
What Do We Learn from Cross-Regional Empirical Estimates in Macroeconomics?

Adam Guren, *Boston University, United States of America, and NBER, United States of America*

Alisdair McKay, *Federal Reserve Board of Minneapolis, United States of America*

Emi Nakamura, *University of California, Berkeley, United States of America, and NBER, United States of America*

Jón Steinsson, *University of California, Berkeley, United States of America, and NBER, United States of America*

I. Introduction

A growing literature uses variation across cities or regions to identify the effects of economic shocks of interest to macroeconomists.¹ What exactly these estimates identify is often complicated by the fact that metropolitan and regional outcomes reflect both the partial equilibrium effects of the shock in question as well as local general equilibrium responses to the shock. In this paper we propose an approach by which applied researchers can isolate the partial equilibrium effect of the shock. The partial equilibrium effect is useful for several reasons. First, it has a clear theoretical interpretation and speaks more directly to specific economic mechanisms. In contrast, estimates that include local general equilibrium effects reflect a combination of several economic mechanisms and are more difficult to interpret as a result. Second, the partial equilibrium effect can more easily be matched with a theoretical counterpart for calibration purposes. The method we propose allows researchers to avoid formulating and solving a multi-region general equilibrium macroeconomic model to be able to compare their empirical results to analogous concepts in a model.

The easiest way to describe our approach is in the context of a concrete application. The application we focus on is the analysis of so-called housing wealth effects. The US housing boom and bust in the 2000s focused attention of economists on the effect of changes in home prices on consumer spending. Prominent recent papers in this literature use regional data to estimate the effect of changes in house prices on outcomes such as spending, car registrations, and employment (e.g., Mian, Rao, and Sufi 2013;

Mian and Sufi 2014). The appropriate interpretation of these estimates is not straightforward. House prices are endogenous at the level of a city and a shock that changes home prices surely alters consumption through other channels. The shock may affect consumption directly. In addition, the increased spending triggered by higher house prices will raise wages and incomes locally, which will lead to more local spending (a local general equilibrium effect). For these reasons, it is not immediately clear what we can learn from the response of city-level consumption to a change in local house prices.

We show how existing empirical estimates of the housing wealth effect on consumption can be decomposed into the partial equilibrium effect of house prices on consumption and local general equilibrium effects. We start by drawing a distinction between prices that are determined nationally and prices that are determined locally. For example, financial markets are highly integrated at a national level, whereas labor markets and markets for nontradable goods are quite local. This distinction is important because variation in national prices will be absorbed by the constant in a cross-sectional regression or the time fixed effect in a panel regression and will therefore not affect cross-regional estimates of the housing wealth effect.

Local general equilibrium effects that operate through local markets will, however, not be captured by the constant or time fixed effects in cross-regional regression analysis. The key insight in our paper is that estimates of the local fiscal multiplier can be used to gauge the strength of these local general equilibrium effects. In particular, we derive conditions under which one can remove local general equilibrium effects from a city-level estimate of the housing wealth effect simply by dividing that estimate by an estimate of the local fiscal multiplier. The logic underlying this result is that the equilibrium response to an increase in local demand will be the same whether that demand comes from private consumption as a result of the housing wealth effect or from a fiscal shock. Dividing by the local fiscal multiplier yields an estimate of the partial equilibrium housing wealth effect that corresponds to the effect of a change in home prices holding fixed wages and other nonhousing prices. This partial equilibrium effect has a simple interpretation and it can be used to discipline a partial equilibrium model of housing and consumption.

In recent complementary work, we estimate the housing wealth effect based on city-level variation in house prices and using retail employment as a proxy for local consumption (Guren et al. 2020). We estimate an elasticity of retail employment with respect to house prices of 0.072. We

furthermore show in that paper that retail employment has approximately a unit elasticity with respect to consumption in the aggregate and across cities, which allows us to interpret the retail employment response as a consumption response. To convert the elasticity we estimate into a marginal propensity to consume out of housing wealth (MPCH), we divide by the housing consumption ratio, which averaged 2.17 from 1985 to 2016. This yields an MPCH of 3.3 cents on the dollar. For the reasons discussed above, this estimate reflects both the partial equilibrium and local general equilibrium responses of consumption to housing. To isolate the partial equilibrium effect, we divide it by an estimate of the local fiscal multiplier. Nakamura and Steinsson (2014) estimate the local fiscal multiplier to be approximately 1.5. Dividing our housing wealth effect estimate by the local fiscal multiplier yields a partial equilibrium MPCH of 2.2 cents on the dollar.

Our approach of combining several reduced-form estimates to identify a structural parameter is an application of the general method of simultaneous equations identification. The identification challenge is that multiple structural systems can give rise to the same reduced-form estimates. For example, the cross-regional housing wealth effect may be large because the partial equilibrium effect of house prices on consumption is large or because local general equilibrium effects are large. Identification requires making restrictions on the system. In our example, the restrictions are exclusion restrictions: for example, shocks to government spending do not directly affect consumption demand (they affect consumption through income). The key point in our argument is that allowing for more sources of exogenous variation (government spending shocks) helps identify the system even though it enlarges the system. We discuss how the identification problem and solution are closely related to structural vector autoregression (VAR) methods. A general lesson is that researchers who use regional variation can benefit from approaching their estimates as components of a system of simultaneous equations. In this context, shocks to government purchases are particularly useful in identifying local general equilibrium effects because the direct effect of the shock on demand is known. We discuss how this simultaneous equations approach is valuable in many situations where the object of interest is the direct effect of a change in demand whether it arises from housing wealth effects, credit supply shocks, foreign demand, or other sources.

We present several refinements of the basic idea of dividing by the local fiscal multiplier. The most important refinement from a quantitative perspective is to allow for the fact that an increase in house prices stimulates

local construction activity. This is a separate channel from the standard consumption multiplier because the initial partial equilibrium increase in demand is not consumption but residential investment. We derive a simple formula for isolating the partial equilibrium effect of house prices on consumption in the presence of effects on local construction activity. To evaluate this formula, we need an estimate of the response of construction activity to house prices. We present such an estimate using an analogous research design to that used in Guren et al. (2020). The refined formula yields an estimate of the partial equilibrium housing wealth elasticity of 0.040 or a partial equilibrium MPC_H of 1.8 cents on the dollar.

A second important refinement is to allow for dynamics. In a dynamic context, there is no single fiscal multiplier. Rather there is an entire impulse response of output to a fiscal shock. For the dynamic case, we develop a matrix version of the simple formula that applies in the static case. For example, the dynamic version of our formula involves the inverse of a matrix with the (i, j) element giving the effect of a fiscal shock at date j on output at date i (a matrix version of dividing by the local fiscal multiplier). Using simulations, we explore how well our static formula performs when the data are generated by a dynamic model. We show that under some conditions it holds almost exactly. In the richest dynamic structural model we consider, our simple static formula accounts for the bulk of the needed adjustment but somewhat underestimates it.

In our analysis, we address Davidoff's (2016) critique of the use of heterogeneous supply constraints—such as those captured by the Saiz (2010) housing supply elasticity estimates—as instruments for home prices. Davidoff points out that if housing markets experience a common demand shock but move along different supply curves, then prices and quantities should be negatively correlated: the most constrained cities should see the largest housing price responses but the smallest housing quantity responses. However, Davidoff shows that there is a positive relationship between housing price growth and the growth of housing units. Our construction employment estimates confirm Davidoff's critique applies at business cycle frequencies.²

We show that a model that allows for differences between short-run and long-run housing supply elasticities can address Davidoff's critique. We consider a version of our dynamic model in which short-run housing supply elasticities are low in all cities, but housing supply elasticities are more heterogeneous in the long run. This distinction reflects both the time it takes to plan and develop new housing units (which makes short-run elasticities low in all cities) and differences in constraints on land supply

that are not binding in the short run but may bind in the long run as in Nathanson and Zwick (2018). House prices, like other asset prices, are forward-looking in nature and, as a result, are primarily determined by the long-run elasticity of housing supply even in the short run. The short-run construction response, on the other hand, reflects the short-run constraints faced by housing developers. Consider two cities that have the same short-run supply elasticity but differ in the long-run supply elasticity. A common shock to expected future housing demand will move prices differently across the two cities due to expectations about future housing supply responses. Furthermore, expectations about larger future capital gains (or smaller future capital losses) will imply that the shock results in a larger change in current housing demand in the city with a less elastic long-run housing supply curve. In terms of construction, both cities move along the same short-run housing supply curve but by different amounts. The upward-sloping short-run supply curve yields a positive correlation between prices and quantities even though the changes in prices are generated primarily by a common demand shock moving the cities along different (long-run) supply curves.

The idea of dividing cross-regional estimates of housing wealth effects by estimates of the local fiscal multiplier to arrive at an estimate of the partial equilibrium effect of changes in house prices on consumption implicitly assumes that existing evidence on the local fiscal multiplier is stronger than for the partial equilibrium effect of house prices on consumption. However, if the reverse is true, our method is just as useful. In this case, we can use our method to infer the size of the local fiscal multiplier from the combination of cross-regional estimates of housing wealth effects and estimates of the partial equilibrium effect of house prices on consumption. More generally, in the context of housing wealth effects, our method implies that cross-regional estimates can provide information about some combination of the local fiscal multiplier and the strength of partial equilibrium effects.

The key idea we explore in this paper is that the general equilibrium adjustment to a change in private consumption is equivalent to the general equilibrium adjustment to a government spending shock. This demand equivalence idea is also explored in several contemporaneous papers. In the context of a two-period model of the stock market wealth effect, Chodorow-Reich, Nenov, and Simsek (2020) derive a demand equivalence result that links the direct spending response to the change in the local wage bill. Wolf (2019a) lays out conditions under which demand equivalence holds exactly for the impulse responses of a dynamic model

and Wolf (2019b) applies those results to cross-regional comparisons and local general equilibrium. In our analysis, the relationship between impulse responses is expressed in terms of the matrix relationship described above.³ Applying this result directly is challenging because it requires that the researcher observe the full dynamic response to the shock of interest and a fiscal spending shock that has the same dynamics as the (as yet unknown) partial equilibrium response of interest. However, we show that the simpler static adjustment works fairly well across several alternative specifications of a fully dynamic model. As each of these papers considers a different application, taken together, they demonstrate that the demand equivalence logic that is common among them is useful in a variety of contexts.⁴

Interpreting cross-regional estimates is tricky because of general equilibrium effects. On the one hand, these estimates include local general equilibrium effects. The focus of this paper is getting from the cross-regional estimates to partial equilibrium effects by removing local general equilibrium effects. On the other hand, cross-regional estimates difference out national general equilibrium effects. This implies that cross-regional estimates do not directly answer questions about aggregate effects (e.g., what is the aggregate effect of fiscal stimulus, the China shock, or the 2000s rise and fall of house prices). A rapidly growing recent literature has used multi-region general equilibrium models to assess what cross-regional estimates imply about these macro questions (e.g., Nakamura and Steinsson 2014; Beraja, Hurst, and Ospina 2019; Herreno 2020). One way to do this type of analysis is to use the cross-regional estimate to distinguish among competing general equilibrium models and then see what the favored general equilibrium model implies about the macroeconomic question of interest (Nakamura and Steinsson 2018).

The paper is organized as follows. Section II lays out the challenge the paper seeks to address. Section III presents adjustments of cross-sectional estimates of the housing wealth effect for local general equilibrium effects in simple static environments. Section IV explains that our adjustment of the housing wealth effect is an application of estimating one equation in a system of simultaneous equations. Section V provides a fully structural, multi-region macro model of the housing wealth effect on consumption. Section VI derives the matrix adjustment for local general equilibrium effects in the dynamic model presented in Section V. Section VII discusses the importance of distinguishing between the short-run and long-run housing supply elasticities. Section VIII conducts a Monte Carlo analysis of the fully structural model to evaluate the accuracy and robustness of the simple adjustment used in Section III. Section IX concludes.

II. Interpreting Cross-Regional Regressions

To address the issue of how to interpret cross-regional regression coefficients, it is useful to consider a concrete example. The example we focus on is the estimation of the housing wealth effect. A relatively standard estimating equation for a regional estimate of the housing wealth effect is

$$\Delta c_{i,t} = \psi_i + \xi_t + \beta \Delta p_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where i indexes cities, t indexes time, $\Delta c_{i,t}$ is the first difference of log consumption in city i , $\Delta p_{i,t}$ is the first difference of log house prices in the city, ψ_i is a city fixed effect, ξ_t is a time fixed effect, and $\varepsilon_{i,t}$ captures unmodeled influences. The coefficient of interest is β , which gives the elasticity of local consumption with respect to local house prices.

To identify the causal effect of local house prices on local consumption, researchers must confront the twin challenges of reverse causation and measurement error. A common approach to overcoming these challenges is premised on the view that house prices differentially respond to aggregate housing demand shocks across cities due to differences in housing supply elasticities. This approach is, for example, used by Mian et al. (2013), Mian and Sufi (2014), and Guren et al. (2020). Let's suppose we have used such a shift-share approach to form a causal estimate of β . The question we ask is, How should we interpret this estimate?

Let's suppose for simplicity that the consumption function of households in the regional economies under consideration is a linear function of household income $y_{i,t}$, house prices $p_{i,t}$, interest rates R_t , and the source of aggregate variation that drives house prices; call it Ω_t

$$c_{i,t} = C_y y_{i,t} + C_p p_{i,t} + C_R R_t + C_\Omega \Omega_t, \quad (2)$$

where C_y , C_p , C_R , and C_Ω are the coefficients in this linearized consumption function. Later in the paper, we describe an example of a model where such an equation arises, although in a dynamic model consumption will depend not only on current income and prices but also on future income and prices. In this consumption function, the coefficient C_p has a straightforward interpretation as the amount consumption changes when house prices rise holding incomes and other prices fixed. In other words, C_p summarizes a partial equilibrium experiment.

Let's now suppose that regression equation (1) is estimated on data generated by equation (2). In this case, what is the interpretation of the coefficient β ? First, the variation in consumption coming from $C_R \Delta R_t + C_\Omega \Delta \Omega_t$ and any other aggregate factors that may enter the consumption

function is common across cities and will be absorbed by the time fixed effect ξ_t . Notice that this means that the direct effect of the shock on consumption and any national general equilibrium effects that are mediated by national prices will not be captured by the coefficient β .

What about effects that the shock may have on local income $\Delta y_{i,t}$? To the extent that $\Delta y_{i,t}$ is correlated with $\Delta p_{i,t}$, our estimate of β in equation (1) will not only reflect the partial equilibrium housing wealth effect, C_p , but also the response of consumption to changes in local income. Unfortunately, there is a compelling reason to expect $\Delta y_{i,t}$ and $\Delta p_{i,t}$ to be correlated even when a sophisticated identification strategy is employed. The partial equilibrium housing wealth effect $C_p \Delta p_{i,t}$ will itself raise demand for locally produced goods and thereby bring about an increase in local incomes. This implies that β will reflect not just the partial equilibrium effect C_p but also local general equilibrium effects induced by the initial partial equilibrium effect.

An alternative, more direct, approach to estimating the partial equilibrium effect of house prices on consumption is to regress an individual-level consumption measure on changes in house prices and a control for individual-level income by ordinary least squares (OLS; e.g., Campbell and Cocco 2007) or instrumenting for changes in house prices (e.g., Aladangady 2017). It is important to note, however, that these studies use city-level house price variation. Without the individual-level income control, the coefficient on house prices in this type of regression will therefore include local general equilibrium effects. Furthermore, as Campbell and Cocco (2007) and Aladangady (2017) stress, there are several reasons why including individual-level income as a control is unlikely to purge the regression coefficient on changes in house prices of the local general equilibrium effects. First, individual-level income is likely measured with error. If so, changes in house prices may be a useful proxy for changes in income and the resulting omitted variable bias will include some component of the local general equilibrium effect. Second, changes in overall individual-level income are likely to have different statistical properties (e.g., less persistence) than changes in income induced by changes in house prices. This means that the marginal propensity to consume (MPC) out of these different sources of income is likely different, which implies that controlling for individual-level income would not accurately capture the local general equilibrium effects of house price changes. Third, there are likely other local general equilibrium channels that would not be controlled for by including individual-level income such as changes in the local price level.

This logic applies more generally outside the context of housing wealth effects. Although there are some situations where individual-level variation allows the direct estimation of a partial equilibrium effect, in many situations the most convincing identification strategies rely on regional variation. In other situations, all variations are by definition regional. In these cases, one must remove the local general equilibrium effects to obtain an estimate of the partial equilibrium effect.

III. A Simple Adjustment for Local General Equilibrium Effects

In this section we derive a simple adjustment of cross-sectional regression estimates of the housing wealth effect for local general equilibrium effects. The idea is that researchers who have constructed estimates of the housing wealth effect using cross-sectional regressions at the metropolitan or state level, can use this adjustment to recover a rough estimate of what their regression results imply about the partial equilibrium effect of house prices on consumption. The simple formulas we derive in this section are based on several approximations and simplifying assumptions. In Sections V and VI, we then present a fully specified multi-region general equilibrium model and an exact adjustment formula for this model. Later in the paper we show that the simple adjustment derived in this section is very close to the exact adjustment in certain cases and reasonably close in others.

A. *The Fiscal Multiplier as a Measure of Local General Equilibrium Effects*

Our central idea is that estimates of the local fiscal multiplier can be used to gauge the strength of local general equilibrium effects. We start by illustrating this in a very simple static case. The economy consists of two equally sized regions: “home” and “foreign.” Each region has three markets: a goods market, a housing market, and a labor market. In the home region, the two relative prices are the real wage w and the price of housing p , both denominated in goods. Suppose goods are produced with labor according to the production function $Y = N$, where Y is goods produced and N is labor supply. Labor supply is given by a function $N(w, p, T)$ that depends on the wage, the price of housing, and taxes T . Household demand for goods is given by $C(w, p, T)$. In addition to this private consumption demand, goods are used for public consumption in amount G , where G is exogenous. We assume that taxes are set at the national level

and satisfy a government budget constraint. It is important to our argument that both regions face the same taxes.

In this simplest case, we ignore trade across regions and across time; these features are added in Subsection III.C and Section V, respectively. The aggregate resource constraint is then $Y = C(w, p, T) + G$. Although this resource constraint is very standard, it embeds the important assumption that an increase in demand from private consumption requires the exact same supply response as an increase in demand from the government. For housing, we specify an excess demand function $H(w, p, T, s)$, where s is an exogenous shock. Because we only use data on the price of housing and not on the quantity of housing, we do not need to specify housing supply and demand separately.

Given these assumptions, the equilibrium level of wages and house prices in this home region is given by the solution to the following two equations:

$$C(w, p, T) + G = N(w, p, T) \quad (3)$$

$$H(w, p, T, s) = 0, \quad (4)$$

taking G and T as given.

The foreign region mirrors the home region with relative prices denoted w^* and p^* . These prices satisfy equations

$$C(w^*, p^*, T) + G^* = N(w^*, p^*, T) \quad (5)$$

$$H(w^*, p^*, T, s^*) = 0. \quad (6)$$

There are two important assumptions here. First, note that these are the same functions C , N , and H as in the home region but with different arguments. Second, in the foreign region the government spending differs from the home region, but the taxes do not.

We make two additional simplifying assumptions. First, there are no wealth effects on labor supply, $N_p \equiv \partial N / \partial p = 0$. This assumption implies that changes in housing wealth are not supply shocks in addition to being demand shocks. We view this assumption as being a reasonable approximation to reality in the short run. Second, we assume that house prices are independent of income, $H_w = 0$. This assumption is less likely to hold in reality. We relax it in Subsection III.D. It is, however, helpful to make this assumption in this first pass to simplify the exposition.

Suppose we observe an instrumental variables (IV) estimate of the housing wealth effect based on regional variation. In our notation, this is

$d\hat{Y}/d\hat{p} = (d\hat{Y}/ds)/(d\hat{p}/ds)$, where s is the shock (instrument) used to estimate $d\hat{Y}/d\hat{p}$ and a hat denotes a cross-regional difference: $\hat{Y} = Y - Y^*$. The trouble is that this IV estimate is the total derivative of consumption with respect to house prices, not the partial derivative C_p . The total derivative includes local general equilibrium effects; for instance, the initial shock may raise wages and lead to an increase in local consumption, which will further raise wages and increase local consumption, and so on. To get from the IV estimate to the partial equilibrium effect of house prices on consumption C_p , we need to adjust for these local general equilibrium effects.

To this end, take the total derivatives of equations (3) and (5) with respect to G and take the difference across regions of the resulting expressions to arrive at

$$C_w \frac{d\hat{w}}{dG} + C_p \frac{d\hat{p}}{dG} + 1 = N_w \frac{d\hat{w}}{dG}. \tag{7}$$

Note that the effect of taxes does not appear because $\hat{T} = 0$ as both regions face the same taxes. This lines up well with the empirical estimates, which include a constant term or time fixed effects and are therefore estimated off of cross-regional differences that omit factors that affect all regions equally. Taking total derivatives of equations (4) and (6) with respect to G and rearranging yields $d\hat{p}/dG = 0$. Then equation (7) implies $d\hat{w}/dG = 1/(N_w - C_w)$.

Taking total derivatives of equations (3)–(6) with respect to s and performing similar manipulations yields the response of prices to s . In summary, we have the following matrix of the price response to the two shocks:

$$\begin{pmatrix} \frac{d\hat{w}}{dG} & \frac{d\hat{w}}{ds} \\ \frac{d\hat{p}}{dG} & \frac{d\hat{p}}{ds} \end{pmatrix} = [N_w - C_w]^{-1} \begin{pmatrix} 1 & -\frac{H_s}{H_p} C_p \\ 0 & \frac{H_s}{H_p} (C_w - N_w) \end{pmatrix}. \tag{8}$$

We similarly differentiate the resource constraint. This yields the response of \hat{Y} to the two shocks:⁵

$$\frac{d\hat{Y}}{dG} = C_w \frac{d\hat{w}}{dG} + C_p \frac{d\hat{p}}{dG} + 1 \tag{9}$$

$$\frac{d\hat{Y}}{ds} = C_w \frac{d\hat{w}}{ds} + C_p \frac{d\hat{p}}{ds}. \tag{10}$$

Combining equations (8) and (10) yields

$$\begin{aligned}\frac{d\hat{Y}}{ds} &= -C_p \frac{H_s}{H_p} \frac{N_w}{N_w - C_w} \\ \frac{d\hat{p}}{ds} &= -\frac{H_s}{H_p},\end{aligned}$$

which in turn yields the regional IV estimate of the housing wealth effect

$$\frac{d\hat{Y}}{d\hat{p}} = \frac{d\hat{Y}/ds}{d\hat{p}/ds} = (1 - C_Y)^{-1} C_p, \quad (11)$$

where $C_Y \equiv C_w/N_w$ is the MPC out of income. From this we see that the regional IV estimate of the housing wealth effect is equal to the partial equilibrium response of consumption to house prices C_p multiplied by a local general equilibrium feedback factor $(1 - C_Y)^{-1}$.

Notice also that equations (8) and (9) imply that the local fiscal multiplier is equal to

$$\frac{d\hat{Y}}{dG} = 1 + \frac{C_w}{N_w - C_w} = (1 - C_Y)^{-1}. \quad (12)$$

The local fiscal multiplier is, thus, exactly equal to the local general equilibrium factor in equation (11). Intuitively, an increase in home prices of one unit spurs an extra C_p of spending, which then triggers local adjustments in wages with accompanying consumption effects. These same local adjustments occur when the initial spending is due to a government spending shock.

An important feature of government spending shocks is that their size is known: the dollar amount of the government spending is observed (or the dollar amount that is explained by whatever instrument we are using to identify exogenous variation in government spending). This is why the "direct effect" in equation (12) is equal to 1 as opposed to some unknown scaling factor. This feature makes the local fiscal multiplier estimates particularly useful as measures of local general equilibrium effects.

Combining these last two equations, we get that

$$C_p = \frac{d\hat{Y}/d\hat{p}}{d\hat{Y}/dG}. \quad (13)$$

In other words, the partial equilibrium effect of house prices on consumption, C_p , is equal to the cross-regional IV estimate of the housing wealth effect, $d\hat{Y}/d\hat{p}$, divided by the local fiscal multiplier, $d\hat{Y}/dG$. An

estimate of the local fiscal multiplier can therefore be used to convert a cross-regional IV estimate of the housing wealth effect into an estimate of the partial equilibrium effect of house prices on consumption.

In Guren et al. (2020), we estimate an MPCH of 3.3 cents on the dollar. This estimate corresponds to the total effect captured by $d\hat{Y}/d\hat{p}$. Nakamura and Steinsson (2014) estimate a local fiscal multiplier of about 1.5.⁶ Equation (13) then implies that the partial equilibrium MPCH is 2.2 cents on the dollar.

B. Residential Investment

A potentially important channel that we abstract from above is the response of residential investment to changes in house prices: an increase in house prices may induce an increase in residential investment, which then induces local general equilibrium effects. We now augment the simple model above to allow for this channel. In this case, the equilibrium level of wages and house prices in the home region are given by the solution to the following two equations:

$$C(w, p, T) + I(p, T) + G = N(w, T) \quad (14)$$

$$H(p, T, s) = 0. \quad (15)$$

Relative to the previous example, we have added demand for local goods coming from residential investment $I(p)$. We maintain the “one good” setup in which output produced with labor can be converted into consumption, government purchases, and residential investment. This implies that an increase in residential investment leads to a supply response and a general equilibrium response on wages and incomes that unfolds in the same way as changes in demand coming from private or public consumption. We will have to account for this response to recover the partial equilibrium response of consumption to house prices. In addition, we assume that residential investment is independent of income conditional on house prices, $I_w = 0$.

The foreign region again mirrors the home region. In an approach similar to the approach used in Subsection III.A, manipulation of the equilibrium conditions and resource constraints shows the regional IV estimate of the housing wealth effect for this case is

$$\frac{d\hat{Y}}{d\hat{p}} = \frac{d\hat{Y}/ds}{d\hat{p}/ds} = (1 - C_Y)^{-1}(C_p + I_p). \quad (16)$$

The difference from the approach used in Subsection III.A is that there are two partial equilibrium effects: one for consumption and another for residential investment.

The addition of residential investment does not affect the regional fiscal multiplier (as we have assumed that $I_w = 0$). The regional fiscal multiplier $(1 - C_Y)^{-1}$ can, therefore, again be used to adjust for local general equilibrium effects, which yields

$$C_p = \left(\frac{d\hat{Y}/d\hat{p}}{d\hat{Y}/dG} \right) - I_p.$$

Because $I_w = 0$, $I_p = d\hat{I}/d\hat{p}$. We must, however, also take account of the fact that the increase in residential investment contributes to $d\hat{Y}/d\hat{p}$. In this case, we have $d\hat{Y}/d\hat{p} = (d\hat{C}/d\hat{p}) + (d\hat{I}/d\hat{p})$. Using these expressions, we can rewrite the above equation as

$$C_p = \left[\frac{(d\hat{C}/d\hat{p}) + (d\hat{I}/d\hat{p})}{d\hat{Y}/dG} \right] - \frac{d\hat{I}}{d\hat{p}}. \quad (17)$$

In this case, we need a regional estimate of the response of residential investment to a change in house prices in addition to the regional housing wealth effect and fiscal multiplier estimates. We use changes in construction and real estate employment as a proxy for residential investment. Using a quarterly panel of core-based statistical area (CBSA)-level employment and house prices covering 1990–2017, we estimate an analogous specification to the housing wealth effect estimate we present in column 2 of table 1 of Guren et al. (2020). Results are reported in table 1 and we focus on our preferred “sensitivity” instrument. Our estimated elasticity is 0.362 (with a standard error of 0.053). To convert this to a simple derivative, we must divide by the ratio of housing wealth to residential investment H/I , which we estimate to be 28.2.⁷ This yields 0.013 as our estimate of $d\hat{I}/d\hat{p}$. Plugging in this estimate along with our estimates of the housing wealth effect and fiscal multiplier into equation (17) yields a partial equilibrium MPCH estimate of 1.8 cents on the dollar.

C. Demand Leakage Due to Trade

The analysis above ignores the fact that the data used to estimate the regional housing wealth effect come from regional economies that are open to trade with other regions. This implies that some of the extra demand induced by higher house prices “leaks out” to other regions. Let ϕ

Table 1
Elasticity of Construction and Real Estate Employment to Home Prices

Sample Period	1990–2017
Ordinary least squares	.470*** (.025)
Sensitivity instrument	.362*** (.053)
Saiz instrument	.500*** (.102)

Note: The specification is the same as the specification used for table 1 in Guren et al. (2020) except that the dependent variable is construction and real estate employment rather than retail employment. “Ordinary least squares” uses no instrument. “Sensitivity instrument” uses the sensitivity instrument described in Guren et al. (2020). Saiz uses an instrument that interacts Saiz’s elasticity with the national change in house prices. All three approaches use the same control variables: two-digit industry shares with date-specific coefficients, the cyclical sensitivity control described in Guren et al. (2020), and the analogously constructed controls for differential city exposure to interest rates and the Gilchrist-Zakrajsek excess bond premium along with CBSA and division-time fixed effects. Standard errors are two-way clustered at the time and CBSA level.

***Indicates statistical significance at the 0.1% level.

be the expenditure share on local goods. Households display home bias when $\phi > 1/2$. Here, again, we maintain the “one good” assumption, which in this case implies that the real exchange rate between the regions is fixed and equal to one. We relax this assumption in Section V. Building on equations (14)–(15), we have the following system:

$$\begin{aligned} \phi[C(w, p, T) + I(p, T)] + (1 - \phi)[C(w^*, p^*, T) + I(p^*, T)] + G &= N(w, T), \\ (1 - \phi)[C(w, p, T) + I(p, T)] + \phi[C(w^*, p^*, T) + I(p^*, T)] + G^* &= N(w^*, T), \\ H(p, s, T) &= 0, \\ H(p^*, s^*, T) &= 0. \end{aligned}$$

In this case, it is important to distinguish between local expenditures and local output. We define $E = C(w, p, T) + I(p, T)$ to be home expenditures. Regional IV estimates of the housing wealth effect measure the response of expenditures to house prices, not the response of output to house prices. With this definition of local expenditures, we can use

similar manipulations of the equilibrium conditions and resource constraints as in Subsections III.A and III.B to derive the regional IV estimate of the housing wealth effect:

$$\frac{d\hat{E}}{d\hat{p}} = \frac{d\hat{E}/ds}{d\hat{p}/ds} = (1 - \Phi C_Y)^{-1}(C_p + I_p), \quad (18)$$

and the local fiscal multiplier:

$$\frac{d\hat{Y}}{dG} = (1 - \Phi C_Y)^{-1}, \quad (19)$$

where $\Phi \equiv 2\phi - 1$. Once again, we see that the cross-regional housing wealth effect $d\hat{E}/d\hat{p}$ is equal to a partial equilibrium effect $(C_p + I_p)$ multiplied by a local general equilibrium effect, which is exactly equal to the local fiscal multiplier $d\hat{Y}/dG$.

The strength of the local general equilibrium effect is tempered by the degree of openness of each regional economy. Trade linkages attenuate the differences in activity across regions because some of the extra spending in the home region spills over onto the foreign region. In our simple model, this is captured by the factor Φ in the denominator on the right-hand side of equation (19). This factor runs from zero to one depending on the degree of home bias of demand. In a more complex model with movements in the relative prices of home and foreign goods, the elasticity of substitution between home and foreign goods would also play a role in determining the size of this attenuation.

Combining the last two equations, using the fact that $d\hat{E}/d\hat{p} = (d\hat{C}/d\hat{p}) + (d\hat{I}/d\hat{p})$, the fact that $I_p = d\hat{I}/d\hat{p}$, and rearranging yields

$$C_p = \left[\frac{(d\hat{C}/d\hat{p}) + (d\hat{I}/d\hat{p})}{d\hat{Y}/dG} \right] - \frac{d\hat{I}}{d\hat{p}}, \quad (20)$$

which is no different from equation (17). In other words, the attenuation due to openness has no impact on our analysis because it only shows up in the size of the local general equilibrium effect, which we measure in the data. There is, however, some subtlety to this outcome. The housing wealth effect is measured in terms of an expenditure response whereas the fiscal multiplier is measured in terms of a production response and normally we would think that production is more attenuated than expenditure. However, in specifying the model we assumed that the government buys a purely local good, not a mix of home and foreign goods. So, the production response to a government spending shock is no more attenuated than the expenditure response to house prices.

D. Income Effects on Housing

We now allow for an income effect on the price of housing. In this case, the housing market equilibrium conditions become

$$\begin{aligned} H(w, p, s, T) &= 0, \\ H(w^*, p^*, s^*, T) &= 0, \end{aligned}$$

where housing demand now responds to wages. The goods market equilibrium conditions are the same as in Subsection III.C. We can again use similar manipulations of the equilibrium conditions and resource constraints as in earlier sections to derive the regional IV estimate of the housing wealth effect:

$$\frac{d\hat{E}}{d\hat{p}} = \frac{d\hat{E}/ds}{d\hat{p}/ds} = (1 - \Phi_{C_Y})^{-1}(C_p + I_p), \tag{21}$$

and the local fiscal multiplier:

$$\frac{d\hat{Y}}{dG} = \left(1 - \Phi_{C_Y} + \Phi \frac{C_p + I_p}{N_w} \frac{H_w}{H_p} \right)^{-1}. \tag{22}$$

Derivations of these equations, which encompass all the derivations in this section as special cases, are presented in appendix A.

Notice that the housing wealth effect is the same in this case as in Subsection III.C: equation (21) is the same as equation (18). However, the local fiscal multiplier is different: equation (22) is different than equation (19). This means that the addition of income effects on house prices breaks the exact equivalence between the local fiscal multiplier and the local general equilibrium effects induced by a change in house prices. The reason for this is that part of the response to the government spending shock comes through house prices and the housing wealth effect, but we would like to isolate only the part of the fiscal multiplier that relates to wage adjustments.

To this end, suppose we observe an estimate of the response of house prices to income $d\hat{p}/d\hat{Y}$.⁸ Some further manipulation of the equilibrium conditions yields

$$C_p = \left[\frac{(d\hat{C}/d\hat{p}) + (d\hat{I}/d\hat{p})}{d\hat{Y}/dG} \right] \left[1 - \Phi \left(\frac{d\hat{C}}{d\hat{p}} + \frac{d\hat{I}}{d\hat{p}} \right) \frac{d\hat{p}}{d\hat{Y}} \right]^{-1} - \frac{d\hat{I}}{d\hat{p}}. \tag{23}$$

This expression, which we derive in appendix A, differs from equation (20) by the term in the large bracket. This term represents the

adjustment to the local fiscal multiplier estimate that is needed to focus on the part of general equilibrium effect that comes from change in wages.

We need two extra pieces of data to be able to evaluate equation (23). First, we need an estimate of $\Phi = 2\phi - 1$. We use Nakamura and Steinsson's (2014) estimate of $\phi = 0.69$, which implies $\Phi = 0.38$.⁹ Second, we need an estimate of the response of house prices to income $d\hat{p}/d\hat{Y}$. Lamont and Stein (1999) provide a set of short-run estimates of the income elasticity of house prices, which implies that it is less than 0.8 and more likely near 0.3. We use 0.3 as our estimate, but our conclusions are little changed by using 0.8. This estimate must be divided by the ratio of housing wealth to total expenditures of 2.02 to yield a value of 0.149 for the derivative $d\hat{p}/d\hat{Y}$ in equation (23).

Together, these numbers imply that the term inside the large bracket in equation (23) is 0.998. In other words, allowing for income effects on housing implies that the simple idea of dividing the housing wealth effect by the local fiscal multiplier is off by only a minuscule amount. Taking this effect into account does not affect the implied partial equilibrium MPCH to the precision we are reporting.

IV. Identification in a System of Equations

The arguments we made in the previous section may seem special. In fact, they are applications of a general method for identification that is used pervasively in economics: simultaneous equations estimation. To see this and to see how cross-regional estimates can be used in other settings to identify partial equilibrium responses, it is useful to recast what we have done as an application of the general method of identification when variables are determined by a system of equations. In doing this, we see that our approach to identifying partial equilibrium responses from cross-regional estimates is mathematically closely related to identification of structural VARs (another prominent example of simultaneous equations identification in macroeconomics).

Using this simultaneous equations perspective, it is easy to verify whether we have the necessary sources of variation to identify the partial equilibrium effect of interest. We see formally that the partial equilibrium housing wealth effect is not identified by cross-regional housing wealth effect estimates alone but is identified with the help of cross-regional responses to government spending shocks.

To keep things as simple as possible, we focus on the example from Subsection III.A. Appendix B shows that the same argument applies for

the more complicated system considered in Subsection III.D. Appendix B furthermore provides other examples of systems of equations—namely, one involving local credit supply shocks and another involving foreign demand shocks—where government spending shocks are again useful to identify the direct (partial equilibrium) effect of the shock.

Consider equations (3)–(6) from Subsection III.A. Linearizing these equations and using $Y = N(w)$ and its inverse, $w = N^{-1}(Y)$, we can write this as a system of three equations

$$\begin{aligned} \hat{C} &= C_p \hat{p} + C_Y \hat{Y} \\ \hat{Y} &= \hat{C} + \hat{G} \\ 0 &= H_p \hat{p} + H_s \hat{s}, \end{aligned}$$

where $C_Y = C_w/N_w$ as before. This system has three endogenous variables \hat{C} , \hat{Y} , and \hat{p} and two exogenous variables \hat{G} and \hat{s} . The exogenous variables \hat{G} and \hat{s} are the “structural shocks” affecting this system.

We can rewrite this system as

$$X\Gamma = z\Delta, \tag{24}$$

where X and z are row vectors given by $X = (\hat{C}, \hat{Y}, \hat{p})$ and $z = (\hat{G}, \hat{s})$; Γ and Δ are matrices of coefficients. Let’s now postmultiply the system by Γ^{-1} to obtain the reduced-form representation of the system:

$$X = z\Delta\Gamma^{-1}. \tag{25}$$

Our goal is to estimate C_p (an element of Γ). However, what we estimate using the reduced-form response of \hat{C} to \hat{s} is an element of $\Delta\Gamma^{-1}$, which is influenced by other aspects of the system such as the response of \hat{C} to \hat{Y} .

Formally, the identification challenge is that multiple structural systems (eq. [24]) can give rise to the same reduced form (eq. [25]). We can see this by postmultiplying equation (24) by an arbitrary 3×3 matrix F . This yields $X(\Gamma F) = z(\Delta F)$. The reduced form of this system is

$$X = z(\Delta F)(\Gamma F)^{-1} = z\Delta\Gamma^{-1}.$$

The F and F^{-1} cancel out implying that this alternative structural system yields the same reduced-form system. For example, the same reduced-form response of C to s can arise if C_p is small and C_Y is large or vice versa.

To identify the system, we need to place restrictions on the system (on Γ and Δ) such that there is no scope to transform the system with an F

without violating one or more of the restrictions. To be more precise, we need restrictions so that the only allowable F is the identity matrix.

Note the analogy to structural VARs. A structural VAR is $A(L)Y_t = u_t$ (with $A(0) = I$) and $Ru_t = \epsilon_t$, where u_t are the reduced-form errors and ϵ_t are the structural errors. Equation (25) in our application is analogous to $Ru_t = \epsilon_t$ in the structural VAR setting. In structural VARs, identification involves placing restrictions on R , just as in our setting identification involves placing restrictions on Δ and Γ . However, there are some differences. In structural VARs, the typical assumption is that ϵ_t is not observed but has the same dimension as u_t (same number of shocks as variables). Researchers place restrictions on R (and use the variance-covariance matrix of u_t) to identify the remainder of R (and ϵ_t). In our case, the structural shocks z are observed, but z may have a smaller dimension than X . We are interested in identifying an element of Γ but can only observe elements of $\Delta\Gamma^{-1}$. Our case is more analogous to identification of structural VARs with observed external instruments.

As our interest is in identifying the partial equilibrium housing wealth effect, we focus on whether the coefficients in the first equation of our system—the linearized consumption function—are identified. This is less demanding than identifying the whole system because we can allow for transformations of the system that change the other equations as long as they do not change the first equation. In other words, we need only be sure that the first column of any transformation F is the first column of the identity matrix.

In general, identification of (parts of) a system of equations requires a certain number of restrictions. Intuitively, the number of observed facts must equal the number of unknown parameters of interest. In the simultaneous equations literature, this is known as the “order condition” (see, e.g., Wooldridge 2010, section 9.2.2). In addition, the restrictions must result in a system that has certain rank. Intuitively, the facts must each provide information about the system that is independent of any linear combination of the other facts. This is known as the “rank condition.” Finally, identification requires a normalization.

We begin by discussing the order condition for our case. As our system contains three endogenous variables, we need two restrictions on the system to satisfy the order condition. In our example, these restrictions are that neither G nor s appears directly in the consumption function. These restrictions are similar to the exclusion condition in an IV regression. Without excluding G from the consumption function, we would not be able to tell if a large fiscal multiplier reflects a large C_Y or a large direct effect of G on

C. But to understand the local general equilibrium amplification of the housing wealth effect we specifically want to know C_Y . Notice that if we dropped \hat{G} from the system, we would have only one restriction and we would not be able to satisfy the order condition for identification of C_p without making a direct assumption about C_Y .

We now turn to the rank condition. Our restrictions can be represented by the matrices

$$R_\Gamma = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad R_\Delta = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

such that

$$\underbrace{(R_\Gamma \Gamma + R_\Delta \Delta)}_{\equiv \mathcal{R}} e_1 = 0, \tag{26}$$

where $e_1 \equiv (1, 0, 0)'$ selects the first column of \mathcal{R} .

Now suppose we transform the system by postmultiplying it by F . As we want the first columns of Γ and Δ to be identified, we want to make sure that only an F with the first column e_1 can satisfy the restrictions on the system. In this case, F does not transform the first column of Γ and Δ , which houses the coefficients of the first equation in the system. To satisfy the restrictions on the system, F must satisfy

$$(R_\Gamma \Gamma F + R_\Delta \Delta F) e_1 = \mathcal{R} F e_1 = \mathcal{R} f_1 = 0, \tag{27}$$

where $f_1 \equiv F e_1$ is the first column of F .

Our goal is that F does not transform the first equation in the system so its first column is the first column of the identity matrix, that is, $f_1 = e_1$. So the question is what must be true of \mathcal{R} such that f_1 can only satisfy equation (27) if $f_1 = e_1$? To answer that, we want to make sure that \mathcal{R} maps any vector other than e_1 to a nonzero vector. This requires that all of the columns of \mathcal{R} other than the first one are linearly independent. Or equivalently, the rank of \mathcal{R} must be one less than the number of equations, which means two in our case. This is the rank condition. It is easily verified for our system.

The final requirement for identification is a normalization. The rank condition only requires that $f_1 = e_1$ up to a scalar multiple and so a normalization is needed to pin down this scalar. In our system, the first equation has a natural normalization: when we write $\hat{C} = C_p \hat{p} + C_Y \hat{Y}$ we have imposed that the coefficient on \hat{C} is 1. Together, these conditions—the

rank condition with a normalization—are necessary and sufficient for identification (Wooldridge 2010, Thm. 9.2).

Clearly if there are fewer than two restrictions, the rank of \mathcal{R} cannot be two because \mathcal{R} has fewer than two rows. (This is the order condition.) Suppose that G did not appear in the second equation and was therefore absent from the system altogether. In this case, we would only have one restriction to place on Δ and Γ and the order condition would fail. Introducing a second exogenous source of variation (G) is therefore crucial for identification in our application.

One potentially useful piece of prior knowledge we have about G is that the coefficient on G in the resource constraint is 1 (i.e., we know the direct demand effect of the G shocks). However, the argument above does not make use of this information. Suppose we introduce an unknown parameter Y_G such that the second equation becomes

$$\hat{Y} = \hat{C} + Y_G \hat{G}.$$

Even though we do not know Y_G , we can still identify C_p by the argument above. The reason is that in our reduced form we are implicitly assuming that we observe both $d\hat{Y}/d\hat{G}$ and $d\hat{C}/d\hat{G}$. Differentiating the resource constraint gives

$$\frac{d\hat{Y}}{d\hat{G}} = \frac{d\hat{C}}{d\hat{G}} + Y_G.$$

The parameter Y_G can then be found from the difference between $d\hat{Y}/d\hat{G}$ and $d\hat{C}/d\hat{G}$.

However, if we instead assume that we are not able to observe the consumption response $d\hat{C}/d\hat{G}$ directly, then we do need to use the prior knowledge that $Y_G = 1$. To see this, eliminate \hat{C} from the system so we have a system of two endogenous variables in two unknowns \hat{Y} and \hat{p} :

$$\begin{aligned} (1 - C_Y)\hat{Y} &= C_p\hat{p} + Y_G\hat{G} \\ 0 &= H_p\hat{p} + H_s\hat{s}. \end{aligned}$$

Our goal is still to identify C_p . So, we are interested in identifying the first equation. We can impose the restriction that \hat{s} does not appear in the first equation. This suffices for the order and rank conditions to be satisfied. The knowledge that $Y_G = 1$ then provides the crucial normalization we need to complete the identification. This shows that one aspect of government spending shocks that can make them more valuable for identification of partial equilibrium responses than other shocks is the fact that we know the size of the shock.

V. Dynamic Model

We now present a microfounded, dynamic model of multiple regions. After presenting the model, we show how the arguments laid out in Sections III and IV can be applied in the context of this model.

Demographics. There are two regions: “home” and “foreign.” The population of the entire economy is 1 with a share n in the home region. All variables are expressed in per capita terms.

Preferences. Households maximize

$$\mathbb{E}_0 \sum_{t=0} \beta^t u(C_t, N_t, H_t; \Omega_t),$$

where the arguments are consumption, labor supply, units of housing H_t chosen at date t and held to date $t + 1$, and an aggregate housing demand shock Ω_t . The period utility function is given by

$$u(C, N, H; \Omega) = \frac{\left\{ \left[\left(C - \frac{N^{1+\nu}}{1+\nu} \right)^\kappa (H - \Omega)^{1-\kappa} \right]^{1-\sigma} - 1 \right\}}{(1 - \sigma)}.$$

Note that consumption and leisure are substitutable in the style of Greenwood, Hercowitz, and Huffman (1988), which eliminates wealth effects on labor supply, an assumption we maintained in Section III. We model the housing demand shock using a Stone-Geary formulation, but this exact specification is unimportant. What matters is that there is a shock that changes the marginal rate of substitution between housing and nondurables.¹⁰

Commodities and technology. Consumption C_t is a Cobb-Douglas bundle of final goods produced in both home and foreign regions:

$$C_t = \phi^{-\phi} (1 - \phi)^{-(1-\phi)} C_{H,t}^\phi C_{F,t}^{1-\phi},$$

where $C_{F,t}$ is the consumption in the home region of the good produced in the foreign region.¹¹ We use * to denote foreign variables. So, $C_{H,t}^*$ is the consumption in the foreign region of the good produced in the home region. We assume

$$C_t^* = \phi^{*\phi} (1 - \phi^*)^{-(1-\phi^*)} C_{F,t}^{*\phi} C_{H,t}^{*(1-\phi^*)}.$$

The parameters $\phi > n$ and $\phi^* > 1 - n$ capture the degree of home bias in demand for goods. The price index for the consumption bundle in the home region is $\mathcal{P}_t = \mathcal{P}_{H,t}^\phi \mathcal{P}_{F,t}^{1-\phi}$, where $\mathcal{P}_{H,t}$ and $\mathcal{P}_{F,t}$ are the prices of the final goods produced in the home and foreign regions, respectively.

Each region produces a final good using a continuum of intermediate inputs. The production of the final good satisfies

$$Y_t = \left(\int_0^1 y_t(z)^{\frac{\eta-1}{\eta}} dz \right)^{\frac{\eta}{\eta-1}}.$$

Each intermediate good is produced linearly with labor according to $y_t(z) = N_t(z)$.

Housing supply. The supply of housing satisfies

$$H_t = (1 - \delta)H_{t-1} + I_t^\alpha M_t^{1-\alpha}. \quad (28)$$

Here we assume that the construction of new residential housing units requires two inputs: residential investment I_t and construction permits M_t , which are sold by the federal government. The construction permits are a tractable way to represent a variety of factors that limit housing supply including zoning regulations and limits to new land supply. The elasticity of supply of the construction permits may differ across regions giving rise to different housing supply elasticities. Residential investment requires a mix of local and imported inputs analogous to the mix used for consumption:

$$I_t = \phi^{-\phi} (1 - \phi)^{-(1-\phi)} I_{H,t}^\phi I_{F,t}^{1-\phi}.$$

Markets. The two regions share the same money, which serves as the numeraire. Final goods markets are competitive and completely integrated across regions. The prices of intermediate-goods firms are sticky. These firms receive an opportunity to change their price each period with $1 - \chi$ as in Calvo (1983). The labor markets are local to each region and competitive with real wages denoted w_t . Units of housing trade at relative price p_t . Households trade a nominal bond that pays interest i_t between t and $t + 1$. Let $\mathcal{P}_t B_t$ be the nominal value of bond holdings in the home region at the end of period t . We consider two cases for asset markets. In the “incomplete markets” economy, there is only trade in risk-free nominal bonds. In the “complete markets” economy, the regions also trade state-contingent assets in quantities A_t at prices $\Xi_{t,t+1}$. In the complete markets economy, the bond is redundant, but it can still be priced and this price will enter our monetary policy rule.

Intermediate-goods firms produce profits, which are rebated to the households in the region. We use D_t to denote the real profits received. We impose a portfolio holding cost in the style of Schmitt-Grohé and Uribe (2003) whereby holding bond position B_t incurs a flow cost ζB_t^2 . This portfolio cost implies that steady-state wealth holdings in each region are determinate. This can be viewed as a crude approximation to precautionary savings motives that decline with wealth.

Government. The government purchases goods, sells construction permits, and sets monetary policy. Let G_t and G_t^* be per capita spending in home and foreign regions, respectively. The government buys local goods in each region. The exogenous process for G_t is

$$G_t = (1 - \rho_G)\bar{G} + \rho_G G_{t-1} + \epsilon_{G,t}. \tag{29}$$

G_t^* is independent of G_t but follows the same process.

The government’s monetary policy may be described by a rule for the nominal interest rate:

$$1 + i_t = \beta^{-1} + \varphi_\pi (\pi_t^n \pi_t^{*1-n} - 1) + \varphi_y \left[n \log \left(\frac{Y_t}{\bar{Y}} \right) + (1 - n) \log \left(\frac{Y_t^*}{\bar{Y}^*} \right) \right], \tag{30}$$

where policy responds to the population-weighted averages of inflation and output.

The government sells construction permits according to the rule

$$M_t = \bar{M} p_t^\gamma. \tag{31}$$

The parameter γ is the elasticity of construction permits granted with respect to the price of housing. This parameter is meant to reflect some combination of the stringency of zoning regulations and the availability of suitable vacant land. The government sets the relative price of a permit, q_t , equal to its marginal product in construction of new housing units. It is fairly standard to model housing supply as combining a flow of new land or permits with residential investment. We assume that the supply of permits is price elastic whereas the literature typically assumes it is constant (Davis and Heathcote 2005; Favilukis, Ludvigson, and Van Nieuwerburgh 2017; Kaplan, Mitman, and Violante 2020). Later, we allow the regions to differ in their permit supply elasticities (i.e., $\gamma \neq \gamma^*$) in the spirit of identification schemes that follow Saiz (2010).

The government imposes lump-sum taxes in nominal amounts $\mathcal{P}_t T_t$ and $\mathcal{P}_t^* T_t^*$. The national government budget constraint is

$$n\mathcal{P}_{H,t} G_t + (1 - n)\mathcal{P}_{F,t}^* G_t^* = n\mathcal{P}_t T_t + (1 - n)\mathcal{P}_t^* T_t^* + n\mathcal{P}_t q_t + (1 - n)\mathcal{P}_t^* q_t^*.$$

We assume that the government taxes each region equally (per capita) in nominal terms.

Market clearing. The market for home goods clears if

$$Y_t = \phi \left(\frac{\mathcal{P}_{H,t}}{\mathcal{P}_{F,t}} \right)^{\phi-1} (C_t + I_t) + \frac{1-n}{n} (1 - \phi^*) \left(\frac{\mathcal{P}_{H,t}}{\mathcal{P}_{F,t}} \right)^{-\phi^*} (C_t^* + I_t^*) + G_t, \quad (32)$$

where $\mathcal{P}_{H,t}/\mathcal{P}_{F,t}$ is the real exchange rate. This expression involves local and home expenditures on the bundles of home- and foreign-produced goods. The cost-minimizing bundle depends on the degree of home bias and the real exchange rate. Similarly, the market for foreign goods clears if

$$Y_t^* = \frac{n}{1-n} (1 - \phi) \left(\frac{\mathcal{P}_{H,t}}{\mathcal{P}_{F,t}} \right)^{\phi} (C_t + I_t) + \phi^* \left(\frac{\mathcal{P}_{H,t}}{\mathcal{P}_{F,t}} \right)^{1-\phi^*} (C_t^* + I_t^*) + G_t^*. \quad (33)$$

Bond market clearing requires

$$nB_t + (1-n)B_t^* = 0.$$

Decision problems. Under incomplete markets and assuming certainty equivalence so that the real return on bonds is treated as known, $R_t \equiv (1 + i_t)/\pi_{t+1}$, the household maximizes

$$\sum_{t=0}^{\infty} \beta^t u(C_t, N_t, H_t; \Omega_t),$$

subject to the budget constraint

$$p_t H_t + C_t + B_t + \zeta B_t^2 = W_t N_t + D_t + R_{t-1} B_{t-1} + p_t H_{t-1} (1 - \delta),$$

where R_t is the gross real interest rate between t and $t + 1$.¹² The intra-temporal optimality conditions of the household's problem imply

$$\begin{aligned} \tilde{H}_t &= x_t \tilde{C}_t \\ N_t^p &= w_t \end{aligned}$$

where

$$\begin{aligned} \tilde{C}_t &\equiv C_t - \frac{N_t^{1+\nu}}{1+\nu} \\ \tilde{H}_t &\equiv H_t - \Omega_t \\ x_t &\equiv \frac{1-\kappa}{\kappa} \left(p_t - \mathbb{E}_t \left[p_{t+1} (1-\delta) \beta \frac{u_{C,t+1}}{u_{C,t}} \right] \right)^{-1}. \end{aligned}$$

Abstracting from the portfolio holding cost, we have

$$\tilde{C}_t = \kappa \frac{R_{t-1}B_{t-1} + p_t H_{t-1}(1 - \delta) + \sum_{\tau=t}^{\infty} R_{t,\tau}^{-1} \left[Y_{\tau} - \frac{1-\kappa}{\kappa X_{\tau}} \Omega_{\tau} - \frac{N_{\tau}^{1+\nu}}{1+\nu} \right]}{\sum_{\tau=t}^{\infty} R_{t,\tau}^{-1} X_{t,\tau}}, \quad (34)$$

where

$$R_{t,\tau} \equiv R_{t,\tau-1} R_{\tau-1} \quad \forall \tau > t$$

$$X_{t,\tau} \equiv \left[\beta^t R_{t,\tau} \left(\frac{x_{\tau}}{x_t} \right)^{(1-\kappa)(1-\sigma)} \right]^{1/\sigma},$$

and $R_{t,t} = 1$. See appendix C.1 (apps. C–E are available online) for the derivation.

Turning to construction, a representative competitive real estate developer maximizes revenue from new houses less material and permit costs:

$$\max_{I_t, M_t} \{ p_t I_t^{\alpha} M_t^{1-\alpha} - I_t - q_t M_t \}.$$

The first-order condition of this problem with respect to I_t and equation (31) implies

$$I_t = \alpha^{\frac{1}{1-\alpha}} \bar{M} p_t^{\frac{\gamma+1}{1-\alpha}}, \quad (35)$$

so the supply of new housing is

$$I^{\alpha} M_t^{1-\alpha} = (\alpha p_t)^{\frac{\alpha}{1-\alpha}} M_t.$$

Finally, intermediate-goods producers set their reset prices, $\tilde{\mathcal{P}}_t$, to solve

$$\max_{\tilde{\mathcal{P}}_t} \mathbb{E}_t \sum_{\tau=t}^{\infty} \chi^{\tau} \lambda_{t,\tau} \left[\left(\frac{\tilde{\mathcal{P}}_t}{\mathcal{P}_{\tau}} - w_{\tau} \right) y_{\tau} \right]$$

where $\lambda_{t,\tau}$ is the discount factor between τ and t and subject to the demand curve for their variety $y_{\tau} = Y_{\tau} (\tilde{\mathcal{P}}_t / \mathcal{P}_{H,\tau})^{-\eta}$. This problem gives rise to a forward-looking inflation response to variations in the real wage and the real exchange rate (see app. D.1).

VI. Adjusting for Local General Equilibrium Effects in the Full Model

We now show how the static relationships derived in Section III relate to dynamic relationships in the context of the dynamic model. We consider a perfect foresight transition lasting \mathcal{T} periods. We assume that the two

regions are equally open to trade. Given their unequal sizes, this implies $1 - \phi^* = n/1 - n(1 - \phi)$. We define $\Phi \equiv \phi + \phi^* - 1$. Furthermore, to keep the expressions in this section as simple as possible, we assume that prices are perfectly rigid. This implies that the real exchange rate is constant at one.

Taking a cross-regional difference of the market-clearing conditions (eqs. [32] and [33]) yields

$$\hat{Y} = \Phi(\hat{C} + \hat{I}) + \hat{G}, \quad (36)$$

where \hat{Y} is a column vector of length T that gives values of $Y_t - Y_t^*$ for all $t \in \{1, \dots, T\}$. \hat{C} , \hat{I} , and \hat{G} are defined similarly.

Linearizing the consumption function shown in equation (34) around a symmetric steady state and taking a cross-regional difference yields

$$\hat{C} = \mathbf{C}_p \hat{p} + \mathbf{C}_Y \hat{Y}, \quad (37)$$

where \mathbf{C}_Y is a $T \times T$ matrix where the $[t, s]$ element gives the coefficient of the response of C_t to Y_s and the matrix \mathbf{C}_p is defined similarly. \mathbf{C}_Y is the intertemporal MPC matrix highlighted by Auclert, Rognlie, and Straub (2018). Notice that Ω , taxes, and interest rates do not appear in this expression because these variables are common across regions and drop out when we take the difference.

Linearizing the residential investment response (eq. [35]) around a symmetric steady state and taking a cross-regional difference yields

$$\hat{I} = \mathbf{I}_p \hat{p}, \quad (38)$$

where we abstract (for now) from regional heterogeneity in land supply.

Combining equations (36), (37), and (38) yields

$$\hat{Y} = \Phi \mathbf{M}(\mathbf{C}_p + \mathbf{I}_p) \hat{p} + \mathbf{M} \hat{G} \quad (39)$$

where $\mathbf{M} \equiv [I - \Phi \mathbf{C}_Y]^{-1}$. Using equations (37), (38), and (39), local expenditure is given by

$$\hat{C} + \hat{I} = \mathbf{M}(\mathbf{C}_p + \mathbf{I}_p) \hat{p} + \mathbf{C}_Y \mathbf{M} \hat{G}, \quad (40)$$

where we have used the definition of \mathbf{M} to note that $I + \mathbf{C}_Y \Phi \mathbf{M} = \mathbf{M}$.

From this last equation, we can calculate the impulse response of expenditures to home prices as

$$\frac{d\hat{E}}{d\hat{p}} = \mathbf{M}(\mathbf{C}_p + \mathbf{I}_p), \quad (41)$$

where $\hat{E} \equiv \hat{C} + \hat{I}$ and $d\hat{E}/d\hat{p}$ is a $\mathcal{T} \times \mathcal{T}$ matrix in which the (t, s) element gives the response of expenditure in period t to a change in house prices in period s . This is the dynamic analog to our static (scalar) IV estimate $d\hat{E}/d\hat{p}$ from the setting discussed in Section III.

Notice that, just as in Section III, the regional impulse response is equal to the partial equilibrium response $\mathbf{C}_p + \mathbf{I}_p$ multiplied by a local general equilibrium feedback factor \mathbf{M} . Our next task is to relate this local general equilibrium feedback factor to observables. To this end, we linearize housing demand, $H(p, Y, T, R, \Omega)$, and housing supply (eq. [28]) and equate them, which yields

$$\mathbf{H}_p p + \mathbf{H}_R R + \mathbf{H}_Y Y + \mathbf{H}_T T + \mathbf{H}_\Omega \Omega = \mathbf{H}_p^S p.$$

Taking a cross-regional difference and rearranging yields

$$\hat{p} = \mathbf{p}_Y \hat{Y}, \tag{42}$$

where

$$\mathbf{p}_Y \equiv (\mathbf{H}_p^S - \mathbf{H}_p)^{-1} \mathbf{H}_Y,$$

and \mathbf{H}_Y is the response of housing demand to income. Substituting equation (42) into equation (39) and rearranging yields

$$\hat{Y} = [I - \Phi \mathbf{M} (\mathbf{C}_p + \mathbf{I}_p) \mathbf{p}_Y]^{-1} \mathbf{M} \hat{G}. \tag{43}$$

From equation (43) the impulse response of output to government spending is:

$$\frac{d\hat{Y}}{d\hat{G}} = [I - \Phi \mathbf{M} (\mathbf{C}_p + \mathbf{I}_p) \mathbf{p}_Y]^{-1} \mathbf{M}. \tag{44}$$

Here again, $d\hat{Y}/d\hat{G}$ is a $\mathcal{T} \times \mathcal{T}$ matrix in which the (t, s) element gives the response of output in period t to a change in government spending in period s . This is the dynamic equivalent of the static (scalar) local fiscal multiplier from the setting discussed in Section III.

Rearranging equation (44) and using the fact that $d\hat{p}/d\hat{Y} = \mathbf{p}_Y$ we have that

$$\mathbf{M} = \left(I - \Phi \frac{d\hat{E}}{d\hat{p}} \frac{d\hat{p}}{d\hat{Y}} \right) \frac{d\hat{Y}}{d\hat{G}}.$$

Plugging this expression for \mathbf{M} into equation (41) and rearranging yields

$$\mathbf{C}_p = \left[\frac{d\hat{Y}}{d\hat{G}} \right]^{-1} \left(I - \Phi \frac{d\hat{E}}{d\hat{p}} \frac{d\hat{p}}{d\hat{Y}} \right)^{-1} \left(\frac{d\hat{E}}{d\hat{p}} \right) - \mathbf{I}_p. \tag{45}$$

This result is a dynamic analog to the main result in Subsection III.D (see eq. [23]). As we note above, in this dynamic setting, the components $d\hat{E}/d\hat{p}$, $d\hat{Y}/dG$, and so on are matrices rather than scalars. If these matrices have important off-diagonal elements, then the logic of our static examples is complicated by dynamic responses of the economy. However, if the contemporaneous responses are large relative to the dynamic responses (i.e., the matrices are close to diagonal), then the logic of the static economy goes through because in that case equation (45) reduces to the same scalar relationship as equation (23).

VII. Long-Run Heterogeneity in Housing Supply

To provide a convincing assessment of the accuracy of the simple adjustment formula that we present in Section III, we need a model that can roughly match our empirical estimates for the observables that enter that adjustment formula: the local housing wealth effect, the local fiscal multiplier, and the local effect of house prices on construction. Our estimates of the local housing wealth effect and local effect of house prices on construction rely on an identification strategy that exploits heterogeneity in housing supply curves across cities interacted with aggregate home price changes as an instrument. Research using this identification strategy has proxied for city-level housing supply elasticities with topographic features of the cities (Saiz 2010) or equilibrium sensitivity of local house prices to regional house price variation (Guren et al. 2020).

We can introduce heterogeneity in housing supply elasticities into the model we present in Section V and assess whether this model can match our empirical estimates. When we do this, we find that the response of residential investment to house prices is far from our empirical estimates. In fact, this response is negative in the model: when we calibrate the model to have γ and γ^* equal to the 10th and 90th percentiles, respectively, of the elasticities estimated by Saiz (2010), the model implies an elasticity of residential investment to house prices of -12 . In contrast, the empirical estimates we present in Section III are positive, ranging from 0.36 to 0.50.

The left panel of figure 1 illustrates the economics behind this counterfactual prediction of the model. The panel plots housing supply curves for two cities with different supply elasticities. In response to a common demand shock, the less elastic city (represented by line S) has a larger price response but a smaller quantity response.¹³ Davidoff (2016) has




Fig. 1. Long-run and short-run housing supply curves. A color version of this figure is available online.

critiqued estimates of the housing wealth effect based on the identification strategy described above on the grounds that quantity growth has been larger for inelastic cities than elastic cities. Our finding that construction employment responds positively to changes in house prices is another piece of evidence on quantities that is at odds with the simple view of heterogeneous housing supply curves discussed above.

However, it is relatively simple to remedy this empirical failing of our model by allowing for differences between short-run and long-run housing supply elasticities. Suppose, in particular, that housing supply is inelastic in all cities in the short run but becomes more elastic in the long run in one of the cities. The two panels of figure 1 illustrate this with the right panel depicting the short run and the left panel depicting the long run. The equilibrium price of housing is forward looking as current housing demand depends on expectations of all future capital gains on housing. As a result, the equilibrium housing price is largely determined by long-run forces in the housing market. However, the incentives to construct and sell new houses depend on the current availability of inputs to construction and the current price of houses. Therefore, the construction response depends much more on short-run forces in the housing market. The short-run equilibrium in the housing market reflects different endogenous changes in housing demand across regions reflecting the different expected capital gains going forward. Short-run housing demand will increase more in the region in which a larger capital gain is expected and this moves that region further up the common short-run supply curve. This logic generates a positive response of construction to house prices even though regional house price fluctuations reflect differences in (long-run) housing supply elasticities.

We use a regime-switching formulation to model short-run and long-run differences. We assume that the economy is currently in the short-run regime and is expected to switch to the long-run regime with 2% probability each period after which it will remain in the long-run regime. In the short run, the supply of land available for construction is fixed in both regions ($\gamma = \gamma^* = 0$). In the long run, the supply of land responds to house prices but differentially in the two regions ($\gamma < \gamma^*$). When we simulate the economy, we assume that the economy is always in the short-run regime and the long-run regime never materializes. Construction in both regions reflects movements along the same short-run supply curve whereas house prices are differentially affected by aggregate changes in housing demand that move expectations along the heterogeneous long-run housing supply curves.

VIII. Monte Carlo Analysis

We are now ready to use the dynamic model developed in Sections V and VII to assess the accuracy of the simple adjustment for local general equilibrium effects that we derived in Section III (eq. [23]). We do this by presenting results from a series of Monte Carlo simulations of different versions of the dynamic model. Recall that discrepancies between the appropriate size of the adjustment in our full model and the simple adjustment formula arise from the nondiagonal elements in equation (45).

We will work with equation (23) rewritten in terms of elasticities:

$$\frac{\partial \log C}{\partial \log p} = \left[\frac{e}{\frac{d\hat{Y}}{dG} \left(1 - \Phi e \frac{\bar{C}}{\bar{Y}} \frac{d \log \hat{p}}{d \log \hat{Y}} \right)} \right] - \frac{\bar{I}}{\bar{C}} \frac{d \log \hat{I}}{d \log \hat{p}} \quad (46)$$

where

$$e \equiv \frac{d \log \hat{C}}{d \log \hat{p}} + \frac{\bar{I}}{\bar{C}} \frac{d \log \hat{I}}{d \log \hat{p}}$$

and we define $\log \hat{X} = \log X - \log X^*$.

The left-hand side of equation (46) is the partial equilibrium response of consumption to changes in house prices. In a dynamic setting, this depends not only on current house prices but also on the entire future expected path of house prices. In our model, we compute the partial equilibrium effect using the analytical consumption function (eq. [34]), and the exact dynamics of house prices implied by the model. Specifically, we first simulate the general equilibrium model. Then for each period

of this simulation, we feed the resulting current house price and expected future path for house prices into the household's analytical consumption function and calculate the household's consumption (holding everything other than the path for house prices constant). This yields a series for consumption in the home and foreign regions. Finally, we regress the cross-regional difference in consumption on the simulated cross-regional difference in house prices and take the regression coefficient from this regression as the partial equilibrium response of consumption to house prices. See appendix C for additional details.

To calculate the right-hand side of equation (46) in our model, we perform two simulations of the model. First, we simulate the model with only aggregate housing demand shocks. The output from this simulation allows us to calculate $d \log \hat{C} / d \log \hat{p}$ and $d \log \hat{I} / d \log \hat{p}$. We do this simply by running an OLS regression of the change in $\log C$ and $\log I$, respectively, on the change in $\log p$ in the model and including time fixed effects. By limiting the shocks to aggregate housing demand shocks in this simulation, we are estimating using the same variation that we isolate with our sensitivity instrument in Guren et al. (2020). Next, we simulate the model with only region-specific government purchases shocks. The output from this simulation allows us to calculate $d \hat{Y} / d \hat{G}$ and $d \log p / d \log Y$ in a similar fashion. Conceptually, when we calculate $d \log \hat{p} / d \log \hat{Y}$ from the output of this simulation, we are implicitly using government purchases as an instrument for changes in local output. Finally, \bar{C} / \bar{Y} and \bar{I} / \bar{C} are steady-state ratios and $\Phi \equiv \phi + \phi^* - 1$ is a measure of home bias all of which we assume are known to the analyst.

A. Model Calibration

We calibrate the model as follows. A time period is a quarter. We set the population share of the home region to be 2% with the interpretation that the home region is a city and the foreign region is the rest of the country. We set the home-bias parameter $\phi = 0.4$ based on the share of shipments in the Commodity Flow Survey that goes to the same metro area. We set the elasticity of substitution between varieties to be $\eta = 6$.

We normalize the steady-state supply of construction permits, \bar{M} , so that the steady-state relative price of a unit of housing is one. We set the depreciation rate on housing to 3% annually. We target a 4.4% share of residential investment to gross domestic product (GDP), which is the average ratio over the period 1970–2019. This implies that the residential investment share in the construction of new houses (i.e., one minus the

land share) is $\alpha = 0.38$. The home and foreign regions differ in their land (construction permit) supply elasticities, γ . We set them to match the 10th and 90th percentiles of the elasticity estimates from Saiz (2010), which are 1.05 and 4.39. The Saiz estimates reflect the response of housing units, which we interpret as the change in H_t . The (long-run) price elasticity of housing supply in the model is $[\alpha/(1 - \alpha)] + \gamma$. Therefore, we set $\gamma = 0.45$ and $\gamma^* = 3.78$.

Turning to preferences, we set the subjective discount factor $\beta = 0.99$, and we set $\kappa = 0.58$ to target a 25% expenditure share on housing, which is the average housing expenditure in the Consumer Expenditure Survey (CEX) in 2018. We set the labor supply elasticity to $\nu = 1$ and the coefficient of risk aversion to $\sigma = 2$. We set a steady-state G/Y ratio of 20% and we use standard interest rate rule parameters $\varphi_\pi = 1.5$ and $\varphi_y = 0.125$.

Regional government purchases follow independent AR(1) processes with quarterly persistence of 0.95. The housing demand shock follows an AR(1) with the same persistence. We set the scale of the portfolio holding cost to $\zeta = 10^{-4}$. We set the quarterly Calvo adjustment probability to 11% to target the point estimate of 0.030 of the inflation response to local government spending shocks reported by Nakamura and Steinsson (2014). We consider the robustness of our results to our parameter choices in Subsection VIII.F.

B. Complete Markets

As a starting point, we consider a complete-markets version of our model. In this version, the consumption response to house prices is a function of the current user cost only. We also start with a specification in which prices are fully rigid ($\chi = 1$), houses are produced entirely from land with no material inputs ($\alpha = 0$), and the land supply elasticities are fixed at their long-run values. The partial equilibrium housing wealth effect is particularly simple to compute in this case. Equating the marginal utility of consumption between regions yields

$$\frac{C_t - \psi \frac{L_t^{1+\nu}}{1+\nu}}{C_t^* - \psi \frac{L_t^*}{1+\nu}} = \left(\frac{p_t - \mathbb{E}_t[\omega_{t+1} p_{t+1} (1 - \delta)]}{p_t^* - \mathbb{E}_t[\omega_{t+1} p_{t+1}^* (1 - \delta)]} \right)^{\frac{(1-\alpha)(\sigma-1)}{\sigma}}, \quad (47)$$

where ω_{t+1} is the stochastic discount factor for payoffs at $t + 1$. The region with a higher user cost will consume more nondurables and less housing. With mean-reverting house price dynamics, a higher price of housing is

associated with a larger user cost, which induces a positive relationship between nondurable consumption and house prices.

Results for this version of our model are shown in the first column of table 2. The first four rows of this table report the components of the right-hand side of equation (46). We describe above how these are calculated from the Monte Carlo simulation. The first row reports the measured housing wealth effect ($d \log \hat{C} / d \log \hat{p}$) in this version of the model is 0.022. The second row reports the local fiscal multiplier ($d \hat{Y} / dG$) is 1.477. The third row reports the construction response ($d \log \hat{I} / d \log \hat{p}$), which in this version of the model we assume to be zero. Finally, the fourth row reports the income elasticity of house prices ($d \log \hat{p} / d \log \hat{Y}$) is small with complete markets (-0.004).

Using equation (46), we can combine these four responses from the model to calculate the partial equilibrium housing wealth effect implied by our simple formula. The fifth row in the table reports that for the complete-markets model this turns out to be 0.015. Row (6) then reports the actual partial equilibrium housing wealth effect for this version of the model, which also turns out to be 0.015. In other words, in this version of the model, our simple adjustment formula is very accurate. Because there is no construction response and the income elasticity of house prices is small, our formula boils down to dividing the measured housing wealth effect by the local fiscal multiplier. Finally, row (7) reports the magnitude of the error associated with the implied housing wealth effect relative to the error associated with the measured housing wealth effect defined as $| \text{Row 5} - \text{Row 6} | / | \text{Row 1} - \text{Row 6} |$.

Table 2
Monte Carlo Analysis of Housing Wealth Elasticity

	(i)	(ii)	(iii)	(iv)
Complete markets	✓			
Rigid prices	✓	✓		
Construction				✓
Long-run housing supply heterogeneity				✓
(1) Measured housing wealth effect	.022	.121	.151	.144
(2) Local fiscal multiplier	1.477	1.499	1.451	1.465
(3) Construction response	0	0	0	1.613
(4) Income elasticity of home prices	-.004	.032	.151	.167
(5) Implied partial equilibrium housing wealth effect	.015	.081	.104	.072
(6) Actual partial equilibrium housing wealth effect	.015	.063	.078	.041
(7) Relative error	.000	.306	.368	.295

C. *Incomplete Markets*

The second column of table 2 reports results for a version of our model with incomplete markets but maintains the assumption that prices are rigid and that no resources are used in constructing houses. With incomplete markets the determination of the consumption response to house prices is more complicated as it is affected by expectations of future incomes and user costs (see eq. [34]). In this case, the off-diagonal elements of the matrices in equation (45) become more important and it need not be the case that these dynamic relationships can be summarized by simple regressions. We find that the measured housing wealth effect is 0.121. Using equation (46) to adjust for local general equilibrium effects yields an implied partial equilibrium housing wealth effect of 0.081. However, the actual partial equilibrium housing wealth effect is a bit lower at 0.063. In this case, the adjustment implied by our formula goes in the right direction but does not go far enough. Here the relative error is about 1/3 as large after applying our formula.

D. *Incorporating Price Responses*

The assumption of rigid prices implies that both regions face the same real interest rate. When prices respond differentially in the two regions, real interest rates differ across regions. Suppose that the home region experiences a larger increase in activity. It will then also experience a larger increase in inflation, which reduces the real interest rate and further stimulates demand in the home region. However, the differential price response changes the real exchange rate, which reduces the demand for goods from the home region and increases the demand for goods from the foreign region. On one hand, these differential price responses affect the fiscal multiplier, so in principle, adjusting the measured housing wealth effect by the fiscal multiplier may fully account for these effects. On the other hand, the price responses further complicate the dynamics of the responses in ways that may not be fully captured by our approach.

The third column in table 2 allows for some degree of price flexibility. Specifically, we set the quarterly Calvo adjustment probability to 11% to target the response of inflation to local government spending shocks of 0.030 reported by Nakamura and Steinsson (2014). In this case, we find that the measured housing wealth effect rises to 0.151. The local fiscal multiplier is reduced due to the expenditure switching effect. The

combination of these two changes raises the implied housing wealth effect to 0.104. Even though real interest rates do not change in the partial equilibrium calculation, the partial equilibrium housing wealth effect depends on the particular dynamics of house prices that we feed into the calculation and the house price dynamics change as a result of price adjustments. As a result, the actual partial equilibrium housing wealth effect is somewhat larger in this case rising to 0.078. Again, our simple adjustment formula somewhat underestimates the needed adjustment. The relative error is 0.368 in this case.

E. Adding Construction

The fourth column of table 2 presents results for a version of the model in which we allow for resources to be used in the construction of housing; that is, we set α equal to its calibrated value of $\alpha = 0.38$. To generate a positive response of residential investment to changes in house prices, we also introduce the regime-switching dynamics for γ and γ^* described in Section VII. This version of the model yields a construction response of 1.613, which is somewhat higher than the value we estimate in the data (0.363). If we instead assume that the housing supply elasticities γ and γ^* differ in the short run, the model generates a construction response of -12 . The short-run/long-run distinction reverses the sign of the construction response even though the heterogeneous response of house prices still reflects differences in supply curves across regions.

The measured housing wealth effect in this case is 0.144. Taking account of the construction response and using the local fiscal multiplier to adjust for the local general equilibrium effect yields an implied housing wealth effect of 0.072. The actual partial equilibrium housing wealth effect is 0.041. As in the previous two cases, our simple adjustment somewhat understates the needed adjustment. The relative error is 0.295 in this case.

F. Robustness

Table 3 reports results for several variants of the model specification in column (iv) of table 2. We focus our robustness analysis on this model specification because it is the richest one in table 2 and the one that comes closest to the magnitudes of the measured housing wealth effect and construction response. Each column in table 3 reports results for a version of the model in which we vary one or two parameters and leave

Table 3
Monte Carlo Analysis of Housing Wealth Elasticity: Robustness

	(i)	(ii)	(iii)	(iv)
	$\rho_G = .9$	$\rho_\Omega = .9$	$\chi = .75$ $\varphi = .7$	$\sigma \rightarrow 1$
Measured housing wealth effect	.144	.167	.341	.262
Local fiscal multiplier	1.449	1.465	2.318	1.466
Construction response	1.613	1.613	1.613	1.613
Income elasticity of home prices	.082	.167	.208	.165
Implied partial equilibrium housing wealth effect	.073	.088	.104	.154
Actual partial equilibrium housing wealth effect	.041	.041	.074	.110
Relative error	.305	.372	.111	.287

the remaining parameters at their baseline values. Column (i) reduces the quarterly persistence of the government spending shocks from its baseline value of 0.95 to 0.90. This change has little effect on our results. Column (ii) reduces the persistence of the aggregate housing demand shock from the baseline of 0.95 to 0.90. This change slightly raises the measured housing wealth effect and implied partial equilibrium housing wealth effects, while leaving the actual housing wealth effect virtually unchanged. Overall, columns (i) and (ii) are reassuring that the exact details of the dynamics of the changes in house prices and changes in government spending are not crucial to the performance of our adjustment.

Column (iii) of table 3 considers a case with more price flexibility and less openness. Making prices more flexible leads to a smaller fiscal multiplier due to expenditure switching after a government purchases shock while making the economies less open raises the fiscal multiplier. The combination of parameters considered here is close to those used by Nakamura and Steinsson (2014). In this case, we see a large increase in the measured and actual housing wealth effect and fiscal multiplier relative to our baseline. Our simple adjustment still yields an implied housing wealth effect much closer to the true partial equilibrium effect with the relative error now falling to about 1/9 on account of the larger role of the fiscal multiplier.

Column (iv) considers a case where $\sigma \rightarrow 1$. In our baseline model with $\sigma = 2$, housing demand shocks (changes in Ω_t) raise the marginal utility of consumption due to the interaction in the utility function between consumption and housing. This leads to a strong direct response of consumption to Ω shocks. As Ω is an aggregate shock, this force affects both regions equally and is differenced out in the cross section, but it implies

a very strong response of aggregate consumption to Ω in the time series. This does not occur in the case with $\sigma \rightarrow 1$ as the utility function becomes additively separable between consumption and housing and the time-series housing wealth effect is more comparable in magnitude to the cross-sectional housing wealth effect. Despite the very different behavior of aggregate consumption with $\sigma \rightarrow 1$, the cross-sectional results are similar to our baseline.

G. Summary

The analysis we present in this section shows that the simple static formula we derive in Section III to adjust estimates of housing wealth effects for local general equilibrium effects tends to somewhat underestimate the needed adjustment. In our richest specification, column (iv) of table 2, the adjustment yields an estimate of the housing wealth effect that has an error that is 1/3 as large as that associated with the measured housing wealth effect. More fully accounting for the dynamics of the responses to house prices as in the matrix relationships discussed in Section VI may yield a more accurate estimate of the partial equilibrium effect. However, the much simpler approach of using our static formula seems to account for the bulk of the needed adjustment (roughly 2/3).

IX. Conclusion

Cross-regional empirical estimates have become part of the macroeconomist's toolkit, but the appropriate interpretation of these estimates can be difficult as they often blend together partial equilibrium responses with local general equilibrium effects. We argue that researchers can benefit from approaching cross-regional estimates as part of a system of simultaneous equations that also integrates other sources of evidence on the magnitude of local general equilibrium effects, such as evidence on the magnitude of fiscal multipliers.

This approach allows researchers to decompose the effects they estimate into the components arising from the direct partial equilibrium effect and those arising from local general equilibrium effects. We apply this methodology to analyzing regional estimates of housing wealth effects and show that an important part of the estimated regional effects likely arise from local general equilibrium effects. Gauging the relative importance of the partial and local general equilibrium components of

these effects is crucial to developing appropriate microfoundations for the housing wealth effect.

In our analysis of housing, we also address a recent critique by Davidoff (2016), who points out that if housing markets experience a common demand shock but move along different supply curves, then prices and quantities should be negatively correlated. The relationship in the data is the opposite. We show, however, that a model that allows for differences between short-run and long-run housing supply elasticities can address Davidoff's critique. In such a model, expectations about future capital gains imply that the common aggregate shock results in a larger change in current housing demand in the city with a less elastic long-run housing supply curve.

Appendix A

Derivation of Key Equations in Section III

Here we derive the key equations in Subsection III.D. The derivations of the corresponding equations in earlier subsections of Section III are special cases of this derivation. To get the results in Subsection III.C, set $H_w = 0$. To get the results in Subsection III.B, set $\Phi = 1$. Finally, to get the results in Subsection III.A, set $I_p = 0$.

We start by taking derivatives of the first two equations in Subsection III.C with respect to G —these equations represent the goods market equilibrium conditions in Subsection III.D as well as Subsection III.C. This yields

$$\begin{aligned} & \phi(C_p + I_p) \frac{dp}{dG} + \phi C_w \frac{dw}{dG} + (1 - \phi)(C_p + I_p) \frac{dp^*}{dG} + (1 - \phi)C_w \frac{dw^*}{dG} \\ & + (C_T + I_T) \frac{dT}{dG} + 1 = N_w \frac{dw}{dG}, \\ & (1 - \phi)(C_p + I_p) \frac{dp}{dG} \\ & + (1 - \phi)C_w \frac{dw}{dG} + \phi(C_p + I_p) \frac{dp^*}{dG} + \phi C_w \frac{dw^*}{dG} \\ & + (C_T + I_T) \frac{dT}{dG} = N_w \frac{dw^*}{dG}. \end{aligned}$$

Subtracting the second of these equations from the first and rearranging, we get

$$(N_w - \Phi C_w) \frac{d\hat{w}}{dG} - \Phi(C_p + I_p) \frac{d\hat{p}}{dG} = 1, \quad (\text{A1})$$

where $\Phi = 2\phi - 1$ and hatted variables denote cross-regional differences, for example, $\hat{p} = p - p^*$.

Next, we take derivatives of the first two equations in Subsection III.D with respect to G . This yields

$$\begin{aligned} H_w \frac{dw}{dG} + H_p \frac{dp}{dG} + H_T \frac{dT}{dG} &= 0, \\ H_w \frac{dw^*}{dG} + H_p \frac{dp^*}{dG} + H_t \frac{dT}{dG} &= 0. \end{aligned}$$

Subtracting the second of these equations from the first yields

$$H_w \frac{d\hat{w}}{dG} + H_p \frac{d\hat{p}}{dG} = 0. \tag{A2}$$

Similarly, we now take derivatives of the goods market and housing market equilibrium conditions with respect to s and then take cross-regional differences to arrive at

$$(N_w - \Phi C_w) \frac{d\hat{w}}{ds} - \Phi(C_p + I_p) \frac{d\hat{p}}{ds} = 0, \tag{A3}$$

$$H_w \frac{d\hat{w}}{ds} = H_p \frac{d\hat{p}}{ds} = -H_s. \tag{A4}$$

We can now solve equations (A1)–(A4) for the relative response of house prices and wages to the two exogenous shocks:

$$\begin{pmatrix} \frac{d\hat{w}}{dG} & \frac{d\hat{w}}{ds} \\ \frac{d\hat{p}}{dG} & \frac{d\hat{p}}{ds} \end{pmatrix} = M \begin{pmatrix} 1 & -\Phi(C_p + I_p) \frac{H_s}{H_p} \\ -\frac{H_w}{H_p} & -(N_w - \Phi C_w) \frac{H_s}{H_p} \end{pmatrix}, \tag{A5}$$

where

$$M = \left(N_w - \Phi C_w + \Phi(C_p + I_p) \frac{H_w}{H_p} \right)^{-1}.$$

Define total home expenditures as $E = C(w, p, T) + I(p, T) + G$. The derivative of total home expenditures with respect to s is

$$\frac{dE}{ds} = (C_p + I_p) \frac{dp}{ds} + C_w \frac{dw}{ds} + C_T \frac{dT}{ds}.$$

Taking the difference between this equation and the foreign version of this equation yields

$$\frac{d\hat{E}}{ds} = (C_p + I_p) \frac{d\hat{p}}{ds} + C_w \frac{d\hat{w}}{ds}.$$

Using equation (A5) we then get that

$$\frac{d\hat{E}}{ds} = -(C_p + I_p)N_wM \frac{H_s}{H_p}$$

and

$$\frac{d\hat{E}}{d\hat{p}} = \frac{d\hat{E}/ds}{d\hat{p}/ds} = \frac{C_p + I_p}{1 - \Phi C_Y}, \quad (\text{A6})$$

where, again, $C_Y \equiv C_w/N_w$. Next, we differentiate home output and foreign output with respect to G and take the difference to get

$$\frac{d\hat{Y}}{dG} = \Phi(C_p + I_p) \frac{d\hat{p}}{dG} + \Phi C_w \frac{d\hat{w}}{dG} + 1.$$

Using equation (A5) we then get that

$$\frac{d\hat{Y}}{dG} = N_wM. \quad (\text{A7})$$

Taking the ratio of equations (A6) and (A7) yields

$$\frac{d\hat{E}/d\hat{p}}{d\hat{Y}/d\hat{p}} = (C_p + I_p) \frac{1 - \Phi C_Y + \Phi \frac{C_p + I_p}{N_w} \frac{H_w}{H_p}}{1 - \Phi C_Y}.$$

Manipulation of this equation yields

$$\frac{d\hat{E}/d\hat{p}}{d\hat{Y}/d\hat{p}} = (C_p + I_p) \left(1 - \Phi \frac{d\hat{E}/d\hat{p}}{d\hat{Y}/d\hat{p}} \frac{d\hat{p}}{dG} \right),$$

which can be manipulated further to yield equation (23) in the main text.

Appendix B

Additional Examples of Identification in Systems of Equations

This appendix presents three example systems of equations and demonstrates that in each of them the government spending shock is needed to identify the coefficient of interest.

B.1. Example from Subsection III.D

We start by showing that the argument regarding identification in Section IV applies to the richer setting considered in Subsection III.D.

Linearize the equations of the model and take the difference across regions to yield

$$\begin{aligned}\hat{C} &= C_p \hat{p} + C_Y \hat{Y} \\ \hat{I} &= I_p \hat{p} \\ \hat{Y} &= \Phi(\hat{C} + \hat{I}) + \hat{G} \\ 0 &= H_Y \hat{Y} + H_p \hat{p} + H_s \hat{s}\end{aligned}$$

where we have substituted out for wages using $\hat{w} = N_w^{-1} \hat{Y}$ and defined $C_Y = C_w N_w^{-1}$ and $H_Y = H_w N_w^{-1}$ as in Section IV. There are four endogenous variables $(\hat{C}, \hat{I}, \hat{Y}, \hat{p})$ and two exogenous variables (\hat{G}, \hat{s}) . We will write the system as $(\hat{C}, \hat{I}, \hat{Y}, \hat{p})\Gamma = (\hat{G}, \hat{s})\Delta$, where the coefficient matrices are

$$\Gamma = \begin{pmatrix} -1 & 0 & \Phi & 0 \\ 0 & -1 & \Phi & 0 \\ C_Y & 0 & -1 & H_Y \\ C_p & I_p & 0 & H_p \end{pmatrix} \quad \Delta = \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -H_s \end{pmatrix}.$$

The restrictions on the first equation are

$$R_\Gamma = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad R_\Delta = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

such that $(R_\Gamma \Gamma + R_\Delta \Delta)e_1 = 0$. The three restrictions are exclusion restrictions that exclude I , G , and s , respectively, from appearing directly in the consumption function. To check the rank condition, form $\mathcal{R} = R_\Gamma \Gamma + R_\Delta \Delta$:

$$\mathcal{R} = \begin{pmatrix} 0 & -1 & \Phi & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -H_s \end{pmatrix},$$

which has rank 3 provided that $H_s \neq 0$ so the rank condition is satisfied. The first equation has a natural normalization as the coefficient on C is

-1. If we drop G from the system, we lose one restriction and the order condition fails.

B.2. Local Credit Supply Example

The spirit of this example is that credit markets for investment spending are (partially) segmented across space perhaps due to banking relationships. The credit supply shock could represent a shock to the banks serving a given region that affects the supply of credit to the region leading to a change in local interest rates. We first argue that fiscal shocks are needed to identify the investment demand curve, and then we solve for the coefficients from the investment demand curve as a function of the reduced form estimates. Here, we focus on a single regional economy and the variables can be interpreted as deviations from national averages that are absorbed by time fixed effects.

The system is

$$Y = C + I + G$$

$$I = I_r r + I_y Y$$

$$C = C_y Y$$

$$b = b_l I$$

$$b = b_s s + b_r r$$

where the endogenous variables (Y, I, C, b, r) correspond to output, investment, consumption, credit quantity, and interest rates. The exogenous variables (s, G) correspond to a credit supply shifter and government purchases. In order, the equations are the resource constraint, an investment demand curve, a consumption function, credit demand, and credit supply.

The consumption and investment equations are both identified in that they satisfy the rank condition and have a natural normalization. Without the shock to G , the investment demand curve is not identified as the order condition fails: there are three excluded variables and five endogenous variables. Intuitively, the problem is that the reduced form shows that investment reacts after a shock to s , but we do not know if I_r is small and I_y is large or vice versa. The fiscal shocks provide information about the magnitude of I_y .

With the reduced form estimates in hand, we can solve for I_r as follows. Define the following IV estimates:

$$\begin{aligned} \frac{dY}{dG} &= \frac{-b_i I_r + b_r}{b_i C_y I_r - b_i I_r - b_r C_y - b_r I_y + b_r} \\ \frac{dr}{dY} &\equiv \frac{dr/dG}{dY/dG} = -\frac{b_i I_y}{b_i I_r - b_r} \\ \frac{dC}{dY} &\equiv \frac{dC/dG}{dY/dG} = C_y \\ \frac{dI}{dr} &\equiv \frac{dI/ds}{dr/ds} = \frac{I_r(C_y - 1)}{C_y + I_y - 1}. \end{aligned}$$

The reduced forms are expressions involving structural parameters given by (ratios of) the appropriate elements of $\Delta\Gamma^{-1}$. Manipulating these expressions yields

$$\begin{aligned} I_r &= \frac{dI}{dr} \\ M &\equiv \frac{dY}{dG} \left(1 - \frac{dr}{dY} \frac{dI}{dr} - \frac{dC}{dY} \right). \end{aligned} \tag{B1}$$

Here M is our “adjusted” fiscal multiplier as we need to adjust the fiscal multiplier in a similar manner to what we did in Section IV. The IV regression dI/dr measures how much investment responds to a given change in local interest rates. When we measure the local fiscal multiplier, we capture the channels operating through I_y , C_y , and the effect on local interest rates as credit demand rises. We need to remove the effect on interest rates just as we removed the effect of fiscal shocks on house prices in Section IV. We also need to adjust for the role of consumption in the fiscal multiplier. Specifically, we are interested in how a unit movement in investment is amplified. The direct effect, first round, second round, and so on are given by

$$\begin{aligned} \text{Output response} & \quad 1, I_y + C_y, (I_y + C_y)^2, \dots \\ \text{Investment response} & \quad 1, I_y, I_y(C_y + I_y), \dots \end{aligned}$$

The investment response is therefore $1 + [I_y/(1 - C_y - I_y)] = M$. It is then straightforward to solve for I_y . Rearranging M yields

$$I_y = \frac{(M - 1) \left(1 - \frac{dC}{dY} \right)}{M}.$$

B.3. Foreign Demand Example

Consider the system

$$Y = C + G + X_s s$$

$$C = C_Y Y$$

where there are two endogenous variables Y and C and two exogenous variables G and s where s is an instrument for foreign demand. X_s is an unknown coefficient that gives the direct output effect of a change in s . For example, s could be China exposure, but we do not know how China exposure translates to a change in local demand because there may be more or less scope for local production to transition to other goods for which demand is still strong after the China shock. The goal is to estimate X_s . Substitute the second equation into the first,

$$Y = MG + MX_s s,$$

where $M \equiv 1/(1 - C_Y)$. In the absence of the fiscal shock, we only identify the product MX_s . With the fiscal shock, we can determine M . One way to relate this to the discussion of identification in Section IV is to write the equation with the coefficient of interest on s ,

$$(C_Y - 1)Y + G + X_s s = 0,$$

and note that without G we lack a normalization.

Endnotes

Authors' email addresses: Steinsson (jsteinsson@berkeley.edu), Nakamura (enakamura@berkeley.edu). An earlier version of this paper previously circulated as an appendix to "Housing Wealth Effects: The Long View" by the same authors. We thank Joao Fonseca Rodrigues for excellent research assistance. We thank Adrien Auclert, Gabe Chodorow-Reich, Erik Hurst, and Valerie Ramey for helpful comments and discussions. Guren thanks the National Science Foundation (grant SES-1623801) and the Boston University Center for Finance, Law, and Policy. Nakamura thanks the National Science Foundation (grant SES-1056107). Nakamura and Steinsson thank the Alfred P. Sloan Foundation for financial support. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/nber-macroeconomics-annual-2020-volume-35/what-do-we-learn-cross-regional-empirical-estimates-macroeconomics>.

1. Prominent examples include Autor, Dorn, and Hanson (2013), Chodorow-Reich (2014), Mian and Sufi (2014), Nakamura and Steinsson (2014), and Martin and Philippon (2017). See Nakamura and Steinsson (2018) and Chodorow-Reich (2019) for further discussion of this literature.

2. Davidoff (2016) shows that the change in housing units from 1980 to 2010 was negatively correlated with Saiz's housing supply elasticity across cities. The long-horizon

quantity response is less worrisome than the business cycle quantity response because differential demand trends across cities can be absorbed by city-level fixed effects in a panel specification (Guren et al. 2020).

3. Wolf describes his results in terms of addition and subtraction of impulse response functions. To understand the connection, a simplified version of our result is $C_{PE} = F^{-1}E$, where F is the matrix in which the (i, j) element gives the response of output at horizon i to a change in government spending at horizon j and the column vector E is the measured housing wealth effect impulse response. Wolf expresses his result as $C_{PE} + (F - I)C_{PE} = E$ where $(F - I)C_{PE}$ is the private consumption response to the fiscal shock that has the same dynamic profile as the partial equilibrium housing wealth effect.

4. Groundwork for these papers was provided by Auclert and Rognlie (2020) who show that the general equilibrium effects of a shock to consumption can be separated into a partial equilibrium path of consumption and a general equilibrium multiplier matrix that does not depend on the shock that perturbs consumption.

5. Here we use the demand side of the economy to form the quantity responses (eqs. [9] and [10]). As we are analyzing equilibrium changes in quantities, we can use either the demand response or supply response to the equilibrium prices and arrive at the same answer.

6. Nakamura and Steinsson find larger multipliers in regional data than in state data. As the analysis of the housing wealth effect is undertaken at the city (CBSA) level, it may be appropriate to use a fiscal multiplier somewhat below 1.5.

7. We construct this estimate as $(H/C)(C/I)$. Between 1985 and 2016, the average ratio of H/C was 2.17, where H is measured as the market value of owner-occupied real estate from the Flow of Funds and C is measured as total personal consumption expenditures (PCE) less PCE on housing services and utilities from the National Income and Product Accounts. Over this same period, I/C was 0.077, where I is residential investment from the National Income and Product Accounts.

8. Think of this as an IV estimate. For example, in the current model $d\hat{p}/d\hat{Y}$ can be estimated as a ratio of responses to exogenous variation in government spending: $(d\hat{p}/dG)/(d\hat{Y}/dG)$.

9. Nakamura and Steinsson estimate the local fiscal multiplier based on state-level data, whereas Guren et al. estimate the housing wealth effect on consumption at the CBSA level. Incorporating trade linkages into our analysis is accounting for the fact that the local fiscal multiplier is attenuated by trade linkages so it makes sense to use a value of φ consistent with the geographic unit used to estimate the fiscal multiplier.

10. In addition to changing the marginal rate of substitution between housing and nondurables, a shock to Ω_t also affects the marginal utility of consumption as the extra demand for housing is not fully satisfied in equilibrium. This can lead to a strong consumption response to Ω_t in the aggregate time series, but it equally affects both regions and therefore does not influence our cross-sectional analysis.

11. Including the term $\phi^{-\varphi}(1 - \phi)^{-(1-\varphi)}$ in the definition of the bundle simplifies the expression for the price index.

12. This case applies to the steady-state, perfect foresight transitions and first-order accurate solutions to stochastic economies.

13. Appendix E describes this negative relationship between quantities and prices in econometric terms and shows that our adjustment formula from Section III remains valid with heterogeneous housing supply curves.

References

- Aladangady, A. 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107:3415–46.
- Auclert, A., and M. Rognlie. 2020. "Inequality and Aggregate Demand." Working paper, Stanford University.
- Auclert, A., M. Rognlie, and L. Straub. 2018. "The Intertemporal Keynesian Cross." Working paper, Stanford University.

- Autor, D. H., D. Dorn, and G. H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103:2121–68.
- Beraja, M., E. Hurst, and J. Ospina. 2019. "The Aggregate Implications of Regional Business Cycles." *Econometrica* 87:1789–833.
- Calvo, G. A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12:383–98.
- Campbell, J. Y., and J. F. Cocco. 2007. "How Do House Prices Affect Consumption?" *Journal of Monetary Economics* 54:591–621.
- Chodorow-Reich, G. 2014. "The Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008–9 Financial Crisis." *Quarterly Journal of Economics* 129:1–59.
- . 2019. "Geographic Cross-Sectional Fiscal Multipliers: What Have We Learned?" *American Economic Journal: Policy* 11:1–34.
- Chodorow-Reich, G., P. T. Nenov, and A. Simsek. 2020. "Stock Market Wealth and the Real Economy: A Local Labor Market Approach." *American Economic Review*, forthcoming.
- Davidoff, T. 2016. "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated with Many Demand Factors." *Critical Finance Review* 5:177–206.
- Davis, M. A., and J. Heathcote. 2005. "Housing and the Business Cycle." *International Economic Review* 46:751–84.
- Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh. 2017. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium." *Journal of Political Economy* 125:140–223.
- Greenwood, J., Z. Hercowitz, and G. W. Huffman. 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review* 78 (3): 402–17.
- Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson. 2020. "Housing Wealth Effects: The Long View." *Review of Economic Studies*, forthcoming.
- Herreno, J. 2020. "The Aggregate Effects of Bank Lending Cuts." Working paper, Columbia University.
- Kaplan, G., K. Mitman, and G. L. Violante. 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy* 128 (9): 3285–3345.
- Lamont, O., and J. Stein. 1999. "Leverage and House-Price Dynamics in US Cities." *RAND Journal of Economics* 30:498–514.
- Martin, P., and T. Philippon. 2017. "Inspecting the Mechanism: Leverage and the Great Recession in the Eurozone." *American Economic Review* 107:1904–37.
- Mian, A., K. Rao, and A. Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128:1687–726.
- Mian, A., and A. Sufi. 2014. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82:2197–223.
- Nakamura, E., and J. Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *American Economic Review* 104:753–92.
- . 2018. "Identification in Macroeconomics." *Journal of Economic Perspectives* 32:59–86.
- Nathanson, C. G., and E. Zwick. 2018. "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market." *Journal of Finance* 73:2587–633.
- Saiz, A. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125:1253–296.

- Schmitt-Grohé, S., and M. Uribe. 2003. "Closing Small Open Economy Models." *Journal of International Economics* 61:163–85.
- Wolf, C. 2019a. "The Missing Intercept: A Demand Equivalence Approach." Working paper, Princeton University.
- . 2019b. "The Missing Intercept in Cross-Regional Regressions." Working paper, Princeton University.
- Wooldridge, J. M. 2010. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.