The Missing Intercept in Cross-Regional Regressions*

Christian K. Wolf
Princeton University

October 23, 2019

Abstract: I generalize the results in Wolf (2019) to estimate the aggregate effects of generic consumption demand shifters on the basis of cross-*regional* micro regressions. Cross-regional regressions do not estimate macro counterfactuals since they (i) are contaminated by local general equilibrium effects and (ii) difference out aggregate general equilibrium feedback. I show that, in a broad class of business-cycle models, evidence on regional government spending shocks is informative about (i), and aggregate fiscal multipliers can pin down (ii). I apply my results to study household deleveraging due to a credit tightening, and find that the large cross-regional responses estimated in Mian et al. (2013) map into significant declines of *aggregate* consumption and output. Any calibrated structural model that reaches a different conclusion either lies outside of my model class or is inconsistent with evidence on regional and aggregate fiscal multipliers.

Keywords: general equilibrium transmission, consumption, fiscal policy, credit tightening, household deleveraging. *JEL codes:* E2, E3, H3.

*Email: `ckwolf@princeton.edu`. I received helpful comments from Ben Moll, Chris Tonetti, Gianluca Violante and Tom Winberry. I also acknowledge support from the Alfred P. Sloan Foundation and the Macro Financial Modeling Project, as well as the Washington Center for Equitable Growth.

1 Introduction

Much recent work in applied macroeconomics tries to learn about the effects of policy stimulus and other shocks through variation in shock exposure across households, firms, or regions. In Wolf (2019), I argue that aggregate fiscal multipliers are informative about the "missing general equilibrium intercept" of many popular micro regressions run at the level of individual firms and households (e.g. Johnson et al., 2006; Zwick & Mahon, 2017). However, some of the most convincing evidence on the response of consumption to macro shocks is at the *regional* level (Mian et al., 2013; Guren et al., 2018), not the level of individual households. Mapping cross-regional estimates into aggregate counterfactuals is likely to be more challenging, for two reasons. First, cross-regional estimates are presumably contaminated by local general equilibrium effects. Second, it is not clear to what extent the potentially heterogeneous incidence of shocks across regions matters for macro outcomes.

In this note I generalize the results of Guren et al. (2019) and Wolf (2019) to characterize the connection between cross-regional regression estimates and macro counterfactuals for a general family of consumption demand shocks. My analysis builds on a rich multi-region macro model. Regions are heterogeneously exposed to a generic consumption demand shock; this heterogeneous exposure is exploited to estimate cross-regional regressions of relative consumption responses on relative shock exposure. I use the model to establish the following. First, cross-regional regression estimates invariably pick up *local* general equilibrium effects. Under relatively mild assumptions, the response of regional consumption to regional government spending shocks is informative about the strength of these local feedback effects. Second, for macro counterfactuals, it does not matter by how much private consumption spending changed in each individual region, just how much it changed *on average* across regions. With regional heterogeneity thus irrelevant for aggregate general equilibrium effects, the results of Wolf (2019) apply; in particular, the response of aggregate consumption to fiscal policy shocks is informative about the strength of aggregate general equilibrium feedback. Applied to household deleveraging due to tighter borrowing conditions, my results suggest that the large cross-regional estimates of Mian et al. (2013) also correspond to significant declines in *aggregate* consumption and output.

I interpret cross-regional household spending regressions through the lens of a rich quantitative business-cycle model. Similar to Beraja et al. (2019), the model consists of a unit continuum of interdependent regions, subject to common and region-specific shocks. In this model, I study the propagation of generic "consumption demand" shifters. The model is rich enough to allow for almost all frictions familiar from standard quantitative business-cycle analysis, including uninsurable household income risk (Kaplan et al., 2018), real and financial investment frictions (Khan & Thomas, 2013), and nominal rigidities (Christiano et al., 2005; Smets & Wouters, 2007). I prove that, in this model, the *estimand* of a regression of regional consumption on measures of regional exposure to the demand shock is equal to the impulse response of regional consumption to a purely region-specific shock. As such, it differs from aggregate impulse responses for two reasons: (i) the presence of local general equilibrium adjustment and (ii) the absence of aggregate general equilibrium feedback effects.

I first learn about the strength of local general equilibrium adjustment. I show that, to strip out this feedback, it suffices to take cross-regional micro regression estimates and *subtract* the response of local consumption to an appropriately scaled local government spending shock. The logic is identical to that in Wolf (2019): If private and public spending shocks share identical local general equilibrium propagation, then the response of local consumption to a local fiscal shock can simultaneously proxy for local general equilibrium effects associated with changes in private spending. Importantly, such "demand equivalence" obtains under mild and empirically plausible restrictions on my benchmark class of models. A related correction is discussed in Guren et al. (2019), who *divide* local spending responses by measures of local fiscal output multipliers. Relative to them, my approach affords somewhat greater generality – in particular, it remains valid without full home bias and with investment –, but it does so only at the cost of higher informational requirements: I require evidence on local consumption responses, rather than the more readily estimable output multipliers.

With local general equilibrium effects stripped out, it remains to identify the aggregate general equilibrium feedback associated with a collection of regional demand shocks. A key building block result is the following: To first order, it does not matter for macro aggregates *where* spending demand changed, just how much it changed overall.[1] This irrelevance result relies on symmetry in tastes and technologies across regions, but does not require regions to be vanishingly small (or more generally, equal in size). The results of Wolf (2019) thus apply: Under mild assumptions, the response of aggregate consumption to a change in aggregate fiscal spending proxies for the general equilibrium feedback of the regional demand shocks.

I apply my identification results to study the aggregate effects of changes in household borrowing conditions. Through cross-regional regressions, Mian et al. (2013) estimate that contractions in borrowing capacity during the Great Recession, mainly driven by falling house

[1] The geographical incidence of spending shocks of course invariably matters for relative prices and quantities, but for aggregates it turns out to be sufficient to look at cross-regional averages.

prices, account for substantial cross-regional differences in consumption. The standard interpretation of their estimates – which simply ignores any possible regional or aggregate general equilibrium feedback – suggests a drop in *aggregate* consumption expenditure of around 3 per cent. There is, however, little consensus on the differenced-out general equilibrium effects. For example, in the models of Guerrieri & Lorenzoni (2017) and Jones et al. (2018), partial equilibrium spending contractions are largely reversed in general equilibrium.

My identification results instead suggest that the cross-regional estimates of Mian et al. map into significant declines in aggregate consumption and output. I begin with a simple back-of-the-envelope calculation. First, local consumption is usually estimated to moderately increase after a rise in regional government spending (Dupor et al., 2018). This suggests that the cross-regional spending response of 3 per cent is only a mild over-estimate of the underlying aggregate partial equilibrium drop in consumption. Second, given aggregate government spending multipliers of around 1 (Ramey, 2018), this implied partial equilibrium drop in demand should translate roughly one-for-one into declines in aggregate consumption and output. I then move on to compute full general equilibrium counterfactuals in a rich structural model, disciplined to be jointly consistent with (i) estimates of cross-regional spending responses and (ii) evidence on regional and aggregate government spending multipliers. Consistent with the static back-of-the-envelope calculation, the model predicts that cross-regional regression estimates are only a slight over-estimate of the aggregate consumption decline.[2] By my identification results, any structural analysis inconsistent with this conclusion either breaks the demand equivalence underlying my identification results or is inconsistent with empirical evidence on local and aggregate fiscal multipliers.

LITERATURE. This note complements Wolf (2019) in that it provides a structural framework for interpreting and aggregating cross-*regional* estimates of (consumption) spending responses to macro shocks. Influential previous empirical work has produced cross-regional estimates of spending effects out of changes in housing wealth (Mian et al., 2013; Mian & Sufi, 2014; Guren et al., 2018), informing the "household deleveraging" view of the Great Recession. Similarly, quasi-random variation at the regional level has also allowed researchers to produce credible estimates of the cross-region effects of various kinds of policy stimulus, notably the cash-for-clunkers program (Mian & Sufi, 2012) and consumer bankruptcy pro-

[2]The drop in consumption does *not* rely on a binding zero lower bound (ZLB) constraint. Evidence on whether the propagation of macro shocks is affected by the ZLB is mixed (Ramey & Zubairy, 2018; Debortoli et al., 2019); to the extent that, as predicted by standard theory, a binding ZLB implies further amplification, my results are best regarded as a *lower bound* on true macro counterfactuals.

tection (Auclert et al., 2019). My results apply to all of these shocks to consumer spending. The overarching objective of my analysis – to map region-level regression estimates into aggregate counterfactuals – is also shared in much of the recent trade literature, notably Adao et al. (2019). In both trade and macro models, cross-regional regression estimands are not interpretable as aggregate counterfactuals due to *spatial linkages* – integrated labor and product markets in their case, home bias and macro policy in mine. They directly estimate these linkages in a particular structural model, while I indirectly proxy for their effects through regional and macro-level sufficient statistics.

The proposed mapping of cross-regional spending effects into aggregate counterfactuals requires reliable estimates of the transmission of local and aggregate government spending shocks. One of the first papers to estimate the local consumption multipliers necessary to strip out local general equilibrium effects is Dupor et al. (2018). Under more stringent assumptions, local *output* multipliers, as estimated in Nakamura & Steinsson (2014), can serve the same purpose (Guren et al., 2019). The much larger literature on aggregate government spending multipliers is reviewed, for example, in Hall (2009) or Ramey (2018).

Finally, my main application relates to the "household deleveraging" view of the Great Recession, proposed in Mian et al. (2013) and Mian & Sufi (2014, 2015). While it is firmly established that counties with greater house price declines saw greater declines in local spending, there is conflicting evidence on the extent to which the cross-regional regression estimates actually map into significant aggregate consumption and output declines. Indeed, most notable theoretical contributions to the literature (Eggertsson & Krugman, 2012; Guerrieri & Lorenzoni, 2017; Jones et al., 2018) find that aggregate outcomes depend sensitively on the general equilibrium model closure, particularly the presence of a binding ZLB constraint. My quantitative analysis – which implicitly embeds "normal-time" monetary feedback – instead suggests that cross-regional spending declines of a magnitude as documented by Mian et al. (2013) should translate into large aggregate consumption and output losses. These findings also rationalize the similarity between cross-regional regression estimates and model-based counterfactuals documented in Beraja et al. (2019, Figure 4), at least at short horizons.

Outline. Section 2 uses a general class of macro models with regional heterogeneity to (i) characterize the estimand of cross-regional regressions and (ii) connect cross-regional regression estimands and macro counterfactuals. Section 3 applies these insights to study the aggregate effects of a household credit tightening. Section 4 concludes, and supplementary details as well as all proofs are relegated to several appendices.

2 Regional regressions in macro models

This section develops the main theoretical results on the interpretation and aggregation of cross-regional regression estimates. Section 2.1 introduces the general class of macro models with regional heterogeneity that underlies my arguments. Section 2.2 uses the model to offer a structural interpretation of cross-regional regression estimands, and to identify measurable sufficient statistics that connect the regression estimates with the desired macro counterfactuals. Finally Section 2.3 discusses the generality and robustness of my identification results, and briefly considers various extensions.

2.1 A structural model framework

I consider a model economy composed of a unit continuum $k \in [0,1]$ of ex-ante identical islands. Each island is populated by households and firms; fiscal and monetary policy for all regions are set by common fiscal and monetary authorities. All islands are subject to idiosyncratic as well as aggregate structural shocks. Throughout, I study perfect foresight transition dynamics, which to first order are identical to the familiar first-order perturbation solution of macroeconomic models with aggregate risk.[3] Anticipating my identification results and empirical application, I study transition paths after changes in (local and aggregate) household borrowing conditions as well as government spending.

The environment merges three modeling traditions in business-cycle macroeconomics. First, I allow for essentially all of the frictions studied in canonical (representative-agent) New Keynesian DSGE models (Smets & Wouters, 2007; Justiniano et al., 2010). Second, households face incomplete insurance markets (Kaplan et al., 2018), and capital re-allocation across firms is subject to various real and financial constraints (Khan & Thomas, 2013). Third, the modeling of regional heterogeneity itself closely follows the earlier contributions of Gali & Monacelli (2005), Nakamura & Steinsson (2014) and Beraja et al. (2019). Since the model is a combination of several familiar ingredients, I only sketch the individual blocks here, with most emphasis on the household and fiscal sides. A detailed characterization of agents' problems and the formal equilibrium definition are relegated to Appendix A.1.

Notation. The realization of a variable x at time t along the perfect foresight transition path will be denoted x_t, and the entire time path will be denoted $\mathbf{x} = \{x_t\}_{t=0}^{\infty}$. Hats denote

[3] See for example Fernández-Villaverde et al. (2016), Boppart et al. (2018) or Wolf (2019).

deviations from the deterministic steady state, bars denote steady-state values, and tildes indicate region-specific deviations from aggregates, so $\tilde{x}_{k,t} \equiv x_{k,t} - x_t$.

I study the transmission of two structural shocks $s \in \{b, g\}$ – changes in household borrowing conditions b and in government spending g. The shock faced by region k, $\varepsilon_{s,k,t}$, has an aggregate component $\varepsilon_{s,t}$ and a purely regional component $\xi_{s,k,t}$, with $\int_0^1 \xi_{s,k,t} dk = 0$ at all times. I write the aggregate path for shock s as $\boldsymbol{\varepsilon}_s$, and use subscripts $\boldsymbol{\varepsilon}$ for transitions after a path $\boldsymbol{\varepsilon} \equiv (\boldsymbol{\varepsilon}_b', \boldsymbol{\varepsilon}_g')'$. I reserve the simpler s subscripts for one-time single aggregate shocks – that is, shock paths with $\varepsilon_{s,0} = 1$ and $\varepsilon_{u,\tau} = 0$ for $(u, \tau) \neq (s, 0)$. Analogously, I use s_k subscripts for transition paths following one-time region-specific shocks $\xi_{s,k,0} = 1$.

Households. Each region k is populated by a unit continuum of households $i \in [0, 1]$. Households consume the consumption goods of all regions k, but have a preference for the local good, governed by the home bias parameter ϕ. Specifically, the index of consumption goods for a generic household i in region k satisfies

$$c_{i,k,t} \equiv \left[\phi^{\frac{1}{\eta_r}} (c_{i,k,t}^H)^{\frac{\eta_r - 1}{\eta_r}} + (1 - \phi)^{\frac{1}{\eta_r}} (c_{i,k,t}^F)^{\frac{\eta_r - 1}{\eta_r}} \right]^{\frac{\eta_r}{\eta_r - 1}}$$

$c_{i,k,t}^H$ is her consumption of the local good, $c_{i,k,t}^F$ is an index of her consumption of all foreign goods, and η_r governs the substitutability between local and foreign goods. The foreign consumption index is

$$c_{i,k,t}^F \equiv \left[\int_0^1 (c_{i,k,t}^{k'})^{\frac{\eta_f - 1}{\eta_f}} dk' \right]^{\frac{\eta_f}{\eta_f - 1}}$$

where η_f governs the substitutability between different foreign goods. I choose as my numeraire the price index for the foreign bundle, which is identical for all households i and all regions k. Letting $p_{k,t}$ denote the *real* relative price of the consumption good of region k, we thus have $1 = \left(\int_0^1 p_{k,t}^{1-\eta_f} dk \right)^{\frac{1}{1-\eta_f}}$. The (common) overall price index for all households in region k is then

$$q_{k,t} \equiv \left[\phi p_{k,t}^{1-\eta_r} + (1 - \phi) \right]^{\frac{1}{1-\eta_r}}$$

Given the real price $q_{k,t}$ of the consumption index for households in region k, it is straightforward to state the consumption-savings and labor supply problems of all households $i \in [0, 1]$. Households have preferences over the consumption index $c_{i,k,t}$ and labor supply $\ell_{i,k,t}$, and are subject to idiosyncratic productivity risk $e_{i,k,t}$. They can self-insure through investments in liquid bonds $b_{i,k,t}^h$ at nominal rate i_t^b; borrowing in the liquid asset is subject to an additional penalty $\kappa_{k,t}^b$ and only possible up to some maximum $\underline{b}_{k,t}$. Households earn income from labor, $w_{k,t}e_{i,k,t}\ell_{i,k,t}$, and receive a share of dividends, $d_{i,k,t}$. Labor earnings are taxed by the government, but households also receive lump-sum rebates $\tau_{i,k,t}$. Given a path of supplied labor $\boldsymbol{\ell}^h_{i,k}$, the consumption-savings problem of household i is then as follows:[4]

$$\max_{\{c_{i,k,t},b^h_{i,k,t}\}} \quad \mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t u(c_{i,k,t},\ell_{i,k,t})\right] \tag{1}$$

such that

$$q_{k,t}c_{i,k,t} + b^h_{i,k,t} = (1-\tau_\ell)w_{k,t}e_{i,k,t}\ell^h_{i,k,t} + \frac{1+i^b_{t-1}+\kappa^b_{k,t}\mathbf{1}_{b^h_{i,t-1}<0}}{1+\pi_t}b^h_{i,k,t-1} + \tau_{i,k,t} + d_{i,k,t}$$

and

$$b^h_{i,k,t} \geq \underline{b}_{k,t}$$

as well as a (stochastic) law of motion for household-specific productivity, with $\int_0^1 e_{i,k,t} = 1$ at all times t and for all regions k. Anticipating my main application to household deleveraging, and similar to Guerrieri & Lorenzoni (2017), I allow the tightness of borrowing constraints $\underline{\mathbf{b}}_k$ as well as the borrowing wedge $\boldsymbol{\kappa}^b_k$ to follow (region-specific) time-varying paths, indexed by a shock path $\boldsymbol{\varepsilon}_{b,k}$. For simplicity I abstract from explicit modeling of a housing sector; the model can however equivalently be interpreted as one with fully rigid individual housing, exogenous house prices, and borrowing conditions mechanically tied to housing values.

Finally, household labor supply is intermediated by price-setting labor unions; details on their problem are relegated to Appendix A.1. Aggregating across households, we obtain regional consumption aggregates $(\mathbf{c}_k, \mathbf{c}^H_k, \mathbf{c}^{k'}_k)$, savings $\mathbf{b}_k$ and labor supply $\boldsymbol{\ell}^h_k$.

Firms. The region-k local good is produced by a region-specific corporate sector, consisting of perfectly competitive intermediate goods producers, monopolistically competitive retailers with nominal price rigidities, and a final goods aggregator. Intermediate goods producers $j \in [0,1]$ accumulate capital, hire local labor, issue debt, and sell their composite intermediate good, all possibly subject to a rich set of real and financial frictions. Retailers purchase the intermediate good, costlessly differentiate, set prices, and sell the differentiated good to the aggregator. The aggregated local final good is used for home and foreign consumption, local government consumption, and local investment, and the overall dividends of intermediate

[4] To streamline the presentation, I abstract from household habit formation and savings in illiquid assets. These extensions would work exactly as in the more general model of Wolf (2019).

goods producers and retailers are returned to households in region k.[5]

Government. The monetary authority sets the nominal rate on bonds, i^b, in accordance with a conventional (Taylor) rule. The fiscal authority decides how much to consume of each regional final good, $g_{k,t}$. Overall fiscal consumption $g_t \equiv \int_0^1 p_{k,t} g_{k,t} dk$ and total lump-sum transfers $\tau_t \equiv \int_0^1 \int_0^1 \tau_{i,k,t} didk$ are financed through debt issuance and labor income. The government budget flow constraint is

$$\frac{1+i_{t-1}^b}{1+\pi_t} b_{t-1} + g_t + \tau_t = \tau_\ell \int_0^1 \int_0^1 w_{k,t} e_{i,k,t} \ell_{i,k,t} didk + b_t$$

Anticipating my use of government spending shock impulse responses as key sufficient statistics for model identification, I allow region-specific government spending $\mathbf{g}_k$ to follow potentially region-specific paths, indexed by a shock path $\{\boldsymbol{\varepsilon}_{g,k}\}$. Given paths for spending targets $\mathbf{g}_k$, initial nominal debt b_{-1} and a path of prices and quantities $(\{\mathbf{w}_k\}, \{\boldsymbol{\ell}_k\}, \mathbf{i}^b, \boldsymbol{\pi})$, a government debt financing rule is a path $\boldsymbol{\tau}$ such that the flow government budget constraint holds at all periods t, and $\lim_{t \to \infty} \left(\prod_{s=0}^t \frac{1+\pi_s}{1+i_{s-1}^b} \right) b_t = 0$. For simplicity, and in keeping with most previous work, I assume that regional transfers $\boldsymbol{\tau}_k$ are identical across regions in each period.

Equilibrium. An equilibrium of this economy is a set of prices and aggregates such that, given exogenous (shock) processes and government policies, all agents are optimizing and all markets – local output and labor markets, as well as the aggregate bond market – clear. A formal equilibrium definition is stated in Appendix A.1.

2.2 From regional regressions to macro counterfactuals

The model allows me to offer a structural interpretation of the cross-regional regressions routinely run in applied work. For regression estimands to be well-defined I of course need a notion of aggregate risk. I thus, exactly as in Wolf (2019), consider the linear vector moving-average representation induced by the first-order perturbation solution of the model, assuming that the shocks $s \in \{b, g\}$ follow $\varepsilon_{s,k,t} = \varepsilon_{s,t} + \xi_{s,k,t}$, where the aggregate and regional shock components are mutually i.i.d. and $N(0,1)$.

[5]The assumption of fully local factors of production – while made in previous work (e.g Nakamura & Steinsson, 2014) – is stark. My results, however, are not sensitively tied to this assumption. For example, all results continue to go through in a model with roundabout production, where intermediate goods producers in each region use as an additional production input an aggregator of the intermediate goods of all regions.

The micro regression. I assume that the econometrician estimates cross-regional regressions of the form

$$c_{k,t+h} = \alpha_k + \delta_t + \beta_{b,h} \times \varepsilon_{b,k,t} + u_{k,t+h} \tag{2}$$

where α_k is a region fixed effect and δ_t is a time fixed effect. Regressions like (2) are run, for example, in Mian & Sufi (2012), Mian et al. (2013), or Auclert et al. (2019). Region-specific fixed effects absorb any time-invariant heterogeneity across regions; more importantly, time fixed effects ensure that the regression exploits variation *across* regions, thus lessening the requirements for exogeneity of the shock measure $\varepsilon_{b,k,t}$ at the aggregate level.[6]

The rest of this section achieves two objectives. First, I offer a structural interpretation of the $\beta_{b,h}$'s, and in particular discuss how they relate to the objects of interest – the aggregate consumption impulse response function $\hat{\mathbf{c}}_b$. Second, I identify a set of (in principle estimable) sufficient statistics that – under some further structural assumptions – allow researchers to map the $\beta_{b,h}$'s into $\hat{\mathbf{c}}_b$.

Interpretation. It turns out that the estimand $\beta_{b,h}$ has a simple interpretation as the horizon-h impulse response of consumption in region k to a deleveraging shock in region k.

Proposition 1. *Consider the linear vector moving average representation induced by the structural model of Section 2.1. Then the ordinary least-squares estimand of* $\boldsymbol{\beta}_b = (\beta_{b,0}, \beta_{b,1}, \ldots)'$ *satisfies*

$$\boldsymbol{\beta}_b = \hat{\mathbf{c}}_{b_k,k} \tag{3}$$

where $\hat{\mathbf{c}}_{b_k,k}$ *is the impulse response function of region-k consumption to a one-off region-k deleveraging shock* $\varepsilon_{b,k,0} = \xi_{b,k,0} = 1$.

$\boldsymbol{\beta}_b$ differs from $\hat{\mathbf{c}}_b$ – the true object of interest – for two reasons. First, as an estimate of the regional response to a regional shock, it captures regional general equilibrium effects. Thus, and contrary to the simpler across-household or across-firm regressions studied in Wolf (2019), it is not interpretable as a proper partial equilibrium spending response. Second, and exactly like across-household or across-firm regressions, it does not account for *aggregate* general equilibrium effects.

In practice, local general equilibrium effects are unlikely to be informative about aggregate feedback; for example, the strength of local effects depends on the openness of regional

[6] To be clear, in my controlled model environment, availability of the shock measure $\varepsilon_{b,k,t}$ is enough to trivially recover $\mathbf{c}_b$ through direct local projections (Jordà, 2005). In practice, however, such shock measures are rarely available, necessitating the use of time fixed effects.

economies, but is independent of aggregate policy rules, exactly opposite to aggregate effects. I thus do not try to exploit information on regional general equilibrium effects, but instead treat it as a nuisance: In the following, I first *remove* any local general equilibrium amplification to obtain pure partial equilibrium effects, and then try to account for the missing *aggregate* general equilibrium feedback.

Removing local general equilibrium effects. The regional consumption impulse response $\hat{\mathbf{c}}_{b_k,k}$ has two components: the *direct* response of consumption to the shock (i.e., the solution to the consumption-savings problem (1), changing *only* the borrowing wedge and constraint), and the *indirect* effects of local price feedback. To remove local general equilibrium effects, I invert the demand equivalence logic of Wolf (2019). The argument is simple. First, I give sufficient conditions for regional government spending and consumption demand shocks to propagate identically in general equilibrium. Since the response of local consumption to a local government spending shock, $\hat{\mathbf{c}}_{g_k,k}$, will purely reflect these common general equilibrium effects, the difference $\hat{\mathbf{c}}_{b_k,k} - \hat{\mathbf{c}}_{g_k,k}$ should recover the direct (partial equilibrium) consumption response to the deleveraging shock. Formally, this logic goes through under the following assumption.

Assumption 1. *There are either no (short-run) wealth effects in labor supply, or wages are perfectly sticky.*

We then arrive at the following result.[7]

Proposition 2. *Let $\hat{\mathbf{c}}_{b_k,k}^{PE}$ denote the direct (partial equilibrium) response of consumption in region k to a regional consumption demand shock $\varepsilon_{b,k,0} = \xi_{b,k,0} = 1$. Then, under Assumption 1 and to first order,*

$$\hat{c}_{b_k,k,h}^{PE} = \beta_{b,h} - \beta_{g,h} \tag{4}$$

where $\beta_{g,h}$ – the estimand of a cross-regional regression like (2) for local government spending – is also the impulse response function of region-k consumption to a region-k government spending shock, with government spending in region k changing by $\hat{g}_{g_k,k,h} = \phi \times \hat{c}_{b_k,k,h}^{PE}$.

Proposition 2 shows that regressions exactly analogous to (2) – but with local government spending shocks $\varepsilon_{g,k,t}$ on the right-hand side – contain, under relatively weak structural

[7]Recall that I assume identical transfers across all regions. I thus need no additional assumption on the financing of the local spending shock.

assumptions, all the information needed to strip the regression estimand in (2) of its confounding local general equilibrium effects. Encouragingly, this result promises to be practically useful because its informational requirements are manageable: First, estimates of the degree of home bias, necessary to align the demand pressures on *local* production, are readily obtainable. Second, given the increasingly widespread availability of regional consumption data, the required $\beta_{g,h}$ are estimable (Dupor et al., 2018).

Similarly to my analysis here, Guren et al. (2019) also try to rid cross-regional regression estimates of local general equilibrium effects. To this end, they propose to *divide* estimates of regional consumption spending effects by estimates of the regional fiscal (output) multiplier. For the impact response, their proposed correction is exactly identical to mine in a model without capital and with full home bias.

Proposition 3. *Consider a variant of the benchmark model without capital ($\alpha = 0$) and with complete home bias ($\phi = 1$), and let $m^y_{g_k,k} \equiv \frac{y_{g_k,k,0}}{g_{g_k,k,0}}$ denote the impact local government spending multiplier. Then, under Assumption 1 and to first order,*

$$\hat{c}^{PE}_{b_k,k,0} = \frac{\beta_{b,0}}{m^y_{g_k,k}} \tag{5}$$

In my model, local output is used for local and foreign consumption, local investment, and local government spending. Without investment and with full home bias, local output is used only for local private and public consumption, so local consumption impulse responses and local output multipliers contain the same information, and (4) and (5) are equivalent.

To see more clearly how the two corrections are likely to differ in practice, it will be instructive to re-write my proposed correction (4) in the static multiplier notation of Guren et al.. Letting $m^c_{g_k,k} \equiv \frac{c_{g_k,k,0}}{g_{g_k,k,0}}$ denote the impact regional consumption multiplier, it is straightforward to show that, under the assumptions of Proposition 2,

$$\hat{c}^{PE}_{b_k,k,0} = \frac{\beta_{b,0}}{1 + \phi \times m^c_{g_k,k}} \tag{6}$$

Two differences are noteworthy: First, to account for the unobserved local investment feedback, (6) looks at local *consumption* multipliers. Second, the local consumption multiplier is scaled by the degree of home bias, $\phi \in (0, 1)$. Intuitively, for the exact demand equivalence result of Proposition 2, the underlying government spending shock must be strictly smaller than the overall partial equilibrium consumption response, since some of that consumption response leaks to other regions. Dividing by the full regional consumption multiplier would thus strictly over-state local general equilibrium amplification.[8]

Overall, the trade-off between the two proposed corrections is best phrased in terms of applicability versus informational requirements: The additive correction (4) applies in a broader class of models, but requires information on local consumption responses, rather than the more readily estimable local output multipliers.

Aggregation. It remains to aggregate the region-specific partial equilibrium spending responses $\{\hat{\mathbf{c}}_{b_k,k}^{PE}\}$ into a general equilibrium impulse response for aggregate consumption $\hat{\mathbf{c}}_b$. My main result is that aggregate consumption impulse responses to aggregate government spending shocks can provide a consistent estimate of the missing general equilibrium feedback. The argument proceeds in two steps.

First, I establish a particular *regional invariance* result: It does not matter where precisely (partial equilibrium) consumption spending demand changed, just by how much it changed overall, averaging across all regions. Intuitively, the incidence of spending shocks clearly matters for relative prices but, given enough symmetry across regions, it will not matter for macroeconomic aggregates. This result allows me to treat the collection of region-specific spending shocks $\{\hat{\mathbf{c}}_{b_k,k}^{PE}\}$ as equivalent to an aggregate shock that increases partial equilibrium spending in each individual region by $\int_0^1 \hat{\mathbf{c}}_{b_k,k}^{PE} dk$.

Second, I proxy for the missing general equilibrium intercept of this aggregate consumption demand shock through the response of aggregate consumption to aggregate government spending shocks. As in Wolf (2019), the underlying demand equivalence result requires further restrictions on the household block of the economy and on the deficit financing rule of the government.

Assumption 2. *There is no penalty for liquid borrowing ($\kappa_t^b = 0$ for all t).*

This assumption ensures that households and government borrow and lend at identical rates at all times.[9] In particular, this restriction guarantees that private and public demand shocks can be financed using identical paths of taxes and transfers. The next assumption ensures that this is indeed the case.

---

[8] In Appendix A.4, I show that the implied static correction is slightly more involved with non-atomistic (and heterogeneously-sized) regions. In a simple two-region economy with relative sizes $\zeta \in (0,1)$ and $1-\zeta$, the correction term on the consumption multiplier is $\phi \times \frac{\phi-\zeta}{\phi(1-\zeta)}$. Intuitively, the further adjustment reflects the fact that leaking local spending to other regions now also has aggregate price effects.

[9] A deleveraging shock is thus a pure shock to household borrowing constraints, exactly as in Guerrieri & Lorenzoni (2017) or Korinek & Simsek (2016).

Assumption 3. *The path of taxes and transfers used to finance a given government spending shock* $\{\varepsilon_g\}$ *depends only on the present value of the expenditure, not its time path. A spending path with 0 net present value is purely deficit-financed, and so it elicits no direct tax response.*

A formal statement of the identification result now follows.

Proposition 4. *Consider an aggregate deleveraging shock* $\{\varepsilon_{b,k}\}$ *that changes partial equilibrium consumption demand in region* k *by* $\hat{\mathbf{c}}_{b_k,k}^{PE}$*. Then the aggregate effects of this shock are, under Assumptions 1 to 3 and to first order, given as*

$$\hat{\mathbf{c}}_b = \int_0^1 \hat{\mathbf{c}}_{b_k,k}^{PE} dk + \hat{\mathbf{c}}_g \tag{7}$$

where $\hat{\mathbf{c}}_g$ *is the response of aggregate consumption to an aggregate government spending shock with* $\hat{\mathbf{g}}_g = \int_0^1 \hat{\mathbf{c}}_{b_k,k}^{PE} dk$*.*

The additive decomposition in (7) is exactly analogous to that in Wolf (2019), and valid under exactly the same assumptions. The only novelty of Proposition 4 is the underlying regional invariance result.

Summary. My identification results connect cross-regional regression estimates of consumption demand to the associated full general equilibrium counterfactuals. Such identification results are useful for two reasons.

First, under somewhat stringent assumptions and informational requirements, Propositions 2 and 4 allow the construction of semi-structural aggregate counterfactuals for *particular* shocks, without ever requiring the solution of a structural model. This is the approach taken in Wolf (2019). Second, my results suggest that the combination of (i) cross-regional consumption spending regression estimates and (ii) evidence on the effects of local and aggregate government spending changes should be highly informative as target moments in structural model estimation. Section 3 illustrates this point through estimation of model-based counterfactuals for an aggregate shock to household borrowing conditions.

2.3 Generalizations and robustness

I briefly conclude this section with several comments on the generality and limitations of my identification results.

First, I consider a multi-region model with infinitesimal islands, rather than a two-region model, as in Nakamura & Steinsson (2014). I do so for notational simplicity and (arguably) empirical relevance (Chodorow-Reich, 2017; Beraja et al., 2019), but my identification results extend almost without change to models with finitely many regions and heterogeneous size, as I show in Appendix A.4. Second, while my formal identification results in Propositions 2 and 4 rely on Assumptions 1 to 3, they apply approximately in even more general quantitative structural models. In particular, because of the regional invariance result, all conclusions about accuracy of the demand equivalence approximation in Wolf (2019) apply without change to my aggregation results in Proposition 4. The local identification result in Proposition 2 is likely to be accurate for the same reasons: As is well-known, even moderate degrees of wage and price stickiness materially dampen the effects of shifts in labor supply (Christiano et al., 2011; Christiano, 2011), so Assumption 1 – which is all that is needed for *regional* demand equivalence – is likely to hold at least approximately. I provide concrete illustrative evidence through the analysis in Section 3. Third, while stated here in the context of household deleveraging shocks, it is immediate that my identification results apply without change to generic "consumption demand" shifters. Furthermore, proceeding exactly as in Wolf (2019), it is possible to generalize all results to cross-regional regressions for *investment* demand counterfactuals. However, since such regressions are rare in applied work, I do not pursue this extension here.

3 Application: household deleveraging

This section uses my identification results to provide estimation targets for a structural model of household deleveraging. As such, my analysis is similar in spirit to Guerrieri & Lorenzoni (2017) and Jones et al. (2018), but differs in its approach to model estimation.

Section 3.1 begins with a static back-of-the-envelope calculation, using the identification results of Section 2 together with existing empirical evidence to construct a simple approximation to the desired full macro counterfactual. Section 3.2 then moves to estimation of a rich heterogeneous-household, multi-region model of the macro-economy. Finally, Section 3.3 constructs aggregate counterfactuals in the estimated model and shows that the simple back-of-the-envelope calculation provides an excellent approximation, illustrating the power of my identification results.

3.1 Back-of-the-envelope calculation

Existing empirical work on cross-regional spending responses mostly estimates a single (often impact) response, rather than the entire dynamic paths required in Propositions 2 and 4.

While not strictly speaking sufficient in rich dynamic models, these single responses nevertheless allow a simple back-of-the-envelope calculation.

First, household deleveraging due to tightening borrowing conditions is well-known to be associated with substantial *cross-regional* responses in consumption (Mian et al., 2013). Recall that the aggregated cross-regional spending response to a tightening in household borrowing conditions in the model of Section 2.1 is defined as

$$\int_0^1 \beta_{b,0} \times \varepsilon_{b,k,0} \, dk \quad = \quad \beta_{b,0} \times \varepsilon_{b,0} \tag{8}$$

Mian et al. (2013) estimate the county-level response of consumption to changes in housing wealth. Consistent with the results in Berger et al. (2017), I interpret this estimated housing wealth effect as acting mostly through collateral effects – in other words, a shock to household borrowing conditions.[10] Mian et al. estimate that, for a dollar decline in housing wealth, relative regional consumption drops by around 6 cents; multiplying this $\beta_{b,0}$ by aggregate declines in housing wealth during the Great Recession, we arrive at a full cross-regional spending response of around 3 per cent of total aggregate consumption.[11]

Second, Dupor et al. (2018) estimate that county-level consumption rises by around 25-30 cents for every dollar of local government spending. Given a standard estimate of county-level home bias of around $\phi = 0.6$,[12] the simple static approximation (6) implies that the estimated cross-regional spending decline of around 3 per cent corresponds to a pure partial equilibrium spending contraction of around 2.6 per cent.

Third, transitory changes in government spending are usually estimated to change output one-for-one, and lead to little general equilibrium crowding-in or -out of private spending (Ramey, 2018; Wolf, 2019). The 2.6 per cent estimated partial equilibrium contraction in consumption should thus translate into a similarly large general equilibrium drop, as well as around a 1.7 per cent contraction in aggregate output (assuming a consumption share of around 65 per cent). The next two sections show that this simple static approximation accurately predicts aggregate counterfactuals in a rich estimated dynamic model.

[10] This is consistent with an interpretation of the model in Section 2 featuring fully rigid individual housing, exogenous house prices, and borrowing conditions mechanically tied to housing values.

[11] Beraja et al. (2019) implement a similar back-of-the-envelope calculation, and arrive at almost identical numbers (see their Figure 4).

[12] Home bias at the MSA level can be calibrated directly from data on across-MSA shipments of goods (which is readily obtainable from the Commodity Flow Survey). Adjusting for the size of counties relative to MSAs gives the county-level home bias.

3.2 Model estimation

I consider the structural model outlined in Section 2.1. Consistent with my focus on household deleveraging, I allow for uninsurable income risk and so rich heterogeneity at the household level (Kaplan et al., 2018), but keep the rest of the economy close to standard business-cycle models (Justiniano et al., 2010). In particular, and as in Wolf (2019), I assume neither the absence of wealth effects in labor supply nor fully rigid prices (violating Assumption 1), and furthermore allow household borrowing to be subject to a time-varying borrowing wedge (violating Assumption 2). My demand equivalence results thus apply only approximately.

Many of the targets for the parameterization of the structural model are relatively standard, and thus relegated to Appendix A.3. Crucially, however, I complement these standard estimation targets with various moments identified by the theory in Section 2.2 as highly informative about the desired macro counterfactuals.

Credit Tightening. I implement the shock to household borrowing conditions as a temporary tightening in borrowing constraints $\underline{b}_{k,t}$ as well as a temporary increase in the borrowing wedge $\kappa_{k,t}$, similar to Guerrieri & Lorenzoni (2017). For simplicity, I tighten both kinds of financial constraints by the same percentage amount, and choose this amount to match the peak *cross-regional* consumption decline of around 3 per cent estimated in Mian et al. (2013). I assume that borrowing conditions remain tight for two years, and then return to baseline relatively quickly, consistent with the length of the most severe phase of the Great Recession. Further details on the chosen shock processes are provided in Appendix A.3.[13]

Fiscal multipliers. I select parameters governing model dynamics – in particular the degree of nominal price rigidity as well as the monetary feedback rule – to ensure consistency of the model with extant empirical evidence on the cross-regional and aggregate transmission of government shocks.

Figure 1 compares model-implied and VAR-estimated consumption and output impulse responses to aggregate government spending shocks.[14] Consistent with most of the existing literature, I estimate that a short-lived increase in government spending today – financed with

[13] All subsequent results are robust to changing the fractions of the consumption demand drop attributable to tightening borrowing constraints and higher borrowing wedges. In particular, and consistent with the results in Wolf (2019), inaccuracies in demand equivalence due to heterogeneous private and public discount rates are only moderate. Ultimately, what matters most is the overall consumption drop, not its source.

[14] Details for the VAR estimation are provided in Wolf (2019).

an increase in the public deficit today, and higher taxes on households over the following years – roughly moves output one-for-one, and only elicits a weak response of private spending. The structural model is then parameterized to be consistent with these results: The orange line in the right panel shows that, following a transitory and deficit-financed increase in government spending, output increases by around 0.2 per cent – a unit multiplier. The left panel gives the corresponding flat consumption response, consistent with the reasonably tightly estimated zero response in the recursive VAR.

IRF Matching, Aggregate G

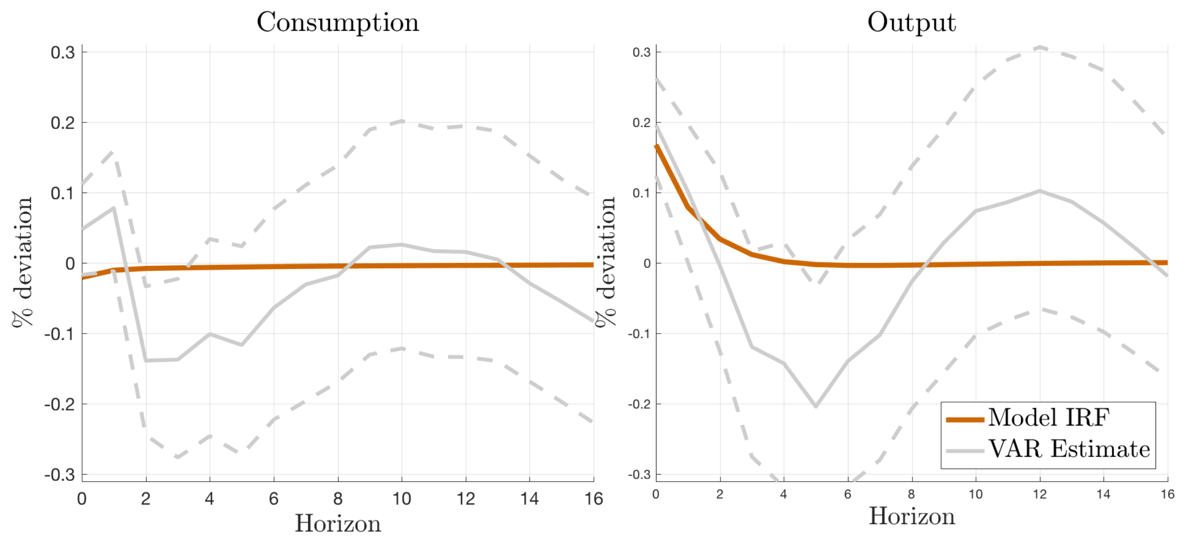

Figure 1: Estimated (grey) and model-implied (orange) impulse responses to a transitory deficit-financed expansion in government spending. The dotted lines are 16th and 84th percentile posterior bands of the estimated VAR. For details on the empirical specification see Wolf (2019).

For the effects of local government spending, I match the preferred point estimates of Dupor et al. (2018). They find that a transitory (hump-shaped) relative increase in local government spending is associated with a four-year cumulative relative consumption multiplier of 0.29. In my model, the exact same hump-shaped government spending path is associated with a discounted cumulative relative multiplier of 0.26.[15]

Overall, my model is consistent with *positive regional* but mildly *negative aggregate* consumption responses to increases in fiscal spending. Local multipliers are positive because, first, high MPCs set in motion a standard Keynesian feedback loop (as in Auclert et al., 2018), and second, neither fiscal nor monetary feedback moderate the expansion in demand.

[15]Specifically, I replicate the model-simulated spending path displayed in Figure 3 of Dupor et al. (2018), and compute cumulative multipliers exactly as in their paper.

Aggregate multipliers in contrast are close to 0 precisely because of aggregate tax financing and the moderating effects of imperfectly accommodative monetary policy.

3.3 Deleveraging in general equilibrium

Previous work has not reached a consensus on the aggregate effects of a household credit tightening and the associated deleveraging. Guerrieri & Lorenzoni (2017) show that, in a neoclassical model, large *partial equilibrium* consumption demand contractions are mostly reversed in general equilibrium.[16] In a sticky-wage variant of the same model, the declines in aggregate consumption and output instead turn out to be large. Jones et al. (2018) proceed differently and estimate a rich structural model using likelihood techniques, decomposing cross-regional and aggregate fluctuations into various explicitly modeled structural shocks. In their model, deleveraging shocks have small effects when monetary policy follows a conventional Taylor rule, and so explain little of the aggregate consumption drop in the Great Recession. In yet another example of a structural analysis disciplined using likelihood-based estimation, Beraja et al. (2019) find aggregate consumption declines that are, at least over the short run, extremely similar to those implied by naive aggregation of cross-regional micro regressions like (8).

In the remainder of this section I present counterfactuals for a household credit tightening in the structural model of Section 3.2. The model has a similar structure to previous work, but the approach to identification is entirely different – I do not rely on a *full information* approach giving structural decompositions of all cross-regional and macro aggregates, but instead directly target the sufficient statistics identified in Section 2.2. Since this "identified moments" approach is valid across a large class of structural models, it is arguably more robust to mis-identification than decompositions that are sensitive to all features of the researcher's model (Nakamura & Steinsson, 2018).

Deleveraging in the estimated model. Results for deleveraging shocks in the estimated structural model are summarized in Figure 2. Building on the analysis in Section 2, the figure shows three sets of impulse responses: (i) the response of regional consumption to a regional deleveraging shock (in purple), (ii) the aggregated average partial equilibrium consumption demand response $\int_0^1 \hat{\mathbf{c}}_{b_k,k}^{PE} dk$ (in green), and (iii) aggregate general equilibrium

[16] In fact, in a flexible-price ($\phi_p, \phi_w \to 0$), labor-only ($\alpha \to 0$) version of my benchmark model, any partial equilibrium consumption demand change is *fully* offset in general equilibrium.

feedback and the full macro counterfactual (in orange and grey, respectively).

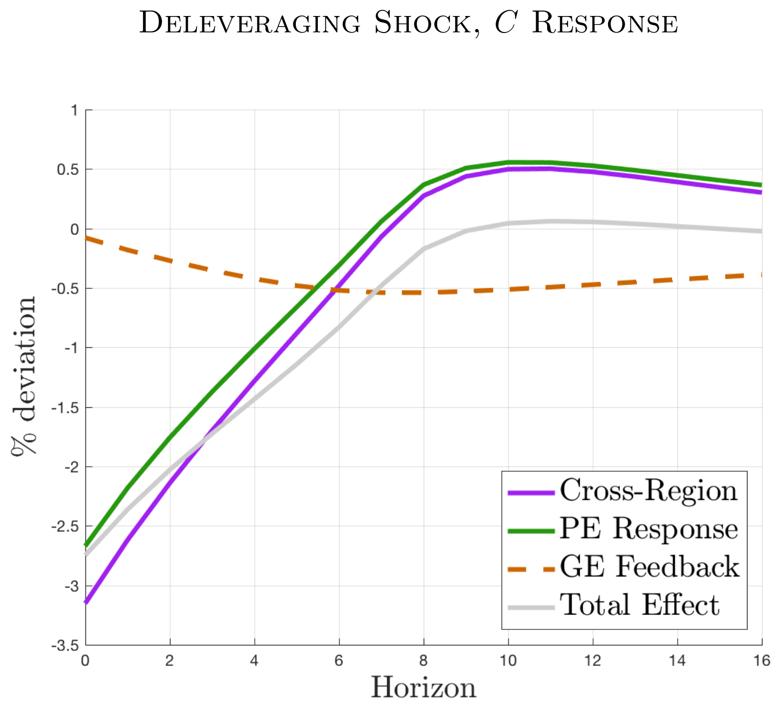

Figure 2: Decomposition of model-implied impulse responses of consumption to a transitory deleveraging shock.

The key take-away is that the static back-of-the-envelope calculation of Section 3.1 provides an excellent approximation to the true model-implied aggregate counterfactual: First, the impact cross-regional response – which by Proposition 1 is identical to the estimand of cross-regional regressions like (2) – is around 3 per cent, as in the estimates of Mian et al. (2013). Second, given a cross-regional impact response of around 3 per cent, the back-of-the-envelope calculation suggested a pure partial equilibrium demand contraction of around 2.6 per cent. The true model-implied partial equilibrium response is a very close 2.67 per cent. Third, the orange and grey lines show that the model features relatively limited general equilibrium amplification, consistent with the matched VAR evidence. Finally, aggregate output declines by around 1.9 per cent (not shown), also as predicted above.[17] This accuracy is not surprising: As I discuss in Wolf (2019), the assumptions required for demand

[17] Given the model-implied estimate of the pure partial equilibrium demand response, I could exploit the richness of estimated aggregate government spending shocks to directly match this demand path, and provide a truly *semi*-structural counterfactual, exactly as in Wolf (2019). Results are available upon request, and unsurprisingly closely agree with Figure 2. I do not do so here because arriving at the full partial equilibrium spending path already required some model structure.

equivalence and so for the identification results of Section 2.2 are robustly *nearly* satisfied in quantitatively relevant structural macro models, including the model considered here.

Discussion. A model that is jointly consistent with (i) large cross-regional spending responses to household deleveraging, (ii) moderately positive local consumption multipliers and (iii) a unit aggregate fiscal multiplier for deficit-financed spending will invariably imply large general equilibrium effects of household deleveraging. Importantly, my analysis shows that these conclusions apply in a structural model with a conventional – in fact quite aggressive – monetary policy feedback rule. As such, and contrary to previous work (e.g. Jones et al., 2018), my results imply that the deleveraging shock studied in Mian et al. (2013) is likely to have large aggregate effects even in the absence of a binding ZLB constraint.

Of course, with constrained monetary policy, aggregate general equilibrium responses may be even larger. In principle, there are two ways of estimating such crisis-time counterfactuals. First, the aggregate government spending impulse responses $\hat{\mathbf{c}}_g$ required for my methodology may be estimated over a sample period with binding ZLB constraint. Perhaps surprisingly, existing empirical evidence suggests that the propagation of structural shocks was not affected particularly strongly by the ZLB (Ramey & Zubairy, 2018; Debortoli et al., 2019). Second, the estimated structural model for normal times can be used for crisis-time counterfactuals, with the monetary authority fixing the nominal rate. In this case, the general equilibrium drop of consumption is (unsurprisingly) larger, at around 3.8 per cent. The results presented here are thus best interpreted as an informative *lower bound* for the aggregate effects of household deleveraging.

Overall, my findings have a "sufficient statistics" interpretation similar to Chetty (2009) and Arkolakis et al. (2012): *Any* fully calibrated structural model that is inconsistent with the results presented here must either break demand equivalence or is inconsistent with empirical evidence on local and aggregate fiscal multipliers.

4 Conclusion

I have extended the results in Wolf (2019) to connect cross-regional regression estimates and macro counterfactuals. The proposed approach, formally justified in a large class of multi-region, heterogeneous-agent structural macro models, only requires the researcher to have estimates of local and aggregate *public* spending shocks. I apply my identification results to study household deleveraging and conclude that the cross-regional regression estimates of

Mian et al. (2013) correspond to large *aggregate* declines in consumption and production.

Future work should try to further improve measurement of the identified sufficient statistics. First, estimates of full impulse response *paths* following local government spending shocks will allow researchers to strip out local general equilibrium effects in a fully semi-structural way. Second, extending estimates of aggregate fiscal multipliers to alternative macroeconomic regimes – notably the ZLB episode – should facilitate the construction of *state-dependent* counterfactuals.

References

Adao, R., Arkolakis, C., & Esposito, F. (2019). Spatial linkages, global shocks, and local labor markets: Theory and evidence. Technical report, National Bureau of Economic Research.

Arkolakis, C., Costinot, A., & Rodriguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, *102*(1), 94–130.

Auclert, A., Dobbie, W. S., & Goldsmith-Pinkham, P. (2019). Macroeconomic effects of debt relief: Consumer bankruptcy protections in the great recession. Technical report, National Bureau of Economic Research.

Auclert, A., Rognlie, M., & Straub, L. (2018). The Intertemporal Keynesian Cross. Working Paper.

Beraja, M., Hurst, E., & Ospina, J. (2019). The aggregate implications of regional business cycles. Technical report, National Bureau of Economic Research.

Berger, D., Guerrieri, V., Lorenzoni, G., & Vavra, J. (2017). House prices and consumer spending. *The Review of Economic Studies*, *85*(3), 1502–1542.

Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting mit shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, *89*, 68–92.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, *1*(1), 451–488.

Chodorow-Reich, G. (2017). Geographic cross-sectional fiscal multipliers: What have we learned? *American Economic Journal: Economic Policy*.

Christiano, L., Eichenbaum, M., & Rebelo, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, *119*(1), 78–121.

Christiano, L. J. (2011). Comment on" what fiscal policy is effective at zero interest rates?". In *NBER Macroeconomics Annual 2010, Volume 25* (pp. 113–124). University of Chicago Press.

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, *113*(1), 1–45.

Debortoli, D., Galí, J., & Gambetti, L. (2019). On the empirical (ir) relevance of the zero lower bound constraint. Technical report, National Bureau of Economic Research.

Dupor, B., Karabarbounis, M., Kudlyak, M., & Mehkari, M. (2018). Regional consumption responses and the aggregate fiscal multiplier. FRB St. Louis Working Paper.

Eggertsson, G. B. & Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach. *The Quarterly Journal of Economics*, *127*(3), 1469–1513.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., & Schorfheide, F. (2016). Solution and estimation methods for dsge models. In *Handbook of Macroeconomics*, volume 2 (pp. 527–724). Elsevier.

Gali, J. & Monacelli, T. (2005). Monetary policy and exchange rate volatility in a small open economy. *The Review of Economic Studies*, *72*(3), 707–734.

Grigsby, J., Hurst, E., & Yildirmaz, A. (2019). Aggregate nominal wage adjustments: New evidence from administrative payroll data. Technical report, National Bureau of Economic Research.

Guerrieri, V. & Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, *132*(3), 1427–1467.

Guren, A., McKay, A., Nakamura, E., & Steinsson, J. (2019). What do we learn from cross-sectional empirical estimates in macroeconomics? Working Paper.

Guren, A. M., McKay, A., Nakamura, E., & Steinsson, J. (2018). Housing wealth effects: The long view. Technical report, National Bureau of Economic Research.

Hall, R. E. (2009). By how much does gdp rise if the government buys more output? *Brookings Papers on Economic Activity*.

Johnson, D. S., Parker, J. A., & Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, *96*(5), 1589–1610.

Jones, C., Midrigan, V., & Philippon, T. (2018). Household leverage and the recession. Technical report, National Bureau of Economic Research.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, *95*(1), 161–182.

Justiniano, A., Primiceri, G. E., & Tambalotti, A. (2010). Investment shocks and business cycles. *Journal of Monetary Economics*, *57*(2), 132–145.

Kaplan, G., Moll, B., & Violante, G. (2018). Monetary Policy according to HANK. *American Economic Review*, *108*(3), 697–743.

Khan, A. & Thomas, J. K. (2013). Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity. *Journal of Political Economy*, *121*(6), 1055–1107.

Korinek, A. & Simsek, A. (2016). Liquidity trap and excessive leverage. *American Economic Review*, *106*(3), 699–738.

Mian, A., Rao, K., & Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, *128*(4), 1687–1726.

Mian, A. & Sufi, A. (2012). The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. *The Quarterly journal of economics*, *127*(3), 1107–1142.

Mian, A. & Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, *82*(6), 2197–2223.

Mian, A. & Sufi, A. (2015). *House of debt: How they (and you) caused the Great Recession, and how we can prevent it from happening again*. University of Chicago Press.

Nakamura, E. & Steinsson, J. (2014). Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review*, *104*(3), 753–92.

Nakamura, E. & Steinsson, J. (2018). Identification in macroeconomics. *Journal of Economic Perspectives*, *32*(3), 59–86.

Ramey, V. A. (2018). Ten years after the financial crisis: What have we learned from the renaissance in fiscal research? Working Paper.

Ramey, V. A. & Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from us historical data. *Journal of Political Economy*, *126*(2), 850–901.

Schmitt-Grohé, S. & Uribe, M. (2003). Closing small open economy models. *Journal of international Economics*, *61*(1), 163–185.

Smets, F. & Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, *97*(3), 586–606.

Wolf, C. K. (2019). The missing intercept: A sufficient statistics approach to general equilibrium effects. Working Paper.

Zwick, E. & Mahon, J. (2017). Tax policy and heterogeneous investment behavior. *The American Economic Review*, *57*(3), 391–414.

A Model details

This appendix provides further details on the structural model that underlies the identification results of Section 2 and the aggregate counterfactuals of Section 3. Appendix A.1 provides the missing details on the rich structural model sketched in Section 2.1, Appendix A.2 then characterizes the model equilibrium, and Appendix A.3 contains all details on the particular parameterization used in Section 3. Finally, Appendix A.4 extends my identification results to a model with non-infinitesimal regions.

A.1 Model outline

Recall that the model is populated by households, firms, and the government.

Households. It remains to specify the problem of wage-setting unions. Each region is populated by a unit continuum of unions, which intermediates labor services and sells them to a competitive labor packer. I assume that the union satisfies labor demand by demanding a common amount of labor input from all of its members. The labor packer then aggregates union-specific labor to composite labor services provided to firms at the region-specific wage index $w_{k,t}$. Proceeding exactly as in Wolf (2019), it can be shown that the problem of labor unions induces the following non-linear wage-NKPC:

$$\pi^w_{k,t}(1+\pi^w_{k,t}) = \frac{\varepsilon_w}{\theta_w}\ell^h_{k,t}\left[\int_0^1 \left\{-u_\ell(c_{i,k,t},\ell^h_{k,t}) - \frac{\varepsilon_w - 1}{\varepsilon_w}(1-\tau_\ell)w_{k,t}e_{i,k,t}u_c(c_{i,k,t},\ell^h_{k,t})\right\} di\right] + \beta\pi^w_{k,t+1}(1+\pi^w_{k,t+1}) \quad \text{(A.1)}$$

where $1+\pi^w_{k,t} = \frac{w_{k,t}}{w_{k,t-1}} \times (1+\pi_t)$, ε_w denotes the elasticity of substitution between different kinds of labor, and θ_w denotes the Rotemberg adjustment cost.

Together, the household consumption-savings problem and the wage-NKPC (A.1) characterize aggregate household behavior. I assume that the solutions to each problem exist and are unique, and summarize the solution in terms of aggregate consumption, saving and union labor supply functions $\mathbf{c}_k = \mathbf{c}_k(\mathbf{s}^h_k;\boldsymbol{\varepsilon}_k)$, $\mathbf{c}^H_k = \mathbf{c}^H_k(\mathbf{s}^h_k;\boldsymbol{\varepsilon}_k)$, $\mathbf{c}^{k'}_k = \mathbf{c}^{k'}_k(\mathbf{s}^h_k;\boldsymbol{\varepsilon}_k)$, $\boldsymbol{\ell}^h_k = \boldsymbol{\ell}^h_k(\mathbf{s}^u_k)$, and $\mathbf{b}_k = \mathbf{b}_k(\mathbf{s}^h_k;\boldsymbol{\varepsilon}_k)$, where $\mathbf{s}^h = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}_k, \mathbf{d}_k, \boldsymbol{\tau}_k, \boldsymbol{\ell}_k, \mathbf{p})$ and $\mathbf{s}^u = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}_k, \mathbf{c}_k)$. For all identification results, I will impose the high-level assumption that all of those functions are at least once differentiable in their arguments.

FIRMS. Since region-k firms are owned by region-k households, I – with the same justification as in Wolf (2019) – assume that all firms in a given region k discount at the common rate $1 + r^b_{k,t} = \frac{1+i^b_{t-1}}{1+\pi_{k,t}}$, where $\pi_{k,t} = \frac{q_{k,t}}{q_{k,t-1}} \times (1 + \pi_t)$.

1. *Intermediate goods producers.* The problem of intermediate goods producer j is to

$$\max_{\{d^I_{j,k,t}, y_{j,k,t}, \ell^f_{j,k,t}, k_{j,k,t}, i_{j,k,t}, u_{j,k,t}, b^f_{j,k,t}\}} \quad \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \left(\prod_{q=0}^{t-1} \frac{1}{1+r^b_{k,q}} \right) d^I_{j,k,t} \right]$$

such that

$$\begin{aligned}
d^I_{j,k,t} &= \underbrace{p^I_{k,t} y_{j,k,t} - w_{k,t} \ell^f_{j,k,t}}_{\pi_{j,k,t}} - \xi_{j,k,t} \times \mathbf{1}_{i_{j,k,t} \neq 0} - (1 - \mathbf{1}_{i_{j,k,t}<0} \times \varphi) i_{j,k,t} \\
&\quad - \phi(k_{j,k,t}, k_{j,k,t-1}, i_{j,k,t}, i_{j,k,t-1}) - b^f_{j,k,t} + \frac{1+i^b_{t-1}}{1+\pi_t} b^f_{j,k,t-1} \\
y_{j,k,t} &= y(e_{j,k,t}, u_{j,k,t} k_{j,k,t-1}, \ell^f_{j,k,t}) \\
i_{j,k,t} &= k_{j,k,t} - [1 - \delta(u_{j,k,t})] k_{j,k,t-1} \\
-b^f_{j,k,t} &\leq \Gamma(k_{j,k,t-1}, k_{j,k,t}, \pi_{j,k,t}) \\
d^I_{j,k,t} &\geq \underline{d}
\end{aligned}$$

Adjustment costs have a convex and continuously differentiable part ϕ, a potentially firm-specific fixed adjustment cost $\xi_{j,k,t}$ (distributed according to some cdf $F(\xi)$ over support $\mathbb{R}^+$), and may feature partial irreversibility, with $\varphi \in [0,1]$. For all numerical experiments I will assume a simple Cobb-Douglas production technology, with $y_{j,k,t} = e_{j,k,t}(u_{j,k,t} k_{j,k,t-1})^\alpha (\ell^f_{j,k,t})^{1-\alpha}$. Aggregating across firms, the solution to this problem gives optimal production $\mathbf{y}_k(\bullet)$, labor demand $\boldsymbol{\ell}^f_k(\bullet)$, investment $\mathbf{i}_k(\bullet)$, intermediate goods producer dividends $\mathbf{d}^I_k(\bullet)$, capital utilization rates $\mathbf{u}_k(\bullet)$ and liquid corporate bond savings $\mathbf{b}^f_k(\bullet)$ as a function of nominal returns $\mathbf{i}^b$, inflation $\boldsymbol{\pi}$, local wages $\mathbf{w}_k$, the cost of the local consumption basket $\mathbf{q}_k$, and the local intermediate goods price $\mathbf{p}^I_k$.

2. *Retailers.* A unit continuum of retailers purchases the intermediate good at price $p^I_{k,t}$, costlessly differentiates it, and sells it on to a final goods aggregator. Price setting is subject to a Rotemberg adjustment cost. As usual, optimal retailer behavior gives rise to a standard NKPC as a joint restriction on the paths of inflation and the intermediate

goods price. In log-linearized form:

$$\widehat{\log(\pi_{k,t})} = \underbrace{\frac{\varepsilon_p}{\theta_p} \frac{\varepsilon_p - 1}{\varepsilon_p}}_{\kappa_p} \times (\widehat{\log(p^I_{k,t})} - \widehat{\log(p_{k,t})}) + \beta \times \widehat{\log(\pi_{k,t+1})}$$

where $\widehat{\log(\pi_{k,t})} = \widehat{\log(\pi_t)} + (\widehat{\log(p_{k,t})} - \widehat{\log(p_{k,t-1})})$, ε_p denotes the substitutability between different kinds of retail goods, and θ_p denotes the Rotemberg adjustment cost. In an equivalent (to first-order) Calvo formulation, the slope of the NKPC instead is given as

$$\kappa_b = \frac{(1 - \frac{1}{1+\bar{r}}\phi_p)(1 - \phi_p)}{\phi_p}$$

where $1 - \phi_p$ is the probability of a price re-set. Finally, total dividend payments of retailers are

$$d^R_{k,t} = (1 - p^I_{k,t})y_{k,t}$$

3. *Aggregators.* Aggregators purchase retail goods and aggregate them to the composite final good. They make zero profits.

Total dividend payments by the corporate sector are given as

$$d_{k,t} = d^I_{k,t} + d^R_{k,t}$$

With some algebra, it is straightforward to show that in fact

$$d_{k,t} = y_{k,t} - w_{k,t}\ell_{k,t} - i_{k,t}$$

Assuming unique solutions to the firm sector decision problems, we can – along a perfect foresight transition path – summarize the aggregate firm sector simply through a set of optimal regional production, labor hiring, investment, dividend payment and bond demand functions, $\mathbf{y}_k = \mathbf{y}_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$, $\boldsymbol{\ell}^f_k = \boldsymbol{\ell}^f_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$, $\mathbf{i}_k = \mathbf{i}_k(\mathbf{s}^f; \boldsymbol{\varepsilon}_k)$, $\mathbf{d}_k = \mathbf{d}_k(\mathbf{s}^f; \boldsymbol{\varepsilon}_k)$ and $\mathbf{b}^f_k = \mathbf{b}^f_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$, where $\mathbf{s}^f_k = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}_k, \mathbf{p}_k)$. As before, I will assume that these aggregate firm sector-level functions are at least once differentiable in their arguments.

Government. The rate on nominal bonds i^b is set by the monetary authority. In line with standard empirical practice I assume that

$$\widehat{\log(i_t^b)} = \rho_m \widehat{\log(i_{t-1}^b)} + (1-\rho_m)\left(\phi_\pi \widehat{\log(\pi_t)} + \phi_y \widehat{\log(y_t)} + \phi_{dy}\widehat{\log(y_{t-1})}\right)$$

where $y_t \equiv \int_0^1 y_{k,t} dk$. For impulse response matching on government spending shock transmission, I assume that the fiscal authority follows the debt financing rule

$$\hat{\tau}_t = -(1-\rho_\tau) \times \hat{b}_{t-1}$$

Market-Clearing. The economy has three sets of markets. First, output market-clearing dictates that

$$y_{k,t} = c_{k,t}^H + c_{k,t}^A + i_{k,t} + g_{k,t}$$

where $c_{k,t}^H \equiv \int_0^1 c_{i,k,t}^H di$ and $c_{k,t}^A \equiv \int_0^1 \int_0^1 c_{i,k',t}^k di dk'$ are home and (aggregate) foreign consumption of the region-k output good. Second, region-specific labor markets clear if

$$\ell_{k,t}^h = \ell_{k,t}^f$$

Third, the aggregate bond market clears if

$$\int_0^1 \left\{ b_{k,t}^h + b_{k,t}^f \right\} dk = b_t$$

where $b_{k,t}^h \equiv \int_0^1 b_{i,k,t}^h di$ denotes total savings of region-k households.

Equilibrium definition. The equilibrium definition is standard. I assume that there exists a unique and region-symmetric deterministic steady state. To allow interpretation of perfect foresight transition paths as conventional first-order perturbation solutions, I impose that the economy is indeed initially in steady state, and then study perfect foresight transition equilibria back to the initial deterministic steady state.[18]

[18] To ensure that all individual regions return to steady state, additional assumptions may be required. As I show in Appendix A.2, the equilibrium can be characterized as that of a single one-region economy, with individual regions relative to the aggregate isomorphic to small open economies (Beraja et al., 2019). It is, however, well-known that standard small open economy models often do not admit the existence of a stationary distribution of bonds (Schmitt-Grohé & Uribe, 2003). In my quantitative model in Section 3, this is not a problem, since uninsurable income risk leads to a downward-sloping demand for liquidity. For models without income risk, a simple – and for my results inconsequential – fix is to assume that bond

Definition 1. *Given initial distributions $\mu^h_{k,0} = \bar{\mu}^h$ and $\mu^f_{k,0} = \bar{\mu}^f$ of households and intermediate goods producers over their idiosyncratic state spaces, initial real wages $w_{k,-1} = \bar{w}$, nominal aggregate prices p_{-1}, and real government debt $b_{-1} = \bar{b}$, as well as exogenous shock paths $\{\varepsilon_t\}_{t=0}^{\infty}$, a regional recursive competitive equilibrium is a sequence of regional quantities $\{c_{k,t}, c^H_{k,t}, c^{k'}_{k,t}, \ell^h_{k,t}, \ell^f_{k,t}, b^h_{k,t}, b^f_{k,t}, b_t, d_{k,t}, y_{k,t}, i_{k,t}, k_{k,t}, g_{k,t}, \tau_{k,t}\}_{t=0}^{\infty}$ and prices $\{\pi_t, i^b_t, w_{k,t}, p_{k,t}\}_{t=0}^{\infty}$ such that:*

1. Household Optimization. *Given prices and government rebates, the paths of home and foreign consumption $\mathbf{c}_k = \mathbf{c}_k(\mathbf{s}^h_k; \boldsymbol{\varepsilon}_k)$, $\mathbf{c}^H_k = \mathbf{c}^H_k(\mathbf{s}^h_k; \boldsymbol{\varepsilon}_k)$, $\mathbf{c}^{k'}_k = \mathbf{c}^{k'}_k(\mathbf{s}^h_k; \boldsymbol{\varepsilon}_k)$, total labor supply $\boldsymbol{\ell}^h_k = \boldsymbol{\ell}^h_k(\mathbf{s}^u_k)$, and asset holdings $\mathbf{b}^h_k = \mathbf{b}^h_k(\mathbf{s}^h_k; \boldsymbol{\varepsilon}_k)$ are consistent with optimal household and wage union behavior in every region k.*

2. Firm Optimization. *Given prices, the paths of total regional production $\mathbf{y}_k = \mathbf{y}_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$, investment $\mathbf{i}_k = \mathbf{i}_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$, capital $\mathbf{k}_k$, labor demand $\boldsymbol{\ell}^f_k = \boldsymbol{\ell}^f_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$, and asset holdings $\mathbf{b}^f_k = \mathbf{b}^f_k(\mathbf{s}^f_k; \boldsymbol{\varepsilon}_k)$ are consistent with optimal firm behavior.*

3. Government. *The liquid nominal rate is set in accordance with the monetary authority's Taylor rule. The government spending, rebate, and government debt issuance paths are jointly consistent with the government's budget constraint, its exogenous laws of motion for spending, and its financing rule.*

4. Market Clearing. *The regional goods and labor markets as well as the aggregate bond market all clear.*

A.2 Equilibrium characterization

I now show that the regional economy admits a particular kind of aggregation: All aggregates behave as in an analogous single-region economy, and all individual regions relative to the aggregate evolve as small open economies. This result builds on and extends the equilibrium characterization in Beraja et al. (2019) to a larger class of models.

Macro aggregates. The benchmark model of Wolf (2019) is the natural one-region analogue of the regional economy presented here. Households and firms in the single region have the same preferences and face the same constraints as households and firms in a region k of the regional model; the sole difference is that now there is a single composite consumption

holdings enter household discount factors, as in Beraja et al. (2019).

good c, obviating the need to keep track of real relative prices $p_{k,t}$ and region-specific price indices $q_{k,t}$.

My first result is that, at least to first order, aggregates in the regional economy behave like aggregates in this analogous one-region economy, buffeted by a common *average* aggregate shock $\boldsymbol{\varepsilon} \equiv \int_0^1 \boldsymbol{\varepsilon}_k dk$.

Lemma A.1. *Consider the structural model of Section 2.1. Suppose that, for each one-time single shock* $\{b, g\}$*, the equilibrium transition path exists and is unique. Then, to first order, all aggregate quantities* $\{c_t, \ell_t, b_t, y_t, i_t, k_t, g_t, \tau_t\}$ *and all prices* $\{\pi_t, i_t^b, w_t\}$ *evolve exactly as in a perfect foresight transition path of the analogous one-region model, subject to the exogenous aggregate shock path* $\boldsymbol{\varepsilon} \equiv \int_0^1 \boldsymbol{\varepsilon}_k dk$.

Regional economies. The second result is that, again to first order, all regional quantities evolve, relative to economy-wide averages, exactly as if each region was a small open economy, facing exogenous paths of nominal rates $\mathbf{i}^b$ in a foreign currency with inflation $\boldsymbol{\pi}$, lump-sum rebates $\boldsymbol{\tau}_k$, total demand from abroad equal to $\mathbf{g}_k + \mathbf{c}_k^F(\mathbf{p}_k)$, and a local shock $\boldsymbol{\xi}_k$.

Lemma A.2. *Consider the structural model of Section 2.1. Suppose that, for each one-time single shock* $\{b, g\}$*, the equilibrium transition path exists and is unique. Then, to first order, all relative regional quantities* $\{\tilde{c}_{k,t}, \tilde{\ell}_{k,t}, \tilde{b}_{k,t}, \tilde{y}_{k,t}, \tilde{i}_{k,t}, \tilde{k}_{k,t}\}$ *and prices* $\{\tilde{w}_{k,t}, p_{k,t}\}$ *evolve exactly as in a perfect foresight transition path of the analogous small open economy hit by one-time single regional shocks* $\{b, g\}$.

A.3 Estimation & parameterization

The structural model underlying the analysis of Section 3 is a special case of the rich model class outlined in detail in Appendix A.1. By Lemma A.1, the model aggregates to a conventional one-region economy. I choose all functional forms so that this implied aggregate representation is identical to the benchmark HANK model of Wolf (2019, Section A.2.2). In particular, I allow for rich household heterogeneity, but consider a simple firm block, close to canonical New Keynesian models (Smets & Wouters, 2007; Justiniano et al., 2010).

Steady State. I calibrate all steady-state parameters exactly as in Wolf (2019, Table 1). The only missing parameters are those governing the degree of home bias and the substitutability between different goods in household consumption baskets. Since I interpret a region in my economy to correspond to a county, I set $\phi = 0.61$, consistent with Dupor et al. (2018). For goods substitutability between regions, I follow Nakamura & Steinsson (2014) and set $\eta_r = 2$; for simplicity I also set $\eta_f = 2$.

Dynamics. Wolf (2019) estimates the parameters of the model governing dynamics by matching the time series properties of U.S. macroeconomic data. I do not follow such a likelihood approach, but instead directly match empirical evidence on local and aggregate government spending transmission – the moments that my theory has identified as near sufficient statistics for the desired macro counterfactuals. The only parameter not disciplined in this way is the degree of nominal wage stickiness; given its centrality for my equivalence results, I directly discipline it from micro data. In particular, I show in Wolf (2019) that the slope of the wage-NKPC (A.1) can be equivalently written as

$$\kappa_w = \frac{(1 - \frac{1}{1+\bar{r}}\phi_w)(1-\phi_w)}{\phi_w(\varepsilon_w \frac{1}{\varphi} + 1)}$$

where $1 - \phi_w$ is the probability of wage adjustment in the quarter. I set $\phi_w = 0.6$, consistent with the micro evidence in Grigsby et al. (2019) and Beraja et al. (2019). For all other parameters, I fix the same prior distributions as in Wolf (2019), but then estimate parameters to match as well as possible aggregate and regional government spending impulse responses (as in Christiano et al., 2005). Estimated posterior mode values are displayed in Table 1.

Dynamics Parameter Values, HANK Model

Coefficient	**Description**	**Value**
ϕ_p	Price Calvo Parameter	0.72
ζ	Capacity Utilization	3.21
κ	Investment Adjustment Cost	1.73
ρ_m	Taylor Rule Persistence	0.71
ϕ_π	Taylor Rule Inflation	2.23
ϕ_y	Taylor Rule Output	0.22
ϕ_{dy}	Taylor Rule Output Growth	0.31
ρ_b	Debt Persistence	0.77

Table 1: HANK model, parameters governing dynamics, estimated using moment-matching.

Relative to the estimates in Wolf (2019), the monetary authority is somewhat more aggressive in leaning against increases in aggregate output, and prices are somewhat less sticky. Given high average MPCs, more aggressive responses and less nominal rigidity are needed to keep aggregate consumption multipliers around 0 and aggregate fiscal multipliers around 1.

It remains to set the path for the shock to household borrowing conditions. For simplicity, I assume that both the borrowing limit and the wedge follow the same path in log deviations from steady state. In particular, both measures of financial constraints tighten by 50 per cent for one year, and then gradually return to baseline, with a quarterly AR(1) persistence of 0.6. The size of the shock is matched to give cross-regional consumption decline estimates of around 3 per cent, while the persistence is designed to match the length of the Great Recession. Results for other choices of the time profile of the shock are qualitatively similar and available upon request.

A.4 A two-region model

All results survive almost without change in a model with finitely many and heterogeneously-sized regions. Without loss of generality, I illustrate the argument in a variant of the two-region model studied in Nakamura & Steinsson (2014).

Model sketch. I index the two regions as H and F. The relative size of region H is $\zeta \in (0,1)$. The index of consumption for a generic household i in region H satisfies

$$c_{i,H,t} \equiv \left[\phi_H^{\frac{1}{\eta}} (c_{i,H,t}^H)^{\frac{\eta-1}{\eta}} + (1-\phi_H)^{\frac{1}{\eta}} (c_{i,H,t}^F)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$$

Preferences of households in region F are defined analogously, with $\phi_F \equiv \phi_H \times \frac{\zeta}{1-\zeta}$. All within-region preference and production primitives are exactly as in my multi-region benchmark economy; across-region aggregates are defined as the straightforward aggregate

$$x_t \equiv \zeta x_{H,t} + (1-\zeta) x_{F,t}$$

for a generic quantity or price x. Finally, the output market-clearing condition for region H becomes

$$y_{H,t} = c_{H,t}^H + \frac{1-\zeta}{\zeta} c_{F,t}^H + i_{H,t} + g_{H,t}$$

where all region-wide aggregates are defined as in the benchmark economy. The market-clearing condition for the foreign output good is analogous.

Results. Propositions 1, 2 and 4 generalize almost without change; I only state results and sketch proof strategies here.

First, the cross-regional regression (2), with observations weighted by the relative size of the region, now estimates across-region relative impulse responses of local shocks:

$$\beta_{b,h} = \hat{c}_{b_H,H,h} - \hat{c}_{b_H,F,h}$$

To see this, note that, by the properties of the standard fixed effects estimator, the left-hand side of the regression (2) is equal to

$$c_{b_H,H,h} - (\zeta \times c_{b_H,H,h} + (1-\zeta) \times c_{b_H,F,h}) = (1-\zeta)(c_{b_H,H,h} - c_{b_H,F,h})$$

Similarly, the right-hand side shock covariate is $(1-\zeta)(\varepsilon_{b_H,H,h} - \varepsilon_{b_H,F,h})$, and so the conclusion follows.

Second, local general equilibrium effects can still be removed by differencing out the effects of local government spending shocks. The argument is now slightly more involved: We have

$$\hat{c}^{PE}_{b_k,k,h} = \beta_{b,h} - \beta_{g,h}$$

where the government spending shock is normalized to affect government spending as $\frac{\phi-\zeta}{\phi(1-\zeta)} \times \phi \times \hat{c}^{PE}_{b_k,k,h}$. The logic underlying the slightly adjusted scaling factor is as follows. In partial equilibrium, the consumption demand shock increases spending on both the home and the foreign good. To proxy for (relative) local general equilibrium effects, we need to look at a local government spending shock that increases *relative* spending on the home good by the same amount. The proposed correction achieves exactly that. Also note that, for either $\zeta = 0$ (a vanishingly small region) or $\phi = 1$ (full home bias) the correction is again identical to that of the benchmark case.

Third, the economy again admits a regional invariance result; the proof strategy is almost exactly identical to that of Lemma A.1. We can thus recover aggregate counterfactuals via

$$\hat{\mathbf{c}}_b = \left[\zeta \hat{\mathbf{c}}^{PE}_{b,H} + (1-\zeta)\hat{\mathbf{c}}^{PE}_{b,F}\right] + \hat{\mathbf{c}}_g$$

where $\hat{\mathbf{g}}_g = \zeta \hat{\mathbf{c}}^{PE}_{b,H} + (1-\zeta)\hat{\mathbf{c}}^{PE}_{b,F}$.

B Proofs and auxiliary lemmas

B.1 Proof of Lemma A.1

The proof proceeds in two steps. First, I show that, to first order, the aggregated solutions of the regional household and firm problems are identical to the solution of household and firm problems in the analogous single-region economy. Second, I show that market-clearing in all regions implies that the corresponding aggregate markets must clear. This verifies all conditions in the equilibrium definition for a single-region economy (see Wolf (2019)), completing the proof.

1. Recall that, for all regions k, optimal regional consumption satisfies

   $$\mathbf{c}_k = \mathbf{c}_k(\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}_k, \mathbf{d}_k, \boldsymbol{\tau}_k, \mathbf{p}_k; \boldsymbol{\varepsilon}_k)$$

   Thus, to first order, we have

   $$\hat{\mathbf{c}}_k = \frac{\partial \mathbf{c}_k}{\partial \mathbf{i}^b} \times \hat{\mathbf{i}}^b + \frac{\partial \mathbf{c}_k}{\partial \boldsymbol{\pi}} \times \hat{\boldsymbol{\pi}} + \frac{\partial \mathbf{c}_k}{\partial \mathbf{w}_k} \times \hat{\mathbf{w}}_k + \frac{\partial \mathbf{c}_k}{\partial \mathbf{d}_k} \times \hat{\mathbf{d}}_k + \frac{\partial \mathbf{c}_k}{\partial \boldsymbol{\tau}_k} \times \hat{\boldsymbol{\tau}}_k + \frac{\partial \mathbf{c}_k}{\partial \mathbf{p}_k} \times \hat{\mathbf{p}}_k + \frac{\partial \mathbf{c}_k}{\partial \boldsymbol{\varepsilon}_k} \times \boldsymbol{\varepsilon}_k$$

   By the symmetry of regions, all partial derivative matrices are the same. With $\hat{\mathbf{w}} \equiv \int_0^1 \hat{\mathbf{w}}_k dk$, $\hat{\mathbf{d}} \equiv \int_0^1 \hat{\mathbf{d}}_k dk$, $\hat{\boldsymbol{\tau}} \equiv \int_0^1 \hat{\boldsymbol{\tau}}_k dk$, $\int_0^1 \hat{\mathbf{p}}_k dk = 0$ and $\boldsymbol{\varepsilon} \equiv \int_0^1 \boldsymbol{\varepsilon}_k dk$, we thus conclude that, to first order,

   $$\int_0^1 \hat{\mathbf{c}}_k(\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}_k, \mathbf{d}_k, \boldsymbol{\tau}_k \mathbf{p}_k; \boldsymbol{\varepsilon}_k) dk \quad = \quad \hat{\mathbf{c}}(\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \mathbf{d}, \boldsymbol{\tau}; \boldsymbol{\varepsilon})$$

   where the function on the right-hand side is the optimal aggregate consumption function in the one-region analogue economy. Exactly analogous results apply for the aggregate labor supply function and the aggregate household liquid savings function. Similarly, symmetry across producer blocks in the different regions also implies identical partial derivatives, so the same argument as above applies to aggregate output supply, investment demand, labor demand, dividend pay-out, and corporate savings functions.

   Overall, we have thus verified that aggregate quantities and prices in the regional economy are also consistent with optimal firm and household behavior in an analogous one-region economy. Since fiscal and monetary aggregates agree by construction, it remains to verify that all markets clear.

2. From definition of the regional equilibrium, we know that all regional goods and labor

markets as well as the aggregate bond market clear. Aggregating all across regions, we overall have

$$\begin{aligned} y_t &= c_t + i_t + g_t \\ \ell_t^h &= \ell_t^f \\ b_t^h + b_t^f &= b_t \end{aligned}$$

for all t. But all aggregates agree (to first order) with optimally demanded and supplied quantities in the single-region economy, so markets also clear (to first order) in the aggregate economy.

□

B.2 Proof of Lemma A.2

The proof proceeds in two steps. First, I show that, to first order, relative regional quantities are identical to the solution of household and firm problems in the analogous small open economy. Second, I show that equilibrium relative quantities satisfy all equilibrium market-clearing conditions of the small open economy. Thus all requirements of a small open economy equilibrium are satisfied, completing the argument.

1. Consider relative regional consumption $\tilde{\mathbf{c}}_k \equiv \mathbf{c}_k - \mathbf{c}$. Following the same steps as in the proof of Lemma A.1 we can establish that, to first order,

$$\tilde{\mathbf{c}}_k = \tilde{\mathbf{c}}_k(\tilde{\mathbf{w}}_k, \tilde{\mathbf{d}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k)$$

where $\tilde{\mathbf{c}}_k(\bullet)$ is the optimal consumption function for the small open economy. Exactly analogous arguments apply to local consumption of the local good as well as local labor supply. Applying the same logic to relative foreign consumption of the region-k output good, we conclude that, again to first order,

$$\tilde{\mathbf{c}}_k^A = \tilde{\mathbf{c}}_k^A(\mathbf{p}_k)$$

where $\tilde{\mathbf{c}}_k^A(\bullet)$ is the foreign consumption demand function for the small open economy k. Finally, a similar argument establishes that relative investment demand, production, labor demand, and dividend pay-outs are given as functions only of $(\tilde{\mathbf{w}}_k, \mathbf{p}_k)$, and in particular equal to optimal firm decisions in a small open economy.

We have thus verified that relative quantities in the regional economy are also consistent with optimal household and firm behavior in the analogous small open economy. It remains to verify that all markets of the small open economy clear.

2. Since local markets clear by assumption, and since the corresponding aggregate markets clear by Lemma A.1, it is immediate that the local labor and output markets also clear in deviations from aggregates:

$$\begin{aligned} \tilde{\mathbf{c}}_k^H(\tilde{\mathbf{w}}_k, \tilde{\mathbf{d}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) + \tilde{\mathbf{c}}_k^A(\mathbf{p}_k) + \tilde{\mathbf{i}}_k(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) + \tilde{\mathbf{g}}_k(\tilde{\boldsymbol{\varepsilon}}_k) &= \tilde{\mathbf{y}}_k(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) \\ \tilde{\boldsymbol{\ell}}_k^h(\tilde{\mathbf{w}}_k, \tilde{\mathbf{d}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) &= \tilde{\boldsymbol{\ell}}_k^f(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) \end{aligned}$$

where

$$\tilde{\mathbf{d}}_k = \tilde{\mathbf{d}}_k(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k)$$

But these markets are all that needs to clear for equilibrium in the small open economy.

□

B.3 Proof of Proposition 1

By the standard properties of fixed-effects regression, we can re-write regression (2) as

$$\tilde{c}_{k,t+h} = \beta_{b,h} \times \xi_{b,k,t} + \tilde{u}_{k,t+h}$$

The desired conclusion is then immediate by Lemma A.2. □

B.4 Proof of Proposition 2

For all of the following, I use the shorthand notation I use the notation $\frac{\partial}{\partial \boldsymbol{\varepsilon}_k}$ to refer to derivatives for shock sequences where only shocks in region k are non-zero, and $\frac{\partial}{\partial \boldsymbol{\varepsilon}_{s_k,k}}$ for specific shocks s_k in region k.

By Lemma A.2, impulse responses to local deleveraging and government spending shocks are characterized via the system

$$\begin{aligned} \tilde{\mathbf{c}}_k^H(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) + \tilde{\mathbf{c}}_k^A(\mathbf{p}_k) + \tilde{\mathbf{i}}_k(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) + \tilde{\mathbf{g}}_k(\tilde{\boldsymbol{\varepsilon}}_k) &= \tilde{\mathbf{y}}_k(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) \\ \tilde{\boldsymbol{\ell}}_k^h(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) &= \tilde{\boldsymbol{\ell}}_k^f(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k) \end{aligned}$$

where I have dropped dependence on dividend payments, since $\tilde{\mathbf{d}} = \tilde{\mathbf{d}}(\tilde{\mathbf{w}}_k, \mathbf{p}_k; \tilde{\boldsymbol{\varepsilon}}_k)$. Now let $\tilde{x}_{k,t} = (\tilde{w}_{k,t}, \tilde{p}_{k,t})$. Then, to first order, the equilibrium is characterized by the linear system

$$\begin{pmatrix} \frac{\partial \tilde{\mathbf{y}}_k}{\partial \tilde{\mathbf{x}}_k} - \frac{\partial \tilde{\mathbf{c}}_k^H}{\partial \tilde{\mathbf{x}}_k} - \frac{\partial \tilde{\mathbf{i}}_k}{\partial \tilde{\mathbf{x}}_k} \\ \frac{\partial \tilde{\boldsymbol{\ell}}_k^f}{\partial \tilde{\mathbf{x}}_k} - \frac{\partial \tilde{\boldsymbol{\ell}}_k^h}{\partial \tilde{\mathbf{x}}_k} \end{pmatrix} \times \tilde{\mathbf{x}}_k = \begin{pmatrix} \frac{\partial \tilde{\mathbf{c}}_k^H}{\partial \tilde{\boldsymbol{\varepsilon}}_k} + \frac{\partial \tilde{\mathbf{g}}_k}{\partial \tilde{\boldsymbol{\varepsilon}}_k} \\ \frac{\partial \tilde{\boldsymbol{\ell}}_k^h}{\partial \tilde{\boldsymbol{\varepsilon}}_k} \end{pmatrix} \times \tilde{\boldsymbol{\varepsilon}}_k \tag{B.1}$$

where I have already used the fact that neither government spending nor deleveraging shocks affect firm decisions in partial equilibrium.

Under Assumption 1, $\frac{\partial \tilde{\boldsymbol{\ell}}_k^h}{\partial \tilde{\boldsymbol{\varepsilon}}_k} = 0$. Next, since $\frac{\partial \tilde{\mathbf{c}}_k^H}{\partial \tilde{\boldsymbol{\varepsilon}}_k} = \phi \frac{\partial \tilde{\mathbf{c}}_k}{\partial \tilde{\boldsymbol{\varepsilon}}_k}$, the assumption that $\hat{g}_{g_k,k,h} = \phi \times \hat{c}^{PE}_{b_k,k,h}$ implies that

$$\frac{\partial \tilde{\mathbf{c}}_k^H}{\partial \tilde{\boldsymbol{\varepsilon}}_{b_k,k}} \times \tilde{\boldsymbol{\varepsilon}}_{b_k,k} = \frac{\partial \tilde{\mathbf{g}}_k}{\partial \tilde{\boldsymbol{\varepsilon}}_{g_k,k}} \times \tilde{\boldsymbol{\varepsilon}}_{g_k,k}$$

It is thus immediate from (B.1) that both shocks induce the same prices responses, and so indeed

$$\tilde{\mathbf{c}}_{b_k,k} = \tilde{\mathbf{c}}^{PE}_{b_k,k} + \tilde{\mathbf{c}}_{g_k,k}$$

where $\tilde{\mathbf{c}}^{PE}_{b_k,k} \equiv \frac{\partial \tilde{\mathbf{c}}_k}{\partial \tilde{\boldsymbol{\varepsilon}}_{b_k,k}} \times \tilde{\boldsymbol{\varepsilon}}_{b_k,k}$. Finally, since a k-specific shock has no aggregate effects, we can immediately replace tildes (indicating deviations from aggregates) with hats (indicating deviations from steady state). $\square$

B.5 Proof of Proposition 3

With full home bias and without investment, we have

$$\frac{\hat{c}_{b_k,k,0}}{m_{g_k,k}} = \hat{c}_{b_k,k,0} \backslash \frac{\hat{c}_{g_k,k,0} + \hat{g}_{g_k,k,0}}{\hat{g}_{g_k,k,0}}$$

Under Assumption 1 and if $\hat{g}_{g_k,k,0} = \hat{c}^{PE}_{b_k,k,0}$, then $\hat{c}_{g_k,k,0} = \hat{c}_{b_k,k,0} - \hat{c}^{PE}_{b_k,k,0}$ and so

$$\frac{\hat{c}_{b_k,k,0}}{m_{g_k,k}} = \hat{c}_{b_k,k,0} \backslash \left[1 + \frac{\hat{c}_{b_k,k,0} - \hat{c}^{PE}_{b_k,k,0}}{\hat{c}^{PE}_{b_k,k,0}}\right] = \hat{c}^{PE}_{b_k,k,0}$$

as claimed. $\square$

For completeness I also provide the static (impact) identification argument in more general models, as summarized in (6). We have

$$\hat{c}_{b_k,k,0} = \hat{c}^{PE}_{b_k,k,0} + \hat{c}_{g_k,k,0}$$

where

$$\hat{c}_{g_k,k,0} = m^c_{g_k,k} \times \hat{g}_{g_k,k,0} = m^c_{g_k,k} \times \phi \times \hat{c}^{PE}_{b_k,k,0}$$

Re-arranging, (6) follows.

B.6 Proof of Proposition 4

By Lemma A.1, the response of aggregate consumption to the composite shock $\{\boldsymbol{\varepsilon}_{b,k}\}$ is identical to the response of aggregate consumption to an aggregate shock $\boldsymbol{\varepsilon}_b \equiv \int_0^1 \boldsymbol{\varepsilon}_{b,k} dk$ in the analogous one-region economy. By the demand equivalence result in Wolf (2019), we know that

$$\hat{\mathbf{c}}_b = \hat{\mathbf{c}}_b^{PE} + \hat{\mathbf{c}}_g$$

where $\hat{\mathbf{c}}_b^{PE}$ is the partial equilibrium response to the shock $\boldsymbol{\varepsilon}_b$, and $\hat{\mathbf{c}}_g$ is the response of aggregate consumption to a spending shock with $\mathbf{g}_g = \hat{\mathbf{c}}_b^{PE}$. But, following the same steps as in the proof of Lemma A.1, we also know that, to first order, $\hat{\mathbf{c}}_b^{PE} = \int_0^1 \hat{\mathbf{c}}^{PE}_{b_k,k}$, completing the argument. □