The Dynamics of Internal Migration: A New Fact and its Implications*

Greg Howard[†] Hansen Shao[‡]

2026

Abstract

We propose a new model of internal migration, based on persistent and spatially-correlated idiosyncratic utility. The model is motivated by a new fact in the data that simple moving cost models struggle to match: the t-year interstate migration rate is proportional to the square root of t. The new model maintains the tractability and flexibility of standard migration models, but better matches the dynamics of migration, including the new fact. It has substantially different welfare implications and makes different counterfactual predictions, especially in terms of dynamic adjustment and long-run responses.

Keywords: regional evolution, misallocation, gravity equation, labor mobility, moving costs

JEL Codes: R23, R13, J61

*This paper was previously circulated as "Internal Migration and the Microfoundations of Gravity." We would like to thank the editor Thomas Chaney, three anonymous referees, Nelson Lind, Javier Quintana, Kyle Mooney, Yao Wang, Vivek Bhattacharya, Gregor Schubert, Andrii Parkhomenko, Eduardo Morales, Marieke Kleemans, Kyuri Park, Cihang Wang, Jaysa Rafi, Humberto Martinez-Garcia and participants at the Illinois Macro Lunch, Illinois Young Applied Faculty Lunch, the European UEA meetings, Midwest Macro, the AREUEA National Meetings, the Stanford Institute for Theoretical Economics Housing & Urban Economics Conference, the Online Spatial and Urban Seminar, the Chicago Fed Urban Conference, the Cleveland Federal Reserve, and the North American UEA meetings for constructive feedback. We also would like to thank Jialan Wang, Julia Fonseca, and Peter Han for creating the Gies Consumer and Small Business Credit Panel and the Gies College of Business for supporting this dataset. For the Panel Survey of Income Dynamics, the collection of data used in this study was partly supported by the National Institutes of Health under grant number R01 HD069609 and R01 AG040213, and the National Science Foundation under award numbers SES 1157698 and 1623684.

[†]University of Illinois Urbana-Champaign, glhoward@illinois.edu.

[‡]University of Illinois Urbana-Champaign.

1 Introduction

We document a new empirical regularity: the t-year interstate migration rate, defined as the share of people living in a different state than they did t years ago, scales quite closely with the square root of t. This new fact is a puzzle for the widely used moving cost model, which typically implies a linear relationship.

The main contribution of our paper is a simple but novel model that can match the new fact by assuming that idiosyncratic utility is correlated across time and space. The model we propose features Spatially and Persistently Autocorrelated Epsilons, so we refer to it as the "SPACE" model.[1] Unlike the standard model, which interprets low migration rates as the result of large moving costs, our model rationalizes low migration rates as a result of high persistence in idiosyncratic determinants of location. Our model is able to replicate bilateral one-year migration flows from the data, maintaining the flexibility of the standard models, while also featuring more realistic dynamics.

We compare the implications of the new model to the moving cost model and find that, for many but not all questions, the models draw different conclusions. These findings reshape our understanding of the causes of low migration, the dynamics of population adjustment, the long-run population elasticities to local changes, and the changes in implied utilities across space in recent years.

The first part of our paper documents the new square root fact using data from the Gies Consumer and Small Business Credit Panel (GCCP), a 15-year panel recording the location of approximately 1 percent of all Americans with a credit report every year. While the square root fact is related to the well-known fact that return migration is common (Kennan and Walker, 2011), we show that the $\sqrt{t}$ fact is not a simple result of return migration but captures richer dynamics. We also show that this fact does not naturally occur in standard moving cost models. In those models, location choice is a Markov process, which, combined with low rates of migration, implies that the t-year migration rate has an approximately linear relationship with t.[2]

We then build a model that can reconcile this new fact while maintaining the tractability and flexibility of the standard model. Building on the model of McFadden (1978), we

[1] ϵ is the common notation for the random component in a random utility model.

[2] With flexible enough moving costs that depend on the past history of location choices, a model could match almost any relationship between the t-year migration rate and t. Nonetheless, we still think the square root fact is a puzzle for two reasons: first, the models that people actually use do not explain it, possibly due to the large state-space of very flexible models; and second, because when a model can match any relationship, there is no reason for it to match this particular relationship. Our model provides a rationale for this particular relationship. We discuss this further in Section 2.3.

consider a generalized-extreme-value discrete-choice model to introduce correlation over space and time. The SPACE model leads to closed-form solutions for state populations and interstate migration. One important result of this model is that the cross-state population elasticities—a key statistic for quantitative spatial modeling—are directly proportional to the bilateral migration rate. We show a way to calibrate the model that allows the use of simple formulas for population changes, i.e. "exact hat" algebra, and which also allows computationally feasible simulations of individuals' location choices over time. This tractability allows for the SPACE model to easily serve as the migration block of more complex quantitative dynamic spatial models.[3]

Next, we compare the implications of the SPACE model to those of the moving cost model. For many questions, we show that choosing how to model migration is not innocuous but leads to important differences in how economists answer central questions about location choice and migration.

First, we demonstrate that the SPACE model is better at predicting future locations of individuals. We compare the forecasting performance of each model using the Kullback-Leibler divergence, and demonstrate that the SPACE model does better at predicting out-of-sample locations, especially at longer horizons. This is consistent with the idea that the SPACE model is able to match realistic dynamics of location choice.

Second, the SPACE model and the moving cost model have very different perspectives on why people do not move. Moving cost models estimate large moving costs (e.g. Kennan and Walker, 2011; Bryan and Morten, 2019; Giannone, Li, Paixao and Pang, 2020; Zerecero, 2021). In contrast, the SPACE model does not need moving costs at all to rationalize observed levels of migration in the data.

Third, we turn to macroeconomic questions, which typically depend on the elasticity of local populations to local utility. We show that both models feature similar short-run population cross-elasticities, in that the elasticity of population in state i with respect to utility in state j is approximately proportional to the gross migration rate between the two states.[4] In other words, if the purpose of a model with migration is to predict short-run effects on populations, both models deliver similar results.

[3] For example, in Appendix E, we embed the SPACE model into a model with local housing production to study transition dynamics.

[4] Even though the models have similar elasticities, the rationale for why the elasticity is related to migration is a bit different. In the SPACE model, the rationale is that migrants are close to indifferent between living in each state, so the mass of people who will move in response to a small shock is proportional to the number of migrants. In the moving cost model, the rationale is that the extreme value function has a functional form such that the number of people on the margin is proportional to the number of people who make that choice.

However, in the long run, population cross-elasticities are quite different across the two models. In the SPACE model, population elasticities are the same in the short-run and the long-run, meaning that long-run elasticities are still proportional to the migration between the two places. But in the moving cost model, the long-run population elasticities are approximately the same as a static discrete choice logit model, i.e. the elasticity is proportional to the population share of the shocked region. In the data, population shares and gross migration rates have little correlation, so the long-run elasticities of the two models are also uncorrelated.

A fourth difference follows naturally from the previous one: the dynamics of regions' population changes are quite different in the two models. In the SPACE model, the dynamics are simple. In response to a permanent utility change, the population adjusts fully, contemporaneous to the utility shock. But in the moving cost model, the dynamics are relatively slow and can be unintuitive. In that model, every period, a new set of people receives a sufficiently large enough shock to move, so a permanent utility shock raises the migration rate and population adjusts slowly. Furthermore, because the long-run elasticities are related to population shares, not migration shares, states distant from the shock adjust particularly slowly, while nearby states adjust quickly and sometimes overshoot the new long-run steady state.

Finally, the SPACE model and the moving cost model interpret the data very differently in terms of which states have become higher-utility over time. With the SPACE model, we use standard exact-hat techniques to map observed population changes onto implied utility changes across time. We can also do the same with the moving cost model. When we use U.S. population data to infer which states are gaining in terms of relative utility, we draw substantially different conclusions depending on the chosen model. This is critical if we want to estimate the role of policy or economic shocks on welfare.

We finish the paper by discussing the importance of the differences between the SPACE and moving cost model in the context of the literature. We show how the differences we highlight are central to some of the questions that are asked in the dynamic spatial literature. We discuss whether various approaches to enrich the moving cost model would deliver similar results to the SPACE model.

The specific features of the SPACE model may at first appear to be reverse engineered to match the new square root fact while lacking a basis in reality. However, we argue that the SPACE model has two very realistic features of preferences. First is that the match-specific utility for location is persistent over time. Surveys suggest that people primarily cite family and employment considerations as reasons for interstate moves (Jia, Molloy, Smith and Wozniak, 2023).[5] People's feelings about these networks are surely correlated over time, and it is an empirical fact that each of these incentives is persistent in terms of location. Second is that match-specific utility is spatially correlated. Considering people's stated preferences, the ability to live near family is highly-correlated across space. If state i is close to family, then states near i are also close to family. Likewise, the jobs available to people in specific industries or with specific skills are geographically concentrated. Natural amenities or regional cultures—other possible sources of idiosyncratic utility—are also spatially correlated. The precise distributions we use to represent these correlations in the SPACE model are indeed convenient mathematically, but it is hard to support the argument that spatially and autocorrelated idiosyncratic utility is less realistic than the independent and identically distributed (i.i.d.) utilities of a moving cost model.

To clarify the contributions of the model, we wish to be specific regarding the difference between moving costs and persistence in match-specific utility. While mathematically straightforward to specify (as we do here), it is important to understand what each term means when mapped onto the real world. A typical moving cost model involves a one-time irreversible cost borne by people who leave one area for another. In contrast, persistent match-specific utility means that the change in utility when a person moves from one location to the other is both persistent over time and partially reversible should the person move back to the original location.[6] A moving truck and the psychological cost of throwing a goodbye party clearly are moving costs. But many factors described as "costs" in the literature are easily reversible, although the forgone benefits may decay with time. Living far from friends or a particularly amenity is a persistent and ongoing burden rather than a one-time cost. Even though such factors are often called "moving costs" in the literature, we think that terminology is used because existing models have not been able to distinguish persistence in match-specific utility from moving costs. The rest of this paper will give many reasons why this distinction is important.

[5]Koenen and Johnston (2024) documents that social networks have causal effects on migration behavior.

[6]In Appendix D, we estimate a model that includes both persistence in unobserved heterogeneity and moving costs. The model is not nearly as tractable as either model alone. The estimated model is able to match dynamics of migration even better than the SPACE model without moving costs. However, on several important dimensions, the differences between the SPACE model and the SPACE model with moving costs are quite small.

1.1 Literature

How individuals choose where to live is a classic question in the urban economics literature. Many urban models assume utility is equalized across space in the tradition of Rosen (1979) and Roback (1982). Other more quantitative models assume a discrete choice framework for locations to answer a variety of questions, such as the role of endogenous amenities on location choice (Diamond, 2016) or spatial misallocation on aggregate output (Hsieh and Moretti, 2019).

A growing share of this literature has explicitly looked at the dynamics of location choice, that is, migration. Since at least Blanchard and Katz (1992), migration has been recognized as a key feature in how regions adjust to economic shocks. In this vein, papers studying the rise or decline of regional economies place significant emphasis on migration (Caliendo, Dvorkin and Parro, 2019; Allen and Donaldson, 2020; Morris-Levenson and Prato, 2022), and especially the speed of net migration (Glaeser and Gyourko, 2005; Kleinman, Liu and Redding, 2023; Amior and Manning, 2018; Davis, Fisher and Veracierto, 2021). Similarly, when aggregating to the macroeconomic level, migration is critical to the speed of adaption to changing technologies or external shocks (Tombe and Zhu, 2019; Hao, Sun, Tombe and Zhu, 2020; Eckert and Peters, 2022; Giannone, 2017; Heise and Porzio, 2021; Bryan and Morten, 2019). A growing literature has emphasized how migration, including internal migration, plays an important role in adapting to global warming (Rudik, Lyn, Tan and Ortiz-Bobea, 2021; Cruz and Rossi-Hansberg, 2024; Oliveira and Pereda, 2020). Migration also represents an important margin when analyzing housing markets in particular (Schubert, 2021).[7] Central to many of these questions is the elasticity of local populations to different shocks over various time horizons. One of the contributions of this paper is to examine how robust those conclusions are to alternative ways of modeling migration.

Corresponding to the growth of interesting questions related to migration, there have also been advances in methods of modeling migration. Kennan and Walker (2011) formulated the canonical model of migration using the dynamic logit formulation. Kaplan and Schulhofer-Wohl (2017), Giannone et al. (2020), Porcher (2020), Mangum and Coate (2019), Zerecero (2021), and Monras (2020) have built on this formulation to incorporate additional realistic features of migration, such as richer information frictions, migrant wealth, home bias, and nested decision making.[8] Other approaches, such as Coen-Pirani

[7] Howard and Liebersohn (2021) and Howard, Liebersohn and Ozimek (2023) also study the effects of changing location choice on housing markets, but model location choice in a static discrete choice framework rather than explicitly having a notion of migration.

[8] In particular, Kaplan and Schulhofer-Wohl (2017) argues that changes in information frictions can

(2010) and Davis et al. (2021) do not use the dynamic logit framework, but have similar discrete choice models that improve the tractability in a way specific to their goals. All of these models use moving costs to explain the low rates of migration, and potentially adjust those moving costs to explain the high rates of return migration. In contrast, only one paper to our knowledge uses persistence in unobservable match-specific utility to explain low migration rates: Bayer and Juessen (2012). However, its model is too complex to tractably scale beyond two regions, limiting its use in many empirical applications.

One type of persistent match-specific utility has been modeled by Zabek (2024), Mangum and Coate (2019) and Zerecero (2021), by incorporating a preference for living in one's birthplace.[9] These models share similarities with the SPACE model, in that they also tend to feature smaller moving costs (Zerecero, 2021) and would intuitively feature more return migration than the standard moving cost model (Mangum and Coate, 2019). However, adding birthplace preferences to a moving cost model does not reproduce the square root fact and does not substantially alter many of the distinctions between the two models.

At the same time that models of internal migration have become more popular, there has also been new empirical evidence on the determinants of and constraints to migration. For example, Saks and Wozniak (2011) shows that migration is cyclical; Kleemans (2015) studies the income shocks that cause migration; Farrokhi and Jinkins (2024) examines the attachment hypothesis using a policy change among Danish refugees; Koşar, Ransom and Van der Klaauw (2021) uses a survey experiment to study how people make location decisions; and Fujiwara, Morales and Porcher (2022) proposes a methodology for uncovering information frictions in location choice. Our paper also contributes to this literature by establishing a new stylized fact that the t-year migration rate is proportional to $\sqrt{t}$.

help explain the decline in interstate migration, along with decreases in the different returns to various occupations across space. Giannone et al. (2020) builds a rich model of migration that incorporates wealth and borrowing, to analyze how credit and savings can affect if and where people choose to move. Porcher (2020) builds a tractable model of rational inattention in the dynamic migration context to argue that information frictions are one of the main reasons people do not move. Mangum and Coate (2019) models biases for birthplace and long-tenured locations, and they use their model to argue that a shift of the American population to the West and to the South is responsible for slowing labor mobility. Zerecero (2021) also examines a model that includes a preference for birthplace. Monras (2020) looks at the asymmetric response of inmigration and outmigration to local shocks, and builds a dynamic nested logit model to better capture the phenomenon.

[9]The canonical model in Kennan and Walker (2011) also includes a premium for birthplace.

2 New fact

In this section, we present the new square root fact and argue that it is a puzzle for existing models.

2.1 Data

Throughout the paper, we primarily measure migration using the Gies Consumer and Small Business Credit Panel (2004-2018), which is credit data from one of the leading credit report providers. It is suited to this study due to its large sample size and its panel dimension. When we can, we verify the empirical patterns using the IRS migration data, the American Community Survey, or the Panel Survey of Income Dynamics (IRS Migration Data, 2004-2018; Ruggles, Genadek, Goeken, Grover and Sobek, 2015; Panel Survey of Income Dynamics, 1969-1997).[10] The credit data is a 15-year panel of individuals making up a 1 percent random sample of the United States. It records the state of residence in each year, allowing us to calculate migration rates at longer horizons. In appendix C.1, we compare migration patterns in the GCCP and other well-known datasets, verifying the level of migration and the gravity pattern are similar. The GCCP has a one-year migration rate similar to the IRS data, but slightly higher. One contributing factor to the higher migration rate in the credit data may be that coverage of credit reports does not extend to all individuals. In particular, lower income people are less likely to have credit reports and are also less likely to move.[11]

[10]For other papers using the GCCP, see Fonseca (2022), Fonseca and Wang (2022), and Han (2023). DeWaard, Johnson and Whitaker (2019) analyzes a similar credit dataset (the Federal Reserve Bank of New York/Equifax Consumer Credit Panel) on how it can be used to study migration.

[11]While there are some well-known drawbacks to the IRS data, e.g. it is based only on tax filers, it is one of the most comprehensive administrative datasets keeping track of migration. The reasons why migration rates are particularly low in 2014 and high in 2016 are not well understood, as these anomalous values did not show up in other datasets measuring migration (see DeWaard, Hauer, Fussell, Curtis, Whitaker, McConnell, Price, Egan-Robertson, Soto and Castro (2022)).

Similarly, while credit data are not designed as a dataset to study migration, they do have location information, and the bureau gets the addresses from a person's financial accounts. The biggest concern with credit data is that moves may show up with a lag, as people do not always immediately change their addresses with their financial institutions. For our square root fact, we check the robustness to using the Panel Survey of Income Dynamics (1969-1997). The GCCP is an unbalanced panel, with yearly observations occurring in May. For matching migration patterns and rates, we focus on the 2004-2005 period, so we only observe data if they had a credit report in both of those years. For some of the dynamics, we address the unbalanced nature of the panel depending on the moments of the data that we are interested in.

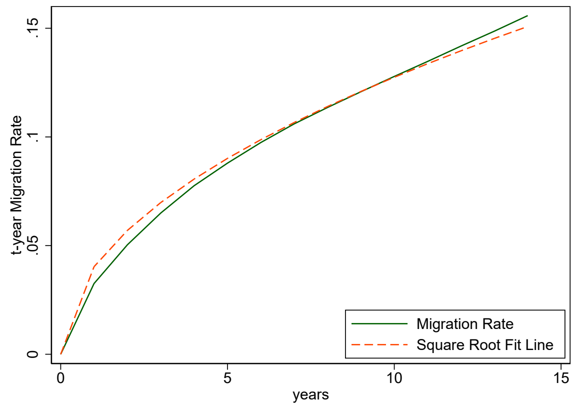

(a) GCCP

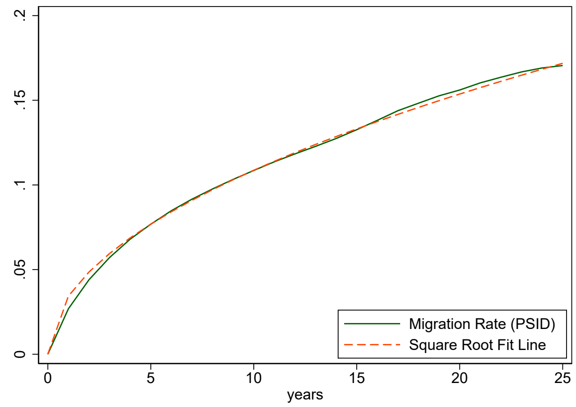

(b) PSID

Figure 1: Migration Rates at Different Horizons. Migration rate at year t is calculated as the percentage of people living in a different state than they did t years ago. Both datasets are unbalanced panels and use any observations in which the state of residence is recorded t years apart. Source: GCCP and PSID.

2.2 Square root fact

Define the t-year migration rate to be the share of people living in a different state than they did t years ago. The new fact is that the t-year migration rate is approximately proportional to $\sqrt{t}$. In Figure 1a, the solid line is the t-year migration rate in the GCCP. The dashed line is a constant times the square root of t, with the constant chosen to match the level of migration. As is apparent from the figure, the shape of the migration rate is very similar to the square root line.[12]

Since we cannot measure dynamic migration moments in the IRS or ACS data, one might wonder if the square root fact is driven by some sort of mismeasurement in the GCCP. In Figure 1b, we show that the square root fact is also present in data from the Panel Survey of Income Dynamics (1969-1997). In fact, when we extend the horizon to 25 years, the square root pattern holds through that longer time period as well.[13]

This fact is not the consequence of averaging across many origins and destinations or many cohorts or many years. In Appendix C.3, we show the distribution of the square root fact across pairs of origins and destinations, ages, cohorts, and starting year. For each of these figures, most of the distribution is concentrated close to the square root line.

[12]Each point is the mean of a binary variable with millions of observations, so if we tried to put standard errors on the graph, they would not be visible. Interested readers can find the standard deviation in Appendix C.2.

[13]Since mismeasurement in the GCCP may be a particular concern for young people, we show in Appendix C.4 that it holds for both people under 45 and people over 45.

We further show that there is not meaningful heterogeneity whether we look at distance between states, the time-frame of our sample, or if we focus on young or old people.

This new fact relates to the more-well-known fact that return migration is common. Many papers in the literature show a significant fraction of workers return to their previous location (e.g. Kennan and Walker, 2011; Kaplan and Schulhofer-Wohl, 2017). One consequence of this fact is that the two-year migration rate is significantly less than twice the one-year migration rate. However, we believe we are the first to document this specific relationship.

2.3 The new fact is a puzzle

The square root fact is notable not only because it is an empirical regularity in need of an explanation, but also because it is at odds with simple existing models. This section shows that the most common model of internal migration in fact leads to an approximately linear relationship between the t-year migration rate and t.

First we outline a standard moving cost migration model which we will then use to show the linear relationship. There are a continuum of individuals of mass 1, denoted by n, who can choose to live in locations denoted by i, j, or k. An agent who lived in i at $t-1$ has utility:

$$V_t(i) = \max_j \{u_{jt} - \delta_{ij} + \epsilon_{jnt} + \beta \mathbb{E} V_{t+1}(j)\} \tag{1}$$

where u_{jt} represents the common utility for everyone living in j at time t, δ_{ij} represents the bilateral moving cost between i and j, and ϵ_{jnt} represents an i.i.d. random variable with an extreme-value distribution. We assume ϵ_{jnt} has a Gumbel distribution with scale parameter 1. If we define $v_{jt} \equiv u_{jt} + \beta \mathbb{E} V_{t+1}(j)$, then the migration probability for an individual living in i is given by:

$$\frac{m_{i\to j,t}}{p_{it}} = \frac{e^{v_{jt}-\delta ij}}{\sum_k e^{v_{kt}-\delta_{ik}}} \tag{2}$$

What does this model predict for the dynamics of migration, especially for the shape of the t-year migration rate? When moving costs are high—which is required to match the low amounts of interstate migration in the data—then the following proposition shows that the t-year migration rate is approximately linear in t.

To set up the proposition, we assume that moving costs are given by $\delta_{ij} = \delta'_{ij} + \Delta$ when $i \neq j$, and $\delta_{ii} = 0$. For $i \neq j$, migration costs consist of δ'_{ij}, a pair-specific component that governs the relative amount of migration to j, and Δ, a common component which governs the overall amount of migration in the economy. This way, when we change Δ, we are not changing the relative amount of migration from i to j versus i to k. In addition, let us define a steady-state to be when the u_{it} are constant over t, and each location's population, p_{it}, is constant over t as well.

Proposition 1. *In the steady-state of a moving cost model, as the common component of moving costs tends to infinity, the t-year migration is proportional to t.*

$$\lim_{\Delta \to \infty} \frac{m^t_{i \to j}}{m^1_{i \to j}} = t$$

where $m^t_{i \to j}$ is the t-year migration from i to j.

The proof can be found in Appendix A.1. This proposition establishes that the square root fact is not a natural consequence of our standard models. In the standard model, we infer high moving costs based on the fact that migration is low, and this proposition establishes that high moving costs imply a linear relationship between the t-year migration rate and t.

In many calibrated models, moving costs are large but less than infinity, so in Appendix C.5, we also check that the model estimated to the data would generate a relationship that looks linear. In simulations, it is also straightforward to add in complications to the model, such as including state variables for the person's location from the prior year, the person's home state, or the age of the person. In Appendix C.5, we explore whether these state variables could generate a square root relationship, but find that the relationship is still quite linear.

Of course, a flexible enough model of moving costs could match the square root pattern. For example, one could match many of the dynamic facts about migration by assuming moving costs increase with the time spent in a particular location and that moving costs are lower when returning to a past location. Similarly, if allow many different unobservable types with heterogeneity in moving costs, we could also generate the square root fact. We have four comments about this line of thought. First, these more flexible models are not the models commonly implemented in the literature. Recent papers that have estimated dynamic effects in spatial models (Caliendo et al., 2019; Kleinman et al., 2023) do not include features that would generate the square root fact. Therefore, much of what the literature knows about the macro dynamics of internal migration is built on models that do not match an important aspect of the micro dynamics.[14] Second, and related to

[14]People do write down models which have excess return migration in the first period (Kennan and

the first point, these more-flexible models lose the tractability of the standard moving cost model. If moving costs depend on a long history of locations, then calculating the macro elasticities depends on keeping track of the size of the population with each of those histories, which is computationally expensive. Third, these models can replicate any relationship between the t-year migration rate and t. To hit the square root fact, you have to parameterize these models just right to not have too much or too little concavity. The model that we propose will always approximately generate a square root, and it does so without the complexity or the flexibility of a model with lots of moving cost heterogeneity. To use the language of Fudenberg, Gao and Liang (2023), these models are not "restrictive" compared to the model presented in the next section.[15] Finally, even if the previous three points are not sufficient to convince the reader to prefer the model in this paper to a very-flexible moving cost model, remaining skeptics should still know about a new model that can fit the micro data equally well and makes different predictions for welfare and counterfactuals. This should inform us about the robustness of the conclusions drawn with the standard model.

3 A new model reconciling the square root fact

3.1 The SPACE model

This section introduces a new model of internal migration which can resolve the square root puzzle from the previous section. Rather than depend on moving costs, it assumes that the match-specific utility (the ϵ's) are spatially correlated and persistent. Because the model features Spatially and Persistently Autocorrelated Epsilons, we call it the SPACE model.

As in the moving cost model, there is a continuum of individuals with mass 1, a finite number of discrete locations, and discrete time. We keep the same notation where n denotes the individual, i the location, and t the year. Individuals choose their location to

Walker, 2011; Kaplan and Schulhofer-Wohl, 2017). But these are also not flexible enough to match the square root fact.

[15]Fudenberg et al. (2023) explains why more restrictive models are desirable compared to less restrictive: "A potential reason for this preference is that models are often meant to capture behavior in related but not identical domains. Given enough data, models that are very unrestrictive will fit any specific dataset well, but may do so by learning idiosyncratic details of those datasets that do not in fact transfer across settings. In contrast, if a highly specific and structured model happens to fit a dataset well, this may generate more confidence that the model's structure extends to other settings."

maximize utility:

$$V_t(\vec{\epsilon}_{nt}) = \max_i \{u_{it} + \epsilon_{nit}\} + \beta \mathbb{E}[V_{t+1}(\vec{\epsilon}_{nt+1})|\vec{\epsilon}_{nt}] \tag{3}$$

where u_{it} is a common flow utility for location i and ϵ_{nit} (the ith element of vector $\vec{\epsilon}_{nt}$) is a person-location-match-specific utility.[16] Note that the choice of location i does not affect the continuation value ($\mathbb{E}[V_{t+1}(\vec{\epsilon}_{t+1})|\vec{\epsilon}_t]$) because there are no moving costs, so the choice is made sequentially each period, and it has no effect on future choices. In other words, $v_{it}(\vec{\epsilon}_{nt})$, which we defined as $u_{it} + \beta \mathbb{E}[V_{t+1}(\vec{\epsilon}_{nt+1})|\vec{\epsilon}_{nt}]$, differs from u_{it} by just a constant that depends on n but not i.

To generate spatial correlation, we assume that $\epsilon_{nt} \equiv (\epsilon_{1nt}, ..., \epsilon_{Int})$ is distributed as a generalized extreme value distribution, where the marginal distribution of ϵ_{int} is a Gumbel distribution, but they are not necessarily independent of one another:[17]

$$\vec{\epsilon}_{nt} \sim F(\cdot)$$

where

$$F(\epsilon_{1nt}, ..., \epsilon_{Int}) = \exp(-G(e^{-\epsilon_{1nt}}, ..., e^{-\epsilon_{Int}})) \tag{4}$$

where G is a correlation function in the sense of McFadden (1978). To be specific, G is defined over the range of I non-negative real numbers, and it must satisfy the following properties: it is non-negative; it is homogenous of degree 1; the limit when any one of its arguments approaches infinity is infinity; and the cross-partial with respect to any k distinct arguments is nonnegative if k is odd and nonpositive if k is even.

Under these assumptions, the probability of an agent choosing location i is:

$$p_i = e^{u_i} \frac{G_i(e^{u_1}, ..., e^{u_I})}{G(e^{u_1}, ..., e^{u_I})}$$

where G_i is the partial derivative of G with respect to its ith argument. See McFadden (1978) for the derivation.

We also wish to make the ϵ correlated not just over space but also over time. To do that, we assume that the joint distribution of $\vec{\epsilon}_{nt}$ and $\vec{\epsilon}_{nt+1}$ is given by:

$$(\vec{\epsilon}_{nt}, \vec{\epsilon}_{nt+1}) \sim F_2(\cdot, \cdot)$$

[16]We do not specify the source of the u_is, so the reader can think of the SPACE model as being the migration block of a spatial model, and that the u_is would originate in the housing, production, and amenities blocks.

[17]See Lind and Ramondo (2023) as an example of a generalized extreme value distribution in trade.

where

$$F_2(\epsilon_{1nt}, ..., \epsilon_{Int}, \epsilon_{1nt+1}, ..., \epsilon_{Int+1}) = \exp\Big(- G\big(H(e^{-\epsilon_{1nt}}, e^{-\epsilon_{1nt+1}}), ..., H(e^{-\epsilon_{Int}}, e^{-\epsilon_{Int+1}})\big)\Big)$$

and

$$H(x_1, x_2) = \left(x_1^{\frac{1}{1-\rho}} + x_2^{\frac{1}{1-\rho}} \right)^{1-\rho}$$

where $\rho < 1$.[18] We will further assume that $G(H(\cdot,\cdot), ..., H(\cdot,\cdot))$ is also a correlation function under the criteria above.[19]

Note that the cumulative distribution function of $\vec{\epsilon}_{nt}$ can be calculated by taking the limit as each element of $\vec{\epsilon}_{nt+1}$ goes to infinity. In this case, $\lim_{\epsilon_{t+1}\to\infty} H(e^{-\epsilon_t}, e^{-\epsilon_{t+1}}) = e^{-\epsilon_t}$. So the marginal distribution of $\vec{\epsilon}_{nt}$ is given by F from equation (4). A similar argument applies so that the marginal distribution of $\vec{\epsilon}_{nt+1}$ is also given by F. Hence, the distribution over n of $\vec{\epsilon}_{nt}$ is time-invariant, even though an individual n's realization of $\vec{\epsilon}_{nt}$ will not be the same as $\vec{\epsilon}_{nt+1}$.

The joint distribution F_2 implies a conditional distribution $\tilde{F}(\vec{\epsilon}_{nt+1}|\vec{\epsilon}_{nt})$. This can be iterated as a Markov chain to calculate distributions of the future ϵ.

Migration occurs when the locations i that maximize $u_{it} + \epsilon_{nit}$ and $u_{i,t+1} + \epsilon_{ni,t+1}$ are different.

Proposition 2. *In the SPACE model when the u_is are fixed over time, migration from i to j is given by:*

$$m_{i\to j} = (1-\rho)e^{u_i+u_j}\left(\frac{G_i(e^{u_1}, ..., e^{u_I})G_j(e^{u_1}, ..., e^{u_I})}{G(e^{u_1}, ..., e^{u_I})^2} - \frac{G_{ij}(e^{u_1}, ..., e^{u_I})}{G(e^{u_1}, ..., e^{u_I})}\right) \tag{5}$$

where G_{ij} is the second derivative of G with respect to its ith and jth elements.

[18] ρ is not the correlation between ϵ_{int} and ϵ_{int+1}. The correlation is $1-(1-\rho)^2$ (Cascetta, 2001).

[19] Suppose G has a cross-nested structure as follows:

$$G(x_1, ..., x_I) = \sum_k \left(\sum_{i\in I_k} \alpha_{ik} x_i^{\frac{1}{1-\gamma_k}} \right)^{1-\gamma_k}$$

for nests k and weights $\alpha_{ik} > 0$. Each I_k is a subset of the set of locations $\{1, ..., I\}$. Then, as long as $\rho > \gamma_k$ for all k, G will be a correlation function. This can be shown because G is the sum of several nested logit correlation functions, and the sum of correlation functions is a correlation function. This claim is important because then the function F_2 is indeed a proper cumulative distribution function. Lind and Ramondo (2023) prove that cross-nested formulations of G can approximate any correlation function, so combined with this proposition, then any correlation function can be approximated by a correlation function that permits some ρ that induces the persistence we desire to add to the model.

Alternatively,

$$m_{i\to j} = (1-\rho)p_i p_j \left(1+\tau_{ij}(e^{u_1},...,e^{u_I})\right) \tag{6}$$

where $\tau_{ij} = -\frac{G_{ij}(e^{u_1},...,e^{u_I})G(e^{u_1},...,e^{u_I})}{G_i(e^{u_1},...,e^{u_I})G_j(e^{u_1},...,e^{u_I})}$.

Please refer to Appendix A.2 for the proof. Equation (6) resembles a gravity equation. τ_{ij}, which is non-negative, is related to the correlation between two alternatives. More correlated locations—in the sense that they have higher τ_{ij}—will have more migration between them. Conversely, if migration is more correlated over time (higher ρ), there will be less migration.

Corollary 1. *The cross-elasticity of population to utility is given by the migration rate times a constant.*

$$\frac{\partial p_i}{\partial u_j} = -\frac{1}{1-\rho} m_{i\to j}$$

when $i \neq j$.

See Appendix A.3 for the proof. This corollary is important because in many applications, we are interested in the population elasticity to local shocks. This corollary tells us that migration is a sufficient statistic to know these elasticities up to a constant.

The migration in Corollary 1 is a steady-state migration between i and j, which in the steady-state of the model is equal to the migration from j to i. In the data, these two migration rates are highly related, even conditional on origin, destination, and distance (see Appendix C.6). However, migration is rarely exactly balanced in the data. In subsequent parts of the paper, we will adjust the data to be balanced by using the logarithmic average i.e. $(m_{i\to j,t} - m_{j\to i,t})/\log(m_{i\to j,t}/m_{j\to i,t})$, because in a calibrated parametric version of the model that we introduce in the next section, this is a precise approximation of the steady-state migration we need for Corollary 1 (see Appendix B.4 for details).

Having set up the SPACE model, we now present a proposition analogous to Proposition 1 from the moving cost model, to see whether the SPACE model can match the square root fact. In this case, the limit we consider is for the persistence parameter ρ to approach 1.

When ρ is close to 1, ϵ_{int} will resemble a random-walk with logistic innovations. It is well known that the standard deviation of a random walk grows with the square root of time. As the standard deviation grows, the odds of crossing a threshold—which corresponds to moving in this model—grow roughly proportionally to the standard deviation, hence generating the square root fact. We formalize this intuition in the following lemma:

Lemma 1. *Define Λ^x to be the sum of x i.i.d. mean-zero logistic random variables. Then in steady-state of the SPACE model,*

$$\lim_{\rho \to 1} \frac{m^t_{i \to j}}{m^1_{i \to j}} = \frac{\mathbb{E}[|\Lambda^{2t}|]}{\mathbb{E}[|\Lambda^2|]}$$

where $\mathbb{E}[|\cdot|]$ denotes the mean absolute deviation (MAD).

Please refer to Appendix A.4 for the proof. This proposition establishes that for ρ close to 1, relative migration over t years will be proportional to the mean absolute deviation of the sum of $2t$ independent logistic random variables, following the intuition from above. Importantly, we can calculate bounds on this ratio:

Proposition 3. *In the steady-state of the SPACE model, as $\rho \to 1$, the ratio of t-year migration to 1-year migration for any state-pair is bounded below by $\sqrt{t}$ and bounded above by $\sqrt{\pi/3}\sqrt{t}$:*

$$\sqrt{t} \leq \lim_{\rho \to 1} \frac{m^t_{i \to j}}{m^1_{i \to j}} \leq \sqrt{\frac{\pi}{3}}\sqrt{t}$$

Please refer to Appendix A.5 for the proof. This proposition shows the SPACE model is able to naturally match the square root fact—albeit approximately. There are several senses in which the fit is only approximate. First, the proposition only establishes bounds. Fortunately, $\sqrt{\pi/3}$ is approximately 1.023, so the bound is tight.[20] Second, the proposition relies on $\rho \to 1$, while in practice we will calibrate a $\rho < 1$. Our calibration in the next section calibrates a high persistence, and we check the square root fact quantitatively under that calibration. Finally, the proposition assumes the economy is in steady-state, whereas in the data, the utility of each state (u_i) is changing. We check that the changes are not large enough to change the square root fact in simulations in Appendix B.4.

The SPACE model has a closed-form solution for aggregate welfare. Aggregating over the ϵ's, welfare is given by:

$$W_t = \mathbb{E}_{\vec{\epsilon}_{nt}}[V_t(\vec{\epsilon}_{nt})] = \sum_{s=0}^{\infty} \beta^s \left(\bar{\gamma} + \log G(\exp(u_{1t+s}), ..., \exp(u_{It+s}))\right)$$

[20] A tight band in Proposition 3 is intuitive because, if the innovations were normally distributed and Lemma 1 referenced normal distributions instead of logistic ones, the ratio would be exactly the square root, without the 2.3 percent deviation. Similarly, if the lemma were based on the standard deviation rather than the mean absolute deviation, the ratio would also be precisely the square root. Although this exact relationship does not hold for the MAD of logistic distributions, logistics are approximately normal, and the MAD is approximately proportional to the standard deviation for many distribution families. Consequently, it is not surprising that the relationship holds approximately.

where $\bar{\gamma}$ is the Euler-Mascheroni constant. This follows directly from McFadden (1978). It follows that

$$\frac{\partial W_t}{\partial u_{it+s}} = \beta^s p_{is} \qquad \text{and} \qquad \frac{\partial^2 W_t}{\partial u_{it+s} \partial u_{jt+s}} = -\beta^s \frac{1}{1-\rho} m_{i \to j,s}$$

where $m_{i \to j,s}$ represents the amount of steady-state migration from i to j that would occur when the location utilities are given by u_{is}.[21] So to second-order, the welfare effects of changes in local utility are given by population and migration, adjusted for parameters β and ρ.

3.2 Calibration

Many functions G will exactly match the populations and migration in the data. And in many settings where the outcome of interest is the change in populations, the choice of a specific G makes no first-order difference, due to the previous result that $\frac{\partial p_i}{\partial u_j} = \frac{1}{1-\rho} m_{i \to j}$.

Nonetheless, in some applications, it will be helpful to take a stand on a specific functional form for G. It is sometimes useful to have a functional form that allows for analytical formulas for population changes and the ability to simulate an individual's migration path.[22] We will also be interested in finding a G that is consistent with a value of ρ near 1, because that is the parameterization of ρ that leads to the square root fact.

Consider the following G functional form, which generates a cross-nested logit with nests that can contain any number of locations:

$$G(x_1, ..., x_I) = \sum_{q=1}^{Q} \left(\sum_{i \in \mathcal{N}_q} (w_{iq} x_i)^{\frac{1}{1-\gamma}} \right)^{1-\gamma}, \tag{7}$$

where nests are indexed by $q = 1, \ldots, Q$, each nest q corresponds to a subset $\mathcal{N}_q \subseteq \{1, \ldots, I\}$, and w_{iq} are nonnegative weights that we will calibrate to match populations and migration (with $w_{iq} = 0$ if $i \notin \mathcal{N}_q$). This choice of functional form leads to a cross-nested logit, in which nests can contain an arbitrary number of locations. The within-nest elasticity of substitution is $\frac{1}{1-\gamma}$ and the across-nest elasticity of substitution is 1.

In order to do individual-level simulations, it will be helpful to parameterize γ and the nest weights w_{iq} as a function of ρ, and consider the limit as $\rho \to 1$. In particular, we

[21] A good approximation of the this migration would be the logarithmic mean of $m_{i \to j}$ and $m_{j \to i}$ in the data (see Appendix B.4).

[22] Another minor benefit to a calibration is a closed-form solution for migration outside of steady-state (see Appendix B.4).

will parameterize

$$1 - \gamma = \frac{1-\rho}{1-\tilde{\rho}}$$

where $\tilde{\rho}$ is fixed, so as $\rho \to 1$, $\gamma \to 1$ as well. This will allow us to consider the nest for each person as fixed over time, but allow for agents to move within nests. In order to approximately match the square root fact, we will aim to choose a $\tilde{\rho}$ as close to 1 as possible while also matching migration rates in the data.

Rather than calibrating the w_{iq}, it is more intuitive to calibrate

$$\tilde{w}_{iq} \equiv (w_{iq} e^{u_i})^{\frac{1}{1-\gamma}} \left(\sum_{k \in \mathcal{N}_q} (w_{kq} e^{u_k})^{\frac{1}{1-\gamma}} \right)^{-\gamma}$$

Note that these $\tilde{w}$ are also a function of ρ since they depend on γ. From there, the w's can be backed out as $w_{iq} = e^{-u_i}\, \tilde{w}_{iq}^{1-\gamma} \left(\sum_{k \in \mathcal{N}_q} \tilde{w}_{kq} \right)^{\gamma}$. One intuition is that these adjusted weights represent the mass of people that choose a location–nest pair: $\tilde{w}_{iq}$ choose i via nest q. Under the above definitions,

$$p_i = \sum_{q\colon i \in \mathcal{N}_q} \tilde{w}_{iq} \tag{8}$$

and

$$m_{i \to j} = (1-\tilde{\rho}) \sum_{q\colon \{i,j\} \subseteq \mathcal{N}_q} \frac{\tilde{w}_{iq} \tilde{w}_{jq}}{\sum_{k \in \mathcal{N}_q} \tilde{w}_{kq}}. \tag{9}$$

See Appendix B for the derivation.

This calibration is useful for two reasons. First, it leads to a simple formula for populations and straightforward exact-hat algebra. Define $\hat{p}_i = p_i'/p_i$, the growth in population. Similarly, define $\hat{u}_i = \exp(\frac{u_i' - u_i}{1-\rho})$, the change in local utility, normalized such that the elasticity of location choice with respect to utility does not diverge as $\rho \to 1$. Then, the change in populations as a function of the change in utility can be expressed as

$$\hat{p}_i = \hat{u}_i^{1-\tilde{\rho}} \sum_{q=1}^{Q} \frac{\tilde{w}_{iq}}{p_i} \frac{\sum_{k \in \mathcal{N}_q} \tilde{w}_{kq}}{\sum_{k \in \mathcal{N}_q} \tilde{w}_{kq} \hat{u}_k^{1-\tilde{\rho}}} \tag{10}$$

See Appendix B for the derivation.

Second, this calibration is useful because each person is drawn into choosing between $|\mathcal{N}_q|$ locations and never considers anywhere else. Furthermore, within a nest, one can simulate the $|\mathcal{N}_q|$ locations as independent of each other, since the nesting structure accounts for all of the correlation. This is helpful because drawing from high-dimensional generalized extreme value distributions is computationally-intensive. By only having to draw extreme value distributions that are correlated from one period to the next, the computational burden is greatly alleviated.

Of course, there are still many ways to choose the $\tilde{w}_{iq}$ and $\tilde{\rho}$. A natural way might be to maximize $\tilde{\rho}$ as to match the square root fact, while treating equations (8) and (9) as constraints. Unfortunately, this is not a problem on which we could make significant progress. However, if we restrict $\mathcal{N}_q$ to be sets of locations with two elements, the problem becomes tractable. This restriction has the added advantage of only having to simulate a correlated GEV for two locations at a time, making the computational burden of doing simulations lighter. In particular, the solution to this constrained optimization problem is that the $1-\tilde{\rho}$ is the largest eigenvalue of an $I \times I$ matrix M, defined by

$$M_{ij} = \begin{cases} \frac{m_{i\to j}}{p_i} & \text{if } i \neq j \\ \frac{m_i}{p_i} & \text{if } i = j \end{cases}$$

where $m_i = \sum_{k\neq i} m_{i\to k}$.[23] The calibrated $\tilde{w}_{ij} = \frac{1}{1-\tilde{\rho}} m_{i\to j} \left(1 + \ell_j/\ell_i\right)$, where ℓ_i is the ith element of the corresponding eigenvector, where we shorten the notation so that $\tilde{w}_{ij}$ is the share of people that choose i from the nest that includes i and j, i.e. $\tilde{w}_{ij} \equiv \tilde{w}_{iq}$ where $q = \{i, j\}$. The calibrated value of $\tilde{\rho}$ is 0.892, which, as we will show in the next section, is close enough to 1 to generate the square root fact in simulations.[24] Details on this calibration can be found in Appendix B.[25]

3.3 Dynamics of migration in the calibrated SPACE model

Under the calibration from the previous section, each agents' actions are as if they were drawing two Gumbel-distributed ϵ's that are independent between the two locations, but correlated over time with correlation parameter $\tilde{\rho}$. Compared to a generalized extreme value distribution with many dimensions, it is computationally efficient to simulate se-

[23]For the calibration, we use the migration matrix summing all one-year migration from 2004-2018 in the GCCP, and using the logarithmic average to symmetrize it. (See Appendix B.4 for the rationale of using the logarithmic average.)

[24]The autocorrelation implied by this parameter is $1 - (1-\tilde{\rho})^2 = 0.988$.

[25]The downside of assuming two elements per nest is that it means that simulations will never have anyone that chooses to live in more than two locations over time, whereas in the data people move to their third or fourth state somewhat commonly. In Appendix B.2, we propose a calibration that has a reasonably high $\tilde{\rho}$ and which would generate people moving to many potential states.

quences of conditional Gumbels where ϵ_{it} and ϵ_{it+1} have joint distribution

$$\exp(-(e^{-\epsilon_{it}/(1-\tilde{\rho})} + e^{-\epsilon_{it+1}/(1-\tilde{\rho})})^{1-\tilde{\rho}})$$

because there is a closed-form formulation of the CDF of ϵ_{it+1} given ϵ_{it}. We can easily draw a sequence for location i and another sequence for location j, and then that person will choose i when $\epsilon_{it} > \epsilon_{jt} + \log \frac{\tilde{w}_{ji}}{\tilde{w}_{ij}}$.

In this section, we simulate 100,000 agents using the above algorithm for each state-pair, and record their location choice for 15 years.[26] Each observation is weighted by $\tilde{w}_{ij} + \tilde{w}_{ji}$ to make the simulation representative, and we can use the simulated panel to measure dynamic moments of migration.

In Figure 2a, we show that the t-year migration rate does follow a square root pattern in both the data and in a simulation of the SPACE model. The model does not match the data perfectly, with the lines diverging over time. In part, this is because the model is calibrated to match the one-year migration rate, and the 14-year migration rate, which in the simulation is about $\sqrt{14}$ times the one-year migration rate, is sensitive to that choice. This figure also presents the same exercise for a simulation of the moving cost model, which is much more linear and diverges much more from the data.[27,28]

Of course, the t-year migration rate is not the typical way the dynamic moments of migration are presented in the data, so it is worth examining whether the model captures more commonly examined moments. A natural moment is the conditional probability of moving given a previous move, i.e. the hazard rate of migration.

Figure 2b shows the probability of another migration at different time horizons after an interstate move.[29] Since we do not target these statistics in the calibration, the simulated statistics do not match the data perfectly. But the general pattern is similar, especially its decay as the person has lived in the state for longer.

[26]The ϵ_{it} draws are reused for each state-pair, but the threshold for which state to live in changes by state-pair.

[27]For the moving cost model simulation, we use the observed 1-year migration probabilities for the entire GCCP panel, 2004-2018, and simulate 15 years of location choices for 300 million people.

[28]A model that features both persistent unobserved heterogeneity and moving costs can improve upon the fit of all three dynamics moments mentioned in this section. Intuitively, because moving costs and persistence generate different patterns for the t-year migration rate, they can be separately identified off of short-run and long-run migration rates (e.g. 1-year and 10-year migration). However, such a model is much less tractable than either the SPACE model or the moving cost model. In practice, a model that combines both features calibrates significant persistence and small moving costs, and leaving the moving costs out of the model completely makes little quantitative difference. See Appendix D for details.

[29]To be included in this analysis, a person must show up for the number of years that would be necessary to calculate the statistic, but we do not use a balanced panel.

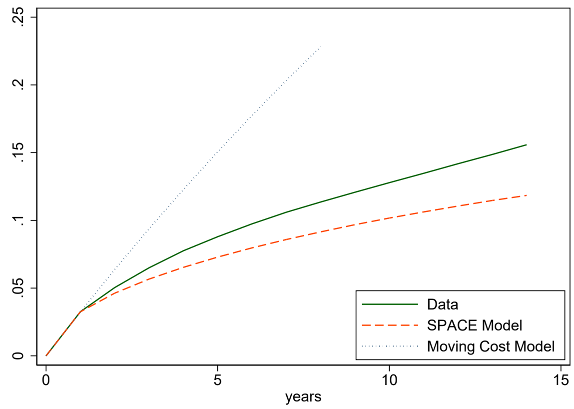

(a) Migration rate in the data and models

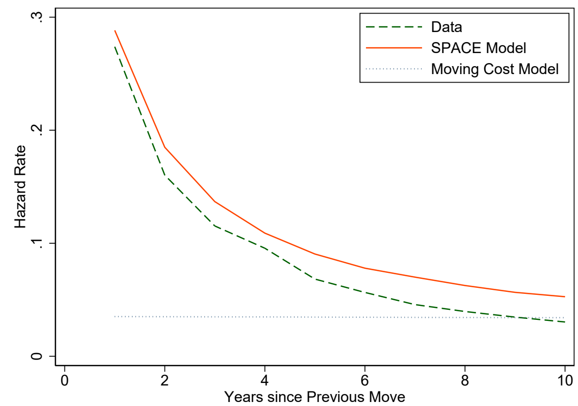

(b) Hazard rate of migration

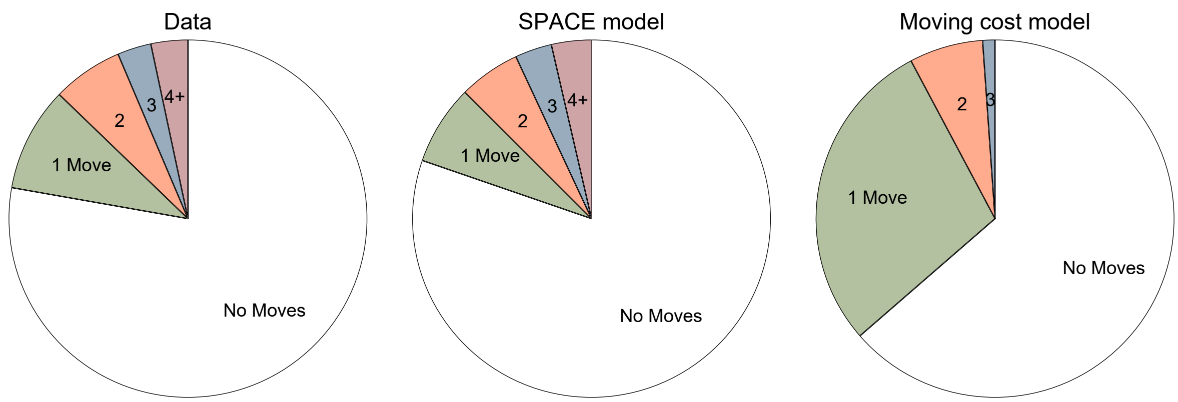

(c) Number of moves in 14 years

Figure 2: Dynamic moments. In panel (a), the t-year migration rate is calculated as the percent of people living in a different state than they were t years ago. Data are from an unbalanced panel, and included any observations from 2004-2018 for which the state of residence is observed t years apart. In panel (b), the conditional probability of migration is plotted. For the value at x years, the probability of migration is conditional on the person having migrated x years previously and remained in the same state ever since. In panel (c), the number of moves in 14 years is calculated for people whose state is observed in every year from 2004-2018.

The intuition for the decreasing hazard rate over time is simple. Conditional on having moved recently, the agents are likely relatively indifferent between the two regions and are likely to move back. The longer they have stayed in one region, the more likely that their accumulated utility shocks have drawn them further away from being indifferent, so the probability of migration decreases over time. The literature has typically focused on the concept of "attachment" to explain this phenomenon (Mangum and Coate, 2019; Farrokhi and Jinkins, 2024). In the SPACE model, people who have lived in a location for longer are more attached, but it is because their repeated decision not to move has revealed that they like the location, not due to an economic force that increases their utility by staying there longer.

In addition to hazard rates, another readily observable statistic is the distribution of the number of interstate moves over time. Figure 2c looks at how many moves are made over a 14 year period. In the data, a large majority of people make zero moves, but some people make many moves. Here, we include in this chart only people for whom we have data in all 15 years (for up to 14 possible moves). The model is able to capture the large fraction of people that never move. It also comes close to the data on the number of people that move once or twice. Importantly, it captures the fact that a few percent of people move four or more times over the 14 years. The figure also includes similar statistics for a moving cost model, which fits the data substantially less well.

Together, these moments reinforce that the SPACE model captures essential features of migration dynamics in the data, especially compared to the standard moving cost model.

4 Does the new model matter?

The previous section introduced a new model and showed that it did a better job at matching dynamic moments in the micro data. In this section, we explore the implications of the new model by comparing it to the workhorse model. For some questions, we find the differences are minimal, while for others we find substantial differences.

4.1 Micro forecasting accuracy

An important distinction between the two models is their implications for forecasting individuals' locations. Suppose we observe the agent's location in 2004 and wish to forecast where they will live in every subsequent year until 2018. We use the calibrated versions of each model to perform the forecasting. We judge the performance of the models using the mean Kullback-Leibler divergence. Specifically, we use the same simulations as before, and then for each initial location, we calculate the simulated probability that a person who was in state i in 2004 ends up in state j in year t.[30] Then using people's true locations in the data, we calculate the Kullback-Liebler divergence:

$$KL_t = \frac{1}{N}\sum_n \log\left(\frac{\mathbb{P}_{\text{data}}(\text{lives in } j \text{ in } t|\text{lived in } i \text{ in } 2004)}{\mathbb{P}_{\text{model}}(\text{lives in } j \text{ in } t|\text{lived in } i \text{ in } 2004)}\right)$$

for each year. We plot KL_t in Figure 3.

Both models initially have a low Kullback-Leibler divergence, as they are about equally good at predicting locations in 2005 since they were both parameterized to match the migration data in that year. But over time, the moving cost model's Kullback-Leibler divergence grows sharply, suggesting that the log likelihood is on average about 0.12 log-points worse per observation than the maximum possible performance of any model by 2018. In contrast, the SPACE model has a Kullback-Leibler divergence of 0.014 log-points in 2018, suggesting limited room for further improvement.

The SPACE model provides more accurate forecasts over longer horizons because it matches the dynamic moments. In particular, many more people end up moving away from their initial state in the moving cost model because moving probabilities are independent over time, whereas the SPACE model is better able to match the total number of people who leave.

4.2 Moving costs need not be large

Kennan and Walker (2011) estimate an average moving cost of \$312,146 (in 2010 dollars).[31] This estimate exceeds six times the median household income in that year, which was \$49,445 (Census, 2011).[32] Such a large cost is one of the main reasons that most people do not move, in their model. Economists can argue about whether that number is reasonable, and even within Kennan and Walker (2011), there is substantial heterogeneity in moving costs. In contrast, the SPACE model can match the main facts about internal migration without any moving costs. In other words, the fact that most people do not move is not sufficient evidence to conclude that moving costs are large.

[30] This exercise is why we chose to simulate hundreds of millions of agents, as we had to ensure that every origin-destination-timeframe that shows up in the data is also in the simulations.

[31] They also include an analysis of moving costs conditional on moving, but they include the payoff shocks in the moving costs, and consequently find that the average moving cost is actually very negative.

[32] Many papers estimate moving costs of a similar magnitude. See Table 2 of Howard (2026a).

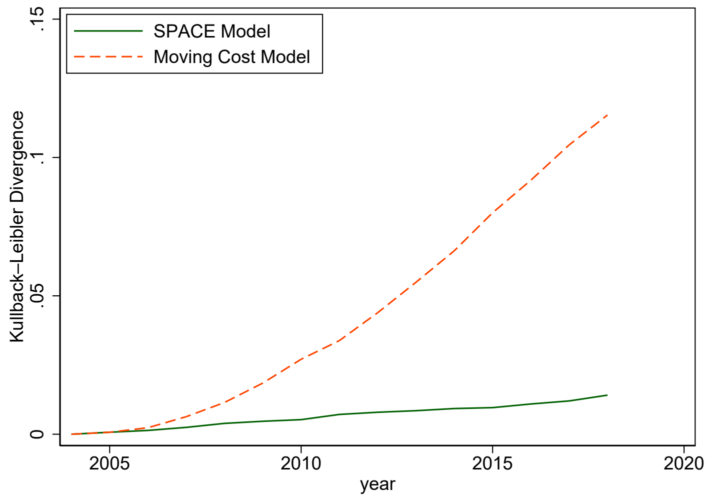

Figure 3: Kullback-Leibler divergence of the SPACE and Moving Cost models. For the SPACE model, 100,000 agents per state-pair are simulated over 15 years, and KL_t is calculated based on the simulated data using the formula in Section 4.1, weighted to be nationally representative. For the moving cost model, 300,000,000 observations are simulated over 15 years, and KL_t is calculated based on the simulation. The data is from the GCCP.

In Appendix D, we calibrate an augmented version of the SPACE model that also features moving costs to match the one-year and 10-year rates of migration. The moving costs are about two orders of magnitude smaller than had we calibrated a model that had only moving costs.

A common counterfactual in the literature is to consider changes in moving costs (Kennan and Walker, 2011; Schubert, 2021; Zerecero, 2021), which is also an actual policy used by some localities.[33] For example, Kennan and Walker (2011) finds that a moving subsidy could substantially increase the gross migration rate. In a moving cost model, a temporary incentive to move to location i has a very persistent impact on the population of i. In contrast, in our model, a moving subsidy would encourage people to relocate, but only for as long as the subsidy lasts. After the subsidy expires, they are no more likely to remain in the place they moved to than they would be to live there had the subsidy never occurred.[34]

One particular way of lowering "moving costs" may be improving infrastructure such as

[33] A handful of cities around the United States offer monetary incentives to relocate (Cornerstone Home Lending, 2021).

[34] One could imagine other economic reasons that populations remain higher after a population expansion, such as the accumulation of housing capital or agglomeration in that place.

roads, which increases migration (Morten and Oliveira, 2024). The SPACE model would interpret the increased migration as an increase in the correlation of idiosyncratic utility across the locations that the infrastructure connected.[35] In the moving cost model, the increased migration must reflect lower moving costs, and so welfare would have increased, net the cost of the roads. In the modified SPACE model, though, the welfare effects are much more muted.[36]

4.3 Macro population elasticities

Another common use for migration models is to calculate population elasticities to changes in a location's utility, v_i, both in the short run and the long run. In the standard moving cost model, the elasticity of population with respect to v_j is given by:

$$\frac{\partial \log p_i}{\partial v_j} = -\sum_k \frac{m_{k\to i}}{p_i}\frac{m_{k\to j}}{p_k}$$

This is approximately proportional to the migration rate between i and j, since the two terms where $k = i$ or $k = j$ are much larger than the remaining terms.

The elasticities in the two models are similar in that they are approximately proportional to the migration rate when migration rates are low. Given that the scale of v_j is not specified in either model, the constant terms are ignorable without loss of generality. However, in the long-run, the similarities of population elasticities between the SPACE model and the moving cost model break down. Consider a one-time permanent change in v_{it} for the SPACE model or the moving cost model. In the SPACE model, the population elasticity is still exactly the same, since it is given by corollary 1. This is not the case with the moving cost model. In fact, when migration is symmetric and moving costs are high, we show that the long-run elasticities converge to a static logit.

Proposition 4. *In the long-run steady state of a moving cost model, if migration is symmetric, in the sense that* $\lim_{\Delta\to\infty} m_{i\to j}/m_{j\to i} = 1$*, then*

$$\lim_{\Delta\to\infty} \frac{\partial \log p_i}{\partial v_j} = \begin{cases} -2p_j & \text{if } i \neq j \\ 2 - 2p_j & \text{if } i = j \end{cases} \tag{11}$$

[35] This increased correlation is a natural consequence of infrastructure that reduces travel time, as it makes social networks and amenities in the other city more accessible, thereby increasing the correlation of the idiosyncratic utility experienced in the newly connected places. Hence, new infrastructure can increase migration even if the infrastructure itself is not the primary means of migration.

[36] Of course, other welfare benefits, such as those that come through increased trade, are not contingent on the migration model.

See Appendix A.6 for the proof. Recall that the total population is mass 1, so p_i is both the population of i and its population share. These steady-state elasticities are the same as a static logit, and a key difference from the SPACE model is that they have no relationship to migration data. In Appendix C.8, we show equation (11) is a good numerical approximation to the long-run of a calibrated moving cost model, when moving costs lead to realistic migration rates, rather than being infinite.

In the long-run, the moving cost model has no notion that closer states are better substitutes or that states with higher migration are likely to be more impacted by a change in the other state. The moving cost model does not imply that a state with a high migration rate will exhibit a greater long-run elasticity than a state with a low migration rate. Rather, the only pertinent factor is the population share of the state receiving the shock to calculate all the relevant elasticities (approximately).[37]

Figure 4 summarizes the conclusions of this section. In the short-run (Panel a), the SPACE model and the moving cost model predict practically the same cross-elasticities of population. But the in the long-run (Panel b), there is almost no relationship between the two.

4.4 Macro adjustment dynamics

Given the differences in long-run population elasticities, it follows that the intermediate dynamics must also be different across the two models. We illustrate by considering a one-time permanent shock to Louisiana utility, $v_{\text{Louisiana}}$, to see how populations respond in each model. In Figure 5 panels (a) and (b), we show that the population dynamics after a one-time permanent shock are starkly different.[38]

In the SPACE model (Panel a), the population adjustment in Louisiana is immediate, and the population stops adjusting after the first period. In contrast, in the moving cost model (Panel b), the population adjustment takes many years, with the model finally

[37] In the moving cost model, the migration rates govern the speed of adjustment (Kleinman et al., 2023), but not the long-run effects.

Including birthplace as a state variable in the moving cost model would mean that adjustments would depend on the population shares of individuals by birthplace. To this extent, migration between states would be correlated to the population cross-elasticities (Zabek, 2024).

Other features of the model can also change the population elasticities. For example, Monte, Redding and Rossi-Hansberg (2018) adds commuting to a static model of location choice to generate variation in population elasticities.

[38] We chose Louisiana because in Appendix C.9, we do further analysis on the Hurricane Katrina episode. The size of the shocks in each model is normalized to have a long-term effect of about 5 percent of the population for Louisiana. However, the scale is irrelevant to the focus of this exercise, which is the dynamics.

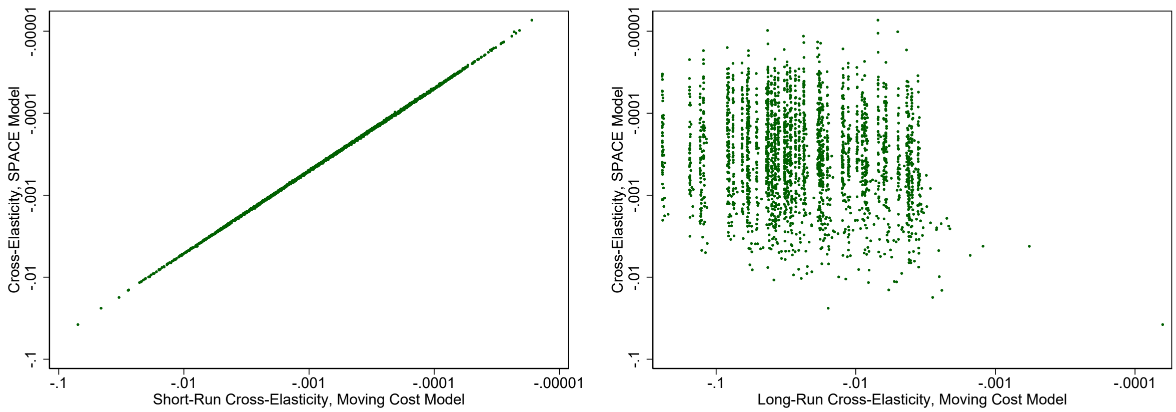

(a) Short-run

(b) Long-run

Figure 4: A comparison of the population cross-elasticities between the SPACE and moving cost models. For both figures, each dot represents a pair of states. The point is located at the population cross-elasticity between the two states in each of the two models. The constant multiplicative terms are ignored, since each model is subject to a normalization of utility. Data for migration and population shares comes from the GCCP. Both axes in both graphs have log scales.

approaching a steady-state after almost 200 years.[39]

In Panels (c) and (d) we illustrate the dynamics for other states in response to the same shock to Louisiana's utility. In the SPACE model (Panel c), there is a bigger population effect on Mississippi than there is on New York, as one would expect due to the geography. Again, the adjustment occurs immediately. But in Panel (d), the dynamics follow interesting and perhaps counterintuitive patterns. In New York, the population adjustment is particularly slow because of low migration between Louisiana and New York. In contrast, for Mississippi, the population dramatically overshoots its long-run steady-state because there is so much migration between Louisiana and Mississippi.

This exercise is intended to highlight difference between the models, not argue in favor of one over the other. While the time series of Louisiana's population share after Hurricane Katrina does resemble that of the SPACE model (see Figure A10a), we do not wish to judge whether Katrina represented a one-time permanent disutility shock. Nonetheless, the differences in the transition dynamics imply that the models will interpret empirical facts differently. We will elaborate on this point in Section 5.1.

[39]In an estimated model with both persistent unobserved heterogeneity and moving costs, we still see a very fast population adjustment (see Appendix D). It is much closer to the macro dynamics of the SPACE model than the moving cost model.

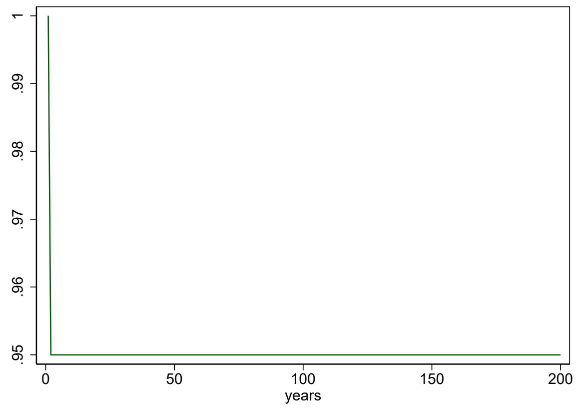

(a) Louisiana, SPACE model

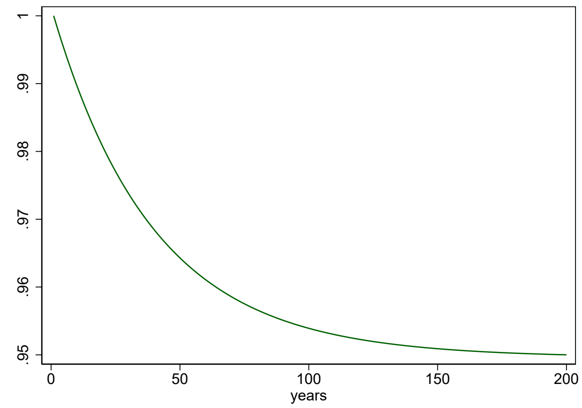

(b) Louisiana, moving cost model

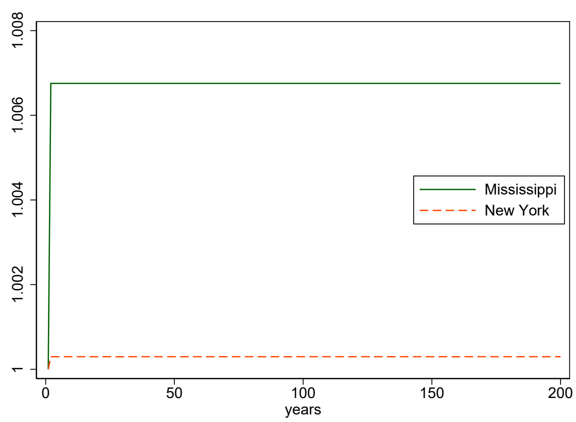

(c) Mississippi and New York, SPACE model

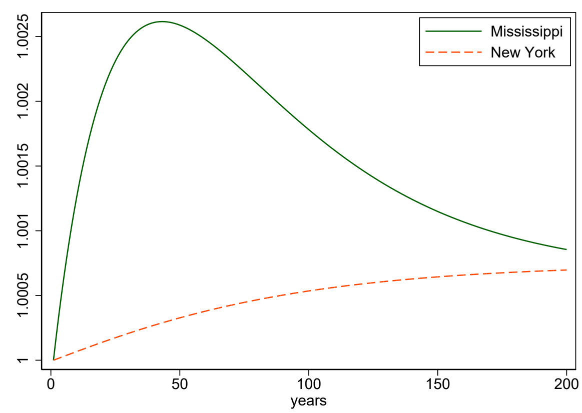

(d) Mississippi and New York, moving cost model

Figure 5: Population Dynamics after a one-time permanent change in $v_{\text{Louisiana}}$, in the SPACE model and the moving cost model, for Louisiana, Mississippi, and New York. The shock size is chosen to cause a long-run five percent decline in Louisiana's population. The y-axis in each figure measures the population in each state, relative to the population in that state at time $t = 0$, before the shock. Mississippi and New York were chosen to represent two states for which there is high gross migration with Louisiana, and low gross migration with Louisiana, respectively.

4.5 Implied utility changes

Given the differences in population elasticities, it must be the case that the models will imply different things about changes in utility over time. This is important for papers that wish to estimate the welfare effect of some policy or event that varies across space. The population changes in the SPACE model can be expressed with exact hat algebra. Under our previous calibration, the population changes are given by:

$$\hat{p}_i = \hat{v}_i^{1-\tilde{\rho}} \sum_{j \neq i} \frac{\tilde{w}_{ij}}{p_i} \frac{\tilde{w}_{ij} + \tilde{w}_{ji}}{\tilde{w}_{ij}\hat{v}_i^{1-\tilde{\rho}} + \tilde{w}_{ji}\hat{v}_j^{1-\tilde{\rho}}}$$

Recall that $\hat{v}_i = \hat{u}_i$ due the continuation value differing from the current value by only a constant. This equation can be numerically inverted to calculate the $\hat{v}_i$ based on the $\hat{p}_i$. Note that while this exact formula depends on our calibration, it will give the same answer to first-order of any choice of G, since the population elasticities to v are pinned down by migration rates.

For the moving cost model, the exact hat algebra is given by:

$$\hat{p}_{it} = \hat{v}_i \sum_j \frac{\frac{m_{j\to i}}{p_i} \hat{p}_{jt-1}}{\sum_k \frac{m_{j\to k}}{p_j} \hat{v}_k}$$

This formula is similar to the dynamic hat algebra derived by Caliendo et al. (2019) (see Appendix B.3 for the derivation). Note that it depends on $\hat{p}_{it-1}$ as well as $\hat{p}_{it}$. For this exact hat algebra, we need to know about the change in how populations are changing, rather than simply the change in population. Again, this is straightforward to invert numerically, given data on $\hat{p}_{it}$ and $\hat{p}_{it-1}$.

To illustrate this, we consider the utility changes implied by the SPACE model and moving cost model from 2005-2018, the span of our data. We show the results in Figure 6. In the SPACE model, the places with the largest increases in relative utility are in the South and West, places that have seen large growth in population. New England and the Rust Belt have some of the largest relative decreases. In the moving cost model, the utility changes are almost the opposite. New York and New England have increased in relative utility, while the South and the West have mostly had relative declines.[40] Overall, the correlation between the log of utility changes implied by the two models is -0.497. This has important implications for estimating spatial models. For example, if one wanted

[40] The utility changes in the moving costs model are driven by changes in the net migration rate. While the South and West have had higher net migration rates over much of this period, the rate of that growth has slowed, leading us to infer declining utility.

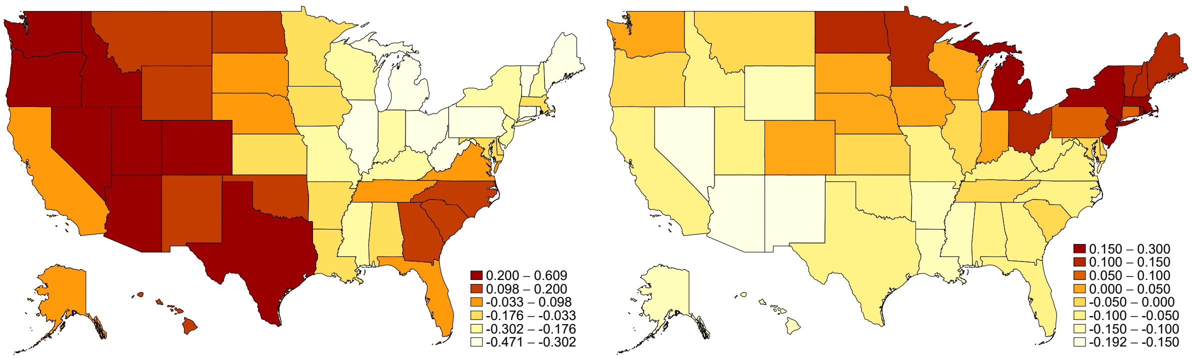

(a) SPACE model

(b) Moving cost model

Figure 6: Change in utilities v_j, 2005-2018, implied by the SPACE model and the moving cost model. Implied utility changes are calculated by inverting the exact hat algebra in Section 4.5. Population changes are from the Bureau of Economic Analysis (2004-2018).

to estimate the effects of a wage, rent, or amenity change on utility, you would get very different answers using the implied utilities from the SPACE model versus a moving cost model.

5 Discussion

5.1 Relation to literature

Given the major differences between the two models on several of these questions, it is worth emphasizing why these differences matter. For the micro forecasting and for the interpretation, the reasons to care about differences are readily apparent, but for the population elasticities and the dynamics, it is important to consider the context of the literature.

One big question in the literature is to what extent does population adjust to shocks? For example, if one particular location has a shock that permanently increases the utility of living there, how will that affect the distribution of the population around the country? This is a question that is asked by Caliendo et al. (2019) with respect to the China shock of Autor, Dorn and Hanson (2013), by Giannone (2017) with respect to skill-biased technological change, and by Cruz and Rossi-Hansberg (2024), Oliveira and Pereda (2020), and Rudik et al. (2021) with respect to climate change.[41] In the short run, both

[41] A reader may wonder why we do not replicate one of these papers to highlight the differences. However, doing such a replication would erroneously indicate that the SPACE model is less good at hitting the medium-run dynamics of population adjustment. The reason for this is that even if these

models agree that places with high gross migration, such as D.C., have highly elastic population responses to local shocks, compared to places with little migration. However, in the long-run, the SPACE model continues to make this prediction, while the moving cost model predicts similar elasticities for all locations. Similarly, in the short-run, both models agree that the population effects are felt in states that have lots of migration between them and the state with the shock. A shock to D.C. will affect Maryland and Virginia more in the short run than it will affect Arizona. This is consistent with the "donut" phenomenon during the recent COVID-19 crisis, as areas around major cities have experienced population and house price growth in recent years (Ramani and Bloom, 2021). Again, this holds in the long-run too for the SPACE model, but migration would not generate a long-run donut phenomenon in the moving cost model.

Another key question in this literature is how quickly the migration adjustment takes place (Kleinman et al., 2023; Amior and Manning, 2018; Caliendo et al., 2019). In the data, migration is sometimes quite persistent. For example, the Rust Belt has had low inmigration for decades, and the Sun Belt has had high inmigration for decades. In the moving cost model, much of this persistence is due to the fact that migration is inherently persistent (Kleinman et al., 2023). For example, the moving cost model might infer that the Rust Belt had a large negative utility shock a long time ago and the process of moving out has been very slow.

The SPACE model interprets this fact as being about the utility of a location adjusting slowly. It could be that the underlying shocks to utility are slow. Another possibility is that there was a large initial shock, but some equilibrium mechanism causes utility to decline slowly. For example, housing is durable, and so housing becomes cheap as people move out, keeping utility from falling too quickly (Glaeser and Gyourko, 2005). In Appendix E, we show the effects of a productivity shock to Louisiana in a model that has SPACE migration and durable housing. Indeed, the population adjustments are slow once housing is introduced. Similarly, there may be similar mechanisms through the labor market that make structural transformation slow.[42] The SPACE model would emphasize

models are not explicitly targeting the medium-run dynamics, they do get to choose what features of the world to add and can choose to include or not include features that will get the dynamics right. For example, Glaeser and Gyourko (2005) argues that the reason declining cities decline slowly is because the housing stock is slow to depreciate. Many of the quantitative papers do not have this feature. Therefore, replacing the migration block from those models would lead to unrealistic dynamics, if we did not also add in a feature like the one in Glaeser and Gyourko (2005). A reader may think of the moving costs as representing these other features, such as housing depreciation or labor market frictions. In that case, it would be preferable to explicitly model those features. For example, we add a model of housing depreciation in Appendix E.

[42]Kleinman et al. (2023) discuss how having location specific durable capital can keep wages high after

these forces as explanations for persistence in migration, whereas the moving cost model would attribute the persistence to an inherent property of migration itself, and find less explanatory power for these forces.

Conversely, some adjustments in the data are quite fast, such as Hurricane Katrina. The SPACE model can easily accommodate these fast adjustments by assuming a fast adjustment in utility. Even in a model with housing, utility can adjust quickly because a hurricane destroys much of the housing stock, forcing people out immediately. In contrast, the moving cost model requires implausible assumptions regarding utility to generate the fast decline in population followed by a small rebound the next year. The moving cost model implies that utility in Louisiana two years after Hurricane Katrina is higher than it was before the hurricane (see Appendix C.9).

Other papers are concerned with the effects of location-specific shocks on aggregate outcomes such as welfare or output (Tombe and Zhu, 2019; Eckert and Peters, 2022; Hsieh and Moretti, 2019). The degree to which people are able to move is an important factor for these outcomes. Because the second order effects are determined by these population elasticities, it shows that a shock that affects a higher-gross-migration place will have larger total effects on welfare if the shock is positive, and smaller effects if the shock is negative. Again, this holds in the short-run for both models, but in the long-run only for the SPACE model.

Finally, the spatial correlation of a shock is an important determinant of its welfare consequences. If a negative shock is extremely localized, it may be easy to move away from it, and there will be lots of insurance. If shocks are correlated across space, then the welfare effects may be much less insurable. Of course, in the long-run of a moving cost model, this effect will no longer hold.

5.2 Richer moving cost models

In Section 4, we compared the SPACE model to a simple version of a moving cost model. Yet as mentioned in the literature review, there is significant research that enriches the moving cost model to match a variety of facts (Kaplan and Schulhofer-Wohl, 2017; Giannone et al., 2020; Porcher, 2020; Mangum and Coate, 2019; Zerecero, 2021; Monras, 2020). Kennan and Walker (2011) includes features to increase home bias and return migration. As we demonstrated in Section 2.3, neither return migration, home bias, nor age is a sufficient feature to match the square root fact. So the puzzle that motivated

a negative productivity shock.

the model is not dependent on us having considered a simple version of the moving cost model. Nonetheless, in this section we ask whether the differences we highlighted between the moving cost model and the SPACE model are robust to considering richer versions of the moving cost model. In the rest of this section, we consider each of the differences we highlighted before.

For the prediction of individuals' locations, the reason that the SPACE model outperformed the moving cost model was because it could hit the square root fact. So if extensions of the moving cost model still do not hit the square root fact, they might be an improvement at predicting locations, but are not going to make the same predictions as the SPACE model.

For the interpretation of why people rarely move, some of the additional features imply lower estimated moving costs (Zerecero, 2021; Giannone et al., 2020), but never by orders of magnitude. So the difference between the SPACE model—which has no moving costs—and any moving cost model will remain substantial.

For population elasticities, more complex moving cost models feature long-run elasticities which may not necessarily be the same as a static logit model. For example, models with home bias have more similar elasticities to the SPACE model than the baseline moving cost model does.[43] However, except by coincidence, none of the additional features would generate the property that the long-run and short-run population elasticities are the same. So the marked difference between the SPACE model and the moving cost model will remain, both for long-run elasticities and for adjustment speeds.

For the implied utility changes, the exact implied utilities will certainly change with a richer model. However, it does not change the fundamental fact that the implied utility changes are primarily correlated to changes in migration in the moving cost model, and primarily correlated to changes in populations in the SPACE model.

6 Conclusion

This paper documents a new empirical regularity in internal migration: the t-year migration rate scales approximately with the square root of t. This pattern poses a challenge to the standard moving cost framework, which implies a linear relationship under com-

[43]With home bias, the elasticities are given by $\lim_{\Delta \to \infty} \partial \log p_j / \partial v_k = -\sum_i w_{ij} p_{ik}$, where $w_{ij} = \frac{p_{ij}}{p_j}$ is the share of people in j who are from i, and p_{ik} is the share of people from i living in k. These population shares are likely correlated to the amount of bilateral migration. However, it is a different formula, and the population share levels are likely different than the migration rates, so the dynamics in the two models will still be different.

monly used assumptions. To account for this fact, we propose a new model of location choice, the SPACE model, that introduces persistent and spatially correlated unobserved heterogeneity. The model is analytically tractable, matches key features of the data including the square root relationship and return migration rates, and provides a better out-of-sample forecast of location dynamics than the canonical moving cost alternative.

The SPACE model makes substantially different predictions from existing models. The SPACE model generates population elasticities that are proportional to bilateral migration flows, leading to markedly different long-run responses to local shocks than a standard moving cost model, as well as different dynamics. It also changes conclusions about the size of moving costs and which parts of the country have become better places to live.

The tractability of the SPACE model, especially its closed-form expressions for population and compatibility with exact hat algebra, makes it a practical foundation for embedding into broader spatial equilibrium frameworks. It is our hope that that future researchers will incorporate the SPACE framework into more complex models, much as researchers today use the moving cost framework in models that also include trade frictions and capital accumulation. We also hope that the SPACE model will be used in static models of location choice to generate realistic cross-elasticities of populations that are disciplined by migration data.

Data availability statement

This paper uses a combination of publicly available data and confidential data that cannot be publicly shared. Some of the data and all of the code to replicate this paper can be found on Zenodo at `https://doi.org/10.5281/zenodo.19240923` (Howard and Shao, 2026). The PSID data must be separately downloaded from `https://doi.org/10.3886/ICPSR304488.V1` (Howard, 2026b). Data from the GCCP must be applied for and typically requires a University of Illinois affiliation. Please see replication files for details.

References

Allen, Treb and Dave Donaldson, "Persistence and path dependence in the spatial economy," 2020. National Bureau of Economic Research Working Paper.

Amior, Michael and Alan Manning, "The persistence of local joblessness," *American Economic Review*, 2018, *108* (7), 1942–70.

Autor, David, David Dorn, and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Bayer, Christian and Falko Juessen, "On the dynamics of interstate migration: Migration costs and self-selection," *Review of Economic Dynamics*, 2012, *15* (3), 377–401.

Blanchard, Olivier Jean and Lawrence F Katz, "Regional evolutions," *Brookings Papers on Economic Activity*, 1992, *1992* (1), 1–75.

Bryan, Gharad and Melanie Morten, "The aggregate productivity effects of internal migration: Evidence from Indonesia," *Journal of Political Economy*, 2019, *127* (5), 2229–2268.

Bureau of Economic Analysis, "Regional Economic Accounts, Annual Personal Income by State," 2004-2018. https://apps.bea.gov/regional/downloadzip.htm.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro, "Trade and labor market dynamics: General equilibrium analysis of the china trade shock," *Econometrica*, 2019, *87* (3), 741–835.

Cascetta, Ennio, *Transportation Systems Engineering: Theory and Methods*, Vol. 49 of *Applied Optimization*, New York: Springer Science & Business Media, 2001.

Census, U.S., "Income, Poverty and Health Insurance Coverage in the United States: 2010," 2011.

Coen-Pirani, Daniele, "Understanding gross worker flows across US states," *Journal of Monetary Economics*, 2010, *57* (7), 769–784.

Cruz, José-Luis and Esteban Rossi-Hansberg, "The economic geography of global warming," *Review of Economic Studies*, 2024, *91* (2), 899–939.

Davis, Morris A, Jonas DM Fisher, and Marcelo Veracierto, "Migration and urban economic dynamics," *Journal of Economic Dynamics and Control*, 2021, *133*, 104234.

DeWaard, Jack, Janna Johnson, and Stephan Whitaker, "Internal migration in the United States: A comprehensive comparative assessment of the Consumer Credit Panel," *Demographic research*, 2019, *41*, 953.

_ , Mathew Hauer, Elizabeth Fussell, Katherine J Curtis, Stephan D Whitaker, Kathryn McConnell, Kobie Price, David Egan-Robertson, Michael Soto, and Catalina Anampa Castro, "User beware: Concerning findings from the post 2011–2012 US internal revenue service migration data," 2022.

Diamond, Rebecca, "The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000," *American Economic Review*, 2016, *106* (3), 479–524.

Eckert, Fabian and Michael Peters, "Spatial Structural Change," 2022. NBER Working Paper.

Farrokhi, Farid and David Jinkins, "Root growing and path dependence in location choice: Evidence from Danish refugee placement," *Regional Science and Urban Economics*, 2024, *105*, 103975.

Fonseca, Julia, "Less Mainstream Credit, More Payday Borrowing? Evidence from Debt Collection Restrictions," *Journal of Finance*, 2022.

_ **and Jialan Wang**, "How Much do Small Businesses Rely on Personal Credit?," 2022.

Fudenberg, Drew, Wayne Gao, and Annie Liang, "How flexible is that functional form? Quantifying the restrictiveness of theories," 2023.

Fujiwara, Thomas, Eduardo Morales, and Charly Porcher, "A Revealed Preference Approach to Measuring Information Frictions in Migration Decisions," 2022.

Giannone, Elisa, "Skill-Biased Technical Change and Regional Convergence," 2017.

_ **, Qi Li, Nuno Paixao, and Xinle Pang**, "Unpacking Moving," 2020.

Gies Consumer and Small Business Credit Panel, 2004-2018. Accessed July 20, 2021.

Glaeser, Edward L and Joseph Gyourko, "Urban decline and durable housing," *Journal of Political Economy*, 2005, *113* (2), 345–375.

Han, Peter, "The Portfolio Lending Premium in the Mortgage Market," 2023.

Hao, Tongtong, Ruiqi Sun, Trevor Tombe, and Xiaodong Zhu, "The effect of migration policy on growth, structural change, and regional inequality in China," *Journal of Monetary Economics*, 2020, *113*, 112–134.

Heise, Sebastian and Tommaso Porzio, "The aggregate and distributional effects of spatial frictions," 2021. National Bureau of Economic Research Working Paper.

Howard, Greg, "JUE Insight: Moving Cost Magnitudes in Moving Cost Models," *Journal of Urban Economics*, 2026.

_ , "Supplementary Data for "The Dynamics of Internal Migration: A New Fact and Its Implications."," 2026. `https://doi.org/10.3886/ICPSR304488.V1`.

_ **and Hansen Shao**, "Replication Package for "The Dynamics of Internal Migration: A New Fact and its Implications"," 2026. `https://doi.org/10.5281/zenodo.19240923`.

_ **and Jack Liebersohn**, "Why is the rent so darn high? The role of growing demand to live in housing-supply-inelastic cities," *Journal of Urban Economics*, 2021, *124*, 103369.

_ **,** _ **, and Adam Ozimek**, "The Short-and Long-Run Effects of Remote Work on US Housing Markets," *Journal of Financial Economics*, 2023.

Hsieh, Chang-Tai and Enrico Moretti, "Housing constraints and spatial misallocation," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 1–39.

IRS Migration Data, "Internal Revenue Service Statistics of Income Tax Statistics - Migration Data," 2004-2018. https://www.irs.gov/statistics/soi-tax-stats-migration-data.

Jia, Ning, Raven Molloy, Christopher Smith, and Abigail Wozniak, "The economics of internal migration: Advances and policy questions," *Journal of Economic Literature*, 2023, *61* (1), 144–180.

Kaplan, Greg and Sam Schulhofer-Wohl, "Understanding the long-run decline in interstate migration," *International Economic Review*, 2017, *58* (1), 57–94.

Kennan, John and James R Walker, "The effect of expected income on individual migration decisions," *Econometrica*, 2011, *79* (1), 211–251.

Kleemans, Marieke, "Migration choice under risk and liquidity constraints," 2015.

Kleinman, Benny, Ernest Liu, and Stephen Redding, "Dynamic spatial general equilibrium," *Econometrica*, 2023.

Koenen, Martin and Drew Johnston, "Social ties and residential choice: Micro evidence and equilibrium implications," Technical Report, Technical report, Harvard University 2024.

Koşar, Gizem, Tyler Ransom, and Wilbert Van der Klaauw, "Understanding migration aversion using elicited counterfactual choice probabilities," *Journal of Econometrics*, 2021.

Lending, Inc. Cornerstone Home, "These 11 U.S. Cities Will Pay You to Live There," 2021. Accessed August 9, 2022. https://www.houseloanblog.net/get-paid-to-relocate/.

Lind, Nelson and Natalia Ramondo, "Trade with Correlation," *American Economic Review*, February 2023, *113* (2), 317–53.

Mangum, Kyle and Patrick Coate, "Fast locations and slowing labor mobility," 2019.

McFadden, Daniel, "Modelling the choice of residential location," 1978.

Monras, Joan, "Economic shocks and internal migration," 2020.

Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg, "Commuting, migration, and local employment elasticities," *American Economic Review*, 2018, *108* (12), 3855–90.

Morris-Levenson, Joshua and Marta Prato, "The Origins of Regional Specialization." PhD dissertation, The University of Chicago 2022.

Morten, Melanie and Jaqueline Oliveira, "The effects of roads on trade and migration: Evidence from a planned capital city," *American Economic Journal: Applied Economics*, 2024, *16* (2), 389–421.

Oliveira, Jaqueline and Paula Pereda, "The impact of climate change on internal migration in Brazil," *Journal of Environmental Economics and Management*, 2020, *103*, 102340.

Panel Survey of Income Dynamics, 1969-1997. https://psidonline.isr.umich.edu/.

Porcher, Charly, "Migration with costly information," 2020.

Ramani, Arjun and Nicholas Bloom, "The Donut effect of COVID-19 on cities," 2021. National Bureau of Economic Research Working Paper.

Roback, Jennifer, "Wages, rents, and the quality of life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Rosen, Sherwin, "Wage-based indexes of urban quality of life," *Current Issues in Urban Economics*, 1979, pp. 74–104.

Rudik, Ivan, Gary Lyn, Weiliang Tan, and Ariel Ortiz-Bobea, "The Economic Effects of Climate Change in Dynamic Spatial Equilibrium," 2021.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek, "Integrated Public Use Microdata Series: Version 6.0 [dataset]," `http://doi.org/10.18128/D010.V6.0` 2015. Minneapolis: University of Minnesota.

Saks, Raven E and Abigail Wozniak, "Labor reallocation over the business cycle: New evidence from internal migration," *Journal of Labor Economics*, 2011, *29* (4), 697–739.

Schubert, Gregor, "House Price Contagion and US City Migration Networks," 2021.

Tombe, Trevor and Xiaodong Zhu, "Trade, migration, and productivity: A quantitative analysis of china," *American Economic Review*, 2019, *109* (5), 1843–72.

Zabek, Michael A, "Local ties in spatial equilibrium," *American Economic Journal: Macroeconomics*, 2024.

Zerecero, M, "The Birthplace Premium," 2021.