Missing Intercept, Biased Slope: Identification in General Equilibrium*

Ed Manuel†

February 21, 2026

Abstract

Many studies use cross-sectional data to identify the "direct" (or "partial equilibrium") effects of macroeconomic shocks. I consider identification of these parameters within a general class of linear data-generating processes that feature agent-heterogeneity, dynamics and general equilibrium effects. I focus on identification via panel shift-share regressions and present various results, both negative and positive. My negative results relate to identification via "exogenous shocks". I show this approach fails to identify objects of interest due to heterogeneous general equilibrium effects across the population. Instead, I demonstrate a robust and simple approach to identification by adding control variables that capture the targeting rule of the exposure-shares across units. I illustrate my results in a Two-Agent New Keynesian Model. In an application estimating 'local' fiscal multipliers, I find that adjustments to regression specifications in line with my theoretical results substantially alter the conclusions from previous studies around the effectiveness of military spending at stimulating activity.

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 effects of macroeconomic shocks are often key parameters in macroeconomic models, providing 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, 2020).

In this paper, I take as given the usefulness of these objects and study methods for reliably achieving identification. 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 then consider identification via cross-sectional methods. I focus on panel shift-share regressions (see e.g. Borusyak et al., 2025, for a review), departing from previous work along a number of important dimensions. First, I depart from the "Stable Unit Treatment Value Assumption" (SUTVA) – an assumption that underpins classic results on identification across the applied micro literature but which explicitly rules out the possibility of general equilibrium effects arising from the intervention under study. I show that departures from this assumption – alongside incorporating other key features of my environment – has material consequences for identification. Second, my framework maps directly to (a class of) heterogeneous-agent dynamic stochastic general equilibrium models, exactly the macroeconomic models which cross-sectionally identified moments are typically intended to be informative about. Third, my framework jointly specifies processes for macroeconomic aggregates alongside unit-level outcomes, and is cast within the Slutzky-Frisch paradigm commonly used to study dynamic causal effects in macroeconomics (Stock and Watson, 2018). This allows me to present a unified analysis of identification in macroeconomics using either time-series or cross-sectional data.

Identification Results I present a series of results on identification via shift-share panel regressions within my environment – both negative and positive. To fix ideas, I focus on identification of the direct effects of government stimulus – e.g. household marginal propensities to consume out of government transfers or regional government-spending multipliers. My negative results relate to identification via "exogenous shocks". I show that even if the researcher has access to a series of essentially random adjustments in government stimulus at the national level over time, shift-share regressions fail to identify the direct effects of the stimulus since they are contaminated by the differential response of treatment and control groups to the general equilibrium effects of the shock. These negative results echo some previous work on identification in macroeconomic settings which I then extend along numerous dimensions. One novel result is that, even if the researcher has access directly to the differential general equilibrium effects of the shock on each unit, and adds control variables to capture this, the resulting estimand can put negative weights on some treatment effects. Hence the estimand can suffer from sign reversals – taking on a negative value when all treatment effects are positive in the population – a similar issue that has been shown to arise within the difference-in-differences literature.

My main contribution is to demonstrate an approach to identification that is immune to these concerns. Specifically, I show that identification can be achieved by adding control variables that capture the underlying 'targeting rule' of the government stimulus across units. In a panel shift-share setting, this requires interacting these control variables with the time-series shock of interest. This approach remains robust in the face of various features that typically pose significant challenges for identification e.g. heterogeneous and dynamic treatment effects, forward-looking expectations, general equilibrium effects, etc. The control-based approach I propose is simple to implement and does not require the researcher to have any knowledge of the underlying general equilibrium effects of the shock under study. Nevertheless, the informational requirements are generally strong since the control variables must span the various dimensions of unit-level heterogeneity that simultaneously determine treatment and drive heterogeneous responses to the general equilibrium (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) and regressions include a constant term, although it is generally harder to satisfy in empirical settings where variation in stimulus received by each unit lacks a clear assignment mechanism (as in e.g. Nakamura and Steinsson, 2014).

I take this result further by adding the assumption that impulse responses to all shocks in my environment have an appropriate "sequence-space" representation (Auclert et al., 2021). This allows me to move away from the traditional goal in macroeconometrics of identifying impulse responses to particular macroeconomic *shocks* and towards an alternate goal of identifying the effects of particular treatment *paths*. I demonstrate that, even when the researcher lacks access to a series of government stimulus shocks, under an appropriate "no anticipation" assumption on the aggregate shifter, panel local-projection shift-share coefficients have the interpretation of tracing out dynamic causal effects in response to dynamic treatment paths for the stimulus policy under study (e.g. paths for household transfers, paths for regional government spending etc.). The idea that estimated impulse response functions are best interpreted in this way has precedence both in applied work (see e.g. Fukui et al. (2025) for a recent example), and in a more recent literature in macroeconomics on the mapping between the effects of policy shocks and policy rules (McKay and Wolf, 2023; Barnichon and Mesters, 2023). I build on this by utilising this interpretation constructively to formalise novel results on identification. In particular, this re-characterisation of the object of interest allows me to formalise exactly the usefulness of employing cross-sectional data to answer questions of interest in macroeconomics (vis-à-vis using time-series data) since it demonstrates how cross-sectional identification can proceed when researchers are unable to credibly identify time-series shocks of interest.

My baseline environment considers an infinite population of agents who respond heterogeneously to a set of aggregate variables that adjust in general equilibrium following a shock. I extend this to consider identification with a finite number of agents who respond directly to treatment of each individual agent, covering a class of (linearised) network models. I show that a similar bias arises due to heterogeneous spillover effects from aggregate shocks in these settings. The bias can be large when unit-level shocks are "clustered", such that units who receive large stimulus shocks simultaneously experience larger spillover effects arising from stimulus shocks to other units. I demonstrate that an appropriate correction is again to control for the underlying targeting rule that generates the clustering of the unit-level shocks.

Quantitative Example I demonstrate my results within a simple calibration of a Two-Agent New Keynesian 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 marginal propensities to consume 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 shift-share regressions that interact the shock of interest with the share of the stimulus going to each unit can be heavily biased for the MPC. For example, when the transfers are targeted entirely to Ricardian households, 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.

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. In this case, the resulting estimate retains the desired interpretation as capturing a weighted average of MPCs of recipient households for essentially any assignment of transfers across the population.

I then show that, with these control variables, the researcher can recover an estimate of (intertemporal) MPCs by projecting on any combination of time-t shocks in the model. This holds even when aggregate transfers follow an endogenous policy rule (e.g. responding to adjustments in aggregate inflation). In this case, a shift-share regression which includes a technology shock interacted with the stimulus shares (alongside appropriate controls) retains its interpretation as capturing household MPCs, even though the shifter in this case captures variation in transfers that are endogenous to business cycle developments.

Empirical Applications I then consider the practical importance of my econometric results through various empirical applications. I first revisit the empirical application of Nakamura and Steinsson (2014) who estimate local fiscal multipliers out of military spending using a shift-share research design. Their chosen shock reflects changes in national military spending, similar in spirit to canonical shocks employed in the time-series literature (Ramey, 2011). My results imply that identification fails even when this shock captures genuinely random adjustments in national government spending over time. I then show that the exposure-shares in the shift-share regression from Nakamura and Steinsson (2014) are highly predictable with just a handful of variables that standard macroeconomic models suggest should be important determinants of heterogeneity in regional responses to GE objects – e.g. measures of regional openness and MPCs (Bellifemine et al., 2023). Including combinations of these variables as additional controls (interacted with the shock of interest) greatly reduces the estimated local fiscal multiplier in the shift-share regression – from 2.5 in the original study to 0.5-0.9 depending on the specification, where the latter are tightly estimated (e.g. rejecting multipliers above 1.5 at conventional levels of significance). This essentially reverses the headline results of the paper that regional multiplier estimates are more in line with those from a "New Keynesian" model than a "Neoclassical" model, and underscores the practical significance of my econometric results.

I also revisit Chodorow-Reich (2019) who combines three studies estimating local fiscal multipliers out of the 2009 stimulus package in the US - Chodorow-Reich et al. (2012), Wilson (2012), Dupor and Mehkari (2016). I show all three instruments exhibit significant regional clustering. I demonstrate explicitly with a simple quantitative model that such clustering tends to overstate local fiscal multipliers when spending stimulates external demand and trade is governed by a gravity equation. I use discussion from the original studies regarding the underlying targeting rule to propose various control variables in order to clean each instrument of the clustering. With these additional control variables I find weaker evidence of above-1 fiscal multipliers than in Chodorow-Reich (2019), although, taken together, the three studies continue to jointly estimate a multiplier out of the 2009 stimulus package of around 2.0.

Literature This paper relates to various strands of literature.

First, and most directly, I relate to a literature on (panel) shift-share research designs. The majority of this literature considers identification in settings without dynamics or general equilibrium effects (Adao et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022b; Arkhangelsky and Korovkin, 2020; Majerovitz and Sastry, 2023). I show that incorporating these features has important implications for identification – and demonstrate a route to identification within a comprehensive framework that incorporates all of: heterogeneity, dynamics and general equilibrium effects. My work is complementary to this previous literature in that I highlight that the key to identification in this environment is to compare across cross-sectional units that are appropriately "balanced" (as in e.g. Goldsmith-Pinkham et al., 2020; Arkhangelsky and Korovkin, 2020) rather than solely by relying on exogeneity of the time-series shocks (as in e.g. Adao et al., 2019; Borusyak et al., 2022b; Majerovitz and Sastry, 2023). My environment of interest is closely related to Almuzara and Sancibrián (2024), and, like them, I begin with a general characterisation of shift-share regression estimands in that setting. I then employ this characterisation to answer a different set of questions from Almuzara and Sancibrián (2024) – they are concerned with inference (not identification), and consider cases where the researcher is interested only in recovering an estimate of unit-level heterogeneity in impulse responses (rather than the "direct" effect of shocks per se).[1] Thus, while they provide a general framework for inference when employing microeconomic data in macroeconomic settings, I work within the same environment to provide a general framework for identification.

In recent work, Donaldson (2026) also considers identification via shift-share panel regressions in a dynamic stochastic general equilibrium setting. As in this paper, she notes that typical shift-share specifications are not robust when agents respond heterogeneously to GE objects. The two works are complementary and jointly offer applied researchers a menu of options for both assessing and correcting for biases stemming from heterogeneous responses to GE objects. Donaldson (2026) focuses on a setting where agents respond homogeneously to their own treatment but are heterogeneously exposed to changes in a single GE variable (the interest rate), and shows that 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)). In contrast, I consider a general setting where agents respond heterogeneously to their own treatment as well as to multiple GE objects and develop an approach to identification that is robust to this possibility. This is crucial since macroeconomic shocks typically exhibit multiple GE channels of transmission – 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' discount factors – simultaneously drives heterogeneous responses to the other channels. I demonstrate a route to identification that is robust to this concern and which succeeds even when the researcher is unable to credibly identify any individual macroeconomic shock (typically the motivation for employing cross-sectional data to achieve identification). Various other work also considers identification via panel data in macroeconomic settings – although their focus is on achieving identification of the aggregate effects of shocks e.g. via SVAR-based identification schemes (Sarto, 2025; Matthes et al., 2025).

My discussion of biases that can arise for shift-share regressions when agents respond heterogeneously to GE effects also closely relates to recent discussion of identification in the presence of spatial spillovers in Adão et al. (2019) and Borusyak et al. (2022a), who focus on trade

[1] Specifically, they are primarily concerned with studies such as Ottonello and Winberry (2020) who estimate the heterogeneous effects of monetary policy shocks across firm-types, while I am concerned with studies such as Nakamura and Steinsson (2014) who seek to recover an estimate of local fiscal multipliers.

and migration models respectively. Both papers propose to achieve identification by first writing down a fully-specified GE model and then estimating a model-consistent reduced-form equation. While motivated by similar concerns to this paper, the setting of interest and proposed solutions are very different – since my solution 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.

Further, I relate to a literature in macroeconomics on the appropriate interpretation of cross-sectional regressions from the perspective of macroeconomic models (Nakamura and Steinsson, 2014, 2018; Chodorow-Reich, 2019; Koby and Wolf, 2020; Chodorow-Reich, 2020; Guren et al., 2021; 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. (2021) – in which case the intercept term in cross-sectional regressions (or likewise the time fixed effect in panel regressions) captures all relevant general equilibrium effects of policy interventions (and so the issues for identification discussed in this paper do not arise).[2] The fact that identification of key model parameters is complicated by the presence of heterogeneous agents and general equilibrium effects appears to have first been noted in Koby and Wolf (2020) who employ a fully calibrated structural model to circumvent the problem. Wolf (2023) also considers identification in a setting with heterogeneous agents but assumes that the researcher has access to a series of well-identified stimulus shocks and that shock-exposure across agents is by construction randomly assigned – where my key contribution is to demonstrate a route to identification when both these assumptions fail.

I also relate to a literature analysing the estimands of cross-sectional research designs in settings with heterogeneous treatment effects – 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). I depart from this literature by considering identification in a (semi-parametric) framework that nests a class of heterogeneous-agent DSGE models. My key result leverages the insights from Borusyak and Hull (2024) who show that design-based (but not model-based) regression-control specifications identify a convex weighted-average of treatment effects – where I show that this insight can still be used constructively to achieve identification of objects of interest in this environment.

Throughout, the discussion is focused on the macroeconomics literature (and on identifying the effects of fiscal stimulus in particular). However, my results are also informative for applications across other areas of economics – e.g. in finance (e.g. Chodorow-Reich, 2014), labour (e.g. Acemoglu and Restrepo, 2020), trade (e.g. Autor et al., 2013) and development

[2]In these settings, cross-sectional regressions identify direct effects of interventions, and the intercept term soaks up the GE effects alongside the average effect of any other shocks, which generates the classic "missing intercept" problem.

(e.g. Gerard et al., 2021; Egger et al., 2022; Galego Mendes et al., 2023), – that employ cross-sectional identification in settings with potentially sizable spillover effects via general equilibrium 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 Two-Agent New Keynesian model. Sections 5 and 6 contain empirical applications. Section 7 concludes.

2 Simple Setting

I begin with a simple setting to demonstrate how bias can arise in settings with heterogeneous-agents and GE effects and to describe heuristically my proposed solution. Consider some aggregate shock (or "treatment"), denoted g, distributed across a population $i \in N$. Let g_i denote the treatment received by each unit where $g = \int_i g_i di$. For simplicity, consider the case with binary treatments: $g_i \in \{0, 1\}$. Suppose units are heterogeneous, with outcomes determined by the following simple (linear) causal model with heterogeneous treatment effects:

$$y_i = \beta_i g_i + \zeta_i + u_i \tag{1}$$

where β_i captures the effect of unit is own treatment (holding fixed treatment for all other units in the population), ζ_i captures the effect of general equilibrium spillovers arising from the aggregate shock (i.e. how unit-i responds to g, even absent own treatment: $g_i = 0$) and u_i captures the counterfactual outcome that would have occurred for each unit absent the shock. In this setting, ζ_i could capture unit-specific responses to aggregate variables that adjust in general equilibrium in response to the aggregate shock (e.g. monetary policy) or spillovers across units that form part of a 'network'.

Object of Interest In this setting the effect of the shock on the aggregate outcome variable $y = \int_i y_i di$ is given by:

$$\int_i (\beta_i g_i + \zeta_i)\, di$$

Although this is typically the *ultimate* object of interest in macroeconomics, the purpose of this paper is to discuss specifically how cross-sectional methods can be effectively employed to recover an estimate of the "direct" effects of the intervention on individual units. As has been argued elsewhere, these objects serve as "identified moments" across a broad class of models and thereby provide powerful evidence on key questions in macroeconomics by providing effective indirect evidence on the overall effects of policy interventions (Nakamura and Steinsson, 2018). When agents are assumed to be heterogeneous, the goal of typical regression analysis is to learn about some average measure of these direct effects:

$$\int_i \tilde{\omega}_i \beta_i di \tag{2}$$

for $\int_i \tilde{\omega}_i = 1$ and $\forall i : \tilde{\omega}_i \geq 0$. Note that the requirement that the $\tilde{\omega}_i$'s sum to be one and be non-negative ensures that the object is indeed interpretable as a weighted average of underlying treatment effects.[3] The typical approach then uses cross-sectional regressions to recover this object of interest e.g. to run:

$$y_i = \alpha + \beta^{ols} g_i + e_i \tag{3}$$

Identification Under what assumptions does this regression recover the object of interest? Treating the outcome variable as the *change* in an outcome relative to the period before the shock, the OLS coefficient in (3) reduces to the difference-in-differences estimator, where it is well-known that identification is typically achieved under the following 'parallel trends' assumption:

Assumption 1. *Absent the treatment, outcomes for treated and control groups would have evolved in parallel:*

$$E[u_i|g_i = 1] = E[u_i|g_i = 0] \tag{4}$$

Note in this setting however, this assumption does not suffice for identification of the object of interest, as the following clarifies:

Lemma 1. *Under Assumption 1, the OLS estimand from* (3) *can be decomposed as:*

$$\beta^{ols} = \underbrace{E[\beta_i|g_i = 1]}_{\text{ATT}} + \underbrace{E[\zeta_i|g_i = 1] - E[\zeta_i|g_i = 0]}_{\text{Differential responses to G.E. effects of shock}} \tag{5}$$

The first term is exactly the object of interest and so the second term can be interpreted as generating bias. This bias term does not arise either when there are no spillovers arising from the treatment ($\zeta_i = 0$) – when the Stable Unit Treatment Value Assumption (SUTVA) holds – or when these spillovers affect all units exactly the same ($\zeta_i = \alpha$). While SUTVA may be a reasonable assumption in some settings – as with 'small-scale interventions' such as those studied in Fagereng et al. (2021) and Boehm et al. (2025) – it is clearly inappropriate in settings

[3] In particular, requiring $\tilde{\omega}_i \geq 0$ ensures that the resulting estimate does not suffer from a risk of 'sign reversals', an issue that has been discussed extensively in the microeconometrics literature on identification with heterogeneous treatment effects – see e.g. de Chaisemartin and D'Haultfœuille (2023).

where the researcher themselves claim that the intervention under study generated potentially sizeable general equilibrium effects.

Why would this bias term be problematic in practice? Consider a case where a researcher estimates a high-β^{ols} coefficient following a government spending shock across regions. it is tempting to conclude that this implies government spending is effective at stimulating consumption at the *local* level – a result which is more consistent with traditional 'Keynesian' accounts around the effect of stimulus. But another interpretation of this finding is that the treatment and control groups responded very differently to the GE effects of the shock. This could occur, if for example, the additional spending was financed via distortionary labour taxes, and the control group were more sensitive to the this higher taxation (e.g. had a higher Frisch elasticity of labour supply). In this case, the result is consistent with a 'neoclassical' story in which government spending is not effective at stimulating consumption (at the local or national level), but distortionary taxes used to finance such spending has differential (negative) effects across regions.[4]

How might researchers ensure that their empirical specification is robust to this concern? The spirit of my proposed solution relies on constructing treatment and control groups that are appropriately 'balanced' with respect to their responsiveness (or exposure) to the general equilibrium effects from the shock. Even when the treatment is targeted in such a way that this balance does not hold, bias can be avoided so long as researchers include controls that capture this targeting rule.

3 Identification Results

The 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 all proofs from this section to Appendix A.

3.1 Environment

I begin with a general characterisation of a (linearised) DSGE model with agents who are heterogeneous across a range of (non-time-varying) characteristics. I consider an economy with a unit-continuum of agents $i \in [0, 1]$ and an infinite horizon $t = 0, 1, 2, \ldots$. Agents are heterogeneous across some n_x vector of (non-time-varying) variables, denoted by $x_i = [x_{1,i}, x_{2,i}, \ldots, x_{n_x,i}]'$. The economy is subject to n_ε independent, mean-zero, unit-variance stochastic shocks denoted by $\varepsilon_t = [\varepsilon_{1,t}, \ldots, \varepsilon_{n_\varepsilon,t}]'$. I treat the first shock as a particular shock of interest

[4]The discussion here is deliberately stylised. I present a formal example in Section 4 where I demonstrate that standard cross-sectional identification methods mistakenly imply that transfers to Ricardian households in a Two-Agent New Keynesian model have very large stimulative effects on household consumption.

which I refer to as a government stimulus shock – I label this as $\varepsilon_{1,t} \equiv \varepsilon_t^g$ and collect remaining shocks in the $(n_\varepsilon - 1)$-length vector ε_t^{-g}.

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 $(n_Y - 2)$-length 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 unit-level outcome variable maps to the aggregate outcome variable in the standard way: $y_t = \int_i y_{i,t} di$.

I assume the treatment of interest is distributed arbitrarily across agents, with z_i denoting the "share"[5] of government stimulus going to each agent, with their own treatment denoted $g_{i,t} = z_i g_t$. Aggregate and unit-level variables are linear functions of all current and past shocks (as well as steady-state values):

$$Y_t = \sum_{l=0}^{\infty} \Theta_l \varepsilon_{t-l} + \bar{Y} \tag{6}$$

$$Y_{i,t} = \sum_{l=0}^{\infty} \Theta_{i,l} \varepsilon_{t-l} + \bar{Y}_i \tag{7}$$

where Θ_l and $\Theta_{i,l}$ are $n_Y \times n_\varepsilon$ matrices that map shocks to aggregate and unit-level outcomes respectively, and $\bar{Y}$ and $\bar{Y}_i$ are n_Y-length vectors denote steady-state outcomes. It is well-known that the Structural Vector-Moving Average (SVMA) model (6) encompasses all discrete-time, linearized DSGE models, where (7) then extends this structure to the unit-level outcomes (as in Almuzara and Sancibrián (2024)). I focus on the dynamic causal effect of government stimulus shocks on each units' outcome $y_{i,t}$, which I assume admits the following "sequence-space" decomposition:

$$\boldsymbol{\Theta}_{i,1}^{y} = \mathcal{M}_i^{yg} \boldsymbol{\Theta}_1^{g} z_i + \mathcal{M}_i^{yw} \boldsymbol{\Theta}_1^{w} \tag{8}$$

where $\boldsymbol{\Theta}_{i,1}^{y} = \{\Theta_{i,h,1}^{y}\}_{h=0}^{\infty}$, $\boldsymbol{\Theta}_1^{g} = \{\Theta_{h,1}^{g}\}_{h=0}^{\infty}$ and $\boldsymbol{\Theta}_1^{w} = \{\Theta_{h,1}^{w}\}_{h=0}^{\infty}$ are sequences collecting the dynamic causal effects of ε_t^g on $y_{i,t}$, g_t and w_t, and $\mathcal{M}_i^{yg}$ and $\mathcal{M}_i^{yw}$ are conformable and map bounded sequences into bounded sequences, and depend only on underlying sources of agent-heterogeneity x_i.[6]

Decompositions of this form are common within macroeconomic models cast in the sequence-space to distinguish between the notion of "direct" (or "partial equilibrium") and "indirect"

[5] I do not restrict z_i's to be literal shares in the sense of all being positive and integrating to one over the population.

[6] Formally, I impose that the mappings are measurable with respect to x_i in that, for any bounded sequence $\boldsymbol{v}$, $E[\mathcal{M}_i v | x_i] = \mathcal{M}_i v$.

(or "general equilibrium") effects of shocks. In my setting, the first term can be understood as capturing the "direct effect" of the shock – since it captures unit is responses to the path of their own stimulus $g_{i,t}$, 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 ($z_i = 0$). Since decompositions of the form (8) arise naturally in linear settings, this imposes relatively minor additional structure on the transmission of the shock to each-unit – allowing agents to respond heterogeneously to an essentially arbitrary set of transmission channels in a forward-looking way.

As well as the standard assumption that shocks are mutually independent and i.i.d. over time, I assume they are independent of unit-level variables:

$$\forall i \, \forall t : \varepsilon_t \perp (z_i, x_i, \bar{y}_i) \tag{9}$$

where this is a natural assumption given the unit-level variables are non-time-varying.

3.2 Object of Interest

The goal is to study the direct effect of government stimulus on each agent (holding fixed treatment for all other agents). To ease notation, I focus on the effect at the h'th horizon, $[\Theta^y_{i,h,1}]$, and following (8), decompose this as:

$$\Theta^y_{i,h,1} = \underbrace{\beta^y_{i,h}}_{\text{Direct Effect}} z_i + \underbrace{\zeta^y_{i,h}}_{\text{GE Effect}} \tag{10}$$

where $\beta^y_{i,h}$ and $\zeta^y_{i,h}$ correspond to the h'th term of the sequences $\mathcal{M}^{yg}_i \mathbf{\Theta}^g_1$ and $\mathcal{M}^{yw}_i \mathbf{\Theta}^w_1$ respectively. My object of interest is then simply some weighted average of my notion of direct effects – i.e. $E[\omega_i \beta^y_{i,h}]/E[\omega_i]$ for some non-negative weights ω_i. To tighten the connection between $\beta_{i,h}$ and typical objects of interest in macroeconomic models, it is helpful to consider various examples. [7]

Transfers to Households Consider government transfers $g_{i,t}$ to households in a setting where household consumption is determined according to: $\mathbf{c}_{i,t} = \mathbf{c}_i(\mathbf{g}_{i,t}; \mathbf{w_t})$ where w_t collects aggregate variables that the household treats as given (e.g. the real interest rate, taxes etc.) and boldface denotes sequences: $\mathbf{x}_{i,t} = \{x_{i,s}\}_{s=t}^{\infty}$ for all variables-x. To first-order, given $g_{i,t} = z_i g_t$, the response of household consumption to a shock to transfers ε^g_t is given by:

[7] Note that my notion of "direct effects" differs from Wolf (2023) who uses this terminology exclusively for the partial equilibrium effects of interventions holding *all* prices fixed. In this paper, I'm only concerned with recovering an estimate that is interpretable as the effect of *own*-treatment holding fixed treatment for all other units. As I discuss below, this corresponds directly to typical objects of interest in studies that employ cross-sectional data – e.g. household marginal propensities to consume, firm-level responses to tax cuts/subsidies, regional fiscal multipliers.

$$\boldsymbol{\Theta}^y_{i,1} = \left[\frac{\partial \boldsymbol{c}_i}{\partial \boldsymbol{g_{i,t}}}\tilde{\boldsymbol{g}}_t\right] z_i + \frac{\partial \boldsymbol{c}_i}{\partial \boldsymbol{w_t}}\tilde{\boldsymbol{w}}_t$$

where $\tilde{\boldsymbol{x}}_t$ denotes the effect of the shock on paths for aggregate variables $\boldsymbol{x}_t$. In models with atomistic households, the second-term does not depend on z_i, and so the first term captures my object of interest. This corresponds exactly to the notion of households' intertemporal marginal propensities to consume (Auclert et al., 2024): how households adjust their dynamic consumption path over time in response to news about a dynamic transfer path (holding fixed any general-equilibrium objects).

Government Spending in Regions Consider government purchases $g_{i,t}$ across regions in settings where output in each region is determined according to: $\mathbf{y}_{i,t} = \mathbf{y}_i(\mathbf{g}_{i,t}, \mathbf{p}_{i,t}; \mathbf{w_t})$ where boldface again denotes sequences. In this case, the response of regional output to a shock to government purchases ε^g_t is given by:

$$\boldsymbol{\Theta}^y_{i,1} = \left[\frac{\partial \boldsymbol{y}_i}{\partial \boldsymbol{g}_{i,t}}\tilde{\boldsymbol{g}}_t\right] z_i + \frac{\partial \boldsymbol{y}_i}{\partial \boldsymbol{p}_{i,t}}\tilde{\boldsymbol{p}}_{i,t} + \frac{\partial \boldsymbol{y}_i}{\partial \boldsymbol{w}_t}\tilde{\boldsymbol{w}}_t$$

As before, with atomistic regions, the final term typically does not depend on z_i. But in regional models the response of prices in each region typically depends on region-specific shocks. Assuming this term can be further decomposed into the adjustment that occurs as a result of the shock in region-i vs shocks in "foreign" regions $g^F_{i,t}$, we have:

$$\tilde{\boldsymbol{p}}_{i,t} = \left[\frac{\partial \boldsymbol{p}_{i,t}}{\partial \boldsymbol{g}_{i,t}}\tilde{\boldsymbol{g}}_t\right] z_i + \frac{\partial \boldsymbol{p}_{i,t}}{\partial \boldsymbol{g}^F_{i,t}}\tilde{\boldsymbol{g}}^F_{i,t}$$

And so, in this case, my notion of the direct effect includes the adjustment of local prices that occurs as a result of the shock – i.e. the conventional notion of a local fiscal multiplier.[8]

Subsidies in a Production Network Consider government production subsidies $g_{i,t}$ across firms in settings with k industries that form a production network, where the output of each firm is determined according to: $\mathbf{y}_{i,t} = \mathbf{y}_i(\mathbf{g}_{i,t}; \mathbf{p}_t)$ where $p_t = [p_{1,t}, ..., p_{k,t}]$ is a vector of prices such that $p_{l,t}$ denotes the price of the the good from industry-l, where this is constructed as an appropriate aggregator of prices from each individual firm in that industry. The response of each firm's output to a subsidy shock ε^g_t is given by:

$$\boldsymbol{\Theta}^y_{i,1} = \left[\frac{\partial \boldsymbol{y}_i}{\partial \boldsymbol{g}_{i,t}}\tilde{\boldsymbol{g}}_t\right] z_i + \frac{\partial \boldsymbol{y}_i}{\partial \boldsymbol{p}_{l,t}}\tilde{\boldsymbol{p}}_{l,t}$$

[8]The fact that regional regressions typically include the effect of any adjustment in local prices is discussed extensively in e.g. Guren et al. (2021) and Wolf (2019). If researchers are concerned with recovering a "pure" partial equilibrium effect in settings with regional data then additional steps are typically required and discussed in both Guren et al. (2021) and Wolf (2019). My results are still useful in this case since they clarify how to achieve identification in the initial regression (i.e. the first step for the application of the methods proposed in these other papers).

With atomistic firms – as in e.g. Rubbo (2023) – the second term again does not depend on z_i and so the first term captures my notion of direct effects – the response of output for each firm to a government subsidy (holding industry prices fixed).

3.3 Identification

My focus is on identification via shift-share panel regressions. I first characterize the population estimand of a generic shift-share regression:

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

$$y_{i,t+h} = \beta^{ss} s_i * d_t + e_{i,t} \tag{11}$$

where $\forall i \forall t : s_i \perp \varepsilon_t$ *and* $d_t = \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}$ *for some* $1 \times n_\varepsilon$*-length vector* Ψ_l*. Then:*

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

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

Proof See Appendix A.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 closely 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 observable aggregate variable.[9]

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 underlying structural shocks that underpin the shifter d_t and "weights" these up according to the contribution of each shock in d_t. Note this combination of impulse responses functions reflect 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 – see Corollary 1, Appendix A.2. The shift-share coefficient is then constructed by regressing these weighted 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, $d_t = \varepsilon_{k,t}$, then so long as the shift-share regression includes time fixed effects, the result-

[9] One small caveat to this is that I restrict d_t to be zero-mean in Proposition 1. I do so for expositional purposes although it is immediate that a version of Proposition 1 goes through allowing d_t to be non-zero mean so long as unit fixed effects are included.

ing estimate simply takes the difference of impulse responses between treatment and control groups:

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

Besides providing some underlying intuition for shift-share regressions, Proposition 1 is useful since – in combination with applications of Frisch-Waugh-Lovell – it allows me to study directly the population estimand of essentially any possible shift-share specification within my environment.

I now consider routes to identification. The previous literature has stressed two distinct routes for shift-share identification – either via exogenous shares (Goldsmith-Pinkham et al., 2020) or exogenous shocks (Adao et al., 2019; Borusyak et al., 2022b). Within my environment, these two distinct routes can be summarised as follows:

Condition 1 **(Identification From Shocks)** $d_t = \varepsilon^g_t$

Condition 2 **(Identification From Shares)** $z_i \perp x_i$

The first condition assumes the econometrician has direct access to a series of government stimulus shocks. Since these are independent of all other structural shocks, this mimics an experimental ideal where the researcher is able to directly observe a series of essentially random adjustments in government spending at the national level over time. The second condition instead assumes that the exposure of each unit to changes in government stimulus z_i is independent of any unit-characteristics x_i. This mimics an alternate experimental ideal where the exposures z_i are essentially randomly assigned across the population. I now consider the validity of each of these routes within my environment.

3.3.1 Identification From Shocks

I start by considering identification when the researcher has direct access to a series of government spending shocks (i.e. Condition 1 holds). The following Proposition clarifies that this is generally not sufficient to achieve identification of the object of interest:

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

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon^g_t + e_{i,t} \tag{12}$$

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

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

Proof See Appendix A.3

Following the discussion of Proposition 1, in this case the shift-share coefficient is constructed by differencing impulse responses for treatment and control groups. However, absent any restriction on z_i, the resulting coefficient is contaminated by the differential responses of treatment and control groups to the general equilibrium effects of the shock. Previous macroeconomics literature has derived a similar bias term for shift-share regressions with identified shocks in particular settings with heterogeneous agents (see e.g. Koby and Wolf, 2020; Donaldson, 2026) – where the above result extends this to my general environment. As I come on to demonstrate in the next section, the bias in macroeconomic settings can be severe, dramatically altering conclusions around the effectiveness of particular government stimulus policies.

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.[10] 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:*

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon_t^g + \lambda' z_i * w_t + e_{i,t} \tag{13}$$

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

$$\beta^{ss} = \lambda_{0,1} \underbrace{E[\beta^y_{i,h}|z_i = 1]}_{\text{ATT}} + \underbrace{\lambda_{0,1} \left(E[\zeta^y_{i,h}|z_i = 1] - E[\zeta^y_{i,h}|z_i = 0] \right)}_{\text{Bias from G.E. responses}}$$
$$+ \underbrace{\sum_{k=2}^{n_\varepsilon} \lambda_{0,k} \left(E[\Theta^y_{i,0,k}|z_i = 1] - E[\Theta^y_{i,0,k}|z_i = 0] \right)}_{\text{Bias from other contemp. shocks}} + \underbrace{\sum_{l=0}^{\infty} \sum_{k=1}^{n_\varepsilon} \lambda_{l,k} \left(E[\Theta^y_{i,h,k}|z_i = 1] - E[\Theta^y_{i,h,k}|z_i = 0] \right)}_{\text{Bias from other lagged shocks}}$$

where the λ-loadings are given by:

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

and:

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

[10] This approach to dealing with this concern is common in applied work – see e.g. Nakamura and Steinsson (2014), Pennings (2021), Chodorow-Reich et al. (2021) and Fukui et al. (2025). Donaldson (2026) refers to regressions of the form (13) as the "control function" approach and similarly demonstrates bias within their environment.

Proof See Appendix A.4

Proposition 3 characteristics the bias in this case purely in terms of the (macro) and (micro) structural impulse responses Θ_l^w and $\Theta_{i,l}^y$. By Frisch-Waugh-Lovell, the shift-share coefficient in this case is constructed by first orthogonalising the shock with respect to w_t, and then using this orthogonalised shock $(\varepsilon_t^g)^{\perp w_t}$ as the shifter. Since w_t generally reflects an outcome variable of the shock ε_t^g, the additional bias terms can be understood as stemming from the "bad control" problem, where conditioning on outcome variables introduces bias (even when the treatment is randomly assigned). In particular, orthogonalising ε_t^g with respect to w_t recovers a linear combination of all shocks up to time-t, and so the resulting shift-share coefficient now additionally captures the differential response of treatment and control groups to all (past and contemporaneous) shocks, where the λ-loading on each term depends on the particular combination of shocks captured by $(\varepsilon_t^g)^{\perp w_t}$. It is of course possible that these additional bias terms offset each other, although there is no reason for this to hold in general. I demonstrate this with a quantitative example in later sections where I show the bias from this approach can be severe.

One way to understand the bias here is that the controls fail to effectively "hold fixed" the relevant GE variables that drive heterogeneous responses across units. Returning to the sequence-space decomposition (16), 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 $\mathbf{\Theta}_k^w$ - which in general is not captured by the additional control variables in (13). Indeed, it follows from Proposition 1 that researchers can test directly whether their controls successfully recover a linear combination of shocks that serves to "hold fixed" these GE variables via (time-series) regressions of the form:

$$w_{t+h} = \beta_h^w(\varepsilon_t^g) + \delta' w_t + u_{t+h}$$

where generically, although $\beta_h^w = 0$ by construction, we have: $\beta_h^w \neq 0$ for $h > 0$, and even $\beta_h^w \neq 0$ for $h < 0$ – i.e. the researcher will additionally now be capturing differential responses to GE variables that occur both after *and* before time-t. One way to remove the bias term from Proposition 2 is to follow the logic of the solution proposed in Donaldson (2026) and instead project on a linear combination of time-t shocks that produce opposite paths for GE variables than ε_t^g – i.e. projecting on a set of time-t shocks that effectively hold fixed the entire dynamic path for w_t. Clearly, orthogonalising ε_t^g with respect to w_t does not in general recover the necessary combination of time-t shocks to implement this strategy.

I now consider an alternate strategy to directly control for units' differential responses to GE variables. To investigate this, I consider a situation where the researcher has direct access to the differential GE responses of agents to the shock – i.e. to $\zeta_{i,h}^y$ in equation (10). The informational requirements in this case are exceptionally high – since $\zeta_{i,h}^y$ is the h'th element of the sequence $\mathcal{M}_i^{y,w}\mathbf{\Theta}_k^w$ from equation (16), it requires the researcher to have knowledge both of how GE variables adjusted to the shock (captured by $\mathbf{\Theta}_k^w$), and how units responded differentially to such changes (captured by $\mathcal{M}_i^{y,w}$). Nevertheless, one might imagine a situation where these heterogeneous spillover terms take on a simple form – e.g. with spillovers only operating between "groups" or "regions" of units or via some known "network" structure – such that it is possible to construct a plausible set of controls that capture this. This speaks to a large literature that seeks to identify the direct effect of shocks on each unit alongside identifying heterogeneous spillover effects (see e.g. Huber, 2022). Somewhat surprisingly, I now show that even with this knowledge, controlling for $\zeta_{i,h}^y$ does not recover the object of interest:

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

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon_t^g + \lambda \zeta_{i,h}^y * \varepsilon_t^g + e_{i,t} \tag{14}$$

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

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

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 $\zeta_{i,h}^y$.

Proof See Appendix A.5

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 have an interpretation as capturing some (properly) weighted average of the underlying direct effects. In particular, the estimate has the potential to suffer from "sign reversals" – a negative value for β_{ss} even when the underlying direct effects of the shock are all positive. The result follows from Proposition 1, combined with results in Borusyak and Hull (2024) around negative weights with "model-based" controls. Problems associated with negative weights have been discussed extensively within the difference-in-differences literature in the context of two-way fixed-effect estimation – the above result then highlights that this problem also arises in settings where researchers attempt to control for heterogeneous spillover effects (while leaving unmodeled any heterogeneity in units' responses to their *own* treatment). While it is possible that the issue of negative weights from regressions of the form (14) may be avoided by using alternate (non-OLS) estimators, since this route to identification necessarily requires exceptionally high informational requirements (requiring knowledge of $\zeta_{i,h}$), I instead go on to consider alternate 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 Condition 2 holds), then regression (12) no longer suffers from bias since, when the shares are orthogonal to unit characteristics, we have:

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

Condition 2 is strong however and so I consider a weakening that permits identification via control variables even when z_i is correlated with x_i. Letting a_i denote some observable n_a-vector of control variables, consider:[11]

Condition 3 (Identification From Shares With Controls) $E[z_i|x_i, a_i] = \alpha + \delta' a_i$

In this case, controlling for a_i interacted with the shock of interest is sufficient for identification:

Proposition 5 *Suppose Condition 3 holds. Consider shift-share regressions of the form:*

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon^g_t + \lambda' a_i * \varepsilon^g_t + e_{i,t} \tag{15}$$

Then:

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

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

Proof See Appendix A.6

The controls here serve to remove the bias term while ensuring the weights ω_i are guaranteed to be non-negative. As above, the result follows from Proposition 1, combined with results in Borusyak and Hull (2024) around the avoidance of negative weights with "design-based" controls. 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 essentially unrestricted GE effects. The identifying assumptions are strong in that it requires researchers to model correctly the assignment of the shares z_i and to include controls that capture this. Specifically the controls

[11] I again assume $\forall i, \forall t : a_i \perp \varepsilon_t$ such that Proposition 1 still applies.

must span any dimensions of unit-heterogeneity x_i that determine z_i. 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. In practice, in order to implement this procedure, researchers must first take a stance on the unit-level characteristics that are likely to drive heterogeneous responses to GE effects in their setting – x_i in my framework – and then, to the extent that the exposure-shares z_i are likely to be correlated with x_i, to control directly for x_i in the shift-share regression (or argue that they have some control variables a_i that effectively 'span' x_i). I demonstrate how to implement this procedure explicitly with examples via a quantitative model and an empirical application in later sections.

Proposition 5 continues to rely on the researcher having access to a single well-identified shock ε_t^g (i.e. that Condition 1 holds). 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 this assumption can be meaningfully weakened. To that end first consider a weaker version of Condition 1:

Condition 4 (No Anticipation) $d_t = q'\varepsilon_t$ for $q \in R^{n_\varepsilon \times 1}$

Note that this condition ensures that the shifter d_t is a function of only contemporaneous shocks. Within this environment, that allows the shifter to be essentially any (endogenous) variable so long as it is not predictable prior to time-t – hence "no anticipation". I now show that this alternate condition is sufficient to achieve identification of an object closely related to the object of interest discussed thus far:

Proposition 6 *Assume Condition 3 and 4 hold. Suppose in addition to (8), the effect of all shocks can be decomposed as:*

$$\forall k \in (1, ..., n_\varepsilon) : \boldsymbol{\Theta}_{i,k}^y = \mathcal{M}_i^{yg} \boldsymbol{\Theta}_k^g z_i + \mathcal{M}_i^{yw} \boldsymbol{\Theta}_k^w \tag{16}$$

Consider the pair of shift-share regressions of the form:

$$y_{i,t+h} = \alpha_{i,y} + \delta_{t,y} + \beta_{y,h}^{ss} z_i * d_t + \lambda' a_i * d_t + e_{i,t} \tag{17}$$

$$g_{i,t+h} = \alpha_{i,g} + \delta_{t,g} + \beta_{g,h}^{ss} z_i * d_t + \lambda' a_i * d_t + u_{i,t} \tag{18}$$

Then:

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

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

Proof See Appendix A.7

The Proposition 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}.

The result relies on extending the sequence-space decomposition (8) 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 (16) 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).

This result then formalises exactly the usefulness of employing cross-sectional data in macroeconomic settings. In particular, it clarifies 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. Note further that, in this case, 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 (so long as it continues to satisfy the "no anticipation" assumption). This is potentially a fruitful avenue of exploration for applied work. As Auclert et al. (2024) note, the full intertemporal MPC (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 to different macroeconomic shocks could in theory be used to identify the full iMPC matrix $\mathcal{M}_i^{yg}$. I provide a simple illustration of this point via a Two-Agent New Kenynesian model in the next section.

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.

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 B, 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 GE variables. As discussed in the simple setting from Section 2, the exogeneity requirements are conceptually distinct, and generally somewhat stronger, than the canonical parallel trends assumptions employed in the difference-in-differences literature.

Finite-n I also consider settings with a finite-n population of agents, where spillovers occur directly between units. This speaks to a general class of (linearised) "network" models, which do not directly map to the main environment from Section 3.1. In Appendix C, I show that, in this environment, a similar bias arises as in Proposition 1, and that a similar correction can be employed, following Proposition 5. I supplement these results in Appendix E 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 regionally clustered and trade follows a "gravity" equation. The bias arises since units with high levels of government spending are simultaneously more exposed to higher government spending in neighbouring regions via increased demand for exports than the control group. A simple correction in this case is to control for the underlying treatment assignment that generated such clustering, where this can be as simple as including region fixed-effects. This discussion relates to my empirical application in Section 6.

4 Quantitative Example

To demonstrate my key results in a familiar setting, I consider the use of shift-share regressions to recover an estimate of household consumption responses to transfers in a Two-Agent New Keynesian (TANK) model. I relegate all details of the model to Appendix D – I provide a high-level overview of the model in Section 4.1 and then discuss identification in Section 4.2.

4.1 Setting

Households The model is populated by a unit mass of households $i \in [0,1]$ with preferences over consumption $c_{i,t}$ and labour $n_{i,t}$. Households receive income from their labour via the real wage w_t, alongside real dividend income d_t, real taxes τ_t and transfers $\mathcal{T}_{i,t} = z_i \mathcal{T}_t$. A fraction $\alpha \in [0,1]$ of households are unconstrained (U) and can choose to hold bonds b_t with (ex-post) real return r_t, while the remaining fraction are hand-to-mouth (H), unable to borrow or save. Each unconstrained household $i \in [0,\alpha]$ solves a familiar dynamic consumption-savings problem:

$$\max_{c_{i,t},n_{i,t},b_{i,t}} E_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{(c_{i,t})^{(1-\sigma)}}{1-\sigma} - \theta \frac{(n_{i,t})^{(1+\chi)}}{1+\chi} \right] \\ \text{s.t. } c_{i,t} + b_{i,t} \leq w_t n_{i,t} + d_t - \tau_t + r_t b_{u,t-1} + \mathcal{T}_{i,t} \tag{19}$$

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

$$\max_{c_{i,t},n_{i,t}} \left[\frac{(c_{i,t})^{(1-\sigma)}}{1-\sigma} - \theta \frac{(n_{i,t})^{(1+\chi)}}{1+\chi} \right] \\ \text{s.t. } c_{i,t} \leq w_t n_{i,t} - \tau_t + \mathcal{T}_{i,t} \tag{20}$$

Firms Production is identical to the standard New-Keynesian model with a representative competitive final goods firm alongside monopolistically competitive intermediate goods producers, the latter facing sticky prices.

Policy Monetary policy is set according to a Taylor Rule. The government's budget constraint is:

$$G_t + \int_i z_i \mathcal{T}_t = \tau_t \tag{21}$$

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 $\varepsilon_t^{\mathcal{T}}$, the economy is also subject to shocks to government spending ε_t^G, monetary policy ε_t^m and technology ε_t^A.

Shares I define the binary variable:

$$HTM_i = \mathcal{I}(i \in [\alpha, 1])$$

where $\mathcal{I}(.)$ denotes the indicator function. I consider binary exposures to the transfer $z_i \in \{0,1\}$ where the probability of treatment varies with hand-to-mouth status:

$$Pr(z_i = 1|HTM_i = 0) = \delta_1 \text{ and } Pr(z_i = 1|HTM_i = 1) = \delta_2$$

For simplicity, I restrict δ_1 to be binary: $\delta_1 \in \{0, 1\}$, and allow δ_2 to take on values in the range $[0, 1]$. In this case the model essentially reduces to a "Three"-Agent model consisting of: Ricardian households (who either receive the transfer or not); (ii) "Treated" hand-to-mouth households (who receive the transfer); (iii) "Control" hand-to-mouth households (who do not receive the transfer), where the relative size of groups (ii) and (iii) depends on the choice of δ_2. This keeps the model tractable while allowing me to demonstrate the full range of results from the previous section.

4.2 Identification

I consider an econometrician who seeks to recover an estimate of households' marginal propensities to consume out of government transfers. Given the household problems (19) and (20), individual household consumption functions – that determines each household's consumption choices over time – can be written as a function of sequences in: r_t, w_t, d_t, τ_t, $\mathcal{T}_{i,t}$. That is, for all i, t, we have:

$$\boldsymbol{c}_{i,t} = \boldsymbol{c}_i(\boldsymbol{\mathcal{T}}_{i,t}, \boldsymbol{w}_t, \boldsymbol{d}_t, \boldsymbol{\tau}_t)$$

And so, to first-order, the effect of government transfer shocks on household consumption can be decomposed as:

$$\frac{d\boldsymbol{c}_{i,t}}{d\varepsilon_t^{\mathcal{T}}} = \left[\mathcal{M}_i^{c\mathcal{T}} \tilde{\boldsymbol{\mathcal{T}}}_t\right] z_i + \mathcal{M}_i^{cr} \tilde{\boldsymbol{r}}_t + \mathcal{M}_i^{cd} \tilde{\boldsymbol{d}}_t + \mathcal{M}_i^{c\tau} \tilde{\boldsymbol{\tau}}_t + \mathcal{M}_i^{cw} \tilde{\boldsymbol{w}}_t$$

where $\mathcal{M}_i^{cx}$ denotes $\dfrac{\partial \boldsymbol{c}_i}{\partial \boldsymbol{x_t}}$ and $\tilde{x}_t$ denotes $\dfrac{\partial \boldsymbol{x}_t}{\partial \boldsymbol{\varepsilon_t}}$. My interest is in the first term – which here captures household intertemporal marginal propensities to consume (iMPC) – which I label:

$$\boldsymbol{\beta}_i = \mathcal{M}_i^{c\mathcal{T}} \tilde{\boldsymbol{\mathcal{T}}}_t$$

Since this varies only depending on whether households are hand-to-mouth or not, I define:

$$\mathcal{M}_U^{c\mathcal{T}} = \forall i \in [0, \alpha] : \mathcal{M}_i^{c\mathcal{T}} \text{ and } \mathcal{M}_{HTM}^{c\mathcal{T}} = \forall i \in [\alpha, 1] : \mathcal{M}_i^{c\mathcal{T}} \tag{22}$$

$$\boldsymbol{\beta}_U = \{\forall i \in [0, \alpha] : \beta_{i,s}\}_{s=0}^{\infty} \text{ and } \boldsymbol{\beta}_{HTM} = \{\forall i \in [\alpha, 1] : \beta_{i,s}\}_{s=0}^{\infty} \tag{23}$$

as the values of $\mathcal{M}_i^{c\mathcal{T}}$ and $\boldsymbol{\beta}_i$ for Ricardian and hand-to-mouth households respectively. As in the previous section I consider different routes to identification.

4.2.1 Identification With Exogenous Shocks

I assume the econometrician has access to the true transfer shocks ε_t^g, and begin by considering shift-share regressions of the form:

$$c_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon_t^g + e_{i,t} \tag{24}$$

Figure 1 shows my results for different values of δ_1 and δ_2. Figure 1a plots the path for transfers following the shock. Figures 1b and 1c then compare the shift share coefficients, shown in the orange lines, to the true iMPC of the treated group (i.e. the ATT), shown in the black dots, for the extreme cases where transfers are targeted entirely to Ricardian households ($\delta_1 = 1, \delta_2 = 0$) or entirely to hand-to-mouth households ($\delta_1 = 0, \delta_2 = 1$). These plots serve to highlight the potential severity of the bias term from Proposition 2. For example, when transfers are targeted to the Ricardian households (Figure 1b), the true iMPC 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.

Figure 1: Targeted transfers in TANK (Exogenous Shocks, Endogenous Shares)

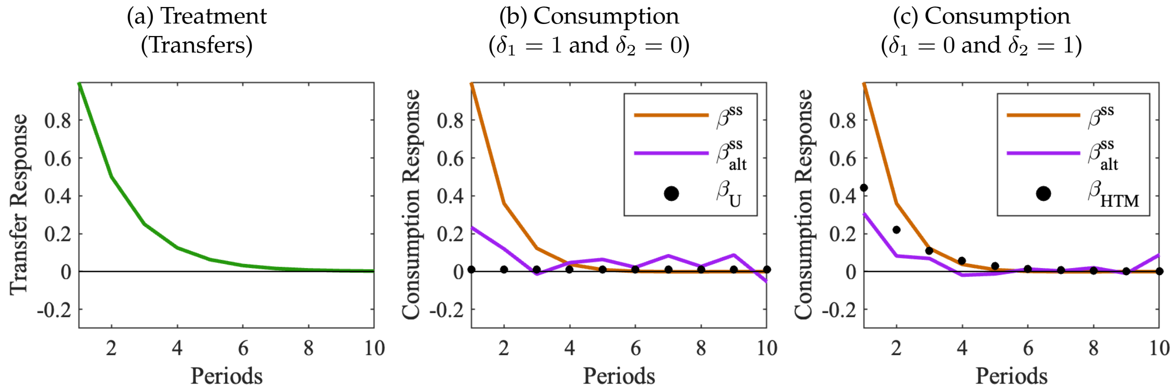

Notes: Green line in (a) plots the impulse response of transfers following the shock ε_t^g – i.e. $\tilde{\mathcal{T}}_t$. Orange (Purple) lines in (b) and (c) denote β_{ss} (β^{ss}_{alt}) from (24) for the case with: $\delta_1 = 1$, $\delta_2 = 0$ and $\delta_1 = 0$, $\delta_2 = 1$ respectively. Black dots in (b) and (c) denote the 'direct' response of unconstrained and hand-to-mouth households to the shock respectively, defined by (22). All responses are scaled to the contemporaneous rise in transfers i.e. to $\tilde{\mathcal{T}}_0$.

Next, I consider a correction in line with Proposition 3. Specifically, I construct estimates of β^{ss}_{alt} from a regression that includes as additional controls (interacted with the shares z_i):

$$c_{i,t+h} = \alpha_i + \delta_t + \beta^{ss}_{alt} z_i * \varepsilon_t^g + \beta^{ss} z_i * m_t + e_{i,t} \tag{25}$$

where $m_t = [r_t, d_t, \tau_t, w_t]$ – i.e. exactly the aggregate GE variables that determine differential responses across treated and control groups in this setting. Figure 1 shows my results, where I additionally report estimates of β^{ss}_{alt} from (25) in the purple line. This alternate estimator 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).

4.2.2 Identification With Exogenous Shares

Now I demonstrate that identification can instead proceed via exogenous shares. Following Proposition 5, I do not assume that the shares z_i are directly exogenous, but rather that the econometrician is aware that treatment status varied depending on whether households were hand-to-mouth, and includes appropriate control variables to account for this. Specifically, I consider regressions of the form:

$$c_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon^g_t + \lambda HTM_i * \varepsilon^g_t + e_{i,t} \tag{26}$$

I consider the case where: $\delta_1 = 0$ and $\delta_2 = 0.5$. In this case, using Proposition 1, the shift-share coefficient is essentially recovered from a "saturated" regression of each units' impulse response to the shock on z_i controlling for a constant and HTM_i. Hence the coefficient on z_i simply captures the difference in impulse responses for hand-to-mouth agents in "treatment" and "control" groups. Figure 2 displays my results. As before, Figure 2a plots the path for transfers following the shock. Figure 2b then demonstrates that the shift share coefficient, shown in the orange lines, recovers exactly an estimate of hand-to-mouth households' marginal propensities to consume, shown in the black dots.

Figure 2: Targeted transfers in TANK (Exogenous Shocks, Exogenous Shares)

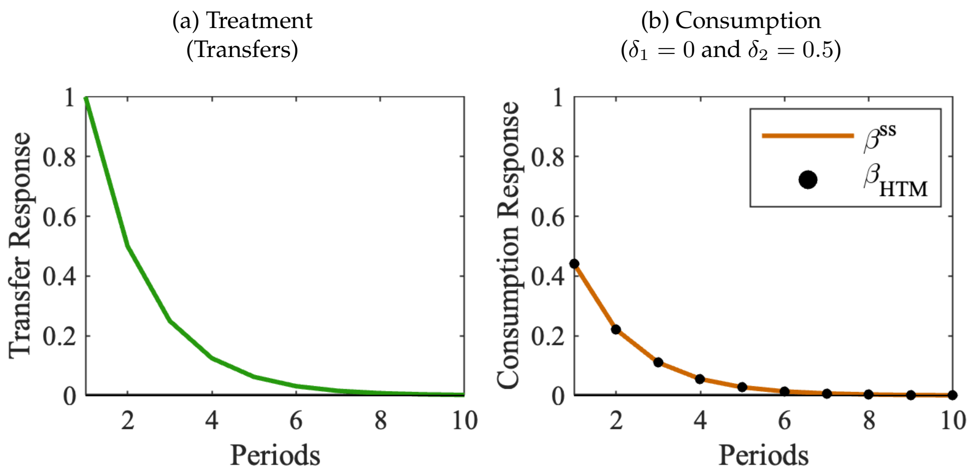

Notes: Green line in (a) plots weighted impulse responses of transfers to shocks to $\varepsilon_{\mathcal{T}}$ and $\varepsilon_{\mathcal{A}}$— i.e. $\tilde{\mathcal{T}}_t$. Orange (Purple) lines in (b) and (c) denote β_{ss} (β^{ss}_{alt}) from (24) for the case with: $\delta_1 = 1$, $\delta_2 = 0$ and $\delta_1 = 0$, $\delta_2 = 1$ respectively. Black dots in (b) and (c) denote the 'direct' response of unconstrained and hand-to-mouth households to the shock respectively, defined by (22). All responses are scaled to the contemporaneous rise in transfers i.e. to $\tilde{\mathcal{T}}_0$.

Finally, I consider an application of Proposition 6. To do so I tweak slightly the baseline model such that transfers additionally respond to other macroeconomic shocks.[12] Specifically I assume that transfers also follow a Taylor-style rule of the form:

$$\mathcal{T}_t = (1-\rho_\mathcal{T})\mathcal{T} + \rho_\mathcal{T}\mathcal{T}_{t-1} + (1-\rho_\mathcal{T})\phi_\mathcal{T}(\log \Pi_t - \log \Pi) + s_\mathcal{T}\varepsilon_t^\mathcal{T}. \tag{27}$$

where Π_t denotes aggregate inflation and $\phi_\mathcal{T} > 0$ such that the transfers rise in response to higher inflation. I now assume the econometrician does not have access to the transfer shock $\varepsilon_\mathcal{T}$, but does have access to the technology shock $\varepsilon_\mathcal{A}$:[13]

$$d_t = \varepsilon_t^\mathcal{A} \tag{28}$$

I then assume they run pairs of regressions of the form

$$c_{i,t+h} = \alpha_{i,c} + \delta_{t,c} + \beta_c^{ss} z_i * d_t + \lambda_c HTM_i * d_t + e_{i,t} \tag{29}$$

$$\mathcal{T}_{i,t+h} = \alpha_{i,\mathcal{T}} + \delta_{t,\mathcal{T}} + \beta_\mathcal{T}^{ss} z_i * d_t + \lambda_\mathcal{T} HTM_i * d_t + u_{i,t} \tag{30}$$

Figure 3 below shows my results. The green line in figure 3a traces out the path for transfers following the technology shock. As in Proposition 6, the shift-share coefficient $\beta_\mathcal{T}^{ss}$, shown in the purple dots, traces out exactly this path. Figure 3b then shows the path for β_y^{ss} traced out by the shift-share regression in the orange line, where this corresponds to HTM households' iMPCs out of a transfer shock that produces the dynamic path for transfers from Figure 3a.

In this case, the variation in aggregate transfers underpinning the regressions stems from an endogenous response of transfers to business cycle developments. In particular, note that regressing *aggregate* consumption C_{t+h} on the shifter d_t would not recover a valid estimate of the causal effect of transfers on the macroeconomy. Nevertheless, identification of MPCs can still proceed via regressions (29) and (30) – where this information can still be informative about the effect of particular transfer policies at stimulating consumption. This serves to demonstrate exactly the usefulness of employing cross-sectional methods to achieve identification in macroeconomic settings.

5 Empirical Application 1: Nakamura and Steinsson (2014)

I now illustrate the practical implications of my econometric framework. In line with the framing in the previous sections, I focus on studies that seek to estimate the 'direct' effects of fiscal

[12] Note that this adjustment is not necessary to apply Proposition 6, although it serves to demonstrate how projecting on different macroeconomic shocks can produce different estimates of consumption responses to different paths for transfers.

[13] This example is only illustrative. Proposition 6 clarifies that the econometrician does not require *any* of the individual shocks from the model to achieve identification in this case, but can instead use any linear combination of time-t shocks – including say, reduced-form Wold innovations from a correctly-specified VAR.

Figure 3: Targeted transfers in TANK (No Anticipation, Exogenous Shares)

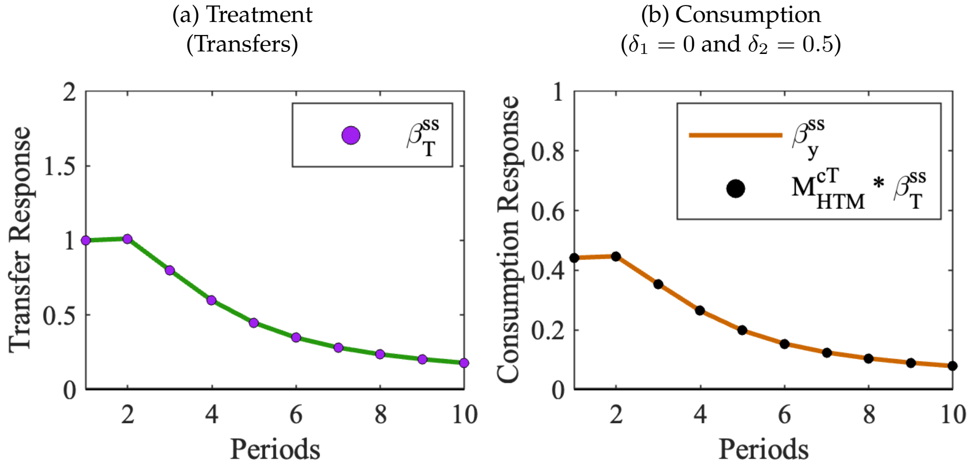

Notes: Green line in (a) plots impulse response of transfers to technology shock $\varepsilon_{\mathcal{A}}$. Purple dots in (a) trace out the shift-share coefficient $\beta_{\mathcal{T}}^{ss}$ from (30). Orange line in (b) denotes β_{ss} from (30) for the case with: $\delta_1 = 0$ and $\delta_2 = 0.5$. Black dots in (b) and (c) denote the HTM iMPC matrix multiplied by $\beta_{\mathcal{T}}^{ss}$. All responses are scaled to the contemporaneous rise in transfers i.e. to $\tilde{\mathcal{T}}_0$.

stimulus. In this section, I revisit Nakamura and Steinsson (2014), who estimate local (state-level) fiscal multipliers using a shift-share research design with a national shifter. I describe the empirical setup, then discuss issues related identification.

5.1 Setting

Nakamura and Steinsson (2014) seek to identify local fiscal multipliers through a shift-share research design. They use annual panel data on US states for 1966–2006, and estimate regressions of the form:[14]

$$\frac{y_{i,t} - y_{i,t-2}}{y_{i,t-2}} = \alpha_{i,y} + \delta_{t,y} + \beta_y^{ss} z_i * \varepsilon_t^g + u_{i,t} \tag{31}$$

$$\frac{g_{i,t} - g_{i,t-2}}{g_{i,t-2}} = \alpha_{i,g} + \delta_{t,g} + \beta_g^{ss} z_i * \varepsilon_t^g + e_{i,t} \tag{32}$$

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, ε_t^g denotes national military spending in year t, and z_i denotes the average level of military spending in state-i relative to state output in the first five years of the sample. The "local multiplier" is then constructed as the ratio of the estimated coefficients:

$$\beta_m = \frac{\beta_y^{ss}}{\beta_g^{ss}} \tag{33}$$

[14]The original study additionally considers a specification where state military spending is instrumented using total national military spending interacted with a state dummy. I focus discussion on the shift-share specification but also consider the sensitivity of results from this additional specification in Appendix F.

Nakamura and Steinsson (2014) argue that identification in this case comes from exogeneity of the shocks ε_t^g:

"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 is a claim about the implicit mechanism that lies behind the assignment of the shocks ε_t^g over time (i.e. the motivation for national military buildups), and not about the assignment of the (time-invariant) shares z_i across states. In particular, nothing in the above description suggests that states with high-exposure shares would be 'similar' to states with low-exposure shares across a range of characteristics.[15] Rather, the authors implicitly argue that any such differences do not pose an issue for identification, so long as the military spending shocks are appropriately exogenous. This claim is correct but only if there were no GE effects from military spending shocks (i.e. SUTVA holds), or if all agents respond the same to any GE effects. The authors implicitly assume the latter through the following claim:

"By including time fixed effects, we control for aggregate shocks and policy that affect all states at a particular point in time – such as changes in distortionary taxes and aggregate monetary policy."

Proposition 2 from Section 3 highlights that these claims about identification break down when agents respond heterogeneously to GE effects. Even if national military buildups occurred randomly over time, regressions of the form (31) in general fail to identify local fiscal multipliers since they are contaminated by differential effects of treatment and control groups to any GE effects of the shock.

Proposition 4 clarifies that one solution is to include additional control variables (interacted with ε_t^g) that capture the (implicit) assignment mechanism of the shares z_i across states. Specifically, the controls must span exactly the variables that determine heterogeneous responses to GE objects that move in response to the shock. What control variables would be appropriate in this case? I take as a starting point Bellifemine et al. (2023) who construct a multi-region Heterogeneous-Agent New Keynesian model of US counties and show that the response of regions to aggregate shocks are shaped by two key dimensions of heterogeneity: MPCs and trade openness. In their model, these two variables are "sufficient statistics" for regional heterogeneity, capturing differences in regional responses to shocks regardless of whether the underlying differences across states are driven by e.g. heterogeneity in discount factors, income risk or borrowing limits.[16]

[15]The authors themselves suggest that it is implausible that military spending shares are assigned randomly across states, writing: "Military spending is notoriously political and thus likely to be endogenous to regional economic conditions". As an example, and as I go on to show explicitly, one would expect military spending shares to correlate highly with states' industry composition (e.g. the balance of tradable vs non-tradable goods production).

[16]One key message of this model 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 Proposition 5: so long as

5.2 Share-Assignment

I begin by demonstrating that the shares z_i indeed correlate highly with variables that plausibly drive heterogeneity in response to GE variables. 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 be richer (i.e high-income, low MPCs) and more open (a higher fraction of employment in tradable goods) than the average state. To test this, I run regressions of the form:

$$z_i = \alpha + \lambda' a_i + u_i \tag{34}$$

where I include in a_i various state-characteristics. Specifically I include measures of states' MPCs (mpc_i) and openness ($open_i$), which I take directly from Bellifemine et al. (2023), as well as states' income per capita (inc_i), which I take from the The Correlates of State Policy database (Grossmann et al., 2021).[17] The first two variables link directly to the discussion above, while income per capita can be viewed as an additional variable that plausibly captures differences in military spending across states while also being an important determinant of regional heterogeneity in response to shocks.

Table 1: Predictability of Shares

	Dependent Variable: Shares (z_i)
MPC	-10.927^{***}
	(1.882)
Openness	2.061^{*}
	(1.223)
Income per capita	0.220^{***}
	(0.079)
Observations	51
R^2	0.431
Adjusted R^2	0.395

Note: Table displays coefficient estimates alongside (heteroskedasticity-robust) standard errors and p-values from regression (34). Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

the *assignment rule* is linear then controlling linearly for appropriate variables is sufficient for identification.

[17] The openness and MPC measures vary over time so I use the first year from which both are available (2011). The appropriate controls would capture drivers of z_i and so ideally would capture these measures from the early-sample period when the shares are constructed (1966-1971). I find there is relatively little variation in these two measures over time so expect this to be of limited importance for my results. I use a measure of states' income per capita from 1975.

I find that these three variables jointly explain over 40% of the variation in the shares z_i, with each variable significant at least at the 10% level (Table 1). The signs of the coefficients are in line with priors: states with high shares of military spending tend on average to be more open, have higher incomes and lower MPCs. To further validate the use of these three control variables specifically I consider a purely data-driven exercise where I include 13 variables in a_i that capture various dimensions of 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 Urban workers).[18] I estimate (34) via LASSO and, using 10-fold cross-validation with a λ_{1se} penalty, I find that the regression selects two variables: mpc_i and inc_i.

5.3 Shift-Share Identification

Guided by my theoretical results I then estimate the following pair of regressions:

$$\frac{y_{i,t} - y_{i,t-2}}{y_{i,t-2}} = \alpha_{i,y} + \delta_{t,y} + \beta_y^{ss} z_i * \varepsilon_t^g + \lambda_y' a_i * \varepsilon_t^g + u_{i,t} \tag{35}$$

$$\frac{g_{i,t} - g_{i,t-2}}{g_{i,t-2}} = \alpha_{i,g} + \delta_{t,g} + \beta_g^{ss} z_i * \varepsilon_t^g + \lambda_g' a_i * \varepsilon_t^g + e_{i,t} \tag{36}$$

including variously mpc_i, $open_i$ and inc_i in a_i, with the multiplier then defined appropriately as $\beta_y^{ss}/\beta_g^{ss}$. Table 2 displays my results. To maintain direct comparability with the original study I compute standard errors by clustering at the state-level. The first column replicates the baseline results from Nakamura and Steinsson (2014) with an estimated multiplier of 2.48 (significant at the 1% level). The second column adds controls for mpc_i and $open_i$ in which case the multiplier drops to 1.52 (significant at the 5% level). The third column adds controls for mpc_i and inc_i – the variables selected by LASSO as the key determinants of z_i – in which case the multiplier drops to 0.9 (significant at the 10% level). The final column adds all three control variables in which case the multiplier drops to 0.5 and is no longer significantly different from zero at conventional levels.

Although the significance levels drop as more control variables are added, the coefficients become more precisely estimated. In the final column for example, one can reject a multiplier above 1.3 at the 5% level. In Appendix F, I additionally show results for an alternate specification employed by Nakamura and Steinsson (2014) where I instrument for state military procurement using total national procurement interacted with a state dummy. Including these same additional control variables (interacted with the shock as above) produces qualitatively similar, but quantitatively less stark, results: in that case the multiplier is lower in all specifications than in the original study, but falls from a baseline of 1.43 to 1.02 with all three additional

[18] I take additional variables from the Correlates of State Policy database, where I look specifically for variables that are available in the earlier part of the sample (i.e. pre-1975) when the "shares" z_i were determined.

control variables.

Table 2: Replication and Extension of Nakamura and Steinsson (2014)

	Dependent Variable: Output			
	(1)	(2)	(3)	(4)
Prime military contracts	2.48***	1.52**	0.90**	0.52
	(0.97)	(0.81)	(0.38)	(0.40)
Additional Controls	None	mpc_i, $open_i$	mpc_i, inc_i	mpc_i, $open_i$, inc_i
Observations	1989	1989	1989	1989

Note: Table displays coefficient estimates alongside standard errors and p-values (clustered at the state level) from the IV regressions described in the main text. Significance at the 10%, 5%, and 1% levels are denoted by *, **, and *** respectively.

6 Empirical Application 2: Chodorow-Reich (2019)

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). Specifically, following Chodorow-Reich (2019), I first estimate pairs of regressions of the form:

$$y_i = \alpha_y + \beta_y \varepsilon_i^g + \delta_y' x_i + e_i \tag{37}$$

$$g_i = \alpha_g + \beta_g \varepsilon_i^g + \delta_g' x_i + u_i \tag{38}$$

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 over 2008:12 and 2010:12, x_i collects lagged state-employment (in changes and levels) and lagged growth in Gross State Production as controls for pre-treatment economic conditions, and ε_i^g is an instrument that captures exogenous shocks to ARRA spending. Local fiscal 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 related 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).

Since all instruments are constructed using pre-recession formulae, it is plausible that the resulting shocks are uncorrelated with any other shocks that drove states' outcomes over the same period. Nevertheless, bias can still arise if government spending shocks ε_i^g are distributed across units such that the spillover effects from other units' shocks differentially affected treatment and control groups. I demonstrate this explicitly with a simple quantitative model in Appendix E. In the model, each state is subject to spillovers from government spending in other states via trade channels. In this case, when demand for imports is governed by a gravity equation, bias can arise when government spending shocks are geographically clustered since the treatment and control groups are heterogeneously exposed to changes in external demand from shocks to other units.

Shock-Assignment To consider the potential for this bias to arise, I begin by visualising the shocks $\varepsilon^g = [\varepsilon^g_1, ..., \varepsilon^g_n]$. Figure 4 displays heatmaps for each measure, which all display significant regional clustering.

Figure 4: Heatmaps of Instruments Across Studies

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

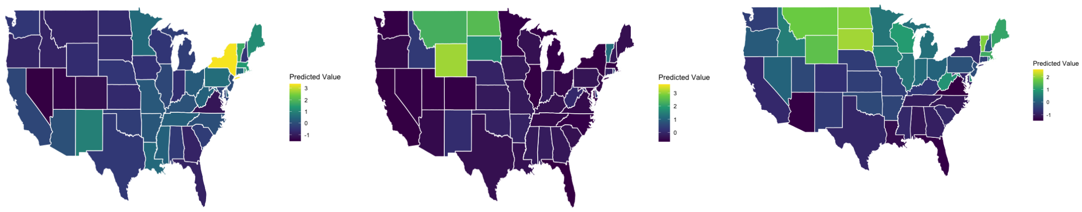

To assess this more formally I run regressions of the form:

$$\varepsilon^g_i = \alpha + \gamma \left(\sum_{j \neq i} w_{i,j} \varepsilon^g_j \right) + \theta' x_i + e_i \tag{39}$$

where, I set $w_{i,j}$ as an appropriate proxy for the relative spillover effects of the shock in state-j on state-i. I proxy for these spillover effects using geographical distance between states, consistent with various evidence of sizable local spillovers from the ARRA package proportional to geographical distance (Dupor and McCrory, 2018; Auerbach et al., 2020).[19] Specifically I set:

$$w_{i,j} = \left(\frac{d_{i,j}}{\sum_{k \neq i} d_{i,k}} \right)^{-\delta} \tag{40}$$

where $d_{i,j}$ is the geographical distance between state-i and state-j and δ controls how local spillovers decay with distance. I set $\delta = 5$ – broadly consistent with the relationship between distance and trade from the gravity literature (see Adão et al. (2019)). Panel A of Table 3 displays my results. Each of the columns represents estimates for the instruments – from

[19]Note that my identification strategy does not rely on these capturing the 'true' spillovers - I use them here only as a way of conducting a plausibility test to establish the potential for bias in this setting.

Chodorow-Reich et al. (2012) ('FMAP'), Wilson (2012) ('DOT') and Dupor and Mehkari (2016) ('DM') – where γ is positive and highly significant in each case.

Following the logic of the previous section, a way to adjust for this bias is to directly control for the implicit assignment mechanism that generated the clustering. To this end, I then estimate regressions of the form:

$$\varepsilon_i^g = \alpha + \lambda' a_i + \theta' x_i + \eta_i \tag{41}$$

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[20] – 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.[21] Panel B of Table 3 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. Finally, I confirm that these residualised instruments are plausibly successful at removing the bias by re-estimating equation (39) using the residualised shocks $\eta = [\eta_1, ..., \eta_n]$ in place of $\varepsilon_i^g = [\varepsilon_1^g, ..., \varepsilon_n^g]$. I find γ becomes small and statistically insignificant in each case (Appendix Table 8). I additionally present plots of these residualised shocks in Appendix Figure 6.

Results Finally, I estimate regressions of the form:

$$y_i = \alpha_y + \beta_y \varepsilon_i^g + \gamma_y' a_i + \delta_y' x_i + e_i \tag{42}$$

$$g_i = \alpha_g + \beta_g \varepsilon_i^g + \gamma_g' a_i + \delta_g' x_i + u_i \tag{43}$$

[20]This is directly visible from Figure 4b, where the largest shocks all occur in four very sparsely populated neighbouring Western states.

[21]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 costal navigation" which motivates the use of additional regional fixed effects.

Table 3: Predictability of Shocks

	Dependent Variable: Shocks (ε_i^g)		
	(1)	(2)	(3)
A: Model 1			
Distance-Weighted Foreign Shocks	0.68***	0.58*	0.50**
	(0.19)	(0.23)	(0.20)
Adjusted R^2	0.35	0.51	0.33
B: Model 2			
Region 2	-5.12***	-	-0.04
	(1.38)	-	(0.11)
Region 3	-6.81***	-	-0.26***
	(1.29)	-	(0.10)
Region 4	-7.37***	-	-0.23*
	(1.45)	-	(0.13)
Pop. Density	-	-0.79***	-0.16
	-	(0.22)	(0.11)
Pop. Density Sq.	-	0.18***	0.04**
	-	(0.04)	(0.02)
Adjusted R^2	0.47	0.63	0.43
Instrument	FMAP	DOT	DM
Observations	50	50	50

Note: Table displays coefficient estimates alongside (heteroskedasticity-robust) standard errors and p-values from regressions described in the main text. Panel A displays estimates of γ from (39) and Panel B displays estimates of λ from (41). 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.

with multipliers defined as the ratio: β_y/β_g. Table 4 displays my results. The first panel 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 instrument 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).

Table 4: 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)	(0.91)	(0.91)
F-stat.	26.07	10.75	49.21	36.29
Adj. R^2	0.44	0.44	0.45	0.45
B: + Assignment Controls				
ARRA Spending	1.46	3.77	2.40*	2.01*
	(1.36)	(2.36)	(1.22)	(1.01)
F-stat.	12.74	4.74	25.96	17.61
Adj. R^2	0.47	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 p-values, F-statistics and R-squareds from IV regressions with reduced-form (**??**) with different control variables a_i. Panel A displays estimates with no additional controls, and Panel B displays estimates with the controls from Panel B of Table 3. 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.

7 Conclusion

The use of micro data to answer key questions of interest in macroeconomics has risen sharply in recent years. In this paper I consider the issues that arise for cross-sectional identification in a comprehensive framework that incorporates key features of modern macroeconomic models. I show that otherwise credible research designs from the applied microeconometrics literature generally fail to identify meaningful economic parameters in these settings. I propose a simple solution that ensures cross-sectional regressions identify exactly the direct effects of policy interventions.

References

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

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

ADÃO, R., C. ARKOLAKIS, AND F. ESPOSITO (2019): "General Equilibrium Effects in Space: Theory and Measurement," NBER Working Papers 25544, National Bureau of Economic Research, Inc.

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 (2020): "On Policy Evaluation with Aggregate Time-Series Shocks," CERGE-EI Working Papers wp657, The Center for Economic Research and Graduate Education - Economics Institute, Prague.

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.

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.

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.

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," Unpublished manuscript.

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.

FLYNN, J. P., C. PATTERSON, AND J. STURM (2021): "Fiscal Policy in a Networked Economy," NBER Working Papers 29619, National Bureau of Economic Research, Inc.

FUKUI, M., E. NAKAMURA, AND J. STEINSSON (2025): "The Macroeconomic Consequences of Exchange Rate Depreciations," *The Quarterly Journal of Economics*, 140, 3015–3065.

GALEGO MENDES, A., W. MIYAMOTO, T. L. NGUYEN, S. M. PENNINGS, AND L. FELER (2023): "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 (2021): "Cash Transfers and Formal Labor Markets : Evidence from Brazil : Cash Transfers and the Local Economy: Evidence from Brazil," Policy Research Working Paper Series 9778, The World Bank.

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 (2021): "What Do We Learn from Cross-Regional Empirical Estimates in Macroeconomics?" *NBER Macroeconomics Annual*, 35, 175–223.

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

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.

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.

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.

OTTONELLO, P. AND T. WINBERRY (2020): "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 88, 2473–2502.

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

RUBBO, E. (2023): "Networks, Phillips Curves, and Monetary Policy," *Econometrica*, 91, 1417–1455.

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

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. (2019): "The Missing Intercept in Cross-Regional Regressions," Working paper.

——— (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 Section 2 Proofs

A.1 Proof of Proposition 1

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

$$y_{i,t+h} = \beta^{ss} s_i * d_t + e_{i,t}$$

where $s_i \perp \varepsilon_t$ and $d_t = \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}$ for $(1 \times n_\varepsilon)$ length (row-)vector Ψ_l.

Then:

$$\beta^{ss} = \frac{E\left[\left(\sum_{l=0}^{\infty} \Theta_{i,h+l}^{y} w_l\right) s_i\right]}{E[s_i^2]}$$

where η_{ll} *is* n_ε*-length vector of 'weights' with k'th element equal to:* $\eta_{k,l} = \frac{\psi_{k,l}}{\sum_{k=1}^{n_\varepsilon} \sum_{l=0}^{\infty} \psi_{k,l}^2}$.

Proof: The shift-share coefficient is equivalent to:

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

Plugging in equation (7):

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

Since $\forall i \forall t : \bar{y}_i, s_i \perp \varepsilon_t$ and shocks are mean-zero:

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

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

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

Switching terms:

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

Note that $\forall i, \forall t : x_i, z_i \perp \varepsilon_t \implies \forall i, \forall t, \forall h : \Theta_{i,h+l}^{y} \perp \varepsilon_t$. And also $\forall i, \forall t : s_i \perp \varepsilon_t$, so:

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

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

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

We have the desired expression:

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

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

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

A.2 Proof of Corollary 1

Corollary 1 *Consider regressions of the form:*

$$Y_{t+h} = \beta^{ols} d_t + e_t$$
$$\forall i \in [0,1] : Y_{i,t+h} = \beta_i^{ols} d_t + e_{i,t}$$

for $d_t = \sum_{l=0}^{\infty} \Psi_l \varepsilon_{t-l}$. *Then:*

$$\beta^{ols} = \sum_{l=0}^{\infty} \Theta_{h+l} \eta_l$$
$$\forall i \in [0,1] : \beta_i^{ols} = \sum_{l=0}^{\infty} \Theta_{i,h+l} \eta_l$$

where η_l *has the same form as Proposition 1.*

Proof: The proof follows the same steps as Proposition 1 setting $s_i = 1$. For completeness, for β^{ols} we have:

$$\beta^{ols} = \frac{E\left[\left(\sum_{l=0}^{\infty}(\Theta_{h+l}\varepsilon_{t-l} + \bar{Y})\right)\left(\sum_{m=0}^{\infty} \Psi_m \varepsilon_{t-m}\right)\right]}{E\left[d_t^2\right]}.$$

Since shocks are mean-zero and serially uncorrelated, the numerator reduces to:

$$\sum_{l=0}^{\infty} E[(\Theta_{h+l}\varepsilon_{t-l})(\Psi_l \varepsilon_{t-l})]$$

And using unit-variance:

$$E[(\Theta_{h+l}\varepsilon_{t-l})(\Psi_l \varepsilon_{t-l})] = \Theta_{h+l}\, E[\varepsilon_{t-l}\varepsilon_{t-l}']\, \Psi_l' = \Theta_{h+l}\Psi_l'.$$

Therefore

$$\beta^{ols} = \frac{\sum_{l=0}^{\infty} \Theta_{h+l}\Psi_l'}{E\left[d_t^2\right]} = \sum_{l=0}^{\infty} \Theta_{h+l}\eta_l,$$

where the (vector) weights are

$$\eta_l := \frac{\Psi_l'}{E[d_t^2]} = \frac{E[\varepsilon_{t-l}(\Psi_l \varepsilon_{t-l})]}{E[d_t^2]}.$$

Noting again $E[d_t^2] = \sum_{m=0}^{\infty} \Psi_m \Psi_m' = \sum_{m=0}^{\infty} \sum_{k=1}^{n_\varepsilon} \psi_{k,m}^2$, the kth element is

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

which this matches η_l from Proposition 1. it is straightforward to repeat this argument for $Y_{i,t+h}$. □

A.3 Proof of Proposition 2

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

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon_t^g + e_{i,t}$$

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

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

Proof: By FWL, the shift-share coefficient is equivalent to:

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

where $\ddot{g}_{i,t}$ denotes the residual from a (population) regression of $g_{i,t} \equiv z_i \varepsilon_t^g$ on unit- and time-fixed effects. Given shocks are mean-zero we have:

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

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

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

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

$$\beta^{ss} = E[\theta_{i,h}^{yg} | z_i = 1] - E[\theta_{i,h}^{yg} | z_i = 0]$$

The desired conclusion then follows by plugging in (10):

$$\beta^{ss} = E[\beta^y_{i,h}|z_i = 1] + E[\zeta^y_{i,h}|z_i = 1] - E[\zeta^y_{i,h}|z_i = 0] \quad \square$$

A.4 Proof of Proposition 3

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

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon^g_t + \lambda' z_i * w_t + e_{i,t}$$

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

$$\beta^{ss} = \lambda^g_0 \underbrace{E[\beta^y_{i,h}|z_i = 1]}_{\text{ATT}} + \underbrace{\lambda^g_0 \Big(E[\zeta^y_{i,h}|z_i = 1] - E[\zeta^y_{i,h}|z_i = 0]\Big)}_{\text{Bias from G.E. responses}}$$
$$+ \underbrace{\sum_{k=2}^{n_\varepsilon} \lambda_{0,k} \Big(E[\Theta^y_{i,0,k}|z_i = 1] - E[\Theta^y_{i,0,k}|z_i = 0]\Big)}_{\text{Bias from other contemp. shocks}} + \underbrace{\sum_{l=0}^{\infty}\sum_{k=1}^{n_\varepsilon} \lambda_{l,k} \Big(E[\Theta^y_{i,h,k}|z_i = 1] - E[\Theta^y_{i,h,k}|z_i = 0]\Big)}_{\text{Bias from other lagged shocks}}$$

Proof: By FWL, the shift-share coefficient is equivalent to:

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

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

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

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

$$(\varepsilon^g_t)^{\perp w_t} = \underbrace{\left(1 - b'\Theta^w_{0,1}\right)}_{\equiv \psi_{0,1}} \varepsilon^g_t + \sum_{k=2}^{n_\varepsilon} \underbrace{b'\Theta^w_{0,k}}_{\equiv \psi_{0,k}} \varepsilon_{k,t} + \sum_{l \geq 1} \underbrace{b'\Theta^w_l}_{\equiv \psi_l} \varepsilon_{t-l}$$

where:

$$b' = -(\Theta^w_{0,1})' \Sigma^{-1}_w, \qquad \Sigma_w = \sum_{l=0}^{\infty} \Theta^w_l (\Theta^w_l)', \qquad \psi_l = [\psi_{l,1} ... \psi_{l,n_\varepsilon}]$$

By Proposition 1:

$$\beta^{ss} = \frac{E[\sum_{l=0}^{\infty} \Theta^y_l \lambda_l \tilde{z}_i]}{E[\tilde{z}^2_i]}$$

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

$$\beta^{ss} = \sum_{l=0}^{\infty}\sum_{k=1}^{n_\varepsilon} \lambda_{k,l} E[\Theta^y_{l,k}|z_i = 1] - E[\Theta^y_{l,k}|z_i = 0]$$

The finally applying the decomposition (10):

$$\beta^{ss} = \lambda_{0,1} E[\beta^y_{i,h}|z_i = 1] + \lambda_{0,1}\left(E[\zeta^y_{i,h}|z_i = 1] - E[\zeta^y_{i,h}|z_i = 0]\right)$$
$$+ \sum_{k=2}^{n_\varepsilon} \lambda_{0,k}\left(E[\Theta^y_{i,0,k}|z_i = 1] - E[\Theta^y_{i,0,k}|z_i = 0]\right) + \sum_{l=0}^{\infty}\sum_{k=1}^{n_\varepsilon} \lambda_{l,k}\left(E[\Theta^y_{i,h,k}|z_i = 1] - E[\Theta^y_{i,h,k}|z_i = 0]\right) \quad \square$$

A.5 Proof of Proposition 4

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

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon^g_t + \lambda\zeta^y_{i,h} * \varepsilon^g_t + e_{i,t}$$

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

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

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

Proof By FWL, the shift-share coefficient is equivalent to:

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

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

$$\tilde{g}_{i,t} = \tilde{z}_i(\varepsilon^g_t)$$

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

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

Plugging in (10):

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

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

$$E\left[(\tilde{z}_i)^2\right] = E[z_i - a - b\zeta^y_{i,h})\tilde{z}_i]$$
$$= E[z_i\tilde{z}_i]$$

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

$$\frac{E\left[\beta^y_{i,h} z_i \tilde{z}_i\right]}{E\left[z_i\tilde{z}_i\right]} \quad \square$$

A.6 Proof of Proposition 5

Proposition 5 *Suppose Condition 3 holds. Consider shift-share regressions of the form:*

$$y_{i,t+h} = \alpha_i + \delta_t + \beta^{ss} z_i * \varepsilon^g_t + \lambda' a_i * \varepsilon^g_t + e_{i,t}$$

Then:

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

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

Proof By FWL, the shift-share coefficient is equivalent to:

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

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

$$\tilde{g}_{i,t} = \tilde{z}_i\varepsilon^g_t$$

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

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

Plugging in (10):

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

Note that under Condition 3, since the conditional expectation function is linear:

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

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

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

Finally since:

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

it then follows that:

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

A.7 Proof of Proposition 6

Proposition 6 *Assume Condition 3 and 4 hold. Consider the pair of shift-share regressions of the form:*

$$\begin{aligned} y_{i,t+h} &= \alpha_{i,y} + \delta_{t,y} + \beta^{ss}_{y,h} z_i * d_t + \lambda' a_i * d_t + e_{i,t} \\ g_{i,t+h} &= \alpha_{i,g} + \delta_{t,g} + \beta^{ss}_{g,h} z_i * d_t + \lambda' a_i * d_t + u_{i,t} \end{aligned}$$

Then:

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

for $\boldsymbol{\beta}^{ss}_y = \{\beta^{ss}_{y,h}\}^{\infty}_{h=0}$ and $\boldsymbol{\beta}^{ss}_g = \{\beta^{ss}_{g,h}\}^{\infty}_{h=0}$, where $\omega_i = E[\tilde{z}^2_i|x_i, a_i]$

Proof: Start with $\boldsymbol{\beta}^{ss}_y$. By Proposition 1:

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

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 (16):

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

Then by Condition 3:

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

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

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

Next:

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

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

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

Now turn to β_g^{ss}. By FWL, letting $x_{i,t} = z_i d_t$:

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

where $\tilde{g}_{i,t+h}$ and $\tilde{x}_{i,t}$ denote $g_{i,t+h}$ and $x_{i,t}$ orthogonalised wrt unit- and time-fixed effects and $a_i * d_t$. Given $g_{i,t+h} = z_i g_t$, we have:

$$\tilde{g}_{i,t+h} = \tilde{z}_i(g_t - E[g_t]), \qquad \tilde{x}_{i,t} = \tilde{z}_i(d_t - E[d_t]),$$

Then since the OLS estimand satisfies scaling invariance, $\beta_{g,h}^{ss}$ is equivalent to a population regression of $(g_t - E[g_t])$ on $(d_t - E[d_t])$, which by Corollary 1 gives:

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

And so finally:

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

B Extension 1: 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 treat the shocks as fixed (i.e. non-stochastic) in this case. Specifically I suppose the econometrician has access to only a single government spending shock that occurs at some fixed point in time s and uses cross-sectional local-projections to estimate the direct (dynamic) causal effect of the shock. Consider the following two conditions in this setting:

Condition 1a : $\{\varepsilon_t^g\}_{t\neq s} = 0$ and $\forall k \neq 1 : \Theta_k^g = 0$

Condition 2a : $E[z_i|x_i] = 0$

Condition 1a ensures that the the researcher observes some one-off stimulus shock ε_s^g, where the aggregate path for the stimulus is pinned down purely by that one-off shock. This is a natural condition and merely ensures that the estimand of interest is interpretable as recovering the effect of some specific (unanticipated) stimulus policy.[22] Condition 2a is then familiar and is the key exogeneity requirement, ensuring that exposure to the stimulus is (mean) independent of underlying agent characteristics.[23]

Proposition A1 *Assume* (16) *holds alongside Conditions 1a and 2a. Consider a cross-sectional regression for some single time-series observation ($t = s$) of the form:*

$$y_{i,s+h} - y_{i,s-1} = \beta^{ols} z_i \varepsilon_s^g + e_i$$

Then:

$$\beta^{ols} = \frac{E[\beta_{i,h}^y \omega_i]}{E[\omega_i]}$$

where $\omega_i = E[(z_i)^2]$.

Proof: First note that the path for the outcome variable can be written as:

$$y_{i,s+h} = \beta_{i,h} z_i \varepsilon_s^g + \zeta_{i,h} \varepsilon_s^g + \bar{y}_i + u_{i,s+h}$$

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

$$y_{i,s+h} - y_{i,s-1} = \beta_{i,h} z_i \varepsilon_s^g + \zeta_{i,h} \varepsilon_s^g + \tilde{u}_{i,s+h}$$

where $\tilde{u}_{i,s+h}$ can again be written as a linear combination of all shocks other than ε_s^g. Given condition 1a, $\tilde{u}_{i,s+h}$ depends only on x_i (not on z_i), since no other shocks operate via the g-channel. Turning then to the OLS estimand:

[22]Note that g can be defined narrowly in this case – e.g. not as aggregate e.g. transfers or government spending, but just the portion of transfers or government spending that was adjusted at time s 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.

[23]I work with zero-mean exposure shares in this case purely for simplicity.

$$\begin{aligned}\beta^{ols} &= \frac{E[(\beta_{i,h} z_i \varepsilon_s^g + \zeta_{i,h}\varepsilon_s^g + \tilde{u}_{i,s+h})\,(z_i \varepsilon_s^g)]}{E[(z_i\varepsilon_s^g)^2)]} \\ &= \frac{E[\beta_{i,h}\,(z_i)^2]}{E[(z_i)^2]} + \frac{E[\zeta_{i,h} z_i]}{E[(z_i)]} + \frac{\varepsilon_s^g E[\tilde{u}_{i,s+h} z_i]}{(\varepsilon_s^g)^2\, E[(z_i)^2)]} \\ &= \frac{E[\beta_{i,h}\,(z_i)^2]}{E[(z_i)^2]}\end{aligned}$$

where the final line follows uses: $E[\zeta_{i,h} z_i] = E[\zeta_{i,h} E[z_i|x_i]] = 0$ and $E[\tilde{u}_{i,s+h} z_i] = E[\tilde{u}_{i,s+h} E[z_i|x_i]] = 0$. □

C Extension 2: Finite-n

I extend the results in Section 3 to cover settings with a finite population. This requires adjusting the environment somewhat.

C.1 Environment

I now consider a setting with a finite population $i \in \{1, 2, ..., n\}$, observed for periods $t = [1, 2, ...]$. I keep the rest of the environment broadly the same, but now define ε_t^g as an $n-$length vector of random variables: $\varepsilon_t^g = [\varepsilon_{1t}^g, ..., \varepsilon_{nt}^g]'$, and likewise for other shocks so that ε_t^{-g} is $(n_\varepsilon - 1) * n$ vector. I continue to assume shocks are (vector-wise) mutually independent – so e.g. $\varepsilon_t^g \perp \varepsilon_t^{-g}$ – and $i.i.d.$ over time but allow them to be non mean-zero. The unit-level outcome of interest $y_t = [y_{1t}, ..., y_{nt}]'$ evolve as before according to:[24]

$$y_t = \sum_{l=0}^{\infty} \Theta_l^y \varepsilon_{t-l} \tag{44}$$

where ε_t now denotes a $(n_\varepsilon * n)$- vector of shocks and Θ_l^y is a $n \times (n_\varepsilon * n)$ matrix.

C.2 Object of Interest

Mirroring the previous section, to focus on the object of interest, the effect of an aggregate government stimulus shock ε_t^g on each unit at some horizon h can be decomposed as:

$$\Theta_{i,h,1}^y \varepsilon_t^g = \underbrace{\beta_{i,h}}_{\text{Direct Effect}} \varepsilon_{i,t}^g + \underbrace{\sum_{j \neq i} \zeta_{i,j,h}}_{\text{Spillover Effect}} \varepsilon_{j,t}^g \tag{45}$$

where $\Theta_{i,h,1}^y$ is a $1 \times n_\varepsilon$ length (row) vector that picks the relevant entries of Θ_l^y, capturing the effect of government stimuls shocks ε_t^g on unit-i. As before the object of interest is the direct effect $\beta_{i,h}$. 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

[24]Since I now allow non mean-zero shocks I abstract from any additional stead-state term

other regions). Models of this form are fairly common in the analysis of fiscal spending – see e.g. Flynn et al. (2021).

C.3 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.

Proposition A2 *Consider a setting where outcomes can be represented by (44) and government-spending shocks for each unit $\varepsilon^g_{i,t}$ are mean-zero. Consider regressions of the form:*

$$y_{i,t} = \beta^{ols}\varepsilon^g_{i,t} + e_{i,t} \tag{46}$$

Then:

$$\beta^{ols} = \underbrace{\frac{\sum_i \beta_{i,h} E[w_i]}{\sum_i E[w_i]}}_{\text{ATT}} + \underbrace{\frac{\sum_i \sum_{j\neq i} \zeta_{i,j,h} E[\varepsilon^g_{i,t}\varepsilon^g_{j,t}]}{\sum_i E[w_i]}}_{\text{Bias from GE spillovers}}$$

where $w_i = (\varepsilon^g_{i,t})^2 > 0$.

Proof: The OLS estimand in this case is:

$$\begin{aligned}\beta^{ols} &= \frac{E[\sum_i y_{i,t}\varepsilon^g_{i,t}]}{E[\sum_i(\varepsilon^g_{i,t})^2]} \\ &= \frac{E[\sum_i(\beta_{i,h}\varepsilon^g_{i,t} + \sum_{j\neq i}\zeta_{i,j,h}\varepsilon^g_{j,t} + u_{i,t+h})\varepsilon^g_{i,t}]}{E[\sum_i(\varepsilon^g_{i,t})^2]}\end{aligned}$$

where $u_{i,t+h}$ collects the effect of all other shocks other than ε^g_t. Since shocks are mututally independent and i.i.d, $E[u_{i,t+h}\varepsilon^g_{i,t}] = 0$. And so:

$$\beta^{ols} = \frac{\sum_i \beta_{i,h} E[(\varepsilon^g_{i,t})^2]}{\sum_i E[(\varepsilon^g_{i,t})^2]} + \frac{\sum_i \sum_{j\neq i}\zeta_{i,j,h} E[\varepsilon^g_{i,t}\varepsilon^g_{j,t}]}{\sum_i E[(\varepsilon^g_{i,t})^2]} \quad \square$$

The bias arises since, even though ε^g_t is orthogonal to *other* shocks, government-spending shocks may be 'clustered' 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-$\varepsilon^g_{i,t}$ 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. 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):

Proposition A3 *Consider a setting where the shocks are assigned across units based on unit-characteristics according to the following 'rule':*

$$\varepsilon^g_{i,t} = \delta' a_i + \eta_{i,t}$$

where a_i is a n_a-length vector of observables and $\eta^g_t = [\eta_{1,t}, ..., \eta_{n,t}]$ denotes a mean-zero, i.i.d vector of random variables where:

$$\forall j \neq i : \eta_{i,t} \perp \eta_{j,t}$$
$$\forall l : \eta_t \perp \varepsilon_{t-l}$$

Consider the following regression controlling for a_i:

$$y_{i,t+h} = \beta^{ols} \varepsilon^g_{i,t} + \lambda' a_i + e_{i,t}$$

Then:

$$\beta^{ols} = \frac{\sum_i \beta_{i,h} E[\omega_i]}{\sum_i E[\omega_i]}$$

where $\omega_i = (\eta_{i,t})^2$.

Proof: Consider first the OLS estimand from a regression of $\varepsilon^g_{i,t}$ on a_i:

$$\begin{aligned}
\delta^{ols} &= E\left[\sum_i a_i a_i'\right]^{-1} E\left[\sum_i a_i \varepsilon^g_{i,t}\right] \\
&= \left[\sum_i a_i a_i'\right]^{-1} \sum_i a_i E\left[\varepsilon^g_{i,t}\right] \\
&= \left[\sum_i a_i a_i'\right]^{-1} \sum_i a_i \left(\delta' a_i + E\left[\eta_{i,t}\right]\right) \\
&= \delta
\end{aligned}$$

And so by FWL:

$$\begin{aligned}
\beta^{ols} &= \frac{E[\sum_i y_{i,t} \eta_{i,t}]}{E[\sum_i (\eta_{i,t})^2]} \\
&= \frac{E[\sum_i (\beta_{i,h} \varepsilon^g_{i,t} + \sum_{j \neq i} \zeta_{i,j,h} \varepsilon^g_{j,t} + u_{i,t+h}) \eta_{i,t}]}{E[\sum_i (\eta_{i,t})^2]} \\
&= \frac{E[\sum_i (\beta_{i,h} \left(\delta' a_i + \eta_{i,t}\right) + \sum_{j \neq i} \zeta_{i,j,h} \left(\delta' a_j + \eta_{j,t}\right) + u_{i,t+h}) \eta_{i,t}]}{E[\sum_i (\eta_{i,t})^2]}
\end{aligned}$$

where the first line substitutes in the DGP for $y_{i,t+h}$, collecting the effect of all other shocks other than ε_t^g in $u_{i,t+h}$ and the second line uses the targeting rule. Since η_t are mutually independent and mean-zero $\forall j \neq i : E[\eta_{j,t}\eta_{i,t}] = 0$. And since η_t are independent of other shocks: $\forall l : E[\eta_t u_{i,t+h}] = 0$. So:

$$\beta^{ols} = \frac{\sum_i \beta_{i,h} E\left[(\eta_{i,t})^2\right]}{\sum_i E\left[(\eta_{i,t})^2\right]} \quad \square$$

Similar in spirit to Proposition 5, the implication of this result is that, when bias arises due to 'clustering' of government spending shocks across regions, this can be removed by controlling for the underlying unit-observables that feature in the targeting rule that drives the clustering. I demonstrate this with a quantitative model of fiscal spending in US states in Appendix E.

D Two-Agent New Keynesian (TANK) Model

This appendix describes the Two-Agent New Keynesian (TANK) model used as an illustrative laboratory in the paper. The model comes directly from a graduate course taught by Eric Sims, with full details helpfully provided online: `https://sites.nd.edu/esims/files/2024/04/notes_tank_sp2024.pdf`. I reproduce the full model here for completeness and and to make explicit the key modifications relative to the baseline framework.

D.1 Overview

The economy is a standard New Keynesian model augmented with household heterogeneity. There is a unit mass of households. A fraction $\alpha \in (0, 1)$ are unconstrained ("Ricardian") households with full access to financial markets, while the remaining fraction $1 - \alpha$ are constrained or hand-to-mouth households who consume their current disposable income each period. In order to implement the exercise from Section 4, I further decompose the hand-to-mouth agents such that a fraction $\delta_2 \in (0, 1)$ receive some government cash transfer, while $1 - \delta_2$ do not. Firms operate under monopolistic competition with Calvo-style price rigidity, monetary policy follows a Taylor rule, and fiscal policy consists of government spending financed through lump-sum taxes and transfers.

D.2 Households

D.2.1 Unconstrained households

A representative unconstrained household maximizes expected lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_{u,t}^{1-\sigma}}{1-\sigma} - \theta \frac{N_{u,t}^{1+\chi}}{1+\chi} \right), \tag{47}$$

where $C_{u,t}$ denotes consumption, $N_{u,t}$ labor supply, β is the discount factor, σ is the inverse elasticity of intertemporal substitution, and χ is the inverse Frisch elasticity.

The budget constraint (in real terms) is

$$C_{u,t} + b_{u,t} = w_t N_{u,t} + d_t + \delta_1 \mathcal{T}_t + (1 + i_{t-1})\Pi_t^{-1} b_{u,t-1} - \tau_t, \tag{48}$$

where $b_{u,t}$ denotes real bond holdings, w_t the real wage, d_t firm profits rebated lump-sum, g_t denotes lump-sum transfers, τ_t taxes, i_t the nominal interest rate, and Π_t gross inflation. The additional parameter $\delta_1 \in \{0, 1\}$ determines whether or not the unconstrained households receive the government transfer $\mathcal{T}_t$.

The optimality conditions are:

$$\theta N_{u,t}^{\chi} = C_{u,t}^{-\sigma} w_t, \tag{49}$$

$$1 = \mathbb{E}_t \left[(1 + i_t) \Lambda_{u,t,t+1} \Pi_{t+1}^{-1} \right], \tag{50}$$

$$\Lambda_{u,t,t+1} = \beta \left(\frac{C_{u,t+1}}{C_{u,t}} \right)^{-\sigma}, \tag{51}$$

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

D.2.2 Constrained households

Constrained households have identical preferences but cannot borrow or save and do not own firms. I distinguish between 'treated' and 'control' constrained households – where the former receive a government transfer while the latter do not. The representative 'treated' household solves:

$$\max_{C_{h,t}^d, N_{h,t}^d} \frac{(C_{h,t}^d)^{1-\sigma}}{1-\sigma} - \theta \frac{(N_{h,t}^d)^{1+\chi}}{1+\chi} \tag{52}$$

subject to

$$C_{h,t}^d = w_t N_{h,t}^d + \mathcal{T}_t - \tau_t \tag{53}$$

And the representative 'control' household solves:

$$\max_{C_{h,t}^c, N_{h,t}^c} \frac{(C_{h,t}^c)^{1-\sigma}}{1-\sigma} - \theta \frac{(N_{h,t}^c)^{1+\chi}}{1+\chi} \tag{54}$$

subject to

$$C_{h,t}^c = w_t N_{h,t}^c - \tau_t \tag{55}$$

Aggregate consumption and labor for constrained households are then

$$C_{h,t} = \delta_2 C_{h,t}^d + (1 - \delta_2) C_{h,t}^c, \tag{56}$$

$$N_{h,t} = \delta_2 N_{h,t}^d + (1 - \delta_2) N_{h,t}^c. \tag{57}$$

D.3 Firms and Price Setting

The model features a representative competitive final goods firm alongside monopolistically competitive intermediate goods producers. A fraction $1 - \phi$ of intermediate goods producers can reset prices each period. Optimal price setting for intermediate producers yields:

$$p_t^{\#} = \frac{\varepsilon}{\varepsilon - 1}\frac{x_{1,t}}{x_{2,t}}, \tag{58}$$

$$x_{1,t} = mc_t Y_t + \phi \mathbb{E}_t \left[\Lambda_{u,t,t+1}\Pi_{t+1}^{\varepsilon} x_{1,t+1}\right], \tag{59}$$

$$x_{2,t} = Y_t + \phi \mathbb{E}_t \left[\Lambda_{u,t,t+1}\Pi_{t+1}^{\varepsilon-1} x_{2,t+1}\right], \tag{60}$$

where $p_t^{\#}$ denotes the optimal relative reset price, Y_t denotes aggregate production and mc_t denotes real marginal cost:

$$mc_t = \frac{w_t}{A_t}. \tag{61}$$

where A_t denotes aggregate productivity. Price dispersion v_t evolves according to:

$$v_t = (1-\phi)(p_t^{\#})^{-\varepsilon} + \phi \Pi_t^{\varepsilon} v_{t-1}, \tag{62}$$

$$1 = (1-\phi)(p_t^{\#})^{(1-\varepsilon)} + \phi \Pi_t^{(\varepsilon-1)}; \tag{63}$$

Finally, aggregate production satisfies:

$$Y_t v_t = A_t N_t. \tag{64}$$

D.4 Monetary and Fiscal Policy

Monetary policy follows a Taylor rule with interest rate smoothing:

$$i_t = (1-\rho_i)i + \rho_i i_{t-1} + (1-\rho_i)\phi_\pi(\log \Pi_t - \log \Pi) + s_i \varepsilon_t^i. \tag{65}$$

There is no borrowing and so the government budget constraint is:

$$G_t + \alpha\delta_1 T_t + (1-\alpha)\delta_2 T_t = \tau_t \tag{66}$$

where G_t denotes government spending on the production good.

D.5 Aggregation and Market Clearing

Aggregate labor and consumption are:

$$N_t = \alpha N_{u,t} + (1-\alpha)N_{h,t}, \tag{67}$$

$$C_t = \alpha C_{u,t} + (1-\alpha)C_{h,t}. \tag{68}$$

The resource constraint is:

$$Y_t = C_t + G_t. \tag{69}$$

D.6 Exogenous Processes

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

$$\log A_t = \rho_A \log A_{t-1} + s_A \varepsilon_t^A, \tag{70}$$

$$\log G_t = (1-\rho_G)\log G + \rho_G \log G_{t-1} + s_G \varepsilon_t^G, \tag{71}$$

$$\mathcal{T}_t = (1-\rho_h)\mathcal{T} + \rho_h \mathcal{T}_{t-1} + s_{\mathcal{T}} \varepsilon_t^{\mathcal{T}} \tag{72}$$

D.7 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: $\mathcal{T} = 0$.

D.8 Calibration

Table 5 reports the parameter values used in the simulations.[25] I keep these parameter values constant and then display results for different values of δ_1 and δ_2 in the main text. I solve for impulse responses via first-order perturbation using Dynare.

Table 5: 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 labor
α	0.6	Share of unconstrained households
ϕ	0.75	Calvo price stickiness
$\varepsilon/(\varepsilon-1)$	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_i	0.0025	Std. dev. monetary shock
$s_{\mathcal{T}}$	0.01	Std. dev. transfer shock

D.9 Extension

I consider an extension to the model where transfers instead follow a Taylor rule of the form:

$$\mathcal{T}_t = (1-\rho_{\mathcal{T}})\mathcal{T} + \rho_{\mathcal{T}}\mathcal{T}_{t-1} + (1-\rho_{\mathcal{T}})\phi_{\mathcal{T}}(\log \Pi_t - \log \Pi) + s_{\mathcal{T}}\varepsilon_t^{\mathcal{T}}. \tag{73}$$

For this exercise, I keep parameter values the same as above and set $\phi_{\mathcal{T}} = 2.0$.

[25] I follow closely the "standard" values for parameters employed in the illustrative simulations from Eric Sims - see https://sites.nd.edu/esims/files/2024/04/notes_tank_sp2024.pdf

E Fiscal Spending in Network Model

This Appendix describes a simple quantitative model to illustrate the results from Appendix C. The discussion links closely to the empirical application from Section 6.

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 in fixed next supply: $\bar{s}$. There is a representative agent in each region with Cobb-Douglas preferences:

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

where c_{ji} denotes the consumption of region i of the jth good and s_i denotes savings of region i. The government spends g_i on each region's good and raises taxes τ_i in each region subject to the budget constraint:

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

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

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

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 $A_{ij} = \alpha_{ij} * mpc_j$. 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 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 linkage between i and j (governed by α_{ij}), as well as all indirect linkages (e.g. $\alpha_{ik} * \alpha_{kj}$ 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:

$$y = \mathcal{M}^g g$$

where $\mathcal{M}^g = (I - A)^{-1} - (1/n)(I - A)^{-1} A J_n$. I assume the econometrician has access to repeated observations $t = [1, 2, ..., T]$ 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 $u_t = [u_{1,t}, ..., u_{n,t}] \sim i.i.d.\ N(0, \sigma_u I_n)$, such that:

$$\tilde{y}_t = \mathcal{M}^g g_t + u_t$$

Calibration I consider a simple parametrization of the model that satisfies rotational symmetry. I assume there are 50 regions, with each spending 60% of their income on their own good ($\forall i : \alpha_{ii} = 0.6$) and MPCs of 50% ($\forall i : \alpha_{si} = 0.5$). The proportions spent on other regions are inversely proportional to the 'distance' between each region, as if the regions were evenly spaced around a unit-circle i.e. –

$$\forall i \neq j : \alpha_{ij} = (1 - \alpha_{ii}) \cdot \frac{w_{ij}}{\sum_{k \neq i} w_{ik}}$$

where w_{ij} captures the 'distance' between each region:

$$w_{ij} = \frac{1}{\sqrt{2(1 - \cos(\theta_i - \theta_j))}},$$

and θ_i captures the 'location' of each unit on the circle:

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

In this calibration, regions satisfy rotational symmetry ensuring all diagonal elements of the $\mathcal{M}^g$ matrix are the same:

$$\beta_{i,h} = [\mathcal{M}^g]_{ii} = \beta^*$$

where, for this calibration, $\beta^* \approx 1.4$ (broadly consistent with local fiscal multipliers estimated in the literature). I consider a simple assignment rule for spending shocks which targets spending towards regions implicitly 'located' in one-half of the unit-circle - i.e.:

$$\varepsilon^g_{i,t} = \delta D_i + \eta_{i,t} = \delta \mathcal{I}\{i \in [1, ..., 25]\} + \eta_{i,t} \tag{74}$$

where $\eta_t = [\eta_{1,t}, ..., \eta_{n,t}] \sim i.i.d.\ N(0, \sigma_\eta I_n)$.

Identification I consider cross-sectional regressions of the form:

$$y_{i,t} = \beta_s \varepsilon^g_{i,t} + e_{i,t} \tag{75}$$

$$y_{i,t} = \beta_l \varepsilon^g_{i,t} + \gamma D_i + u_{i,t} \tag{76}$$

And consider consider the large-sample bias of the estimators across simulations i.e. $E[\hat{\beta} - \beta^*]$ for $T = 5000$.

Results Figure 5 shows intuition for the bias that can arise in this case and the correction. Figure 5a shows the distribution of shocks across regions for a single draw of $[\varepsilon_t^g, u_t]$. In this case, $\hat{\beta}_s$ is much larger than the true local multiplier since regions who experience high-ε_i^g simultaneously experience a greater rise in external demand for their exports since their close-trading partners also tend to be experience large shocks. Figure 5b shows the distribution of the residualised shocks (i.e. residualised with respect to D_i). Since the residualisation removes the regional clustering, $\hat{\beta}_l$ is much closer to the true local multiplier.

Figure 5: Bias-correction with regionally clustered shocks

(a) Original Shocks ($|\text{Error}| \approx 23\%$) (b) Residualised Shocks ($|\text{Error}| < 1\%$)

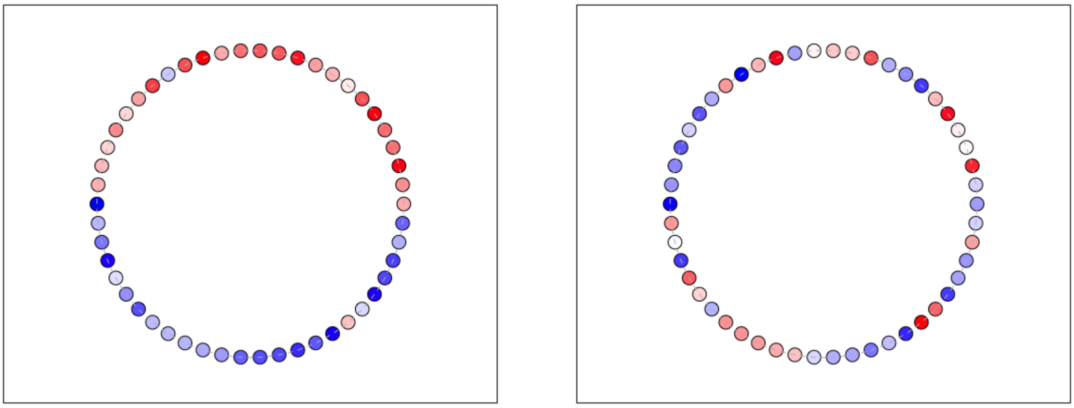

Notes: Blue denotes lower values of the shock while red denotes higher values. Left hand-side chart plots $\varepsilon_{i,t}^g$ for a single draw of the shocks and computes the error as: $\hat{\beta}_s - \beta^*$. Right hand-side chart plots $\varepsilon_{i,t}^g$ residualised with respect to D_i and computes the error as: $\hat{\beta}_l - \beta^*$.

Table 6 shows the bias of different estimators in simulations. The 'short' regression over-estimates the local multiplier by over 20%, while the 'long' regression that controls for D_i has negligible bias.

Estimator	Bias (%)
β_s	22.55%
β_l	-0.21%

Table 6: Mean bias of estimators in Monte Carlo simulation.

F Additional Empirical Results

This Appendix collects additional results from the empirical application in Sections 5.

F.1 Nakamura and Steinsson (2014)

I show additional empirical results for the empirical application from Nakamura and Steinsson (2014). Specifically, following Nakamura and Steinsson (2014), I estimate IV-regressions of the

Figure 6: Heatmaps of Residualised Instruments Across Studies

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

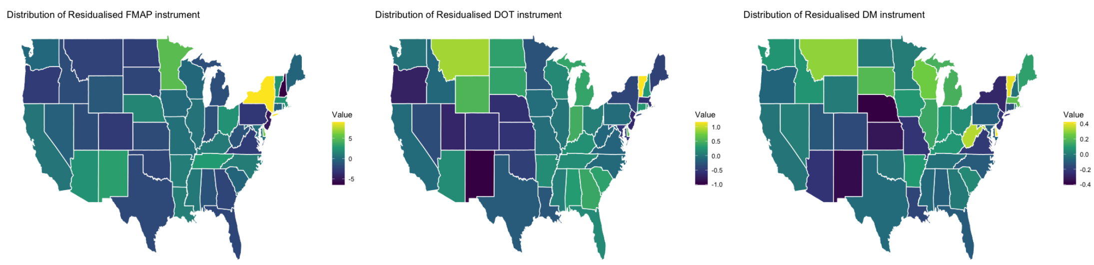

form:

$$\frac{y_{i,t} - y_{i,t-2}}{y_{i,t-2}} = \alpha_i + \delta_t + \beta \frac{g_{i,t} - g_{i,t-2}}{g_{i,t-2}} + u_{i,t} \tag{77}$$

where I instrument the right-hand side variable using total national procurement, ε_t^g interacted with a state dummy. I then additionally estimate IV-specifications of the form:

$$\frac{y_{i,t} - y_{i,t-2}}{y_{i,t-2}} = \alpha_i + \delta_t + \beta \frac{g_{i,t} - g_{i,t-2}}{g_{i,t-2}} + \lambda' a_i \varepsilon_t^g + u_{i,t} \tag{78}$$

where the control variables are intended to additionally purge any variation in the instrument of regional heterogeneity in response to objects that adjust in GE in response to the shock ε_t^g. Table 7 displays my results for different control variables a_i.

Table 7: Replication and Extension of Nakamura and Steinsson (2014) - Alternate Specification

	Dependent Variable: Output			
	(1)	(2)	(3)	(4)
Prime military contracts	1.43***	1.22***	1.22***	1.02***
	(0.36)	(0.37)	(0.33)	(0.34)
Additional Controls	None	mpc_i, $open_i$	mpc_i, inc_i	mpc_i, $open_i$, inc_i
Observations	1989	1989	1989	1989

F.2 Chodorow-Reich (2019)

I show additional empirical results for the empirical application from Chodorow-Reich (2019).

F.2.1 Clustering of residualised shocks

Table 8 reports results estimating regional clustering of the residualised shocks employed in the main text. Figure 6 plots heatmap for these residualised shocks.

Table 8: Predictability of Residualised Shocks

Dependent Variable: Residualised Shocks (η_i^g)			
	(1)	(2)	(3)
Distance-Weighted Residualised-Shocks	-0.34	0.12	-0.02
	(0.20)	(0.19)	(0.19)
Adjusted R^2	-0.02	-0.01	-0.02
Instrument	FMAP	DOT	DM
Observations	50	50	50

Note: Table displays coefficient estimates alongside (heteroskedasticity-robust) standard errors and p-values for estimates of γ from (39) with the original shocks $\varepsilon^g = [\varepsilon_1^g, ..., \varepsilon_n^g]$ replaced with the residualised shocks from (41), $\eta = [\eta_1, ..., \eta_n]$. 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.