

Gross Migration, Housing and Urban Population Dynamics*

Morris A. Davis
University of Wisconsin
mdavis@bus.wisc.edu

Jonas D.M. Fisher
Federal Reserve Bank of Chicago
jfisher@frbchi.org

Marcelo Veracierto
Federal Reserve Bank of Chicago
mveracierto@frbchi.org

December 2, 2013

Abstract

Cities experience significant, near random walk productivity shocks, yet population is slow to adjust. In practise local population changes are dominated by variation in net migration, and we argue that understanding gross migration is essential to quantify how net migration may slow population adjustments. Housing is also a natural candidate for slowing population adjustments because it is difficult to move, costly to build quickly, and a large durable stock makes a city attractive to potential migrants. We quantify the influence of migration and housing on urban population dynamics using a dynamic general equilibrium model of cities which incorporates a new theory of gross migration motivated by patterns we uncover in a panel of US cities. After assigning values to the model's parameters with an exactly identified procedure, we demonstrate that its implied dynamic responses to productivity shocks of population, gross migration, employment, wages, home construction and house prices strongly resemble those we estimate with our panel data. The empirically validated model implies that costs of attracting workers to cities drive slow population adjustments. Housing plays a very limited role.

JEL Classification Numbers: E0, O4, R0

Keywords: Housing, house prices, labor reallocation

*An earlier draft of this paper circulated under the title “The Role of Housing in Labor Reallocation.” We thank Marco Bassetto, John Kennan, Erzo Luttmer, Robert Shimer and seminar participants at various institutions and conferences for helpful comments. The views expressed herein are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1 Introduction

As we document in this paper, cities experience significant, random-walk-like productivity shocks, yet population is slow to adjust. In the light of [Blanchard and Katz \(1992\)](#)'s empirical evidence that internal migration is integral to equilibrating the US labor market, explaining population's slow adjustment should inform our understanding of macroeconomic labor reallocation. Ultimately migration to and from cities is the main driver of a city's population adjustments. Migration frictions associated with leaving and attracting workers to a city naturally impede population adjustments. Housing is another natural candidate because it is difficult to move, requires time to build, and a large durable housing stock makes a city attractive to potential migrants.

To understand the quantitative importance of migration and housing in urban population dynamics we develop a dynamic general equilibrium model of cities with endogenous migration and local housing and compare it to panel data on 365 US cities over the period 1985-2007. Our model is a version of the [Lucas and Prescott \(1974\)](#) islands economy in which islands are interpreted as cities. We propose a new theory of migration between cities interpreted as population movements between the islands. Population adjustments involve net migration, but we argue that it is essential to model the underlying gross flows. Our argument builds on new evidence from our panel data. We find that gross in- and out-migration are strikingly linear in net migration, evidence that both the decisions to leave and move to a city drive changes in net migration, and that migration clearly involves directed search.

In the model workers face idiosyncratic shocks to their taste for where they currently live and this influences the decision to leave a city. After this decision has been made a worker chooses between directed and undirected search for a new city. Workers understand the distribution of city characteristics, but must use costly directed search to find a city with specific labor and housing markets. Undirected search leaves a worker randomly assigned to a city. Including both directed and undirected migration coincides with evidence that moves involve decisions about where to work and enjoy amenities like housing but also intangible factors such as to be near family members. Increases in employment of the existing population are a obvious alternative to net migration for accommodating local fluctuations in labor demand and so labor supply is endogenous in our model as well.

We introduce this theory of migration and labor supply into an otherwise familiar generalization of the neoclassical growth model. The employed population in each city produces intermediate goods that are imperfectly substitutable in the production of the tradeable fi-

nal goods equipment and consumption. Local production combines employment with freely mobile, durable capital, augmentable by equipment investment, and subject to local total factor productivity (TFP) shocks. Individuals have preferences for consumption and housing services, but only enjoy housing in the same city they work or rest. Housing services are derived from locally produced, immobile and durable residential structures and local residential land.

The model is calibrated to aggregate statistics familiar from other studies that work with the neoclassical growth model and features of the data that are specific to our model’s environment. For the latter, we use our new evidence on the relationship between gross and net migration and [Kennan and Walker \(2011\)](#)’s microeconomic estimates of migration costs to obtain the key migration parameters. In addition, we estimate the idiosyncratic TFP process using our panel data thereby pinning down the model’s exogenous source of persistence and variability. Our estimation of the TFP process facilitates estimation of the dynamic responses of key variables to TFP shocks. We use the estimated elasticity of the employment to population ratio with respect to wages from the impact period of a TFP shock to identify the model’s labor supply elasticity. Finally, we calibrate the substitutability of city-specific intermediate goods so that our model matches the empirical cross-section distribution of population. In so doing we confirm that our model is consistent with Zipf’s law, that in its upper tail city population is distributed exponentially with an exponent close to unity. In turn the idiosyncratic process that we estimate and introduce into the model is able to generate a similar law for TFP that we uncover also in our panel data.

We validate the model by studying several of its over-identifying restrictions, observations not used to calibrate its parameters. Specifically, we compare the model’s dynamic responses to TFP shocks of population, gross in- and out-migration, employment, wages, home construction and house prices to those we estimate from our panel data. The model does surprisingly well along this dimension and importantly it is consistent with the slow response of population to TFP shocks that motivates this study even though this evidence is not directly targeted in our calibration. With only TFP shocks driving within-city dynamics we also find that the model is broadly consistent with the unconditional volatility, persistence, and contemporaneous co-movement of the key variables, although there are some interesting shortcomings.

Having established the empirical relevance of our model, we use it to examine how migration and housing influence population adjustments. We find that the process of attracting workers to cities through costly directed search is the prime determinant of slow population

adjustments to TFP shocks. Housing plays a surprisingly limited role, lowering the amplitude of population’s response to a TFP shock but having very little influence over its persistence. We also investigate our model’s implications for the persistence of urban decline seen in cities like Pittsburgh and Detroit. [Glaeser and Gyourko \(2005\)](#) explore this phenomenon both theoretically and empirically, emphasizing the essential technological characteristics of housing that it cannot be moved, takes time to build, and depreciates slowly, finding that housing is a significant source of persistent urban decline. We find that attracting workers to cities through costly directed search also contributes significantly to this persistence, which is a mechanism not considered by [Glaeser and Gyourko \(2005\)](#).¹

Our model builds on an extensive empirical and theoretical microeconomic literature on migration, surveyed by [Greenwood \(1997\)](#) and [Lucas \(1997\)](#). An important recent contribution is [Kennan and Walker \(2011\)](#) who analyze individual migration decisions in the face of wage shocks and moving costs, but without explicit housing, directed search, or equilibrium interactions. One of their key findings is that there are substantial average net benefits to those who migrate away from cities. By calibrating our model to their estimate, we demonstrate how to map microeconomic estimates of migration costs into the pace of macroeconomic migration flows.

The classic references for systems-of-cities models like ours are [Roback \(1982\)](#) and [Rosen \(1979\)](#). These authors consider static environments in which individuals allocate themselves across cities so that they are indifferent to where they live. Recent contributions using this approach include [Albouy \(2009\)](#) and [Diamond \(2012\)](#). Because it is static, the Roback-Rosen model does not inform our understanding of migration and local population adjustments. [Van Nieuwerburgh and Weil \(2010\)](#) introduce dynamics to this framework and therefore their model speaks to migration. It has implications for net population flows, but not for gross flows. [Coen-Pirani \(2010\)](#) also constructs a dynamic Roback-Rosen model. He studies gross population flows among US states in an environment similar to that used by [Davis, Faberman, and Haltiwanger \(2011\)](#) and others to model gross worker flows among firms. Our empirical work demonstrates that gross population flows in a city are very different from gross worker flows in a firm so we introduce a new theory.

Our model also contributes to the literature by introducing a city’s dynamic response to an identified TFP shock as a model validation tool and by estimating the underlying stochas-

¹In fact, the irreversibility constraint on housing, which is the key mechanism in [Glaeser and Gyourko \(2005\)](#), is never binding in our simulations. Incidentally, this constraint appears not to bind in the data as well as new building permits are always strictly positive in our panel of cities.

tic process for TFP.² Model validation in the existing literature emphasizes unconditional cross-sectional and time-series patterns. Even so, the papers that focus on cities abstract from Zipf’s law, perhaps the most notable feature of the cross-section of cities.³ While the literature relies on idiosyncratic TFP shocks to drive variation, it does not provide evidence on the nature of these shocks as we do.⁴

The recent housing boom and bust has prompted a growing literature that seeks to quantify how frictions in housing may impede migration and labor reallocation and possibly give rise to persistent high unemployment. [Karahan and Rhee \(2012\)](#), [Lloyd-Ellis and Head \(2012\)](#) and [Nenov \(2012\)](#) study how the recent collapse in house prices may have limited labor reallocation through disincentives to migrate arising from home ownership and within-location search frictions in housing and labor markets. We abstract from these within-location labor and housing market frictions and instead introduce between-location migration frictions and focus on housing’s essential technological characteristics.⁵

The rest of the paper is organized as follows. Section 2 describes new empirical evidence on migration and population’s response to TFP shocks based on our panel of cities. After this we use two stripped down versions of our quantitative model to describe our approach to modeling migration and the possible role for housing in slowing population adjustments. Section 5 introduces the complete dynamic quantitative model economy and Section 6 describes how we calibrate its parameters. Section 7 validates the quantitative model by comparing its predictions for within city dynamics we estimate from our panel and quantifies the roles of housing and migration in labor reallocation. The last section concludes.

2 Empirical Evidence

In this section we introduce the empirical evidence that motivates our analysis and guides our modeling of migration. We work with an annual panel data set covering 1985 to 2007 that

²[Lloyd-Ellis, Head, and Sun \(2011\)](#) study the within-city responses of population, residential construction and house prices to personal income shocks identified using a panel VAR and a Choleski decomposition of the variance-covariance matrix of the residuals. These authors abstract from migration decisions and equilibrium interactions among cities.

³See for example [Gabaix \(1999\)](#) and [Eeckhout \(2004\)](#).

⁴[Karahan and Rhee \(2012\)](#) estimate an auto-regressive process in the level of GDP per worker using a short panel of cities.

⁵There is also an empirical literature that investigates the effects of housing related financial frictions on mobility. See for example [Ferreira, Gyourko, and Tracy \(2011\)](#), [Modestino and Dennett \(2012\)](#) and [Schulhofer-Wohl \(2012\)](#). We abstract from financial friction in this paper.

includes population, net and gross migration, employment, wages, residential construction, and house prices for 365 Metropolitan Statistical Areas (MSAs) comprising 83% of the aggregate population.⁶ An MSA is a geographical region with a relatively high population density at its core and close economic ties throughout the area measured by commuting patterns. Such regions are not legally incorporated as a city or town would be, nor are they legal administrative divisions like counties or sovereign entities like states. A typical MSA is centered around a single large city that wields substantial influence over the region, *e.g.* Chicago. However, some metropolitan areas contain more than one large city with no single municipality holding a substantially dominant position, *e.g.* the Dallas–Fort Worth metroplex or Minneapolis–Saint Paul. With these caveats, for convenience we refer to our MSAs as cities.

We use IRS data to calculate city-level net and gross migration rates. [Kaplan and Schulhofer-Wohl \(2012\)](#) suggest that IRS-based state-level migration rates are biased.⁷ Our analysis is based on removing time and city fixed effects so any secular or cross-section biases are accounted for. We work with the IRS data because of its wide coverage of US cities and because we view cities to be the natural unit of analysis for studying migration between geographically distinct labor markets. Due to limited sample sizes gross migration rates can only be calculated for a small number of cities using the other main data sources, the Current Population Survey and the American Community Survey. State-level migration rates can be calculated using these surveys. In our context, these data yield very similar results to those we obtain with city-level and state-level migration rates calculated from the IRS data.

2.1 Gross Versus Net Migration

Let a_{it} and l_{it} denote the number of people flowing into and out of city i in year t and p_{it} the population of that city at the end of the same year. For an individual city the *arrival* (in-migration) rate is $a_{it}/\tilde{p}_{it}$ and the *leaving* (out-migration) rate is $l_{it}/\tilde{p}_{it}$, where $\tilde{p}_{it} = (p_{it-1} + p_{it})/2$. These measures of gross migration mirror the measures of gross job flows defined in [Davis, Haltiwanger, and Schuh \(1998\)](#).⁸ The difference between the arrival

⁶See [Davis, Fisher, and Veracierto \(2011\)](#) for a detailed description of our panel data.


⁷They emphasize three sources of bias: migration rates are based on tax filings and so under-represent the poor and elderly; migration is calculated using addresses on tax forms which are not necessarily home addresses; and tax returns for a given year are not all filed in the same year. The magnitude and direction of these biases is unclear.

⁸In practice we approximate $\tilde{p}_{it}$ as the average of the beginning of year t and end of year t IRS-based population. For additional details see the appendix.

and leaving rates is the net migration rate. Gross migration rates fluctuate over the business cycle and have been falling over our sample period.⁹ To abstract from these dynamics we subtract from each city’s gross rate in a year the corresponding cross section average in that year. The net migration rate calculated from the difference between these gross rates is equivalent to subtracting from each city’s raw net migration rate the corresponding cross-section average net migration rate in each year.

Figure 1 contains plots of gross and net migration rates by population decile with only time effects removed. Net migration is essentially unrelated to city size. This finding reflects Gibrat’s law for cities, that population growth is independent of city size. However, the arrival and leaving rates are clearly diminishing in city size. While we think this is an interesting finding worthy of further study, its presence confounds across-city variation with the within-city dynamics we are interested in. Therefore, after removing time fixed effects, for every city we subtract from each year’s arrival and leaving rate the time series average of the sum of the arrival and leaving rates for that city. This removes city fixed effects in gross migration without affecting net migration rates.

Figure 1: Gross and Net Migration Rates by Population Decile



⁹See [Molloy, Smith, and Wozniak \(2011\)](#) and [Kaplan and Schulfofer-Wohl \(2012\)](#) for studies of the trend in gross migration rates.

Figure 2: Gross Migration Rates by Net Migration




Figure 2 displays mean arrival and leaving rates against mean net migration for each net migration decile, after removing both time and city fixed effects and adding back the corresponding unconditional mean to the gross migration rates. Notice first that gross migration is far in excess of the amount necessary to account for net migration. For example, when net migration is zero an average of 11% of the population either moves in or out of a city in any given year.

Second, the arrival rate is monotonically increasing (and the leaving rate is monotonically decreasing) in net migration. The rising arrival rate suggests that migration involves directed search. Otherwise gross arrivals would be independent of net migration. The fact that the arrival rate rises and the departure rate falls with increases in net migration suggests both margins are important when a city's population adjusts to shocks.

Third, and most striking, the gross migration rates all fall almost exactly on the corresponding regression lines.¹⁰ This evidence sharply contrasts with the non-linear relationships for worker flows at firms described by [Davis, Faberman, and Haltiwanger \(2006\)](#). They find a kink at zero for hires and separations as functions of net worker flows. For negative net flows

¹⁰We obtain virtually identical regression lines when we use all the data rather than first taking averages of deciles. We also find qualitatively similar results when we regress gross on net migration separately for each city in our sample.

hires are flat and close to zero while for positive net flows they are linearly increasing; separations as a function of net flows are essentially the mirror image. The linear relationships between gross and net migration displayed in Figure 2 motivate how we specify migration decisions in our model.

The clear negative relationship between the arrival and leaving rates evident in Figure 2 may be surprising given Coen-Pirani (2010)’s focus on a positive correlation between the two gross migration rates at the state level. This difference does not arise because we consider cities rather than the states considered by Coen-Pirani (2010). It arises from our removal of city-specific fixed effects from the gross migration rates. As suggested by Figure 1, when we do not remove these effects the gross migration rates are strongly positively correlated.¹¹

2.2 Responses of Population and Gross Migration to TFP Shocks

We now describe how we estimate dynamic responses of city-level variables to local TFP shocks and report estimates for population and the gross migration rates. To proceed we exploit the first order conditions of final good producers and intermediate good in the quantitative model described in Section 5. These conditions can be used to derive an equation involving TFP, employment and wages. Using this equation and data on employment and wages we obtain a measure of TFP from which we estimate a stochastic process for its growth. We estimate the dynamic response of a variable to TFP shocks by regressing it on current and lagged values of the TFP innovations derived from the estimated TFP growth process. Later we compare these estimated responses to ones calculated using the same procedure from data simulated from our model.

There are N cities that each produce a distinct intermediate good used as an input into the production of final goods. The production function for a representative firm producing intermediate goods in city i at date t is

$$y_{it} = s_{it} n_{y,it}^{\theta} k_{y,it}^{\gamma}, \quad (1)$$

where s_{it} is exogenous TFP for the city, $n_{y,it}$ is employment, $k_{y,it}$ is capital, hereafter referred to as equipment, $\theta > 0$, $\gamma > 0$, and $\theta + \gamma \leq 1$.¹² The output of the final good at date t , Y_t

¹¹Coen-Pirani (2010) removes cross-sectional variation in the occupational characteristics of states prior to his analysis, but not state fixed effects.

¹²The additional subscripts on employment and equipment are used later to distinguish between employment and equipment used in the production of intermediate goods and residential construction.

is produced using inputs of city-specific intermediate goods according to

$$Y_t = \left[\sum_{i=1}^N y_{it}^\chi \right]^{\frac{1}{\chi}}, \quad (2)$$

where $\chi \leq 1$.

Our measurement of city-specific TFP relies on the following definition. For any variable x_{it} :

$$\hat{x}_{it} \equiv \ln x_{it} - \frac{1}{N} \sum_{j=1}^N \ln x_{jt}. \quad (3)$$

Subtracting the mean value of $\ln x_{jt}$ in each period eliminates variation due to aggregate shocks, allowing us to focus on within-city dynamics. Under the assumption of perfectly mobile equipment the rental rate of equipment is common to all cities. It then follows from the first order conditions of competitively behaving final good and intermediate good producers that

$$\Delta \hat{s}_{it} = \frac{1 - \gamma\chi}{\chi} \Delta \hat{w}_{it} + \frac{1 - \theta\chi - \gamma\chi}{\chi} \Delta \hat{n}_{it}, \quad (4)$$


where Δ is the first difference operator and w_i denotes the wage in city i .¹³ Applying the first difference operator eliminates permanent differences in TFP among the cities. Assuming values for χ , θ and γ , and substituting data on wages and employment for $\Delta \hat{w}_{it}$ and $\Delta \hat{n}_{it}$, we use this equation to measure $\Delta \hat{s}_{it}$, the growth rate of city-specific TFP.

Below we calibrate θ and γ using traditional methods and find a value for χ to match the model to Zipf's law. With calibrated values $\chi = 0.9$, $\theta = 0.66$ and $\gamma = .235$ we estimate a first order auto-regression in $\Delta \hat{s}_{it}$ with an auto-correlation coefficient equal to 0.24 and the standard deviation of the error term equal to 0.015. [Wooldridge \(2002\)](#)'s test of the null of no first order serial correlation in the residuals is not rejected, suggesting that this specification is a good fit for the data.

A natural concern about measuring TFP with (4) is that it ignores agglomeration. [Davis, Fisher, and Whited \(2013\)](#) find statistically significant agglomeration effects with a similar approach to measurement in which they model agglomeration as affecting TFP through an externality in output per acre of land as in [Ciccone and Hall \(1996\)](#). It is straightforward to modify equation (4) to include agglomeration in this way and it leads to the same measurement equation for the exogenous component of TFP except that the coefficients on wage and output growth also include the parameter governing the magnitude of the externality. When

¹³See the technical appendix for more details.

Figure 3: Responses of TFP and Population



Note: Point estimates along with 2 standard error bands.

we re-estimate the TFP process using the estimate of the externality parameter in [Davis et al. \(2013\)](#) we find the serial correlation coefficient and the innovation standard deviation fall to 0.20 and 0.013. While we do not include agglomeration in our model, we conjecture that doing so would reconcile the two sets of estimates but have little impact on our other results.¹⁴


We now show how to use the estimated TFP process (without agglomeration) to identify the dynamic responses of variables to exogenous local TFP shocks. Let e_{it} denote the residual from the estimated TFP growth auto-regression. Then, we estimate the dynamic response to a TFP shock of variable $\Delta \hat{x}_{it}$ as the coefficients $b_0, b_1, \dots, b_4$ from the following panel regression:

$$\Delta \hat{x}_{it} = \sum_{l=0}^4 b_l e_{it-l} + u_{it} \quad (5)$$

where u_{it} is an error term which is orthogonal to the other right-hand-side variables under the maintained hypothesis that the process for TFP growth is correctly specified. The dynamic

¹⁴Verifying this conjecture is beyond the scope of this paper. However, in the model considered by [Davis et al. \(2013\)](#) the externality amplifies the response of TFP to an exogenous TFP shock and makes it more persistent.

Figure 4: Responses of Arrival and Leaving Rates



Note: Point estimates along with 2 standard error bands.

response of $\hat{x}_{it}$ is obtained by summing the estimated coefficients appropriately. For the gross migration rates we replace $\Delta\hat{x}_{it}$ with the rates themselves (transformed as described above) in (5) and identify the dynamic responses with the estimated coefficients directly.

Figure 3 displays the percentage point deviation responses of TFP and population to a 1 standard deviation impulse to measured TFP. This plot establishes the claim made in the introduction that productivity, that is TFP, responds much like a random walk, rising quickly to its new long run level, and that population responds far more slowly. Figure 4 shows that the adjustment of population occurs along both the arrival and leaving margins, as suggested by our earlier discussion of Figure 2. On impact the arrival rate jumps up and the leaving rate jumps down and then both slowly returns to their long run levels. The indicated sampling uncertainty suggests that the arrival and leaving margins are about equally important in the adjustment of population to a TFP shock. In particular it is the improvement in local prospects encouraging workers not to move as much as the affect those prospects have on attracting workers to the city through which population adjusts to persistent improvements in local TFP.

3 Modeling Migration

The previous section documents evidence confirming a role for both gross migration margins in population adjustments. We now introduce our theory of migration that is motivated by this evidence. To do so we employ a simple, static model which abstracts from housing, equipment, and labor supply. We use this simplified approach to develop intuition about migration choices, to describe how and why we can reproduce the relationships depicted in Figure 2, and to establish that modeling gross migration is essential for understanding population adjustments. All of the results in this section extend to our more general quantitative model.

3.1 A Static Model of Migration

The economy consists of a large number of geographically distinct cities with initial population x . In each city there are firms which produce identical, freely tradeable consumption goods with the technology sn^θ , where s is a city-wide TFP shock, n is labor and $0 < \theta < 1$. There is a representative household with a unit continuum of members that are distributed across city types $z = (s, x)$ according to the measure μ . Each household member enjoys consumption, C , and supplies a unit of labor inelastically. After the TFP shocks have been realized, but before production takes place, the household decides how many of its members leave each city and how many of those chosen to leave move to each city. Once these migration decisions have been made, production and consumption take place.

The leaving decision is based on each household member receiving a *location-taste* shock ψ , with measure μ_l , that subtracts from their utility of staying in the city in which they are initially located. This kind of shock is used by [Kennan and Walker \(2011\)](#) in their measurement of migration costs. To help us match the empirical evidence on the relationship between gross and net migration we make a parametric assumption for the distribution of individual location-taste shocks in a city of type z :

$$\int_{-\infty}^{\bar{\psi}(l(z)/x)} \psi d\mu_l = -\psi_1 \frac{l(z)}{x} + \frac{\psi_2}{2} \left(\frac{l(z)}{x} \right)^2$$

where the parameters ψ_1 and ψ_2 are both non-negative and $\bar{\psi}(l(z)/x)$ is defined by

$$\frac{l(z)}{x} = \int_{-\infty}^{\bar{\psi}(l(z)/x)} d\mu_l.$$

This parameterization is U-shaped starting at the origin. Initially benefits accrue to increasing the number of leavers from a city, and eventually individuals find it very costly to leave. These features are consistent with evidence in Kennan and Walker (2011) that individuals who move receive substantial non-pecuniary benefits and that non-movers would find it extremely costly if they were forced to move. For example, many individuals move to be near family members or find it very costly to move because they are already near family members. As more people leave a city the remaining inhabitants are those who have a strong preference for living in that city. Subject to these shocks, the household determines how many of its members from each city must find new cities in which to work. Household members chosen to find new cities are called *leavers*.

When deciding where to send its leavers the household understands the distribution of city types μ but does not know the location of any specific type z . However, it can find a particular type of city by obtaining a *guided trip*, a form of *directed search*. To match the evidence on gross and net migration, we adopt a particular functional form for producing guided trips as well. Specifically, by giving up u units of utility each individual household member can produce $\sqrt{2}A^{-1/2}u^{1/2}$ guided trips to the city in which they are initially located, where the parameter A is non-negative. Therefore, to attract $a(z)$ workers to a city of the indicated type the household must incur a total utility cost of $(A/2)(a(z)/x)^2x$.

The production of guided trips encompasses the many ways in which workers are attracted to specific cities, including via informal contacts between friends and family, professional networks, specialized firms like head-hunters, advertising that promotes cities as desirable places to live and work, firms' human resource departments, and via recruiting by workers whose primary responsibility is some other productive activity.¹⁵ Clearly some of these activities are part of recruiting workers within a local labor market and as such would be included in any measurement of the vacancy costs typically assumed in models of labor market search and matching. Our approach can be thought of as capturing the portion of these activities devoted to attracting workers to a local labor market.

If a household member does not obtain a guided trip it can migrate to another city using *undirected search*. Specifically, by incurring a utility cost τ a leaver is randomly allocated to another city in proportion to its initial population. Including undirected search captures the idea that choosing to move to a particular city is often the outcome of idiosyncratic

¹⁵For convenience we have modeled the cost of attracting workers to a city as a direct loss of utility. Our results do not rest on this assumption.

factors other than wages or amenities that are difficult to model explicitly.¹⁶ Furthermore, it is natural to let people move to a location without forcing them to find someone to guide them.

We characterize allocations in this economy by solving the following planning problem:

$$\max_{\{C, \Lambda, a(z), l(z), p(z)\}} \left\{ \ln C - \int \left[\frac{A}{2} \left(\frac{a(z)}{x} \right)^2 x + \left(-\psi_1 \frac{l(z)}{x} + \frac{\psi_2}{2} \left(\frac{l(z)}{x} \right)^2 \right) x \right] d\mu - \tau \Lambda \right\} \quad (6)$$

subject to

$$p(z) \leq x + a(z) + \Lambda x - l(z), \forall z \quad (7)$$

$$\int [a(z) + \Lambda x] d\mu \leq \int l(z) d\mu \quad (8)$$

$$C \leq \int sp(z)^\theta d\mu \quad (9)$$

and non-negativity constraints on the choice variables. The variable Λ is the fraction of the household that engages in undirected search. Since these workers are allocated to cities in proportion to their initial populations, Λ also corresponds to the share of a city's initial population that migrates to that city within the period. Constraint (7) states that population in a city is no greater than the initial population plus arrivals through guided trips and undirected search minus the number of workers who migrate out of the city. Constraint (8) says that total arrivals can be no greater than the total number of workers who migrate out of cities and (9) restricts consumption to be no greater than total production, taking into account that each individual supplies a unit of labor inelastically, $n(z) = p(z)$, $\forall z$.

3.2 Why Both Gross Migration Frictions are Necessary

We now explain why it is necessary to include frictions on both gross migration margins in order to match the evidence depicted in Figure 2. Suppose $A = 0$ so that guided trips can be produced at no cost, but that household members continue to be subject to location-taste shocks, $\psi_1 > 0$ and $\psi_2 > 0$. Then it is straightforward to show

$$\begin{aligned} \frac{a(z)}{x} &= \max \left\{ \frac{p(z) - x}{x} - \Lambda + \frac{\psi_1}{\psi_2}, 0 \right\}; \\ \frac{l(z)}{x} &= \max \left\{ \frac{\psi_1}{\psi_2}, - \left(\frac{p(z) - x}{x} \right) + \Lambda \right\}. \end{aligned}$$

¹⁶Kennan and Walker (2011) model all migration as resulting from undirected search.

Observe that as long as the net population growth rate, $(p(z) - x)/x$, is not too negative, the planner sets the leaving rate, $l(z)/x$ at the point of maximum benefits, ψ_1/ψ_2 , and adjusts population using the arrival rate, $a(z)/x$, only. In this situation the leaving rate is independent of net population adjustments, contradicting the evidence presented in Figure 2.

Now suppose that there are no location-taste shocks, $\psi_1 = \psi_2 = 0$, but it is costly to create guided trips, $A > 0$. In this case we find

$$\begin{aligned}\frac{l(z)}{x} &= \max \left\{ - \left(\frac{p(z) - x}{x} - \Lambda \right), 0 \right\}; \\ \frac{a(z)}{x} &= \max \left\{ \frac{p(z) - x}{x} - \Lambda, 0 \right\}.\end{aligned}$$

Without taste shocks the planner always goes to a corner: when net population growth is positive the leaving rate is set to zero, and when net population growth is negative the arrival rate is set to zero. Clearly the relationship between gross and net migration in this situation also contradicts the evidence depicted in Figure 2. We conclude that to be consistent with the relationship between gross and net migration, it is necessary to include frictions on both gross migration margins.

3.3 Migration Trade-offs and Reproducing Figure 2

For the model to be consistent with the gross flows data as depicted in Figure 2, it also must be true (almost everywhere) that the number of workers leaving a city and the number arriving to the same city using guided trips are both strictly positive, $l(z) > 0$ and $a(z) > 0$. The reason we require $l(z) > 0$ is that gross out-migration is always positive in Figure 2. The reason we require $a(z) > 0$ is that otherwise there would be intervals of net migration in which arrival rates are constant, equal to Λ , which is also inconsistent with Figure 2. Therefore, unless otherwise noted, from now on we assume that $a(z) > 0$ and $l(z) > 0$.

The planner's first order conditions for Λ , $a(z)$, $l(z)$ and $p(z)$ are

$$\tau = \int \lambda \xi(z) x d\mu - \lambda \eta \tag{10}$$

$$\lambda \xi(z) - A \frac{a(z)}{x} = \lambda \eta \tag{11}$$

$$\lambda \xi(z) = \psi_1 - \psi_2 \frac{l(z)}{x} + \lambda \eta \tag{12}$$

$$\xi(z) = s \theta p(z)^{\theta-1} \tag{13}$$

where λ is the marginal utility of consumption and $\lambda\xi(z)$ and $\lambda\eta$ are the Lagrange multipliers corresponding to (7) and (8). The multipliers measure the value of an additional worker in a particular city and the cost of pulling an additional worker from the pool of available migrants. We use (10)–(13) to illustrate the trade-offs between the decision to leave a city and how many workers to allocate to guided trips or undirected search.

Combining (10) with (11) we find

$$\tau = \int Aa(z) d\mu.$$

This equation describes the trade-off between using guided trips and undirected search. The marginal cost of raising the fraction of household members engaged in undirected search is equated to the average marginal cost of allocating those household members using guided trips. The averaging reflects the fact that undirected search allocates workers in proportion to each city's initial population.

Combining (11) and (12) we see that

$$A \frac{a(z)}{x} = \psi_1 - \psi_2 \frac{l(z)}{x}.$$

Intuitively, migration out of a city increases to the point where the marginal benefits of doing so (recall that the location-taste shocks initially imply benefits to leaving a city) are equated with the marginal cost of attracting workers into the city.

Finally, we see from (13) that the shadow value of bringing an extra worker to a city equals the marginal product of labor in that city. It follows from (11) and (12) that absent migration frictions, $A = \psi_1 = \psi_2 = 0$, the efficient allocation of workers across cities involves equating cities' marginal products of labor. This contrasts with the classic [Roback \(1982\)](#) and [Rosen \(1979\)](#) model of a system of cities with free mobility in which the level of utility is equated across workers in different cities. The difference arises from the fact that we have assumed perfect consumption insurance. Migration frictions drive a wedge between marginal products of labor because heterogeneity in initial populations implies differential costs of moving workers around.

We now derive how gross migration relates to net migration. From the first order conditions for $a(z)$ and $l(z)$, (11) and (12), and the population constraint, (7), it is straightforward to show that

$$\frac{a(z)}{x} + \Lambda = \frac{\psi_1}{A + \psi_2} + \frac{A}{A + \psi_2} \Lambda + \frac{\psi_2}{A + \psi_2} \left(\frac{p(z) - x}{x} \right). \quad (14)$$

Clearly, the arrival rate is a linear function of the net migration rate $(p(z) - x)/x$ with the linear coefficient satisfying $0 < \psi_2/(\psi_2 + A) < 1$. Similarly the leaving rate is given by:

$$\frac{l(z)}{x} = \frac{\psi_1}{A + \psi_2} + \frac{A}{A + \psi_2}\Lambda - \frac{A}{A + \psi_2} \left(\frac{p(z) - x}{x} \right). \quad (15)$$

This is also a linear function of the net migration rate with the linear coefficient satisfying $-1 < -A/(\psi_2 + A) < 0$.

Equations (14) and (15) establish that gross migration in the model can be made consistent with the linear relationships depicted in Figure 2. This result is the underlying reason for the quadratic specifications of the location-taste shocks and the production of guided trips. In other words, the relationship between gross and net migration depicted in Figure 2 places strong restrictions on the nature of migration frictions.

3.4 Gross Migration and Population Adjustments

Modeling both gross migration margins is important for replicating Figure 2, but it also plays a crucial role in determining the speed of population adjustments. This can be seen by substituting for $a(z)$ and $l(z)$ in the original planning problem using (14) and (15), which simplifies it to

$$\max_{\{p(z), \Lambda\}} \left\{ \ln \int sp(z)^\theta d\mu - \int \left[\Phi(\Lambda) + \frac{1}{2} \frac{A\psi_2}{A + \psi_2} \left(\frac{p(z) - x}{x} \right)^2 x \right] d\mu - \tau\Lambda \right\}$$

subject to:

$$\int p(z) d\mu = \int x d\mu \quad (16)$$

with non-negativity constraints on the choice variables and where $\Phi(\Lambda)$ is a quadratic function in Λ involving the underlying structural parameters ψ_1, ψ_2 and A . In deriving this simplified planning problem we have used the fact that (7) and (8) reduce to (16) and that this constraint holds with equality at the optimum. Similarly we have used (9) to substitute for consumption in the planner's objective function.

When the planning problem is written in this way we see that population adjustments do not involve the gross migration decisions $a(z)$ and $l(z)$. Nevertheless modeling these decisions matters for understanding population adjustments because the coefficient that determines the cost of net population adjustments, $A\psi_2/(A + \psi_2)$, involves gross migration parameters. The gross in-migration decision matters through the parameter A and gross

out-migration matters through ψ_2 . Also notice that the reduced form costs of adjusting population are quadratic. This is a direct consequence of specifying the location-taste shocks and production of guided trips to reproduce Figure 2. That is, the relationship between gross and net migration displayed in Figure 2 implies quadratic adjustment costs in net population adjustments.

Finally, notice that as long as $a(z) > 0$ and $l(z) > 0$, a maintained assumption in the statement of the simplified planning problem, population adjustments are independent of the undirected search decision. Undirected search is determined by the solution to

$$\tau = \Phi'(\Lambda).$$

An implication of this property is that as long as arrivals are always positive undirected search plays no role in net population adjustments. In the more general quantitative model arrivals will be set to zero in especially undesirable cities. Still, for most cases arrivals are strictly positive so that the amount of equilibrium undirected search is relatively unimportant for our results. This is a useful property given that there is little evidence on the share of in-migration that is a result of undirected versus directed search. Nevertheless we include undirected search in the model because, as emphasized above, otherwise workers would have no way to move other than to obtain a guided trip and we think this is implausible.

3.5 One Possible Decentralization

The challenge for decentralizing the planning problem is how to treat guided trips. One valid approach is to have guided trips allocated entirely within the household through home production without any market interactions. We view guided trips in the model as an amalgam of both market and non-market activities and so we think a more natural decentralization is one that involves both market transactions and home production. We now consider such a decentralization.

Markets are competitive. Firms in a city of type z hire labor at wage $w(z)$ and produce consumption goods to maximize profits. Household members initially located in a type- z city produce $a_m(z)$ guided trips to that city which they sell to prospective migrants at price $q(z)$. The household also home produces guided trips for use by its own members and we denote these by $a_h(z)$. Let $m(z)$ denote the total number of guided trips to z -type cities purchased by household members in the market.

The representative household solves the following optimization problem

$$\max_{\substack{\{C, \Lambda, m(z), \\ a_m(z), a_h(z), \\ l(z), p(z)\}}} \left\{ \ln C - \int \left[\frac{A}{2} \left(\frac{a_m(z) + a_h(z)}{x} \right)^2 x + \left(-\psi_1 \frac{l(z)}{x} + \frac{\psi_2}{2} \left(\frac{l(z)}{x} \right)^2 \right) x \right] d\mu - \tau \Lambda \right\} \quad (17)$$

subject to:

$$C + \int q(z) m(z) d\mu = \int q(z) a_m(z) d\mu + \int w(z) p(z) d\mu + \int \Pi(z) d\mu \quad (18)$$

$$p(z) = x + m(z) + a_h(z) + \Lambda x - l(z), \forall z \quad (19)$$

$$\int [m(z) + a_h(z) + \Lambda x] d\mu = \int l(z) d\mu \quad (20)$$

along with non-negativity constraints on the choice variables. Equation (18) is the household's budget constraint where Π denotes profits from owning the firms. Equation (19) states that the population of a city after migration equals the initial population plus migrants from guided trips and undirected search less the initial population that migrates out of the city. Finally, equation (20) states that the household members that migrate to cities must equal the number of household members that migrate out of cities.

A competitive equilibrium is defined in the usual way with the market clearing conditions

$$\begin{aligned} m(z) &= a_m(z), \forall z \\ n(z) &= p(z), \forall z \\ C &= \int s n(z)^\theta d\mu \end{aligned}$$

which correspond to the markets for guided trips and labor in each city and for consumption. Using $m(z) = a_m(z)$ and the first order conditions of the household's problem we verify that a competitive equilibrium only determines the total number of guided trips into a city $a_m(z) + a_h(z)$; the composition of these guided trips between market and non-market activities is left undetermined.¹⁷

This particular decentralization makes it possible to calculate the total value of guided trips. In particular, as long as there are some guided trips purchased in the open market the total value of these trips is $q(z)a(z)$, with $a(z) = a_m(z) + a_h(z)$ and $q(z) = C A a(z)/x$. We use the total value of guided trips to help calibrate our model to the estimate of average moving costs in Kennan and Walker (2011). Since the split of guided trips between market and non-market activities is unknowable it is ambiguous how to include them when measuring

¹⁷For details see the technical appendix.

employment, wages and aggregate output in the quantitative model. Therefore another advantage of this decentralization is that we can use it to bound the impact of guided trips on our calibration.

4 The Potential Role for Housing

We expect housing to influence population adjustments for the reasons discussed in the introduction: it is costly to build quickly, durable and immobile. This section describes a simplified version of the quantitative model developed below to help understand why housing is an independent source of slow population adjustments. The model borrows the geography and structure of consumption good production from the previous section. There are three differences with that model. First, individuals have a preference for housing services derived from durable, immobile and locally produced structures. Second, to emphasize the role of housing the model excludes migration frictions. Third, because housing is durable the model introduces dynamics in the form of infinitely lived households.

To analyze this model it is convenient to exploit the fact that a competitive equilibrium can be obtained as the solution to a city planning problem that maximizes local surplus taking economy-wide variables as given, where these economy-wide variables must satisfy certain side conditions in an equilibrium. We discuss this property further in the context of the quantitative model below. For simplicity here we focus on the city-planner's problem taking the aggregate variables as given and ignoring the side conditions. In particular we assume that the shadow price of populating the city with an additional individual is exogenous and equal to $\eta > 0$ and the price of the consumption good is normalized to unity.

The recursive formulation of the city-planners problem is

$$V(h, s) = \max_{p, p_y, p_h, h'} \left\{ s p_y^\theta + H \ln\left(\frac{h^\zeta}{p}\right) p - \eta p + \beta V(h', s') \right\} \quad (21)$$

subject to:

$$p_y + p_h = p \quad (22)$$

$$h' = (1 - \delta_h)h + p_h^\alpha \quad (23)$$

with non-negativity constraints on the choice variables and the law of motion for s and the prime symbol denotes next period's value of a variable. The value function V depends on the current stock of housing, h , and productivity, s . It does not depend on population

because of the absence of migration frictions. The current surplus in the Bellman equation (21) equals the sum of local consumption good production and housing services enjoyed, less the cost to the planner of populating the city with p individuals. The planner allocates its choice of population either to consumption goods production, p_y , or housing construction, p_h , indicated by constraint (22). The housing constraint (23) embodies the assumptions that housing takes one year to build, is durable with depreciation rate $\delta_h \in (0, 1)$ and is immobile. TFP in the construction sector is independent of that in the consumption sector. We maintain this assumption throughout because we expect variation in construction TFP across cities is much smaller than in other industries. Consumption goods production and construction is subject to diminishing returns to labor, parameterized by $\theta, \alpha \in (0, 1)$. The parameter $\beta \in (0, 1)$ is the discount factor.

An individual's preference for housing is logarithmic with scale parameter $H > 0$ and housing services are shared equally across the population. The logarithmic assumption is not innocuous as we highlight below. We work with logarithmic preferences for two reasons. First, we want our framework to be compatible with macroeconomic analysis and hence consistent with balanced growth. Second, logarithmic preferences imply a constant share of housing services in households' expenditures and are therefore consistent with the empirical evidence reported in [Davis and Ortalo-Mangé \(2011\)](#). Housing services are subject to diminishing returns in the stock of structures, $\zeta \in (0, 1)$, consistent with land being an input into housing services.

We now consider the equations that characterize a city's dynamics given η . Combining the first order conditions for population and the allocation of labor to the two production sectors with the population constraint (22) yields

$$s\theta f(p, \psi; \alpha, \theta)^{\theta-1} + H \ln\left(\frac{h^\zeta}{p}\right) - H = \eta \quad (24)$$

where $p_y = f(p, \psi; \alpha, \theta)$ and $f(\cdot)$ describes the allocation of labor to the consumption sector as a function of total population and the shadow price of housing in the city ψ , the latter being the Lagrange multiplier on constraint (23). Equation (24) shows that the planner brings people into the city up to the point where the shadow cost of the last individual is equated to the marginal surplus derived from having that individual in the city. This marginal surplus equals the marginal product of their labor plus the housing services they enjoy less the reduction in housing services enjoyed by the population already in the city because the same housing must be allocated among more individuals.

The shadow price ψ satisfies the Euler equation

$$\psi = \beta \mathcal{E}_{h,s} \left[\zeta H \frac{p'}{h'} + \psi' (1 - \delta_h) \right] \quad (25)$$

where $\mathcal{E}_{h,s}$ denotes expectations conditional on the current stock of housing, productivity and the law of motion for productivity. The current price of an extra unit of housing depends on the discounted value of the services the housing will provide plus the undepreciated value of that housing going forward. Equations (24) and (25) along with the constraints (22) and (23) characterize housing and population dynamics in the city. Clearly these dynamics depend on all of the model's parameters and so finding plausible values for these parameters is essential for determining housing's role in population dynamics.

To highlight the potential for housing to influence population adjustments we consider a city starting from steady state, subject it to a one time permanent increase in productivity s , and then follow the paths of population and housing to the new steady state. Suppose housing services are not valued, $H = 0$. In this case $\psi = 0$ and it follows from (24) that population adjusts instantaneously to the permanent increase in productivity. Next, suppose housing is valued $H > 0$ but that $\alpha = 0$. In this case housing is a constant equal to $1/\delta_h$ and all labor is allocated to the consumption sector. With $h = 1/\delta_h$ equation (24) reduces to

$$s\theta p^{\theta-1} - H \ln(\delta_h p) - H = \eta.$$


This equation shows that population is a function of productivity only so that after a permanent increase in productivity population adjusts instantaneously. We conclude that without housing or with constant housing population adjusts immediately to permanent shocks to productivity.

Finally, consider the intermediate case where $H = .205$ and $\alpha = .785$, values taken from the calibration of the quantitative model reported in Table 1 below.¹⁸ The dynamics of population and the housing stock after the one time increase in productivity are displayed in Figure 5. For $\alpha > 0$ the planner trades off taking advantage of the new higher level of productivity in the consumption sector with building more housing to accommodate additional workers. This trade-off affects the speed of adjustment of housing directly and population indirectly depending on the values of the other parameters.

From Figure 5 we see that housing slows population's convergence to the new steady state compared to the previous two cases. However housing is much slower than population

¹⁸The other parameters are also set to their calibrated values: $\delta = .064$, $\zeta = .785$ and $\theta = .66$. We use a log-linear approximation of equations (22)–(25) evaluated at steady state to compute the equilibrium dynamics.

Figure 5: Housing and Population Adjustments



Source: Author's calculations.

to converge. It is difficult to draw a definitive conclusion from this finding because of the role played by the other model parameters and in the quantitative model additional interactions and parameters influence the dynamics. Finally, it should be clear that the logarithmic housing preferences play an important role in this example. In the extreme case typical in the literature, for example [Lloyd-Ellis et al. \(2011\)](#), individual housing demand is completely inelastic so that population and housing follow identical dynamics.

5 The Quantitative Model

This section describes the model we use to quantify housing and migration's influence on urban population dynamics. It combines the models described in the previous two sections so that the migration decisions are now dynamic and incorporates additional features to bring the model closer to the data and facilitate its calibration. The section begins with a description of the economy's environment followed by a characterization of the stationary competitive equilibrium as the solution to a representative city social planning problem with side conditions.

5.1 The Environment

As before the economy consists of a continuum of geographically distinct locations called cities that are subject to idiosyncratic TFP shocks. Cities are distinguished by their stock of housing, h , initial population, x , and the current and lagged TFP, s and s_{-1} , with measure over these state variables given by μ .¹⁹ Within cities there are three production sectors corresponding to intermediate goods, construction and housing services. The representative firm of each sector maximizes profits taking prices as given. Intermediate goods are distinct to a city and imperfectly substitutable in the production of the freely tradeable final goods consumption and equipment. The technologies for producing intermediate and final goods are identical to those underlying our estimates of TFP, described in equations (1) and (2).²⁰ The representative firm in the construction sector augments the local stock of residential structures, h , using labor, n_h , equipment, k_h , and land, b_h , according to

$$h' = (1 - \delta_h) h + n_h^\alpha k_h^\vartheta b_h^{1-\alpha-\vartheta} \quad (26)$$

where $\alpha > 0$, $\vartheta > 0$ and $\alpha + \vartheta < 1$. This specification replaces (23) from the previous section and retains its implications that housing is locally produced, durable, immobile and that changes in local TFP s do not affect residential construction. We assume that equipment used in production and construction is homogenous. Residential structures are combined with land, b_r , to produce housing services $h^{1-\zeta} b_r^\zeta$.

The representative household faces the same migration choices as discussed in Section 3 but being infinitely lived it takes into account the affects of current migration decisions on the allocation across cities of its members' in future periods. In particular, it is now bound by the constraint

$$x' = p \quad (27)$$

in each city where p continues to denote the post-migration population of a city. Its members have logarithmic preferences for consumption and housing services and face a non-trivial labor supply decision. For the latter we assume that each period, after the migration decisions have been made, but before production and construction take place, individual household members receive a labor disutility shock φ with measure μ_n . Similar to our treatment of migration costs we make a parametric assumption for the average disutility of working.

¹⁹Current and lagged TFP both appear in this list to accommodate the estimated TFP process described in Section 2.2. This is discussed further below.

²⁰Equations (1) and (2) are written in terms of the location of a city, indexed by i , but here it is convenient to index them by the type of the city as represented by its state vector (h, x, s, s_{-1}) .

Specifically, if the household decides n of its members in a city will work for a year these costs are specified as

$$\int_{-\infty}^{\bar{\varphi}(n/p)} \varphi d\mu_n = \phi \left(\frac{n}{p} \right)^\pi,$$

where $\phi > 0$, $\pi \geq 1$ and $\bar{\varphi}(n/p)$ satisfies

$$\frac{n}{p} = \int_0^{\bar{\varphi}(n/p)} d\mu_n.$$

The parameter π governs the elasticity of a city's labor supply with respect to the local wage. We include a labor supply decision because a natural alternative to adjusting a city's population to prevailing conditions is for the employment to population ratio to change.

5.2 Steady State Equilibrium

We consider a steady state competitive equilibrium. Since the model is a convex economy with no distortions, the welfare theorems apply. As a consequence the equilibrium allocation can be obtained by solving the problem of a social planner that maximizes the expected utility of the representative household subject to resource feasibility constraints. However, it is more useful to characterize the equilibrium allocation as the solution to a representative city social planner's problem with side conditions. This approach to studying the equilibrium allocation follows [Alvarez and Shimer \(2011\)](#) and [Alvarez and Veracierto \(2012\)](#).

The city planner enters a period with the state vector $z = (h, x, s, s_{-1})$. Taking as given aggregate output of tradeable final goods, Y , the marginal utility of consumption, λ , the shadow value of adding one individual to the city's population exclusive of the arrival and leaving costs, $\lambda\eta$, the shadow value of equipment, λr_k , the arrival rate of workers through undirected search Λ , and the transition function for TFP, $Q(s'; s, s_{-1})$, the representative city planner solves

$$\begin{aligned} V(z) = \max_{\substack{\{n_y, n_h, k_y, k_h, \\ h', b_r, b_h, p, a, l\}}} & \left\{ \lambda \frac{1}{\chi} Y^{1-\chi} [s n_y^\theta k_y^\gamma]^\chi + H \ln \left(\frac{h^{1-\zeta} b_r^\zeta}{p} \right) p - \phi (n_y + n_h)^\pi p^{1-\pi} \right. \\ & \left. - \lambda r_k (k_y + k_h) - \lambda \eta (a + \Lambda x - l) \right. \\ & \left. - \frac{A}{2} \left(\frac{a}{x} \right)^2 x - \left[-\psi_1 \frac{l}{x} + \frac{\psi_2}{2} \left(\frac{l}{x} \right)^2 \right] x + \beta \int_{s'} V(z') dQ(s'; s, s_{-1}) \right\} \end{aligned}$$

subject to

$$\begin{aligned} p &= x + a + \Lambda x - l \\ n_y + n_h &\leq p \\ b_r + b_h &= 1 \end{aligned} \tag{28}$$

plus (26), (27), and non-negativity constraints on the choice variables.

The objective of the optimization problem is to maximize the expected present discounted value of local surplus. To see this note that the first two terms are the value of intermediate good production and the amount of housing services consumed in the city. The next five terms comprise the contemporaneous costs to the planner of obtaining this surplus: the disutility of sending the indicated number of people to work, the shadow cost of equipment used in the city, and the disutility of net migration inclusive of guided trip production and taste-for-location shocks. The last term is the discounted continuation value given the updated state vector. Constraining the achievement of the city planner's objective are the local resource constraints, the housing and population transition equations and the non-negativity constraints on the choice variables. Note that in the statement of the land constraint we have normalized the local endowment of residential land to unity and used the fact that land being used for current housing services cannot be used to build new structures on. We have also normalized commercial land to unity and assumed that it cannot be converted into residential land and vice versa.

Let $\lambda\xi(z)$ denote the Lagrange multiplier corresponding to constraint (28) in the city planner's problem. This function represents the shadow value of bringing an additional individual to a type- z city. From the first order conditions of the city-planner's problem it is easy to show that

$$\lambda\xi(z) = \begin{cases} A \left[\frac{a(z)}{x} \right] + \lambda\eta, & \text{if } a(z) > 0, \\ \left[\psi_1 - \psi_2 \left(\frac{l(z)}{x} \right) \right] + \lambda\eta, & \text{if } l(z) > 0. \end{cases} \tag{29}$$

which takes into account the fact that $a(z) = l(z) = 0$ will never occur in equilibrium. Notice that if both gross migration rates are positive then the shadow value of a worker is the same as in the static case.

The steady state allocation is the solution to the city planners problem that satisfies particular side conditions. To begin, let $(n_y, n_h, k_y, k_h, h', b_r, b_h, p, a, l)$ be the optimal decision rules for the city planner's problem that takes $(Y, \lambda, \eta, r_k, \Lambda)$ as given and μ be the invariant

distribution generated by the optimal decision rules (h', p) and the transition function Q . In addition define

$$\begin{aligned} K &= \int (k_y + k_h) d\mu \\ C &= Y - \delta_k K \end{aligned}$$

These two equations define the aggregate equipment stock and per capita consumption, where $0 < \delta_k < 1$ denotes the equipment depreciation rate. Now suppose the following equations are satisfied

$$Y = \left\{ \int \left[s n_y(z)^\theta k_y(z)^\gamma \right]^\chi d\mu \right\}^{\frac{1}{\chi}} \quad (30)$$

$$\lambda = \frac{1}{C} \quad (31)$$

$$\int a(z) d\mu + \Lambda = \int l(z) d\mu \quad (32)$$

$$r_k = \frac{1}{\beta} - 1 + \delta_k \quad (33)$$

$$\lambda \int [\xi(z) - \eta] x d\mu - \tau \leq 0, (= 0 \text{ if } \Lambda > 0) \quad (34)$$

Then $\{C, K, n_y, n_h, k_y, k_h, h', b_r, b_h, p, \Lambda, a, l\}$ is a steady state allocation.²¹

In the steady state the variables taken as given in the city planner's problem solve the side conditions given by (30)–(34). Equation (30) expresses aggregate output in terms of intermediate good production in each city. This equation is the theoretical counterpart to equation (2) used to estimate city-specific TFP. The marginal utility of consumption is given by equation (31). Equation (32) states that total in-migration equals total out-migration. Equation (33) defines the rental rate for equipment. The last side condition (34) is equivalent to (10) in the static model and similarly determines steady state undirected search.

²¹We prove this result in the technical appendix where we also outline how we solve the model. We take a traditional dynamic programming approach to solving the city planner's problem. This is complicated substantially by the fact that there are four state variables in the city planner's problem, two of them endogenous. Furthermore the TFP process has a large domain. We overcome the computational challenges of a large dimensional and high variance state space in two main ways. First we exploit a parsimonious spline method to approximate the planner's value function and one-period return function. Second we take advantage of the large number of processors contained in graphics cards.

The function $\xi(z)$ in (34) can be shown to satisfy

$$\begin{aligned}\xi(z) = & C\phi[n_y(z) + n_h(z)]^\pi (\pi - 1)p(z)^{-\pi} + CH \ln \left(\frac{h(z)^\varsigma b_r(z)^{1-\varsigma}}{p(z)} \right) - CH \\ & + \beta \int \left(CA \left[\frac{a(z')}{p(z)} \right]^2 + C\psi_2 \left[\frac{l(z')}{p(z)} \right]^2 + \Lambda [\xi(z') - \eta] \right) dQ(s'; s, s_{-1}) \\ & + \beta \int \xi(z') dQ(s'; s, s_{-1}).\end{aligned}\quad (35)$$

The value of bringing an additional individual to a city is the expected discounted value of four terms: the benefits of obtaining a better selection of worker disutilities given the same amount of total employment $n_y + n_h$; the benefits of the local housing services that the additional person will enjoy; the costs of reducing the amount of housing services that everybody else in the city will enjoy when an additional person is brought in; and the expected discounted value of starting the following period with an additional person. This last term includes the benefits of having an additional person producing guided trips to the city, the benefits of obtaining a better selection of location-taste shocks (given the same number of individuals leaving the city), and the benefits of attracting additional people to the city through the undirected search technology.

When there are no migration frictions, $A = \psi_1 = \psi_2 = 0$, equation (29) implies that the marginal value of bringing an additional individual to a city is equated across cities as in the static case, $\xi(z) = \eta, \forall z$. However, unlike the static case this does not imply that wages are equated across cities. Instead, equation (35) says that the marginal savings in worker disutility plus the marginal impact on the utility of housing services is equated. When in addition to $A = \psi_1 = \psi_2 = 0$ housing structures are made perfectly mobile across cities, the same condition is obtained because land remains immobile. However, when land is also made mobile, then the marginal savings in work disutility and the marginal utility of housing services are each equated across cities.

6 Calibration

We now calibrate the steady state competitive equilibrium to U.S. data.²² Our calibration has two important characteristics. First, the city-specific TFP process is chosen to match our estimates presented in Section 2.2 thereby pinning down the model's exogenous source

²²Except where noted the aggregate data used to calibrate our model is obtained from Haver Analytics.

of persistence and volatility. Second, the calibration targets for the remaining parameters involve features of the data that are not primary to our study. So, for instance, we do not choose parameters to fit our estimated response of population to a TFP shock. The model's response of population to a TFP shock is the consequence of the estimated TFP process and the remaining parameters that are chosen to fit other features of the data.

In addition to specifying the stochastic process for TFP we need to find values for 16 parameters:

$$\theta, \gamma, \alpha, \vartheta, \delta_k, \delta_h, \beta, H, \zeta, \pi, \phi, \psi_1, \psi_2, A, \tau, \chi.$$

These include the factor shares in production of intermediate goods and residential structures, depreciation rates for equipment and residential structures, the discount factor, the coefficient on housing services in agents' preferences and land's share in housing services, the parameters governing labor disutility and migration costs, and intermediate goods' substitutability in final goods production. We calibrate these parameters conditional on a given quantity of undirected search Λ determined by τ . For larger values of τ undirected search is relatively small so that $a(z) > 0, \forall z$. In these cases the behavior of the model is invariant to the specific value of τ . For smaller values of τ undirected search is large and $a(z) = 0$ for some z . In these cases the behavior of the model is affected. It turns out that even for seemingly large steady state Λ corner solutions for $a(z)$ are either non-existent or extremely rare. We set our baseline so that undirected search is 3.8% of the population, roughly 70% of all moves.²³

The baseline calibration for the assumed value of τ is summarized in Table 1. There we indicate for each parameter the proximate calibration target, the actual value for the target we obtain in the baseline calibration, and the resulting parameter value.²⁴ Measuring GDP plays a key role in our calibration and so we begin by discussing this.

In the model GDP is measured as

$$GDP = Y + I, \tag{36}$$

where Y is output of non-construction final goods and I is residential investment. Residential investment is measured as the value in contemporaneous consumption units of the total

²³The specific value is $\tau = 1$. For this value the baseline calibration has 0.3% of city-year observations involving zero arrivals.

²⁴As in most similar studies there is not a one-to-one mapping between targets and parameters. See the appendix for details on the underlying source data for calculation of the empirical values of the targets.

Table 1: Baseline Calibration

Parameter	Parameter Description	Calibration Target	Target Value	Actual Value	Parameter Value
θ	Labor's share in intermediate goods	$\int w [n_y + n_h] d\mu / GDP$	0.64	0.64	0.66
α	Labor's share in construction	$\int n_h d\mu / \int [n_y + n_h] d\mu$	0.042	0.042	0.41
β	Discount factor	Real interest rate	0.04	0.04	0.9615
γ	Intermediate goods' equipment share	K_y / GDP	1.63	1.63	0.235
δ_k	Depreciation rate of equipment	$\delta_k K / GDP$	0.16	0.16	0.104
ϑ	Equipment's share in construction	K_h / GDP	0.022	0.022	0.05
δ_h	Depreciation rate of structures	I / GDP	0.064	0.064	0.045
ζ	Land's share in housing services	$\int q^b b_r d\mu / [\int q^b h d\mu + \int q^b b_r d\mu]$	0.37	0.36	0.215
H	Housing coefficient in preferences	$\int q^b h d\mu / GDP$	1.55	1.50	0.205
ϕ	Labor disutility	$\int [n_y + n_h] d\mu / \int p d\mu$	0.63	0.63	1.61
π	Labor supply elasticity	$\partial \ln [n_y + n_h / p] / \partial \ln w$	0.24	0.25	5.0
ψ_1	Taste shock slope	Mean arrivals	5.5	5.5	6.07
ψ_2	Taste shock curvature	Slope of arrivals versus net	0.55	0.55	43.9
A	Guided trip cost	Average moving costs/average wages	-1.9	-1.9	35.6
χ	Intermediate goods' complementarity	Zipf's law for population	-1.0	-1.3	0.9
g	Drift in technology	Zipf's law for TFP	-3.5	-3.4	-0.0017
ρ	TFP lag coefficient	Serial corr. of TFP growth	0.24	0.24	0.35
σ	TFP innovation std. err.	TFP growth innovation std. err.	0.015	0.015	0.019

Note: The calibration is based on a particular value for the undirected search parameter τ . See the text for details. The actual values for the serial correlation and innovation of TFP growth are based on simulations. The underlying parameter values for the TFP process are somewhat different due because of the discrete approximation we use.

additions to local housing in a year. Specifically,

$$I = \int \left[\beta \int q^h(z') dQ(s'; s, s_{-1}) \right] n_h(z)^\alpha k_h(z)^\vartheta b_h(z)^{1-\alpha-\vartheta} d\mu$$

where q^h denotes the price of residential structures. This price is obtained as the solution to the following no arbitrage condition

$$q^h(z) = r_h(z) + (1 - \delta_h) \beta \int q^h(z') dQ(s'; s, s_{-1})$$

where r_h denotes the rental price of residential structures which equals the marginal product of structures in the provision of housing services. The National Income and Product Accounts (NIPA) measure of private residential investment is the empirical counterpart to I . Our empirical measure of Y is the sum of personal consumption expenditures less housing services, non-residential fixed investment and private business inventory investment. Because our model does not include government expenditures and net exports we exclude these from our empirical concept of GDP.

Our measure of model GDP excludes the value of guided trip services produced in the model, which is a questionable assumption. The presence of guided trips also has implications for how we measure wages because workers produce guided trips and in principle they should be compensated for this. We now address how these issues affect our calibration. Using the decentralization discussed in Section 3.5, we calculate the total value of guided trips in our baseline calibration to be 1.8% of model GDP as defined above. Recall that we interpret guided trips as encompassing many market and non-market activities. Some of these activities appear in the national accounts as business services and therefore count as intermediate inputs that do not end up directly in measured GDP. Others do not appear anywhere in the national accounts because they are essentially home production or are impossible to measure. However, given its small size including the total value of guided trips in our model-based measures of wages and GDP does not change our baseline calibration.

Measuring employment also is complicated by the fact that all household members participate in generating guided trips. We count those agents engaged in intermediate good production, n_y , and residential construction, n_h , as employed and measure their wages by their marginal products excluding the value of guided trips. The non-employed who also produce guided trips are assumed to be engaged in home production and so are not included in our accounting of employment. In Table 1 the labor share parameters are chosen to match total labor compensation as a share of GDP (the target is borrowed from traditional real

business cycle studies) and our estimate of the share of construction employment in total private non farm employment.

We fix the discount rate so the model's real interest rate is 4%. Combined with this target the equipment-output ratio in the non-construction sector, K_y/Y , identifies equipments's share in that sector's production. Our empirical measure of equipment for this calculation is the Bureau of Economic Analysis' (BEA) measure of the stock of non-residential fixed capital. Equipment's depreciation rate is identified using the investment to GDP ratio, where we measure investment using the NIPA estimate of non-residential fixed investment. Equipment's share in residential construction is identified by the ratio of capital employed in the residential construction sector, K_h , to GDP where the empirical counterpart to capital in this ratio is the BEA measure of non-residential fixed capital employed in residential construction. The depreciation rate of residential structures is identified using the residential investment to GDP ratio.

We identify the housing service parameters as follows. First the housing coefficient H is chosen to match the residential capital to GDP ratio, where the measurement in the data is as described above. Land's share in housing services, ζ , is chosen to match the estimate of land's share of the total value of housing in [Davis and Heathcote \(2007\)](#). To measure this object in the model we need the price of land, q_b . We obtain this variable as the solution to the arbitrage condition

$$q^b(z) = r_b(z) + \beta \int q^b(z') dQ(s'; s, s_{-1}),$$

where r_b denotes the rental price of land which equals the marginal product of land in the provision of housing services. Land's share of the economy-wide value of housing is then given by $\int q^b b_r d\mu / [\int q^h h d\mu + \int q^b b_r d\mu]$.

The labor disutility parameters are based on statistics involving employment to population ratios. The multiplicative parameter ϕ is identified using the ratio of aggregate civilian employment to population obtained from the Census Bureau. The representative household equates the disutility of putting an additional household member to work in a city with that city's wage. Therefore we find π by equating the change in the log employment to population ratio divided by the change in the log wage in the period of a TFP shock estimated with our data to the value of this object estimated in the same way with data simulated from the model. We use the methods described in [Section 2.2](#) to estimate the dynamic responses of the logs of employment, population and wages in a city to a local TFP shock we need for these calculations.

The migration parameters are chosen to match Figure 2 and average costs of moving as a fraction of average wages estimated by Kennan and Walker (2011). For this we take Λ as given since it is determined by our pre-selected value for τ . To reproduce Figure 2 we require

$$\begin{aligned}\frac{\psi_1}{A + \psi_2} + \frac{A}{A + \psi_2} \Lambda &= 5.5; \\ \frac{\psi_2}{A + \psi_2} &= 0.55.\end{aligned}$$

These conditions set the constant and slope coefficients in equation (14) to their empirical counterparts displayed in Figure 2. Kennan and Walker (2011) estimate the average costs of migration for those who move to be about -1.9 times average wages.²⁵ We measure the average cost of actual moves in our model to be as close as possible to Kennan and Walker (2011)'s concept of moving costs. Specifically, average costs include the total value of guided trips, the total consumption value of the location-taste shocks of those who move, and the difference in the wages and the consumption value of the housing services in the cities migrated from and to:

$$\begin{aligned}& \frac{C \int \left(-\psi_1 \frac{l(z)}{x} + \psi_2 \left(\frac{l(z)}{x} \right)^2 \right) x d\mu}{\int l(z) d\mu} + \frac{\int q(z) a(z) d\mu + C\tau\Lambda}{\int a(z) d\mu + \Lambda} \\ & + \frac{\int \left[w(z) + CA \ln \left(\frac{h(z)^\epsilon b_r(z)^{1-\epsilon}}{p(z)} \right) \right] l(z) d\mu}{\int l(z) d\mu} \\ & - \frac{\int \left[w(z) + CA \ln \left(\frac{h(z)^\epsilon b_r(z)^{1-\epsilon}}{p} \right) \right] (a(z) + \Lambda x(z)) d\mu}{\int a(z) d\mu + \Lambda},\end{aligned}$$

where wages in a type- z city, $w(z)$, equal the marginal product of labor in intermediate goods production in that type of city. Average wages are measured as

$$\frac{\int w(z) [n_y(z) + n_h(z)] d\mu}{\int [n_y(z) + n_h(z)] d\mu}.$$

The last part of the baseline calibration is to assign a value to χ , the parameter that determines intermediate goods' substitutability, and to specify the stochastic process driving TFP fluctuations. These choices are interconnected. The TFP process is estimated using the procedure described in Section 2.2. When we apply this methodology for plausible values of

²⁵The value -1.9 equals the ratio -\$80,768/\$42,850. The numerator is the wage income of the median AFQT scorer aged 30 in 1989 reported in Table III and the denominator is the entry in the row and columns titled 'Total' in Table V.

θ , γ and χ the growth rate of technology is well-represented as a first order auto-regression. This suggests considering the following process for city-specific TFP:

$$\ln s_{t+1} - \ln s_t = g + \rho (\ln s_t - \ln s_{t-1}) + \varepsilon_{t+1}, \quad (37)$$

where $\varepsilon_{t+1} \sim N(0, \sigma^2)$. This process is non-stationary in levels and therefore is inconsistent with a steady state. We address this problem by adopting a reflecting barrier process for TFP. Specifically:

$$\ln s_{t+1} = \max \{g + (1 + \rho) \ln s_t - \rho \ln s_{t-1} + \varepsilon_{t+1}, \ln s_{\min}\}.$$

According to this process s_t is reflected at the barrier $\ln s_{\min}$ (which we normalize to zero).²⁶

The case $\rho = 0$ has been studied thoroughly in the context of cities by [Gabaix \(1999\)](#). In this case $g < 0$ generates a stationary process in levels where the invariant distribution has an exponential tail given by

$$\Pr [s_t > b] = \frac{d}{b^\omega}$$

for scalars d and b . A striking characteristic of cities is that when s measures a city's population one typically finds that $\omega \simeq 1$. Equivalently a regression of log rank on log level of population yields a coefficient close to -1. This property is called Zipf's law. For convenience we refer to ω as the Zipf coefficient.


The case $\rho > 0$, which applies under our estimates for serial correlation in TFP growth, has not been studied. Our simulations suggest this process behaves similarly to the $\rho = 0$ case in that the invariant distribution also has an exponential tail. We verify below that a version of Zipf's law holds for TFP as we measure it in [Section 2.2](#) and so using the reflecting barrier process seems appropriate.

We choose the drift parameter g and the substitution parameter χ to match the Zipf coefficients for TFP, s_t , and population, p_t . Since our measurement of TFP depends on χ (as well as θ and γ) we use the following procedure. For given χ we estimate (37) and find the value of g which reproduces the Zipf coefficient for TFP in the data. Using the resulting process in our model we generate a Zipf coefficient for population. We then repeat this procedure for different values of χ and choose the value that implies a Zipf coefficient for population that is as close to its empirical value of 1 as possible. We arrive at $\chi = 0.9$ and a Zipf coefficient for population equal to 1.3. The corresponding values of g , ρ and σ are in [Table 1](#).

²⁶[Coen-Pirani \(2010\)](#) considers a stationary AR(2) process for TFP, calibrating it to match serial correlation in net worker flows.

To demonstrate how well our model does at replicating the dual Zipf's laws for population and TFP, Figure 6 displays plots of log rank versus log level for population and TFP from the data and our calibrated model.²⁷ Notice how in the data the Zipf coefficient is larger for TFP than population. This arises naturally in the model because population tends to be allocated away from lower toward higher TFP cities. ? finds a similar relationship between employment and TFP in an equilibrium model of firm size.

Figure 6: Zipf's Laws for Population and TFP



7 Quantitative Analysis

We now consider the model's empirical predictions. First we examine how well the model is able to reproduce our finding that cities' populations are slow to adjust to TFP shocks. We confirm that the model's success along this dimension does not come at the expense of strongly counterfactual predictions for gross migration or the behavior of the labor and housing markets. This analysis leads us to conclude that despite choosing parameters to match

²⁷The scales for the data plot differ from the model plot. This is because we use the cumulative distribution functions to measure rank in the model and we restrict the domain of TFP and hence population because it is extremely costly to solve the model over a grid that is wide enough to encompass the empirical distributions of population and TFP.

evidence not directly related to the dynamics of interest our model nonetheless generally excels in replicating them.

Next we investigate how migration and housing influence slow population adjustments. We find that costs of directed search through the model’s guided trip technology is the principle source of slow population adjustments. We interpret this finding as demonstrating that the myriad ways individuals get informed about desirable locations to live and work represent significant barriers to rapid labor reallocation. The fact that we identify the model’s migration parameters without consideration of within-city dynamic responses to TFP shocks lends substantial credibility to this interpretation. Interestingly, housing plays only a small role slowing labor reallocation. This is despite having several characteristics that make it a natural candidate for slowing population adjustments.

Finally, we study the implications of our model’s successful accounting of slow population adjustments for urban decline. Specifically, we examine the decline of a city receiving a long lasting reduction in TFP. We find that our estimates of migration costs translate to very slow adjustment with a city taking multiple decades to converge to its new long run population. This finding suggests that costly migration is a major factor determining the surprising persistence of urban decline.

7.1 Model Validation

A traditional way to validate a model like our’s is to compare unconditional model-generated statistics with those we have estimated. While this is a worthwhile endeavor (which we pursue below), it is an unreasonably strong test of the model. We are confident cities are subject to more than just TFP shocks such as to amenities, taxes and demand for the city’s output. These shocks affect variation of the variables of interest and will influence estimates of unconditional moments.

Instead of studying unconditional dynamics, we focus on conditional dynamics to validate our model. In particular the dynamic response of variables to TFP shocks. To do this we estimate the dynamic responses of population, arrival and leaving rates, employment, wages, home construction and house prices using the procedure described in Section 2.2) and compare these responses to ones estimated using the same procedure with model-simulated data.

Figure 7 displays model and estimated responses of population to a one standard devia-

Figure 7: Responses of Population





Figure 8: Responses of Arrival and Leaving Rates



tion innovation to TFP along with plus and minus 2 standard error bands for the estimates.²⁸ The model estimates are statistically and economically close to the ones from the data. That is, using migration frictions calibrated using microeconomic evidence on gross migration and estimates of average migration costs, the model closely replicates the empirical dynamic response of population to a TFP shock. In other words our model accounts for population's slow response to TFP shocks. Figure 8 shows accounting for slow population adjustments involves also replicating quite closely the dynamic responses of the arrival and leaving rates. The goodness of fit is not as good as with population, for example both responses are more persistent than in the data and the arrival rate response is a little too strong. Nevertheless the model does surprisingly well.

Figure 9: Responses of Labor and Housing Markets




Figure 9 shows the dynamic responses of employment, wages, residential investment and house prices. We define house prices as the total value of structures and land used to produce housing services per unit of housing services provided. This corresponds to a price of housing per square foot, q^{sf} under the assumption that every square foot of built housing yields the

²⁸These standard errors do not take into account the sampling uncertainty in our estimates of the underlying TFP process.

same quantity of housing services:

$$q^{sf}(z) = \frac{q^h(z) h(z) + q^b(z)}{h(z)^{1-\varsigma} b_r(z)^\varsigma}.$$

The labor market responses are a good fit, but the model is less successful accounting for housing. The model’s residential investment response misses the hump shape and the house price response is too fast. One explanation for the discrepancy between the model and data of the housing variables is that our model does not include any search frictions in the housing market. [Lloyd-Ellis et al. \(2011\)](#) demonstrate that search frictions show promise in generating serially correlated responses of construction and house price growth to productivity shocks.

7.2 Unconditional Statistics

Table 2 displays unconditional standard deviations and contemporaneous correlations of the variables just discussed in the model and in our data. The patterns for the model can be inferred from the previously discussed figures since TFP shocks are the only source of variation. Except for population, the standard deviations are expressed relative to the standard deviation of population. The correlations are all with population. The statistics are based on the levels of the gross migration rates and on the growth rates of the other variables. The variables have been transformed as described in Section 2 prior to the analysis.

The first thing to notice is that TFP shocks generate about two thirds of the overall variation in population. As emphasized previously, we do not expect TFP shocks to explain all the variation because there are other shocks to cities. In terms of relative volatilities the model generates variation similar to the data for the gross migration rates and the labor market variables although wages are a little too volatile. The model is consistent with residential construction being the most volatile variable, but it fluctuates much less in the model than in the data. This mirrors the inability of the model to reproduce the amplitude of the response to TFP shocks alone. The relative volatility of house prices in the model is high too, but not quite as high in the data, again similar to the conditional correlation. The model is qualitatively consistent with all the correlations with population growth. The largest discrepancies with the data involve the arrival and leaving rates being perfectly positively and negatively correlated with population growth. Perhaps [Coen-Pirani \(2010\)](#)’s mechanism inducing a positive correlation between the gross migration rates, absent from our model, could overcome this deficiency.

Table 3 addresses within-city serial correlation of the variables in Table 2. Population,

Table 2: Volatility and Co-movement Within Cities

Variable	Standard Deviation		Correlations	
	Data	Model	Data	Model
Population	1.33	0.87	–	–
Arrival Rate	0.65	0.53	0.59	1.00
Leaving Rate	0.58	0.48	-0.42	-1.00
Employment	1.58	1.23	0.56	0.93
Wages	1.23	1.81	0.16	0.32
Construction	19.7	4.27	0.14	0.40
House Prices	3.76	2.32	0.29	0.47

Note: The statistics are based levels of the gross migration rates and on the growth rates of the other variables. The latter variables have been transformed as described in Section 2.2 prior to calculating growth rates. Standard deviations of all variables except population are expressed relative to the standard deviation for population. Correlations are with population.

gross migration, employment and wages all display similar persistence to that in the data, although the model’s variables are more persistent. Construction in the model and data are similarly random-walk like, although this feature of the unconditional moments clearly is due to the affects of other shocks given the TFP responses. House prices display the greatest differences. In the data house price growth displays substantial serial correlation while in the model house prices are more like random-walks. [Lloyd-Ellis et al. \(2011\)](#) demonstrate how within-city search frictions which are not in our model can generate serial correlation in house price growth.

7.3 The Source of Slow Population Adjustments

We now address the sources of slow population adjustments in our model. Figure 10 displays impulse responses to TFP shocks implied by several different versions of the model for this purpose. The different versions consist of perturbations relative to the baseline, calibrated version of the model, holding parameters not involved in the perturbation fixed at their baseline values. The “Free Guided Trips” case sets $A = 0$. This case has the same implications as assuming all the migration parameters are set to zero, because when guided trips are free the city-planner sets the leaving rate in each city to the constant value that minimizes leaving costs and adjusts population by changing the arrival rate at zero cost. “No

Table 3: Serial Correlation Within Cities

Variable	Lag			
	1	2	3	4
Population				
Data	0.81	0.74	0.67	0.63
Model	0.93	0.87	0.81	0.75
Arrival Rate				
Data	0.82	0.70	0.58	0.47
Model	0.93	0.87	0.81	0.75
Leaving Rate				
Data	0.77	0.74	0.67	0.60
Model	0.93	0.87	0.81	0.75
Employment				
Data	0.52	0.29	0.21	0.15
Model	0.73	0.63	0.58	0.54
Wages				
Data	0.15	0.04	0.05	0.07
Model	0.20	0.02	-0.02	-0.02
Construction				
Data	0.12	0.10	-0.02	-0.11
Model	-0.09	0.02	0.04	0.04
House Prices				
Data	0.73	0.31	-0.06	-0.25
Model	0.07	0.06	0.05	0.05


Note: The variables are have been transformed as described in Section 2.2 prior to calculating the statistics. The gross migration rates are levels and all other variables are growth rates.

Location-Taste Shocks” is the case where $\psi_1 = \psi_2 = 0$ so that costly guided trips are the only migration friction. Under “Mobile Housing” housing can be re-allocated across cities in the same way as equipment. In this case a city’s dynamics are not influenced by the durability or the size of the local housing stock nor the city’s ability to produce houses to accommodate new workers because housing can be moved without cost from less desirable cities. “Full Flexibility” combines all the perturbations, the version of the model without mobility costs and with mobile housing. The left plot in Figure 10 displays the levels of the responses and the right one shows the responses after first dividing them by the value attained in the last period of the response to more clearly show the speed of adjustment.

Figure 10 shows that under Full Flexibility the population dynamics essentially follow the

path of TFP with roughly 90% of the long run adjustment occurring after 2 years compared to 85% for TFP (see Figure xx) – absent frictions the model has essentially no internal mechanism to propagate TFP shocks. The No Location-Taste Shocks and Mobile Housing cases are very close to the baseline. In other words removing from the model costly out-migration or immobile housing, leaving costly guided trips as the only model friction, leaves the population response essentially as slow as it is in the baseline economy. Making guided trips free moves the response closer to the full-flexibility case, but does not take the model all the way there. Recall that making guided trips free leads to the same model responses as when migration is completely costless. Therefore in the Free Guided Trip case the only friction is that housing is immobile. The discrepancy with Full Flexibility arises from a property of adjustment costs highlighted by [Abel and Eberly \(1994\)](#). The first adjustment cost introduced to an otherwise frictionless model always has a relatively large impact on dynamics. In other words, introducing immobile housing into an otherwise frictionless model has large effects, although this is not sufficient to deliver the amplitude and persistence of the population response in the data. But, the dynamics of population with migration costs but mobile housing, the Mobile Housing case, are essentially the same as the baseline. We conclude that the prime driver of slow population adjustment is the costly guided trip technology.

Figure 10: Impact of Model Features on Population Adjustment



This finding confirms results in [Kennan and Walker \(2011\)](#) who arrive at their findings with a very different methodology. Using individual-level data from the National Longitudinal Survey of Youth and a very different approach to identification that does not account for equilibrium interactions or housing they find that migration costs are a significant source of slow population adjustments at the state level. Our findings are not independent given that we use their migration cost estimates to calibrate the migration parameters in our model. Nevertheless they confirm that [Kennan and Walker \(2011\)](#)’s findings are robust to the presence of housing and equilibrium interactions.

7.4 Migration and Urban Decline


This section describes how the slow population adjustments in the calibrated model translate to the persistence of urban decline in cities like Detroit. We conduct the following experiment for this. First, we simulate an individual city for a long time at the highest TFP level so that the housing stock and population are near the steady state of a city with a permanently high TFP level. Next we suppose there is a one time drop of TFP which implies a 50% drop in the city’s population if the TFP level were to remain lower forever. We then plot the dynamics of population over time. These dynamics are displayed in Figure 11 show that it takes roughly 50 years for the city to reach its long run population after the TFP shock.

Since the dominant source of slow population adjustment in the model is the cost of attracting workers to a city, these costs drive the persistence of urban decline in this experiment. Housing is not very important at all, which contrasts with [Glaeser and Gyourko \(2005\)](#) who argue that durable and immobile housing underly persistent urban decline. These authors do not consider the costs of attracting workers to a city in their analysis. We consider both, but housing turns out to be relatively unimportant.

8 Conclusion

This paper documents that population adjusts slowly to near random-walk TFP shocks and proposes an explanation for why. The explanation is that the incentive to reallocate population after a TFP shock is limited by the costs of attracting workers to desirable cities, that is adjustment costs to increasing population through in-migration are the dominant source of slow population adjustments. Our model of migration that delivers this result is not arbitrary, but is dictated by the nature of the relationship between gross and net

Figure 11: Persistent Urban Decline



population flows in cities that we uncover in our panel of 365 cities from 1985 to 2007.

Our model has left out other interesting model features that are undoubtedly important for understanding the full range of adjustments to shocks within cities. Chief among these omissions are search frictions in local labor and housing markets. We think it would be interesting to add these features to our framework. Doing so would help disentangle the contributions to labor reallocation of traditional search frictions from the migration frictions we have introduced in this paper.

Taken together our findings point to a heretofore ignored mechanism in the determination of macroeconomic adjustment. While we have not done this experiment, our results strongly suggest that a mis-allocated housing stock due to overbuilding in the recent housing boom would have little impact on macroeconomic labor adjustment and the sluggish economy since the housing bust is unlikely to have been driven by such a mis-allocation. This conclusion of course derives from a model without any frictions in the financing of housing. If housing is to be important for macroeconomic labor adjustment it must be through these or other frictions.

References

- Abel, A. B. and J. C. Eberly (1994). A unified model of investment under uncertainty. *American Economic Review* 84(1), 1369–1384.
- Albouy, D. (2009). The unequal geographic burden of federal taxation. *Journal of Political Economy* 117, 635–667.
- Alvarez, F. and R. Shimer (2011). Search and rest unemployment. *Econometrica* 79(1), 75–122.
- Alvarez, F. and M. Veracierto (2012). Fixed-term employment contracts in an equilibrium search model. *Journal of Economic Theory*.
- Blanchard, O. J. and L. F. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity* 1992(1), 1–75.
- Ciccone, A. and R. E. Hall (1996). Productivity and the density of economic activity. *American Economic Review* 86(1), 54–70.
- Coen-Pirani, D. (2010). Understanding gross worker flows across us states. *Journal of Monetary Economics* 57, 769–784.
- Davis, M., J. Fisher, and M. Veracierto (2011). The role of housing in labor reallocation. Revision to Chicago Fed working paper no. 2010-18.
- Davis, M., J. Fisher, and T. Whited (2013). Macroeconomic implications of agglomeration. Forthcoming, *Econometrica*.
- Davis, M. and J. Heathcote (2007). The price and quantity of residential land in the United States. *Journal of Monetary Economics* 54(8), 2595–2620.
- Davis, M. and F. Ortalo-Mangé (2011). Household expenditures, wages and rents. *Review of Economic Dynamics* 14(2), 248–261.
- Davis, S., J. Faberman, and J. Haltiwanger (2006). The flow approach to labor markets: New data sources and micro-macro links. *Journal of Economic Perspectives* 20(3).
- Davis, S., J. Faberman, and J. Haltiwanger (2011). Labor market flows in the cross section and over time. *Journal of Monetary Economics* 59, 1–18.

- Davis, S., J. Haltiwanger, and S. Schuh (1998). *Job Creation and Destruction*. Cambridge, MA: MIT Press.
- Diamond, R. (2012). The determinants and welfare implications of us workers' diverging location choices by skill: 1980–2000. Harvard University working paper.
- Eeckhout, J. (2004). Gibrat's law for (all) cities. *American Economic Review* 94(5), 1429–1451.
- Ferreira, F., J. Gyourko, and J. Tracy (2011). Housing busts and household mobility: An update. NBER Working Paper No. 17405.
- Gabaix, X. (1999). Zipf's law for cities: An explanation. *Quarterly Journal of Economics* 114(3), 739–767.
- Glaeser, E. L. and J. Gyourko (2005). Urban decline and durable housing. *Journal of Political Economy* 113(2), 345–375.
- Greenwood, M. (1997). Internal migration in developed countries. In M. Rosenzweig and O. Stark (Eds.), *Handbook of Population and Family Economics*, Volume 1B. North Holland.
- Kaplan, G. and S. Schulhofer-Wohl (2012). Understanding the long-run decline in interstate migration.
- Karahan, F. and S. Rhee (2012). Geographical reallocation and unemployment during the great recession: The role of the housing bust. University of Pennsylvania working paper.
- Kennan, J. and J. R. Walker (2011). The effect of expected income on individual migration decisions. *Econometrica* 79(1), 211–251.
- Lloyd-Ellis, H. and A. Head (2012). Housing liquidity, mobility, and the labour market. *Review of Economic Studies*.
- Lloyd-Ellis, H., A. Head, and H. Sun (2011). Search and the dynamics of house prices and construction. Queen's University working paper.
- Lucas, R. (1997). Internal migration in developing countries. In M. Rosenzweig and O. Stark (Eds.), *Handbook of Population and Family Economics*, Volume 1B. North Holland.

- Lucas, R. and E. Prescott (1974). Equilibrium search and unemployment. *Journal of Economic Theory* 7, 188–209.
- Modestino, A. and J. Dennett (2012). Are american homeowners locked into their houses? the impact of housing market conditions on state-to-state migration. Federal Reserve Bank of Boston Working Paper No. 12-1.
- Molloy, R., C. L. Smith, and A. Wozniak (2011). Internal migration in the united states. *Journal of Economic Perspectives* 25(2), 1–42.
- Nenov, P. (2012). Regional mismatch and labor reallocation in an equilibrium model of migration. Norwegian Business School working paper.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy* 90(6), 1257–1278.
- Rosen, S. (1979). Wagebased indexes of urban quality of life. In P. N. Miezkowski and M. R. Straszheim (Eds.), *Current Issues in Urban Economics*, pp. 74–104. Baltimore, MD: Johns Hopkins University Press.
- Schulhofer-Wohl, S. (2012). Negative equity does not reduce homeowners’ mobility. NBER Working Paper No. 16701.
- Van Nieuwerburgh, S. and P.-O. Weil (2010). Why has house price dispersion gone up? *Review of Economic Studies* 77, 1567–1606.
- Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.