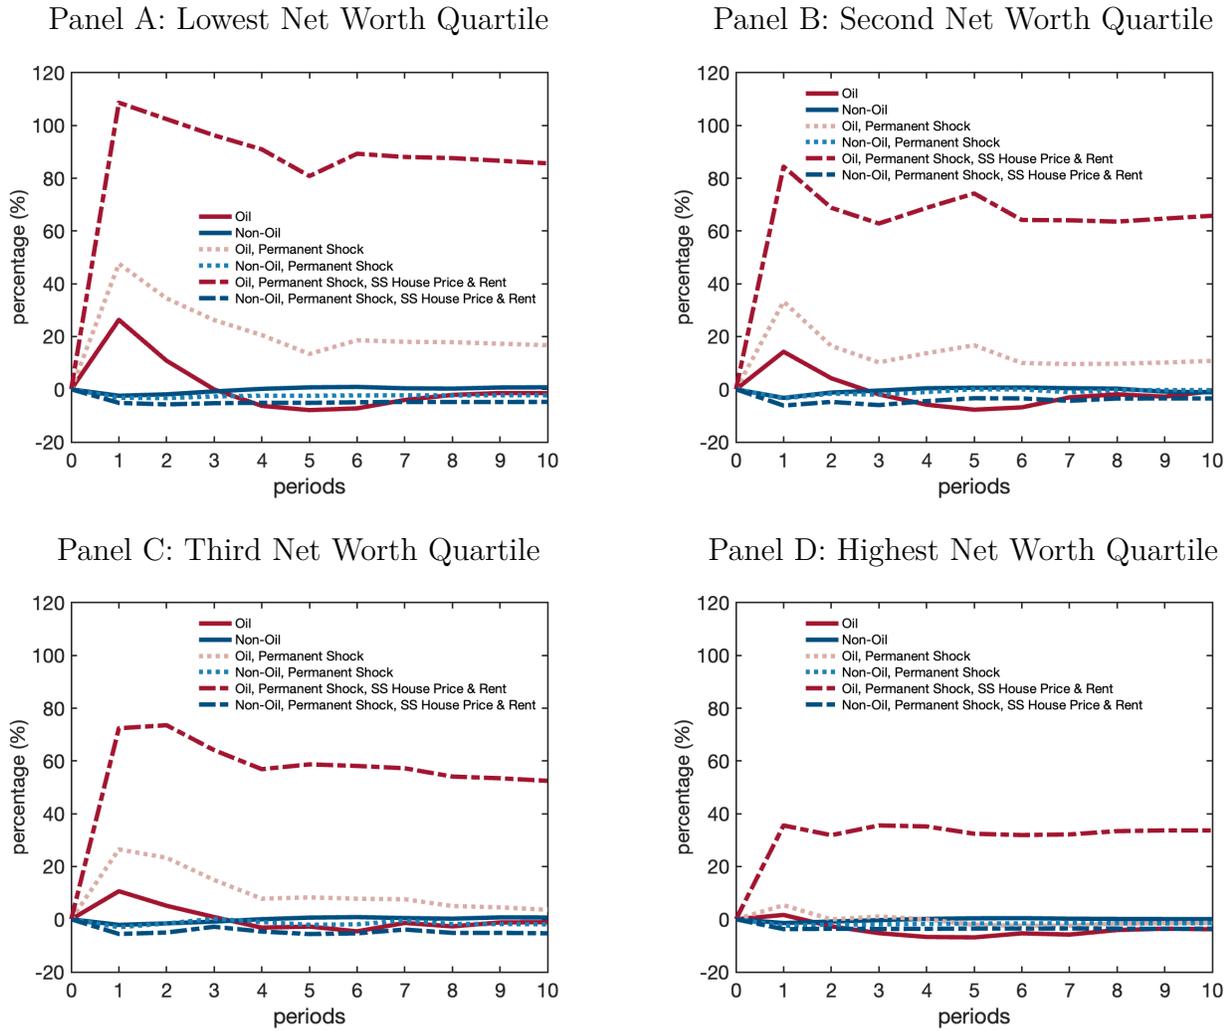
Note: Figure J.10 plots the model-implied out-migration response to a permanent local downturn for the aggregate population (Panel A) and for homeowners (Panel B). In each panel, the red line reports oil-exposed cities and the blue line non-oil cities. Out-migration responses are reported as percent deviations from the stationary pre-shock equilibrium.

Figure J.11: Out-Migration Response to a Permanent Local Downturn by Age Group



Note: Figure J.11 plots the out-migration response to a permanent local downturn by age group: Panel A ages 25–35, Panel B ages 36–45, Panel C ages 46–55, and Panel D ages 56–65. In each panel, the red line reports oil-exposed cities and the blue line non-oil cities. Out-migration responses are reported as percent deviations from the stationary pre-shock equilibrium.

Figure J.12: Out-Migration Response to a Permanent Local Downturn by Net Worth Group



Note: Figure J.12 plots the out-migration response to a permanent local downturn by net-worth quartile: Panel A (lowest), Panel B (second), Panel C (third), and Panel D (highest). In each panel, the red line reports oil-exposed cities and the blue line non-oil cities. Out-migration responses are reported as percent deviations from the stationary pre-shock equilibrium.

Table J.2: The Welfare Effect of a Permanent Local Downturn

Demographics	Baseline (level,% of lifetime C)		Permanent shock ($\Delta\%$ vs baseline)	
	Oil Region	Canada	Oil Region	Canada
All	-1.52	-0.23	-144.08	-147.83
Homeowners	-1.76	-0.27	-177.27	-144.44
Renters	-1.17	-0.18	-75.21	-155.56
Age 25–34	-1.58	-0.27	-113.92	-159.26
Age 35–44	-1.69	-0.23	-163.91	-143.48
Age 45–64	-1.26	-0.16	-190.48	-131.25
Age > 64	-0.93	-0.10	-235.48	-110.00
Net worth – Q1	-1.31	-0.21	-83.21	-133.33
Net worth – Q2	-1.61	-0.23	-147.83	-156.52
Net worth – Q3	-1.78	-0.25	-169.10	-144.00
Net worth – Q4	-1.39	-0.24	-197.12	-154.17

Note: Table J.2 reports the welfare effects of the decline in productivity generated by the oil crisis under the baseline (temporary) shock and under a counterfactual permanent shock. Columns 1–2 show consumption-equivalent welfare losses (in percent of lifetime consumption) in the baseline economy for households in the oil regions and in Canada as a whole. Columns 3–4 report the percentage change in the magnitude of these welfare losses under the permanent-shock counterfactual, computed as $100 \times (\Delta W^{\text{perm}} - \Delta W^{\text{baseline}})/|\Delta W^{\text{baseline}}|$; negative values indicate larger (more negative) welfare losses than in the baseline. Welfare changes are computed using equation (22).

References

- Auclert, Adrien, Hannes Malmberg, Frederic Martenet, and Matthew Rognlie.** 2021. “Demographics, Wealth, and Global Imbalances in the Twenty-first Century.” National Bureau of Economic Research.
- Boar, Corina, and Virgiliu Midrigan.** 2022. “Efficient redistribution.” *Journal of Monetary Economics*, 131: 78–91.
- Davidoff, T.** 2016. “Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors.” *Critical Finance Review*, 5(2): 177–206.
- Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson.** 2021. “Housing wealth effects: The long view.” *The Review of Economic Studies*, 88(2): 669–707.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante.** 2020. “The Housing Boom and Bust: Model Meets Evidence.” *Journal of Political Economy*, 128(9): 3285–3345.
- Kennan, John, and James R Walker.** 2011. “The effect of expected income on individual migration decisions.” *Econometrica*, 79(1): 211–251.
- Kilian, Lutz, and Xiaoqing Zhou.** 2022. “The propagation of regional shocks in housing markets: Evidence from oil price shocks in Canada.” *Journal of money, credit and banking*, 54(4): 953–987.
- Molloy, Raven, Christopher L. Smith, and Abigail K. Wozniak.** 2014. “Declining Migration within the U.S.: The Role of the Labor Market.” *National Bureau of Economic Research*, 20065.
- Rouwenhorst, K. G.** 1995. “Asset Pricing Implications of Equilibrium Business Cycle Models.” *Frontiers of business cycle research*, Princeton University Press: 294–330.
- Saiz, A.** 2010. “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics*, 125(3): 1253–1296.
- Sinai, Todd M.** 2012. “House price moments in boom-bust cycles.” National Bureau of Economic Research.