Missing Intercept, Biased Slope: Identification in General Equilibrium*

Ed Manuel†

September 14, 2026

Abstract

Many studies use panel regressions leveraging units' heterogeneous exposures to aggregate shocks to identify the "direct" (partial equilibrium) effects of macroeconomic shocks. I consider identification of these parameters within an environment that features agent-heterogeneity, dynamics and general equilibrium effects. I show that identification in these settings is complicated by the presence of agents' heterogeneous responses to general equilibrium effects of the shock under study. I demonstrate a route to identification by adding control variables that capture the targeting rule of the exposures across units. I further show that this approach can be used to recover estimates of *intertemporal* direct effects – key objects for modern macroeconomic models. I illustrate my approach in applications estimating local fiscal multipliers and housing wealth effects.

Key Words: Identification; General Equilibrium; Spillovers

JEL Codes: C21, C23, E10, E20, E60.

*I thank Marco Bellifemine, Gabriel Chodorow-Reich, Jonathon Hazell, Ethan Ilzetzki, Xavier Jaravel, Ben Moll, Ricardo Reis, Silvana Tenreyro and Christian Wolf for helpful comments. All errors remain my own.

†London School of Economics. Email Address: `e.manuel@lse.ac.uk`.

1 Introduction

An increasingly popular approach in empirical macroeconomics uses cross-sectional data to identify the direct (or "partial equilibrium") effects of macroeconomic shocks. This includes classic studies estimating households' marginal propensities to consume (Parker et al., 2013), regional fiscal multipliers (Nakamura and Steinsson, 2014) or firms' responses to tax cuts (Zwick and Mahon, 2017). It is well-understood that this approach in general suffers from the so-called "missing intercept" problem: cross-sectional data typically does not contain sufficient information to recover the total (or "general equilibrium") effect of the shock under study. However, the approach remains popular since the direct response of individual units to macroeconomic shocks can provide information on key parameters in macroeconomic models. Estimates of these "direct" effects thus serve as a powerful source of indirect evidence on questions of central interest such as the effectiveness of fiscal stimulus (see e.g. Nakamura and Steinsson, 2018; Chodorow-Reich, 2019).

In this paper, I study methods for reliably achieving identification of these objects. I do so in a comprehensive setup that incorporates key features of modern macroeconomic models: dynamics, forward-looking expectations, heterogeneous agents and general equilibrium effects. I focus discussion on identification via panel "shift-share" regressions that seek to leverage units' heterogeneous exposures to aggregate shocks via:

$$y_{i,t+h} = \alpha_i + \delta_t + \beta(s_i \cdot d_t) + \text{controls} + u_{i,t} \tag{1}$$

where $y_{i,t+h}$ is the outcome variable of interest for unit i (such as household consumption or regional output) h periods from time-t, d_t is a time-t macroeconomic shock, s_i is a measure of the "exposure" of unit-i to the shock under study, and α_i and δ_t denote unit- and time-fixed effects respectively.[1] To fix ideas, I focus on identification of the direct effects of government stimulus – e.g. households' marginal propensities to consume (MPCs) out of government transfers or regional government-spending multipliers. My key contribution is to provide conditions on the exposure-shares s_i and the shock-variable d_t such that, alongside an appropriate choice of control variables, β captures the (average) direct response of individual units to fiscal stimulus shocks. Although shift-share identification has been well-studied in the microeconometrics literature (see e.g. Borusyak et al., 2025, for a practical discussion), I depart from this previous work by studying identification within an environment that maps directly to (a class of) heterogeneous agent dynamic stochastic general equilibrium (HA-DSGE) models – exactly the macroeconomic models which cross-sectionally identified moments are typically intended to be informative about.

[1] Such specifications are common in macroeconomics – e.g. Nakamura and Steinsson (2014); Pennings (2021); Guren et al. (2021b); Chodorow-Reich et al. (2021)– as well as other fields – e.g. Dube and Vargas (2013); Nunn and Qian (2014).

Identification Results I present a series of results on identification via panel regressions of the form (1) within my environment. Previous literature highlights two distinct routes to identification: either via exogeneity of the shocks d_t, or via exogeneity of the shares s_i. My initial results are negative. I show that, absent any exogeneity restriction on the shares s_i, the shift-share regression (1) generally fails to identify the direct effects of government stimulus shocks since β is contaminated by the differential response of treatment and control groups to the GE effects of the shock. This holds even in the experimental ideal where the researcher has direct access to the structural shock of interest – e.g. even if they are able to observe a series of essentially random, unanticipated adjustments in government stimulus at the national level over time. The problem arises since, in GE, changes in government spending impact individuals via various channels – not just via the stimulus directly but also via changes in e.g. interest rates, taxes, wages – and so (1) generally picks up the heterogeneous responses of units to *all* of these transmission channels simultaneously. Crucially, heterogeneity across units implies that inclusion of time fixed effects in (1) does not serve to ensure these GE responses simply "wash out" in the cross-section. These negative results echo some previous discussion on identification in macroeconomic settings (e.g. Koby and Wolf (2020); Donaldson (2026)). I stress that various approaches to correct for this issue through additional control variables that seek to capture these GE spillovers directly do not overcome the problem – either because they are not robust to the presence of *dynamic* channels of GE transmission, or because they are not robust to the presence of heterogeneity in units' responses to their *own* treatment.

I have two key contributions. First, I demonstrate a route to identification that is robust to these concerns. My approach views the exposure shares s_i as being implicitly targeted by the policymaker based on some set of underlying unit-level characteristics. For example, government social security transfers are naturally targeted to households with lower levels of income; and likewise national military spending is naturally targeted to regions with larger high-tech manufacturing sectors, and thus with a larger proportion of employment in the tradable sector. I then show that identification can be achieved by controlling appropriately for any unit-level characteristics that feature in the underlying "targeting rule" and determine units' responsiveness to GE objects. In the shift-share regression (1), this requires adding as controls the interaction of these unit-level variables with the time-series shock of interest. Intuitively, this approach ensures that β from regression (1) compares only across units that (on average) respond the same to any GE effects arising from the aggregate shock d_t. My key result demonstrates that this approach successfully recovers a (convex-) weighted average of agents' direct effects, the typical target of applied work in settings with heterogeneous treatment effects.

Second, I show how to leverage this strategy to recover agents' responses to different "dynamic treatments". For example, I demonstrate how to recover estimates of households' *intertemporal* marginal propensities to consume (iMPCs), key objects for macroeconomic models (Auclert et al., 2024). Intuitively, after partialling out the targeting rule, the exposure shares in (1) pin down the particular "channel of transmission" of interest – and so any underlying macroeconomic shock that operates through that channel generates useful variation that can be exploited for identification. And so running (1) with the "treatment variable" on the left-hand side recovers the "dynamic treatment", and then running with the outcome variable recovers the associated effects. Since different macroeconomic shocks generally produce different paths for the treatment variable of interest, including various aggregate time-t disturbances as the shifter in (1) traces out the effects of various dynamic treatments – e.g. various "slices" of the iMPC matrix.

I extend these results to various settings. While my main results focus on identification via panel shift-share regressions, I demonstrate a similar route to identification in settings with purely cross-sectional data (a common feature of applied work). Further, while my baseline environment is motivated by HA-DSGE models where a continuum of agents respond heterogeneously to aggregate variables that adjust in GE, I extend discussion to settings with a finite number of agents who respond directly to the treatment of each other individual unit, as in (linearised) dynamic network models.

I stress that this route to identification remains robust in the face of various features that typically pose significant challenges for identification, e.g. heterogeneity, dynamics, forward-looking expectations, GE effects, etc. Further, this approach does not require the researcher to have any direct knowledge of the underlying GE effects – e.g. by how much interest rates or taxes adjust to the intervention under study – or even access to any well-identified structural shock. Nevertheless, the informational requirements are generally demanding, and crucially rely on the researcher having access to control variables that effectively span the various dimensions of unit-level heterogeneity that simultaneously determine treatment and drive heterogeneous responses to the GE effects of the policy intervention. This requirement is trivially satisfied in settings where stimulus was randomised across units (as in e.g. Parker et al., 2013), although it is generally harder to satisfy in empirical settings where variation in stimulus received by each unit lacks a clear "targeting rule" based on underlying unit characteristics

Quantitative Example I demonstrate my results within a Two-Agent New Keynesian (TANK) model. I start by considering a setting where the researcher has direct access to a series of government transfer shocks, and seeks to estimate household MPCs out of the transfer. I assume the probability of a household receiving the transfer depends on their hand-to-mouth status, and consider results for various targeting rules. I show that a shift-share regression of the form (1) that interacts the transfer shock with an indicator for whether the household receives the transfer can be heavily biased for MPCs. For example, when the transfers are targeted entirely to Ricardian households, for a standard parameterisation of the model, the shift-share regression estimates an on-impact MPC out of the transfer of around 1.0, even though the true MPC is around 0.01. In this case, the shock leads to a significant contraction in aggregate demand –

hand-to-mouth agents thus cut their consumption (given lower wages), while Ricardian agents increase consumption (given their intertemporal smoothing motive and lower interest rates). An econometrician running (1) in this setting picks up the difference in consumption responses across Ricardian and hand-to-mouth agents, and so may (incorrectly) conclude that targeted transfers to Ricardian agents are highly effective at stimulating consumption (at least directly). Within this model, the bias is avoided by simply controlling for the transfer shock interacted with an indicator for whether households are Ricardian or hand-to-mouth.

I then demonstrate how to use this same approach to trace out household iMPCs. Intuitively, after controlling appropriately for hand-to-mouth status, any macroeconomic shock that triggers an endogenous response of transfers provides sufficient variation to identify the MPC of the treated group. Thus, even lacking any well-identified structural shock, a researcher can learn about MPCs simply by comparing consumption responses to various 'news" about macroeconomic aggregates – e.g. news about future activity or inflation. Projecting on different news measures in the shift-share regression – e.g. news at different horizons or about different macroeconomic aggregates – is thus informative about an additional slice of households' iMPCs.

Empirical Applications I illustrate my framework through two applications. My first application demonstrates the quantitative importance of my results – in particular, the need for research designs that are robust to heterogeneous responses to GE effects. I revisit Nakamura and Steinsson (2014) who estimate local fiscal multipliers out of military spending using a shift-share research design, interacting a measure of national military spending with state-level exposure-shares. My results highlight that identification fails even when adjustments in national military spending are exogenous to economic developments (the common assumption from the time-series literature, e.g. Ramey, 2011). I then show that the exposure-shares in the shift-share regression from Nakamura and Steinsson (2014) are highly predictable with a set of variables that macroeconomic theory posits as important determinants of heterogeneity in regional responses to changes in aggregate conditions – e.g. measures of regional openness and MPCs (Bellifemine et al., 2023). Various approaches to selecting controls to purge this endogeneity in the exposure shares greatly reduce the estimated local fiscal multiplier. Without additional controls, I estimate a two-year multiplier of 1.5, rising to over 2.5 after four years. My preferred specifications instead yield a two-year multiplier below 1.0, rising to around 1.0 after four years. These latter estimates are more tightly estimated, significantly different from the no-controls estimate, and relatively stable across various sensitivity analyses.[2]

My second application highlights the practical usefulness and novelty of my framework. I revisit the application of Guren et al. (2021b) who estimate housing wealth effects – the re-

[2]The original study reports a range of static multipliers across different specifications, with a preferred estimate "of approximately 1.5". I build off the authors' shift-share specification for which they report a multiplier of 2.5 – towards the upper end of the range presented in the paper.

sponse of local retail employment to an increase in local house prices. The shares in this case are constructed based on the Saiz (2010) housing supply elasticity instrument. The measure is determined by a city's topographical features (e.g. the presence of bodies of water) that are plausibly exogenous to economic fundamentals (e.g. regional MPCs), as required by my identification results. I present a series of balance tests broadly consistent with this interpretation, and further demonstrate the robustness of my main results across a battery of sensitivity tests. As in Guren et al. (2021b), I find significant evidence of housing wealth effects. I then demonstrate how this same research design can be leveraged to identify the effects of different house price "treatment paths" on employment. I identify the effect of a steady, gradual decline in house prices alongside a sharper decline then reversal – tracing out the employment effects of each.

Literature This paper relates to various strands of literature. First, and most directly, I relate to a literature in macroeconomics on the appropriate interpretation of cross-sectional regression estimands from the perspective of macroeconomic models (Nakamura and Steinsson, 2014, 2018; Chodorow-Reich, 2019; Koby and Wolf, 2020; Chodorow-Reich, 2020; Guren et al., 2021a; Wolf, 2023). Traditional discussion of this issue considers settings where agents are homogeneous – see e.g. Nakamura and Steinsson (2014); Chodorow-Reich (2020); Guren et al. (2021a) – in which case the intercept term in cross-sectional regressions (or likewise the time fixed effect in panel regressions) partials out all relevant GE effects of policy interventions (and so the issues for identification discussed in this paper do not arise).

The fact that identification that leverages cross-unit variation in exposures to aggregate shocks is complicated by the presence of heterogeneous agents and GE effects has been discussed in prior work (Koby and Wolf, 2020; Wolf, 2023; Canova, 2024; Donaldson, 2026).[3] Wolf (2023) demonstrates a route to identification of MPCs in a HA-DSGE setting where the researcher has direct access to the interaction of randomly-assigned exposure-shares with a series of well-identified aggregate stimulus shocks – my key contribution is to demonstrate a route to identification when both these assumptions fail.[4] In recent work, Donaldson (2026) also considers identification via shift-share panel regressions in a setting where agents respond heterogeneously to GE effects. The two works are highly complementary, highlighting two distinct routes to identification. Donaldson (2026) focuses on an environment where agents are heterogeneously exposed to changes in a single GE variable (the interest rate), and shows that

[3]The core issues around identification discussed in this paper relate to settings where researchers seek to exploit variation stemming from *macroeconomic* shocks for identification. An alternative literature in macroeconomics instead exploits "*micro*" experiments for identification – as in e.g. classic studies estimating household MPCs from lottery wins (Fagereng et al., 2021) and randomised control trials (Boehm et al., 2025). Since the GE effects of such interventions are likely negligible, the issues for identification discussed in this paper do not arise – a point I make explicitly later.

[4]The core contribution in Wolf (2023) is to demonstrate an approach to *recover* the missing intercept, which requires combining evidence on MPCs with time-series evidence on the effects of a "demand-equivalent" shock. I am instead primarily concerned with the prior question of how to identify MPCs.

this bias can be stripped out by using time-series evidence on the effects of monetary policy shocks (following the method in McKay and Wolf (2023)). The two approaches can thus be understood as "cleaning" the shift-share interaction term in (1) along different margins: the approach in Donaldson (2026) seeks to "clean" the aggregate shock d_t, by offsetting its aggregate GE effects, while the route I discuss instead seeks to "clean" the exposure shares s_i. Unlike Donaldson (2026), I demonstrate a route to identification within an environment where agents respond heterogeneously to their own treatment as well as to multiple GE objects. This is an important feature since macroeconomic shocks typically transmit through multiple GE channels – e.g. not just via interest rates, but also prices, wages, taxes, etc. – and the forces that drive heterogeneous responses to any one of these channels – e.g. heterogeneity in units' MPCs – simultaneously drive heterogeneous responses to the other channels (including the researcher's particular channel of interest).[5]

This paper also closely relates to recent work on estimating heterogeneity in "micro" responses to macroeconomic shocks (Almuzara and Sancibrián, 2024). My environment of interest – a structural vector-moving average model with heterogeneity in unit-level impulse responses – is similar to Almuzara and Sancibrián (2024), and, as in that paper, I begin with a general characterisation of the estimand from shift-share regressions within that environment. I then use this characterisation to answer a different set of questions, focusing on identification rather than inference. Various other work also considers identification via panel data with heterogeneous units in macroeconomic settings – although their focus is on recovering the "missing intercept", i.e. on achieving identification of the total (not direct) effects of shocks on macroeconomic aggregates (Sarto, 2025; Matthes et al., 2025). My work is complementary to Matthes et al. (2025) in particular since it clarifies the appropriate approach for leveraging cross-sectional identification to recover the "direct" effects of macroeconomic shocks, which can then be used in conjunction with the authors' proposed method to recover the "total" effect of the shock by combining with time-series evidence.

I also relate to a broad literature on (panel) shift-share research designs (Adão et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022b; Arkhangelsky and Korovkin, 2024; Majerovitz and Sastry, 2023). My key contribution is to consider identification in a setting that explicitly incorporates key features of modern macroeconomic models – dynamics, forward-looking behaviour, general equilibrium effects and the presence of heterogeneous agents – which I show has important implications for identification.[6] My approach is closer in spirit to

[5] In Appendix C.1, I directly discuss and extend the approach from Donaldson (2026), including the informational requirements to implement this strategy within my environment.

[6] My proposed approach essentially leverages exogeneity of units' exposure shares, as in Goldsmith-Pinkham et al. (2020). More recently, Arkhangelsky and Korovkin (2024) propose a synthetic control method, which involves using a pre-sample period to construct a combination of units with high exposure shares that display similar patterns in outcomes to units with low exposure shares. The approach seeks to partial out units' heterogeneous exposures to aggregate confounders that co-move with the aggregate shock, which can be viewed as similar to the core issues for identification studied here. As the authors note, the method is not valid in the presence of dynamic treatment effects – I demonstrate this explicitly in a HA-DSGE setting in Appendix C.2.

Adão et al. (2025) and Borusyak et al. (2022a), who focus on the interpretation of shift-share regressions within the context of trade and migration models respectively. While motivated by similar concerns to this paper, the setting of interest and proposed solutions are very different – since the approach I explore is "design-based" (i.e. relying only on the researcher having knowledge of the treatment assignment across units), it: (a) does not rely on knowledge of the underlying GE model that generates the spillovers; and so (b) allows researchers to recover "identified moments" that are valid across a broad class of models (as in Nakamura and Steinsson (2018)).

A key building block underlying my analysis is a series of findings on cross-sectional identification in settings with heterogeneous treatment effects – see e.g. classic discussion of OLS and IV (Imbens and Angrist, 1994; Angrist, 1998), as well as more recent discussion on DiD (see e.g. de Chaisemartin and D'Haultfœuille (2023) for a review). In particular, my key result leverages insights from Borusyak and Hull (2024) who show that design-based regression-control specifications identify a convex weighted-average of treatment effects. I show how to leverage this insight to achieve identification of objects of interest within a HA-DSGE setting.

Finally, although discussion throughout is focused on the macroeconomics literature (and on identifying the effects of fiscal stimulus in particular), my results may also be informative for applications across other areas of economics. For example, applications across finance (e.g. Chodorow-Reich, 2014), labour (e.g. Acemoglu and Restrepo, 2020), trade (e.g. Autor et al., 2013) and development (e.g. Egger et al., 2022; Galego Mendes et al., 2024; Gerard et al., 2025) similarly employ cross-sectional identification in settings with potentially sizable spillover effects via GE forces, and seek to interpret resulting estimates as capturing the "direct", "local", or "partial equilibrium" effect.

Outline The rest of the paper is outlined as follows. Section 2 presents the key issues discussed in the paper through the lens of a simple setting. Section 3 sets out my general environment and contains my key results on identification. Section 4 provides an illustration of my results in a TANK model. Section 5 contains empirical applications. Section 6 concludes.

2 Simple Example

I begin with an example to highlight the intuition underlying the core results in this paper. I consider an econometrician who inhabits a TANK model and is interested in estimating household MPCs. I relegate all details of the model to later sections, providing only a very high-level overview for now.

Setting The economy is populated by a unit-continuum of households: $i \in [0, 1]$. A fraction $\alpha \in (0, 1)$ are unconstrained ($\mathcal{U}$), able to borrow and save in a risk-free bond at rate r_t; the remaining fraction are hand-to-mouth ($\mathcal{H}$), consuming all their income each period. Households receive labour income at wage w_t and dividend income d_t, pay lump-sum taxes τ_t and receive transfers $\mathcal{T}_{i,t}$. I assume taxes are levied evenly across all agents, while dividend income is received only by the unconstrained agents.[7] Each unconstrained household $i \in [0, \alpha]$ solves a familiar dynamic consumption-savings problem:

$$\begin{gathered} \max_{\{c_{i,t}, n_{i,t}, b_{i,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{(c_{i,t})^{(1-\sigma)}}{1-\sigma} - \theta \frac{(n_{i,t})^{(1+\chi)}}{1+\chi} \right] \\ \text{s.t. } c_{i,t} + b_{i,t} \leq w_t n_{i,t} + d_t - \tau_t + r_t b_{i,t-1} + \mathcal{T}_{i,t}, \forall t \end{gathered} \tag{2}$$

Each hand-to-mouth household $i \in [\alpha, 1]$ solves the static problem each period:

$$\begin{gathered} \max_{c_{i,t}, n_{i,t}} \left[\frac{(c_{i,t})^{(1-\sigma)}}{1-\sigma} - \theta \frac{(n_{i,t})^{(1+\chi)}}{1+\chi} \right] \\ \text{s.t. } c_{i,t} \leq w_t n_{i,t} - \tau_t + \mathcal{T}_{i,t} \end{gathered} \tag{3}$$

The rest of the economy is as-standard in the New Keynesian framework: monopolistically competitive intermediate good producers face sticky prices, and monetary policy is set according to a Taylor Rule. The model is linearised around a zero-transfers steady state.

Experiment Consider an econometrician who observes steady-state outcomes for all units at period $t-1$, then observes outcomes over time following an unexpected one-unit shock to aggregate transfers, $\varepsilon_t^s = 1$, with no other (contemporaneous or future) shocks hitting the economy. Further, suppose the transfer received by each agent following the shock depends on their "exposure" to aggregate transfers $\mathcal{T}_t$ following:

$$\mathcal{T}_{i,t} = z_i \mathcal{T}_t$$

where z_i denotes a binary exposure measure. Suppose these exposures are distributed across the population according to some "targeting rule":

$$Pr(z_i = 1) = \delta_{\mathcal{U}} \cdot \mathcal{I}(i \in \mathcal{U}) + \delta_{\mathcal{H}} \cdot \mathcal{I}(i \in \mathcal{H}) \tag{4}$$

where $\mathcal{I}(i \in \mathcal{U})$ and $\mathcal{I}(i \in \mathcal{H})$ are binary indicators taking on the value of one when household-i is unconstrained and hand-to-mouth respectively, and $\delta_{\mathcal{U}}$ and $\delta_{\mathcal{H}}$ capture respective probabilities of receiving the transfer for each group.

[7] None of the substantive discussion here depends on these assumptions.

Identification Suppose the econometrician runs difference-in-differences regressions of the form:

$$\Delta c_{i,t+h} = \alpha_h + \beta_h^{DiD} z_i + u_{i,t+h} \tag{5}$$

where $\Delta c_{i,t+h} = c_{i,t+h} - c_{i,t-1}$ so that β_h^{DiD} computes the difference in consumption responses to the shock across treatment and control households. What does β_h^{DiD} identify in this case? It is well-known that, in general, β_h^{DiD} does not capture the *total* response of consumption to the shock (inclusive of all general equilibrium effects). Nevertheless, the typical goal of cross-sectional regressions is to recover an estimate of the *direct* (or partial equilibrium) effect of the shock. To make these notions precise, following e.g. Wolf (2023), it is useful to first express households' optimal consumption choices over time as a function of time variation in prices and quantities faced by the household. In this model, the variables taken as given by the household include their own transfer $\mathcal{T}_{i,t}$, alongside aggregate wages w_t, interest rates r_t, dividends d_t and taxes τ_t. Letting boldface denote sequences of outcomes: $\boldsymbol{x}_{i,t} = \{x_{i,t+h}\}_{h=0}^{\infty}$, household consumption choices over time can thus be expressed as:

$$\boldsymbol{c_{i,t}} = \boldsymbol{\mathcal{C}}_i(z_i \boldsymbol{\mathcal{T}}_t, \boldsymbol{r}_t, \boldsymbol{d}_t, \boldsymbol{\tau}_t, \boldsymbol{w}_t)$$

where $\boldsymbol{\mathcal{C}_i}$ is the consumption function for household-i, mapping sequences of prices and quantities to a sequence of consumption choices. Next, let $\Theta_{i,h}^{xs}$ denote the impulse response of outcome variable x for unit i at horizon h to the transfer shock ε_t^s, and let boldface denote sequences of impulse responses: $\boldsymbol{\Theta}_i^{xs} = \{\Theta_{i,h}^{xs}\}_{h=0}^{\infty}$. The *total* effect of the shock on *aggregate* consumption can then be defined as:

$$\boldsymbol{\Theta}^{cs} = \int_i \boldsymbol{\Theta}_i^{cs} di \tag{6}$$

where $\boldsymbol{\Theta}_i^{cs} = \dfrac{d\boldsymbol{\mathcal{C}_i}}{d\varepsilon_t^s}$, with the derivative evaluated at unit-i's steady-state. The *direct* effect on *each unit-i* can then be defined from the following decomposition:

$$\boldsymbol{\Theta}_i^c s = \underbrace{\mathcal{M}_i^{\mathcal{T}} \boldsymbol{\Theta}^{\mathcal{T}s} z_i}_{\text{Direct Effect}} + \underbrace{\mathcal{M}_i^{w} \boldsymbol{\Theta}^{ws} + \mathcal{M}_i^{r} \boldsymbol{\Theta}^{rs} + \mathcal{M}_i^{d} \boldsymbol{\Theta}^{ds} + \mathcal{M}_i^{\tau} \boldsymbol{\Theta}^{\tau s}}_{\text{GE Effect}} \tag{7}$$

where $\mathcal{M}_i^x$ denotes the *partial* derivative $\dfrac{\partial \boldsymbol{\mathcal{C}_i}}{\partial \boldsymbol{x}_t}$ evaluated at steady-state.[8] The "direct effect" captures how consumption for unit-i responds at each point in time to a "treatment path" of transfers traced out by $\boldsymbol{\Theta}^{\mathcal{T}s}$, where $\mathcal{M}_i^{\mathcal{T}}$ captures the intertemporal MPC matrix for unit-i (Auclert et al., 2024). The "GE effect" instead captures the response of unit-i to other aggregate variables that adjust in general equilibrium to the shock – in this case, summarised by their combined response to wages w_t, interest rates r_t, dividends d_t and taxes τ_t. Given MPCs are

[8]I provide closed-form expressions for the various $\mathcal{M}$-mappings in Appendix A.

heterogeneous in this setting, the typical goal of regressions of the form (5) is to identify some average measure of direct effects – i.e. to recover:

$$\int_i \omega_i \mathcal{M}_i^{\mathcal{T}} \mathbf{\Theta}^{\mathcal{T}s} di, \quad \text{for:} \quad \int_i \omega_i di = 1 \;\text{ and }\; \forall i : \omega_i \geq 0 \tag{8}$$

where the requirement that the ω_i's sum to one and are non-negative ensures that the object is indeed interpretable as a (convex)-weighted average of underlying MPCs. Note the key "parallel trends" and "no anticipation" assumptions that underpin differences-in-difference analysis hold exactly in this case: the shock is unanticipated by construction and, absent the shock (i.e. setting $\varepsilon_t^s = 0$), all units would simply remain at their respective steady states. Nevertheless, β_h^{DiD} does not in general have an interpretation as capturing household MPCs. For example, when the transfers are perfectly targeted towards the unconstrained households ($\delta_{\mathcal{U}} = 1$, $\delta_{\mathcal{H}} = 0$), we have the following:[9]

$$\boldsymbol{\beta}^{DiD} = \underbrace{\mathcal{M}_{\mathcal{U}}^{\mathcal{T}} \mathbf{\Theta}^{\mathcal{T}s}}_{\text{ATT}} + \underbrace{[\mathcal{M}_{\mathcal{U}}^{w} - \mathcal{M}_{\mathcal{H}}^{w}]\, \mathbf{\Theta}^{ws} + [\mathcal{M}_{\mathcal{U}}^{r} - \mathcal{M}_{\mathcal{H}}^{r}]\, \mathbf{\Theta}^{rs} + \left[\mathcal{M}_{\mathcal{U}}^{d} - \mathcal{M}_{\mathcal{H}}^{d}\right] \mathbf{\Theta}^{ds} + [\mathcal{M}_{\mathcal{U}}^{\mathcal{T}} - \mathcal{M}_{\mathcal{H}}^{\mathcal{T}}]\, \mathbf{\Theta}^{\mathcal{T}s}}_{\text{Bias from differential GE responses}}$$

where $\boldsymbol{\beta}^{DiD} = \{\beta_h^{DiD}\}_{h=0}^{\infty}$, and $\mathcal{M}_{\mathcal{U}}^{x}$ and $\mathcal{M}_{\mathcal{H}}^{x}$ denote partial derivatives for unconstrained and hand-to-mouth agents respectively. The first term is exactly the object of interest – the MPC out of the transfer for recipient households – and so the second term can be interpreted as generating bias. In this case, $\boldsymbol{\beta}^{DiD}$ not only fails to capture the *total* response of aggregate consumption to the shock, inclusive of all GE channels, but it also fails to capture the *direct* effect of the shock on household consumption, i.e. household MPCs. It thus suffers from a "missing intercept" as well as a "biased slope" – hence the title.

The bias arises due to a failure of the Stable Unit Treatment Value Assumption (SUTVA). In particular, note that the consumption function of each unit depends not only on their *own* treatment status (z_i), but also on the treatment status of all other households (indirectly, via GE forces). Note the bias does not arise in settings where agents are homogeneous: i.e. when, for all aggregate variables x and households i, $\mathcal{M}_i^x = \mathcal{M}^x$. Similarly, the bias does not arise when there are no GE adjustments to the aggregate shock – i.e. when, for all aggregate variables x, $\mathbf{\Theta}^x = \mathbf{0}$. From the perspective of the model, this latter assumption holds when the transfers are targeted towards a sufficiently small fraction of households – i.e. when the econometrician has access to an appropriate "*micro*" experiment.[10] Note this assumption plausibly holds in classic studies exploiting "small-scale" interventions for identification – e.g. estimates of household MPCs from lottery wins (Fagereng et al., 2021) and randomised control trials (Boehm et al.,

[9]Note that, with no other shocks hitting the economy, the difference-in-difference estimand is simply the difference in impulse responses across treated (unconstrained) and control (hand-to-mouth) households: $\boldsymbol{\beta}^{DiD} = \mathbb{E}[\mathbf{\Theta}_i^{cs} \mid z_i = 1] - \mathbb{E}[\mathbf{\Theta}_i^{cs} \mid z_i = 0]$. The result then follows plugging in the decomposition (7).

[10]I come on to demonstrate this later with a quantitative example.

2025). The purpose of this paper however is to discuss cross-sectional identification in settings where variation stems from *macroeconomic* shocks, a common feature of applied work.

Why would this bias term be problematic in practice? As I come on to demonstrate, under a standard parameterisation, when the shock is targeted entirely to unconstrained households: $\beta_0^{DiD}/\Theta_0^{\mathcal{T}s} \approx 1.0$. For the econometrician, it is tempting to conclude that this implies the transfer shock was effective at directly stimulating household consumption – i.e. that, on average, recipient households spent all of the transfer they received. This is the wrong conclusion since unconstrained households' on-impact MPCs are instead: $\left[\mathcal{M}_{\mathcal{U}}^{\mathcal{T}}\boldsymbol{\Theta}^{\mathcal{T}s}\right]_0/\Theta_0^{\mathcal{T}s} \approx 0.01$.[11]

How might researchers ensure that their empirical specification is robust to this concern? Suppose instead that transfers are imperfectly targeted – i.e. $\delta_g \in (0,1)$ for at least one $g \in \{\mathcal{U},\mathcal{H}\}$ – and consider a regression specification that additionally controls for households' hand-to-mouth status:

$$\Delta c_{i,t+h} = \alpha_h + \beta_h^{DiD} z_i + \lambda_h \mathcal{I}(i \in \mathcal{U}) + u_{i,t+h} \tag{9}$$

Then it can be shown that β_h^{DiD} now recovers the object of interest:[12]

$$\boldsymbol{\beta} = \omega_{\mathcal{U}}\mathcal{M}_{\mathcal{U}}^{\mathcal{T}}\boldsymbol{\Theta}^{\mathcal{T}s} + \omega_{\mathcal{H}}\mathcal{M}_{\mathcal{H}}^{\mathcal{T}}\boldsymbol{\Theta}^{\mathcal{T}s}, \tag{10}$$

for: $\omega_{\mathcal{H}}, \omega_{\mathcal{U}} \geq 0$ and $\omega_{\mathcal{H}} + \omega_{\mathcal{U}} = 1$. In this case, $\boldsymbol{\beta}^{DiD}$ is constructed by comparing consumption within household-types, then taking variance-weighted averages of such differences – the classic result of Angrist (1998). In turn, this ensures the GE effects appropriately "wash out" in the cross-section, since $\boldsymbol{\beta}^{DiD}$ compares across agents who respond the same to the GE effects of the shock. The key for identification is that (9) effectively controls for the characteristics that appear in the "targeting rule" from (4). In this simple model, doing so is straightforward since agents differ only on a single dimension, their hand-to-mouth status. As I come on to demonstrate, the result goes through with richer forms of heterogeneity, so long as the econometrician continues to control for any unit characteristics that appear in (4). Appealingly, controlling for the targeting rule in this way does not require direct knowledge of the propagation of the shock via GE channels – e.g. how much wages and interest rates actually adjusted to the shock, or how units' characteristics shape their responsiveness to these changes. Further, since identification is "design-based" – leveraging knowledge of the variables that appear in the "targeting

[11] As I discuss in Appendix A, the MPC of hand-to-mouth agents in this calibration lies strictly below one due to income effects on labour supply (as hand-to-mouth agents choose to "spend" some of the transfer on leisure time). So $\beta_0^{DiD} \approx 1.0$ lies strictly outside the range of any households' MPC in this case. Note that a high-β^{DiD} could even be consistent with an alternate model in which *all* households behave according to the permanent income hypothesis – and transfers are financed via distortionary (not lump-sum) taxes, which differentially affect treatment and control groups. This could occur if, for example, the control group were more sensitive to higher labour taxation (e.g. had a higher Frisch elasticity of labour supply) or bore a higher burden of the tax.

[12] This claim is nested by more general results provided later in the paper. I provide closed-form expressions for the ω-weights in Appendix A.

rule" rather than direct knowledge of units' responses to the GE effects – the resulting estimate is robust to the presence of heterogeneous treatment effects, as in Borusyak and Hull (2024).

Note further that, after effectively controlling for the targeting rule via (9), the actual source of the macroeconomic shock underpinning the variation in transfers is essentially irrelevant for the identification argument to go through. Suppose instead the econometrician observed some *other* shock occurring at time-t, say a cost-push shock ε_t^π, and let $\boldsymbol{\Theta}^{\mathcal{T}\pi}$ denote the impulse response of aggregate transfers to this other shock. In this case, since agents' consumption functions $\boldsymbol{\mathcal{C}}_i(.)$ depend only on (future) prices and quantities – and not the actual *source* of the shock – β_h^{DiD} in (9) captures the same weighted average of MPCs, but now for a different *path* for aggregate transfers, i.e.:[13]

$$\boldsymbol{\beta} = \omega_{\mathcal{U}} \mathcal{M}_{\mathcal{U}}^{\mathcal{T}} \boldsymbol{\Theta}^{\mathcal{T}\pi} + \omega_{\mathcal{H}} \mathcal{M}_{\mathcal{H}}^{\mathcal{T}} \boldsymbol{\Theta}^{\mathcal{T}\pi}$$

This is potentially useful, for two reasons. First it substantially weakens the informational requirements, since the econometrician can now learn about household MPCs by leveraging variation from any aggregate shock. Second, leveraging variation across a variety of aggregate shocks allows the econometrician to learn about additional dimensions of MPCs – i.e. different "slices" of agents' intertemporal MPC matrices. I later demonstrate this explicitly with a quantitative model and an empirical application.

Discussion The discussion in this section is special in a number of regards. First, it relates to a particular structural environment – a TANK model with only a single (binary) dimension of heterogeneity. Second, I abstract from the presence of other aggregate shocks simultaneously affecting the macroeconomy. The next section is devoted to illustrating these issues for identification in a more general dynamic environment with heterogeneous agents.

3 Identification Results

This section presents my key results on cross-sectional identification in macroeconomic settings. Section 3.1 introduces the environment, Section 3.2 discusses the object of interest and Section 3.3 contains the identification results. I relegate proofs in this section to Appendix B.

3.1 Environment

I consider a setting with a population of agents, indexed by i observed over time $t = 0, 1, 2, ...$ Agents are (ex-ante) heterogeneous across some n_x-dimensional vector of (non-time-varying)

[13] Note this final step in the argument essentially leverages the notion of "instrument sufficiency", in the sense of McKay and Wolf (2023). Since households do not care *per se* whether transfers adjusted due to a transfer shock, or due to an endogenous response of the policymaker to some other shock hitting the economy, from the perspective of the econometrician, any underlying aggregate shock potentially provides useful variation for estimating MPCs.

characteristics, denoted by $x_i = [x_{1,i}, x_{2,i}, ..., x_{n_x,i}]'$, as well as some additional scalar "exposure share" measure z_i (discussed shortly). The economy is subject to n_ε independent, mean-zero, unit-variance stochastic shocks denoted by $\varepsilon_t = [\varepsilon_{1,t}, ..., \varepsilon_{n_\varepsilon,t}]'$. I treat the first shock as a particular shock of interest which I refer to as a government stimulus shock – I label this as $\varepsilon_{1,t} \equiv \varepsilon_t^s$ and collect remaining shocks in the $(n_\varepsilon - 1)$-dimensional vector ε_t^{-s}.

I assume there are n_Y aggregate (potentially observable) variables which I partition as $Y_t = [y_t, g_t, w_t]'$ where y_t is the outcome variable of interest, g_t denotes the aggregate 'treatment' variable – some form of government stimulus such as government spending or transfers – and w_t is a $(n_Y - 2)$-dimensional vector collecting all other aggregate variables. I denote corresponding unit-level variables: $Y_{i,t} = [y_{i,t}, g_{i,t}, w_{i,t}]'$.

I assume the treatment of interest is distributed arbitrarily across agents, with z_i denoting the "share" of government stimulus going to each agent, with their own treatment denoted $g_{i,t} = z_i g_t$.[14] Aggregate and unit-level variables are linear functions of all current and past shocks (as well as steady-state values):

$$Y_t = \sum_{l=0}^{\infty} \Theta_l \varepsilon_{t-l} + \bar{Y} \tag{11}$$

$$Y_{i,t} = \sum_{l=0}^{\infty} \Theta_{i,l} \varepsilon_{t-l} + \bar{Y}_i \tag{12}$$

where Θ_l and $\Theta_{i,l}$ are $n_Y \times n_\varepsilon$-dimensional matrices that map shocks to aggregate and unit-level outcomes respectively, and $\bar{Y}$ and $\bar{Y}_i$ are n_Y-dimensional vectors denoting steady-state outcomes. It is well-known that the Structural Vector-Moving Average (SVMA) model (11) encompasses all discrete-time, linearised DSGE models, where (12) then extends this structure to the unit-level outcomes allowing agents to respond heterogeneously to the aggregate shocks.[15] I assume the unit-level impulse responses $\Theta_{i,l}$ and steady-states $\bar{Y}_i$ are measurable with respect to (x_i, z_i). As well as the standard assumption that shocks are mutually independent and independent of all past and future shocks, I assume they are independent of unit-level variables:

$$\forall k, l : \varepsilon_{k,t-l} \perp (\{\varepsilon_{j,t-l}\}_{j \neq k}, \{\varepsilon_{t-m}\}_{m \neq l}, z_i, x_i)$$

where this aligns with the interpretation of the shocks as primitive macroeconomic disturbances, as well as the unit-level variables as fixed characteristics that do not respond to these disturbances over time. As previewed in the previous sections, I am interested in the

[14] I do not restrict z_i's to be literal shares i.e. to be positive and integrate to one over the population. I later relax the requirement that $g_{i,t} = z_i g_t$, allowing for essentially arbitrary mappings between unit-level treatments and the aggregate treatment.

[15] This structure maps directly to the TANK model introduced in the previous section – as well as other DSGE models where unit-level heterogeneity is not time-varying (see e.g. Herreño and Pedemonte (2022) and Bellifemine et al. (2023)). I later extend the framework to settings with time-varying heterogeneity (as in e.g. canonical "HANK" models).

identification of the "direct" (or partial equilibrium) – as opposed to the "indirect" (or general equilibrium) – effects of macroeconomic shocks. To distinguish between these concepts, I thus assume the unit-level impulse response functions are decomposable into different channels. I focus on the dynamic causal effect of government stimulus shocks on the outcome variable of interest for each unit, which I label $\Theta^{ys}_{i,h} \equiv \Theta_{i,h,1,1}$. I assume these effects admit the following "sequence-space" (e.g. Auclert et al. (2021)) decomposition:

$$\boldsymbol{\Theta}^{ys}_i = \underbrace{\mathcal{M}^{yg}_i \boldsymbol{\Theta}^{gs} z_i}_{\text{Response to own treatment path}} + \underbrace{\mathcal{M}^{yw}_i \boldsymbol{\Theta}^{ws}}_{\text{Response to GE channels}} \tag{13}$$

where $\boldsymbol{\Theta}^{ys}_i = \{\Theta^{ys}_{i,h}\}_{h=0}^{\infty}$, $\boldsymbol{\Theta}^{gs} = \{\Theta^{gs}_h\}_{h=0}^{\infty}$ and $\boldsymbol{\Theta}^{ws} = \{\Theta^{ws}_h\}_{h=0}^{\infty}$ are sequences collecting the dynamic causal effects of ε^s_t on $y_{i,t}$, g_t and w_t, and $\mathcal{M}^{yg}_i$ and $\mathcal{M}^{yw}_i$ are conformable and map bounded sequences into bounded sequences, and are measurable with respect to x_i only:

$$\mathcal{M}^{yg}_i = \mathcal{M}^{yg}(x_i), \;\; \mathcal{M}^{yw}_i = \mathcal{M}^{yw}(x_i)$$

Decompositions of this form are common within macroeconomic models cast in the sequence-space to distinguish between different channels of transmission. In my setting, the first term can be understood as capturing the "direct effect" of the shock – since it captures unit i's responses to the path of their own stimulus $g_{i,t} = z_i g_t$, and the second term as the "GE effect" since it captures units' responses to other variables and is equal to the effect of the shock for units who have no direct exposure ($z_i = 0$). The substantive assumption underpinning this decomposition is that the 'direct' effects of the shock on each unit are separable from the 'GE effect', i.e. that receipt of the treatment does not also change units' responsiveness to the GE transmission channels of the shock. This assumption is satisfied in general in macroeconomic models linearised around their steady state, so long as changes in exposure shares leave unchanged units' deterministic steady states. Beyond this, the decomposition imposes relatively minor additional structure – allowing agents to respond arbitrarily to a range of transmission channels in a forward-looking way.

3.2 Object of Interest

The goal is to study the direct effect of government stimulus on each agent (i.e. holding fixed treatment for all other agents). To ease notation, I focus on the effect at the h'th horizon, $\Theta^{ys}_{i,h}$, and following (13), decompose this as:

$$\Theta^{ys}_{i,h} = \underbrace{\beta^{y}_{i,h}}_{\text{Direct Effect}} z_i + \underbrace{\zeta^{y}_{i,h}}_{\text{GE Effect}} \tag{14}$$

where $\beta^{y}_{i,h}$ and $\zeta^{y}_{i,h}$ correspond to the h'th term of the sequences $\mathcal{M}^{yg}_i \boldsymbol{\Theta}^{gs}$ and $\mathcal{M}^{yw}_i \boldsymbol{\Theta}^{ws}$

respectively. As in the previous section, my object of interest is some weighted average of my notion of direct effects – i.e. $\mathbb{E}[\omega_i \beta^y_{i,h}]/\mathbb{E}[\omega_i]$ for some non-negative weights ω_i.

As highlighted in Section 2, treating individual units as households, the "treatment" $g_{i,t}$ as government transfers, and the outcome variable $y_{i,t}$ as consumption, my notion of the direct effect $\beta^y_{i,h}$ corresponds to a particular entry of household i's intertemporal MPC matrix. Alternatively, in a setting with government purchases $g_{i,t}$ across regions, $\beta^y_{i,h}$ captures the response of each region to a "local" government spending shock, holding fixed aggregate variables (e.g. interest rates, taxes and aggregate demand for tradables) – the conventional notion of a local fiscal multiplier.[16] Instead, in a setting with government subsidies $g_{i,t}$ across firms, $\beta^y_{i,h}$ can be interpreted as e.g. the response of firm output to the subsidy, holding fixed aggregate conditions that do not vary with firm-specific shocks (e.g. industry-level price indices and aggregate demand conditions).

3.3 Identification

My focus is on identification via shift-share panel regressions.

Proposition 1 *Consider shift-share regressions of the form:*

$$y_{i,t+h} = \beta^{ss}(s_i \cdot d_t) + e_{i,t} \tag{15}$$

where s_i is independent of the shocks (jointly with other unit-level variables) and $d_t = \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}$ for some $1 \times n_\varepsilon$ vector Ψ_l. Then:

$$\beta^{ss} = \frac{\mathbb{E}\left[\left(\sum_{l=0}^{\infty} \Theta^y_{i,h+l} \eta_l\right) s_i\right]}{\mathbb{E}[s_i^2]}$$

where η_l is an $n_\varepsilon \times 1$ vector with k'th element equal to: $\eta_{l,k} = \dfrac{\psi_{l,k}}{\sum_{k=1}^{n_\varepsilon} \sum_{l=0}^{\infty} \psi_{l,k}^2}$.

Proof See Appendix B.1.

The proposition provides a formula for the shift-share coefficient for an arbitrary choice of "shares" s_i and arbitrary "shifter" d_t. This result follows the characterisation of the population estimand from shift-share regressions in a similar environment in Almuzara and Sancibrián (2024) but extended to allow for an arbitrary choice of shifter. The only requirement on the shifter d_t is that it reflect a linear combination of contemporaneous and past shocks – equivalent in this environment to allowing the shifter to be any (de-meaned) observable aggregate variable determined by time-t.

[16] Note in this case my notion of "direct effects" also includes the adjustment of *local* prices that occur in response to shocks to individual regions. In regional settings my definition thus differs somewhat from Wolf (2023) who uses this terminology exclusively for the partial equilibrium effects of interventions holding *all* prices fixed.

The result serves to demonstrate that the shift-share coefficient can be constructed via a two-step procedure. First, one recovers the unit-level impulse responses to the "shifter" d_t.[17] The shift-share coefficient is then constructed by regressing these unit-level impulse responses on the chosen shares s_i. Note that for binary shares, $s_i \in \{0, 1\}$, and a single structural shock, $d_t = \varepsilon_{k,t}$, when the shift-share regression includes time fixed effects, the resulting estimate simply takes the difference of impulse responses between treatment and control groups:

$$\beta^{ss} = \mathbb{E}[\Theta^y_{i,h,k}|s_i = 1] - \mathbb{E}[\Theta^y_{i,h,k}|s_i = 0]$$

I now consider routes to identification. The previous literature has stressed two distinct routes for shift-share identification – either via exogeneity of the heterogeneous exposure shares (Goldsmith-Pinkham et al., 2020) or exogeneity of the aggregate shock (Adão et al., 2019; Borusyak et al., 2022b). Within my environment, a natural interpretation of the first route involves imposing that the exposure of each unit to changes in government stimulus z_i is independent of any unit characteristics x_i:

$$z_i \perp x_i$$

This mimics an experimental ideal where the exposures z_i are randomly assigned across the population. My initial question is whether identification can be achieved without imposing this condition; in particular, whether identification can be achieved only through access to an appropriate aggregate shock. A natural starting point is thus to consider a case where the researcher has direct access to the structural stimulus shock of interest, ε^s_t. Since ε^s_t is independent of all other structural shocks, this mimics an alternate experimental ideal where the researcher is able to directly observe a series of random (and unanticipated) adjustments in government spending at the national level over time.[18]

3.3.1 Identification From Shocks

I start by considering identification when the researcher has direct access to a series of government spending shocks ε^s_t alongside the underlying exposure share measures z_i, but absent any exogeneity condition on these exposure shares (i.e. allowing z_i to be correlated with x_i).

[17] Since I leave open that d_t is a combination of different shocks, strictly speaking, these are linear combinations of impulse responses to the underlying *structural* shocks that underpin d_t. This combination of impulse responses reflects the same linear combination of IRFs that one would recover from (time-series) projections of any of the aggregate or unit-level outcome variables on d_t.

[18] See also Stock and Watson (2018) on the interpretation of structural shocks as the macroeconomic, time-series counterpart of a randomised controlled experiment.

Proposition 2 *Consider a setting where $z_i \in \{0,1\}$. Consider shift-share regressions of the form:*

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss}(z_i \cdot \varepsilon_t^s) + e_{i,t} \tag{16}$$

Then the shift-share coefficient β^{ss} can be expressed as:

$$\beta^{ss} = \underbrace{\mathbb{E}[\beta_{i,h}^y | z_i = 1]}_{\text{ATT}} + \underbrace{\mathbb{E}[\zeta_{i,h}^y | z_i = 1] - \mathbb{E}[\zeta_{i,h}^y | z_i = 0]}_{\text{Bias from GE responses}}$$

Proof See Appendix B.2

Exactly as in the simple example from Section 2, the shift-share coefficient here takes the difference in impulse responses across treatment and control groups, and so is contaminated by the differential responses of each group to the GE effects of the shock. The result focuses on the case with binary exposure shares. Outside of this case, it is straightforward to show that β^{ss} in regression (16) suffers from an *additional* issue. In particular, even absent any GE effects – e.g. assuming $\mathbf{\Theta}^w = 0$ – the shift-share coefficient is generally not robust to the presence of heterogeneity in the direct effects $\beta_{i,h}^y$ – an issue that has been discussed extensively in the applied microeconometrics literature on two-way fixed effects specifications (see e.g. de Chaisemartin and D'Haultfœuille, 2023).[19]

The bias from Proposition 2 arises because agents respond heterogeneously to aggregate variables w_t that adjust in response to the shock. A natural correction might therefore be to allow agents to load differentially on these aggregate variables w_t by including additional interaction terms in the regression.[20] The following clarifies that this is not in general sufficient for identification:

Proposition 3 *Consider a setting where $z_i \in \{0,1\}$. Consider shift-share regressions of the form:*

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss}(z_i \cdot \varepsilon_t^s) + \lambda'(z_i \cdot w_t) + e_{i,t} \tag{17}$$

[19] Absent GE effects, β^{ss} captures a weighted average of the direct effects $\beta_{i,h}^y$, but with weights given by $z_i(z_i - \mathbb{E}[z_i])$. In cases where z_i is non-binary and strictly greater than zero – as is common in applied work when z_i reflects fractions of an aggregate – $z_i(z_i - \mathbb{E}[z_i])$ generally takes on negative values for at least some units.

[20] This approach to dealing with this concern is common in applied work – see e.g. Nakamura and Steinsson (2014), Pennings (2021) and Chodorow-Reich et al. (2021). Donaldson (2026) refers to regressions of the form (17) as the "control function" approach and similarly demonstrates bias within their environment.

Then the shift-share coefficient β^{ss} has the form:

$$\beta^{ss} = \underbrace{\mathbb{E}[\beta^y_{i,h}|z_i = 1]}_{\text{ATT}} + \underbrace{\Big(\mathbb{E}[\zeta^y_{i,h}|z_i = 1] - \mathbb{E}[\zeta^y_{i,h}|z_i = 0]\Big)}_{\text{Bias from GE responses}}$$

$$+ \underbrace{\sum_{k=2}^{n_\varepsilon} \kappa_{0,k} \Big(\mathbb{E}[\Theta^y_{i,h,k}|z_i = 1] - \mathbb{E}[\Theta^y_{i,h,k}|z_i = 0]\Big)}_{\text{Bias from other contemp. shocks}} + \underbrace{\sum_{l=1}^{\infty}\sum_{k=1}^{n_\varepsilon} \kappa_{l,k} \Big(\mathbb{E}[\Theta^y_{i,h+l,k}|z_i = 1] - \mathbb{E}[\Theta^y_{i,h+l,k}|z_i = 0]\Big)}_{\text{Bias from other lagged shocks}}$$

where the κ-loadings are given by:

$$\kappa_{l,k} = \frac{\psi_{l,k}}{\sum_{k=1}^{n_\varepsilon}\sum_{l=0}^{\infty}\psi^2_{l,k}}, \qquad \psi_{0,1} = 1 + b'\Theta^w_{0,1}, \qquad \forall (l,k) \neq (0,1) : \psi_{l,k} = b'\Theta^w_{l,k}$$

and:

$$b' = -(\Theta^w_{0,1})'\,\Sigma_w^{-1}, \qquad \Sigma_w = \sum_{l=0}^{\infty} \Theta^w_l (\Theta^w_l)'$$

Proof See Appendix B.3

Proposition 3 characterises the bias in this case purely in terms of the (macro) and (micro) structural impulse responses Θ^w_l and $\Theta^y_{i,l}$. By Frisch–Waugh–Lovell, the shift-share coefficient in this case is constructed by first orthogonalising the shock with respect to w_t (and a constant), and then using this orthogonalised shock $(\varepsilon^s_t)^{\perp w_t}$ as the shifter. Since orthogonalising ε^s_t with respect to w_t recovers a linear combination of all shocks up to time-t, the resulting shift-share coefficient now additionally captures the differential response of treatment and control groups to all (past and contemporaneous) shocks, where the κ-loading on each term depends on the particular combination of shocks captured by $(\varepsilon^s_t)^{\perp w_t}$.

One way to understand the bias here is that the controls fail to appropriately "hold fixed" the relevant GE variables that drive heterogeneous responses across units. Returning to the sequence-space decomposition (13), note that the responses to GE variables in this environment are captured by units' responses to dynamic *paths* for aggregate variables triggered by the shock – i.e. by units' responses to $\boldsymbol{\Theta}^{ws}$ – which in general are not captured by the additional control variables in (17). One way to remove the bias term from Proposition 2 is to follow the logic of the solution proposed in Donaldson (2026) and subtract off units' responses to a linear combination of time-t shocks that generates the same path for GE variables as ε^s_t – thus effectively holding fixed the entire dynamic path for w_t. I discuss this procedure explicitly in Appendix C.1. In my environment, the informational requirements to implement this procedure are somewhat stringent: requiring the researcher to have access both to the full set of aggregates that adjust in GE, w_t, and to a sufficiently broad number of macroeconomic "news" shocks that serve to simultaneously hold fixed each of these variables across all horizons.[21]

Before proceeding, I briefly consider another strategy to partial out units' differential responses to GE variables. To investigate this, I consider a situation where the researcher seeks to directly model the differential GE responses of agents to the shock – i.e. $\zeta^y_{i,h}$ in equation (14). This speaks to a large empirical literature that seeks to simultaneously identify the direct effects of aggregate shocks on each unit alongside identifying heterogeneous spillover effects (see e.g. Huber, 2022). In this case, researchers typically run a variant of (16) where they additionally include a measure of units' heterogeneous responses to the shock via spillover channels. To speak to this, I consider the idealised setting where the researcher has direct access to these heterogeneous spillover effects:

Proposition 4 *Consider shift-share regressions of the form:*

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss}(z_i \cdot \varepsilon^s_t) + \lambda(\zeta^y_{i,h} \cdot \varepsilon^s_t) + e_{i,t} \tag{18}$$

Then the shift-share coefficient β^{ss} has the form:

$$\beta^{ss} = \frac{\mathbb{E}[\beta^y_{i,h}\omega_i]}{\mathbb{E}[\omega_i]}$$

where $\omega_i = \tilde{z}_i z_i$ and $\tilde{z}_i$ is the (population)-residual from a regression of z_i on a constant and $\zeta^y_{i,h}$.

Proof See Appendix B.4

In this case the bias term is removed, but the weights are given by: $\omega_i = \tilde{z}_i z_i$ – which can in general take on negative values. Hence the shift-share coefficient does not necessarily have an interpretation as capturing some (properly) weighted average of the underlying direct effects. A vast literature in microeconometrics discusses alternative approaches to identification that are robust to negatively weighted treatment effects. Since this particular route to identification anyway imposes exceptionally high informational requirements in my environment – requiring knowledge of the h'th element of the sequence $\mathcal{M}^{yw}_i \boldsymbol{\Theta}^{ws}$ from equation (13) – I instead consider other routes to identification.

3.3.2 Identification From Shares

I now consider whether identification can be achieved via the exogeneity of shares z_i. It is immediate that, with exogenous shares (i.e. assuming $z_i \perp x_i$), regression (16) no longer suffers

[21] In Appendix C.2, I also consider the synthetic control method recently proposed by Arkhangelsky and Korovkin (2024). As the authors note, the method is generally not appropriate in settings with dynamic treatment effects. I demonstrate this explicitly with an example in the TANK model introduced in the previous section.

from bias since, when the shares are orthogonal to unit characteristics, we have:

$$\mathbb{E}[\zeta^y_{i,h}|z_i = 1] - \mathbb{E}[\zeta^y_{i,h}|z_i = 0] = 0$$

This exogeneity condition is strong however and so I consider a weakening that permits identification via control variables even when z_i is correlated with x_i. In particular, I suppose the econometrician has access to some observable n_a-vector of control variables a_i that is independent of the shocks (jointly with the other unit-level variables) and captures the appropriate "targeting rule" for the exposure shares.

Proposition 5 *Suppose the shares are conditionally exogenous:*

$$\mathbb{E}[z_i|x_i, a_i] = \alpha + \delta' a_i \tag{19}$$

Consider shift-share regressions of the form:

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss}(z_i \cdot \varepsilon^s_t) + \lambda'(a_i \cdot \varepsilon^s_t) + e_{i,t} \tag{20}$$

Then:

$$\beta^{ss} = \frac{\mathbb{E}[\beta^y_{i,h}\omega_i]}{\mathbb{E}[\omega_i]}$$

where $\omega_i = \mathbb{E}[\tilde{z}_i^2|x_i, a_i]$ and $\tilde{z}_i$ is the (population)-residual from a regression of z_i on a constant and a_i.

Proof See Appendix B.5

The controls here serve to remove the bias term. As in Borusyak and Hull (2024), negative weights are avoided since the control variables are "design-based" – controlling for the dependency between z_i and x_i, rather than the actual spillover terms $\zeta^y_{i,h}$ as in Proposition 4. Further, as in Borusyak and Hull (2024), the representation for the weights is non-unique: β^{ss} can also be written as a weighted average of direct effects with weights given as $\tilde{z}_i^2$, in which case the weights can be directly estimated in the data.[22]

While it is well-known in general that exogeneity of the shares is a possible route to identification for shift-share regressions (Goldsmith-Pinkham et al., 2020), Proposition 5 demonstrates that this approach remains robust even in an environment with heterogeneous and dynamic treatment effects, forward-looking agents and numerous (and unknown) GE channels of transmission. The identifying assumptions are generally strong: they require the researcher to model correctly the assignment of the shares z_i and to include controls that capture this. In

[22]This follows from a direct application of the law of iterated expectations. I use this later in the empirical application to estimate these weights directly.

particular, the controls must span any unit characteristics that simultaneously determine treatment and heterogeneous responses to GE effects. Nevertheless, identification can be achieved in this case without any knowledge of the underlying GE channels underpinning the transmission of the shock, and the resulting regression is simple to implement.

Proposition 5 continues to rely on the researcher having access to a single well-identified shock ε_t^s. This is a strong assumption in practice since it requires the researcher to have direct access to a series of essentially random adjustments in government stimulus at the national level over time. I now show how this assumption can be meaningfully weakened.

Proposition 6 *Suppose the shares are conditionally exogenous as in (19) and that the shifter satisfies a "lead/lag exogeneity" condition:*

$$d_t = q'\varepsilon_t \text{ for } q \in \mathbb{R}^{n_\varepsilon \times 1} \tag{21}$$

And suppose the decomposition from (13) extends to all shocks:

$$\forall k \in \{1, ..., n_\varepsilon\} : \boldsymbol{\Theta}_{i,k}^y = \mathcal{M}_i^{yg}\boldsymbol{\Theta}_k^g z_i + \mathcal{M}_i^{yw}\boldsymbol{\Theta}_k^w \tag{22}$$

Consider the pair of shift-share local projections for $h = 0, 1, ...$:

$$y_{i,t+h} = \alpha_{i,y,h} + \delta_{t,y,h} + \beta_{y,h}^{ss}(z_i \cdot d_t) + \lambda_{y,h}'(a_i \cdot d_t) + u_{i,t,y,h}$$
$$g_{i,t+h} = \alpha_{i,g,h} + \delta_{t,g,h} + \beta_{g,h}^{ss}(z_i \cdot d_t) + \lambda_{g,h}'(a_i \cdot d_t) + u_{i,t,g,h}$$

Then:

$$\boldsymbol{\beta}_y^{ss} = \frac{\mathbb{E}[\mathcal{M}_i^{yg}\boldsymbol{\beta}_g^{ss}\omega_i]}{\mathbb{E}[\omega_i]}$$

for $\boldsymbol{\beta}_y^{ss} = \{\beta_{y,h}^{ss}\}_{h=0}^{\infty}$ and $\boldsymbol{\beta}_g^{ss} = \{\beta_{g,h}^{ss}\}_{h=0}^{\infty}$, where $\omega_i = \mathbb{E}[\tilde{z}_i^2|x_i, a_i]$ and $\tilde{z}_i$ is the (population)-residual from a regression of z_i on a constant and a_i.

Proof See Appendix B.6

This result demonstrates that the path traced out by the shift-share coefficient $\beta_{y,h}^{ss}$ can be interpreted as a convex-weighted average of units' responses to a dynamic treatment path for g_t (holding fixed the path for any GE-objects w_t), where that dynamic treatment path is traced out by the shift-share coefficient β_g^{ss}. Returning to a previous example, if $y_{i,t+h}$ denotes household consumption, and g_{t+h} denotes government transfers, then the proposition demonstrates that the shift-share coefficient captures some weighted average of household intertemporal MPCs (iMPCs) – i.e. how consumption responds dynamically to a dynamic transfer path traced out by $\boldsymbol{\beta}_g^{ss}$.

The result requires that the shifter d_t is a function of only contemporaneous shocks, a standard lead/lag exogeneity condition typically required for identification in dynamic settings (see Stock and Watson (2018)).[23] Within my environment, this allows the shifter to be any non-predictable variable determined by time-t – similar in spirit to a "no anticipation" assumption. The result also relies on extending the sequence-space decomposition (13) to *all* shocks. The substantive assumption is that the causal mappings $\mathcal{M}_i$ do not depend on the source of the shocks (i.e. they do not vary by k). This is closely related to the notion of "instrument sufficiency" in the macroeconomics literature on the mapping between policy shocks and policy rules (Barnichon and Mesters, 2023; McKay and Wolf, 2023). As discussed in McKay and Wolf (2023), a broad class of macroeconomic models satisfy instrument sufficiency since agent optimisation problems depend only on (expected) paths for policy variables (and not on the underlying source of the shock). Similarly, equation (22) imposes that individual units respond to changes in expected paths for $g_{i,t}$ and w_t in the same way regardless of the underlying source of the shock (i.e. whether it was a government stimulus shock per se, or an endogenous response of government stimulus to some other shock).

Further, an implicit – and practically important – assumption underpinning Proposition 6 is that the shifter is "relevant" in the sense that (in combination with the exposure measure) it in fact predicts differential responses in $g_{i,t}$ across units.[24] In practice, this speaks to researchers leveraging aggregate shocks that have significant explanatory power for their particular "channel" of interest.

The result is then useful for two reasons. First, it formalises how identification of the "direct" effects of some path for government stimulus can be achieved, even when the researcher lacks a well-identified measure of a stimulus shock – i.e. when identification via time-series methods with aggregate data is not possible. Second, and more importantly, it demonstrates that the researcher can recover IRFs for the outcome variable of interest to different treatment paths for g_t by varying the aggregate shifter d_t. This is potentially a fruitful avenue of exploration for applied work. As Auclert et al. (2024) demonstrate, the full iMPC matrix $\mathcal{M}_i^{yg}$ is a crucial object for heterogeneous agent models. Applied work is typically only informative about particular entries however – e.g. household responses over time to an unanticipated and contemporaneous rise in transfers (i.e. the first column of the iMPC matrix). Proposition 6 clarifies that, if the researcher is able to find units who are differentially exposed to shifts in government transfers for exogenous reasons (conditional on controls a_i), then their responses

[23] Note that one can interpret any of the "shocks" included in the vector ε_t as capturing exogenous "measurement error" by simply imposing that the structural impulse responses to these shocks are zero. It is then immediate that Proposition 6 nests the case where the researcher has access to a valid "proxy" for the stimulus shock, in the sense of Stock and Watson (2018). The result then goes further, essentially dropping the "contemporaneous exogeneity" condition typically required for time-series identification – akin to only requiring that the econometrician have access to valid "Wold innovations" rather than any structural shock.

[24] Absent this assumption, the shift-share regression returns $\boldsymbol{\beta}_g^{ss} = 0$, in which case the result still goes through but loses any practical usefulness.

to different macroeconomic news could in theory be used to identify the full iMPC matrix $\mathcal{M}_i^{yg}$. I provide a simple illustration of this point in the following section via a TANK model.

3.3.3 Extensions

I now discuss extensions to the baseline framework. I relegate all details to the Appendix and provide only a high-level overview of my key results. I begin with various natural extensions of my results then turn to two more substantial extensions to the baseline framework.

Proxy exposure measures My results until now assume that the researcher has direct access to the "true" exposure of each unit z_i to changes in government stimulus. In empirical settings, a more realistic assumption is that researchers have access to some "proxy" s_i that is correlated with the exposure shares z_i, while uncorrelated with other sources of unit-heterogeneity x_i. In Appendix C.3, I consider shift-share regressions with this proxy variable s_i and show that it recovers a measure of the underlying direct effects under similar assumptions as Proposition 5. In this case, negative weights on treatment effects are avoided via a "monotonicity" assumption that echoes those made in the microeconometrics literature on instrumental-variable identification (Imbens and Angrist, 1994), which ensures the proxy shares co-vary with the (true) exposure shares uniformly either positively or negatively across sub-groups in the population.

Dynamic treatment heterogeneity My baseline environment restricts all units to receive the same treatment dynamics, where the path is pinned down by Θ^{gs} then scaled by z_i. In Appendix C.4, I consider the more general case, where I do not restrict the profile of each unit's own treatment path, allowing some units to receive more front-loaded stimulus, while stimulus for other units is more back-loaded. I again suppose the econometrician has access to some "proxy" s_i, correlated with some elements of $\boldsymbol{\Theta}_i^{gs}$, but crucially uncorrelated with x_i. In this case, shift-share regressions are again informative about some weighted average of appropriate entries of $\mathcal{M}_i^{yg}$, with weights that are non-negative under a monotonicity condition that holds at every horizon.

Cross-sectional data My main results focus on identification via panel shift-share regressions. Many studies in macroeconomics employ purely cross-sectional data, and seek to identify the direct effects of a single, one-off shock (Mian and Sufi, 2012; Parker et al., 2013; Zwick and Mahon, 2017). In Appendix C.5, I show that my results extend naturally to this setting, where I focus on cross-sectional difference-in-differences regressions. I demonstrate that identification can be achieved so long as unit-level exposure to the shock under study is uncorrelated with any characteristics that determine heterogeneous responses to changes in macroeconomic aggregates.

Time-varying heterogeneity I now turn to more substantial extensions of the framework. My baseline environment imposes that sources of heterogeneity across units are fixed over time. In Appendix C.6, I extend this to consider a setting where agents face idiosyncratic shocks and also respond heterogeneously to macroeconomic shocks based on a (time-varying) vector of potentially endogenous (but predetermined) set of state variables. This speaks directly to a broad class of macroeconomic models, including (linearised) heterogeneous agent New Keynesian (HANK) models. I demonstrate that my key results go through essentially unchanged in this setting, with the additional assumption that the aggregate shock employed in the shift-share regression is independent of the idiosyncratic (or "local") shocks faced by each unit.

Finite-N I also consider settings with a finite number of agents, with shocks to individual units that "spill over" directly to other units. This speaks to identification in a general class of (linearised) dynamic "network" models. A full treatment of this case is beyond the scope of this paper, but in Appendix C.7, I show that, in this environment, a similar bias for two-way fixed effect regressions arises when individual shocks to each unit are correlated, and that a similar correction can be employed, by controlling for the underlying process that generates this correlation. I supplement these results in Appendix F with an additional quantitative model, in which a collection of finite units (e.g. US states) trade directly with each other and are subject to a set of government spending shocks. In this model, estimates of the local fiscal multiplier are biased when government spending shocks are targeted towards particular regions and trade follows a "gravity" equation. This bias arises even when government spending shocks are "exogenous" in the sense of being independent of *other* shocks hitting the economy. The bias arises since states with higher levels of government spending are simultaneously more exposed to higher government spending in states within the same region via increased demand for exports. A simple correction in this case is to control for the underlying treatment assignment that generated such regional clustering, where this can be as simple as including regional fixed effects as controls (even when spillovers operate beyond regional borders). Finally, I use the model alongside my econometric results to revisit evidence on local fiscal multipliers out of the American Recovery and Reinvestment Act (ARRA). Specifically, I revisit Chodorow-Reich (2019), who combines three studies estimating ARRA multipliers – Chodorow-Reich et al. (2012), Wilson (2012), Dupor and Mehkari (2016). Reassuringly, across a range of methods, I find that there is little evidence of substantial bias in local multiplier estimates across the three studies when taken together.

4 Quantitative Example

To demonstrate my results, I now revisit the Two-Agent New Keynesian model discussed in Section 2. I provide further details on the model in Section 4.1 and then provide quantitative results on identification in Section 4.2.

4.1 Overview

The household block is as described in Section 2 – I now provide details on remaining blocks.

Firms Production is identical to the standard New Keynesian model with a representative competitive final goods firm alongside a continuum, $j \in [0,1]$, of monopolistically competitive intermediate goods producers, the latter facing sticky prices. Each intermediate goods producer j operates a linear production technology that depends on labour $n_{j,t}$ as well as aggregate productivity A_t, where the latter follows an exogenous AR(1) process.

Policy Monetary policy is set according to the following Taylor Rule:

$$i_t = (1-\rho_i)\bar{i} + \rho_i i_{t-1} + (1-\rho_i)\phi_\pi(\log \Pi_t - \log \bar{\Pi}) + s_m \varepsilon_t^m. \tag{23}$$

where i_t denotes nominal interest rates, Π_t denotes inflation, and bars denote steady-state values. The government's budget constraint is:

$$G_t + \mathcal{T}_t \int_i z_i \, di = \tau_t \tag{24}$$

In the baseline model, the level of government spending G_t and aggregate transfers $\mathcal{T}_t$ follow exogenous AR(1) processes and so taxes τ_t are set to balance the budget each period.

Shocks Alongside transfer shocks ε_t^s, the economy is also subject to shocks to government spending ε_t^g, monetary policy ε_t^m and technology ε_t^A.

Shares As before, I consider a targeting rule for the binary exposure shares z_i following:

$$Pr(z_i = 1) = \delta_{\mathcal{U}} \cdot \mathcal{I}(i \in \mathcal{U}) + \delta_{\mathcal{H}} \cdot \mathcal{I}(i \in \mathcal{H})$$

where $\delta_{\mathcal{U}}$ and $\delta_{\mathcal{H}}$ reflect respective treatment probabilities for each household type. For the quantitative analysis, I restrict $\delta_{\mathcal{U}}$ to be binary: $\delta_{\mathcal{U}} \in \{0,1\}$, and allow $\delta_{\mathcal{H}}$ to take on values in the range $[0,1]$.

Linearisation The model is linearised around a zero-inflation, zero-transfers steady-state: $\log \bar{\Pi} = 0$, $\bar{\mathcal{T}} = 0$.

4.2 Identification

I consider an econometrician who seeks to recover an estimate of households' marginal propensities to consume out of a shock to government transfers, defined as in (7):

$$\boldsymbol{\beta}_i = \mathcal{M}_i^{\mathcal{T}} \boldsymbol{\Theta}^{\mathcal{T}s} \tag{25}$$

where $\mathcal{M}_i^{\mathcal{T}}$ denotes the iMPC matrix for unit-i – i.e. the partial derivative of household-i's consumption function with respect to transfers $\mathcal{T}_{i,t}$– and $\boldsymbol{\Theta}^{\mathcal{T}s}$ denotes the impulse response of aggregate transfers to the transfer shock, ε_t^s. Finally, I define $\boldsymbol{\beta}_{\mathcal{U}}$ and $\boldsymbol{\beta}_{\mathcal{H}}$ as the values of $\boldsymbol{\beta}_i$ for unconstrained and hand-to-mouth households respectively.

As in the previous section I consider different routes to identification. First, I consider shift-share regressions of the form:

$$c_{i,t+h} = \alpha_{i,h} + \delta_{t,h} + \beta_h^{ss}(z_i \cdot \varepsilon_t^s) + e_{i,t,h} \tag{26}$$

Figure 1 shows my results for different values of $\delta_{\mathcal{U}}$ and $\delta_{\mathcal{H}}$. Figure 1a plots the path for transfers following the shock. Figures 1b–1d then compare the shift-share coefficients $\boldsymbol{\beta}^{ss}$, shown in the orange line, to the true $\boldsymbol{\beta}_i$ of the treated group (i.e. MPCs for the treated group), shown in the black dots. I consider three cases for the targeting rule: (i) targeting all unconstrained households in Figure 1b ($\delta_{\mathcal{U}} = 1, \delta_{\mathcal{H}} = 0$); (ii) targeting all hand-to-mouth households in Figure 1c ($\delta_{\mathcal{U}} = 0, \delta_{\mathcal{H}} = 1$); (iii) targeting only a very small fraction of hand-to-mouth households in Figure 1d ($\delta_{\mathcal{U}} = 0, \delta_{\mathcal{H}} = 0.01$).

Cases (i) and (ii) serve to highlight the potential severity of the bias term from Proposition 2. For example, when transfers are targeted to the unconstrained households (Figure 1b), the true direct effect is close to zero (< 0.01) and flat, while the shift-share coefficient estimates an on-impact MPC of around 1 – mistakenly suggesting the policy intervention was highly effective at stimulating consumption at the household-level. I plot the paths for all GE variables following this shock in Appendix A.4 to provide intuition for the bias in this case. In Appendix A.4, I also consider applying the correction from Proposition 3, which adds control variables of the form $(z_i \cdot w_t)$. I show this approach suffers from similar bias.

In contrast, there is little bias in case (iii). This case mimics a "micro" experiment: since the treated group are sufficiently small ($< 1\%$ of the total population), the GE effects of the shock are minimal, and so any bias from heterogeneous responses to GE effects across treatment and control groups does not arise. In this case the regression compares hand-to-mouth agents who receive the transfer to some combination of hand-to-mouth and Ricardian agents – since the latter are not affected by the shock, they serve as valid controls, despite not being "balanced" on observable characteristics with the treatment group.

Figure 1: Identification Without Exogenous Shares

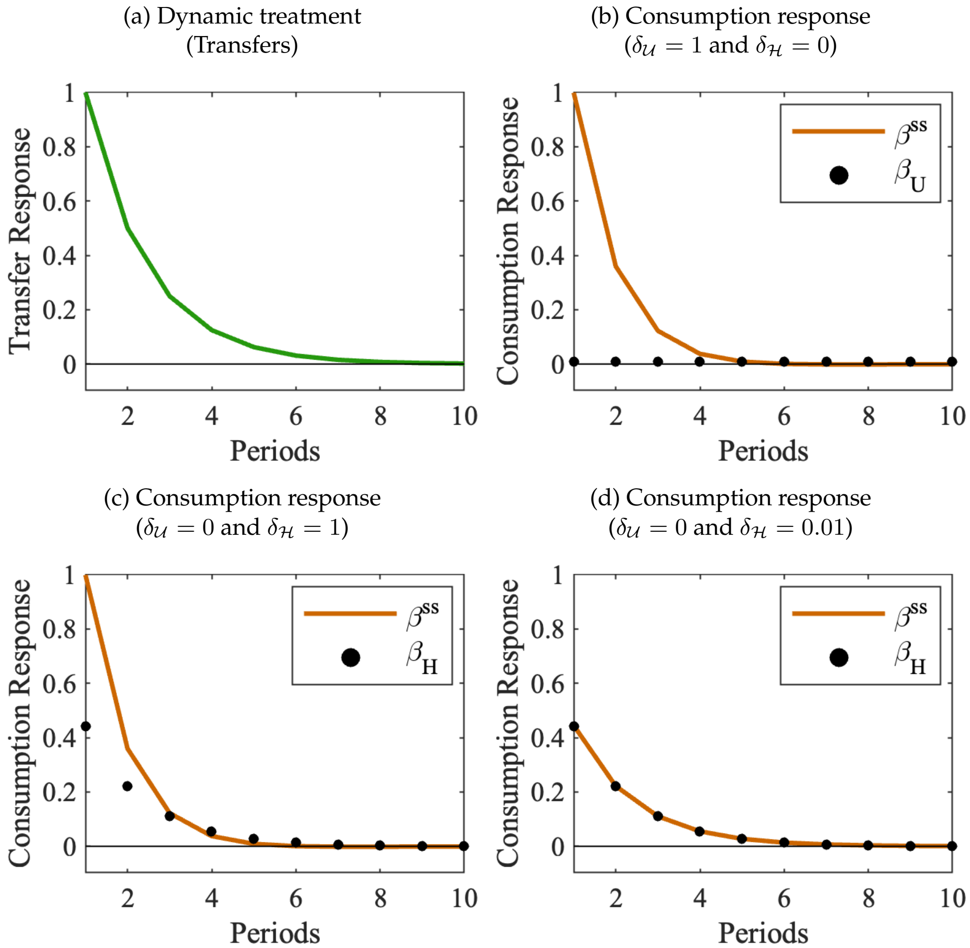

Notes: Green line in (a) plots the impulse response of transfers following the shock. Orange lines in (b)–(d) denote β^{ss} from (26) for the cases with: $\delta_{\mathcal{U}} = 1$, $\delta_{\mathcal{H}} = 0$; $\delta_{\mathcal{U}} = 0$, $\delta_{\mathcal{H}} = 1$; and $\delta_{\mathcal{U}} = 0$, $\delta_{\mathcal{H}} = 0.01$ respectively. Black dots denote the "direct" response of the treated group to the shock, defined by (25). All responses are scaled to the contemporaneous rise in transfers i.e. to $\Theta_0^{\mathcal{T}s}$.

I next demonstrate an application of Proposition 5. I consider regressions of the form:

$$c_{i,t+h} = \alpha_{i,h} + \delta_{t,h} + \beta_h^{ss}(z_i \cdot \varepsilon_t^s) + \lambda_h(\mathcal{I}(i \in \mathcal{U}) \cdot \varepsilon_t^s) + e_{i,t,h} \tag{27}$$

for the case where $\delta_{\mathcal{U}} = 0$ and $\delta_{\mathcal{H}} = 0.5$. In this case, β^{ss} captures the difference in impulse responses for hand-to-mouth agents in "treatment" and "control" groups. This recovers exactly the MPC of hand-to-mouth agents out of the transfer (see Figure 6).

Finally, I demonstrate an application of Proposition 6. I consider a slight tweak to the baseline model such that transfers additionally respond to other macroeconomic shocks. Specifically I assume that aggregate transfers also follow an endogenous policy rule and are subject to a series of "news shocks":

$$\mathcal{T}_t = \underbrace{\rho_{\mathcal{T}}\mathcal{T}_{t-1} + (1-\rho_{\mathcal{T}})\phi_{\mathcal{T}}(\log \Pi_t - \log \bar{\Pi})}_{\text{policy rule}} + \underbrace{\sum_{h=0}^{4} s_{\mathcal{T}}\varepsilon^s_{h,t-h}}_{\text{news shocks}} . \tag{28}$$

The policy rule is deliberately stylised, although it can be viewed as a simple way to capture a number of important features of e.g. social security transfers in the US (see Romer and Romer, 2016). First, large adjustments in transfers often occur as a result of periods of rapid inflation – mechanically so, since the introduction of automatic cost-of-living adjustments in 1975. Second, large adjustments in transfers also occur as a result of (plausibly exogenous) policy reforms, often announced prior to their implementation. I now assume the econometrician does not have access to the underlying transfer shocks ($\varepsilon^s_{h,t}$) but has access to some generic macroeconomic "news" defined as:

$$x^n_{t,h} = \mathbb{E}_t[x_{t+h}] - \mathbb{E}_{t-1}[x_{t+h}] \tag{29}$$

where x_t is some aggregate (observable) macroeconomic variable.[25] I then suppose they run the following pair of regressions

$$c_{i,t+h} = \alpha_{i,c,h} + \delta_{t,c,h} + \beta^{ss}_{c,h}(z_i \cdot x^n_{t,h'}) + \lambda_{c,h}(\mathcal{I}(i \in \mathcal{U}) \cdot x^n_{t,h'}) + e_{i,t,c,h} \tag{30}$$

$$\mathcal{T}_{i,t+h} = \alpha_{i,\mathcal{T},h} + \delta_{t,\mathcal{T},h} + \beta^{ss}_{\mathcal{T},h}(z_i \cdot x^n_{t,h'}) + \lambda_{\mathcal{T},h}(\mathcal{I}(i \in \mathcal{U}) \cdot x^n_{t,h'}) + e_{i,t,\mathcal{T},h} \tag{31}$$

Figure 2 below shows my results for the case where $\delta_{\mathcal{U}} = 0$ and $\delta_{\mathcal{H}} = 0.5$, for news at various horizons h', about either transfers ($x^n_{t,h'} = \mathcal{T}^n_{t,h'}$) in the top panel or inflation ($x^n_{t,h'} = \Pi^n_{t,h'}$) in the bottom panel. The shift-share coefficient $\beta^{ss}_{\mathcal{T}}$ from (31) traces out the green lines in the left-hand side figures – the "treatment path" – while β^{ss}_c from (30) traces out the corre-

[25] Note that news about macroeconomic variables at different horizons is plausibly available in practice either from, e.g., policymaker or professional forecasters, or from moves in asset prices. An alternative way to implement this same identification strategy is to exploit staggered calendar timing across agents in the response of some treatment to news occurring at different points in time – as in e.g. Siani and Zhang (2026).

sponding orange lines on the right-hand side – the corresponding consumption responses of hand-to-mouth agents. Even absent the underlying structural shocks, by projecting on various macroeconomic news, the econometrician generally recovers the consumption response of hand-to-mouth agents to a rich menu of different transfer paths – essentially tracing out additional entries of the intertemporal MPC matrix.

Figure 2: Projecting on Various Macroeconomic News

Panel A: Transfer News

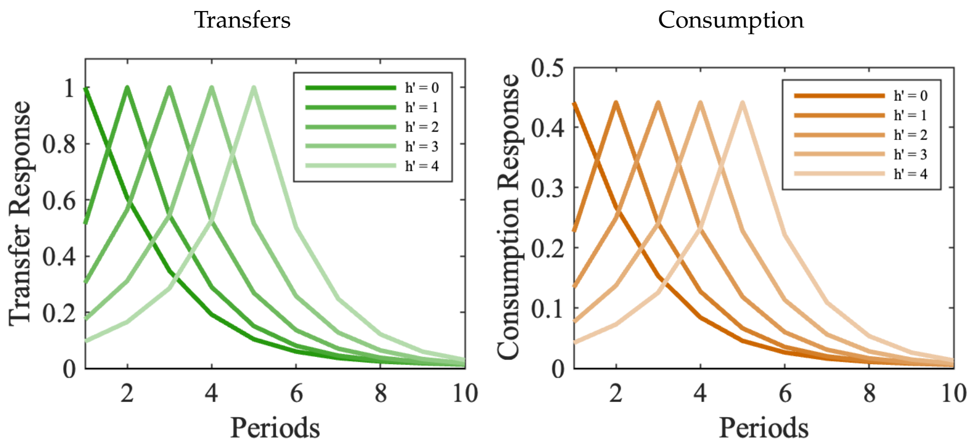

Panel B: Inflation News

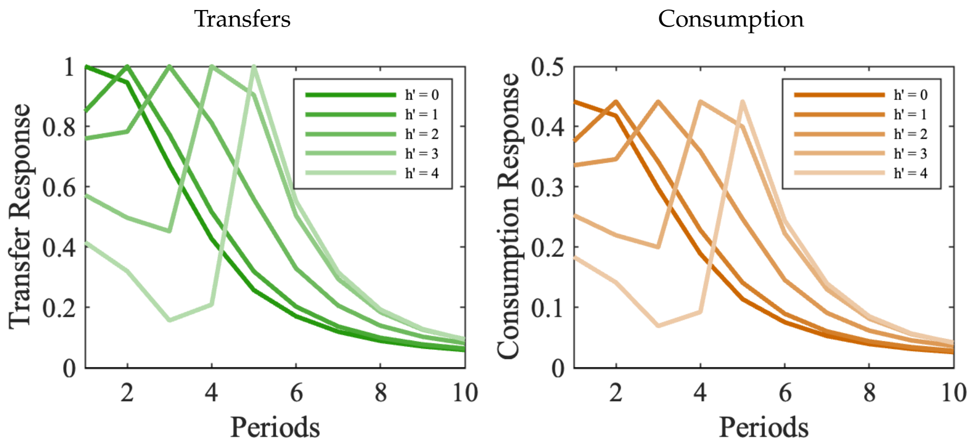

Notes: Green lines in (a) trace out the shift-share coefficients $\beta^{ss}_{\mathcal{T},h}$ from (31). Orange lines in (b) trace out $\beta^{ss}_{c,h}$ from (30). Top panel uses transfer news as the shifter ($x^n_{t,h'} = \mathcal{T}^n_{t,h'}$) while bottom panel uses inflation news ($x^n_{t,h'} = \Pi^n_{t,h'}$). All responses are scaled to imply a one-unit peak response of transfers.

5 Empirical Applications

I now illustrate the practical implications of my econometric framework through two empirical applications based on Nakamura and Steinsson (2014) and Guren et al. (2021b). The primary purpose of my first application is to demonstrate the quantitative importance of my econometric results – in particular, to highlight the need for identification strategies that are robust to heterogeneous responses to GE effects. The primary purpose of my second application is to instead demonstrate the more novel features of my framework.

5.1 Empirical Application 1: Nakamura and Steinsson (2014)

My first application is based on Nakamura and Steinsson (2014) who estimate local (state-level) fiscal multipliers using a shift-share research design with a national shifter. I use this to demonstrate the quantitative importance of my econometric results. I describe the empirical setup, then discuss issues related to identification.

Setting Nakamura and Steinsson (2014) seek to identify local fiscal multipliers through a shift-share research design. They do so using annual panel data on US states for 1966–2006. They regress per capita output and per capita military spending at the state level – their "outcome" and "treatment" variables respectively – on changes in national military spending interacted with a measure of a state's exposure to these aggregate shocks. The exposure shares are constructed as the average level of military spending in each state relative to state output in the early years of the sample.[26] The "local multiplier" is then estimated via IV – i.e. the ratio of the estimated coefficients across regressions. As the authors note, changes in national military spending are driven largely by geopolitical concerns, and thus plausibly exogenous to any aggregate and local economic shocks.[27] Seen through the lens of the results in the previous section, however, this is not sufficient for identification. The key requirement for identification is that the military spending shares are orthogonal to state characteristics that drive heterogeneity in states' responses to aggregate macroeconomic conditions that adjust in response to government spending shocks (e.g. changes in monetary policy, taxes, etc.). A priori, this requirement is unlikely to hold: states that receive high levels of national military spending tend to specialise in high-tech manufacturing sectors (e.g. California) – and as such are likely to differ from the average state across various dimensions (industry structure, average income and wealth etc.). Absent this requirement, my results highlight that an appropriate route to identification is to seek to control directly for the state characteristics that simultaneously determine a state's military spending share as well as their heterogeneous response to macroeconomic aggregates. I take up this challenge below.

[26] The authors also present results from an estimation strategy that constructs instruments (one per state) as the interaction of state fixed effects with the aggregate shifter. I focus on the shift-share specification since this maps directly to my theoretical results (and avoids any concerns around inference with many instruments).

[27] Specifically, they write: *"Our identifying assumption is that the United States does not embark on military buildups—such as those associated with the Vietnam War and the Soviet invasion of Afghanistan—because states that receive a disproportionate amount of military spending are doing poorly relative to other states. This assumption is similar...(to) the common identifying assumption in the empirical literature on the effects of national military spending, that variation in national military spending is exogenous to the United States business cycle."*

Properties of shocks and shares Before turning to my main estimates, I first more formally assess the properties of the aggregate shock and state-level shares from Nakamura and Steinsson (2014), comparing them to the requirements implied by my econometric results. I begin by testing directly whether the military spending shares correlate with other state characteristics that plausibly drive states' responses to macroeconomic conditions. I take as my starting point recent work on regional heterogeneity in responses to macroeconomic shocks. Herreño and Pedemonte (2022) find that regional responses to monetary policy shocks are highly correlated with income per capita, rationalising these findings with a New Keynesian model where regions differ in their fraction of hand-to-mouth households. Bellifemine et al. (2023) construct a multi-region Heterogeneous Agent New Keynesian model of US counties and show that the response of regions to aggregate shocks is shaped by two key dimensions of heterogeneity: MPCs and trade openness (i.e. the fraction of employment in the tradable sector). Motivated by this, I run regressions of the form:

$$s_i = \alpha + \lambda' a_i + u_i \tag{32}$$

where s_i denotes the military spending shares from Nakamura and Steinsson (2014), and where I include in a_i measures of states' MPCs, openness and income per capita.[28] I find that these three variables jointly explain over 45% of the variation in the shares s_i, with the MPC and income per capita variables significant at least at the 10% level (Table 5). The signs of the coefficients are in line with previous discussion: states with high shares of military spending tend on average to be more open, have higher incomes and lower MPCs. I consider various approaches to purge the exposure shares of this endogeneity in my main specification below.

The key requirement for the aggregate shock is simply that it satisfies lead/lag exogeneity – in particular, that it is uncorrelated with any past macroeconomic shocks. Ramey (2020) demonstrates that the shock in Nakamura and Steinsson (2014) – which captures 2-year changes in national military spending – is highly autocorrelated and suggests including lagged variables to account for this. In my main specification below, I use the 1-year change in national military spending, which I find is not significantly autocorrelated. I then further control for 1-year lags of all variables in the main specification to purge any remaining potential serial dependence. As robustness, I consider alternative shocks – changes in excess returns of US defense manufacturers from Fisher and Peters (2010) – which are more plausibly unanticipated.

[28] I construct average income per capita for each state for the years in which the military spending shares are constructed (1966–1970). The openness and MPC measures are taken from Bellifemine et al. (2023) and are only available for a later sample. I take averages for these variables for the years they are available as a proxy for their steady-state values – exactly the objects that drive heterogeneity in regional responses to aggregates in the model from Bellifemine et al. (2023).

Main estimates and sensitivities For my main specification, I follow Ramey and Zubairy (2018) and estimate cumulative multipliers via dynamic shift-share regressions of the form:

$$\sum_{j=0}^{h}\left(\frac{y_{i,t+j}-y_{i,t-1}}{y_{i,t-1}}\right)=\alpha_{y,i,h}+\delta_{y,t,h}+\beta^{ss}_{y,h}(s_i\cdot d_t)+\lambda'_{y,h}(a_i\cdot d_t)+\gamma'_{y,h}x_{i,t-1}+u_{y,i,t,h} \tag{33}$$

$$\sum_{j=0}^{h}\left(\frac{g_{i,t+j}-g_{i,t-1}}{y_{i,t-1}}\right)=\alpha_{g,i,h}+\delta_{g,t,h}+\beta^{ss}_{g,h}(s_i\cdot d_t)+\lambda'_{g,h}(a_i\cdot d_t)+\gamma'_{g,h}x_{i,t-1}+u_{g,i,t,h} \tag{34}$$

where $y_{i,t}$ and $g_{i,t}$ denote per capita output and per capita military procurement spending respectively in state-i in year-t, $d_t \equiv (g_t - g_{t-1})/y_{t-1}$ denotes 1-year changes in national military spending in year t as a fraction of national output, and s_i denotes the average level of military spending in state-i relative to state output in the first five years of the sample. As in Nakamura and Steinsson (2014), I include unit- and time-fixed effects as controls – I then additionally include the interaction of the shock with a vector of state characteristics a_i, and a vector of lagged controls $x_{i,t-1}$. In all specifications, I include one lag of growth in state output and state military spending as well as one lag each of all the right-hand side variables. The multiplier is then defined appropriately as $\beta^{ss}_{y,h}/\beta^{ss}_{g,h}$.[29] Following Almuzara and Sancibrián (2024), I compute standard errors by clustering at the year level. Table 1 displays my results.

The first column mimics the baseline specification from Nakamura and Steinsson (2014) without additional controls for state characteristics, yielding an estimated multiplier of 1.45 at the 2-year horizon, rising to 2.69 at the 4-year horizon (both significant at least at the 10% level). The second column adds a control for state income per capita – a simple proxy for the prevalence of hand-to-mouth agents across regions, as in Herreño and Pedemonte (2022) – in which case the multiplier is lower across all horizons, around 1.04 and 1.45 at the 2- and 4-year horizons respectively (both significant at least at the 10% level). The third column instead adds controls for state MPCs and openness – the two key measures of regional heterogeneity from Bellifemine et al. (2023) – in which case the multiplier drops to 0.5 and 0.95 at the 2- and 4-year horizon (no longer significant at conventional levels). These columns identify local fiscal multipliers if one is willing to believe that the control variables effectively capture the relevant dimensions of regional heterogeneity that drive state responses to macroeconomic aggregates alongside the size of the military sector in each region. This is more plausibly the case in column 3 – in the model from Bellifemine et al. (2023), MPCs and openness are "sufficient statistics" for regional heterogeneity, capturing all relevant heterogeneity in regional responses to aggregate shocks regardless of whether the underlying differences across states are driven by e.g. heterogeneity in discount factors, income risk or borrowing limits.[30]

[29] Since this specification deviates in various ways from the authors' original specification, for full comparability, I also replicate the authors' original shift-share specification in Appendix D.6.

[30] A key message of the model from Bellifemine et al. (2023) is that MPCs and openness shape states' responses to GE variables in highly non-linear ways. These nonlinearities are not directly relevant to the solution proposed in

Rather than relying on economic theory to act as a guide on the appropriate set of controls, I also consider a more data-driven approach to controlling for the targeting rule. I gather data on a range of variables from the Correlates of State Policy database that capture various dimensions of regional heterogeneity across economic (e.g. income and income inequality), fiscal (e.g. tax and welfare policies) and demographic characteristics (e.g. the fraction of Black or Foreign-born residents). Using all these variables (alongside the MPC and openness indicators) in the vector a_i, I estimate the "targeting rule" regression (32) directly. To avoid overfitting, I estimate the regression via LASSO, with the regularisation parameter that governs the degree of shrinkage selected by 10-fold cross-validation. I then include in the vector of controls a_i in regressions (33) and (34) the variables with non-zero coefficients from the LASSO estimation of the "targeting rule".[31] In this case, the multiplier estimate drops to 0.78 and 1.10 at the 2- and 4-year horizon respectively (both significant at least at the 10% level).

In all cases, the additional control variables reduce the estimated multiplier quite substantially. The standard errors around the multiplier are also notably lower – consistent with these additional controls successfully "soaking up" additional variation in states' responses to aggregate military spending shocks. Further, in each case, I find the difference in multiplier estimates at the 4-year horizon relative to the first column of the Table to be significant at least at the 5% level.[32]

How robust are these core findings to the possibility of other omitted interactions? I answer this question in various ways (see Appendix D.4 for details). First, I demonstrate that, conditional on the MPC and openness controls – the candidate "sufficient statistics" for regional heterogeneity – the estimated 4-year multiplier remains relatively stable after further including a variety of additional state characteristics as controls. This includes differences in income, income inequality, demographics and local fiscal policies. Across all nine different specifications, I find the 4-year multiplier estimate remains below 1.0 (Figure 9). Second, I consider including non-linear combinations of variables as controls – e.g. not just the openness and MPC measure, but also their interaction. Repeating each of the specifications in Table 1 with additional non-linear terms as controls, I find similar results (Table 6). Third, I repeat the LASSO analysis discussed above for various choices of the regularisation parameter, effectively altering the total number of control variables included in the final regression. Across various degrees of shrinkage, I find 4-year multiplier estimates consistently around, or even somewhat below, 1.0 (Figure 10).

Another potential threat to identification is that changes in national military spending may not capture genuine macroeconomic news, as required by my identification results. In Ap-

Proposition 5: so long as the *targeting rule* is linear then controlling linearly for appropriate variables is sufficient for identification. This is of course itself a strong assumption – I thus also consider the sensitivity of the main results to including non-linear combinations of controls below.

[31] I provide further details on the LASSO estimation and results in Appendix D.4.

[32] To implement this test, I employ the "stacking trick" from Hull (2022).

Table 1: Cumulative Multiplier Estimates

Horizon (years)	None (1)	Income (2)	MPC + Openness (3)	LASSO (4)
0	0.89	0.82	0.40	0.51
	(0.72)	(0.59)	(0.42)	(0.42)
	$[F = 16.84]$	$[F = 17.23]$	$[F = 17.32]$	$[F = 17.18]$
1	1.15	0.97	0.39	0.69
	(0.80)	(0.69)	(0.43)	(0.47)
	$[F = 10.90]$	$[F = 10.38]$	$[F = 11.08]$	$[F = 10.24]$
2	1.45*	1.04*	0.50	0.78*
	(0.76)	(0.62)	(0.45)	(0.42)
	$[F = 14.89]$	$[F = 16.00]$	$[F = 13.22]$	$[F = 13.92]$
3	1.96**	1.13*	0.70	0.85**
	(0.78)	(0.60)	(0.54)	(0.43)
	$[F = 14.58]$	$[F = 15.82]$	$[F = 12.76]$	$[F = 13.68]$
4	2.69***	1.47**	0.95	1.10**
	(0.84)	(0.68)	(0.62)	(0.46)
	$[F = 14.01]$	$[F = 13.81]$	$[F = 12.35]$	$[F = 12.66]$

Notes: Table displays coefficient estimates alongside (time-clustered) standard errors and F-stats from the IV regressions described in the main text. Column headers refer to additional interaction controls included in the regression. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

pendix D.5, I repeat Table 1 using annual changes in excess returns of US defense manufacturers from Fisher and Peters (2010) as the aggregate shock. Since these capture changes in financial asset prices, they are plausibly unanticipated by economic agents. I find similar, if slightly smaller, multipliers in this case – around 0.6-1.0 at the 4-year horizon in the specifications with additional interaction controls – although standard errors are notably wider.

Why do the estimates change as more control variables are added? My results from Section 3 highlight two reasons why this might be the case. First, the estimates in Table 1 typically capture *different* weighted averages of states' local multipliers. As discussed in Appendix D.7, I find little evidence that this is a significant driver of results. Second, as stressed throughout, the additional interaction terms are necessary to ensure the GE effects appropriately wash out across treatment and control groups. The fall in estimates is therefore consistent with a setting where states that have high military spending shares disproportionately benefit from the GE effects of shocks to national military spending. I explore this further in Appendix D.9. I run time-series local projections of various aggregate variables on the shock d_t. I find that, alongside a protracted rise in national military spending, the shock triggers a tightening in real interest rates, a protracted rise in aggregate consumption of tradables, and a fall in government transfers (with taxes and oil prices broadly flat). One candidate explanation for the change in estimates is that states with high military spending shares – that tend to be both richer and more open – disproportionately benefit from the combination of higher tradable goods demand, higher interest rates and lower transfers.[33]

5.2 Empirical Application 2: Guren, McKay, Nakamura, and Steinsson (2021b)

My second application is based on Guren et al. (2021b) who estimate local housing wealth effects – i.e. the response of local retail employment (a proxy for retail spending) to an increase in local house prices. As above, I describe the empirical setup, then discuss issues related to identification.

Setting Guren et al. (2021b) use quarterly data from US core-based statistical areas (CBSAs) to identify housing wealth effects – regressing CBSA house price growth and retail employment growth on aggregate house price shocks interacted with a measure of each CBSA's exposure to these shocks. The exposure shares are constructed using city-level estimates of housing supply elasticities from Saiz (2010).[34] The measure is determined by a city's topographical features, such as elevation and the presence of bodies of water, that are plausibly exogenous to economic fundamentals. Consistent with this interpretation, the Saiz measure has been employed directly as a cross-sectional instrument in various studies (Mian et al., 2013; Mian and Sufi, 2014).[35] A priori, this makes these exposure shares potentially well-suited to a dynamic shift-share research design with aggregate shocks.

Guren et al. (2021b) interact these shares with changes in aggregate national house prices. Such changes will typically reflect a combination of demand and supply forces stemming from various underlying aggregate shocks. As discussed, this poses no issue for identification: the key requirement is just that these changes are essentially unpredictable. Indeed, my framework clarifies that it would be valid to interact *any* combination of time-t aggregate shocks with an appropriate set of exposure shares. I show how to use this constructively to estimate the response of employment to different dynamic shocks to local house prices.

Properties of shocks and shares Table 10 reports results from a series of balance tests of the exposure shares. Conditional on Census-region fixed effects, I find the Saiz instrument is broadly balanced across a range of characteristics – including the openness and MPC measures from Bellifemine et al. (2023) as well as a broader range of economic and demographic characteristics. A notable exception is income per capita: CBSAs with lower housing elasticity

[33] Bellifemine et al. (2023) and Herreño and Pedemonte (2022) find that low-MPC and high-income states experience smaller declines in local output following a tightening in interest rates. These results are also consistent with findings in Donaldson (2026), who finds that states with higher military spending shares are less responsive to monetary policy shocks.

[34] The authors also report results based on their own measure of each CBSA's responsiveness to regional house price changes – constructed based on a series of regressions of local house prices on regional house prices and various controls. As in the previous application, I focus discussion here instead on the simpler specification with the Saiz (2010) exposure measure.

[35] More recent work has shown the instrument to be correlated with some city characteristics – see Davidoff (2016). I repeat this finding below, but show that it has negligible impact on the headline results.

tend to be significantly richer. They also tend to have a slightly higher fraction of foreign born residents. I control for these variables directly in the main specification below.

For the aggregate shock, I depart slightly from Guren et al. (2021b), who use annual changes in a leave-one-out average of national house prices, and instead use quarterly changes in the Case-Shiller U.S. National Home Price Index. I additionally use the quarterly credit spread measure from Gilchrist and Zakrajšek (2012). In both cases, I control for four lags in the main specification to account for serial dependence.[36]

Main estimates and sensitivities I estimate dynamic shift-share regressions of the form:

$$y_{i,r,t+h} - y_{i,r,t-1} = \alpha_{y,i,h} + \delta_{y,r,t,h} + \beta^{ss}_{y,h}(s_{i,r} \cdot d_t) + \lambda'_{y,h}(a_{i,r} \cdot d_t) + \gamma'_{y,h} x_{i,r,t-1} + u_{y,i,r,t,h} \quad (35)$$

$$p_{i,r,t+h} - p_{i,r,t-1} = \alpha_{p,i,h} + \delta_{p,r,t,h} + \beta^{ss}_{p,h}(s_{i,r} \cdot d_t) + \lambda'_{p,h}(a_{i,r} \cdot d_t) + \gamma'_{p,h} x_{i,r,t-1} + u_{p,i,r,t,h} \quad (36)$$

where $y_{i,r,t}$ and $p_{i,r,t}$ denote retail employment per capita and house prices respectively in CBSA i in region r at time t, $s_{i,r}$ denotes the Saiz instrument and d_t denotes the aggregate shock. Following Guren et al. (2021b), I include controls for CBSA fixed effects, as well as Census region-time fixed effects.[37] I also control for a vector of CBSA characteristics $a_{i,r}$ interacted with d_t and a vector of lagged controls $x_{i,r,t-1}$. In all specifications, I control for four lags of retail employment and house price growth, alongside four lags of all right-hand side variables. In my baseline specification, I include income per capita and the fraction of foreign born residents at the start of the sample in the vector $a_{i,r}$.

Panel A of Figure 3 displays my first set of results. The blue lines denote estimates of $\beta^{ss}_{y,h}$ and $\beta^{ss}_{p,h}$ from (35) and (36) respectively, using the Case-Shiller U.S. National Home Price Index as the aggregate shifter d_t. I find significant evidence of housing wealth effects: a shock that triggers a gradual and persistent rise in local house prices generates a corresponding large and persistent rise in local retail employment. The results are robust to a variety of control choices. First, coefficient estimates are essentially unchanged when removing all CBSA characteristics $a_{i,r}$ as controls (orange dashed line). Second, estimates are similar when controlling for a wider set of economic and demographic characteristics – not just income per capita and the fraction of foreign born residents, but also MPCs, openness, the fraction of Black residents, the home-ownership rate, the size of the construction sector and the average unemployment rate (green dashed line). Results are also largely unchanged when maintaining this larger set of controls and additionally including all (time-varying) CBSA two-digit industry shares interacted with the aggregate shock (yellow dashed line). This is consistent with the Saiz instrument being exogenous with respect to other determinants of regional heterogeneity in response to aggregate

[36] Note that by Frisch–Waugh–Lovell, this is equivalent to first constructing "news" shocks as residuals from an AR(4) regression and then using these shocks instead in the main specification – I provide diagnostics on these news shocks in Appendix E.3.

[37] For full comparability with the original study, I also replicate exactly the original specification from Guren et al. (2021b) in Appendix E.5.

house price shocks.

Panel B of Figure 3 shows results instead using credit spreads as the aggregate shock. In this case, I identify the effects of a more hump-shaped shock to local house prices: prices rise, reaching a peak after around one year, then gradually decline. I find local employment then follows a similar path: rising for about a year then declining. Repeating the same sensitivity analysis, I find that these core findings are again robust to alternative control choices (orange, green and yellow dashed lines).[38]

This application illustrates two points. First, in this setting, the identifying assumptions required by my framework plausibly hold (both a priori and ex post). Second, my framework highlights how to leverage these identifying assumptions to learn new information about the effects of different treatment paths – in this case, the effect of different house price movements on employment.

6 Conclusion

The use of cross-sectional data to answer key questions of interest in macroeconomics has risen sharply in recent years. It is well-understood that identification in these settings generally faces the so-called "missing intercept" problem: there is typically insufficient information in cross-sectional data to recover the "total" (or "general equilibrium") effects of policy interventions. In this paper, I consider a less well-studied problem: if policy interventions indeed propagate via GE forces and agents are heterogeneous, then otherwise credible research designs generally fail to recover *either* the "total" *or* the "direct" effects of these interventions.

The main contribution of this paper is to demonstrate a robust and simple approach to achieve identification of the direct effects of policy interventions using cross-sectional variation. The key informational requirement for this approach is knowledge of the underlying targeting of the exposure to the intervention across units. Crucially, this approach does not require the econometrician to have direct knowledge of the underlying GE transmission mechanisms or to have access to any well-identified time-series shocks – thus respecting the typical motivation for employing cross-sectional variation to achieve identification. Further, armed with appropriately exogenous exposure shares, this approach can be leveraged to learn about the dynamic response of particular outcome variables to a variety of dynamic treatments – a key set of objects for HA-DSGE models. Taken together, my results highlight both the promise and the pitfalls of employing cross-sectional data in macroeconomics.

[38] Note that without any characteristic controls, the local house price shock becomes somewhat more front-loaded, with a correspondingly more front-loaded response of local employment. This is not necessarily a sign of bias in the estimates – the change in control set simply places different weight on regions with more front-loaded house price responses (see the discussion on dynamic treatment heterogeneity in Appendix C.4).

Figure 3: Housing Wealth Effect Estimates And Sensitivity Analysis

Panel A: Aggregate House Price Shock

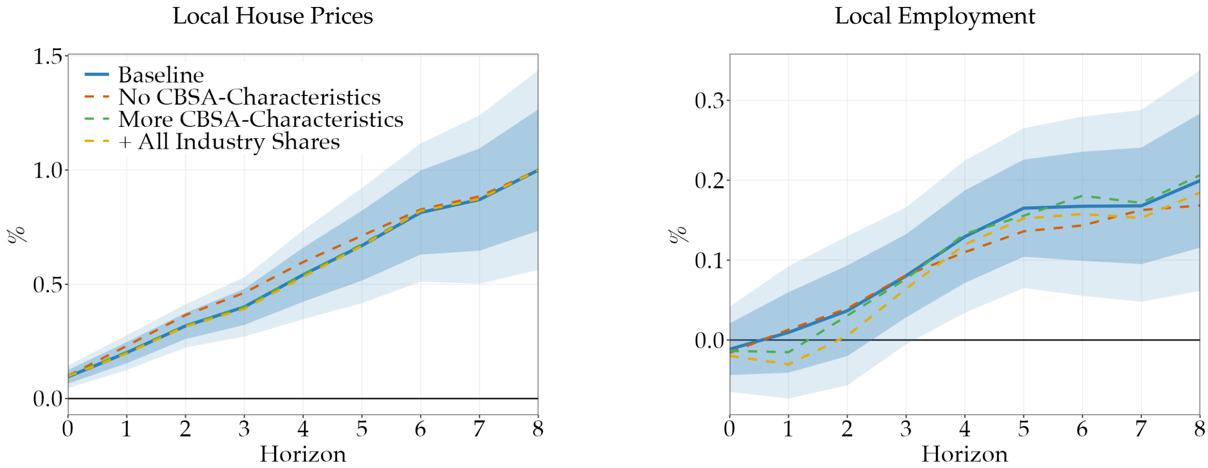

Panel B: Credit Spread Shock

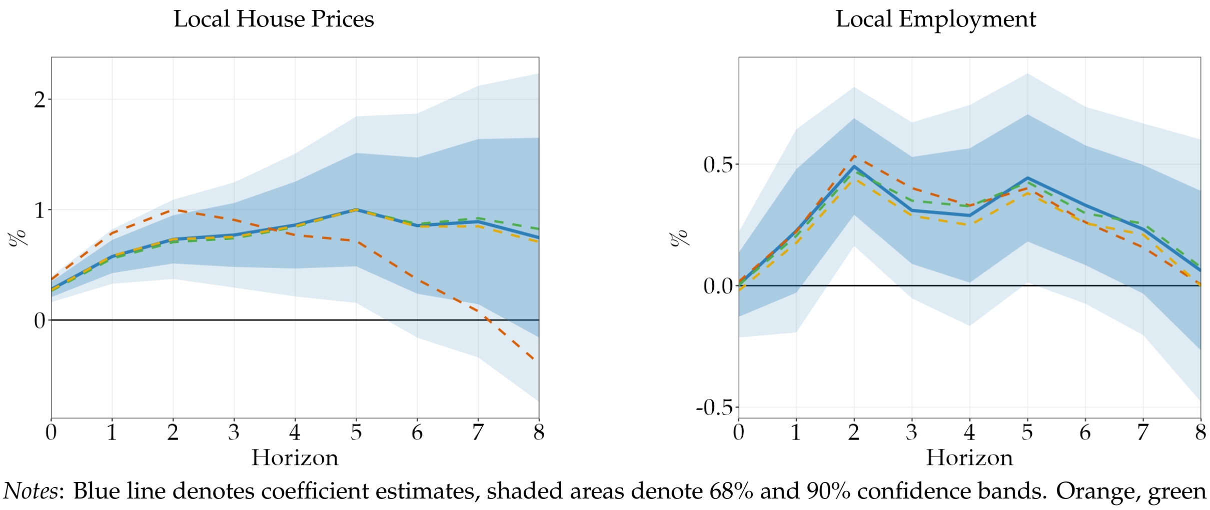

Notes: Blue line denotes coefficient estimates, shaded areas denote 68% and 90% confidence bands. Orange, green and yellow dashed lines report coefficient estimates from sensitivity analyses described in the main text and Appendix E.4. All coefficients are scaled to increase local house prices by 1% at peak.

References

ACEMOGLU, D. AND P. RESTREPO (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244.

ADÃO, R., C. ARKOLAKIS, AND F. ESPOSITO (2025): "General Equilibrium Effects in Space: Theory and Measurement," Working paper.

ADÃO, R., M. KOLESÁR, AND E. MORALES (2019): "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

ALMUZARA, M. AND V. SANCIBRIÁN (2024): "Micro Responses to Macro Shocks," Staff Reports 1090, Federal Reserve Bank of New York.

ANGRIST, J. D. (1998): "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants," *Econometrica*, 66, 249–288.

ARKHANGELSKY, D. AND V. KOROVKIN (2024): "On Policy Evaluation with Aggregate Time-Series Shocks," Working paper.

AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375–2408.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2024): "The Intertemporal Keynesian Cross," *Journal of Political Economy*, 132, 4068–4121.

AUERBACH, A., Y. GORODNICHENKO, AND D. MURPHY (2020): "Local Fiscal Multipliers and Fiscal Spillovers in the USA," *IMF Economic Review*, 68, 195–229.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.

BARNICHON, R. AND G. MESTERS (2023): "A Sufficient Statistics Approach for Macro Policy," *American Economic Review*, 113, 2809–2845.

BARRO, R. AND C. J. REDLICK (2011): "Macroeconomic Effects From Government Purchases and Taxes," *The Quarterly Journal of Economics*, 126, 51–102.

BELLIFEMINE, M., A. COUTURIER, AND R. JAMILOV (2023): "The Regional Keynesian Cross," Discussion Papers 2311, Centre for Macroeconomics (CFM).

BOEHM, J., E. FIZE, AND X. JARAVEL (2025): "Five Facts about MPCs: Evidence from a Randomized Experiment," *American Economic Review*, 115, 1–42.

BORUSYAK, K., R. DIX-CARNEIRO, AND B. KOVAK (2022a): "Understanding migration responses to local shocks," SSRN Working Paper. Available at `https://ssrn.com/abstract=4086847`.

BORUSYAK, K. AND P. HULL (2024): "Negative Weights Are No Concern in Design-Based Specifications," *AEA Papers and Proceedings*, 114, 597–600.

BORUSYAK, K., P. HULL, AND X. JARAVEL (2022b): "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 89, 181–213.

——— (2025): "A Practical Guide to Shift-Share Instruments," *Journal of Economic Perspectives*, 39, 181–204.

CANOVA, F. (2024): "Should we trust cross-sectional multiplier estimates?" *Journal of Applied Econometrics*, 39, 589–606.

CHODOROW-REICH, G. (2014): "The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis," *The Quarterly Journal of Economics*, 129, 1–59.

——— (2019): "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?" *American Economic Journal: Economic Policy*, 11, 1–34.

——— (2020): "Regional data in macroeconomics: Some advice for practitioners," *Journal of Economic Dynamics and Control*, 115.

CHODOROW-REICH, G., L. FEIVESON, Z. LISCOW, AND W. G. WOOLSTON (2012): "Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act," *American Economic Journal: Economic Policy*, 4, 118–45.

CHODOROW-REICH, G., P. T. NENOV, AND A. SIMSEK (2021): "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, 111, 1613–1657.

DAVID, J. M., R. GIACOMINI, X. JIAO, AND W. WANG (2026): "Causal State-Dependent Local Projections," Bristol Economics Discussion Papers 26/829, School of Economics, University of Bristol, UK.

DAVIDOFF, T. (2016): "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors," *Critical Finance Review*, 5, 177–206.

DE CHAISEMARTIN, C. AND X. D'HAULTFŒUILLE (2023): "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey," *The Econometrics Journal*, 26, 1–30.

DONALDSON, P. (2026): "Cross-sectional Identification with Heterogeneous Exposure to General Equilibrium Effects," Working Paper.

DUBE, O. AND J. F. VARGAS (2013): "Commodity Price Shocks and Civil Conflict: Evidence from Colombia," *The Review of Economic Studies*, 80, 1384–1421.

DUPOR, B. AND P. B. MCCRORY (2018): "A Cup Runneth Over: Fiscal Policy Spillovers from the 2009 Recovery Act," *Economic Journal*, 128, 1476–1508.

DUPOR, B. AND M. S. MEHKARI (2016): "The 2009 Recovery Act: Stimulus at the extensive and intensive labor margins," *European Economic Review*, 85, 208–228.

EGGER, D., J. HAUSHOFER, E. MIGUEL, P. NIEHAUS, AND M. WALKER (2022): "General Equilibrium Effects of Cash Transfers: Experimental Evidence From Kenya," *Econometrica*, 90, 2603–2643.

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2021): "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 13, 1–54.

FISHER, J. AND R. PETERS (2010): "Using Stock Returns to Identify Government Spending Shocks," *Economic Journal*, 120, 414–436.

GALEGO MENDES, A., W. MIYAMOTO, T. L. NGUYEN, S. M. PENNINGS, AND L. FELER (2024): "The Macroeconomic Effects of Cash Transfers : Evidence from Brazil," Policy Research Working Paper Series 10652, The World Bank.

GERARD, F., J. NARITOMI, AND J. C. G. SILVA (2025): "Cash Transfers and the Local Economy: Evidence from Brazil," Working paper.

GILCHRIST, S. AND E. ZAKRAJŠEK (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

GROSSMANN, M., M. P. JORDAN, AND J. MCCRAIN (2021): "The Correlates of State Policy and the Structure of State Panel Data," *State Politics Policy Quarterly*, 21, pp. 430–450.

GUREN, A., A. MCKAY, E. NAKAMURA, AND J. STEINSSON (2021a): "What Do We Learn from Cross-Regional Empirical Estimates in Macroeconomics?" *NBER Macroeconomics Annual*, 35, 175–223.

GUREN, A. M., A. MCKAY, E. NAKAMURA, AND J. STEINSSON (2021b): "Housing Wealth Effects: The Long View," *The Review of Economic Studies*, 88, 669–707.

HERREÑO, J. AND M. PEDEMONTE (2022): "The Geographic Effects of Monetary Policy," Working Papers 22-15, Federal Reserve Bank of Cleveland.

HUBER, K. (2022): "Estimating General Equilibrium Spillovers of Large-Scale Shocks," NBER Working Papers 29908, National Bureau of Economic Research, Inc.

HULL, P. (2022): Tweet, `https://twitter.com/instrumenthull/status/1492915860763250691`.

IMBENS, G. W. AND J. D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–475.

KOBY, Y. AND C. K. WOLF (2020): "Aggregation in Heterogeneous-Firm Models: Theory and Measurement," Tech. rep.

MAJEROVITZ, J. AND K. SASTRY (2023): "How Much Should We Trust Regional-Exposure Designs?" Working Papers 2023-018, Federal Reserve Bank of St. Louis.

MATTHES, C., N. NAGASAKA, AND F. SCHWARTZMAN (2025): "Estimating the Missing Intercept," Working Paper 25-12, Federal Reserve Bank of Richmond.

MCKAY, A. AND C. K. WOLF (2023): "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91, 1695–1725.

MERTENS, K. AND J. L. M. OLEA (2018): "Marginal Tax Rates and Income: New Time Series Evidence," *The Quarterly Journal of Economics*, 133, 1803–1884.

MIAN, A., K. RAO, AND A. SUFI (2013): "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*, 128, 1687–1726.

MIAN, A. AND A. SUFI (2012): "The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program," *The Quarterly Journal of Economics*, 127, 1107–1142.

——— (2014): "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, 82, 2197–2223.

NAKAMURA, E. AND J. STEINSSON (2014): "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 104, 753–792.

——— (2018): "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32, 59–86.

NUNN, N. AND N. QIAN (2014): "US Food Aid and Civil Conflict," *American Economic Review*, 104, 1630–1666.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103, 2530–2553.

PENNINGS, S. (2021): "Cross-Region Transfer Multipliers in a Monetary Union: Evidence from Social Security and Stimulus Payments," *American Economic Review*, 111, 1689–1719.

RAMEY, V. (2020): "Comment on "What Do We Learn From Cross-Regional Empirical Estimates in Macroeconomics?"," in *NBER Macroeconomics Annual 2020, volume 35*, National Bureau of Economic Research, Inc, 232–241.

RAMEY, V. A. (2011): "Identifying Government Spending Shocks: It's all in the Timing," *The Quarterly Journal of Economics*, 126, 1–50.

——— (2016): "Macroeconomic Shocks and Their Propagation," Elsevier, vol. 2 of *Handbook of Macroeconomics*, 71–162.

RAMEY, V. A. AND S. ZUBAIRY (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126, 850–901.

ROMER, C. D. AND D. H. ROMER (2016): “Transfer Payments and the Macroeconomy: The Effects of Social Security Benefit Increases, 1952-1991,” *American Economic Journal: Macroeconomics*, 8, 1–42.

SAIZ, A. (2010): “The Geographic Determinants of Housing Supply,” *The Quarterly Journal of Economics*, 125, 1253–1296.

SARTO, A. (2025): “Recovering Macro Elasticities from Regional Data”,” Unpublished manuscript.

SIANI, K. AND J. ZHANG (2026): “The Decision Window: How Budget Timing Shapes Fiscal Policy,” Working paper.

STOCK, J. H. AND M. W. WATSON (2018): “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *Economic Journal*, 128, 917–948.

WILSON, D. J. (2012): “Fiscal Spending Jobs Multipliers: Evidence from the 2009 American Recovery and Reinvestment Act,” *American Economic Journal: Economic Policy*, 4, 251–282.

WOLF, C. K. (2023): “The Missing Intercept: A Demand Equivalence Approach,” *American Economic Review*, 113, 2232–2269.

ZWICK, E. AND J. MAHON (2017): “Tax Policy and Heterogeneous Investment Behavior,” *American Economic Review*, 107, 217–48.

Appendix

A Two-Agent New Keynesian (TANK) Model

This section provides further information on the TANK model.

A.1 Characterising the $\mathcal{M}$-mappings

I first derive closed-form expressions for the $\mathcal{M}$-mappings introduced in equation (7) in Section 2. I use bars to denote steady-state values throughout.

Hand-to-mouth First-order conditions for the hand-to-mouth consumption problem are given by:

$$\theta n_{i,t}^{\chi} = w_t\, c_{i,t}^{-\sigma}, \qquad c_{i,t} = w_t n_{i,t} - \tau_t + \mathcal{T}_{i,t}. \tag{37}$$

Totally differentiating around steady state, we have:

$$\chi \frac{dn_{i,t}}{\bar{n}_i} = \frac{dw_t}{\bar{w}} - \sigma \frac{dc_{i,t}}{\bar{c}_i}, \tag{38}$$

$$dc_{i,t} = \bar{w}\, dn_{i,t} + \bar{n}_i\, dw_t - d\tau_t + d\mathcal{T}_{i,t}. \tag{39}$$

Substituting (38) into (39) and collecting terms in $dc_{i,t}$ gives:

$$dc_{i,t} = \frac{\chi}{\chi + \sigma\zeta_i}\,(d\mathcal{T}_{i,t} - d\tau_t) + \frac{(1+\chi)\,\bar{n}_i}{\chi + \sigma\zeta_i}\, dw_t, \tag{40}$$

where ζ_i is defined as:

$$\zeta_i \equiv \frac{\bar{w}\bar{n}_i}{\bar{c}_i}. \tag{41}$$

Because the problem is static, consumption at $t+h$ depends only on prices and quantities at $t+h$: the $\mathcal{M}$-mappings in this case are diagonal, with common diagonal entries. Further, since the model is linearised around steady state with zero aggregate transfers, steady-state values are identical across all hand-to-mouth agents (regardless of their treatment status). And so we can express each of the $\mathcal{M}$-mappings for all hand-to-mouth agents as:

$$\mathcal{M}_{\mathcal{H}}^{\mathcal{T}} = \frac{\chi}{\chi + \sigma\zeta_{\mathcal{H}}}\, I, \qquad \mathcal{M}_{\mathcal{H}}^{\tau} = -\mathcal{M}_{\mathcal{H}}^{\mathcal{T}}, \qquad \mathcal{M}_{\mathcal{H}}^{w} = \frac{(1+\chi)\,\bar{n}_{\mathcal{H}}}{\chi + \sigma\zeta_{\mathcal{H}}}\, I, \tag{42}$$

where I denotes the identity matrix. The interest-rate and dividend mappings are trivially zero since neither variable enters the hand-to-mouth problem. Note, in this case, the MPC out of a transfer for hand-to-mouth agents is strictly below one, despite the household consuming all of its income each period: a transfer raises consumption, the income effect on labour supply (equation (38), holding w_t fixed) reduces hours, and the lost labour income partially offsets the transfer.

Unconstrained First-order conditions for the unconstrained consumption problem are given by:

$$\theta n_{i,t}^{\chi} = w_t\, c_{i,t}^{-\sigma}, \quad c_{i,t}^{-\sigma} = \beta\, \mathbb{E}_t\left[r_{t+1}\, c_{i,t+1}^{-\sigma}\right], \quad c_{i,t} + b_{i,t} = w_t n_{i,t} + d_t - \tau_t + r_t b_{i,t-1} + \mathcal{T}_{i,t}. \tag{43}$$

Note that the Euler equation implies that in steady state: $\beta\bar{r} = 1$. Linearising the Euler equation around the steady state and iterating forward:

$$dc_{i,t+h} = dc_{i,t} + \frac{\beta\bar{c}_i}{\sigma}\sum_{k=1}^{h} dr_{t+k}, \qquad h \geq 1. \tag{44}$$

And linearising the flow budget constraint:

$$\begin{aligned} dc_{i,t+h} + db_{i,t+h} = \bar{w}\, dn_{i,t+h} + \bar{n}_i\, dw_{t+h} + dd_{t+h} - d\tau_{t+h} \\ + \bar{r}\, db_{i,t+h-1} + \bar{b}_i\, dr_{t+h} + d\mathcal{T}_{i,t+h}. \end{aligned} \tag{45}$$

Linearising the intratemporal FOC yields the same equation as for the hand-to-mouth agent, (38). Combining the linearised FOCs we get:

$$dc_{i,t} = \frac{1-\beta}{\kappa_i}\sum_{h\geq 0}\beta^h\left[\frac{(1+\chi)\,\bar{n}_i}{\chi}\, dw_{t+h} + dd_{t+h} - d\tau_{t+h} + \bar{b}_i\, dr_{t+h} + d\mathcal{T}_{i,t+h}\right] - \frac{\beta\bar{c}_i}{\sigma}\sum_{k\geq 1}\beta^k\, dr_{t+k}, \tag{46}$$

with κ_i defined as:

$$\kappa_i \equiv 1 + \frac{\sigma}{\chi}\frac{\bar{w}\bar{n}_i}{\bar{c}_i}. \tag{47}$$

Again, steady-state values are identical across all unconstrained agents (regardless of their treatment status). And so for all unconstrained agents we have:

$$[\mathcal{M}_{\mathcal{U}}^{\mathcal{T}}]_{h,k} = \frac{(1-\beta)\,\beta^k}{\kappa_{\mathcal{U}}}, \qquad \mathcal{M}_{\mathcal{U}}^{d} = \mathcal{M}_{\mathcal{U}}^{\mathcal{T}}, \qquad \mathcal{M}_{\mathcal{U}}^{\tau} = -\mathcal{M}_{\mathcal{U}}^{\mathcal{T}}, \tag{48}$$

$$[\mathcal{M}_{\mathcal{U}}^{w}]_{h,k} = \frac{(1+\chi)\,\bar{n}_{\mathcal{U}}}{\chi\kappa_{\mathcal{U}}}\,(1-\beta)\,\beta^k, \tag{49}$$

$$[\mathcal{M}_{\mathcal{U}}^{r}]_{h,k} = \frac{(1-\beta)\,\beta^k}{\kappa_{\mathcal{U}}}\,\bar{b}_{\mathcal{U}} + \frac{\beta\bar{c}_{\mathcal{U}}}{\sigma}\left(\mathcal{I}\{1\leq k\leq h\} - \beta^k\,\mathcal{I}\{k\geq 1\}\right), \tag{50}$$

where $\mathcal{I}\{\cdot\}$ denotes the indicator function.

A.2 Characterising the regression weights

I now provide closed-form expressions for the weights in the regression estimand from equation (10) in Section 2. The claim is nested by the discussion of cross-sectional identification in Appendix C.5, which implies we have:

$$\boldsymbol{\beta}^{DiD} = \frac{\mathbb{E}[\mathcal{M}_i^{\mathcal{T}}\boldsymbol{\Theta}^{\mathcal{T}s}\omega_i]}{\mathbb{E}[\omega_i]}, \qquad \omega_i = \tilde{z}_i^2. \tag{51}$$

In this two-agent case:

$$\mathcal{M}_i^{\mathcal{T}} = \mathcal{M}_{g(i)}^{\mathcal{T}}, \qquad \tilde{z}_i = z_i - \delta_{g(i)}, \qquad g(i)\in\{\mathcal{U},\mathcal{H}\} \tag{52}$$

For binary z_i, $\mathbb{E}[\tilde{z}_i^2 \mid g] = \delta_g(1-\delta_g)$. Iterating expectations over $g\in\{\mathcal{U},\mathcal{H}\}$ then yields

$\omega_g \propto \Pr(g)\delta_g(1-\delta_g)$. Thus:

$$\boldsymbol{\beta}^{DiD} = \omega_{\mathcal{U}}\mathcal{M}_{\mathcal{U}}^{\mathcal{T}}\boldsymbol{\Theta}^{\mathcal{T}s} + \omega_{\mathcal{H}}\mathcal{M}_{\mathcal{H}}^{\mathcal{T}}\boldsymbol{\Theta}^{\mathcal{T}s}, \qquad \omega_{\mathcal{U}} = \frac{\alpha\,\delta_{\mathcal{U}}(1-\delta_{\mathcal{U}})}{\alpha\,\delta_{\mathcal{U}}(1-\delta_{\mathcal{U}}) + (1-\alpha)\,\delta_{\mathcal{H}}(1-\delta_{\mathcal{H}})},$$

and $\omega_{\mathcal{H}} = 1 - \omega_{\mathcal{U}}$.

A.3 Rest of the Economy

Firms and Price Setting The model features a representative competitive final goods firm alongside a continuum, $j \in [0,1]$, of monopolistically competitive intermediate goods producers.

The final output is a CES aggregate of intermediates:

$$Y_t = \left(\int_0^1 y_{j,t}^{\epsilon/(\epsilon-1)} dj\right)^{\epsilon/(\epsilon-1)}, \tag{53}$$

and profit maximisation of the final good firm is:

$$\max_{y_{j,t}} P_t Y_t - \int_0^1 p_{j,t} y_{j,t} dj. \tag{54}$$

Given CES production, the aggregate price level can be expressed as:

$$P_t = \left(\int_0^1 p_{j,t}^{(1-\epsilon)} dj\right)^{1/(1-\epsilon)}. \tag{55}$$

Intermediate producers produce output according to:

$$y_{j,t} = A_t n_{j,t} \tag{56}$$

where $y_{j,t}$ and $n_{j,t}$ denote output and labour respectively for firm j, and A_t denotes aggregate productivity. Nominal profit for producer j is given by:

$$P_t d_{j,t} = p_{j,t} y_{j,t} - \frac{P_t w_t}{A_t} y_{j,t} \tag{57}$$

where $d_{j,t}$ denotes real dividends. A fraction $1-\phi$ of intermediate goods producers can reset prices each period. The problem for an updating firm is given by:

$$\max_{p_{j,t}} \quad \mathbb{E}_t \sum_{s=0}^{\infty} \phi^s \prod_{k=1}^{s} (r_{t+k})^{-1} \left(\left(\frac{p_{j,t}}{P_{t+s}}\right)^{1-\epsilon} Y_{t+s} - \frac{w_{t+s}}{A_{t+s}} \left(\frac{p_{j,t}}{P_{t+s}}\right)^{-\epsilon} Y_{t+s} \right) \tag{58}$$

Nominal dividends from the final good firm are given by:

$$P_t d_t^F = P_t Y_t - \int_0^1 p_{j,t} y_{j,t} dj, \tag{59}$$

and from the intermediate firms:

$$P_t d_t^I = \int_0^1 \left(p_{j,t} y_{j,t} - P_t w_t n_{j,t}\right) dj. \tag{60}$$

Total (real) dividends paid to unconstrained households are then:

$$d_t = \left(d_t^I + d_t^F\right)/\alpha \tag{61}$$

Interest Rates Nominal rates set by the central bank and the real rate that enters the household consumption problem are related as:

$$r_t = (1 + i_{t-1})\Pi_t^{-1} \tag{62}$$

where $\Pi_t = P_t/P_{t-1}$.

Aggregation and Market Clearing Aggregate labour and consumption are:

$$N_t = \int_0^1 n_{i,t}di = \int_0^1 n_{j,t}dj \tag{63}$$

$$C_t = \int_0^1 c_{i,t}di \tag{64}$$

The resource constraint is:

$$Y_t = C_t + G_t. \tag{65}$$

Exogenous Processes Productivity, government spending and transfers follow AR(1) processes:

$$\log A_t = \rho_A \log A_{t-1} + s_A\varepsilon_t^A, \tag{66}$$

$$\log G_t = (1 - \rho_G)\log \bar{G} + \rho_G \log G_{t-1} + s_G\varepsilon_t^g, \tag{67}$$

$$\mathcal{T}_t = (1 - \rho_{\mathcal{T}})\bar{\mathcal{T}} + \rho_{\mathcal{T}}\mathcal{T}_{t-1} + s_{\mathcal{T}}\varepsilon_t^s \tag{68}$$

Equilibrium Alongside the household optimisation problems (2)- (3), and the policy block (23)-(24), the above equations provide a complete description of the model. The equilibrium definition is then as standard – sequences of prices and allocations such that: households and firms behave optimally, monetary and fiscal policy follow their respective rules, and markets clear.

Steady State The model is approximated about a zero-inflation steady state: $\Pi = 1$. Government spending is a fixed fraction of output in steady state: $G = \psi_G Y$. And there are no government transfers in steady state: $\mathcal{T} = 0$.

A.4 Quantitative Analysis

For the quantitative analysis, I restrict $\delta_{\mathcal{U}}$ to be binary: $\delta_{\mathcal{U}} \in \{0, 1\}$, and allow $\delta_{\mathcal{H}}$ to take on values in the range $[0, 1]$. In this case, the model reduces to a "Three" Agent New Keynesian model with three representative households: an unconstrained household (who either receives the transfer, or not), a "treated" hand-to-mouth household (who receives the transfer), and a "control" hand-to-mouth household (who does not receive the transfer). In fact, bar the one minor addition of an additional set of hand-to-mouth agents, the model is essentially just the "textbook" TANK model laid out here by Eric Sims: `https://sites.nd.edu/esims/files/2024/04/notes_tank_sp2024.pdf`.

Equilibrium Conditions I summarise the equilibrium conditions for the quantitative model.

The optimality conditions for the unconstrained household give:

$$\theta N_{\mathcal{U},t}^{\chi} = C_{\mathcal{U},t}^{-\sigma} w_t, \tag{69}$$

$$1 = \mathbb{E}_t \left[(1+i_t) \Lambda_{\mathcal{U},t,t+1} \Pi_{t+1}^{-1} \right], \tag{70}$$

$$\Lambda_{\mathcal{U},t,t+1} = \beta \left(\frac{C_{\mathcal{U},t+1}}{C_{\mathcal{U},t}} \right)^{-\sigma}, \tag{71}$$

where $\Lambda_{\mathcal{U},t,t+1}$ denotes the stochastic discount factor.

The optimality conditions for the "treated" hand-to-mouth agent give:

$$C_{\mathcal{H},t}^{T} = w_t N_{\mathcal{H},t}^{T} + \mathcal{T}_t - \tau_t \tag{72}$$

$$\theta (N_{\mathcal{H},t}^{T})^{\chi} = w_t \, (C_{\mathcal{H},t}^{T})^{-\sigma}, \tag{73}$$

and likewise for the representative "control" hand-to-mouth agent:

$$C_{\mathcal{H},t}^{C} = w_t N_{\mathcal{H},t}^{C} - \tau_t \tag{74}$$

$$\theta (N_{\mathcal{H},t}^{C})^{\chi} = w_t \, (C_{\mathcal{H},t}^{C})^{-\sigma}. \tag{75}$$

Aggregate hand-to-mouth consumption and labour are then:

$$C_{\mathcal{H},t} = \delta_{\mathcal{H}} C_{\mathcal{H},t}^{T} + (1-\delta_{\mathcal{H}}) C_{\mathcal{H},t}^{C}, \tag{76}$$

$$N_{\mathcal{H},t} = \delta_{\mathcal{H}} N_{\mathcal{H},t}^{T} + (1-\delta_{\mathcal{H}}) N_{\mathcal{H},t}^{C}, \tag{77}$$

with total labour and consumption:

$$N_t = \alpha N_{\mathcal{U},t} + (1-\alpha) N_{\mathcal{H},t}, \tag{78}$$

$$C_t = \alpha C_{\mathcal{U},t} + (1-\alpha) C_{\mathcal{H},t}. \tag{79}$$

Deriving firms' optimality conditions follows standard steps from the basic New Keynesian model.[39] First note the optimal reset price from the pricing problem (58) gives:

$$p_{j,t}^{*} = \frac{\epsilon}{\epsilon - 1} \frac{\mathbb{E}_t \sum_{s=0}^{\infty} \phi^s \prod_{k=1}^{s} (r_{t+k})^{-1} \frac{w_{t+s}}{A_{t+s}} P_{t+s}^{\epsilon} Y_{t+s}}{\mathbb{E}_t \sum_{s=0}^{\infty} \phi^s \prod_{k=1}^{s} (r_{t+k})^{-1} P_{t+s}^{\epsilon-1} Y_{t+s}} \equiv \frac{\epsilon}{\epsilon - 1} \frac{X_{1,t}}{X_{2,t}}.$$

Using the law of iterated expectations:

$$X_{1,t} = \frac{w_t}{A_t} P_t^{\epsilon} Y_t + \phi \, \mathbb{E}_t \left[r_{t+1}^{-1} X_{1,t+1} \right],$$

$$X_{2,t} = P_t^{\epsilon-1} Y_t + \phi \, \mathbb{E}_t \left[r_{t+1}^{-1} X_{2,t+1} \right].$$

[39] See e.g. `https://sites.nd.edu/esims/files/2024/03/notes_new_keynesian_2024-1.pdf`.

Then define:

$$x_{1,t} \equiv X_{1,t}/P_t^{\epsilon}$$
$$x_{2,t} \equiv X_{2,t}/P_t^{\epsilon-1}$$
$$p_t^{\#} \equiv p_{j,t}^{*}/P_t$$

Equilibrium conditions from the price setting problem can then be summarised as:

$$p_t^{\#} = \frac{\epsilon}{\epsilon-1}\frac{x_{1,t}}{x_{2,t}}, \tag{80}$$

$$x_{1,t} = \frac{w_t}{A_t}Y_t + \phi\mathbb{E}_t\left[\frac{\Pi_{t+1}^{1+\epsilon}}{1+i_t}x_{1,t+1}\right], \tag{81}$$

$$x_{2,t} = Y_t + \phi\mathbb{E}_t\left[\frac{\Pi_{t+1}^{\epsilon}}{1+i_t}x_{2,t+1}\right]. \tag{82}$$

The final good producers' problem (54) implies:

$$y_{j,t} = (p_{j,t}/P_t)^{-\epsilon}Y_t. \tag{83}$$

The production function of intermediates (56) alongside the market clearing condition for labour then implies that aggregate production satisfies:

$$Y_t v_t = A_t N_t. \tag{84}$$

where $v_t \equiv \int_0^1 \left(\frac{p_{j,t}}{P_t}\right)^{-\epsilon} dj$ captures the degree of price dispersion. Finally, the assumption of staggered price-setting for intermediates yields:

$$v_t = (1-\phi)(p_t^{\#})^{-\epsilon} + \phi\Pi_t^{\epsilon}v_{t-1}, \tag{85}$$

$$1 = (1-\phi)(p_t^{\#})^{(1-\epsilon)} + \phi\Pi_t^{(\epsilon-1)}; \tag{86}$$

Monetary and fiscal policy are given by:

$$i_t = (1-\rho_i)\bar{i} + \rho_i i_{t-1} + (1-\rho_i)\phi_\pi(\log\Pi_t - \log\bar{\Pi}) + s_m\varepsilon_t^m \tag{87}$$

$$G_t = \tau_t - \alpha\delta_{\mathcal{U}}\mathcal{T}_t - (1-\alpha)\delta_{\mathcal{H}}\mathcal{T}_t \tag{88}$$

Finally, the resource constraint and the exogenous processes are given by (65)-(68). Using these equations, I solve for impulse responses via first-order perturbation methods in Dynare.

Parameterisation Table 2 reports the parameter values used in the simulations.[40] I keep these parameter values constant and then display results for different values of $\delta_{\mathcal{U}}$ and $\delta_{\mathcal{H}}$ in the main text. I also consider an extension to the model where transfers instead follow a policy rule of the form (28). For this exercise, I keep all parameter values the same and set $\phi_{\mathcal{T}} = 2.0$.

[40] These essentially just mirror the "standard" values for parameters employed in the illustrative simulations from Eric Sims' original TANK model – see `https://sites.nd.edu/esims/files/2024/04/notes_tank_sp2024.pdf`

Table 2: Calibration: TANK Model Parameters

Parameter	Value	Description
β	0.99	Discount factor
σ	1	Inverse EIS
χ	1	Inverse Frisch elasticity
θ	1	Utility weight on labour
α	0.6	Share of unconstrained households
ϕ	0.75	Calvo price stickiness
$\epsilon/(\epsilon-1)$	1.1	Desired markup of price over marginal cost
ψ_G	0.2	Govt. spending share of steady-state output
ϕ_π	1.5	Taylor rule inflation coefficient
ρ_A	0.9	Persistence of productivity
ρ_G	0.8	Persistence of government spending
ρ_i	0.8	Interest rate smoothing
$\rho_{\mathcal{T}}$	0.5	Persistence of transfers
s_A	0.01	Std. dev. productivity shock
s_G	0.01	Std. dev. government spending shock
s_m	0.0025	Std. dev. monetary shock
$s_{\mathcal{T}}$	0.01	Std. dev. transfer shock

Response of GE Variables To complement the discussion in Section 4 and provide intuition for the bias in Figure 1b, I plot the response of GE variables following the shock to unconstrained households (i.e. for the case where $\delta_{\mathcal{U}} = 1, \delta_{\mathcal{H}} = 0$). Figure 4a plots the response of (ex-ante) real rates, defined as:

$$r_t^{ea} = (1 + i_t)/\mathbb{E}_t \Pi_{t+1}$$

Figure 4b plots the response of the various components of income (as a fraction of the increase in aggregate transfers). In order to finance the transfer, taxes rise, generating a contraction in aggregate demand (since resources shift on-net to the Ricardian agents). Thus real rates and wages fall, while dividends rise (since profits are counter-cyclical in this model). Note that the hand-to-mouth agents *cut* their consumption in response to these GE effects (as they face lower wages and higher taxes). In contrast, Ricardian agents *increase* consumption in response to these GE effects (due to the intertemporal smoothing motive, given lower real interest rates). These differential responses to GE effects thus generate severe *upward* bias in β^{ss}.

Controlling for GE Aggregates In addition to the results from the main text, I also consider an estimation procedure in line with Proposition 3. Specifically, I construct estimates of β^{alt} from a regression that includes as additional controls (interacted with the shares z_i):

$$c_{i,t+h} = \alpha_i + \delta_t + \beta^{alt}(z_i \cdot \varepsilon_t^s) + \gamma'(z_i \cdot m_t) + e_{i,t} \tag{89}$$

where $m_t = [r_t, d_t, \tau_t, w_t]$ – i.e. exactly the aggregate GE variables that determine differ-

Figure 4: Response of GE Variables ($\delta_{\mathcal{U}} = 1, \delta_{\mathcal{H}} = 0$)

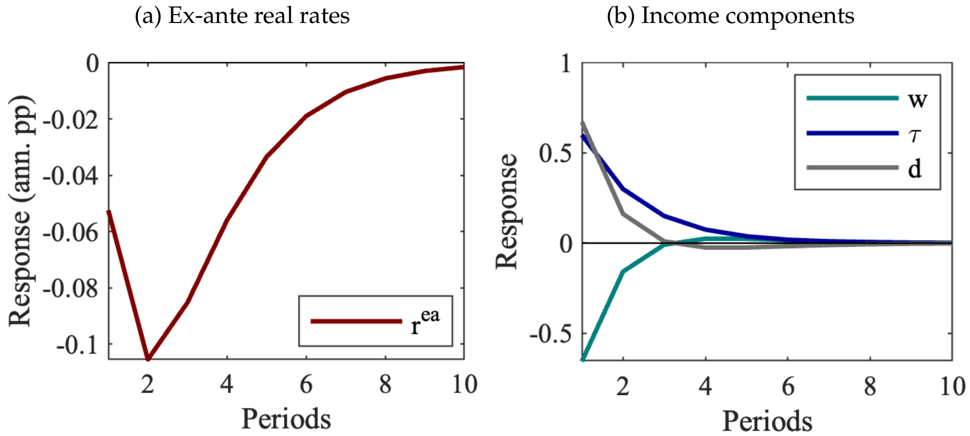

Notes: Maroon line in (a) plots the impulse response of ex-ante real rates, r_t^{ea}. Teal, navy blue and dark grey lines in (b) plot the impulse response of wages w_t, taxes τ_t and dividends d_t respectively (all scaled to the contemporaneous rise in transfers $\Theta_0^{\mathcal{T}}$).

Figure 5: Controlling For GE Variables

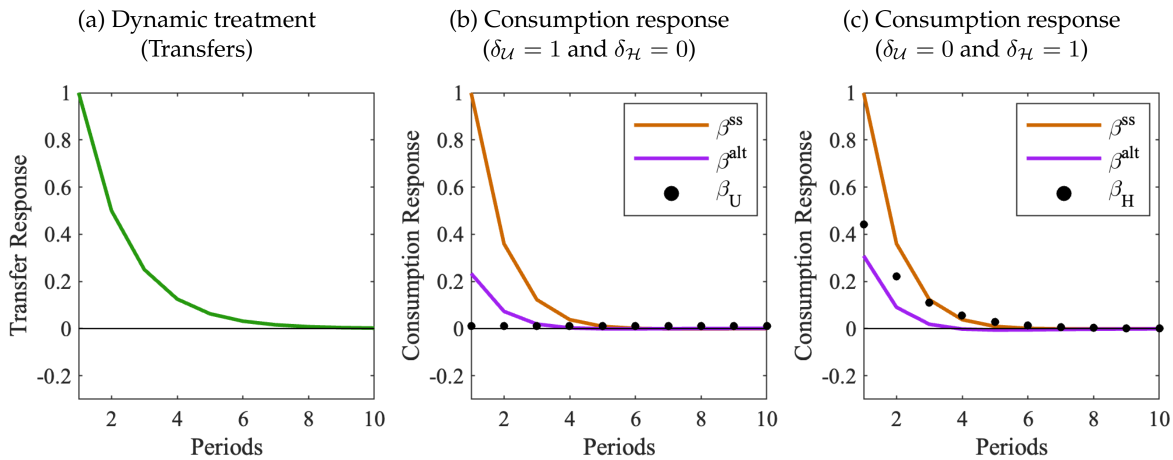

Notes: Green line in (a) plots the impulse response of transfers following the shock. Orange and purple lines in (b) and (c) denote β^{ss} from (26) and β^{alt} from (89) respectively for the case with: $\delta_{\mathcal{U}} = 1$, $\delta_{\mathcal{H}} = 0$ and $\delta_{\mathcal{U}} = 0$, $\delta_{\mathcal{H}} = 1$ respectively. Black dots in (b) and (c) denote the "direct" response of unconstrained and hand-to-mouth households to the shock respectively. All responses are scaled to the contemporaneous rise in transfers i.e. to $\Theta_0^{\mathcal{T}}$.

ential responses across treated and control groups in this setting. Figure 5 shows my results, where I report estimates of $\boldsymbol{\beta}^{ss}$ from (26) (with no controls), alongside estimates of $\boldsymbol{\beta}^{alt}$ from (89) in the purple line. This alternate procedure generally performs poorly – delivering similar estimates regardless of whether the transfers are targeted to Ricardians (with low iMPCs) or hand-to-mouth agents (with high iMPCs).

Additional Results Figure 6 displays results from regression (27). As before, Figure 6a plots the path for transfers following the shock. Figure 6b then demonstrates that $\boldsymbol{\beta}^{ss}$ from (27), shown in the orange lines, recovers exactly the hand-to-mouth households' marginal propensities to consume, shown in the black dots.

Figure 6: Controlling for the Targeting Rule

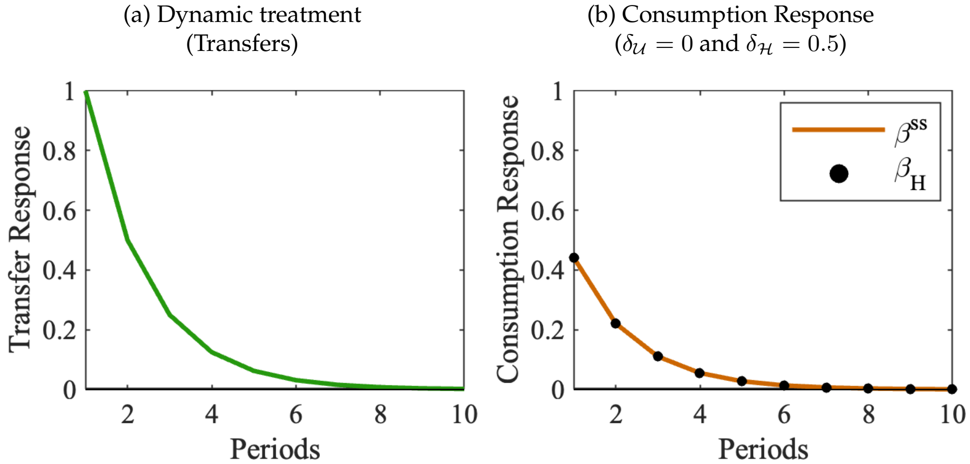

Notes: Green line in (a) plots the impulse response of transfers following the shock. Orange line in (b) denotes β^{ss} from (27) for the case with: $\delta_{\mathcal{U}} = 0$ and $\delta_{\mathcal{H}} = 0.5$. Black dots in (b) denote the "direct" response of treated hand-to-mouth households to the shock, defined by (25). All responses are scaled to the contemporaneous rise in transfers i.e. to $\Theta_0^{\mathcal{T}}$.

B Section 3 Proofs

B.1 Proof of Proposition 1

The shift-share coefficient is equivalent to:

$$\beta^{ss} = \frac{\mathbb{E}[y_{i,t+h} s_i d_t]}{\mathbb{E}\left[(s_i d_t)^2\right]}$$

Plugging in equation (12):

$$\beta^{ss} = \frac{\mathbb{E}\left[\left(\sum_{l=0}^{\infty} \Theta^y_{i,l} \varepsilon_{t+h-l} + \bar{y}_i\right) s_i \sum_{m=0}^{\infty} \Psi_m \varepsilon_{t-m}\right]}{\mathbb{E}\left[\left(s_i \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}\right)^2\right]}$$

Since $(\bar{y}_i, s_i) \perp \{\varepsilon_{t-l}\}_{l=-\infty}^{\infty}$ and shocks are mean-zero:

$$\beta^{ss} = \frac{\mathbb{E}\left[s_i \left(\sum_{l=0}^{\infty} \Theta^y_{i,l} \varepsilon_{t+h-l}\right) \left(\sum_{m=0}^{\infty} \Psi_m \varepsilon_{t-m}\right)\right]}{\mathbb{E}\left[\left(s_i \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}\right)^2\right]}$$

Since $\forall j \neq k : \varepsilon_{t-j} \perp \varepsilon_{t-k}$ and shocks are mean-zero:

$$\beta^{ss} = \frac{\mathbb{E}\left[\left(s_i \sum_{l=0}^{\infty} \left(\Theta^y_{i,h+l} \varepsilon_{t-l}\right) \left(\Psi_l \varepsilon_{t-l}\right)\right)\right]}{\mathbb{E}\left[\left(s_i \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}\right)^2\right]}$$

Switching terms:

$$\beta^{ss} = \frac{\sum_{l=0}^{\infty} \mathbb{E}\left[\left(s_i \left(\Theta^y_{i,h+l}\right) \varepsilon_{t-l} \left(\Psi_l \varepsilon_{t-l}\right)\right)\right]}{\mathbb{E}\left[\left(s_i \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}\right)^2\right]}$$

Note that $(s_i, x_i, z_i) \perp \{\varepsilon_{t-l}\}_{l=0}^{\infty} \implies \left(s_i, \Theta^y_{i,h+l}\right) \perp \{\varepsilon_{t-l}\}_{l=0}^{\infty}$. So:

$$\beta^{ss} = \frac{\sum_{l=0}^{\infty} \mathbb{E}\left[s_i \Theta^y_{i,h+l}\right] \mathbb{E}\left[\varepsilon_{t-l} \left(\Psi_l \varepsilon_{t-l}\right)\right]}{\mathbb{E}\left[s_i^2\right] \mathbb{E}\left[\left(\sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}\right)^2\right]}$$

And so defining η_l as the $n_\varepsilon \times 1$ vector:

$$\eta_l = \frac{\mathbb{E}\left[\varepsilon_{t-l} \left(\Psi_l \varepsilon_{t-l}\right)\right]}{\mathbb{E}\left[\left(\sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}\right)^2\right]}$$

We have the desired expression:

$$\beta^{ss} = \frac{\mathbb{E}\left[\sum_{l=0}^{\infty} \Theta^y_{i,h+l} \eta_l s_i\right]}{\mathbb{E}\left[s_i^2\right]}$$

Note that since the shocks are mutually independent and unit-variance we have:

$$\eta_{l,k} = \frac{\psi_{l,k}}{\sum_{k=1}^{n_\varepsilon} \sum_{l=0}^{\infty} \psi_{l,k}^2}$$

B.2 Proof of Proposition 2

By Frisch–Waugh–Lovell, the shift-share coefficient is equivalent to:

$$\beta^{ss} = \frac{\mathbb{E}[y_{i,t+h}\tilde{g}_{i,t}]}{\mathbb{E}[\tilde{g}_{i,t}^2]}$$

where $\tilde{g}_{i,t}$ denotes the residual from a (population) regression of $g_{i,t} \equiv z_i \varepsilon_t^s$ on unit- and time-fixed effects. We then find:

$$\tilde{g}_{i,t} = (z_i - \mathbb{E}[z_i])(\varepsilon_t^s - \mathbb{E}[\varepsilon_t^s]) = \tilde{z}_i \varepsilon_t^s$$

where $\tilde{z}_i$ denotes de-meaned z_i. By Proposition 1:

$$\beta^{ss} = \frac{\mathbb{E}\left[\Theta_{i,h}^{ys} \tilde{z}_i\right]}{\mathbb{E}\left[(\tilde{z}_i)^2\right]}$$

where note this can be interpreted as the slope coefficient in a population regression of $\Theta_{i,h,1}^{y}$ on z_i and a constant. For binary exposures $z_i \in \{0, 1\}$, this coefficient captures the difference in conditional expectations across groups:

$$\beta^{ss} = \mathbb{E}[\Theta_{i,h}^{ys} | z_i = 1] - \mathbb{E}[\Theta_{i,h}^{ys} | z_i = 0]$$

The desired conclusion then follows by plugging in (14):

$$\beta^{ss} = \mathbb{E}[\beta_{i,h}^{y} | z_i = 1] + \mathbb{E}[\zeta_{i,h}^{y} | z_i = 1] - \mathbb{E}[\zeta_{i,h}^{y} | z_i = 0] \quad \square$$

B.3 Proof of Proposition 3

By Frisch–Waugh–Lovell, the shift-share coefficient is equivalent to:

$$\beta^{ss} = \frac{\mathbb{E}[y_{i,t+h}\tilde{g}_{i,t}]}{\mathbb{E}[\tilde{g}_{i,t}^2]}$$

where $\tilde{g}_{i,t}$ denotes the residual from a (population) regression of $g_{i,t} \equiv z_i \varepsilon_t^s$ on unit- and time-fixed effects and $z_i w_t$. We then find:

$$\tilde{g}_{i,t} = (z_i - \mathbb{E}[z_i])(\varepsilon_t^s)^{\perp w_t} = \tilde{z}_i (\varepsilon_t^s)^{\perp w_t}$$

where $(\varepsilon_t^s)^{\perp w_t}$ denotes the residual from a regression of ε_t^s on a constant and w_t. Regressing ε_t^s on a constant and w_t recovers a linear combination of shocks up to time-t, specifically:

$$(\varepsilon_t^s)^{\perp w_t} = \underbrace{\left(1 + b'\Theta_{0,1}^w\right)}_{\equiv \psi_{0,1}} \varepsilon_t^s \; + \sum_{k=2}^{n_\varepsilon} \underbrace{b'\Theta_{0,k}^w}_{\equiv \psi_{0,k}} \varepsilon_{k,t} \; + \sum_{l \geq 1} \underbrace{b'\Theta_l^w}_{\equiv \psi_l} \varepsilon_{t-l}$$

where:

$$b' = -(\Theta_{0,1}^w)' \, \Sigma_w^{-1}, \qquad \Sigma_w = \sum_{l=0}^{\infty} \Theta_l^w (\Theta_l^w)', \qquad \psi_l = [\psi_{l,1} ... \psi_{l,n_\varepsilon}]$$

By Proposition 1:

$$\beta^{ss} = \frac{\mathbb{E}[\sum_{l=0}^{\infty} \Theta_{i,h+l}^y \kappa_l \tilde{z}_i]}{\mathbb{E}[\tilde{z}_i^2]}$$

where κ_l is n_ε-length vector with k'th element equal to: $\kappa_{l,k} = \frac{\psi_{l,k}}{\sum_{k=1}^{n_\varepsilon} \sum_{l=0}^{\infty} \psi_{l,k}^2}$. Note that for binary, de-meaned z_i the shift-share coefficient captures differences in conditional means:

$$\beta^{ss} = \sum_{l=0}^{\infty} \sum_{k=1}^{n_\varepsilon} \kappa_{l,k} \left(\mathbb{E}[\Theta_{i,h+l,k}^y | z_i = 1] - \mathbb{E}[\Theta_{i,h+l,k}^y | z_i = 0] \right)$$

Then applying the decomposition (14):

$$\begin{aligned}\beta^{ss} &= \kappa_{0,1} \mathbb{E}[\beta_{i,h}^y | z_i = 1] + \kappa_{0,1} \left(\mathbb{E}[\zeta_{i,h}^y | z_i = 1] - \mathbb{E}[\zeta_{i,h}^y | z_i = 0] \right) \\ &+ \sum_{k=2}^{n_\varepsilon} \kappa_{0,k} \left(\mathbb{E}[\Theta_{i,h,k}^y | z_i = 1] - \mathbb{E}[\Theta_{i,h,k}^y | z_i = 0] \right) + \sum_{l=1}^{\infty} \sum_{k=1}^{n_\varepsilon} \kappa_{l,k} \left(\mathbb{E}[\Theta_{i,h+l,k}^y | z_i = 1] - \mathbb{E}[\Theta_{i,h+l,k}^y | z_i = 0] \right) \quad \square\end{aligned}$$

Finally, note that since ε_t^s is an OLS residual, $\mathbb{E}[\varepsilon_t^s (\varepsilon_t^s)^{\perp w_t}] = \mathbb{E}[((\varepsilon_t^s)^{\perp w_t})^2]$, and so $\kappa_{0,1} = 1$.

B.4 Proof of Proposition 4

By Frisch–Waugh–Lovell, the shift-share coefficient is equivalent to:

$$\beta^{ss} = \frac{\mathbb{E}[y_{i,t+h} \tilde{g}_{i,t}]}{\mathbb{E}[\tilde{g}_{i,t}^2]}$$

where $\tilde{g}_{i,t}$ denotes the residual from a (population) regression of $g_{i,t} \equiv z_i \varepsilon_t^s$ on unit- and time-fixed effects and $\zeta_{i,h}^y \varepsilon_t^s$. We then find:

$$\tilde{g}_{i,t} = \tilde{z}_i (\varepsilon_t^s - \mathbb{E}[\varepsilon_t^s])$$

where $\tilde{z}_i$ denotes the residual from a regression of z_i on $\zeta_{i,h}^y$ and a constant. Using Proposition 1:

$$\beta^{ss} = \frac{\mathbb{E}\left[\Theta^{ys}_{i,h}\tilde{z}_i\right]}{\mathbb{E}\left[(\tilde{z}_i)^2\right]}$$

Plugging in (14):

$$\beta^{ss} = \frac{\mathbb{E}\left[\beta^{y}_{i,h}z_i\tilde{z}_i\right]}{\mathbb{E}\left[(\tilde{z}_i)^2\right]} + \frac{\mathbb{E}\left[\zeta^{y}_{i,h}\tilde{z}_i\right]}{\mathbb{E}\left[(\tilde{z}_i)^2\right]}$$

Since $\tilde{z}_i$ is an OLS-population residual, $\mathbb{E}\left[\zeta^{y}_{i,h}\tilde{z}_i\right] = 0$. In addition:

$$\begin{aligned}\mathbb{E}\left[(\tilde{z}_i)^2\right] &= \mathbb{E}[(z_i - c - b\zeta^{y}_{i,h})\tilde{z}_i] \\ &= \mathbb{E}[z_i\tilde{z}_i]\end{aligned}$$

where c and b denote OLS coefficients from the regression of z_i on a constant and $\zeta^{y}_{i,h}$, and the final line follows from the fact OLS-residuals are uncorrelated with regressors. And so:

$$\beta^{ss} = \frac{\mathbb{E}\left[\beta^{y}_{i,h}z_i\tilde{z}_i\right]}{\mathbb{E}\left[z_i\tilde{z}_i\right]} \quad \square$$

B.5 Proof of Proposition 5

By Frisch–Waugh–Lovell, the shift-share coefficient is equal to:

$$\beta^{ss} = \frac{\mathbb{E}[y_{i,t+h}\tilde{g}_{i,t}]}{\mathbb{E}[\tilde{g}^2_{i,t}]}$$

where $\tilde{g}_{i,t}$ denotes the residual from a (population) regression of $g_{i,t} \equiv z_i\varepsilon^s_t$ on unit- and time-fixed effects and $a_i\varepsilon^s_t$. We then find:

$$\tilde{g}_{i,t} = \tilde{z}_i(\varepsilon^s_t - \mathbb{E}[\varepsilon^s_t])$$

where $\tilde{z}_i$ denotes the population residual from a regression of z_i on a constant and a_i. By Proposition 1:

$$\beta^{ss} = \frac{\mathbb{E}\left[\Theta^{ys}_{i,h}\tilde{z}_i\right]}{\mathbb{E}\left[(\tilde{z}_i)^2\right]}$$

Plugging in (14):

$$\beta^{ss} = \frac{\mathbb{E}\left[\beta^y_{i,h} z_i \tilde{z}_i\right] + \mathbb{E}\left[\zeta^y_{i,h} \tilde{z}_i\right]}{\mathbb{E}\left[(\tilde{z}_i)^2\right]}$$

Note that under (19), since the conditional expectation function is linear:

$$\tilde{z}_i = z_i - \alpha - \delta' a_i$$

which implies: $\mathbb{E}[\tilde{z}_i | x_i, a_i] = 0$. And so:

$$\begin{aligned}
\mathbb{E}[\beta^y_{i,h} z_i \tilde{z}_i] + \mathbb{E}[\zeta^y_{i,h} \tilde{z}_i] &= \mathbb{E}[\beta^y_{i,h} \mathbb{E}[\tilde{z}_i z_i | x_i, a_i]] + \mathbb{E}[\zeta^y_{i,h} \mathbb{E}[\tilde{z}_i | x_i, a_i]] \\
&= \mathbb{E}[\beta^y_{i,h} \mathbb{E}[\tilde{z}_i z_i | x_i, a_i]]
\end{aligned}$$

Finally since:

$$\begin{aligned}
\mathbb{E}[\tilde{z}_i z_i | x_i, a_i] &= \mathbb{E}[\tilde{z}_i (\tilde{z}_i + c + \delta' a_i) | x_i, a_i] \\
&= \mathbb{E}[\tilde{z}_i^2 | x_i, a_i]
\end{aligned}$$

it then follows that:

$$\beta^{ss} = \frac{\mathbb{E}\left[\beta^y_{i,h} \mathbb{E}[\tilde{z}_i^2 | x_i, a_i]\right]}{\mathbb{E}[\mathbb{E}[\tilde{z}_i^2 | x_i, a_i]]} \quad \square$$

B.6 Proof of Proposition 6

Start with $\boldsymbol{\beta}^{ss}_y$. Using Proposition 1 and Frisch–Waugh–Lovell, we then find:

$$\boldsymbol{\beta}^{ss}_y = \frac{\sum_k w_k \mathbb{E}\left[\boldsymbol{\Theta}^y_{i,k} \tilde{z}_i\right]}{\mathbb{E}[\tilde{z}_i^2]}$$

where $\tilde{z}_i$ denotes the population residual from a regression of z_i on a constant and a_i, and $w_k = \frac{q_k}{\sum_{k=1}^{n_\varepsilon} (q_k)^2}$. Next recall the decomposition (22):

$$\boldsymbol{\Theta}^y_{i,k} = \mathcal{M}^{yg}_i z_i \boldsymbol{\Theta}^g_k + \mathcal{M}^{yw}_i \boldsymbol{\Theta}^w_k$$

Then by conditional exogeneity of the shares (as in (19)):

$$\forall k : \mathbb{E}[\mathcal{M}^{yw}_i \boldsymbol{\Theta}^w_k \tilde{z}_i] = \mathbb{E}[\mathcal{M}^{yw}_i \boldsymbol{\Theta}^w_k \mathbb{E}[\tilde{z}_i | x_i, a_i]] = 0$$

where the first equality uses LIE and that $\mathcal{M}^{yw}_i$ is measurable with respect to x_i, and the second equality uses that $\mathbb{E}[\tilde{z}_i | x_i, a_i] = 0$. And so returning to the expression for $\boldsymbol{\beta}^{ss}_y$:

$$\boldsymbol{\beta}^{ss}_y = \frac{\sum_k w_k \mathbb{E}[\mathcal{M}^{yg}_i \boldsymbol{\Theta}^g_k z_i \tilde{z}_i]}{\mathbb{E}[\tilde{z}_i^2]}$$

Next:

$$\mathbb{E}[\mathcal{M}^{yg}_i \boldsymbol{\Theta}^g_k z_i \tilde{z}_i] = \mathbb{E}[\mathcal{M}^{yg}_i \boldsymbol{\Theta}^g_k \mathbb{E}[z_i \tilde{z}_i | x_i, a_i]] = \mathbb{E}[\mathcal{M}^{yg}_i \boldsymbol{\Theta}^g_k \mathbb{E}[\tilde{z}_i^2 | x_i, a_i]]$$

where again this uses LIE alongside the fact that $\mathcal{M}_i^{yg}$ is measurable with respect to x_i, then (as in the proof of Proposition 5): $\mathbb{E}[z_i\tilde{z}_i|x_i,a_i] = \mathbb{E}[\tilde{z}_i^2|x_i,a_i]$. Then letting $\omega_i = \mathbb{E}[\tilde{z}_i^2|x_i,a_i]$:

$$\begin{aligned}\boldsymbol{\beta}_y^{ss} &= \frac{\sum_k w_k \mathbb{E}[\mathcal{M}_i^{yg}\boldsymbol{\Theta}_k^g\omega_i]}{\mathbb{E}[\omega_i]} \\ &= \frac{\mathbb{E}[\mathcal{M}_i^{yg}\sum_k w_k \boldsymbol{\Theta}_k^g\omega_i]}{\mathbb{E}[\omega_i]}\end{aligned}$$

Now turn to β_g^{ss}. Using Proposition 1 and Frisch–Waugh–Lovell, we again find:

$$\boldsymbol{\beta}_g^{ss} = \frac{\sum_k w_k \mathbb{E}\left[\boldsymbol{\Theta}_{i,k}^g \tilde{z}_i\right]}{\mathbb{E}[\tilde{z}_i^2]}$$

Next, since $g_{i,t} = z_i g_t$:

$$\boldsymbol{\Theta}_{i,k}^g = z_i \boldsymbol{\Theta}_k^g$$

And so noting that $\mathbb{E}[\tilde{z}_i z_i] = \mathbb{E}[\tilde{z}_i^2]$, we have:

$$\boldsymbol{\beta}_g^{ss} = \sum_k w_k \boldsymbol{\Theta}_k^g$$

And so finally:

$$\boldsymbol{\beta}_y^{ss} = \frac{\mathbb{E}[\mathcal{M}_i^{yg}\boldsymbol{\beta}_g^{ss}\omega_i]}{\mathbb{E}[\omega_i]} \quad \square$$

C Supplementary Theoretical Results

In this Appendix, I present various supplementary theoretical results. I begin by describing two other recent approaches to identification within my environment, and then discuss various extensions to the baseline framework.

C.1 Time-Series Decomposition Method

In this Appendix, I demonstrate an alternate route to identification that dispenses with any exogeneity condition on the shares z_i, building on the logic of Donaldson (2026). The method works by constructing a linear combination of shocks that leaves the entire dynamic path of the GE aggregates w_t unchanged while moving the aggregate treatment g_t. My environment differs somewhat from the setting in Donaldson (2026) – incorporating both heterogeneous "direct effects" as well as multiple channels of GE transmission. The below explanation thus departs slightly from the description of the method in Donaldson (2026) – extending it in a natural way to my environment.[41]

Setup Suppose the sequence-space decomposition (22) holds, and that all impulse responses are truncated at some maximum horizon H, i.e. $\Theta_l = \Theta_{i,l} = 0$ for all $l > H$. Suppose the researcher has access to an n_d-length vector of instruments d_t satisfying the lead/lag exogeneity condition:

$$d_t = Q'\varepsilon_t$$

for an $n_\varepsilon \times n_d$ matrix Q with full column rank (requiring $n_\varepsilon \geq n_d$). Denote the instrument variance matrix $\Sigma_d \equiv \text{Var}(d_t) = Q'Q$, which is non-singular under full column rank.

Procedure The method proceeds in three steps:

Step 1: Run local projections of each element of w_{t+h} on the vector d_t, for $h = 0, 1, \ldots, H$. Collect the coefficients in the $n_w(H+1) \times n_d$ matrix Γ_w. Do the same for g_{t+h}, collecting coefficients in the $(H+1) \times n_d$ matrix Γ_g.

Step 2: Find an n_d-length vector b^* satisfying:

$$\Gamma_w b^* = \mathbf{0} \quad \text{s.t.} \quad \Gamma_g b^* \neq \mathbf{0} \tag{90}$$

[41] My object of interest differs slightly from Donaldson (2026) - see, in particular, the discussion in that paper on biases with "two-way general equilibrium feedback". Donaldson (2026) defines the parameter of interest as the impulse response of $y_{i,t}$ to the particular shock of interest, ε_t^s under a counterfactual monetary policy under which interest rates do not respond to ε_t^s. Below, I'm instead interested in recovering the impulse response of $y_{i,t}$ to some (recoverable) path for g_t hold fixed interest rates (and other relevant aggregates). Note this can instead be interpreted as the effect of some particular (but unknown) combination of government spending news shocks that brings about that estimated path for g_t. This is consistent with the interpretation of time series local projections as the effect of (unknown) combinations of news shocks, as in McKay and Wolf (2023) – exactly the same interpretation employed by Donaldson (2026) in the case of monetary policy shocks.

Step 3: Form the scalar instrument $d_t^* = (b^*)'\Sigma_d^{-1} d_t$ and run the pair of shift-share local projections for $h = 0, 1, \ldots, H$:

$$y_{i,t+h} = \alpha_{i,y,h} + \delta_{t,y,h} + \beta^{ss}_{y,h}(z_i \cdot d_t^*) + e_{i,t,y,h} \tag{91}$$

$$g_{i,t+h} = \alpha_{i,g,h} + \delta_{t,g,h} + \beta^{ss}_{g,h}(z_i \cdot d_t^*) + u_{i,t,g,h} \tag{92}$$

Identification Then, supposing a vector b^* satisfying (90) exists, it can be shown that the shift-share coefficients from (91)-(92) satisfy:

$$\boldsymbol{\beta}^{ss}_y = \frac{\mathbb{E}\left[\mathcal{M}^{yg}_i \boldsymbol{\beta}^{ss}_g \, \omega_i\right]}{\mathbb{E}[\omega_i]}, \qquad \omega_i = z_i(z_i - \mathbb{E}[z_i])$$

for $\boldsymbol{\beta}^{ss}_y = \{\beta^{ss}_{y,h}\}_{h=0}^H$ and $\boldsymbol{\beta}^{ss}_g = \{\beta^{ss}_{g,h}\}_{h=0}^H$.

To see why this is the case, note that, through standard arguments, we have:

$$\beta^{ss}_{y,h} \propto \mathbb{E}\left[y_{i,t+h}\, \tilde{z}_i d_t^*\right] = \mathbb{E}\left[\tilde{z}_i \sum_{l=0}^{\infty} \Theta^y_{i,l}\, \mathbb{E}[\varepsilon_{t+h-l}\varepsilon_t']\, v\right] = \mathbb{E}\left[\tilde{z}_i\, \Theta^y_{i,h} v\right] \tag{93}$$

where $\tilde{z}_i = z_i - \mathbb{E}[z_i]$ and $v = Q\Sigma_d^{-1} b^*$. Then applying the decomposition (22):

$$\Theta^y_{i,h} v = \sum_{k=1}^{n_\varepsilon} v_k\, \Theta^y_{i,h,k} = \underbrace{\left[\mathcal{M}^{yg}_i \left(\boldsymbol{\Theta}^g v\right)\right]_h z_i}_{\text{Direct Effect}} + \underbrace{\left[\mathcal{M}^{yw}_i \left(\boldsymbol{\Theta}^w v\right)\right]_h}_{\text{GE Effect}}$$

where, given the construction of b^* in Step 2, $\boldsymbol{\Theta}^w v = \mathbf{0}$, and so the GE effect term vanishes. Repeating this argument across horizons and then applying it to $\beta^{ss}_{g,h}$ yields the result, as in the proof of Proposition 6.

Subject to concerns around negative weights, the shift-share regression now recovers the object of interest. In fact, negative weights can be avoided in this case through an alternative procedure that replaces Step 3 with a series of local projections (one for each unit-i) of $y_{i,t+h}$ on d_t^*. In this case, it can be shown that the unit-specific local projection coefficients, $\beta^{ols}_{i,h}$, recover:

$$\boldsymbol{\beta}^{ols}_i = \mathcal{M}^{yg}_i \boldsymbol{\beta}^{ss}_g z_i$$

from which weighted averages can then be constructed as desired by the researcher. Intuitively, d_t^* mimics a "micro" experiment, which moves treatment for all agents while simultaneously holding any GE adjustments to the shock fixed. Thus identification can be achieved by directly projecting on d_t^* for each unit individually.

Discussion For this approach to succeed, the researcher must observe the complete list of GE aggregates that respond to the shock, w_t, and must have access to a large number of valid instruments sufficient to span the response of these aggregates across all horizons. In particular, a vector b^* satisfying (90) exists if and only if:

$$\text{rank}\begin{bmatrix}\Gamma_w \\ \Gamma_g\end{bmatrix} > \text{rank}\left(\Gamma_w\right)$$

Economically, this requires the existence of a combination of structural shocks – within the span of the researcher's instruments – that moves the path of g_t while leaving the entire path of w_t unchanged at every horizon. Since Γ_w has $n_w(H+1)$ rows and n_d columns, this condition generally requires $n_d > n_w(H+1)$, so that the researcher requires at least $n_w(H+1)+1$ valid instruments – and correspondingly $n_\varepsilon \geq n_d$ underlying structural shocks for Q to have full column rank. In practice – following the discussion in Donaldson (2026) – this procedure can alternatively be implemented by finding the vector b^* that *minimises* $\Gamma_w b^*$, which generally requires fewer instruments and provides an approximation to d_t^*. The informational requirements are generally strong along two dimensions: the researcher must correctly specify the *complete* list of GE aggregates w_t through which the shock transmits and must have access to a number of valid, contemporaneous (non-anticipated) instruments that approximately shuts down all these GE responses simultaneously at all horizons. Armed with this sufficiently large set of time series instruments, the researcher can then dispense with any exogeneity assumptions regarding the exposure shares z_i.

Seen in this light, the method discussed here is highly complementary to my proposed approach, achieving identification under a very different set of informational requirements. The procedure in Donaldson (2026) is likely well-suited to settings where the researcher is willing to impose that the set of GE aggregates w_t that generate confounding is relatively narrow – and has confidence in a relatively large set of time-series instruments – leaving open that the set of unit-characteristics that determine heterogeneity in response to GE aggregates x_i is relatively large and allowing the implicit "targeting rule" for the exposure shares to be unknown. In contrast, my proposed approach is well-suited to settings where the researcher is able to credibly partial out the appropriate set of characteristics from the targeting rule, without knowledge of the underlying GE effects of the shock under study, and lacking a sufficiently large set of time-series instruments to "subtract off" these effects.

C.2 Synthetic Control Method

In this Appendix, I discuss the synthetic control method from Arkhangelsky and Korovkin (2024). I first introduce the method at a high-level then discuss its application within a HA-DSGE setting. For simplicity – and to ease intuition – I focus discussion on identification within the stylised TANK model employed throughout the paper.

Method The original paper introduces the method in the context of IV estimation. For simplicity, I focus here solely on estimation of the "reduced form".[42] I also abstract from any finite sample estimation, considering directly the population version of the authors' algorithm.[43]

[42] In fact, this is without loss in my baseline environment since the first-stage term anyway contains no information on the w-weights given the assumption that $g_{i,t} = z_i g_t$.

[43] In finite samples, the authors' proposed estimator uses a learning sample to estimate the weights w, where the learning step includes a ridge penalty that vanishes as the learning sample grows. I abstract from this.

To understand the method, consider first the shift-share regression from the main text run at $h = 0$:

$$y_{i,t} = \alpha_i + \delta_t + \beta^{ss} \left(z_i \cdot \varepsilon_t^s\right) + e_{i,t} \tag{94}$$

Rather than running the above OLS regression, the "robust" method proposed by Arkhangelsky and Korovkin (2024) instead first weights the aggregate panel cross-sectionally:

$$Y_t(w) = \mathbb{E}[w_i y_{i,t}] \tag{95}$$

where the weights are chosen to minimise the variance of $Y_t(w)$ after projecting out ε_t^s, subject to: $\mathbb{E}[w_i] = 0$ and $\mathbb{E}[w_i z_i] = 1$. Then the response of interest is recovered via a regression of $Y_t(w)$ on ε_t^s.

Identification in TANK Now consider the version of the TANK model from Appendix A.4, with transfers targeted towards half the hand-to-mouth agents ($\delta_{\mathcal{U}} = 0, \delta_{\mathcal{H}} = 0.5$). In this case the constraints, $\mathbb{E}[w_i] = 0$ and $\mathbb{E}[w_i z_i] = 1$, imply that the weighted aggregate from (95) depends only on some scalar, λ:[44]

$$C_t(w) = C_t(\lambda) = (C^T_{\mathcal{H},t} - C^C_{\mathcal{H},t}) + \lambda(C^C_{\mathcal{H},t} - C_{\mathcal{U},t}) = D_t + \lambda R_t \tag{96}$$

where $C^T_{\mathcal{H},t}$ and $C^C_{\mathcal{H},t}$ denote the consumption of treated and control hand-to-mouth agents respectively, $C_{\mathcal{U},t}$ denotes the consumption of unconstrained households, and we define $D_t \equiv (C^T_{\mathcal{H},t} - C^C_{\mathcal{H},t})$ and $R_t \equiv (C^C_{\mathcal{H},t} - C_{\mathcal{U},t})$. The optimal-$\lambda$ chosen by the "robust" method is then defined as:

$$\lambda^* = \arg\min_{\lambda}\{V(\lambda)\} = \arg\min_{\lambda}\{Var[\tilde{C}_t(\lambda)]\} \tag{97}$$

where $\tilde{C}_t(\lambda)$ is $C_t(\lambda)$ after projecting out ε_t^s. Now let $\Theta^D_{l,k}$ and $\Theta^R_{l,k}$ collect impulse responses for D_t and R_t to shock k at horizon l and order the transfer shock first such that impulse responses to this shock are given by: $\Theta^D_{l,1}, \Theta^R_{l,1}$. Then by the SVMA structure, and since shocks are mutually independent, i.i.d. and unit-variance, we have:

$$V(\lambda) = \sum_{(l,k)\in\Omega} (\Theta^D_{l,k} + \lambda\Theta^R_{l,k})^2 \tag{98}$$

where Ω is the set of (l,k) pairs for $k \in \{1, ..., n_\varepsilon\}$ and $l \geq 0$, other than $(0,1)$. Note that λ^* then admits a simple closed-form solution:

$$\lambda^* = \frac{-b}{c}, \tag{99}$$

where:

$$b \equiv \sum_{(l,k)\in\Omega} \Theta^D_{l,k}\Theta^R_{l,k}, \quad c \equiv \sum_{(l,k)\in\Omega} \left(\Theta^R_{l,k}\right)^2. \tag{100}$$

[44] $\mathbb{E}[w_i z_i] = 1$ pins down the weight on treated hand-to-mouth agents, while $\mathbb{E}[w_i] = 0$ pins down the difference in remaining weights. We then have $\lambda = -\alpha w_{\mathcal{U}}$, where $w_{\mathcal{U}}$ is the weight selected for the unconstrained households, and α is their population share.

And note that finally, this implies the "robust" weighting method returns:

$$\beta^{rob} = \underbrace{\Theta^D_{0,1}}_{\text{ATT}} - \underbrace{\frac{b}{c}\Theta^R_{0,1}}_{\text{Bias}} \tag{101}$$

Note the first term is exactly the object of interest – differencing responses of treated and control hand-to-mouth households to the transfer shock gives the MPC out of the transfer: $\Theta^D_{0,1} = \mathcal{M}^{\mathcal{T}}_{\mathcal{H}}\Theta^{\mathcal{T}}_{0,1}$. The second term can then be interpreted as bias. Since the hand-to-mouth and Ricardian agents generally respond differently to the transfer shock, we have $\Theta^R_{0,1} \neq 0$, and thus we require $b = 0$ for the bias to be removed.

When would this be the case? Note that when transfers are exogenous, the treated and control hand-to-mouth groups respond the same to all shocks *except* the transfer shock, where the difference is given by the MPC out of the transfer path. That is, we have:

$$\Theta^D_{l,k} = 0 \text{ for } k \neq 1, \quad \Theta^D_{l,1} = \mathcal{M}^{\mathcal{T}}_{\mathcal{H}}\Theta^{\mathcal{T}}_{l,1}. \tag{102}$$

When the transfer shock exhibits no persistence ($\forall l \neq 0 : \Theta^{\mathcal{T}}_{l,1} = 0$), we thus have: $b = 0$. More generally, when the transfer shock itself is persistent, or when transfers additionally depend on other shocks hitting the economy, the bias term remains. Intuitively, the bias arises since the synthetic control method forms a control group by searching for households whose consumption behaviour is similar to treated households after partialling out the shock of interest. However, treated and control hand-to-mouth households exhibit different consumption behaviour even after partialling out ε^s_t, since they also respond differently to any other shock that moves transfers (including *past* transfer shocks). The synthetic control algorithm thus seeks to "counterbalance" these differences by placing some weight on Ricardian households when forming $C_t(w)$. This is undesirable precisely because the Ricardian agents respond to the GE effects of the shock under study very differently from the treated group.

Discussion The bias of the synthetic control method in a HA-DSGE setting is not a particularly novel finding – as Arkhangelsky and Korovkin (2024) note, the method is explicitly designed to operate in settings without dynamic treatment effects, and where the confounding force stems purely from other shocks hitting the economy (rather than the GE effects of the shock under study per se).[45]

Relative to the approach explored in this paper, a key advantage of the synthetic control method is that it does not require the researcher to take a stand on the exogeneity of the exposure shares z_i – forming an appropriate control group purely by learning from the data. An interesting open question is thus whether the method could be adjusted to operate within a HA-DSGE setting. For example, in a TANK setting, when transfers are exogenous (i.e. re-

[45]Specifically, they write: "*Our method is designed for applications where the TSLS regression is a priori reasonable, but users worry about potential unobserved confounders. There are multiple reasons why...(TSLS)...might fail, other than omitted variables. For example, the underlying model can be nonlinear, or the dynamic effects of the past treatments can be sizable.*"

sponding to transfer shocks only) it is straightforward to show that the method succeeds so long as the researcher has direct access to the transfer shocks and partials out the entire *history* of transfer shocks from $C_t(w)$ when searching for the optimal weights.[46] When transfers respond to other shocks – likely a more realistic assumption in this setting – there is no clear partialling out step that avoids the bias. Somewhat speculatively then, for the synthetic control method to succeed in a HA-DSGE setting seems to require an appropriate exogeneity condition on the aggregate treatment – g_t in my framework – while leaving open the possibility of correlation between the exposure shares (z_i) and underlying sources of unit-heterogeneity (x_i). This is different from the method proposed in this paper which instead allows arbitrary endogeneity in g_t – and indeed leverages this fact to generate heterogeneity in dynamic treatment paths by exploiting different underlying aggregate shocks – while restricting the endogeneity of the exposure shares z_i. Appropriately adapted, the two methods could thus be useful complements – achieving identification under very different informational requirements. I leave this for future research.

C.3 Extension 1: Imperfect exposure measure

I suppose the econometrician has access to some exposure-measure $s_i = s(x_i, z_i)$. Consider a natural extension of conditional exogeneity in this setting: $\mathbb{E}[s_i|x_i, a_i] = \alpha + \delta' a_i$. Then consider pairs of regressions of the form:

$$y_{i,t+h} = \alpha_{i,y,h} + \delta_{t,y,h} + \beta^{ss}_{y,h}(s_i \cdot \varepsilon^s_t) + \lambda'_{y,h}(a_i \cdot \varepsilon^s_t) + e_{i,t,y,h}$$
$$g_{i,t+h} = \alpha_{i,g,h} + \delta_{t,g,h} + \beta^{ss}_{g,h}(s_i \cdot \varepsilon^s_t) + \lambda'_{g,h}(a_i \cdot \varepsilon^s_t) + e_{i,t,g,h}$$

Following the same steps as Proposition 5, we have:

$$\beta^{ss}_{y,h} = \frac{\mathbb{E}[\beta^y_{i,h}\omega_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}$$

where $\omega_i = \mathbb{E}[\tilde{s}_i z_i|x_i, a_i]$ and $\tilde{s}_i$ denotes the population residual from a regression of s_i on a constant and a_i. Recalling the definition of $\beta^y_{i,h}$, we can write the full sequence of coefficients as:

$$\boldsymbol{\beta}^{ss}_y = \frac{\mathbb{E}[\mathcal{M}^{yg}_i \boldsymbol{\Theta}^{gs}\omega_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}.$$

Also, by Proposition 1, we have for the coefficient from the second regression:

$$\beta^{ss}_{g,h} = \frac{\mathbb{E}[\Theta^{gs}_{i,h}\tilde{s}_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}.$$

And since $g_{i,t} = z_i g_t$, we have:

$$\Theta^{gs}_{i,h} = z_i \Theta^{gs}_h.$$

[46]While achievable in theory, this may be challenging in finite samples when the weights are learned over an initial learning sample.

Following standard steps we then arrive at:

$$\boldsymbol{\beta}_g^{ss} = \boldsymbol{\Theta}^{gs} \frac{\mathbb{E}[\omega_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}.$$

Assuming a standard relevance condition, $\mathbb{E}[\tilde{s}_i z_i] \neq 0$, rearranging for $\boldsymbol{\Theta}^{gs}$ and plugging this into the expression for $\boldsymbol{\beta}_y^{ss}$ we arrive at the same expression as Proposition 6:

$$\boldsymbol{\beta}_y^{ss} = \frac{\mathbb{E}[\mathcal{M}_i^{yg} \boldsymbol{\beta}_g^{ss} \omega_i]}{\mathbb{E}[\omega_i]}.$$

The resulting expression has the desired form (with non-negative weights) under a "monotonicity" assumption, that for all values of (x_i, a_i):

$$\mathrm{Cov}(s_i, z_i | x_i, a_i) \geq 0 \text{ or } \mathrm{Cov}(s_i, z_i | x_i, a_i) \leq 0$$

C.4 Extension 2: Dynamic treatment heterogeneity

I consider the same setting as Appendix C.3, but now assuming the more general decomposition of impulse responses to the shock:

$$\boldsymbol{\Theta}_i^{ys} = \underbrace{\mathcal{M}_i^{yg} \boldsymbol{\Theta}_i^{gs}}_{\text{Response to own treatment path}} + \underbrace{\mathcal{M}_i^{yw} \boldsymbol{\Theta}^{ws}}_{\text{Response to GE channels}}$$

where $\boldsymbol{\Theta}_i^{gs}$ is some (unknown) function of (x_i, z_i). Following the same steps as Appendix C.3 we have:

$$\begin{aligned}\boldsymbol{\beta}_y^{ss} &= \frac{\mathbb{E}[\boldsymbol{\Theta}_i^{ys} \tilde{s}_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]} \\ &= \frac{\mathbb{E}\left[\mathcal{M}_i^{yg} \mathbb{E}[\boldsymbol{\Theta}_i^{gs} \, \tilde{s}_i \mid x_i, a_i]\right]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}.\end{aligned}$$

We then likewise have:

$$\boldsymbol{\beta}_g^{ss} = \frac{\mathbb{E}[\boldsymbol{\Theta}_i^{gs} \tilde{s}_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}.$$

Then defining,

$$\boldsymbol{\beta}_{i,g}^{ss} = \frac{\mathbb{E}[\boldsymbol{\Theta}_i^{gs} \, \tilde{s}_i \mid x_i, a_i]}{\mathbb{E}\left[(\tilde{s}_i)^2\right]}$$

we have:

$$\boldsymbol{\beta}_y^{ss} = \mathbb{E}\left[\mathcal{M}_i^{yg} \boldsymbol{\beta}_{i,g}^{ss}\right].$$

And so finally, assuming an appropriate dynamic relevance condition, $\forall f : \beta_{g,f}^{ss} \neq 0$, we have:

$$\beta_{y,h}^{ss} = \sum_{f=0}^{\infty} \beta_{g,f}^{ss} \frac{\mathbb{E}\left[[\mathcal{M}_i^{yg}]_{h,f} \, \omega_{i,f}\right]}{\mathbb{E}[\omega_{i,f}]}$$

where: $\omega_{i,f} = \beta_{i,g,f}^{ss}$. In this case, then, $\boldsymbol{\beta}_y^{ss}$ and $\boldsymbol{\beta}_g^{ss}$ are jointly informative about some weighted average of appropriate entries of the intertemporal MPC matrix $\mathcal{M}_i^{yg}$. Further, negativity of the weights within each horizon is avoided whenever $\beta^{ss}_{i,g,f}$ has the same sign across units. This amounts to a monotonicity condition that holds individually within each horizon – i.e. the conditional covariance between s_i and each unit's horizon-specific treatment $\Theta^{gs}_{i,f}$ should have the same sign for all values of (x_i, a_i):

$$\forall h : \text{Cov}(s_i, \Theta^{gs}_{i,h}|x_i, a_i) \geq 0 \ \text{ or } \ \text{Cov}(s_i, \Theta^{gs}_{i,h}|x_i, a_i) \leq 0 \tag{103}$$

C.5 Extension 3: Cross-sectional data

I extend the results in Section 3 to cover settings with cross-sectional data. I keep the data generating process the same but treat the shocks as non-stochastic in this case – I denote expectations throughout as $\mathbb{E}_i$ to make this clear. Specifically I suppose the econometrician has access to only a single government spending shock that occurs at some fixed period-t and uses cross-sectional local projections to estimate the direct (dynamic) causal effect of the shock. Consider the following further assumptions in this setting:

$$\{\varepsilon^s_{t'}\}_{t' \neq t} = 0 \ \text{ and } \ \forall k \neq 1 : \boldsymbol{\Theta}^g_k = 0 \tag{104}$$

These assumptions ensure that the researcher observes some one-off stimulus shock where the aggregate path for the stimulus is pinned down purely by that one-off shock. This is a natural condition, ensuring that the estimand of interest is interpretable as recovering the effect of some specific (unanticipated) stimulus policy.[47]

Now consider a cross-sectional regression for some single time-series observation-t of the form:

$$y_{i,t+h} - y_{i,t-1} = \alpha + \beta^{ols}(z_i \cdot \varepsilon^s_t) + \lambda' a_i + e_i.$$

Note that under the standard sequence-space decomposition (22), outcomes can be written as:

$$y_{i,t+h} = \beta^y_{i,h} z_i \varepsilon^s_t + \zeta^y_{i,h}\varepsilon^s_t + \bar{y}_i + u_{i,t+h},$$

where the first term is the "direct" effect, the second term is the "GE" effect and $u_{i,t+h}$ collects the effect of all other shocks. And so taking differences:

$$y_{i,t+h} - y_{i,t-1} = \beta^y_{i,h} z_i \varepsilon^s_t + \zeta^y_{i,h}\varepsilon^s_t + \tilde{u}_{i,t+h}.$$

Further, given (104), and given the decomposition (22), $\tilde{u}_{i,t+h}$ depends only on x_i (not on z_i), since no other shocks operate via the g-channel. Turning then finally to β^{ols}, and letting $\tilde{z}_i$ denote the (population) residual from a regression of z_i on a constant and a_i:

[47] Note that g can be defined narrowly in this case – e.g. not as aggregate transfers or government spending, but just the portion of transfers or government spending that was adjusted at time-t and which the researcher is interested in studying. This is in line with the typical interpretation in applied work which seeks to recover an estimate of the effect of some specific government stimulus shock – e.g. the ARRA stimulus package.

$$
\begin{aligned}
\beta^{ols} &= \frac{\mathbb{E}_i\left[\left(\beta^y_{i,h} z_i \varepsilon^s_t + \zeta^y_{i,h}\varepsilon^s_t + \tilde{u}_{i,t+h}\right)(\tilde{z}_i \varepsilon^s_t)\right]}{\mathbb{E}_i\left[(\tilde{z}_i\varepsilon^s_t)^2\right]} \\
&= \frac{\mathbb{E}_i[\beta^y_{i,h} z_i \tilde{z}_i]}{\mathbb{E}_i[\tilde{z}_i^2]} + \frac{\mathbb{E}_i[\zeta^y_{i,h}\tilde{z}_i]}{\mathbb{E}_i[\tilde{z}_i^2]} + \frac{\varepsilon^s_t \mathbb{E}_i[\tilde{u}_{i,t+h}\tilde{z}_i]}{(\varepsilon^s_t)^2\,\mathbb{E}_i[\tilde{z}_i^2]} \\
&= \frac{\mathbb{E}_i[\beta^y_{i,h} z_i \tilde{z}_i]}{\mathbb{E}_i[\tilde{z}_i^2]}
\end{aligned}
$$

where the final line uses: $\mathbb{E}_i[\zeta^y_{i,h}\tilde{z}_i] = \mathbb{E}_i[\zeta^y_{i,h}\mathbb{E}_i[\tilde{z}_i|x_i, a_i]] = 0$ and $\mathbb{E}_i[\tilde{u}_{i,t+h}\tilde{z}_i] = \mathbb{E}_i[\tilde{u}_{i,t+h}\mathbb{E}_i[\tilde{z}_i|x_i, a_i]] = 0$. Following the same steps as the proof to Proposition 5, we then have:

$$
\beta^{ols} = \frac{\mathbb{E}_i[\beta^y_{i,h}\omega_i]}{\mathbb{E}_i[\omega_i]}
$$

where $\omega_i = \mathbb{E}_i[\tilde{z}_i^2|x_i, a_i]$. As in Proposition 5, the weights can also be expressed as $\tilde{z}_i^2$.

C.6 Extension 4: Time-varying heterogeneity

I now extend results to settings with time-varying heterogeneity.

Environment Suppose outcomes for each agent i at each point in time t can be expressed as:

$$
\mathbb{E}_t[y_{i,t+h}] - \mathbb{E}_{t-1}[y_{i,t+h}] = \Theta^y_{i,t,h}\varepsilon_t + v_{i,t} \tag{105}
$$

where $\mathbb{E}_t[.]$ denotes expectations conditional on the information set generated by the history of structural shocks and the full cross-section of agent-level variables up to t, Ω_t, and $v_{i,t}$ depends on time-t idiosyncratic shocks experienced by each agent. As before assume macro outcomes, $Y_t = [y_t, g_t, w_t]$, can be expressed via the linear SVMA representation:

$$
Y_t = \Theta(L)\varepsilon_t + \bar{Y} \tag{106}
$$

Suppose the impulse responses to shock-s can be decomposed as in the baseline environment:

$$
\boldsymbol{\Theta}^{ys}_{i,t} = \underbrace{\mathcal{M}^{yg}(x_{i,t-1})\boldsymbol{\Theta}^{gs} z_{i,t-1}}_{\text{Direct Effect}} + \underbrace{\mathcal{M}^{yw}(x_{i,t-1})\boldsymbol{\Theta}^{ws}}_{\text{GE Effect}} \tag{107}
$$

where $z_{i,t-1}$ and $x_{i,t-1}$ denote time-varying state variables. I assume all state variables are predetermined at time-t such that all impulse responses depend only on the pre-time-t information set:

$$
\left(z_{i,t-1}, x_{i,t-1}, \boldsymbol{\Theta}^y_{i,t}\right) \in \Omega_{t-1} \tag{108}
$$

Finally, assume shocks are mean-zero and unit-variance, and that standard independence conditions hold:

$$
\varepsilon^s_t \perp \left(\Omega_{t-1}, \varepsilon^{-s}_t, v_{i,t}\right) \tag{109}
$$

where this permits the interpretation of ε^s_t as a structural aggregate shock – independent of other (idiosyncratic and aggregate) shocks and all past information. Note the setup here essentially maintains the core assumptions of the baseline environment, but now allowing unit-level impulse responses to depend on time-varying and potentially endogenous (but pre-determined) state variables – as in e.g. canonical (linearised) HANK models.[48] I now go on to consider identification via shift-share regressions in this environment.

Identification Consider the following regression:

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss}\left(z_{i,t-1} \cdot \varepsilon_t^s\right) + \lambda'\left(a_{i,t-1} \cdot \varepsilon_t^s\right) + u_{i,t} \tag{110}$$

for some control variables $a_{i,t-1} \in \Omega_{t-1}$. By FWL, given shocks are mean-zero, unit-variance and independent of the state-variables:

$$\beta^{ss} = \frac{\mathbb{E}[y_{i,t+h}\tilde{z}_{i,t-1}\varepsilon_t^s]}{\mathbb{E}[(\tilde{z}_{i,t-1})^2]}$$

where $\tilde{z}_{i,t-1}$ is $z_{i,t-1}$ orthogonalised with respect to time-FE as well as $a_{i,t-1}$. Using (105), we then have:

$$\mathbb{E}[y_{i,t+h}\tilde{z}_{i,t-1}\varepsilon_t^s] = \mathbb{E}\left[\left(\Theta_{i,t,h}^{y}\varepsilon_t + v_{i,t} + \mathbb{E}_{t-1}[y_{i,t+h}] + r_{i,t+h}\right)\tilde{z}_{i,t-1}\varepsilon_t^s\right]$$

where $r_{i,t+h} \equiv y_{i,t+h} - \mathbb{E}_t[y_{i,t+h}]$. Then using (109) we have:

$$\begin{aligned}
\mathbb{E}\left[\left(\Theta_{i,t,h}^{ys}\varepsilon_t^s\right)\tilde{z}_{i,t-1}\varepsilon_t^s\right] &= \mathbb{E}\left[\left(\Theta_{i,t,h}^{ys}\tilde{z}_{i,t-1}\right)\mathbb{E}[(\varepsilon_t^s)^2 \mid \Omega_{t-1}, \varepsilon_t^{-s}]\right] = \mathbb{E}\left[\Theta_{i,t,h}^{ys}\tilde{z}_{i,t-1}\right] \\
\mathbb{E}\left[\left(\Theta_{i,t,h}^{y,-s}\varepsilon_t^{-s}\right)\tilde{z}_{i,t-1}\varepsilon_t^s\right] &= \mathbb{E}\left[\left(\Theta_{i,t,h}^{y,-s}\varepsilon_t^{-s}\right)\tilde{z}_{i,t-1}\mathbb{E}[\varepsilon_t^s \mid \Omega_{t-1}, \varepsilon_t^{-s}]\right] = 0 \\
\mathbb{E}\left[\left(v_{i,t}\tilde{z}_{i,t-1}\right)\varepsilon_t^s\right] &= \mathbb{E}\left[\left(v_{i,t}\tilde{z}_{i,t-1}\right)\mathbb{E}[\varepsilon_t^s \mid \Omega_{t-1}, v_{i,t}]\right] = 0 \\
\mathbb{E}\left[\left(\mathbb{E}_{t-1}[y_{i,t+h}]\tilde{z}_{i,t-1}\right)\varepsilon_t^s\right] &= \mathbb{E}\left[\left(\mathbb{E}_{t-1}[y_{i,t+h}]\tilde{z}_{i,t-1}\right)\mathbb{E}[\varepsilon_t^s \mid \Omega_{t-1}]\right] = 0
\end{aligned}$$

And by definition of $r_{i,t+h}$:

$$\mathbb{E}\left[r_{i,t+h}\tilde{z}_{i,t-1}\varepsilon_t^s\right] = \mathbb{E}\left[\left(\tilde{z}_{i,t-1}\varepsilon_t^s\right)\mathbb{E}[r_{i,t+h} \mid \Omega_t]\right] = 0$$

Hence:

$$\beta^{ss} = \frac{\mathbb{E}\left[\Theta_{i,t,h}^{ys}\tilde{z}_{i,t-1}\right]}{\mathbb{E}[(\tilde{z}_{i,t-1})^2]}$$

Then suppose the exposure shares satisfy the following conditional exogeneity condition:

$$\mathbb{E}[z_{i,t-1} \mid x_{i,t-1}, a_{i,t-1}, t] = \mu_t + \eta' a_{i,t-1} \tag{111}$$

Then following the same reasoning as Proposition 5:

$$\beta^{ss} = \frac{\mathbb{E}[\beta_{i,t,h}\omega_{i,t}]}{\mathbb{E}[\omega_{i,t}]} \tag{112}$$

where $\beta_{i,t,h}$ is the h'th element of the sequence $\mathcal{M}^{yg}(x_{i,t-1})\Theta^{gs}$ and $\omega_{i,t} = \mathbb{E}[(\tilde{z}_{i,t-1})^2 \mid x_{i,t-1}, a_{i,t-1}, t] \geq 0$.

[48] See related discussion in David et al. (2026) who study state-dependent local projections in a similar environment.

Discussion The above discussion highlights that identification in this setting can proceed in essentially the same manner as in the main environment. Note that (111) can be understood as considering a setting where, at the start of each new period, the policymaker redistributes the exposure shares $z_{i,t-1}$, targeting them according to units' state variables, $a_{i,t-1}$ – e.g. targeting units with higher levels of income and wealth at the start of the period. In this case, controlling for the time-varying state variables in the targeting rule, $a_{i,t-1}$, is again sufficient for identification.

As in the main environment, it can be shown that the identifying conditions on the aggregate shock variable can be loosened in line with Proposition 6 – allowing correlation with *other* aggregate shocks whose propagation is also governed by a similar decomposition as (107). Thus, as well as being unforecastable – i.e. independent of Ω_{t-1} – the key requirement for the aggregate shifter in this environment is that it is independent of the idiosyncratic (or "local") shocks $v_{i,t}$.

C.7 Extension 5: Finite-N

I extend the results in Section 3 to cover settings with a finite population. I first discuss the adjustments to the environment and then discuss identification.

Environment I now consider a setting with a finite population $i \in \{1, 2, ..., n\}$. I keep the rest of the environment broadly the same, but now define ε_t^s as an $n-$length vector of random variables: $\varepsilon_t^s = [\varepsilon_{1t}^s, ..., \varepsilon_{nt}^s]'$, and likewise for other shocks so that ε_t^{-s} is a $((n_\varepsilon - 1)n)$-dimensional vector. I collect these shocks together in a $(n_\varepsilon \cdot n)$- dimensional vector ε_t, ordering ε_t^s first as before. I assume shocks are mean-zero, (vector-wise) mutually independent – so e.g. $\varepsilon_t^s \perp \varepsilon_t^{-s}$ – and independent over time, $\varepsilon_t \perp \{\varepsilon_{t-l}\}_{l \neq 0}$. Note this allows for correlation between unit-level shocks – i.e. $\mathbb{E}[\varepsilon_{i,t}^s \varepsilon_{j,t}^s] \neq 0$. The unit-level outcome of interest $y_t = [y_{1t}, ..., y_{nt}]'$ evolves as before according to:

$$y_t = \sum_{l=0}^{\infty} \Theta_l^y \varepsilon_{t-l} + \bar{y} \tag{113}$$

where Θ_l^y is a $n \times (n_\varepsilon \cdot n)$ dimensional matrix, and $\bar{y} = [\bar{y}_1, ..., \bar{y}_n]'$ collects steady-state outcomes for each unit.

Object of Interest Mirroring the main text, to focus on the object of interest, the effect of government stimulus shocks ε_t^s on each unit at some horizon h can be decomposed as:

$$\Theta_{i,h,1}^y \varepsilon_t^s = \underbrace{\beta_{i,h}^y \varepsilon_{i,t}^s}_{\text{Direct Effect}} + \underbrace{\sum_{j \neq i} \zeta_{i,j,h}^y \varepsilon_{j,t}^s}_{\text{Spillover Effect}} \tag{114}$$

where $\Theta^y_{i,h,1}$ is a $1 \times n$-length (row) vector that picks the relevant entries of Θ^y_h and captures the effect of government stimulus shocks for all units ε^s_t on unit-i. As before the object of interest is the direct effect $\beta^y_{i,h}$. In regional models this captures how the outcome of interest (e.g. income) in region-i responds to higher government spending in region-i (holding fixed spending in all other regions).

Identification I proceed analogously to Section 3, first highlighting the specific challenges for identification that arise in this setting and then proceeding to my proposed solution. To this end, first consider OLS-regressions including unit and time fixed effects:

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{fe} \varepsilon^s_{i,t} + e_{i,t} \tag{115}$$

Note that the OLS estimand in this case is:

$$\beta^{fe} = \frac{\mathbb{E}[\sum_i y_{i,t+h} \widetilde{\varepsilon}^s_{i,t}]}{\mathbb{E}[\sum_i (\widetilde{\varepsilon}^s_{i,t})^2]}$$

where $\widetilde{\varepsilon}^s_{i,t} = \varepsilon^s_{i,t} - \frac{1}{n}\sum_{k=1}^n \varepsilon^s_{k,t}$ denotes the cross-sectionally demeaned shock. Using (113) and (114), this can be written as:

$$\beta^{fe} = \frac{\mathbb{E}[\sum_i (\beta^y_{i,h} \varepsilon^s_{i,t} + \sum_{j \neq i} \zeta^y_{i,j,h} \varepsilon^s_{j,t} + u_{i,t+h}) \widetilde{\varepsilon}^s_{i,t}]}{\mathbb{E}[\sum_i (\widetilde{\varepsilon}^s_{i,t})^2]}$$

where $u_{i,t+h}$ collects steady-state values plus the effect of all other shocks other than time-t government spending shocks, ε^s_t. Since shocks are mean-zero, mutually independent and independent over time, $\mathbb{E}[u_{i,t+h} \widetilde{\varepsilon}^s_{i,t}] = 0$. Finally, since $\sum_i \widetilde{\varepsilon}^s_{i,t} = 0$, the denominator satisfies $\sum_i (\widetilde{\varepsilon}^s_{i,t})^2 = \sum_i \widetilde{\varepsilon}^s_{i,t} \varepsilon^s_{i,t} = \sum_i \widetilde{w}_{i,t}$. And so:

$$\beta^{fe} = \frac{\sum_i \beta^y_{i,h} \mathbb{E}[\widetilde{w}_{i,t}]}{\sum_i \mathbb{E}[\widetilde{w}_{i,t}]} + \underbrace{\frac{\sum_i \sum_{j \neq i} \zeta^y_{i,j,h} \mathbb{E}[\widetilde{\varepsilon}^s_{i,t} \varepsilon^s_{j,t}]}{\sum_i \mathbb{E}[\widetilde{w}_{i,t}]}}_{\text{Bias from GE spillovers}}$$

where $\widetilde{w}_{i,t} = \widetilde{\varepsilon}^s_{i,t} \varepsilon^s_{i,t}$. As in Proposition 2, the first term captures a weighted average of the direct effects of interest – but where these weights can in general take on negative values. The second term then captures bias stemming from spillovers between units. To understand the bias, it is helpful to first consider the case where units' government spending shocks are mutually independent. In that case, with common shock-variance across units, the bias term reduces to:

$$-\frac{1}{n(n-1)} \sum_i \sum_{j \neq i} \zeta^y_{i,j,h}$$

which captures the average spillover effect from any *single* unit to all other units. This is exactly the "de minimis" bias term in cross-sectional regressions discussed in Chodorow-Reich (2020) when individual units are non-infinitesimal. This bias is unlikely to be a core concern in macroeconomic applications, since the bias vanishes at rate $1/n$. More concerning is that, even though ε_t^s is orthogonal to *other* shocks, government-spending shocks may be correlated across units. In a regional setting, for example, the bias would be large when government spending shocks are distributed across regions in such a way that regions with high-$\varepsilon_{i,t}^s$ simultaneously experienced higher (or lower) spending in regions with particularly close linkages, such that they were simultaneously more (or less) exposed to spillovers arising from the shock. I illustrate this quantitatively in Appendix F.

I now demonstrate that identification can alternatively proceed as in the previous section, by appropriately including controls that capture the *assignment* of government spending shocks across units (without knowledge of the true "linkages" between units). To this end, suppose that the shocks are assigned across units according to the following "targeting rule":

$$\varepsilon_{i,t}^s = \delta' a_{i,t} + \eta_{i,t},$$

where $a_{i,t}$ is a n_a-length vector of observables and $\eta_t = [\eta_{1,t}, ..., \eta_{n,t}]'$ denotes a mean-zero, vector of random variables where:

$$\forall i : \eta_{i,t} \perp \left(\{\eta_{j,t}\}_{j \neq i}, \{a_{j,t}\}_{j=1}^n, \varepsilon_t^{-s}, \{\varepsilon_{t-l}\}_{l \neq 0} \right).$$

Then consider the following regression controlling for $a_{i,t}$ alongside unit- and time- fixed effects:

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{fe} \varepsilon_{i,t}^s + \lambda' a_{i,t} + e_{i,t}$$

By Frisch–Waugh–Lovell, β^{fe} can be obtained by demeaning each variable within periods and partialling out $\widetilde{a}_{i,t} = a_{i,t} - \dfrac{1}{n} \sum_{k=1}^n a_{k,t}$. Demeaning the targeting rule within periods gives $\widetilde{\varepsilon}_{i,t}^s = \delta' \widetilde{a}_{i,t} + \widetilde{\eta}_{i,t}$, where $\widetilde{\eta}_{i,t} = \eta_{i,t} - \frac{1}{n} \sum_{k=1}^n \eta_{k,t}$. Since η_t are mean-zero and independent of the variables in the targeting rule, the pooled regression of $\widetilde{\varepsilon}_{i,t}^s$ on $\widetilde{a}_{i,t}$ thus recovers δ with residual $\widetilde{\eta}_{i,t}$. And so:

$$\begin{aligned} \beta^{fe} &= \frac{\mathbb{E}[\sum_i y_{i,t+h} \widetilde{\eta}_{i,t}]}{\mathbb{E}[\sum_i (\widetilde{\eta}_{i,t})^2]} \\ &= \frac{\mathbb{E}[\sum_i (\beta_{i,h}^y (\delta' a_{i,t} + \eta_{i,t}) + \sum_{j \neq i} \zeta_{i,j,h}^y (\delta' a_{j,t} + \eta_{j,t}) + u_{i,t+h}) \widetilde{\eta}_{i,t}]}{\mathbb{E}[\sum_i (\widetilde{\eta}_{i,t})^2]} \end{aligned}$$

Since $\widetilde{\eta}_{i,t}$ is a linear combination of the time-t assignment shocks η_t, all terms of the form $\mathbb{E}[a_{j,t} \widetilde{\eta}_{i,t}]$ and $\mathbb{E}[u_{i,t+h} \widetilde{\eta}_{i,t}]$ vanish. The denominator here satisfies $\sum_i (\widetilde{\eta}_{i,t})^2 = \sum_i \widetilde{\eta}_{i,t} \eta_{i,t} = \sum_i \widetilde{\omega}_{i,t}$. Finally, since η_t are mutually independent and mean-zero, for $j \neq i$:

$$\mathbb{E}[\eta_{j,t} \widetilde{\eta}_{i,t}] = \mathbb{E}[\eta_{j,t} \eta_{i,t}] - \frac{1}{n} \sum_k \mathbb{E}[\eta_{j,t} \eta_{k,t}] = -\frac{1}{n} \mathbb{E}[(\eta_{j,t})^2]$$

And so collecting terms we get:

$$\beta^{fe} = \underbrace{\frac{\sum_i \beta^y_{i,h} \mathbb{E}[\widetilde{\omega}_{i,t}]}{\sum_i \mathbb{E}[\widetilde{\omega}_{i,t}]}}_{\text{Weighted Average Direct Effect}} - \underbrace{\frac{\frac{1}{n}\sum_i \sum_{j \neq i} \zeta^y_{i,j,h}\, \mathbb{E}\left[(\eta_{j,t})^2\right]}{\sum_i \mathbb{E}[\widetilde{\omega}_{i,t}]}}_{\text{De Minimis Bias}}$$

where the weights for each unit are given by:

$$\mathbb{E}[\tilde{\omega}_{i,t}] = \mathbb{E}[\tilde{\eta}_{i,t}\eta_{i,t}] = \mathbb{E}[\eta^2_{i,t}] - \frac{1}{n}\sum_k \mathbb{E}[\eta_{i,t}\eta_{k,t}] = \left(1 - \frac{1}{n}\right)\mathbb{E}[\eta^2_{i,t}] \geq 0.$$

Similar in spirit to Proposition 5, the implication of this result is that, when bias arises due to "clustering" of government spending shocks across units, this can be removed by controlling for the underlying unit-observables that feature in the targeting rule that drives the clustering. This removes the bias term while simultaneously ensuring non-negative weights. Note that the "de minimis" bias term remains, although again this is unlikely to be a core concern with sufficiently large n. I again demonstrate this with a quantitative model of fiscal spending in US states in Appendix F.

D Additional Empirical Results: Nakamura and Steinsson (2014)

This Appendix collects supplementary information and additional results from the empirical application in Section 5.1.

D.1 Data

The main dataset comes directly from the replication materials from Nakamura and Steinsson (2014), from which I use data on: state-level per capita output and per capita military procurement spending, national military spending, and the average level of military spending in each state relative to state output.

I supplement this with data on state-level MPCs and openness from the replication materials provided by Bellifemine et al. (2023). The authors construct estimates of MPCs at the county level. I convert these to the state-level by applying population weights (with population data taken from the Census Bureau's Population Estimates). The authors also construct county-level estimates of the share of employment in the non-tradable sector. I use population weights to construct state-level estimates, and then construct a measure of openness as one minus the share of employment in the non-tradable sector. I proxy for steady-state values by taking (pre-Covid) averages over the periods in which each series is available (2011–2019 for MPCs, 1986–2019 for openness). I further supplement this with data from the Correlates of State Policy Database (Grossmann et al., 2021), accessed via the *cspp* R-package. Table 4 details the variable names and definitions. All series are averages over the period 1966–1970 (to match the period in which the military shares are constructed), except the final three demographic series which are available only annually and use data from 1970 (or averages over 1960 and 1980 when unavailable).

Finally, I use various aggregate data at annual frequency from FRED: inflation, the Federal Funds rate, goods consumption, government military expenditures, government social benefits to persons and the spot price of West Texas Intermediate crude oil, averaged from monthly to annual frequency – see Table 3. I use the inflation variable to construct an annual CPI index. For the consumption, military expenditure, transfers and oil price series, I first create real variables by deflating by CPI and then take logs. I construct a real Federal Funds rate series by subtracting annual inflation from nominal rates. I supplement this with the extended Barro and Redlick (2011) data on average marginal tax rates from Mertens and Olea (2018).

D.2 Predictability of Shares

Table 5 reports results from regression (32) described in the main text.

Table 3: FRED Series

Variable	FRED code
Consumer price inflation	FPCPITOTLZGUSA
Federal Funds rate	RIFSPFFNA
Goods consumption	DGDSRC1A027NBEA
Federal military expenditures	G160461A027NBEA
Government social benefits to persons	A084RC1A027NBEA
Spot crude oil price (WTI)	WTISPLC

Table 4: Series from Correlates of State Policy Database

Variable Name	Variable Definition
incomepcap	Per capita personal income: personal income received by residents from all sources, divided by midyear population.
incshare_top10	Share of total income earned by the top 10% of earners.
incshare_top1	Share of total income earned by the top 1% of earners.
incshare_top5	Share of total income earned by the top 5% of earners.
incshare_top05	Share of total income earned by the top 0.5% of earners.
incshare_top01	Share of total income earned by the top 0.1% of earners.
gini_coef	State income inequality measured by the Gini coefficient.
x_top_corporateincometaxrate	Top corporate income tax rate.
welfare_spending_percap	Per capita direct payments to individuals through public welfare programs.
percentBlack	Percentage of the population identified as Black by the Census race variable.
percentForeignBorn	Percentage of the population identified as foreign-born by the Census.
percentWorkInManufacturing	Percentage of the population identified as working in manufacturing by the Census.

Table 5: Predictability of Shares

	Dependent Variable: Military Shares
MPC	-12.012^{***}
	(1.849)
Openness	1.296
	(0.889)
Income per capita	0.322^{*}
	(0.173)
R^2	0.478
Adjusted R^2	0.445
Observations	51

Note: Table displays coefficient estimates alongside (heteroskedasticity-robust) standard errors from regression (32). Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

D.3 Aggregate Shock Diagnostics

I report the properties of the various aggregate shocks used in the analysis. Figure 7 reports autocorrelation functions for the 1-year change in aggregate military spending employed in the baseline specification – alongside the 2-year change employed in the Nakamura and Steinsson (2014) specification as well as 1-year changes in the defense returns series from Fisher and Peters (2010) employed as robustness. I find the two-year military spending shock is significantly correlated with its one-year lag (mechanically so, given it captures two-year changes), while there is no significant evidence of serial correlation in the one-year military spending shock and excess defense returns shocks.

Figure 7: Autocorrelation Functions

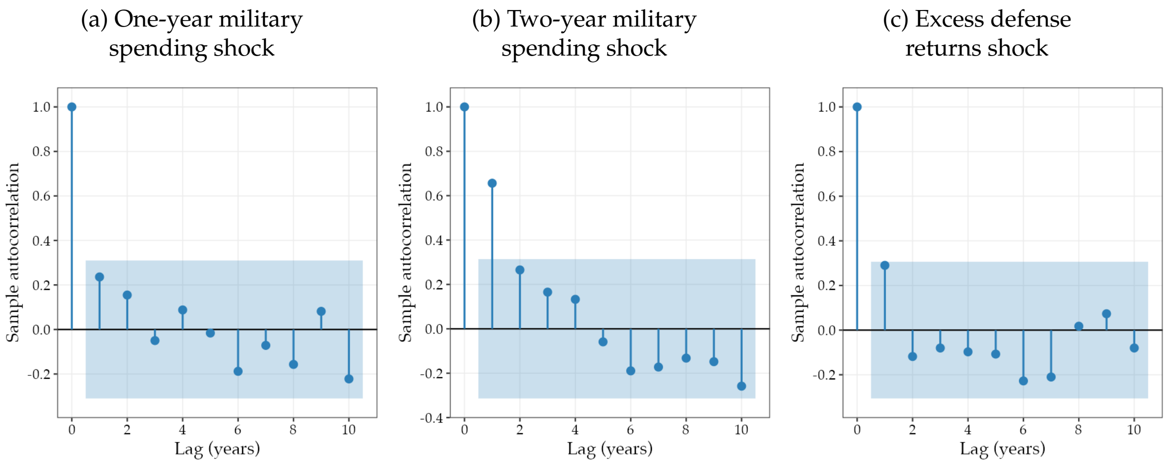

Notes: Autocorrelation function of various aggregate shock series with 95% confidence bands.

D.4 Main Specification: Controlling for Other State Characteristics

LASSO analysis I first consider the predictability of the military spending shares using a broader range of state characteristics. Specifically, I estimate:

$$s_i = \alpha + \lambda' a_i + u_i$$

where I include all the variables from Table 4, alongside the MPC and openness measures in a_i. In order to avoid over-fitting, I estimate coefficients via LASSO with varying levels of the regularisation parameter. Figure 8 displays the estimated coefficient estimates for each variable across values of the regularisation parameter. The exercise highlights the MPC variable alongside income per capita as the two best predictors of the military spending shares (with a non-zero coefficient on both variables at the regularisation parameter selected by 10-fold cross-validation). I then employ these variables as controls in the baseline shift-share regressions (34) and (33), reporting results in the final column of Table 1. I then repeat this exercise where I include in a_i various *higher-order* terms – specifically, each state characteristic squared alongside their interactions. In this case, for the regularisation parameter chosen by 10-fold cross-validation, the variables with a non-zero coefficient are: MPCs, income per capita and the interaction of income per capita with the top rate of corporate income tax. I employ these variable as controls in the baseline regression, reporting results in the final column of Table 6, discussed below.

Figure 8: Predictability of Shares From LASSO

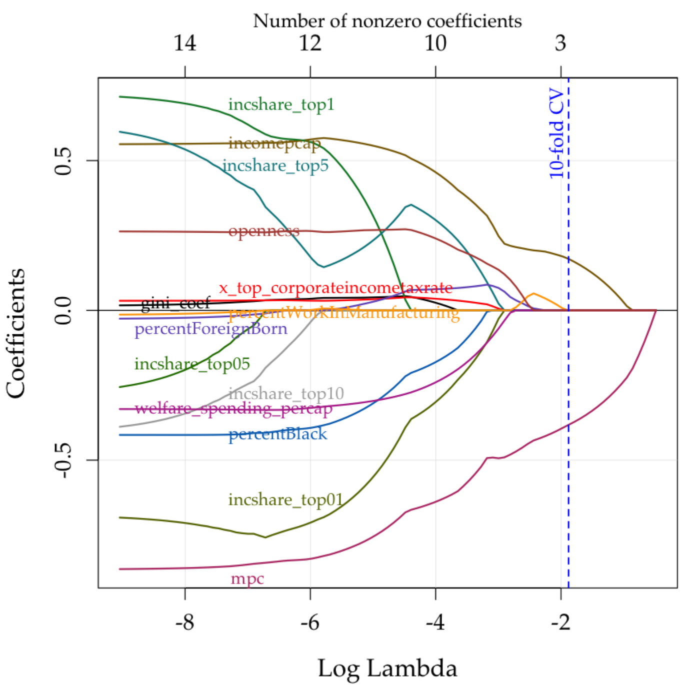

Notes: Coefficient estimates from regression (32), estimated via LASSO for varying regularisation parameters (horizontal axis). Dashed vertical line shows the regularisation parameter chosen by 10-fold cross-validation.

Non-linear controls I re-create Table 1, additionally including non-linear combinations of controls. Specifically, I additionally include all squared and interaction terms – e.g. $(mpc_i)^2$, $mpc_i \cdot open_i$ – in the vector a_i in the main regressions. As in the main specification, I include lags of output and military spending growth, alongside lags of all right-hand side variables as controls. Table 6 displays my results. I find similar results to the main analysis: multiplier estimates drop to around $0.4 - 1.0$ at the 2-year horizons, and $0.8 - 1.4$ at the 4-year horizon with the inclusion of additional higher-order terms as controls.

Table 6: Cumulative Multiplier Estimates (Higher-Order Controls)

	None	Income	MPC + Openness	LASSO
Horizon (years)	(1)	(2)	(3)	(4)
0	0.89	0.58	0.42	0.45
	(0.72)	(0.43)	(0.46)	(0.41)
	$[F = 16.84]$	$[F = 27.15]$	$[F = 19.16]$	$[F = 16.84]$
1	1.15	0.86	0.34	0.69
	(0.80)	(0.54)	(0.51)	(0.48)
	$[F = 10.90]$	$[F = 14.96]$	$[F = 11.84]$	$[F = 10.06]$
2	1.45*	0.99*	0.41	0.81*
	(0.76)	(0.53)	(0.53)	(0.43)
	$[F = 14.89]$	$[F = 19.05]$	$[F = 14.08]$	$[F = 13.72]$
3	1.96**	1.08**	0.61	0.88**
	(0.78)	(0.55)	(0.65)	(0.44)
	$[F = 14.58]$	$[F = 18.24]$	$[F = 13.41]$	$[F = 13.54]$
4	2.69***	1.38**	0.82	1.14**
	(0.84)	(0.59)	(0.75)	(0.47)
	$[F = 14.01]$	$[F = 16.28]$	$[F = 13.01]$	$[F = 12.44]$

Notes: Table displays coefficient estimates alongside (time-clustered) standard errors and F-stats from the IV regressions described in the main text. Column headers refer to additional controls included in the regression. These control variables are included in levels alongside squared and interaction terms (all interacted with the shock). Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

Coefficient stability I consider the sensitivity of the results reported in the main text to the inclusion of additional control variables. The main table demonstrates that the estimated multiplier at the 4-year horizon is substantially lower after including the MPC and openness variables as additional controls. Figure 9 demonstrates that this finding is robust to the inclusion of a variety of further additional control variables. Alongside the MPC and openness controls, I consider various specifications which also control one-at-a-time for measures of income and income inequality, demographics and differences in local state policies. After conditioning on states' MPCs and openness, the estimated multiplier remains around, or even somewhat below, 1.0 when further including various additional controls, suggesting this finding is not sensitive to further omitted sources of regional heterogeneity.

Figure 9: Sensitivity of 4-Year Multiplier Estimate to Additional Controls

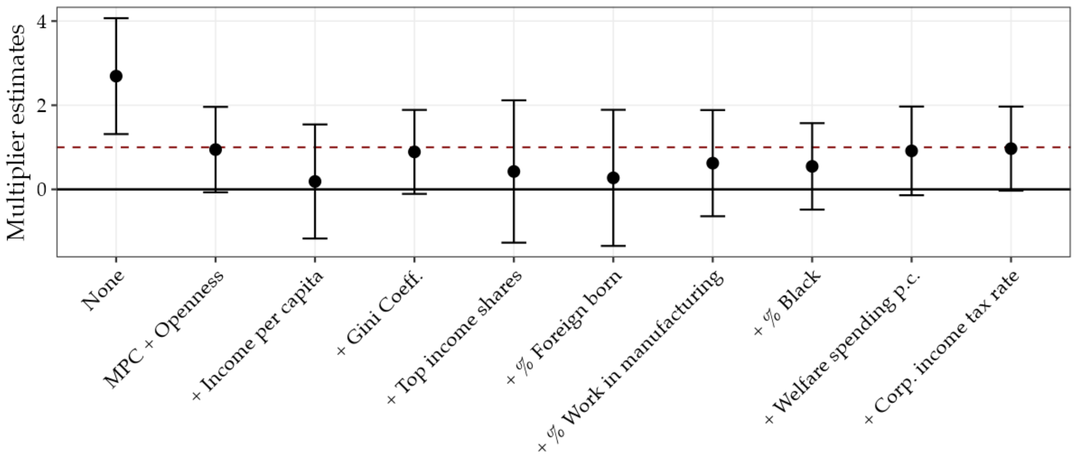

Notes: Multiplier estimates alongside time-clustered 90% confidence intervals from regressions (33) and (34), with various additional controls included in the vector a_i.

The main table also demonstrates that the estimated multiplier at the 4-year horizon is substantially lower after including control variables selected by a first-stage LASSO regression, with the regularisation parameter selected by 10-fold cross-validation. I now demonstrate that this result is robust to varying choices of the regularisation parameter. Figure 10 displays my results. When the regularisation parameter is selected by 10-fold cross-validation, the LASSO regression selects just two variables in the first-stage estimation of the targeting rule, yielding a multiplier estimate of around 1.0. Varying the regularisation parameter to include anywhere up to ten additional controls, the multiplier estimate is consistently around, or even somewhat below, 1.0.

Figure 10: Sensitivity of 4-Year Multiplier Estimate to LASSO regularisation

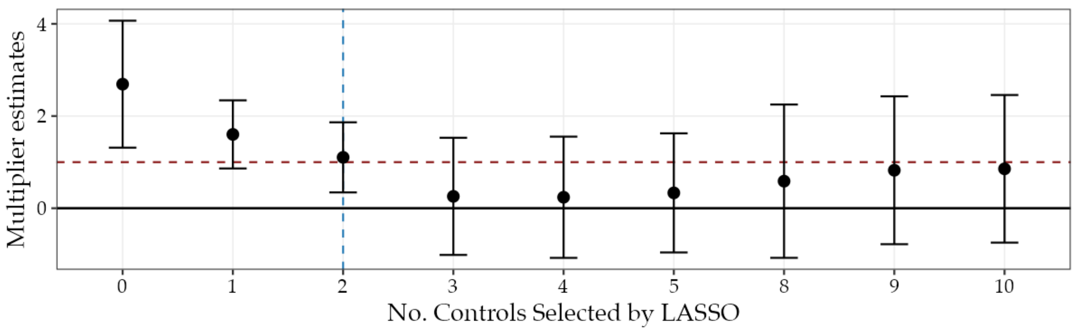

Notes: Multiplier estimates alongside time-clustered 90% confidence intervals from regressions (33) and (34), with increasing number of controls included in the vector a_i selected by a first-stage LASSO regression. Dashed lines denote the number of controls selected by the regularisation parameter at 10-fold cross-validation for the first-stage LASSO.

D.5 Main Specification: Defense Returns Shock

Table 7 displays results for the main specification with the aggregate shock measure d_t instead defined as 1-year changes in the excess defense returns series from Fisher and Peters (2010).[49]

Table 7: Cumulative Multiplier Estimates (Defense Returns Shock)

Horizon (years)	None (1)	Income (2)	MPC + Openness (3)	LASSO (4)
0	1.04	0.54	0.29	0.15
	(1.33)	(0.95)	(1.10)	(0.76)
	$[F = 4.52]$	$[F = 3.75]$	$[F = 4.22]$	$[F = 4.13]$
1	0.86	0.43	0.36	0.33
	(1.47)	(1.00)	(1.09)	(0.78)
	$[F = 4.76]$	$[F = 4.85]$	$[F = 5.77]$	$[F = 5.63]$
2	0.86	0.58	0.55	0.60
	(1.48)	(1.16)	(1.06)	(0.91)
	$[F = 5.49]$	$[F = 4.79]$	$[F = 7.16]$	$[F = 6.34]$
3	1.28	0.85	0.88	0.86
	(1.55)	(1.25)	(1.08)	(1.01)
	$[F = 6.39]$	$[F = 5.66]$	$[F = 7.92]$	$[F = 7.20]$
4	1.35	0.64	1.05	0.64
	(1.38)	(1.07)	(1.02)	(0.89)
	$[F = 8.83]$	$[F = 7.97]$	$[F = 9.31]$	$[F = 8.61]$

Notes: Table displays coefficient estimates alongside (time-clustered) standard errors and F-stats from the IV regressions described in the main text, using the 1-year change in the excess defense returns series from Fisher and Peters (2010) as the aggregate shock d_t. Column headers refer to additional interaction controls included in the regression. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

D.6 Replicating Nakamura and Steinsson (2014)

My main specification deviates from the original specification in Nakamura and Steinsson (2014) along various dimensions. For full comparability, I now also replicate exactly the original specification in Nakamura and Steinsson (2014). Specifically I estimate:

$$\frac{y_{i,t} - y_{i,t-2}}{y_{i,t-2}} = \alpha_{y,i} + \delta_{y,t} + \beta_y^{ss} \, (s_i \cdot d_t) + \lambda_y'(a_i \cdot d_t) + u_{y,i,t} \tag{116}$$

$$\frac{g_{i,t} - g_{i,t-2}}{y_{i,t-2}} = \alpha_{g,i} + \delta_{g,t} + \beta_g^{ss} \, (s_i \cdot d_t) + \lambda_g'(a_i \cdot d_t) + u_{g,i,t} \tag{117}$$

where $d_t = \frac{g_t - g_{t-2}}{y_{t-2}}$. Relative to the baseline specification, growth rates are now all expressed relative to period-$(t-2)$, including in the construction of d_t, and lagged control variables are excluded. Table 8 displays my results for different control variables a_i.

[49] I take the series from the replication package of Ramey (2016) (converted to annual frequency by averaging quarterly values).

Table 8: Original Shift-Share Specification of Nakamura and Steinsson (2014)

	Dependent Variable: Output			
	(1)	(2)	(3)	(4)
Multiplier	2.48**	1.48**	1.25*	0.98**
	(1.01)	(0.71)	(0.68)	(0.47)
Additional Controls	None	inc_i	mpc_i, $open_i$	LASSO
Observations	1989	1989	1989	1989

Notes: Table displays coefficient estimates alongside (time-clustered) standard errors. Top panel shows results from the IV regression defined by (116) and (117). Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

D.7 Estimating State Weights

I estimate the weights on each state underlying the multiplier estimates. To do so, I build on the setting from Appendix C.3, treating the military spending shares s_i as proxies for the sensitivity of state-level military spending to shocks to aggregate spending, given by $z_i\boldsymbol{\Theta}^{\boldsymbol{gs}}$, and treating the shifter as the government spending shock of interest $d_t = \varepsilon_t^s$.

Method Consider the shift-share regressions without additional interaction controls, as in the baseline regression in Nakamura and Steinsson (2014):

$$y_{i,t+h} = \alpha_{i,y,h} + \delta_{t,y,h} + \beta^{ss}_{y,h}(s_i \cdot \varepsilon_t^s) + e_{i,t,h,y}$$
$$g_{i,t+h} = \alpha_{i,g,h} + \delta_{t,g,h} + \beta^{ss}_{g,h}(s_i \cdot \varepsilon_t^s) + e_{i,t,h,g}$$

Following the same steps in Appendix C.3, we then find:

$$\boldsymbol{\beta}^{ss}_y = \frac{\mathbb{E}[\mathcal{M}_i^{yg}\boldsymbol{\beta}_g^{ss}\tilde{s}_i z_i]}{\mathbb{E}[\tilde{s}_i z_i]} + \text{Bias from GE term} \tag{118}$$

where $\tilde{s}_i$ denotes the de-meaned military spending shares and $\mathcal{M}_i^{yg}\boldsymbol{\beta}_g^{ss}$ denotes the direct effect of interest. Putting aside the bias term, β^{ss}_y can thus be written as a weighted average of the direct effects of interest with weights ω_i that satisfy:

$$\omega_i = \tilde{s}_i \cdot z_i \tag{119}$$

Finally, note that z_i can be recovered from the data through a series of regressions, one for each unit i, of the form:

$$g_{i,t+h} = \alpha_{i,g,h} + \beta^{ols}_{i,g,h}\varepsilon_t^s + e_{i,t,g,h} \tag{120}$$

It is immediate by the micro-SVMA structure (alongside the assumption of scaled-up treatments) that the OLS coefficients satisfy: $\beta^{ols}_{i,g,h} = \Theta^{gs}_{i,h} = z_i\Theta^{gs}_h$. And so the weights can be estimated by simply de-meaning the military spending shares s_i and multiplying by the OLS coefficient $\beta^{ols}_{i,g}$. Note these weights will generally take on negative values: even when spending increases in all states in response to higher aggregate spending (so that $\forall i : \beta^{ols}_{i,g} > 0$), the shift-share coefficient negatively weights multipliers for states with below average spending shares (i.e. states for whom $\tilde{s}_i < 0$).

Instead, consider the extended shift-share regression with additional controls:

$$y_{i,t+h} = \alpha_{i,y,h} + \delta_{t,y,h} + \beta^{ss}_{y,h}(s_i \cdot \varepsilon^s_t) + \lambda'_{y,h}(a_i \cdot \varepsilon^s_t) + e_{i,t,h,y}$$
$$g_{i,t+h} = \alpha_{i,g,h} + \delta_{t,g,h} + \beta^{ss}_{g,h}(s_i \cdot \varepsilon^s_t) + \lambda'_{g,h}(a_i \cdot \varepsilon^s_t) + e_{i,t,h,g}.$$

As discussed in Appendix C.3, we have:

$$\boldsymbol{\beta}^{ss}_y = \frac{\mathbb{E}[\mathcal{M}^{yg}_i \boldsymbol{\beta}^{ss}_g \omega_i]}{\mathbb{E}[\omega_i]},$$

with weights that satisfy:

$$\omega_i = \mathbb{E}\left[\left(s_i^{\perp a_i}\right) \cdot z_i \mid x_i, a_i\right], \tag{121}$$

where $s_i^{\perp a_i}$ is the residual from a regression of s_i on a_i and a constant. Note that, by LIE, $\boldsymbol{\beta}^{ss}_y$ can also be written in the same form with weights instead given by:

$$\omega_i = \left(s_i^{\perp a_i}\right) \cdot z_i, \tag{122}$$

which can again be directly recovered from the data (following the same steps as above). Following Borusyak and Hull (2024), I refer to ω_i in (122) as the *ex-post* weights, in contrast to the *ex-ante* weights from (121). Note that the monotonicity assumption discussed in Appendix C.3 refers to the *ex-ante* weights in (121). If one is interested in testing this assumption, the ex-ante weights are therefore the appropriate object to estimate – although doing so requires assumptions on how the expectation of the product, $(s_i^{\perp a_i})\beta^{ols}_{i,g}$, depends on observables. The ex-post weights on the other hand can be directly recovered from the data absent any additional assumptions – and can be easily compared to the weights without controls (119) – thus isolating the potential bias stemming from the GE term alone in (118).

Results I follow the above method to estimate the weights underpinning the 4-year multiplier estimates from Table 1. In practice, I estimate $\beta^{ols}_{i,g,h}$ from (120) at $h = 4$, including the same set of lagged controls as in the baseline (lagged output growth, lagged military spending growth and one lag of the shock). Figure 11 displays estimates of the ex-post weights. The blue dots show the estimated weight for each state for the baseline specification without additional controls from (119) plotted against income per capita. The weights are heavily skewed – the top three states account for over 80% of the weight with Connecticut alone receiving a weight of over 50%. There is also a significant mass of negative weights: the negative weights sum to around -35%. The weights are also positively correlated with income, with the three richest states (Connecticut, District of Columbia and Alaska) all receiving large weights. The green, yellow and red dots show the estimated ex-post weights from (122) for each state for the extended specifications from Table 1. These weights are very similar to those in the specification

Figure 11: Estimated *Ex-Post* State Weights in 4-Year Multiplier Estimates

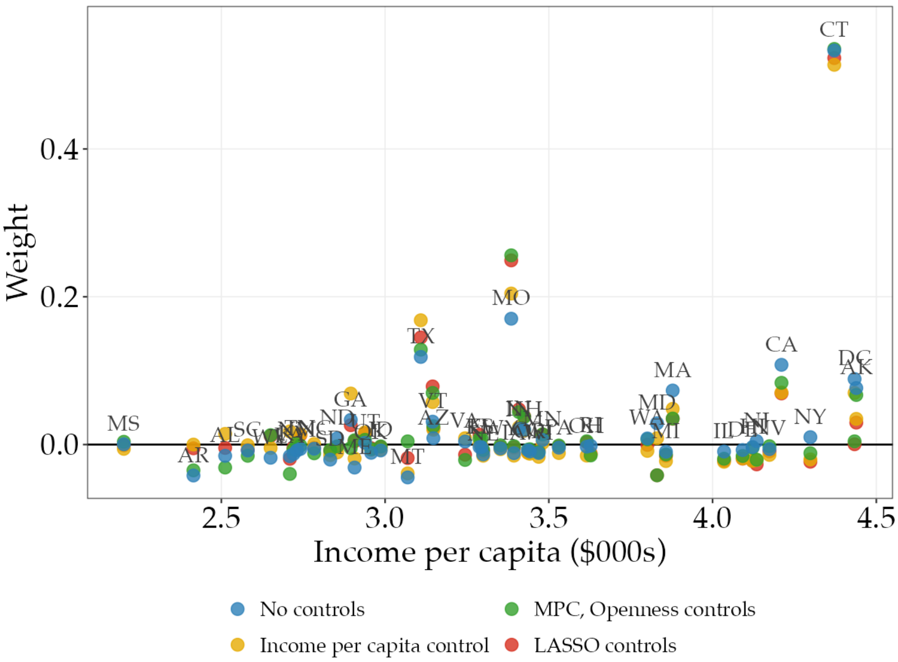

Figure 12: Estimated *Ex-Ante* State Weights in 4-Year Multiplier Estimates

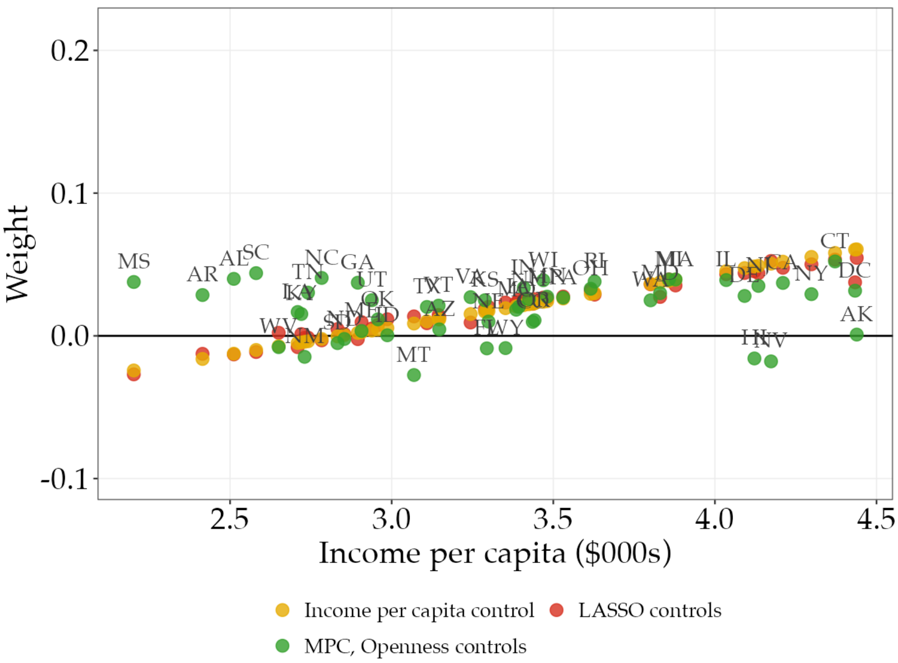

without controls – with a correlation of around 0.95 in each case with the original weights. This suggests any difference in coefficient estimates, relative to the no controls case, can be principally assigned to the change in the bias term from (118), rather than changes in the underlying weights ω_i.

Figure 12 then displays estimates of the ex-ante weights from (121). To do so, I assume the conditional expectation from (121) is linear in the included controls. Somewhat reassuringly, in this case, the weights are much flatter, with less negative mass (around -10%).

D.8 Dynamic First-Stage and Reduced-Form

The main analysis reports dynamic multiplier estimates – i.e. the ratio of $\beta^{ss}_{y,h}$ to $\beta^{ss}_{g,h}$ from (33) and (34). Below I plot separately the first-stage and reduced-form coefficients. For comparability, I normalise all responses to imply a one-unit increase in government spending at peak. Across specifications, the shock triggers a gradual and persistent rise in local government spending. The local output responses to this increase in local spending are then significantly lower in the specifications with additional controls.

Figure 13: Dynamic Responses to Aggregate Military Spending Shock

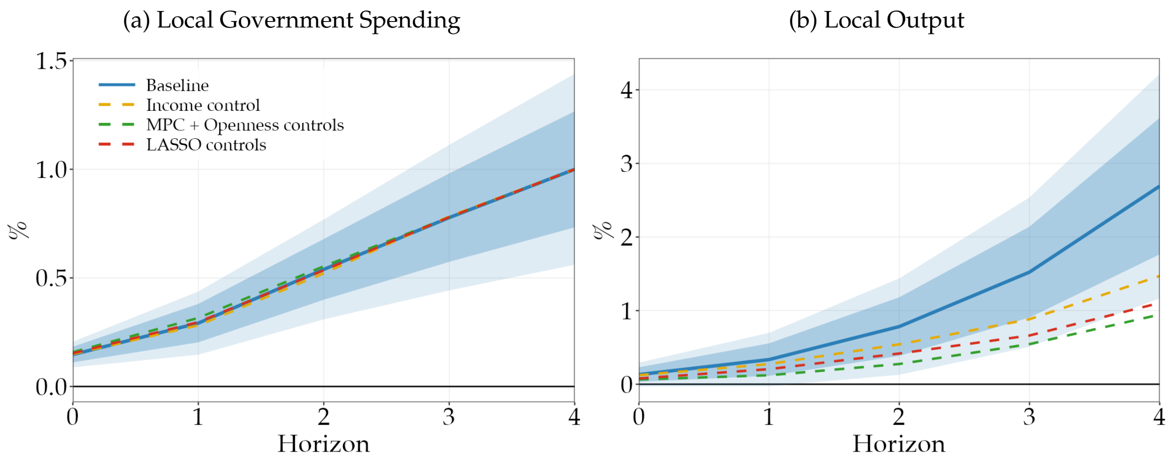

Notes: Blue lines denote estimates of $\beta^{ss}_{y,h}$ and $\beta^{ss}_{g,h}$ from (33) and (34) with no characteristic-interaction controls, with 68% and 90% confidence bands. Dashed lines add the characteristic $\times$ shock interactions for income per capita (yellow), MPC and openness (green), and the LASSO-selected set (red). Each specification is rescaled so that its spending response at $h = 4$ equals 1%.

D.9 Transmission Channels

To investigate the various transmission channels of the shock, I estimate local projections for aggregate variables over the period 1966-2006 of the form:

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h d_t + \lambda'_h x_{t-1} + u_{t+h} \tag{123}$$

where d_t denotes 1-year changes in national military spending relative to national output in year t (as in the main analysis) and x_{t-1} denotes a vector of lagged controls. I include one lag of the change in the outcome variable, $y_{t-1} - y_{t-2}$, alongside one lag of the shock, d_{t-1}, as controls. Figure 14 displays my results. Government military spending rises significantly and persistently following the shock, by around 8% after four years. In response, (real) rates increase, by over 1 p.p. in the first couple of years. Goods consumption also rises, while transfers fall – with oil prices and average tax rates broadly flat.

Figure 14: Response of Aggregate Variables

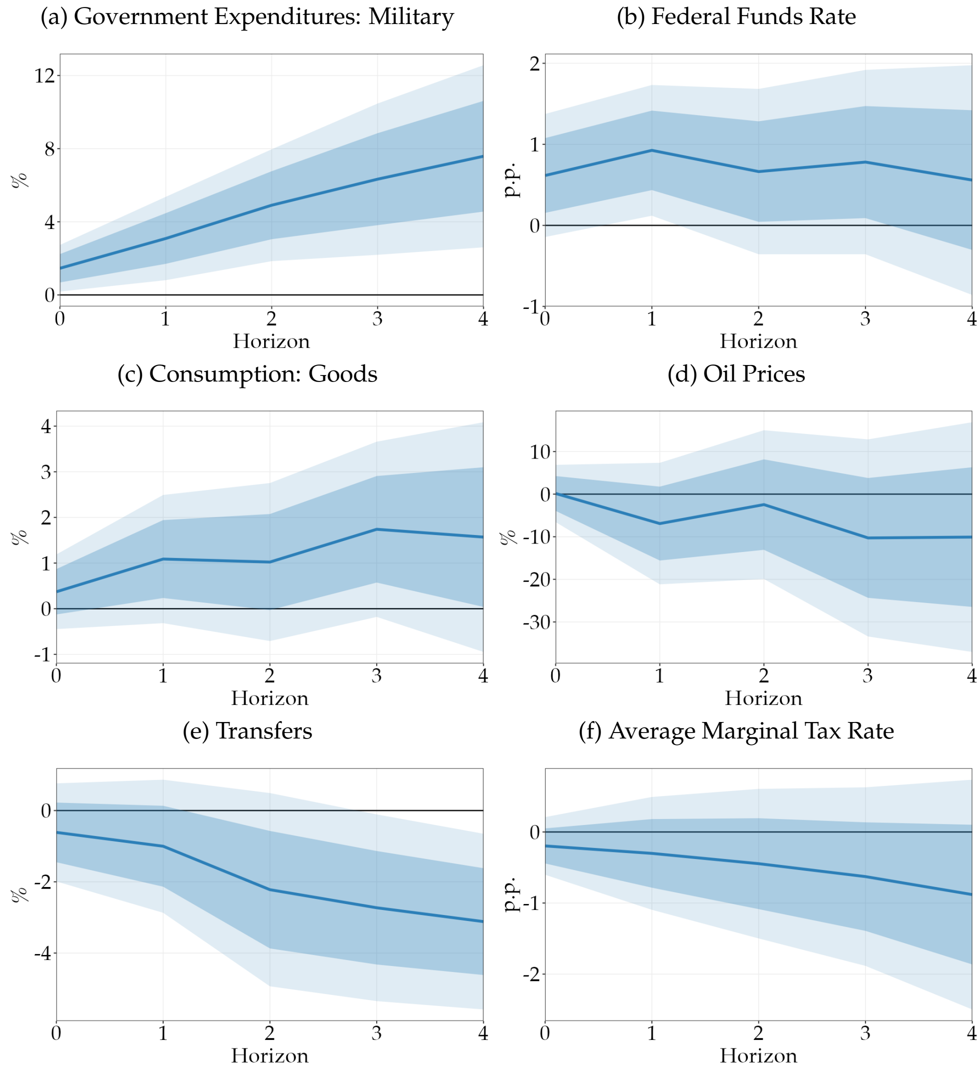

Notes: Response of aggregate variables following a shock to national military spending, estimated via (123). Shaded areas denote 68% and 90% heteroskedastic-robust confidence intervals.

E Additional Empirical Results: Guren et al. (2021b)

This Appendix collects supplementary information and additional results from the empirical application in Section 5.2.

E.1 Data

The main dataset comes directly from the replication materials from Guren et al. (2021b), from which I use the quarterly CBSA panel of house prices, retail employment per capita, population, Census region, and the GDP deflator used throughout to construct real series.[50] The Saiz housing supply elasticity (Saiz, 2010) is taken from the same panel. I follow the authors' sample rule (CBSAs with at least 100 quarters of retail employment data from 1978 onwards) and estimate on the 1990–2017 period.

As aggregate shifter, I use the S&P/Case-Shiller national home price index, taken from FRED (`CSUSHPISA`), averaged from monthly to quarterly frequency, deflated by the GDP deflator and expressed as a one-quarter log change. I also use the Gilchrist–Zakrajšek credit spread (Gilchrist and Zakrajšek, 2012), which I take directly from the replication package of Guren et al. (2021b).

The CBSA characteristics that enter the interaction controls are listed in Table 9. All are county-level series aggregated to CBSAs using the 2013 OMB county-to-CBSA delineation shipped with the replication package. MPCs and openness come from Bellifemine et al. (2023) and are constructed exactly as for the state-level analysis: county estimates are aggregated with population weights (from the Census Bureau's Population Estimates), openness is one minus the share of employment in the non-tradable sector, and steady-state values are proxied by pre-Covid averages over the periods in which each series is available (2011–2019 for MPCs, 1986–2019 for openness). Per capita personal income comes from the BEA's county economic profiles (table CAINC30), summed to CBSA level and deflated by the GDP deflator. The remaining demographic series come from the Census Bureau's *USA Counties* compendium, aggregated with 1990 county population weights. The construction employment share is the SIC 15 share of employment in the Guren et al. (2021b) panel. All characteristics are measured in 1990, the first year of the estimation sample.

E.2 Balance Tests for Saiz Instrument

Table 10 reports balance tests for the Saiz housing supply elasticity measure employed in the main text. I regress the measure on a number of CBSA characteristics, one at a time, reporting estimated coefficients and p-values.

[50] The panel contains SIC-based retail employment through 2000 and NAICS-based employment from 1990 onwards; I chain the NAICS-based series to the SIC-based series from 1990.

Table 9: CBSA Characteristics

Variable	Definition	Source
MPC	Marginal propensity to consume, 2011–2019 average.	Bellifemine et al. (2023)
Openness	One minus the share of employment in the non-tradable sector, 1986–2019 average.	Bellifemine et al. (2023)
Real per capita income	Personal income from all sources divided by population, deflated by the GDP deflator.	BEA CAINC30
Share foreign born	Percentage of the population born outside the United States.	Census *USA Counties*
Share Black	Percentage of the population identified as Black.	Census *USA Counties*
Homeownership rate	Percentage of occupied housing units that are owner-occupied.	Census *USA Counties*
Two-digit industry shares	Shares of employment in the eight two-digit SIC industries (mining; construction; manufacturing; transport and utilities; wholesale; retail; finance, insurance and real estate; services), annual, interpolated to quarterly.	CBP, via Guren et al. (2021b)
Unemployment	Unemployment rate.	BLS LAUS, via Guren et al. (2021b)

Table 10: Balance of the Saiz Instrument Across CBSA Characteristics

	Coefficient	N
Openness	0.072	270
	(0.057)	
MPC	0.063	270
	(0.057)	
Share foreign born, 1990	-0.238^{***}	270
	(0.054)	
Share Black, 1990	0.072	270
	(0.100)	
Homeownership rate, 1990	0.025	270
	(0.061)	
Real per capita income, 1990–2017	-0.311^{***}	270
	(0.061)	
Unemployment, 1990–2017	0.031	269
	(0.049)	
Construction employment share, 1990–2017	-0.072	270
	(0.071)	

Notes: Each row is a separate regression of the Saiz (2010) housing supply elasticity on one CBSA characteristic, both standardised over the sample, so entries are the SD change in the elasticity per SD of the characteristic, with Census-region fixed effects. Heteroskedasticity-robust standard errors in parentheses. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

E.3 Aggregate Shock Diagnostics

Figure 15 reports the autocorrelation function for each of the aggregate shocks orthogonalised with respect to its first four lags.

Figure 15: Autocorrelation Functions

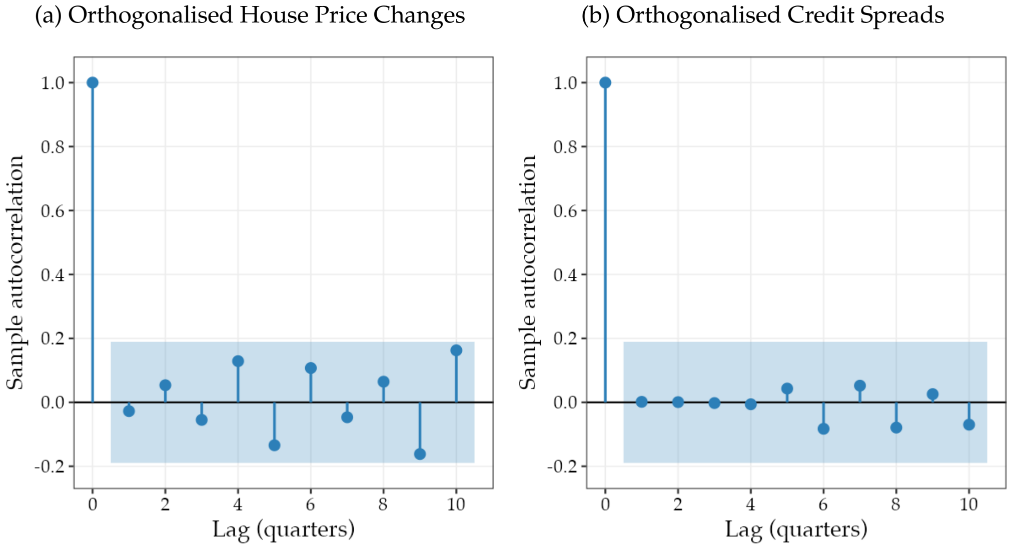

Notes: Autocorrelation function of various aggregate shock series with 95% confidence bands.

E.4 Details on sensitivity analysis

Alongside the main specification, I report results from three additional sensitivity analyses. I describe these in more detail below.

No CBSA Characteristics. The same specification as the baseline, with the vector of CBSA characteristics $a_{i,r}$ empty.

More CBSA Characteristics. In this specification, the set of characteristics $a_{i,r}$ is extended to include the baseline pair plus the other variables from Table 10: the MPC and openness measures, the share of the population that is Black, the homeownership rate, and the averages of the unemployment rate and of the construction share of employment. Each again enters with the contemporaneous shifter and its lags.

+ All Industry Shares. I replace the average construction share of the previous specification with the full set of eight two-digit SIC employment shares from Guren et al. (2021b), of which construction is one, keeping the other seven CBSA characteristics. These are time-varying, so the interaction uses the share in quarter t, $\mu^j_{i,t} \cdot d_t$, together with lags.

E.5 Replicating Guren et al. (2021b)

To ease comparability with the original study, I replicate exactly the results in Guren et al. (2021b), comparing them to the estimates in the main text.

Table 11: Original Specification of Guren et al. (2021b)

	Elasticity
Full controls	0.141***
	(0.038)
No prediction controls	0.163***
	(0.036)
No prediction or industry-share controls	0.163***
	(0.023)
As above, Case-Shiller shifter	0.164***
	(0.025)
Observations	29867
CBSAs	270

Notes: All rows include quarter fixed effects and full-sample CBSA demeaning. The first row adds the two-digit industry shares with quarter-specific coefficients and the leave-out prediction controls for differential CBSA exposure to regional retail employment, the real mortgage rate and the excess bond premium; the second row drops the three prediction controls; the third also drops the industry-share controls; the fourth keeps the third row's controls but replaces the national house price change in the instrument with the one-quarter log change in the real Case-Shiller national index. Standard errors two-way clustered by CBSA and quarter in parentheses. *p<0.1; **p<0.05; ***p<0.01.

The original study runs a variant of regressions (35) and (36). The authors do not include any lagged controls or interactions of the shock with CBSA characteristics, but instead control for industry-share-by-time fixed effects alongside various measures of CBSA-sensitivity to business cycle variables. The original specification also uses annual changes in a leave-one-out average of house prices as the aggregate shock. Finally, the authors report the ratio $\beta^{ss}_{y,0}/\beta^{ss}_{p,0}$ as their estimate of housing wealth effects.

The first row from Table 11 replicates exactly the reported elasticity from the Saiz-instrument from Guren et al. (2021b). Removing the additional controls from Guren et al. (2021b) slightly increases the estimated elasticity – see rows two and three. Switching the aggregate shifter to quarterly changes in the Case-Shiller aggregate house index leaves the elasticity essentially unchanged – see row four. I find similar-sized elasticities in the dynamic specifications reported in the text: e.g. around 0.17-0.19 from $h = 3$ onwards for the comparable "No-CBSA characteristics" specification.

F Estimating Fiscal Multipliers in a Network Model

This Appendix describes a simple quantitative model to illustrate the results from Appendix C.7. I begin by setting out the model and then use it to help revisit estimates of the ARRA package from Chodorow-Reich (2019).[51]

F.1 Model

Setting Each region $i = 1, ..., n$ is endowed with q_i units of a distinct good. All agents are able to save via a bond available at price p_s and in fixed net supply: $\bar{s}$. There is a representative agent in each region with Cobb-Douglas preferences:

$$u_i(c_{1i}, ...c_{ni}, s_i) = \alpha_{1i}\log(c_{1i}) + ... + \alpha_{ni}\log(c_{ni}) + \alpha_{si}\log(s_i)$$

where c_{ji} denotes the consumption of region i on the good endowed to region j and s_i denotes savings of region i. I normalise the preference weights on goods such that $\sum_{j=1}^{n}\alpha_{ji} = 1$, where α_{ji} is region i's goods-expenditure share on the good from region j. The government spends g_i on each region's good and raises taxes τ_i in each region subject to the budget constraint:

$$\sum_i g_i = \sum_i \tau_i$$

I assume taxes are raised evenly across states so that $\forall i : \tau_i = \tau$, and that government spending in each region can be decomposed into a "steady-state" term, $\bar{g}_i$, as well as a shock term, ε_i^g, as: $g_i = \bar{g}_i + \varepsilon_i^g$.

Competitive Equilibrium Define (nominal) income for each region as: $y_i = p_i q_i$. Consumer demand, budget constraints and market clearing imply that for each region i:

$$y_i = \sum_{j=1}^{n}[\alpha_{ij}\underbrace{\frac{1}{1+\alpha_{sj}}}_{\equiv mpc_j}(y_j - \tau_j)] + g_i$$

And so the relationship between income, government spending, and taxes in all regions can be summarised as:

$$\underbrace{y}_{n\times 1} = (I - A)^{-1}\cdot\underbrace{g}_{n\times 1} - (I - A)^{-1}A\cdot\underbrace{\tau}_{n\times 1}$$

where $A_{ij} = \alpha_{ij}\cdot mpc_j$. Note the effects of government spending in this simple model exhibit a Keynesian-cross-style logic: higher government spending in region 1 directly increases incomes in region 1, which increases their spending on all goods, thereby increasing incomes in

[51] Although the setup is somewhat different, the model considered here is essentially equivalent to the "Old Keynesian" model discussed in Chodorow-Reich (2020). As in Chodorow-Reich (2020), I use the model to discuss issues related to identification of regional multipliers. Crucially, I extend the model to allow for both heterogeneity in MPCs and trade linkages, highlighting how heterogeneous spillovers generate bias in typical regional multiplier regressions.

other regions, which further increases spending on all goods (and so on). The effect of higher government spending in region j on region i therefore depends on the direct demand of region j for goods in region i (governed by α_{ij}), as well as all indirect linkages (e.g. $\alpha_{ik} \cdot \alpha_{kj}$ etc.) – where the sum of all these linkages is captured by the appropriate entry in the Leontief-inverse multiplier matrix $(I - A)^{-1}$. Note that given taxes are constrained to be equal across regions, the model can be written:

$$y = \bar{y} + \mathcal{M}^g \varepsilon^g$$

where $\bar{y} = \mathcal{M}^g \bar{g}$ denotes steady state outcomes and $\mathcal{M}^g = (I - A)^{-1} - (1/n)(I - A)^{-1} A J_n$. I assume the econometrician has access to repeated observations $t = 1, 2, ...$ from this static model, and further assume that *observed* incomes $\widetilde{y}_t$ are subject to some idiosyncratic (stochastic) measurement error in each state which I denote $u_t = [u_{1,t}, ..., u_{n,t}] \overset{i.i.d.}{\sim} N(0, \sigma_u^2 I_n)$, such that:

$$\widetilde{y}_t = \bar{y} + \mathcal{M}^g \varepsilon_t^g + u_t$$

Simple Calibration I consider a simple parametrization of the model. There are $N = 50$ regions, located at equally spaced points on a circle. I assume that each region allocates 60% of its goods expenditure to its own good and has an MPC of 50%:

$$\forall i : \alpha_{ii} = 0.6, \qquad mpc_i \equiv 1/(1 + \alpha_{si}) = 0.5.$$

For $i \neq j$, region j's expenditure shares on other regions' goods are given by:

$$\alpha_{ij} = (1 - \alpha_{jj}) \frac{w_{ij}}{\sum_{k \neq j} w_{kj}},$$

where w_{ij} is a function of the distance between regions i and j, d_{ij}. Specifically, I set

$$w_{ij} = [d_{ij}]^{-5}, \;\; d_{ij} = \sqrt{2(1 - \cos(\theta_i - \theta_j))}$$

with locations

$$\theta_i = \frac{2\pi(i-1)}{N}.$$

This particular calibration implies that the expenditure-share matrix is circulant. Hence the resulting multiplier matrix $\mathcal{M}^g$ also satisfies rotational symmetry, so all own-region government-spending multipliers are identical, and it is straightforward to verify that for this calibration, $\forall i : [\mathcal{M}^g]_{i,i} \approx 1.4$ (broadly consistent with local fiscal multipliers estimated in the literature). I assume steady-state government spending is zero, $\bar{g} = [0, ..., 0]$, and consider a simple assignment rule for spending shocks which targets spending towards regions implicitly "located" in one-half of the unit-circle – i.e.:

$$\varepsilon_{i,t}^g = \zeta_t \cdot N_i + \eta_{i,t}$$

where $\zeta_t \overset{i.i.d.}{\sim} N(0, \sigma_\zeta^2)$, $\eta_t = [\eta_{1,t}, ..., \eta_{n,t}] \overset{i.i.d.}{\sim} N(0, \sigma_\eta^2 I_n)$, and N_i denotes whether a unit is implicitly located in the top-half of the unit-circle: $N_i = I(i \in \{1, 2, ..., 25\})$. Finally, I set

$\sigma_\eta^2 = \sigma_u^2 = 1$, and consider results with various values for σ_ζ^2.

Identification I consider identification using a standard two-way fixed effects (TWFE) specification as well as an additional specification that also controls for time-fixed effects interacted with the regional dummy N_i. Table 12 shows the large-sample bias of these two different estimators with $T = 5000$ for various values of σ_ζ^2. As σ_ζ^2 increases, more of the variation in $\varepsilon_{i,t}^g$ is driven by the clustering component, increasing the bias of the TWFE estimator, which reaches as high as around 20%. In contrast, the additional specification performs well across all simulations, with bias close to zero.

Table 12: Bias of different estimators

Estimator	$\sigma_\zeta^2 = 1$	$\sigma_\zeta^2 = 2$	$\sigma_\zeta^2 = 5$	$\sigma_\zeta^2 = 10$
Two-way FE	7.02%	11.36%	18.57%	23.70%
Two-way FE + $N_i \times$ time FE	-0.03%	-0.03%	-0.03%	-0.03%

Figure 16 provides some visual intuition for the difference in performance of the estimators. The left-hand side chart shows the distribution of shocks ε_t^g around the unit-circle for a particular draw of (ζ_t, η_t) that generates particularly large clustering. In this case, the shocks are heavily concentrated in the top half of the circle (with dark red depicting a larger shock for that unit). This generates bias since the states with larger spending shocks simultaneously experience a greater rise in external demand given spending shocks in neighbouring states. The right-hand side chart shows the distribution of shocks after residualising with respect to N_i, which reduces the clustering (and thus reduces the bias).

Figure 16: Geographical Clustering of Shocks

(a) Spending Shocks (b) Residualised Spending Shocks

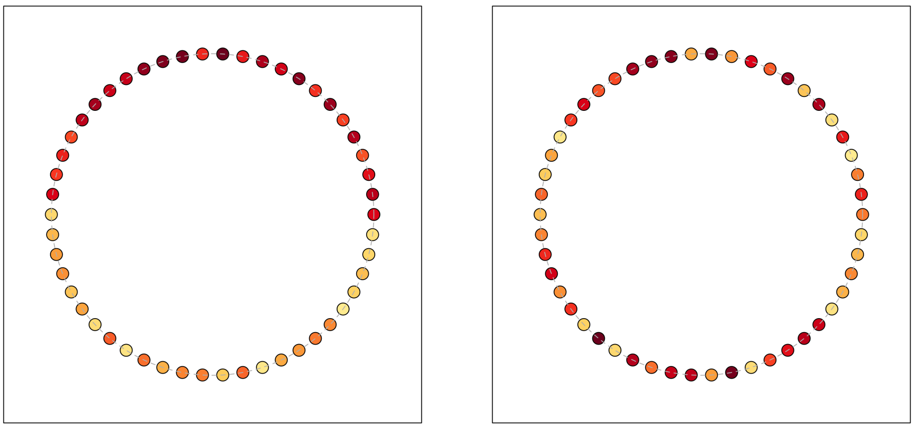

F.2 Estimating Effects of ARRA

I next replicate and extend the analysis in Chodorow-Reich (2019), which itself draws on three studies estimating local employment multipliers from the 2009 American Recovery and Reinvestment Act (ARRA).

Setting Chodorow-Reich (2019) estimates local employment multipliers via pairs of regressions of the form:

$$y_i = \alpha_y + \beta_y \varepsilon_i^g + \delta_y' x_i + e_{y,i} \tag{124}$$

$$g_i = \alpha_g + \beta_g \varepsilon_i^g + \delta_g' x_i + e_{g,i} \tag{125}$$

where y_i is the growth in annualised employment (normalised by the adult population) in state i between 2008:12 and 2010:12, g_i is the total ARRA outlays between 2008:12 and 2010:12, x_i collects lagged state-employment (in changes and levels) and lagged growth in Gross State Product as controls for pre-treatment economic conditions, and ε_i^g is an instrument that captures exogenous shocks to ARRA spending. Local multipliers are then computed by taking the ratios: $\beta_m = \beta_y/\beta_g$. The instruments ε_i^g come from previous studies, exploiting ARRA spending across states that was determined by various pre-recession formulas – from components of Medicaid spending as determined by the Federal Medical Assistance Percentages (FMAP) (Chodorow-Reich et al., 2012), Department of Transport (DOT) spending (Wilson, 2012) as well as other elements of the package (Dupor and Mehkari, 2016).

Identification Since all instruments are constructed using pre-recession formulae, it is plausible that they are uncorrelated with any other shocks that drove states' outcomes over the same period. Nevertheless, as in the model from the previous section, bias can still arise if government spending shocks are geographically clustered and spillovers between states depend on geographical ties. Figure 17 displays heatmaps for each measure, all of which display significant regional clustering.

Figure 17: Heatmaps of Instruments Across Studies

(a) Chodorow-Reich et al. (2012) (b) Wilson (2012) (c) Dupor and Mehkari (2016)

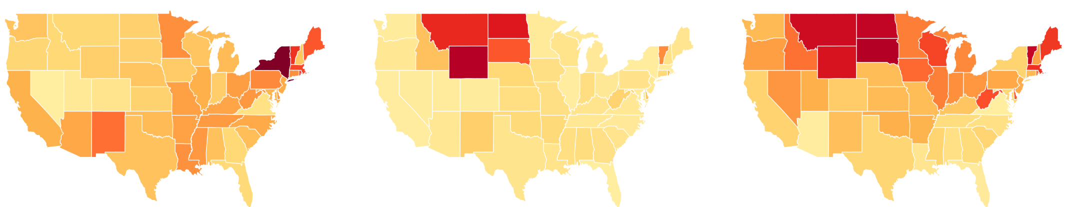

This is a particular concern given various evidence of sizable local spillovers from ARRA that decline with geographical distance (Dupor and McCrory, 2018; Auerbach et al., 2020).

Following the logic of the previous section, one approach to addressing this issue is to directly control for the implicit assignment mechanism that generated the clustering of shocks among states with closer geographical ties. To this end, I first estimate regressions of the form:

$$\varepsilon_i^g = \alpha + \lambda' a_i + \theta' x_i + \eta_i \tag{126}$$

where I include in a_i additional variables that plausibly generate the spatial clustering of shocks. The original studies provide useful information on this point. For example, the instrument from Chodorow-Reich et al. (2012) is based on pre-recession generosity of Medicaid spending, where the authors note this tended to be larger in more liberal coastal and Midwestern states – and so I include regional fixed effects to capture this. The instrument from Wilson (2012) captures differences in department of transport spending – where the original study notes that this in part depends on geographic features of states that in turn determine population density and lane-miles per capita – and so I include controls for population density (and population density squared). Finally, since the instrument from Dupor and Mehkari (2016) collects together various similar sources of variation across different elements of the ARRA package, I include both regional fixed effects and measures of population density.[52] Table 13 shows my results, where I find that these variables are effective at capturing variation in the shocks ε_i^g – coefficients on my selected variables tend to be highly significant and (adjusted) R-squareds are around 0.4-0.6. I present plots of these residualised shocks in Figure 18.

Figure 18: Heatmaps of Residualised Instruments Across Studies

(a) Chodorow-Reich et al. (2012) (b) Wilson (2012) (c) Dupor and Mehkari (2016)

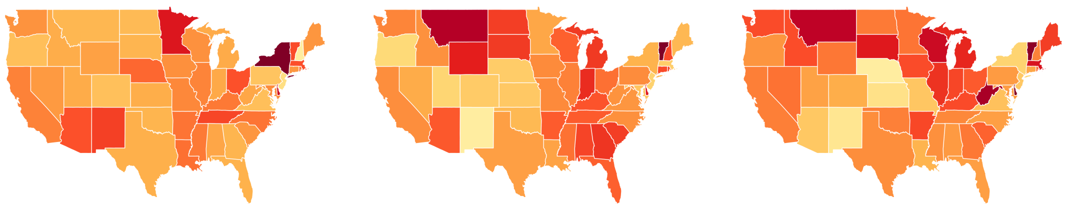

Estimating Fiscal Multipliers Finally, I estimate regressions of the form:

$$y_i = \alpha_y + \beta_y \varepsilon_i^g + \gamma_y' a_i + \delta_y' x_i + e_i \tag{127}$$

$$g_i = \alpha_g + \beta_g \varepsilon_i^g + \gamma_g' a_i + \delta_g' x_i + u_i \tag{128}$$

with multipliers defined as the ratio: β_y/β_g. Table 14 displays my results. The first panel

[52]Dupor and Mehkari (2016) note that population density directly determined variation in their instrument: "the Capital Transit Assistance program allocated roughly $6 billion to fund public transit capital improvements to urbanized areas (UZAs). The apportionment for medium sized UZAs was determined by **population density** and population, and the apportionment for large UZAs (populations greater than 200,000) was determined by factors such as bus revenue vehicle miles, bus passenger miles, fixed guideway route (such as rail) miles and **population density**." Other criteria for the allocation of funds for their instrument include "flood risk" and "the presence of inland and coastal navigation" which motivates the use of additional regional fixed effects.

Table 13: Predictability of ARRA Instruments

	Dependent Variable: ARRA Instrument		
	(1)	(2)	(3)
Region 2	-5.12***	-	-0.04
	(1.62)	-	(0.12)
Region 3	-6.81***	-	-0.26**
	(1.80)	-	(0.13)
Region 4	-7.37***	-	-0.23*
	(1.63)	-	(0.13)
Pop. Density	-	-0.79***	-0.16
	-	(0.26)	(0.13)
Pop. Density Sq.	-	0.18***	0.04*
	-	(0.05)	(0.02)
R^2	0.53	0.67	0.52
Adj. R^2	0.47	0.63	0.43
Instrument	FMAP	DOT	DM
Observations	50	50	50

Note: Table displays estimates of λ from (126) alongside (heteroskedasticity-robust) standard errors from regressions described in the main text. Each column refers to estimates from the FMAP, DOT and DM instrument respectively. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

replicates Chodorow-Reich (2019), where the first three columns report results for each of the instruments, and the final column uses 2SLS using all three instruments simultaneously. As in Chodorow-Reich (2019) I find that, without additional controls, the estimated multiplier is around 1.8-2.3 and highly significant for each study. With additional controls, the multiplier estimate falls for the FMAP instrument (and becomes insignificantly different from zero), and rises but becomes very imprecisely estimated for the DOT instrument (with the instrument no longer "strong" at conventional levels). The DM instrument, which combines variation from various sources, remains significant and positive (2.40), and using all three instruments simultaneously (alongside all control variables) yields an unchanged (albeit less precise) estimate for the multiplier of 2.01. Overall, the evidence for large local fiscal multipliers from ARRA appears somewhat weaker than the discussion in Chodorow-Reich (2019), although jointly the three studies continue to provide significant evidence of sizable jobs multipliers (around 2).

Assessing Bias With The Model To understand these results further, I consider a more realistic calibration of the model from Appendix F.1 that allows me to run various experiments directly related to these ARRA regressions. I assume the model is populated with the same sample of 50 states from Chodorow-Reich (2019). I use data from the Commodity Flow Survey (CFS) of 2007 to calibrate the trade linkages between all (i, j) pairs of states, α_{ij}. I take data

Table 14: Replication and Extension of Chodorow-Reich (2019)

	Dependent Variable: Job-Years per 100k			
	(1)	(2)	(3)	(4)
A: Original				
ARRA Spending	2.29***	2.22*	1.82**	2.01***
	(0.75)	(1.28)	(0.72)	(0.63)
F-stat.	26.07	10.75	49.21	36.29
R^2	0.49	0.49	0.49	0.49
Adj. R^2	0.44	0.44	0.45	0.45
B: + Controls				
ARRA Spending	1.46	3.77	2.40**	2.01**
	(1.36)	(2.36)	(1.22)	(1.01)
F-stat.	12.74	4.66	25.96	17.62
R^2	0.55	0.43	0.56	0.56
Adj. R^2	0.48	0.35	0.46	0.47
Instrument	FMAP	DOT	DM	All
Observations	50	50	50	50

Note: Table displays coefficient estimates alongside (heteroskedasticity-robust) standard errors and F-statistics from IV regressions with reduced-form (127) with different control variables a_i. Panel A displays estimates with no additional controls, and Panel B displays estimates with the controls from Table 13. Each column refers to estimates from the FMAP, DOT and DM instrument respectively, with the final column reporting 2SLS estimates employing all instruments simultaneously. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

on regional MPCs from Bellifemine et al. (2023) to calibrate mpc_i for each state.[53] I consider a static version of the model in which outcomes for all states $y = [y_1, ..., y_n]$, can be expressed as:

$$y = \mathcal{M}^g \varepsilon^g + u \tag{129}$$

where $\varepsilon^g = [\varepsilon^g_1, ..., \varepsilon^g_n]$ and $u = [u_1, ..., u_n] \overset{i.i.d.}{\sim} N(0, I_n)$. I then assume the government spending shocks are targeted based on the following rule:

$$\varepsilon^g_i = \delta' a_i + \eta_i \tag{130}$$

where a_i denotes a vector of (non-stochastic) state characteristics and $\eta = [\eta_1, ..., \eta_n] \overset{i.i.d.}{\sim} N(0, \sigma^2_\eta I_n)$. I calibrate the targeting rule using the regressions from Table 13, taking fitted values as estimates of $\delta' a_i$ and using the variance of the residuals as an estimate of σ^2_η. Mimicking

[53] I follow the same steps as in Appendix D, first constructing state-level estimates by population-weighting the county-level estimates in Bellifemine et al. (2023). I use data from 2011 (roughly corresponding to the ARRA period).

the regressions from Table 14, I then run:

$$y_i = \alpha + \beta g_i + e_i$$
$$y_i = \alpha_l + \beta_l g_i + \gamma_l' a_i + e_i^l$$

In this case, since multipliers are heterogeneous across states (while the variance of the shocks is the same across states, given by σ_η^2), the appropriate estimand is the average multiplier, defined as:

$$\beta^* = \frac{1}{n}\sum_i^n [\mathcal{M}^g]_{i,i}$$

I estimate the bias of $\hat{\beta}$ and $\hat{\beta}_l$ across stochastic draws of (u, η) for each of the three instruments from Chodorow-Reich (2019) – i.e. $\mathbb{E}[\hat{\beta} - \beta^*]$ and $\mathbb{E}[\hat{\beta}_l - \beta^*]$. Table 15 reports my results, reporting the bias as a percent of the estimand. Although non-negligible, I find the bias for the short regression is relatively mild in each case. The bias is largest for the FMAP instrument, although only 5% in this case. The long regression reduces the bias for the FMAP and DOT instruments and leaves it essentially unchanged for the DM instrument.

Table 15: Model-Implied Bias From ARRA Instruments

Estimator	FMAP	DOT	DM
Short regression	4.47%	1.24%	1.26%
Long regression (+ a_i)	-0.43%	-0.41%	-1.28%

To understand why the short-regression bias is generally muted, and also much larger for FMAP than for DOT or DM, note that the bias is driven by the correlation between a state's own targeted spending and the spillover exposure that targeting pattern implies via the trade network in the model. Table 16 reports the correlation between these terms for each instrument. The correlation consistently remains below 0.5, and is much larger for the FMAP instrument relative to the others. Thus, FMAP's regional clustering aligns considerably more with the calibrated CFS trade linkages than DOT's or DM's density-based clustering does – consistent with the fact that states in the same Census region trade disproportionately with one another, whereas the sparsely populated Western states which dominate the DOT and DM instruments engage in relatively little bilateral trade.

Table 16: Correlation Between Targeted Shock and Spillover Term

FMAP	DOT	DM
0.46	0.19	0.33