Fiscal Policy in a Networked Economy

John Sturm Becko* Joel P. Flynn† Christina Patterson‡

Princeton Yale and NBER University of Chicago and NBER

June 20, 2026

Abstract

Fiscal stimulus policies propagate through complex and overlapping economic networks. We study their efficacy and targeting in the presence of input-output linkages, regional trade, and household heterogeneity in employment relationships, marginal propensities to consume (MPCs), and consumption baskets. Theoretically, we derive estimable formulae for fiscal multipliers and characterize how network structures determine their size. Empirically, we estimate that multipliers vary substantially across policies, so targeting is important. However, virtually all variation in multipliers stems from differences in policies' *direct* incidence onto households' MPCs. Thus, while policies' distributional effects depend on network structures, maximally expansionary fiscal policy simply targets households' MPCs.

*Princeton Department of Economics, 20 Washington Road, Princeton, NJ 08540. Email: johnsbecko@gmail.com.

†Yale University Department of Economics, 30 Hillhouse Avenue, New Haven, CT, 06511. Email: joel.flynn@yale.edu.

‡University of Chicago Booth School of Business, 5807 S Woodlawn Ave, Chicago, IL 60637. Email: christina.patterson@chicagobooth.edu.

An earlier version of this paper was entitled "Shock Propagation and the Fiscal Multiplier: the Role of Heterogeneity." We are grateful to Daron Acemoglu, George-Marios Angeletos, Martin Beraja, Olivier Blanchard, Ricardo Caballero, Arnaud Costinot, Dave Donaldson, Andrea Ferrero (Discussant), Ben Golub, Isaac Liu, Jeremy Majerowitz, Andrea Manera, Laura Murphy, Jordan Norris, Elias Papaioannou, Otis Reid, Matthew Rognlie, Karthik Sastry, Lawrence Schmidt, Alp Simsek, Ludwig Straub, Robert Townsend, Ivan Werning, and seminar participants at Harvard, the 2021 NBER Summer Institute Impulse and Propagation Mechanisms Meeting, Stony Brook, Georgetown, the Brookings Institute, the Central Bank of Chile, Oxford, the University of Southern California, the Junior Virtual Macro Conference, the 2022 ASSA Annual Meeting, the 2023 Biennial European Central Bank Conference, the 2021 European Winter Meeting of the Econometric Society, MIT Macro Lunch, and MIT Trade Tea for helpful comments. We also thank RIMS staff members at the Bureau of Economic Analysis for a helpful discussion. First posted version: April 18, 2020.

1. Introduction

Periods of widespread economic slack—when demand is weak and monetary policy is constrained—pose fundamental challenges for policymakers. In such settings, government spending and transfers become the primary tools for stimulating aggregate demand. Yet the optimal design of these policies remains contested. In particular, given a budget for designing a fiscal stimulus program: Should governments rely on broad-based transfers, narrowly targeted support for specific households, or direct purchases from particular sectors? Which of these policies generates the greatest aggregate impact during recessions, and which groups benefit the most?

These questions are complicated by the rich networks that make up present-day economies. Economic linkages through supply chains, regional trade, and heterogeneous employment and consumption relationships prevent a fiscal planner from conducting policy one household, industry, or region at a time. Rather, policymakers must consider the cascades of expenditure they set off, as expenditures in one industry and region reach not only its workers but also others in its supply chain, those at firms where workers spend their marginal income, and so on. While economists have studied these linkages theoretically (see e.g., Miyazawa, 1976; Baqaee and Farhi, 2018; Woodford, 2020; Guerrieri, Lorenzoni, Straub, and Werning, 2020), their quantitative importance and policy relevance is not well understood.

This paper provides a new, semi-structural approach that uses micro-data to quantify the role of these interconnections in shaping the aggregate effects of fiscal policy in a demand-deficient environment. The first part of our paper provides a theory of how two key policy instruments, government purchases and fiscal transfers, propagate through supply chains, employment linkages, and the directed MPCs of heterogeneous households in an environment with fixed prices and a binding zero lower bound. Although these mechanisms interact in complex ways, we show that their effects on fiscal multipliers can be decomposed into three components, each representing a deviation from the benchmark multiplier in a representative firm and agent economy.

The second part of the paper takes this decomposition to the data and documents three key empirical findings. First, fiscal multipliers in this setting vary widely across different government purchase and transfer policies, so targeting is important. Second, there are meaningful spatial and sectoral spillovers from fiscal expenditures, so understanding the network structure of the economy is essential for quantifying the distributional impacts of fiscal stimulus. Third, all observed heterogeneity in multipliers is explained by differences in the incidence of fiscal shocks onto households with varying marginal propensities to consume (MPCs). By contrast, the direction of their consumption spending plays no meaningful role. As a result, maximally expansionary fiscal policy in a widespread recession simply targets high-MPC households, as in much simpler models. This MPC targeting is not only fairly simple, but also quantitatively important: For small stimulus policies, it results in twice as much policy amplification as untargeted, GDP-proportional purchases.

We study our motivating questions in a semi-structural model of a demand-deficient economy. On the household side, we allow for heterogeneity in both the magnitude of households' MPCs and their direction toward different goods. On the firm side, we allow for many sectors and regions, linked to one another through an arbitrary input-output structure. Finally, we allow for any pattern of household employment across the various firms, generating heterogeneous household income processes.

On top of this rich micro-heterogeneity, we impose a more restrictive macro-environment, characterized by fixed prices, a binding zero lower bound, and rationing in the labor market. These features imply that households can lie off their labor supply curves and be involuntarily un(der)employed. This structure makes the model particularly well-suited for studying fiscal policy during deep recessions, when demand is weak and monetary policy is constrained. These assumptions come with limitations: they preclude relative price movements in response to targeted fiscal interventions and abstract from interactions between fiscal and monetary policy that may be important when interest rates are unconstrained (*e.g.,* Bianchi, Faccini, and Melosi, 2023; Ascari, Beck-Friis, Florio, and Gobbi, 2023). However, the key advantage of this abstraction is that it allows us to derive simple analytical results that isolate the role of household heterogeneity and production networks in shaping the income multiplier. The resulting framework is empirically tractable, allowing a transparent calibration to data on consumption patterns, employment relationships, and production linkages. Since the income multiplier we study continues to operate as one block of more complex models with price adjustment, our analysis provides a reference point that may still be relevant in richer settings.

Our first analytical result tracks these linkages in a generalized Keynesian income multiplier matrix that expresses the change in the vector of first period output across industries and regions in response to a shock to final demand, which we define as the change in final demand in response to a change in fiscal policy before incomes have been allowed to adjust. First, the input-output network translates shocks to final demand into changes in the production of each sector. Second, a rationing function captures how these changes in production translate into changes in labor income for each household. Finally, a matrix with the magnitude and direction of household MPCs across goods maps these changes in household income into changes in their demand across industries and regions. This loop repeats ad infinitum, generating our expression for the heterogeneity-adjusted multiplier matrix.

Despite this complexity, we show that the total effect of any fiscal policy on aggregate private and public consumption—or its *fiscal multiplier*—can be decomposed into three distinct effects on top of a baseline Keynesian multiplier that would exist in a model without heterogeneous agents and industrial linkages. First, the *incidence effect* captures that policies with direct incidence onto higher-MPC households change aggregate consumption by more. Second, the *bias effect* captures the additional amplification that occurs when households directly affected by the policy disproportionately direct their marginal spending toward goods produced by higher-than-average-MPC workers. Third, the *homophily effect* captures the additional amplification that occurs when high (low) MPC households direct their spending to other high (low) MPC households, for instance due to geographic concentration. The magnitudes of these three effects are a function of the underlying structure of the economy: we show that economies with higher MPCs and/or labor shares have higher multipliers and economies with less connected input-output networks have more dispersed multipliers.

In order to understand which features of the economy contribute to shock amplification and to gauge the quantitative importance of targeting fiscal policy within this setting, the second part of the paper takes our model to the data. We combine several public-use datasets describing 50 US states (plus DC), 55 sectors, and 80 demographic groups to estimate three key empirical objects: the regional input-output matrix describing the input-use requirements of every industry-region pair; the rationing matrix describing how much each demographic-region pair's income changes in response to a one dollar change in production of each industry-region pair; and the directed MPC matrix describing how much each demographic-region pair consumes from each industry-region pair on the margin.

By combining these estimated matrices with our derived expressions for the generalized income multiplier, we uncover wide variation in government purchases and transfers multipliers depending on where in the economy a fiscal shock originates. Our estimates imply that a dollar of purchases from the sector-region with the highest multiplier leads to twice as much amplification as spending that same dollar on a GDP-proportional basket of goods. However, perhaps surprisingly, we find that virtually all of the difference in multipliers is driven by differences in the incidence of the shock onto households with higher or lower MPCs, and that households' patterns of directed consumption across sectors and regions do not contribute meaningfully to multipliers, implying that the bias and homophily effects are close to zero. We find that the large heterogeneity in shock incidence is driven primarily by dispersion in MPCs in the population and the sorting of workers of different types across sectors and regions. Despite the fact that only the incidence of a shock matters for its effect on *aggregate* consumption, all dimensions of heterogeneity shape the *distribution* of this consumption across households employed in different industries and states. In particular, the structure of our model allows us to quantify the extent to which government purchases and transfers directed toward any one US state spill over across state lines.

These empirical results have sharp implications for a policymaker who aims to maximize aggregate consumption. Our empirical finding on the irrelevance of directed household consumption for multipliers (i.e., the lack of bias and homophily effects) implies that the total multiplier of any fiscal shock depends *only* on its incidence onto households with higher or lower MPC. Thus, MPC targeting is optimal.[1] We illustrate the importance of this targeting by replicating a CARES-like

[1]For fiscal transfers, this amounts to transferring money to those workers with the highest MPCs. Targeting government purchases is more complicated, as the planner hopes to allocate spending so as to affect the workers

transfer policy in the model: Our estimates suggest that government transfers of one thousand dollars to each employed worker would increase aggregate consumption by 68 cents per dollar spent, whereas transfers of two thousand dollars to each worker with above-median MPC would increase consumption by 95 cents per dollar spent.

Related Literature This paper integrates many dimensions of heterogeneity, each studied independently throughout the literature. This allows us to quantify which features of advanced economies matter for which macroeconomic questions.

On the one hand, suppose that a researcher is interested in understanding the response of *aggregate* consumption to fiscal policy. Our results suggest that heterogeneous incidence is the only important margin to consider. Thus, our results echo recent work that stresses amplification of shocks that load more heavily onto households with higher-than-average MPCs (Werning, 2015; Kaplan, Moll, and Violante, 2018; Auclert, 2019; Patterson, 2023; Bilbiie, 2019), and the role of input-output linkages in distributing the incidence of shocks on various industries throughout the economy (Long and Plosser, 1987; Gabaix, 2011; Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi, 2012; Baqaee and Farhi, 2019; Rubbo, 2019; Bigio and La'O, 2020). More recently, Andersen, Vestergaard, Huber, Johannesen, and Straub (2022) use disaggregated economic accounts for Denmark to highlight the role of spending intensity—the share of consumer expenditures that fall on domestic producers—as a determinant of multiplier variation. While our U.S.-based estimates imply that such bias and homophily effects are quantitatively negligible, their results suggest that openness and spending allocations can shape fiscal transmission in other settings.

On the other hand, we find that accounting for regional trade and the direction of consumption (in particular its within-region bias), factors recently emphasized by Farhi and Werning (2017), Caliendo, Parro, Rossi-Hansberg, and Sarte (2018), and Dupor, Karabarbounis, Kudlyak, and Mehkari (2018), does not contribute to the aggregate impact of shocks in a quantitatively meaningful way. However, such features are critical for understanding the distribution of the impact of shocks across households, industries, and space.

Our reduced-form approach to integrating these elements is similar methodologically to Auclert, Rognlie, and Straub (2018), who focus on intertemporal aspects of fiscal stimulus and its financing, while we study the role of heterogeneity within a single period *conditional* on how fiscal policy is financed. This emphasis on cross-sectional heterogeneity echoes earlier work in the regional accounting literature, for which we provide a formal microfoundation (Miyazawa, 1976).

The theoretical part of our paper relates most closely to Baqaee (2015). As we do in this paper, Baqaee emphasizes that shocks to an industry affect not only the factors employed in that industry but also those used in producing its inputs, motivating a "network adjustment" to the labor share of each industry. In more recent work, Baqaee and Farhi (2018) and Bouakez,

with the highest MPCs, which requires knowledge of the input-output network (as in Baqaee, 2015) and the labor rationing process, both of which shape how changes in demand affect the income of workers.

Rachedi, and Santoro (2023) develop rich macroeconomic models featuring these channels as well as endogenous prices and markups. By focusing on rigid prices, we are able to precisely characterize—as well as empirically assess—the channels through which economic linkages affect aggregate shock propagation and how these matter for optimal stimulus policy.[2]

Lastly, this paper also adds to a large empirical literature estimating multipliers from fiscal shocks. Our structural estimates complement reduced-form empirical estimates of open-economy multipliers—we calibrate an aggregate purchases multiplier of 1.28, which is somewhat smaller than, but within the established confidence intervals of, those in Ramey (2011), Nakamura and Steinsson (2014), Chodorow-Reich (2019) and Corbi, Papaioannou, and Surico (2019). We moreover highlight that depending on the sector or state receiving stimulus, purchases multipliers range from roughly 1.1 to 1.6, offering a lens through which disparate estimates in this literature can be rationalized. Recent empirical work has also explored spatial spillovers arising from fiscal policy (see e.g. Feyrer, Mansur, and Sacerdote, 2017; Cox, Müller, Pasten, Schoenle, and Weber, 2019; Auerbach, Gorodnichenko, and Murphy, 2020). Consistent with this evidence, we estimate that, on average, 38% of the total change in value added resulting from a government purchase in one state occurs in other states. Moreover, our structural approach uncovers the distinct channels through which these spillovers operate.

2. Model

To understand the propagation of fiscal policies, we build a semi-structural model with labor rationing. In the model, a continuum of heterogeneous households interact in a competitive, multi-sector, multi-region economy over two periods. Our main assumption is that due to sticky wages and a binding ZLB, labor markets cannot clear on price. Instead, labor is rationed to involuntarily underemployed households in the first period rather than freely supplied at a market-clearing wage.[3] Apart from our focus on the rationing of deficient labor demand, we remain agnostic to much of the economy's structure. We allow for a rich class of borrowing-constrained and even non-optimizing households, a general constant-returns-to-scale input-output structure, and flexible labor rationing protocols. Our model is rich enough to capture many dimensions of household, industrial, and regional heterogeneity, but sufficiently tractable to deliver equations that we later

[2]Our paper relates closely to several concurrent papers, motivated by the Covid-19 pandemic, that explore the effects of fiscal and monetary policy when shocks are heterogeneous across sectors. Woodford (2020) shares the theoretical insight that transfers multipliers vary depending on targeting. Guerrieri et al. (2020) show that the effectiveness of fiscal stimulus in response to supply shocks depends critically on the direction of marginal spending, meaning that government spending is less effective when the sectors employing the higher-MPC workers are shut down. Similarly, Baqaee and Farhi (2020) find, in a model with networks and nominal rigidities, that disparate demand and supply shocks blunt the power of aggregate demand stimulus policies. We differ from these recent papers in our precise decomposition of deviations from the Keynesian case and our focus on calibrating sufficient statistics in order to demonstrate which features of the economy are empirically important for shaping the distribution of multipliers.

[3]In Appendix B, we provide a more explicit microfoundation for rationing equilibrium.

bring directly to the data.

2.1. Model Primitives and Rationing Equilibrium

The economy consists of a unit measure of households with types n drawn from a finite set N, each of mass $\mu_n > 0$, and a finite number of firms $i \in I$. N can capture demographic factors such as age, sex, and race as well as location of residence, and I can capture both the sort of good and its location. Households and firms respectively supply and demand a homogeneous labor factor in order to produce goods over two periods $t \in \{1, 2\}$.

Fixed Wages and Interest Rates. First period nominal wages w^1 are fixed (perhaps due to binding downward wage rigidity) as are expectations of second period nominal wages. Moreover, nominal interest rates ι are fixed (perhaps) due to a binding zero-lower-bound. Thus, the real interest rate, $1 + r = \frac{w^1}{w^2}(1 + \iota)$, is also fixed. Since they are fixed, we normalize both first and second period wages to one, i.e., $w^t \equiv 1$. We denote the vector of prices in period t by $p^t = \{p_i^t\}_{i \in I}$.

Producers. Given prices, a representative producer of each good i in each period t demands labor L_i^t and a vector of inputs $X_i^t = \{X_{ij}^t\}_{j \in I}$ to maximize profits with a continuous and constant returns-to-scale (CRS) production function F_i^t.

$$(X_i^t, L_i^t) \in \arg\max_{X,L} \, p_i^t F_i^t(X, L) - p^t X - L \tag{1}$$

Rationing. Because wages are fixed, the first-period labor market cannot clear on prices. Instead, labor is rationed to households and firms by a non-price mechanism that determines *realized* labor supplies and demands as a function of *preferred* supplies and demands. While we are intentionally agnostic about the fine details of this "rationing protocol," we focus on the case of *deficient demand* for labor. We model this by assuming that firms' labor demands are unconstrained, whereas households take their first-period labor supply as given.

Formally, each household n's labor supply in the first period (ℓ_n^1) is determined as a function of the vector of firm-specific labor demands $L^1 = \{L_i^1\}_{i \in I}$:

$$\mu_n \ell_n^1 = R_n^1(L^1) \tag{2}$$

Although we remain agnostic to the details of labor assignment, we assume the rationing function assigns labor supply so as to exactly meet labor demand, i.e., $\sum_{n \in N} R_n^1(L^1) = \sum_{i \in I} L_i^1$. We offer a detailed microfoundation of this rationing process in Appendix B.

Households. Taking not only prices but also first-period labor supply as given, each household n chooses its second-period labor supply ℓ_n^2 and vectors of first- and second-period consumption $c_n^t = \{c_{ni}^t\}_{i \in I}$. We assume that each household of type n has a continuous and additively separable utility function over consumption u_n^t and labor v_n^t with discount factor $\beta_n \geqslant 0$ and potentially faces a borrowing constraint with a minimum savings level of $\underline{s}_n$. Thus, households' Marshallian consumption demand and labor supply functions solve:

$$
\begin{aligned}
(\ell_n^2, c_n^1, c_n^2) \in \arg\max_{\ell^2, c^1, c^2} \; & u_n^1(c^1) - v_n^1(\ell_n^1) + \beta_n \left(u_n^2(c^2) - v_n^2(\ell^2) \right) \\
\text{s.t. } & p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell_n^1 + \frac{\ell^2}{1+r} \\
& \ell_n^1 - p^1 c^1 - \tau_n^1 \geqslant \underline{s}_n
\end{aligned} \tag{3}
$$

Government. In addition to levying lump-sum taxes τ_n^t on households, the government purchases G_i^t units of each good $i \in I$ subject to running a balanced budget over the two periods.

$$
\sum_{n \in N} \mu_n \left(\tau_n^1 + \frac{\tau_n^2}{1+r} \right) = p^1 G^1 + \frac{p^2 G^2}{1+r} \tag{4}
$$

Market Clearing. Finally, goods markets and the second period labor market clear.

$$
Q_i^t = F_i^t(X_i^t, L_i^t) = \sum_{j \in I} X_{ji}^t + \sum_{n \in N} \mu_n c_{ni}^t + G_i^t, \qquad \sum_{i \in I} L_i^2 = \sum_{n \in N} \mu_n \ell_n^2 \tag{5}
$$

Rationing Equilibrium. A rationing equilibrium is therefore defined as follows:

Definition 1. *Given a profile of government spending and transfers $\{\tau_n^t, G_i^t\}_{t \in \{0,1\}, n \in N, i \in I}$ that satisfies (4), a rationing equilibrium is a profile $\{p_i^t, \ell_n^t, c_n^t, L_i^t, Q_i^t, X_{ij}^t\}_{t \in \{0,1\}, n \in N, i,j \in I}$ that satisfies conditions (1), (2), (3), and (5).*

In Appendix C.1, we provide mild technical assumptions under which a rationing equilibrium exists and prices do not depend on transfers τ or government spending G. This result is an application of the classical non-substitution theorem, which holds in our setting because firms have constant returns to scale (CRS) technologies and produce using a single homogeneous factor (labor). Since firms can scale production up or down without changing marginal costs, changes in the composition of final demand do not affect their relative input use or prices. Because wages in both periods are fixed—due to our assumptions of downward nominal wage rigidity and a binding zero lower bound—and prices are uniquely determined by wages, prices are likewise fixed: they are fully determined by technological primitives, and unaffected by the purchases and transfer policies we study. We therefore normalize the units of each good so that these exogenous prices satisfy $p_i^t = 1$ for all $t \in 1, 2$, $i \in I$, and express consumption in expenditure terms throughout.

The exogeneity of prices also allows us to express households' Marshallian consumption demands as functions of only taxes and their rationed first-period labor supply. Moreover, our assumption that preferences are additively separable implies that first-period taxes and first-period labor supply both affect demand only through households' first-period incomes $h_n^1 \equiv \ell_n^1 - \tau_n^1$. We therefore let $c_n^1(h_n^1, \tau_n^2)$ denote household n's time-1 Marshallian consumption demand.[4]

Finally, several of our results describe the impact of fiscal shocks on the vector of sector-level final output,

$$Y_i^t \equiv Q_i^t - \sum_{j \in I} X_{ji}^t \tag{6}$$

where aggregate output is simply given by:

$$\mathrm{GDP}^t = \mathbb{1}^T Y_i^t \tag{7}$$

Extensions: Imperfect Competition and Dynamics. In Appendices D.1 and D.2, we respectively allow for imperfect competition and arbitrarily many time periods and show that suitably modified versions of all of our main results continue to hold.

2.2. *Interpreting the Model*

Our model is intentionally designed to study fiscal policy in environments characterized by widespread weak demand and constrained monetary policy—settings in which downward nominal wage rigidities bind and price adjustment mechanisms are therefore muted. We adopt a fixed-wage framework with rationing as a tractable approximation to this environment, building on a long tradition in Keynesian macroeconomics that has seen a recent resurgence (Barro, 2025).

In addition to its explicit rigidity assumptions, our baseline model also implicitly rules out relative wage adjustments by assuming that workers supply a single, homogeneous labor factor.[5] While this simplifying assumption facilitates analytical tractability, we recognize that it abstracts from meaningful sectoral variation in labor markets and wage dynamics. To partially address this limitation, we show in Appendix B.3 that our model is equivalent to one with many labor factors differentiated by state and sector, so long as there is downward pressure on all wages, all wages are downwardly rigid, and relative wage-inflation expectations are sticky.

Given that neither intratemporal nor intertemporal prices can adjust, markets must clear

[4] For all of our positive analysis, it is unimportant that households' demands are generated by household optimization. All of our analysis extends to consumption demand and labor supply functions that satisfy:

$$c_n^t = c_n^t(h_n^1, \tau_n^2) \qquad \ell_n^2 = \ell_n^2(h_n^1, \tau_n^2) \qquad \mathbb{1}^T c_n^1 + \frac{\mathbb{1}^T c_n^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell_n^1 + \frac{\ell_n^2}{1+r}$$

[5] By the non-substitution theorem, prices are pinned down by wages, so ruling out relative wage adjustments in turn rules out intra-temporal price adjustments.

through a form of rationing. Our emphasis on labor market rationing rather than product market rationing reflects the empirical reality of widespread underemployment during recessions and follows a rich intellectual tradition in Keynesian macroeconomics stretching back to Patinkin (1949), Clower (1965) and Barro and Grossman (1971). Indeed, the key idea that price rigidities or other frictions may cause a household to lie off its labor supply curve is a staple of many modern macroeconomic approaches to understanding involuntary unemployment and the business cycle, with our exact formulation via a rationing function being closest to that employed by Werning (2015).[6]

Our focus on rationing at fixed prices rules out other, potentially important dynamics related to price adjustment—a force emphasized by recent work including Proebsting (2022), Bouakez, Rachedi, and Santoro (2020), Barattieri, Cacciatore, and Traum (2023), and Cox, Müller, Pasten, Schoenle, and Weber (2024). We view this focus as most appropriate during deep, demand-driven recessions, where interest rates are at a zero-lower bound and all wages are downwardly rigid. To the extent this is an imperfect approximation, price and interest rate adjustments would manifest in our framework as additional, endogenous shocks to demand. Moreover, we focus on marginal changes in policy within this setting, and therefore implicitly assume that the zero lower bound continues to bind when policy changes.

Finally, we emphasize that even in models in which prices and interest rates are partially flexible, the general forces that we uncover and quantify are still operative. In such contexts, our analysis should be understood as providing a "partial-equilibrium" analysis of the income-and-spending block of more general models that feature wage, price, and interest rate adjustments. Insofar as channels are (or are not) quantitatively important in the spending-and-income block, our work may serve as a useful guide as to which features are necessary to include in richer models in future research.

3. The Generalized Multiplier

Within this setting, we explore the impact of shocks to government purchases and transfers. Our goal is to derive an expression for the generalized Keynesian income multiplier that combines the direct effect of shocks on spending, which we define to be the effects of the shock on spending before incomes have adjusted, with their indirect effects, which arise as worker incomes adjust in response to the shock. This is both of independent interest for understanding shock propagation and also a key step in understanding the efficacy of fiscal policy. Importantly, we derive a representation of the generalized multiplier in terms of sufficient statistics that clarify the role of

[6]Among the advantages of our reduced-form approach to rationing and household consumption is that it can accommodate regional migration driven by local underemployment. On the income side of the model, the fact that employment is demand-determined implies that the same total income is rationed to each household type in each region regardless of the size or composition of the demographic group in that region. On the consumption side, migration manifests as a directed MPC away from goods in one region and toward goods in another.

networks in the macroeconomy, and that we later take to the data to study implications for the design of fiscal policy. Proofs of main results are in Appendix A.

3.1. The Generalized Fiscal Multiplier in a Networked Economy

Our main results express the economy's responses to fiscal shocks, breaking them into the direct effect of the shock on goods demand before incomes have been allowed to adjust and the indirect effect that follows as incomes adjust. The first-period direct effect of a shock to government purchases or transfers is given by:

$$\partial Y^1 = dG^1 - \boldsymbol{C^1_{h^1}} \boldsymbol{\mu} d\tau^1 + \boldsymbol{C^1_{\tau^2}} \boldsymbol{\mu} d\tau^2 \tag{8}$$

where here ∂Y^1 and dG^1 are length-I vectors, $\boldsymbol{\mu}$ is a $N \times N$-dimensional diagonal matrix with entries μ_n, and $\boldsymbol{C^1_{h^1}}$ and $\boldsymbol{C^1_{\tau^2}}$ are $I \times N$-dimensional matrices of directed MPCs out of first period net income ($h^1_n \equiv \ell^1_n - \tau^1_n$) and second period transfers, with (i, n) entries corresponding respectively to $\frac{\partial}{\partial h^1_n} c^1_{ni}(h^1_n, \tau^2_n)$ and $\frac{\partial}{\partial \tau^2_n} c^1_{ni}(h^1_n, \tau^2_n)$. These and all other partial derivatives throughout the analysis are assumed to exist and be continuous, and are evaluated at an initial equilibrium before fiscal policy changes.

Note that in defining the direct effect of the policy in this way and focusing our analysis on the indirect effects that these direct effects produce, we abstract from any question of how these policies are financed. Different financing structures would imply different direct effects—for example, under strict Ricardian equivalence the direct effect of the policy could be zero. Much of our analysis will take second period policies as given and focus on changes to first-period policy. In this way, we will take the financing of the policy as given and explore the implications for aggregate consumption and income of alternative ways of directing that spending.

Our first proposition captures the way that an *income multiplier* amplifies the direct effect of the shock ∂Y^1 to generate the complete change in the vector of first-period values added dY^1. To state this result, we let $\boldsymbol{\widehat{X}^1}$ and $\boldsymbol{\widehat{L}^1}$ denote the $I \times I$-dimensional unit-production input-output and labor demand matrices, respectively, in the first-period.[7] Similarly, we denote by $\boldsymbol{\widehat{C}^1}$ the $I \times N$ matrix of marginal propensities to consume each good out of labor income *per unit of expenditure* and we denote by $\boldsymbol{m}$ the diagonal matrix of households' *total MPCs* out of labor income, so that $\boldsymbol{\widehat{C}^1 m} = \boldsymbol{C^1_{h^1}}$. Finally, we denote by $\boldsymbol{R^1_{L^1}}$ the $N \times I$ marginal rationing matrix, with typical element

[7]More formally, define the unit input demands for any firm i as those that solve the following program:

$$\left(\widehat{X}^t_i, \widehat{L}^t_i\right) = \arg \min_{(X^t_i, L^t_i) \text{ s.t. } F^t_i(X^t_i, L^t_i) \geqslant 1} \mathbb{1}^T X^t_i + L^t_i$$

we show that these demands exist and are unique in Proposition 4 and Corollary 1 in the Appendix. $\boldsymbol{\widehat{X}^1}$ is the $I \times I$-dimensional matrix with i^{th} column equal to $\widehat{X}^1_i$; $\boldsymbol{\widehat{L}^1}$ is the diagonal $I \times I$-dimensional matrix with (i, i) entry $\widehat{L}^1_i$.

$\frac{\partial}{\partial L_i^1} R_n^1(L^1)$ corresponding to household type n's change in labor income with respect to a small change in labor demand of i.

Throughout the rest of the paper, we assume that the Leontief-inverse matrix $(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}$ exists.[8] Moreover, in analogy to the assumption that the aggregate MPC is less than one in the simple Keynesian multiplier, we assume the moduli of $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}$ and $\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}$ are less than one, which guarantees that the generalized multiplier is well-defined.[9]

Proposition 1. *Given any rationing equilibrium, the local change in equilibrium first period final output dY^1 following a fiscal shock with direct effect on first-period final output ∂Y^1 is given by:*[10]

$$dY^1 = \left(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}}\ \boldsymbol{m}\ \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}\ \widehat{\boldsymbol{L}}^{\boldsymbol{1}}\ \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\right)^{-1} \partial Y^1 \tag{9}$$

This is the key formula of the paper and can be understood as a generalization of the traditional Keynesian multiplier $(1-MPC)^{-1}$ to the case of input-output networks, heterogeneous households, and arbitrary firm-household employment linkages. The term

$$\underbrace{\widehat{\boldsymbol{C}}^{\boldsymbol{1}}}_{I\times N}\ \underbrace{\boldsymbol{m}}_{N\times N}\ \underbrace{\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}}_{N\times I}\ \underbrace{\widehat{\boldsymbol{L}}^{\boldsymbol{1}}}_{I\times I}\ \left(\boldsymbol{I} - \underbrace{\widehat{\boldsymbol{X}}^{\boldsymbol{1}}}_{I\times I}\right)^{-1} \tag{10}$$

is the analog of the MPC in the traditional multiplier formula. Below each term, we have noted its dimensions, where N is the set of household types and I is the set of goods or firms.

The formula is best understood when read from the right to the left. Following a demand shock to firms (∂Y^1), the term $\left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}$ maps changes in final demand to changes in production via the input-output network. Having pinned down the change in required production, $\widehat{\boldsymbol{L}}^{\boldsymbol{1}}$ maps these to changes in firms' demand for labor. Next, the marginal rationing matrix $\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}$ maps these changes in labor demands to changes in each household's income. The MPC matrix $\boldsymbol{m}$ maps these changes in income into changes in spending. Finally, the spending direction matrix $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}$ maps changes in each household's consumption spending to changes in aggregate consumption of each good. The final generalized multiplier is the Leontief inverse of this object as this loop repeats *ad infinitum.*[11]

In terms of terminology, we refer to the matrix of Equation 9 that maps vectors of direct

[8] In Appendix C.1, we provide sufficient conditions for this to be the case.

[9] We later verify these assumptions empirically. Also note that the modulus is guaranteed to be less than one whenever all households have MPC less than one.

[10] Of course this result presupposes the existence of an equilibrium. In Appendix C.1, we provide basic primitive conditions under which rationing equilibrium is guaranteed to exist.

[11] This same multiplier expression appears in the regional economics literature on social accounting matrices, dating back to Miyazawa (1976). To our knowledge, our result provides the first fully-microfounded justification of this formula, which receives widespread use in the regional economics literature and applied work to compute purchases multipliers (such as the BEA's RIMS II system). The connection to the social accounting literature motivates yet another way to understand the multiplier formula at the zero lower bound. One can think of households as though they are simply additional nodes in the production network, with the restriction that they exchange goods and labor only with firm nodes, and not with other households.

changes in final output to vectors of total changes in final output as the *generalized multiplier*. By contrast, we refer to the effect on aggregate output of a given fiscal policy per unit expenditure simply as its *fiscal multiplier* or *multiplier* for short. Whereas the generalized multiplier is a matrix, the fiscal multiplier of any particular fiscal policy is a scalar. Finally, we refer to the range of fiscal multipliers across the set of all fiscal policies, as the *distribution of multipliers*.

The crucial difference between the generalized multiplier and the traditional Keynesian multiplier is that the structure of production, employment and consumption matters. First, it is important whether shocks load onto low- or high-MPC households, as studied by Patterson (2023). Moreover, the *interaction* between the input-output network and the directed consumption network matters: multipliers are largest when it is not only direct shocks but also higher order responses that load onto high-MPC households, due to those households spending their marginal dollars at firms that hire high-MPC workers or at firms that buy inputs from firms hiring high-MPC workers, and so forth.

In Appendix C.2, we provide formal comparative statics that demonstrate how the distribution of multipliers in the economy depends on the underlying structure of the economy (Proposition 6). Specifically, we show that the multipliers are higher for any initial shock when MPCs rise for all individuals or, at any firm, the share of income rationed to some zero-MPC household decreases and the share rationed to all other households increases (Corollary 3). Less obviously, we also provide conditions under which richer IO linkages between firms contract the distribution of multipliers, i.e., the maximum multiplier falls and minimum multiplier rises (Proposition 7).[12] This occurs because a more connected input-output matrix implies that shocks to any given industry affect a wider array of households, effectively distributing the shock to households with a more diverse set of MPCs.

3.2. Decomposing the Role of Heterogeneity

While the comparative statics discussed above study individual blocks of the model in isolation, the many dimensions of heterogeneity in household and firm characteristics and interconnections also *interact* to produce the generalized multiplier in Proposition 1. In this section, we explain how these many dimensions can be understood through three key channels that lead to greater or lesser amplification relative to the basic Keynesian case. For simplicity, we focus on the case of changes in only *first-period* government purchases and transfers. We consider the general case in Appendix A.2.

Toward decomposing the role of heterogeneity, we define the aggregate spending-to-income network

$$\boldsymbol{\mathcal{G}} \equiv \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1} \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \tag{11}$$

[12] Formally, we compare an economy with no IO linkages to one with an arbitrary IO matrix in a case where, consistent with our later empirical findings, the direction of consumption is irrelevant for amplification.

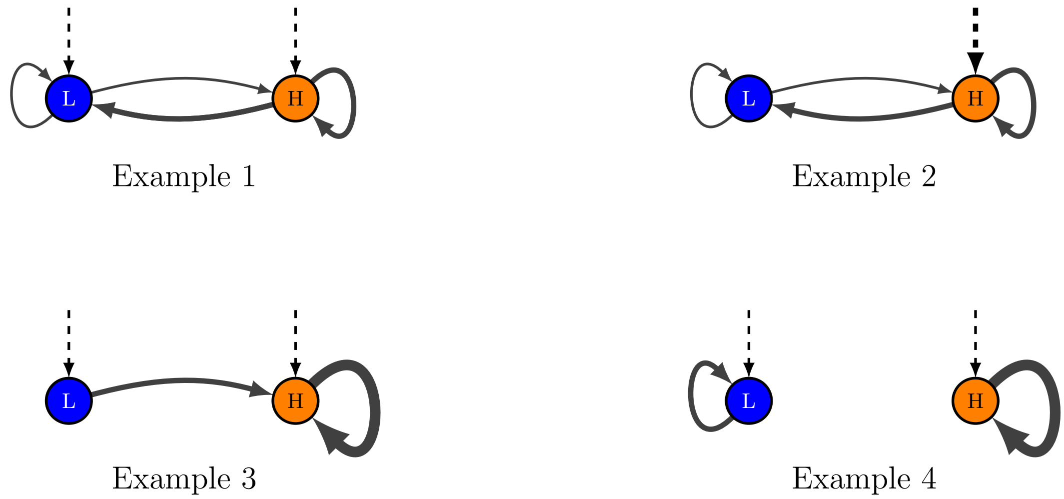

Fig. 1. H and L are nodes corresponding to the households with MPCs m_H and m_L, respectively. The thickness of solid arrows captures the magnitude of households' marginal spending. The thickness of dashed arrows capture the magnitude of the direct incidence of policies onto household income. Example 1: "Neutral" shock and spending-to-income network. Example 2: Shock directed toward high-MPC household ("incidence"). Example 3: Typical HH's marginal spending directed toward HHs with higher than own MPC ("bias"). Example 4: Each HH directs marginal spending toward HHs with same MPC ("homophily").

as the $N \times N$ map from an additional dollar of spending by one household type to the vector of income changes it generates for each other household type. Since every dollar spent eventually becomes income, every column of $\boldsymbol{\mathcal{G}}$ sums to one. Second, we define

$$\partial h^1 \equiv \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \hat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} dG^1 - \boldsymbol{\mu} d\tau^1 \tag{12}$$

as a first period policy shock's $N \times 1$ direct incidence on household first period net incomes.

Before characterizing them formally in Proposition 2, we first illustrate the mechanisms through which household heterogeneity and the structure of the spending-to-income network affect shock propagation with a series of examples. In each of the four examples below, there are two households: one with low MPC $m_L = 0.1$ and one with high MPC $m_H = 0.5$. What differs between the examples is the incidence ∂h^1 that a shock has onto the respective incomes of these households and the structure of their economic interactions through the spending-to-income network $\boldsymbol{\mathcal{G}}$.

Our first example illustrates a neutral case in which—despite the presence of heterogeneous households—the propagation of a shock is "as if" the economy had a single household with MPC $\overline{m} = \frac{m_L + m_H}{2}$. In the example, the shock has incidence $\frac{1}{2}$ on each household, and each household divides its marginal spending equally between itself and the other household (see the top-left panel of Figure 1). As a result, the incidence of spending induced by the income earned in meeting the initial demand shock is exactly $\overline{m}$ times the shock's incidence for each household; similarly for spending induced by income earned in meeting this secondary demand, and so on. Thus, the multiplier is given by $\frac{1}{1-\overline{m}} = 1.43$.

In the second example, the structure of the economy is unchanged, but the shock ∂h^1 is directed entirely to the high-MPC household (see the top-right panel of Figure 1). As a result of this differential incidence, the initial change in household income induces a greater increase in spending. However, since the high-MPC household's divides its spending evenly between household types, subsequent "rounds" of spending still propagate at the baseline, Keynesian multiplier. In this case, the multiplier is given by $1 + \frac{m_H}{1-\overline{m}} = 1.71$, so shocks feature 65% more amplification than the baseline. Thus, a transfer solely to the high-MPC household rather than a uniform transfer of the same size is much more effective at increasing output.

In the third and fourth examples, we return to the neutral income incidence $\partial h^1 = \left(\frac{1}{2}, \frac{1}{2}\right)$ of the first example and instead consider changes to the spending-to-income network $\boldsymbol{\mathcal{G}}$. In the third example, each household directs all of its marginal spending to the sector employing the high-MPC household (see the bottom-left panel of Figure 1). Unsurprisingly, this generates higher amplification, as households' induced spending all propagates at a multiplier corresponding only to the higher-MPC households' MPC. In particular the multiplier is given by $1 + \frac{\overline{m}}{1-m_H} = 1.60$, generating 40% more amplification than the neutral baseline.

In the final example, each household directs all of its marginal spending toward itself (see the bottom-right panel of Figure 1). In this case, each household's share of the shock incidence propagates separately, at $\frac{1}{1-\text{MPC}}$ with that household's MPC. Mathematically speaking, since $\frac{1}{1-MPC}$ is convex in MPC, two isolated economies generate higher average multiplier of $\frac{1}{2}\left(\frac{1}{1-m_L} + \frac{1}{1-m_H}\right) = 1.56$, i.e., they generate 30% more amplification than the integrated economy. Intuitively, since the high-MPC household spends more of its increase in income, it increases GDP more by directing its spending toward its own, high, MPC than the low-MPC household decreases GDP by directing its spending toward its own, low, MPC.

The second, third, and fourth examples illustrate three distinct channels by which the characteristics of and connections between heterogeneous households affect amplification, relative to the representative agent Keynesian benchmark. First, one must account for the *incidence* of a shock onto households of higher or lower MPC. Second, the multiplier is higher when shocked households' marginal spending is *biased* toward households with higher MPCs than the benchmark average MPC. Third, *homophily* in the spending network—in the form of correlation between shocked households' MPCs and MPCs of households on which they spend—also generates amplification.

Proposition 2 establishes that these three channels exactly capture the deviations of shock amplification away from the Keynesian baseline, to second order in MPCs.[13] We define the baseline Keynesian multiplier as the multiplier in an hypothetical economy where households direct their spending across sectors proportional to GDP, following a first-period demand shock that is itself proportional to GDP. We compare this baseline to the multiplier in a general economy, following

[13] We provide an exact decomposition in terms of Bonacich centralities of $\boldsymbol{\mathcal{G}}$ in Appendix A.2. The approximation in Proposition 2 is not only an extremely tight approximation in practice but also provides a simple intuition for the income multiplier as well as a helpful basis for the empirical analysis in Section 5.

a fiscal policy shock with arbitrary incidence across households ∂h^1.[14]

Proposition 2. *The total change in first-period GDP due to a change in first-period fiscal policy with income incidence ∂h^1 that sums to one can be approximated as:*

$$\begin{aligned} dGDP^1 =& \mathbb{1}^T dG^1 + \frac{1}{1-\mathbb{E}_{h*}[m_n]}\Bigg(\underbrace{\mathbb{E}_{h*}[m_n]}_{RA\ Keynesian\ effect} + \underbrace{\mathbb{E}_{\partial h^1}[m_n] - \mathbb{E}_{h*}[m_n]}_{Incidence\ effect} \\ &+ \underbrace{\mathbb{E}_{\partial h^1}[m_n]\left(\mathbb{E}_{\partial h^1}[m_n^{next}] - \mathbb{E}_{h*}[m_n]\right)}_{Biased\ spending\ direction\ effect} + \underbrace{\mathbb{C}ov_{\partial h^1}[m_n, m_n^{next}]}_{Homophily\ effect} \Bigg) + O^3(|m|) \end{aligned} \tag{13}$$

where h^* weights households by the income change they experience following a demand shock proportional to aggregate GDP and $m_n^{next} = \left(m^T \boldsymbol{\mathcal{G}}\right)_n$ is the average MPC of households who receive as income n's marginal dollar of spending.[15]

In Appendix C.3, we discuss conditions under which the incidence, bias, and homophily effects vanish. This case occurs when all firms in the economy employ workers at the margin who have the same average MPC as each other. In this case, the income multiplier for any shock is simply the Keynesian multiplier evaluated at this average MPC (Proposition 8). Note that even in this case, the relevant average MPC need not equal either the population average MPC or the income-weighted average MPC of the population; this is the case only when each firm's marginal employees have that MPC.

Clearly, the conditions required to eliminate the incidence, bias, and homophily effects are knife-edge. In general, the distribution of shocks does affect aggregate responses, and the IO and directed consumption networks affect both the size and direction of these responses. Moreover, one might reasonably expect that each effect formalized by Proposition 2 is empirically relevant based on existing research. Indeed there is wide variation in household MPCs and spending and transfer policies may be disproportionately directed toward certain sectors or households (Lewis, Melcangi, and Pilossoph, 2019; Cox et al., 2019). Meanwhile, Patterson (2023) documents that higher-MPC workers are more exposed to aggregate fluctuation, suggesting an aggregate bias effect in the spending-to-income network. Additionally, Hubmer (2019) documents that higher-income households tend to consume more labor intensive goods, while at the same time a growing regional literature emphasizes that much of consumption is done locally while regions are heterogeneous in both income and wealth levels—both suggesting a sizable role for the (anti-)homophily effect. In the following sections we assess each channel empirically, finding—contrary to the observations above—that only the incidence effect is quantitatively significant.

[14] We also show in Appendix A.2 that this proposition holds for any choice of h^*, although the precise value of the error term will depend on this choice. Our later empirical analysis uses the value of h^* described here.

[15] For any N-length vectors z and x, $\mathbb{E}_z[x_n]$ denotes the average of x_n across household types, weighted by z_n; similarly for $\mathbb{C}$ov.

4. Data and Estimation Methodology

We have so far derived a simple, sufficient-statistics expression for a generalized income multiplier for fiscal shocks. We have also demonstrated how rich household, industry, and regional heterogeneity can interact to potentially amplify or dampen fiscal shocks. We now take our generalized multiplier to the data to understand how various dimensions of heterogeneity in our model shape multipliers and thus the design of fiscal policy in practice. To do this, we directly estimate the sufficient statistics that comprise the generalized multiplier using a variety of datasets. This section describes both the datasets and methodology we use to estimate these sufficient statistics.

First, recall from Proposition 1 that the total response of demand dY^1 to any shock with direct effect ∂Y^1 is given by:

$$dY^1 = \left(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}\right)^{-1}\partial Y^1 \tag{14}$$

Our application assumes there is one representative firm per region-industry pair, and that there is one household type per region-demographic group pair (for a set of demographic groups described below). So to estimate the generalized multiplier, we need estimates of three key objects: the regional input-output matrix $\widehat{\boldsymbol{X}}^{\boldsymbol{1}}$ describing the input use requirements of every region-industry pair, the rationing matrix $\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}$ describing how much each demographic-region pair's income changes in response to a one dollar change in revenue of each region-industry pair, and the directed MPC matrix $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}$ describing how much each demographic-region pair consumes from each region-industry pair when they receive a one dollar income shock.

In moving to the data, we must also account for three empirically-relevant factors left out of our baseline model—capital, profits, and foreign income. At a high level, our strategy, which we will describe in detail below, is to (1) model capital as an intermediate input, (2) model profits by assuming constant markups, as in Appendix D.1, and (3) model foreign factors as a type of "labor" with zero MPC, reflecting that payments leaving the economy do not re-enter through income effects.

The following subsections describe in detail how we estimate each of the three components of our generalized multiplier: the input-output, rationing, and directed consumption matrices. We restrict our attention to the United States in 2012, which is the most recent year for which we have several of the key datasets we use. In Section 5.2.2, we explore the robustness of our results to several assumptions embedded in our estimation methodology.

4.1. The Regional Input-Output Matrix

The regional input-output matrix $\widehat{\boldsymbol{X}}^{\boldsymbol{1}}$ is an $(S \times J) \times (S \times J)$ matrix where S is the number of regions and J is the number of industries/sectors. The (ri, sj) component of this matrix corresponds to the amount of sector i in region r's good required to produce a single unit of sector j in region s's good. To estimate this object, we must first take a stand on the level of granularity at which to model sectors and regions. Guided by the level at which input-output data are available, we work with a slight coarsening of the Bureau of Economic Analysis' (BEA) collapsed input-output sector classification, leaving us with 55 sectors which loosely correspond to the 3-digit NAICS classification. Similarly, as we use the Commodity Flow Survey (CFS) microdata on interstate trade, we set regions at the level of the state (including Washington D.C.), leaving us with 51 regions. This leaves us with 2805 sector-region pairs.

We construct the regional input-output matrix in three steps. First, following others in the literature, we use data from the 2012 BEA make, use, and imports tables to construct the domestic, national input-output matrix, which measures the dollar value of products from industry j that are used by industry i (U.S. Bureau of Economic Analysis, 2012). For commodities produced by multiple industries, we assume that all users of such commodities source them from the various producing industries in the same proportions.[16] We also make an adjustment to account for linkages across industries through capital investment. This is necessary as the standard use table accounts only for changes in intermediate goods usage. To impute each industry's expenditure on investment goods, we assume that all industries invest the same fraction of their gross operating surplus (available in the use table) in capital. To compute the direction of this investment toward different industries, we assume that each firm demands the same investment good and we therefore compute its industrial composition with the same procedure—using the use, make, and import tables—as we use for inputs. We then add this investment correction to the previously constructed input-output matrix.

Second, we use the 2012 public-use microdata from the CFS to construct a matrix describing how much each state imports from all other states. The CFS is a survey conducted by the US Census Bureau and includes data on 4,547,661 shipments from approximately 60,000 establishments (U.S. Census Bureau, 2015). Using this information, we calculate the total value of shipments between each pair of states for each tradable industry using the mapping between commodities and industries outlined in the BEA's make table.[17] For all nontradable industries, we assume that the commodity is sourced entirely within the state.[18]

Finally, we construct the regional input-output matrix by combining the national industry-

[16] While these assumptions are necessary given the limitations of our data, one might be concerned that they play a significant role in our findings. We address this concern in Section 5.2.2 by demonstrating that our main empirical results are robust to a wide range of alternative calibrations.

[17] Caliendo et al. (2018) use a similar methodology to construct their regional input-output matrix.

[18] We assume that the following sectors and sector groups are domestically tradable: agriculture; mining, oil, and gas (but not mining support services); manufacturing; and wholesale. Of our tradable industries, three (two agriculture sectors and one oil and gas sector) are not included in the CFS data. We assume that each destination state sources goods in these industries from origin states in the same proportion as in the sum of CFS data across industries. We assume that all other sectors are domestically non-tradable. This assumption implies that we do not allow for trade in services. We probe the robustness of our results to this assumption in Section 5.2.2.

level input-output with state-by-state trade flows. Specifically, we assume that the amount of industry i in state r used by industry j in state s is the product of the share of industry j's inputs that come from industry i and the fraction of sector i goods flowing to s from r (out of all origin states). This yields a matrix describing, for each industry-region pair, how much of each other industry-region pair's production is used to produce a single unit of output.

4.2. The Directed MPC Matrix

The directed MPC matrix $\widehat{\boldsymbol{C}}^{\mathbf{1}}\boldsymbol{m}$ is a $(S \times J) \times (S \times N)$ matrix where N is the number of demographic groups and recall S and J are the number of regions and industries, respectively. The (ri, sn) component of this matrix maps how a one dollar change in demographic n living in region s's income changes that household's consumption of good i from region r. Again, this first requires us to take a stand on the level of granularity at which to model demographic groups. Guided by the level at which precise estimation of MPCs is possible in the Panel Study of Income Dynamics (PSID), we set the number of demographic groups at 82, comprising 80 baseline groups (five initial income groups, four age groups, two gender groups, two race groups) and two dummy groups for the owners of capital and foreigners.[19]

We construct the directed MPC matrix in three steps. First, we construct MPCs for total consumption expenditure for each of our 80 demographic groups using the PSID, Consumer Price Index (CPI), and Consumer Expenditure Survey (CEX) following the methodology in Patterson (2023). Specifically, we follow the procedure of Gruber (1997), using the panel structure of the PSID to estimate the equation:

$$\Delta C_{Ht} = \sum_x (\beta_x \Delta E_{Ht} \times x_{Ht} + \alpha_x \times x_{Ht}) + \delta_{s(H)t} + \varepsilon_{Ht} \tag{15}$$

where C_{Ht} is household H's consumption at time t, E_{Ht} is household H's labor earnings at time t, x_{Ht} is a demographic characteristic of the individual, and $\delta_{s(H)t}$ is a state by time fixed effect. Consistent with the length of a typical recession, we interpret the first time period in our model as lasting for two years and accordingly estimate (15) using two-year changes for consumption and income. Estimating Equation 15, we obtain the following estimate of the MPC for household H at time t:

$$\widehat{MPC}_{Ht} = \sum_x \hat{\beta}_x x_{Ht} \tag{16}$$

However, there are two challenges in performing this estimation. The first issues arises as there are a wide range of factors that could simultaneously move income and consumption. To address this,

[19]Our five income groups correspond to: less than \$22,000, \$22,000-\$35,000, \$35,000-\$48,000, \$48,000-\$65,000 and more than \$65,000. Our four age groups correspond to those 25-35, 36-45, 46-55 and 56-62. Our race groups are black and non-black. Our gender groups are men and women.

we instrument for changes in labor market earning using transitions into unemployment. This is desirable as such shocks are both large and persistent. Unemployment shocks therefore capture that variation most important to understanding recessions. Indeed, if recessions can be seen as shocks of the same persistence as unemployment, then this MPC is exactly the right object to capture shock propagation in the manner suggested by the model.[20]

The second issue stems from measurement in the PSID: for most of the PSID sample, only expenditure on food consumption is measured. Using only this measure is problematic as food is a necessity and expenditure on food is likely to be distorted by the provision of food stamps (Hastings and Shapiro, 2018). To overcome this issue, we use overlapping information in the PSID and CEX to impute a measure of total consumption expenditure, following the methodology of Blundell, Pistaferri, and Preston (2008) and Guvenen and Smith (2014). Concretely, we use the CEX to estimate demand for food expenditure as a function of durable consumption, non-durable consumption, demographic variables and relative prices from the CPI. Under the assumption of monotone food expenditure, this function can be inverted to predict total consumption as a function of food expenditure and demographics in the PSID. This procedure generates substantial heterogeneity across households in estimated MPCs (see Figure A1 in Appendix E). For discussion of the robustness of these estimates, see Patterson (2023).

Next, we estimate the shares of each of our 55 industries in the consumption baskets of each of our 80 demographic groups using the CEX and CPI. We first deflate consumption over the 54 measured CEX categories using the CPI and then compute the average consumption basket share of each demographic group. Using a concordance between NIPA goods and our industry classifications, we then map consumption at the household level in each category to the 55 industries used in our analysis.

We use these consumption basket shares and our estimated MPCs to construct an estimate of the directed MPC for each of the 80 demographic groups out of each of the 55 industries. We do this by assuming linear Engel curves of households for each category of consumption. Formally, we estimate the directed MPC of household H at time t as:

$$\widehat{MPC}_{n(Ht)i} = \alpha_{n(Ht)i} \widehat{MPC}_{n(Ht)} \tag{17}$$

where $n(Ht)$ is the demographic group of household H at time t—which we from now on suppress when clear from context—and $\alpha_{n(Ht)i}$ is the demographic-specific consumption basket weight of good i. Of course, the imposition of linear Engel curves may be overly restrictive. However, our estimates always lie in the 95% confidence interval of estimates of good-specific MPCs from the PSID in the years in which this is possible (see Figure A2 in Appendix E), suggesting that we are

[20] A more subtle issue is that we use this same MPC when we study negative demand shocks as when we study positive changes in government spending. This choice reflects our interest in the use of fiscal policy *during severe downturns*, where its role is mainly to prevent transitions to unemployment rather than to induce transitions out of unemployment.

capturing reasonable dimensions of heterogeneity with this assumption.

Finally, we use our estimated state-state gross flows in goods to arrive at the regionally-directed MPCs. Formally, for tradable goods, we assume that all households in a state consume from all other states in proportion to the fractions of imports of that good that originate from those states:

$$(\widehat{\boldsymbol{C}}^{\mathbf{1}}\boldsymbol{m})_{risn} = \lambda_{irs}\chi_s\widehat{MPC}_{ni} \tag{18}$$

where λ_{irs} is the fraction of shipments of good i from state s to state r as a function of the total shipments of good i to state r, as we earlier computed to construct the regional input-output matrix,[21] and where χ_s is the share of domestic production (i.e., not imports from other countries) in personal consumption expenditures on sector-s goods, derived from the BEA's use and imports tables.[22] We assume all nontradable goods are consumed within the state.

The procedure above provides the directed MPC entries for the 80 demographic groups. It remains to estimate the directed MPCs for shareholders and foreigners. For foreigners, we simply set all entries to zero.[23] This coincides with the assumption that, of all foreign recipients of income that leaves the US, none spend this income in the US or indirectly cause other spending in the US.[24] For shareholders, we take the MPC out of stock market wealth as estimated by Chodorow-Reich, Nenov, and Simsek (2021) at 0.032.[25] We then allocate consumption across goods according to the average consumption basket of our highest income groups, as we computed in the CEX (with the same construction of the regional direction of consumption).

Finally, recall that we have assumed that the marginal consumption response out of first-period government transfers (i.e., negative taxes) is the same as that out of first period income earned through labor supply. Theoretically, this follows if consumption and labor are additively separable in household utility, as we have assumed. Empirically, the documented cross-sectional patterns in MPCs in response to tax transfers are similar to those we uncover using employment shocks, suggesting that this assumption is not driving the patterns we uncover below (see Parker, Souleles, Johnson, and McClelland, 2013; Fagereng, Holm, and Natvik, 2019).[26]

[21]Considering that the CFS comprises both consumption goods and intermediate goods flows, this method may source too much consumption from outside each region. In Section 5, we explore the robustness of this modeling assumption for how consumption is sourced by considering a model with total consumption autarky where all consumption is sourced within the state. This has a very small impact on our main results.

[22]Note that, since some consumption spending is directed to imports from other countries, each region-demographic (s, n)'s MPC from the perspective of the model, $m_{sn} = \sum_{r,s\in\mathcal{R},i\in\mathcal{J}}(\widehat{\boldsymbol{C}}^{\mathbf{1}}\boldsymbol{m})_{risn}$ is lower than its raw estimated MPC in the PSID.

[23]In later decompositions that use the spending direction of even zero-MPC households, we assume foreigners' spending is directed in proportion to GDP.

[24]In assigning foreigners an MPC of zero, we implicitly hold exports fixed in response to policy.

[25]Since the model assumes that all capital income goes to shareholders, this assumption does not impose that the MPC for labor and capital income are different for the same household.

[26]Our framework is flexible enough that it would be easy to perform our empirical analysis with a different calibration of the MPC out of government transfers. Since the MPC estimates out of tax rebates are noisier than those using unemployment, we maintain this assumption in the current analysis.

4.3. The Rationing Matrix

The rationing matrix $\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}$ is a $(S \times N) \times (S \times J)$ matrix where recall S, N, and J are the number of regions, demographic groups, and industries, respectively. The (rn, si) component of this matrix maps a one dollar change in the production of good i in region s to the resulting change in labor income for demographic n in region r.

We construct this object using several empirical inputs. Intuitively, this object maps output changes in a given industry-region to income changes for workers of a given demographic group in that region, adjusting for their employment shares, local labor shares, and the elasticity of earnings to output. The functional form below summarizes these components, and we explain the role of each term in detail.

$$\left(\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}\right)_{rnsi} = \mathbb{I}[r = s]\frac{y_{inr}}{\sum_n y_{inr}}\alpha_{ir}\beta_i\left(1 + \xi\left(MPC_n - \overline{MPC}_{ir}\right)\right) \tag{19}$$

The term y_{inr} is total earnings of demographic n in industry i in region r, which we measure using the 2000-2017 American Community Survey (ACS, Ruggles, Flood, Sobek, Backman, Cooper, Rivera Drew, Richards, Rogers, Schroeder, and Williams, 2025). The term $y_{inr}/\sum_n y_{inr}$ is therefore the share of total spending in an industry and region that goes to each demographic group n. This term distributes any change in total labor income for industry i in region r across demographic groups based on their observed earnings shares, capturing systematic worker sorting across sectors and regions.

The second term, $\alpha_{ir}\beta_i$, is a proxy for the labor share of gross output in industry i in region r. Since we do not have regional measures of gross output by industry, we instead use the compensation to value-added ratio and convert it to an estimate of the compensation to gross output ratio using the assumption that the ratio of gross output to value-added within an industry is the same in each region. We do this using BEA data on compensation and value added by industry (α_{ir}) and the national value added to gross output ratio in industry i (β_i). The resulting labor share informs how incomes will move when output changes; regions and sectors with high labor shares will see larger changes in worker incomes when output changes than areas with lower labor shares.

The final term in Equation 19 captures the fact that earnings of high-MPC (typically low-income) workers are more sensitive to output shocks than those of low-MPC workers. The adjustment scales the baseline income response of each demographic by $\left(1 + \xi\left(MPC_n - \overline{MPC}_{ir}\right)\right)$, where $\overline{MPC}_{ir}$ is the earnings-weighted average MPC of workers in that industry-region. We set $\xi = 1.33$, with the earnings elasticity–MPC slope estimated by Patterson (2023).[27] The functional form in Equation 19 imposes that this slope holds within each region and industry.[28] The indicator

[27] See Patterson (2023) for more details and discussion.

[28] While the model can in principle accommodate regional migration, our empirical calibration assumes that

function up front imposes the condition that all labor earnings are received within the state where production occurs.

Lastly, not all factor payments go to labor (i.e., the labor share is less than 1). We separately account for this other income with payments to domestic owners of capital and payments made to foreign factors. We compute payments made to domestic owners of capital via the following procedure. We first compute profits in each region-industry pair. To do this, we impute industry profits as a share of production by multiplying national, industry-specific gross operating surplus per unit revenue by the national, cross-industry fraction of gross operating surplus not spent on investment, sourced from from the BEA use table. This gives us a measure of the profits in each industry. We assume that 82.6% of these profits go to domestic equity owners, which was the domestic share of US equity ownership in 2012 (Board of Governors of the Federal Reserve System, 2013). Lastly, we allocate these profits across states according to each state's share of dividend income in the Internal Revenue Service's Statistics of Income (IRS SOI) data (Internal Revenue Service, Statistics of Income Division, 2012). Finally, we compute payments made to foreigners as the residual of payments made to intermediate producers, payments made to labor and payments made to shareholders.[29]

5. Empirical Exploration of Fiscal Multipliers

In this section, we study the propagation of fiscal shocks in our calibrated economy, exploring both changes in aggregate consumption and demand spillovers across regions. We begin by quantifying how the variation in the fiscal multiplier for government purchases or transfers—i.e., the total change in private and public consumption per unit of fiscal spending—depends on how that shock is targeted. We then demonstrate that these differences stem almost exclusively from differences in the initial incidence of shocks on households with different MPCs rather than from variation in which goods these households consume, and we accordingly study how various features of the economy determine the incidence of shocks. Finally, we quantify the extent of geographic spillovers through the generalized multiplier and contextualize our findings in light of recent empirical estimates.

firms draw labor demand proportionally from the existing demographic composition of their region. This rules out migration-induced changes in firm-level workforce composition in response to shocks.

[29]In a small fraction of cases, this leads to a *negative* foreign share of revenues, which is unrealistic. To avoid this, we could alternatively reduce the profit share of revenue in region-industry pairs with high labor shares. Insofar as we use similarly small MPCs for foreigners and shareholders, this alternative calibration would generate similar quantitative results.

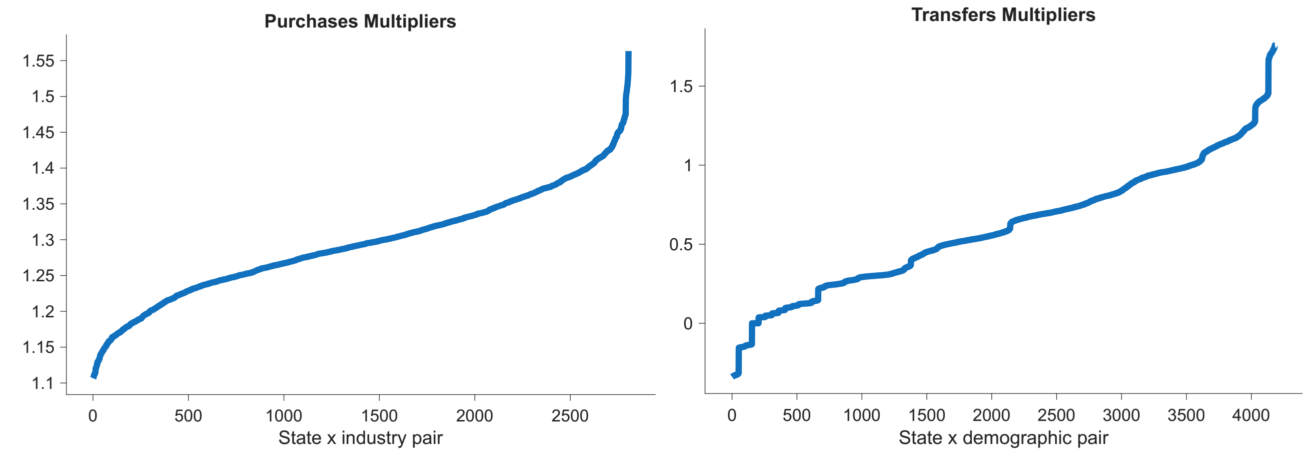

Fig. 2. Left: Change in aggregate consumption resulting from a \$1 government purchase shock to each state-by-industry pair, i.e., distribution of government purchase multipliers. The x-axis sorts state-by-industry pairs by their purchase multipliers and the y-axis plots the estimated multiplier. Right: Change in aggregate consumption from a \$1 transfer shock to each state-by-demographic-group pair, i.e., distribution of transfer multipliers. The x-axis sorts state-by-demographic-group pairs by their transfer multipliers.

5.1. Extent of Heterogeneity in Multipliers

We estimate that the response of aggregate consumption (private consumption plus government spending) to a demand shock which is GDP-proportional across industries and regions, or *aggregate purchases multiplier*, is equal to 1.28, a number consistent with the large literature estimating fiscal multipliers with constrained monetary policy (Ramey, 2011; Chodorow-Reich, 2019). However, the left panel of Figure 2—which shows the effect on consumption of spending a dollar in a given industry within a specified state—documents wide dispersion depending on how a shock is targeted, with the extent of amplification beyond the original purchase varying by a factor of more than five. Indeed, our range of multipliers, from about 1.11 to 1.56, provides one rationalization for the variation in purchases multipliers estimated in the literature. Transfers multipliers, i.e., the effect on aggregate consumption of transferring one dollar to a household of a given demographic within a specified state, vary even more widely. The right panel of Figure 2 shows that the effect on aggregate consumption of transferring a dollar to a household ranges from slightly below zero for some households (some types have negative MPCs) to close to two dollars for others.

Much of the heterogeneity in multipliers remains when targeting is constrained to be more granular: the amplification of government purchases differs by a factor of more than three across industries and more than 1.5 across states (see Figure A16 in Appendix E); transfers multipliers differ by a factor of 1.4 across states and by nearly as much across demographic groups as across demographic-region pairs (See Figure A17 in Appendix E).

5.2. *Sources of Heterogeneity in Multipliers*

Recall from Proposition 2 in Section 3.2 that—for any fiscal shock—three adjustments to the basic Keynesian multiplier capture the effects of heterogeneity. In particular, the dispersion in fiscal multipliers from Figure 2 could derive from differences in (1) the *incidence* effect, wherein some shocks load more heavily on agents with higher MPCs, (2) the *bias* effect, wherein some shocks load onto households who direct their spending to high-MPC households, or (3) the *homophily* effect, wherein some shocks load onto households who direct their spending to other households with similar MPCs to their own. However, empirically, we find that essentially all of the heterogeneity across groups in Figure 2 is driven by the differential direct incidence of those shocks onto agents with different MPCs.

5.2.1. *Importance of Initial Incidence*

To understand why only the incidence effect is empirically large, recall Proposition 2. In order for the bias or homophily terms to be large, there must be significant heterogeneity across households in basket-weighted MPCs m_n^{next}—that is, in the average MPC of the workers ultimately employed in producing n's marginal unit of consumption—and these basket weighted MPCs must differ from the benchmark $\mathbb{E}_{h*}[m_n]$. Indeed, if m_n^{next} is homogeneous and $\mathbb{E}_{\partial h^1}[m_n^{\text{next}}] = \mathbb{E}_{h*}[m_n]$, then both the bias and homophily terms are zero as all households effectively direct their consumption to the same sorts of household targeted by an aggregate shock. The left panel of Figure 3 documents that in the data, there is minimal heterogeneity in basket-weighted MPCs, shown by the very shallow slope between basket-weighted MPCs (y-axis) and household MPCs (x-axis). As a result, the homophily effects are very close to zero. Moreover, the scatterplot demonstrates that basket-weighted MPCs all lie very close to the benchmark average MPC ($\mathbb{E}_{h*}[m_n]$). Consequently, bias effects are also very close to zero. Indeed, for any possible shock, the incidence term accounts for more than 98.8 percent of the multiplier.[30] To drive this point home, the red line in the right panel of Figure 3 shows multipliers from a counterfactual model without heterogeneous consumption in which the bias and homophily effects are identically zero. As one can see, there is effectively no difference in the full distribution of multipliers when we impose this condition, demonstrating that it plays no role in shaping the baseline estimates.[31]

[30]Every feasible ∂h^1 can be obtained as the linear combination of demand shocks to each sector-region pair. We therefore compute the bias and homophily effects from each of these "basis vector" shocks and plot the full distribution of bias and homophily terms (see Figure A3 in Appendix E). Across the full distribution of shocks, the contributions of the bias and homophily terms always comprise less than 1% of the multiplier; they are empirically negligible for all feasible demand shocks. We also compute the full distribution of error terms arising from the approximation in our decomposition result (the right panel of Figure A3 in Appendix E) and find that they are an additional half an order of magnitude smaller than the bias and homophily terms. Our approximation is therefore very tight for any feasible shock.

[31]In Figure A6 of Appendix E we show a scatter plot of the multipliers from these two models. The correlation in multipliers across the two models is nearly perfect.

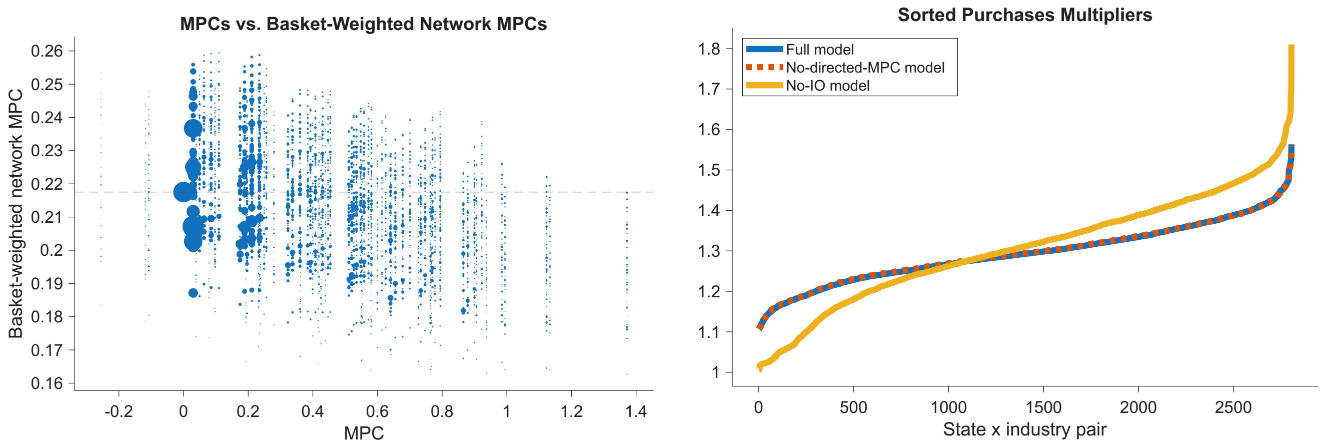

Fig. 3. The left panel shows a scatter of MPCs m_n against basket-weighted MPCs m_n^{next}. The dashed line gives the average MPC $\mathbb{E}_{h*}[m_n]$ for h^* given by the income incidence of a shock to demand proportional to 2012 state-industry GDP. The right panel shows the change in aggregate consumption resulting from a one dollar demand shock to each industry-region pair, sorted by the magnitude of the effect. The full model is the baseline and plotted in blue. The no-directed-MPC model assumes that all households direct their consumption in proportion to aggregate consumption. The no-IO model assumes that there is no use of intermediate goods.

The lack of bias and homophily effects appears to be a real feature of the data, rather than a failure of our estimation approach to capture them. While the bias and homophily terms each operate to second order in the average MPC—which constrains them to be modest in size—it is easy to see from the examples in Section 3.2 that the combination of these terms can, in principle, be quantitatively large. Indeed, our estimates of consumption basket shares in the CEX do display substantial variation across households (see Figure A7 in Appendix E), allowing for the possibility of large bias and homophily effects. The absence of these effects, then, stems from two countervailing empirical observations. First, high-MPC households disproportionately consume goods produced by low-labor-share industries (see Figures A4 and A5 in Appendix E), directing more spending toward capital, the owners of which have low MPCs.[32] Second, our estimates feature substantial within-region non-tradables demand, with around a third of consumption demand reaching labor within the same state (see Figure A8 in Appendix E). Moreover, there is spatial heterogeneity in MPCs, with income-weighted MPCs differing by a factor of 1.7 across states (See Figure A14 in Appendix E). Together, these regional forces generate a modest positive homophily effect whereby higher (lower) MPC workers direct their consumption more toward local labor which similarly features high (low) MPC. However, these labor share and local demand effects are both fairly weak, and they run in opposite directions. When combined, they partially cancel, so that all types spend on goods baskets produced by households of very close to the average MPC.

[32]This finding is consistent with the empirical patterns in Hubmer (2019).

5.2.2. Robustness of Empirical Decomposition

So far, we have established the stark empirical finding of the previous section—namely that bias and homophily effects are quantitatively negligible—under our baseline estimates. In this section, we revisit many of the strong assumptions that underlie this finding by computing bias and homophily effects under a range of alternative calibrations that bound the importance of these assumptions for our conclusions.

To probe the robustness of this finding, we now consider alternative assumptions along the four key dimensions that have the potential to affect the size of this bias and homophily effects: apportionment of trade flows, construction of the IO matrix, heterogeneity in consumption baskets, and income rationing. In each case, we provide alternative assumptions that have the potential to increase the importance of bias and homophily. While many of these alternatives are extreme, this is intentional and serves to emphasize that, even under extreme assumptions, the quantitative magnitude of bias and homophily effects are irrelevant.

First, our main estimates assume that all demand for each tradable good i from within each state s is apportioned across other states r according to trade flows of i from r to s in the Commodity Flow Survey. This rules out the possibility that, for example, two firms in different industries within the same state may source the same input from different states than one another, potentially generating more or less homophily. We also assume that there is no cross-state trade in services. To bound the importance of these assumptions, we consider two extreme alternative calibrations for interstate trade, one in which all buyers of a good (including non-tradables) source it from each state proportionally to the fraction of national production of the good done in the state, and another in which all buyers purchase goods only within their own state.

Second, our main estimates assume that all states have the same cross-sectoral input-output matrix. As this may mask cross-state variation in IO structure that could link industries in a more or less homophilic manner, we consider an extreme calibration in which we shut down connections through the input-output network, instead assuming that firms keep their input spending constant but source all inputs from within their own industry-state-pair. We treat household consumption as in the baseline in this calibration and consider this model of input demand as a fourth option for the regional IO matrix in addition to the three discussed in the previous paragraph, i.e., (i) our baseline, (ii) the case where the source of imports is chosen independent of their destination, and (iii) the case of no interstate trade.

Third, our baseline analysis treats investment identically to intratemporal intermediate inputs. To the extent that capital durability allows investment flows to vary more with gross output than intratemporal inputs, we may understate the importance of the investment component of the input-output network, whose different structure could have different implications bias and/or homophily. We therefore consider an alternative calibration in which we double the investment component of

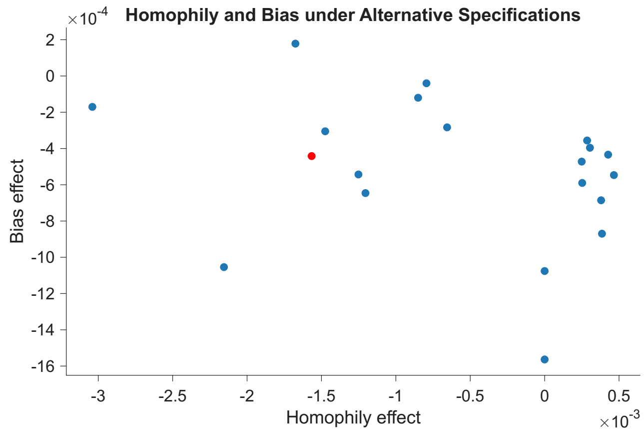

Fig. 4. Scatter plot of bias and homophily effects for a one dollar, GDP-proportional shock in 20 alternative models where we compare all combinations of the modifications described in Section 5.2.2. These modifications adopt extreme alternative assumptions on regional trade, the input-output network, consumption heterogeneity, and income rationing. The red dot corresponds to the baseline model. Here, the reference incidence h^* is that induced by a GDP-proportional shock.

our estimated input-output network while holding fixed the intratemporal component.[33]

Fourth, our baseline estimates allow each demographic group to consume its own bundle of goods, which has an anti-homophilic effect on the income-to-spending network as discussed in the previous section. In order to give our estimates a greater chance of generating homophily, we consider an alternative calibration in which all households consume the same basket of goods, which we set equal to the income-weighted average basket.

Fifth, our baseline estimates assume, following Patterson (2023), that the incomes of higher-MPC workers are more sensitive to aggregate employment, inducing a positive bias effect in the spending-to-income network. We therefore also consider the case where income is rationed purely in proportion to a state-demographic-group's income within a state-industry pair.

Figure 4 shows the size of the bias and homophily effects in the $(3 + 1 + 1) \times 2 \times 2 = 20$ versions of the model described above. In all cases, the bias and homophily effects contribute less than half of a cent to the multiplier of a GDP-proportional shock that changes consumption by around \$1.28 in total. Therefore, while we demonstrate the theoretical possibility of large bias and homophily effects (in Section 3.2), their irrelevance is robust and it is unlikely that they are empirically relevant in advanced economies.

5.2.3. Determinants of Initial Incidence

Since the heterogeneity in shock amplification in Figure 2 does not stem from bias or homophily effects, it must instead come from differences in the incidence of different shocks onto the MPCs of households. For transfers, the initial incidence is immediately apparent and is driven solely

[33] We proportionately reduce all factor payments as a share of gross output so that each firm's spending still equals its revenues.

by heterogeneity in MPCs in the population. However, for government purchases, three distinct factors widen the distribution of multipliers. First, differences in the demographic composition of the workforce across sectors and regions cause large differences in the average MPCs of workers across firms and regions. Second, differences in the share of labor that each sector directly employs cause large differences in the MPC of the ultimate recipients of factor income. In particular, firms employing a lot of capital but little labor pass most factor payments on to the owners of capital who have very low MPC and therefore feature small purchases multipliers. This is shown in Figure A9 in Appendix E, which plots the labor share of each industry-state pair against its purchases multiplier; there is substantial heterogeneity in labor use and low labor use is associated with a small purchases multiplier. Third, differences across firms in the covariance of worker MPC and exposure to changes in firm revenue generate additional widening of the distribution of multipliers. This is shown in Figures A10 and A11 in Appendix E where we compare the baseline model—which features greater rationing to agents with higher MPCs—to a model with rationing to agents uniformly by income; there we observe both an upward shift in the distribution of purchases multipliers as well as an increase in its range.

Conversely, input-output linkages serve an important role in *narrowing* the heterogeneity induced by these differences. This can be seen in the right panel of Figure 3, where the yellow line corresponds to the model without input-output linkages, which features a much more dispersed distribution of multipliers.[34] The role of input-output linkages in reducing dispersion is intuitive. In the absence of inputs, when the firm directly employing the highest-MPC factors gets an additional dollar of revenue, it spends it all on those high-MPC factors. With inputs, this same firm spends a fraction of its revenue on goods produced by other firms, who in turn direct that money to their (by construction) less-than-highest-MPC factors—effectively diluting the MPC of the initial firm. This dilution effect attenuates the heterogeneity in industry multipliers.[35] This same phenomenon explains why the distribution of transfers multipliers in Figure 2 is more dispersed than the distribution of purchases multipliers: A transfer to the highest- or lowest-MPC household reaches it directly, rather than being spread across households with more moderate MPCs.

[34]See Figure A15 in Appendix E for a scatter plot of the multipliers across both the full model and that without input-output linkages.

[35]Our finding that IO linkages reduce heterogeneity in purchases multipliers is distinct from an existing literature that emphasizes the role of IO networks in amplifying economic shocks (Acemoglu et al., 2012; Carvalho, Nirei, Saito, and Tahbaz-Salehi, 2016; Baqaee, 2018; Elliott, Golub, and Leduc, 2020). First and foremost, our finding is not that IO linkages attenuate amplification on *aggregate*, but rather than they reduce the dispersion in amplification across industries. In this sense, we simply have a different focus. Moreover, the key reasons that IO links generate aggregate amplification in the literature—namely, that supply shocks are more powerful when the input share of production is large (a la Hulten) and that supply and demand shocks can cause cascades of firm defaults when production has a fixed cost—play no role in our setting, as we focus on demand shocks and assume production is CRS.

5.3. Regional Demand Spillovers

Finally, we turn our focus away from the effects of fiscal shocks on *aggregate* output and instead consider how income multipliers may propagate across state lines. Such spillovers are of direct policy relevance, as a planner may want to stimulate demand in a particular, depressed, area without "overheating" the economies of other nearby regions. They are also of interest to a recent empirical literature that uses quasi-random cross-regional variation in fiscal spending to estimate local fiscal multipliers (Nakamura and Steinsson, 2014; Chodorow-Reich, 2019). Regional demand spillovers complicate the relationship between these local estimates and the national multiplier, as most research designs only recover the effect of spending on i in GDP in i relative to GDP in j—which is not a suitable control group if the spending indirectly boosts j's GDP.

The regional interlinkages embedded in our model allow us to provide an estimate for the magnitude of these cross-state spillovers. We quantify these spillovers within our model by considering a unit of government purchases in each state, which we assume is distributed across industries within the state in proportion those industries' shares of GDP within the state. Averaging across states, we find that aggregate value added increases by $1.29 in response to $1 of additional government purchases (i.e., public consumption). However, only an average of 79 cents of this increase in value added occurs within the state to which purchases are directed, with the remainder diffusing to other states. These interstate spillovers operate through two channels: First, the final goods purchased in any targeted state are produced, in part, with inputs imported from other states. Second, when workers in the targeted state spend their income generated by the government purchase, they direct a share of their spending toward other states, both directly and by buying goods produced with inputs sourced from other states.[36]

Figure 5 shows these spillovers cartographically, plotting the changes in state-level value added per worker that result from $1-per-worker shock to output in Texas and Michigan, respectively.[37] Consistent with a modest regional bias in trade, we observe slightly greater increases in value added in states that neighbor those shocked directly.

These estimates are in line with recent empirical evidence estimating the magnitude of these spillovers directly. Specifically, Auerbach et al. (2020) use detailed geographic information on local defense spending and find that large positive spillovers across geographies, suggesting the importance of positive demand spillovers through input-output networks and directed MPCs. They also find that the spillovers are decreasing in the distance between cities. Our results are

[36] Appendix Figure A12 shows the results of an analogous exercise for state-specific uniform transfers rather government purchases. We estimate that, out of the 65-cent increase in value added resulting from a $1 uniform transfer to households in an average state, 34 cents occurs within that state. The greater out-of-state share of the value added increase reflects that the spending induced by a transfer shock is partly directed towards goods produced in other states.

[37] Appendix Figure A13 shows versions of these maps that measure spillovers in absolute, rather than per-worker terms.

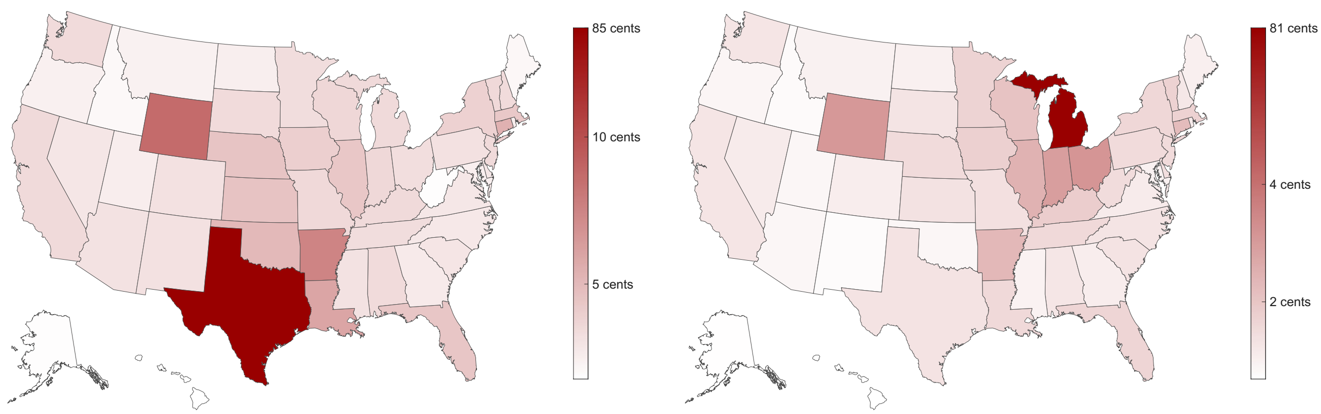

Fig. 5. Changes in state-level value added per worker following a GDP-proportional \$1-per-capita purchases shock to Texas (left panel) and Michigan (right panel).

consistent with this, as our estimated spillovers are largest for the geographically closer states. These estimates suggest that demand spillovers across states are empirically important when evaluating the total effect of localized fiscal spending.

6. Implications for Design of Fiscal Policy

So far, we have studied how fiscal shocks propagate to affect consumption in general equilibrium. We have found that fiscal multipliers vary widely depending on where spending is targeted, and that this variation is driven entirely by the heterogeneous incidence of the shocks on workers with different MPCs. In this section, we explore the implications of these findings for the design of fiscal policy. We consider a setting where the social planner seeks solely to maximize aggregate consumption. This is optimal in the case where everyone has zero marginal labor disutility (i.e., there is widespread unemployment) and the planner is indifferent to redistribution.

When the bias and homophily effects are exactly zero for all possible policies, in Appendix A.3 we show that the change in welfare from first-period fiscal policies is given by:

$$dY \propto \sum_{n \in N} m_n \partial h_n^1 \tag{20}$$

where $\partial h^1 = \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1} dG^1 - \boldsymbol{\mu} d\tau^1$ is the the direct change in household incomes induced by the policy, before spending adjusts, and m_n is the MPC for households of type n. Thus, the change in consumption is proportional to the inner product of household MPCs and how much fiscal policy directly changes household incomes. Intuitively, this implies that the social planner who seeks to maximize aggregate consumption should simply target policies to affect the incomes of households with the highest MPCs. This intuitive result holds because in the absence of bias and homophily effects, all households direct their consumption in the same way for the purposes of amplification. Therefore, the best thing a consumption-maximizing planner can do is simply

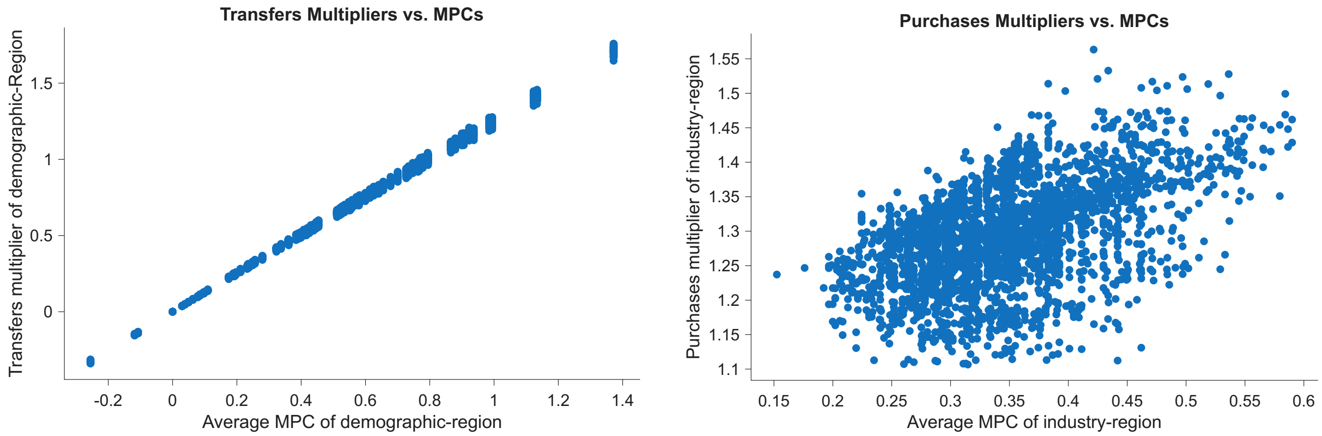

Fig. 6. Left: the effectiveness of targeting transfer stimulus, as measured by the transfers multiplier, by household MPCs at the level of demographic-region pairs. Right: the effectiveness of targeting government purchases stimulus, as measured by the purchases multiplier, by the average MPC of workers at the level of industry-region pairs.

target households with the highest MPCs.

Of course, the bias and homophily terms are not exactly zero in the data, but Figure 6 demonstrates the effectiveness of simple MPC targeting. The left panel scatters household MPCs against the resulting transfers multiplier from giving them a dollar, revealing an effectively perfect relationship between the two.[38] Note that for the design of fiscal transfers, even the IO network and industry labor shares are irrelevant. The social planner simply needs to know the distribution of households MPCs to design policy that maximizes aggregate consumption.

By contrast, the right panel shows that, for government purchases, it is not sufficient to target the sectors *employing* the highest MPC workers: while there is a positive relationship between the MPC of a sector's employees and the purchases multiplier in a sector, the correlation is well below 1. Rather, maximally expansionary purchases policy targets those sectors such that when their production expands, accounting for the intermediates goods they use and the intermediates used by the producers of those intermediates and so on, the resulting change in labor income ends up in the hands of the highest MPC agents. While this requires no knowledge of the direction of household spending, it does rely on an understanding of the structure of production—through the input-output network and labor rationing. The planner must work out the final labor income consequences of their spending and target according to the MPC of the workers receiving that terminal labor income. This echoes results in Baqaee (2015), which emphasizes the need to adjust labor shares for the input-output structure of production. This difference is quantitatively important; the right panel of Figure 6 shows that naively targeting sectors employing the highest

[38]Note that the MPC that we use in Figure 6 is estimated using unemployment as the identifying shock, and therefore captures the consumption response to a potentially persistent shock. The MPC that is better suited for the analysis of fiscal policy would be the MPC out of a transitory shock. If the MPC out of these two shocks are highly correlated across demographic groups, this difference should be less important for the question of which demographic groups to target. While it is hard to test this explicitly, the cross-demographic patterns in MPCs that we utilize here have a correlation of above 0.5 with self-reported MPCs from survey data (Jappelli and Pistaferri, 2014) and have similar patterns as those in response to tax rebates (Parker et al., 2013).

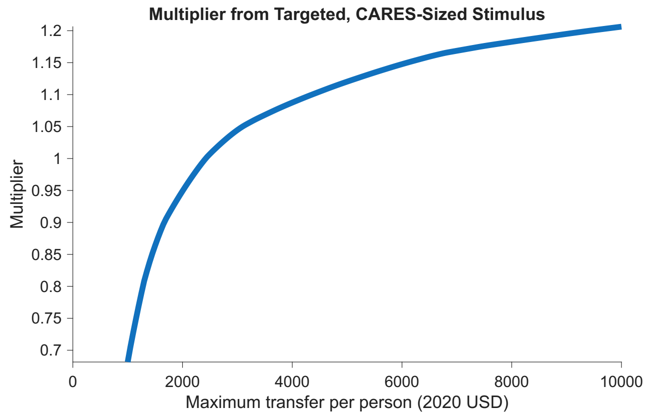

Fig. 7. This figure reports the transfer multiplier (y-axis) out of a policy that allocates a given maximum transfer (x-axis) to all households whose transfer multiplier is greater than a threshold. This threshold is set to make total spending on the policy equal the budget of the CARES stimulus.

MPC workers is effective but leaves much of the gains from targeting on the table.

To the extent that transfer policy bypasses these complications by directly giving income to households, it is easier to target than government purchases. The clear caveat is that government purchases may have direct value. If this is the case, our analysis shows how much stimulus would be sacrificed by that choice.

6.1. Quantifying Gains from MPC-Targeting

The large dispersion in multipliers across sectors in Figure 2 suggests that the potential gains from targeting both transfers and government purchases are quite large. We begin quantifying the potential gains from targeting the highest-multiplier segments of the economy by comparing the maximum purchases and transfers multipliers to those of untargeted purchases and untargeted transfers. For purchases, the change in aggregate private and public consumption due to a shock that targets each industry-region pair proportionally to its value added is 1.28. While this number is sizable, the estimates in Figure 2 demonstrate that had the policymaker instead spent on the state-industry pair with the highest multiplier—which we estimate to be 1.56 in the oil and gas extraction industry in Georgia—the additional consumption induced by the same amount of spending policy would be almost twice as large. For transfer spending, we find that uniformly distributing a dollar to all households would lead to an increase in consumption of 64 cents. In contrast, if the government instead gave that dollar to the group with the highest multiplier—which we estimate is black men in Mississippi aged between 25-35 who earn less than $22,000—it would generate 1.76 additional units of consumption, a 175% increase over the uniform baseline.

While it may be possible for the social planner to achieve the maximum multiplier for small fiscal stimulus programs, for larger programs, the social planner likely will be constrained in the amount that she can transfer to any one segment of the economy. We benchmark the gains from targeting larger transfer schemes by comparing the impact of a CARES-act-like policy—one that transfers $1,200 to each individual making less than $75,000 annually—to more targeted alternatives.[39] Putting these stylized transfers into our model,[40] we find that aggregate consumption increases by 76.4 cents for each dollar spent. Figure 7 shows the multiplier that can be achieved in our model if the government spends the same amount but makes payments of different sizes and targets those payments to households based solely on their multipliers. For example, the value at $2,000 shows the total multiplier that the model predicts if the government gives $2,000 dollars to each worker in reverse order of their multipliers until they exhaust their budget. This calculation shows that, with a maximum transfer size of $1,200, the total multiplier on the income-targeted transfer (0.764) is very close to the multiplier with multiplier-based targeting (0.769), suggesting that income-targeting is effective given the constraint of transferring no more than $1,200 to each individual. However, the government can achieve a higher multiplier by transferring larger amounts to fewer but higher multiplier workers. Indeed, increasing the transfer to $2,500 produces a multiplier of 1.01, more than 30% higher than the benchmark policy, with the same budget (which renders irrelevant questions of how this policy is financed).[41]

7. Conclusion

This paper develops expressions for how fiscal policies affect economic activity in the presence of heterogeneous households and firms and takes these formulae to the data to characterize the dimensions of heterogeneity that affect the efficacy of stimulus policy. We build a Keynesian model with rich household heterogeneity in MPC magnitudes and directions, industrial and spatial linkages, and differential employment sensitivity. All of these elements can be unified into a single, reduced-form network that maps the marginal spending of any given household to the marginal income of factor owners producing the goods the household consumes. We provide a novel decomposition to understand the importance of these rich interconnections, capturing heterogeneity with three corrections to the standard representative-agent Keynesian multiplier.

Empirically, we find that despite a rich regional, input-output and consumption structure, the government can implement maximally expansionary policy by simply targeting either their spending or transfers to households with the highest MPCs. Indeed, we show that concentrating

[39] This is a rough characterization of the CARES act, which included several additional details. Specifically, eligibility depended on household income in the case of married couples and payments depended on the number of dependents. See `https://home.treasury.gov/policy-issues/cares/assistance-for-american-workers-and-families` for the details of the stimulus payments.

[40] We transfer $1,200 (2020 USD) to each employed worker with income below $65,000 (2012 USD).

[41] Of course, an important caveat is that households MPCs could themselves be a function of size of the shock. Using lottery winnings in Norway, Fagereng et al. (2019) find that MPCs fall with the size of the award, with those receiving the equivalent of up to $2,000 US dollars having an average MPC close to 1, those receiving the equivalent of $5,000 having an MPC of around 0.9 and and those receiving the equivalent of $8,000 having an MPC of around 0.5. These estimates loosely suggest that there is substantial scope to increase the size of transfer payments above the $1,200 threshold before substantially altering the magnitude of household MPCs.

transfers among the highest-MPC households can increase the effect of the policy on GDP by up to 175%. Despite the empirical unimportance of the details of the economy for fiscal multipliers, large spillovers make fine-grained modeling of the network structure of the economy necessary for understanding the distributional and spatial impacts of stimulus.

Our results have important implications for the design of fiscal policy. In particular, governments should understand the opportunity costs associated with untargeted fiscal spending. While other important implementation or political constraints may have weighted in favor of uniform stimulus checks, the above analysis suggests that untargeted fiscal policies may leave substantial gains on the table.

A. Omitted Proofs

A.1. Proof of Proposition 1

Proof. Applying the implicit function theorem to the goods market clearing condition, our differentiability assumptions allow us to express

$$\begin{aligned} dQ^1 &= \widehat{\boldsymbol{X}}^{\boldsymbol{1}} dQ^1 + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{h}^{\boldsymbol{1}}} \boldsymbol{\mu} d\ell^1 + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\tau}} d\tau + dG \\ \boldsymbol{\mu} d\ell^1 &= \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1} dY^1 \end{aligned} \tag{A1}$$

Using the definition of ∂Y^1 (Equation 8), recalling the definition $\hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} = \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{h}^{\boldsymbol{1}}}$, and substituting for $d\ell^1$, we have

$$\left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right) dQ^1 = \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1} dY^1 + \partial Y^1 \tag{A2}$$

Finally, recognizing that $dY^1 = (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}) dQ^1$ and solving for dY^1 completes the proof. □

A.2. Proof of Proposition 2

Proof. From Proposition 1 and using that $\mathbb{1}^T \widehat{C}^1 = \mathbb{1}^T$ we have

$$\begin{aligned} dGDP^1 &= \mathbb{1}^T dY^1 \\ &= \mathbb{1}^T \left(\boldsymbol{I} - \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\right)^{-1} \left(dG_1 - \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{\mu} d\tau_1\right) \\ &= \mathbb{1}^T dG_1 + \mathbb{1}^T \left(\boldsymbol{I} - \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\right)^{-1} \left(\hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1} dG^1 - \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{\mu} d\tau_1\right) \\ &= \mathbb{1}^T dG_1 + \mathbb{1}^T \left(\boldsymbol{I} - \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\right)^{-1} \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \partial h^1 \\ &= \mathbb{1}^T dG_1 + \mathbb{1}^T \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \partial h^1 + \underbrace{\mathbb{1}^T \left(\boldsymbol{I} - \hat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\right)^{-1} \hat{\boldsymbol{C}}^{\boldsymbol{1}}}_{b^T \equiv \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1} \hat{\boldsymbol{C}}^{\boldsymbol{1}}\right)^{-1}} \boldsymbol{m} \underbrace{\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1} \hat{\boldsymbol{C}}^{\boldsymbol{1}}}_{\boldsymbol{\mathcal{G}}} \boldsymbol{m} \partial h^1 \\ &= \mathbb{1}^T \Delta G_1 + \left(\mathbb{1}^T + b^T \boldsymbol{m} \boldsymbol{\mathcal{G}}\right) \boldsymbol{m} \partial h^1 \end{aligned} \tag{A3}$$

We now define $\boldsymbol{m}^{\boldsymbol{next}} = \boldsymbol{m} \boldsymbol{\mathcal{G}}$ and $\hat{b}^T = \mathbb{1}^T + b^T \boldsymbol{m}^{\boldsymbol{next}}$. With this, we can write:

$$\begin{aligned} dGDP^1 &= \mathbb{1}^T dG_1 + \mathbb{E}_{\partial h^1}[m_n \hat{b}_n] \\ &= \mathbb{1}^T dG_1 + \mathbb{E}_{\partial h^1}[m_n] \mathbb{E}_{\partial h^1}[\hat{b}_n] + \mathbb{C}\text{ov}_{\partial h^1}[m_n, \hat{b}_n] \\ &= \mathbb{1}^T dG_1 + \mathbb{E}_{\partial h^1}[m_n] \mathbb{E}_{\partial h^1}[1 + b_n m_n^{next}] + \mathbb{C}\text{ov}_{\partial h^1}[m_n, b_n m_n^{next}] \end{aligned} \tag{A4}$$

This exact decomposition in terms of b_n, the Bonacich centralities of $\boldsymbol{m\mathcal{G}}$, can be expressed as an approximate decomposition in terms of MPCs. Concretely, noting that (trivially) $b_n = \frac{1}{1-E_{h*}[m_n]} + O(|m|)$ for any probability vector h^*, we have

$$\begin{aligned} dGDP^1 &= \mathbb{1}^T dG_1 + \mathbb{E}_{\partial h^1}[m_n]\mathbb{E}_{\partial h^1}\left[1 + \frac{m_n^{next}}{1-\mathbb{E}_{h*}[m_n]}\right] + \frac{1}{1-\mathbb{E}_{h*}[m_n]}\mathbb{C}\mathrm{ov}_{\partial h^1}[m_n, m_n^{next}] + O^3(|m|) \\ &= \mathbb{1}^T dG_1 + \mathbb{E}_{\partial h^1}[m_n]\mathbb{E}_{\partial h^1}\left[1 + \frac{m_n^{next} - \mathbb{E}_{h*}[m_n] + \mathbb{E}_{h*}[m_n]}{1-\mathbb{E}_{h*}[m_n]}\right] + \frac{1}{1-\mathbb{E}_{h*}[m_n]}\mathbb{C}\mathrm{ov}_{\partial h^1}[m_n, m_n^{next}] + O^3(|m|) \\ &= \mathbb{1}^T dG_1 + \mathbb{E}_{\partial h^1}[m_n]\left(\frac{1}{1-\mathbb{E}_{h*}[m_n]} + \frac{\mathbb{E}_{\partial h^1}[m_n^{next}] - \mathbb{E}_{h*}[m_n]}{1-\mathbb{E}_{h*}[m_n]}\right) + \frac{1}{1-\mathbb{E}_{h*}[m_n]}\mathbb{C}\mathrm{ov}_{\partial h^1}[m_n, m_n^{next}] + O^3(|m|) \\ &= \mathbb{1}^T dG_1 + \frac{1}{1-\mathbb{E}_{h*}[m_n]}\Bigg[\mathbb{E}_{h*}[m_n] + (\mathbb{E}_{\partial h^1}[m_n] - \mathbb{E}_{h*}[m_n]) + \mathbb{E}_{\partial h^1}[m_n]\left(\mathbb{E}_{\partial h^1}[m_n^{next}] - \mathbb{E}_{h*}[m_n]\right) \\ &\quad + \mathbb{C}\mathrm{ov}_{\partial h^1}[m_n, m_n^{next}]\Bigg] + O^3(|m|) \end{aligned} \tag{A5}$$

Completing the proof. □

A.3. *Proof of Welfare Claims in Section 6*

We first provide a lemma characterizing the marginal change in utilitarian welfare, where the planner places welfare weight $\lambda_n > 0$ on each group. To do this, we define the following objects. First, we define κ_n^t as the marginal value to households of group n of an additional dollar of expenditure at time t. As all prices equal one, this is given by $\kappa_n^t = u_{nc_i}^t$ for some $i \in I$ for which the household has strictly positive consumption. We then define the labor wedge for household n, Δ_n, as $v_n^{1\prime} = \kappa_n^1(1 + \Delta_n)$. A positive wedge indicates that n works as if the wage was higher than it is, i.e., oversupplies labor; a negative wedge represents involuntary un(der)employment. Finally, we define the borrowing wedge for household n, ϕ_n, as

$$\kappa_n^1 = \beta_n \frac{1+r}{1-\phi_n}\kappa_n^2 \tag{A6}$$

If $\phi_n = 0$, then the household lies on the intertemporal Euler equation and the borrowing constraint is slack. If $\phi_n > 0$, then the household is borrowing constrained. Finally, we derive the modified welfare weight $\widetilde{\lambda}_n = \lambda_n \kappa_n^1$, which adjusts the pure preference weight of the planner by the marginal value of expenditure of each group of households n.

Lemma 1. *The change in welfare dW due to a small change in taxes and government purchases in the first period can be expressed as:*

$$dW = \sum_{n\in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 - \left(d\tau_n^1 + (1-\phi_n)\frac{d\tau_n^2}{1+r}\right)\right] \tag{A7}$$

Proof. Differentiating the utilitarian welfare criterion, we have that:

$$\begin{aligned} dW &= \sum_{n\in N} \lambda_n \mu_n \sum_{t=1,2} \beta_n^{t-1} \Big(u_{nc}^t dc_n^t - v_n^{t\prime} d\ell_n^t \Big) \\ &= \sum_{n\in N} \lambda_n \mu_n \sum_{t=1,2} \beta_n^{t-1} \kappa_n^t \left(\mathbb{1}^T dc_n^t - \frac{v_n^{t\prime}}{\kappa_n^t} d\ell_n^t \right) \end{aligned} \tag{A8}$$

Next, note that in the second period, free labor supply implies $v_n^{2\prime} = \kappa_n^2$. In the first period, we have that $v_n^{1\prime} = \kappa_n^1(1+\Delta_n)$. Thus, we have

$$dW = \sum_{n\in N} \lambda_n \kappa_n^1 \mu_n \left[-\Delta_n d\ell_n^1 + \sum_{t=1,2} \frac{\kappa_n^t}{\kappa_n^1} \beta_n^{t-1} \left(\mathbb{1}^T dc_n^t - d\ell_n^t \right) \right] \tag{A9}$$

We note that $\frac{\kappa_n^t}{\kappa_n^1}\beta_n^{t-1} = 1$ for $t = 1$. For $t = 2$, we use the modified Equation A6, and write:

$$dW = \sum_{n\in N} \tilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 + \left(\mathbb{1}^T dc_n^1 - d\ell_n^1 \right) + \frac{1-\phi_n}{1+r} \left(\mathbb{1}^T dc_n^2 - d\ell_n^2 \right) \right] \tag{A10}$$

Differentiating the household's lifetime budget constraint:

$$\mathbb{1}^T dc_n^1 - d\ell_n^1 + \frac{\mathbb{1}^T dc_n^2 - d\ell_n^2}{1+r} = -d\tau_n^1 - \frac{d\tau_n^2}{1+r} \tag{A11}$$

Plugging this in, we have:

$$dW = \sum_{n\in N} \tilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 + \phi_n \left(\mathbb{1}^T dc_n^1 - d\ell_n^1 \right) - (1-\phi_n) \left(d\tau_n^1 + \frac{d\tau_n^2}{1+r} \right) \right] \tag{A12}$$

For households with non-strictly-binding borrowing constraints, $\phi_n = 0$. For households with $\phi_n > 0$, the borrowing constraint binds:

$$\underline{s}_n = l_n^1 - \tau_n^1 - \mathbb{1}^T c_n^1 \implies \mathbb{1}^T dc_n^1 - d\ell_n^1 = -d\tau_n^1 \tag{A13}$$

Thus, we arrive at the final expression:

$$dW = \sum_{n\in N} \tilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 - \left(d\tau_n^1 + (1-\phi_n) \frac{d\tau_n^2}{1+r} \right) \right] \tag{A14}$$

Completing the proof of the lemma. □

We now prove the claim in the text of Section 6. There, we stated that—holding fixed financing (i.e. $d\tau^2 = 0$)—consumption-maximizing planner targets transfers to maximize

$$dY^1 = \sum_{n\in N} m_n \partial h_n^1. \tag{A15}$$

A consumption-maximizing planner is one who is indifferent to redistribution across households—i.e. who sets the modified welfare weights are uniform across households ($\widetilde{\lambda}_n = 1$ for all $n \in N$)—and one who believes that all marginally employed labor comes at no disutility cost—i.e. $\Delta_n = -1$ for all n with $d\ell_n^1 \neq 0$. Under these maintained assumptions, Equation A7 reduces to:

$$dW = \mu^T d\ell^1 - \mu^T d\tau^1 \tag{A16}$$

Moreover, by the formula for the multiplier from Proposition 1, we have that:

$$\boldsymbol{\mu} d\ell^1 = \boldsymbol{\Gamma^1} \left(\boldsymbol{I} - \boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1}\right)^{-1} \left(dG^1 - \boldsymbol{C^1_{h^1}}\boldsymbol{\mu} d\tau^1\right) \tag{A17}$$

where $\boldsymbol{\Gamma^1} = \boldsymbol{R^1_{L^1}}\ \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \ \left(\boldsymbol{I} - \widehat{\boldsymbol{X}^{\boldsymbol{1}}}\right)^{-1}$. Combining these equations and rearranging:

$$\begin{aligned}
dW &= \mathbb{1}^T \boldsymbol{\Gamma^1} \left(\boldsymbol{I} - \boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1}\right)^{-1} \left(dG^1 - \boldsymbol{C^1_{h^1}}\boldsymbol{\mu} d\tau^1\right) - \mu^T d\tau^1 \\
&= \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1}\right)^{-1} dG^1 - \mathbb{1}^T \left[\left(\boldsymbol{I} - \boldsymbol{\Gamma^1}\boldsymbol{C^1_{h^1}}\right)^{-1} \boldsymbol{\Gamma^1}\boldsymbol{C^1_{h^1}} + I\right] \boldsymbol{\mu} d\tau^1 \\
&= \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1}\right)^{-1} dG^1 - \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{\Gamma^1}\boldsymbol{C^1_{h^1}}\right)^{-1} \boldsymbol{\mu} d\tau^1
\end{aligned} \tag{A18}$$

We now use the final maintained assumption of Section 6 that the bias and homophily effects are zero for all possible policies. We first show that, if the bias effects are zero for all purchases and transfers shocks relative to some baseline income incidence h^*, then for each type $n \in N$, either $m_n = 0$ or $m_n^{\text{next}} = \mathbb{E}_{h^*}[m_{n'}]$. To prove this, fix any type $n \in N$ and consider the bias term corresponding to a transfer shock with direct incidence $\partial h^1 = \hat{e}_n$ (i.e., only transfering to n).

$$\text{bias}^{h^*}_{\partial h^1} = \mathbb{E}_{\partial h^1}[m_n] \left(\mathbb{E}_{\partial h^1}[m_n^{\text{next}}] - \mathbb{E}_{h^*}[m_{n'}]\right) = m_n \left(m_n^{\text{next}} - \mathbb{E}_{h^*}[m_{n'}]\right) \tag{A19}$$

The assumption that this is zero then implies that either $m_n = 0$ or $m_n^{\text{next}} = \mathbb{E}_{h^*}[m_{n'}]$.

To apply this fact, recall the definition $m_n^{\text{next}} = (m^T \boldsymbol{\Gamma^1} \widehat{\boldsymbol{C}}^{\boldsymbol{1}})_n$, where $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}$ is the normalized matrix of spending directions, i.e., $\boldsymbol{C^1_{h^1}} = \widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}$. Our previous observation—that for all n, $m_n = 0$ or $m_n^{\text{next}} = \mathbb{E}_{h^*}[m_{n'}]$—then implies that $m^T \boldsymbol{\Gamma^1}\boldsymbol{C^1_{h^1}} = (m^{\text{next}})^T \boldsymbol{m} = \mathbb{E}_{h^*}[m_{n'}] \cdot m^T$.

Applying this fact to the multipliers in Equation A18, we have

$$\begin{aligned}
\mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1}\right)^{-1} &= \sum_{k=0}^{\infty} \mathbb{1}^T \left(\boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1}\right)^k = \mathbb{1}^T + \mathbb{1}^T \boldsymbol{C^1_{h^1}}\boldsymbol{\Gamma^1} + \sum_{k=1}^{\infty} \mathbb{1}^T \boldsymbol{C^1_{h^1}} \left(\boldsymbol{\Gamma^1}\boldsymbol{C^1_{h^1}}\right)^k \boldsymbol{\Gamma^1} \\
&= \mathbb{1}^T + m^T \boldsymbol{\Gamma^1} + \sum_{k=1}^{\infty} \mathbb{E}_{h^*}[m_n]^k m^T \boldsymbol{\Gamma^1} = \left(\mathbb{1} + \frac{1}{1 - \mathbb{E}_{h^*}[m_n]} m\right)^T \boldsymbol{\Gamma^1}
\end{aligned} \tag{A20}$$

Similarly, we have that:

$$\mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{\Gamma^1 C^1_{h^1}}\right)^{-1} = \mathbb{1}^T + \mathbb{1}^T \boldsymbol{\Gamma^1 C^1_{h^1}} + \sum_{k=1}^{\infty} \mathbb{1}^T \boldsymbol{\Gamma^1 C^1_{h^1}} \left(\boldsymbol{\Gamma^1 C^1_{h^1}}\right)^k$$
$$= \mathbb{1}^T + \frac{1}{1 - \mathbb{E}_{h*}[m_n]} m^T = \left(\mathbb{1} + \frac{1}{1 - \mathbb{E}_{h*}[m_n]} m\right)^T \tag{A21}$$

Substituting (A20) and (A21) into Equation A18 shows that:

$$dW = \left(\mathbb{1} + \frac{1}{1 - \mathbb{E}_{h*}[m_n]} m\right)^T \left(\boldsymbol{\Gamma^1} dG^1 - \boldsymbol{\mu} d\tau^1\right) \tag{A22}$$

By budget balance and the assumption that $d\tau^2 = 0$, we moreover have that:

$$\mathbb{1}^T dG^1 - \mathbb{1}^T \boldsymbol{\mu} d\tau^1 = 0 \tag{A23}$$

Thus, as the columns of $\boldsymbol{\Gamma^1}$ sum to 1, we have that:

$$\mathbb{1}^T \boldsymbol{\Gamma^1} dG^1 - \mathbb{1}^T \boldsymbol{\mu} d\tau^1 = 0 \tag{A24}$$

Thus, the change in welfare is simply given by:

$$dW = \frac{1}{1 - \mathbb{E}_{h*}[m_n]} m^T \left(\boldsymbol{\Gamma^1} dG^1 - \boldsymbol{\mu} d\tau^1\right) \tag{A25}$$

yielding the claim given in the text.

References

Acemoglu, D., Azar, P. D., 2020. Endogenous production networks. Econometrica 88(1), 33–82.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80, 1977–2016.

Andersen, A. L., Vestergaard, E. T., Huber, K., Johannesen, N., Straub, L., 2022. Disaggregated economic accounts. Tech. Rep. w30630, National Bureau of Economic Research.

Ascari, G., Beck-Friis, P., Florio, A., Gobbi, A., 2023. Fiscal foresight and the effects of government spending: It's all in the monetary-fiscal mix. Journal of Monetary Economics 134, 1–15.

Auclert, A., 2019. Monetary policy and the redistribution channel. American Economic Review 109, 2333–67.

Auclert, A., Rognlie, M., Straub, L., 2018. The intertemporal keynesian cross. Tech. rep., National Bureau of Economic Research.

Auerbach, A., Gorodnichenko, Y., Murphy, D., 2020. Local fiscal multipliers and fiscal spillovers in the usa. IMF Economic Review 68, 195–229.

Baqaee, D., Farhi, E., 2020. Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. Tech. rep.

Baqaee, D. R., 2015. Targeted fiscal policy. Tech. rep.

Baqaee, D. R., 2018. Cascading failures in production networks. Econometrica 86, 1819–1838.

Baqaee, D. R., Farhi, E., 2018. Macroeconomics with heterogeneous agents and input-output networks. Tech. rep., National Bureau of Economic Research.

Baqaee, D. R., Farhi, E., 2019. The macroeconomic impact of microeconomic shocks: beyond hulten's theorem. Econometrica 87, 1155–1203.

Barattieri, A., Cacciatore, M., Traum, N., 2023. Estimating the effects of government spending through the production network. Tech. rep., National Bureau of Economic Research.

Barro, R. J., 2025. The old keynesian model. Tech. rep., National Bureau of Economic Research.

Barro, R. J., Grossman, H. I., 1971. A general disequilibrium model of income and employment. The American Economic Review 61, 82–93.

Bianchi, F., Faccini, R., Melosi, L., 2023. A fiscal theory of persistent inflation*. The Quarterly Journal of Economics 138, 2127–2179.

Bigio, S., La'O, J., 2020. Distortions in Production Networks. The Quarterly Journal of Economics Qjaa018.

Bilbiie, F. O., 2019. The new keynesian cross. Journal of Monetary Economics .

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. American Economic Review 98, 1887–1921.

Board of Governors of the Federal Reserve System, 2013. Flow of funds accounts of the united states, statistical release z.1: L.213 corporate equities. Release date March 7, 2013; accessed June 8, 2026.

Bouakez, H., Rachedi, O., Santoro, E., 2020. The sectoral origins of the spending multiplier. CIREQ, Université de Montréal.

Bouakez, H., Rachedi, O., Santoro, E., 2023. The government spending multiplier in a multisector economy. American Economic Journal: Macroeconomics 15, 209–239.

Caliendo, L., Parro, F., Rossi-Hansberg, E., Sarte, P.-D., 2018. The impact of regional and sectoral productivity changes on the us economy. The Review of economic studies 85, 2042–2096.

Carvalho, V. M., Nirei, M., Saito, Y., Tahbaz-Salehi, A., 2016. Supply chain disruptions: Evidence from the great east japan earthquake. Tech. rep.

Chodorow-Reich, G., 2019. Geographic cross-sectional fiscal spending multipliers: What have we learned? American Economic Journal: Economic Policy 11, 1–34.

Chodorow-Reich, G., Nenov, P. T., Simsek, A., 2021. Stock market wealth and the real economy: A local labor market approach. American Economic Review 111, 1613–1657.

Clower, R., 1965. The keynesian counter-revolution: a theoretical appraisal. Published in Hahn and Brechling eds., The Theory of Interest Rates pp. 103–125.

Corbi, R., Papaioannou, E., Surico, P., 2019. Regional transfer multipliers. The Review of Economic Studies 86, 1901–1934.

Cox, L., Müller, G. J., Pasten, E., Schoenle, R., Weber, M., 2019. Big g. Tech. rep.

Cox, L., Müller, G. J., Pasten, E., Schoenle, R., Weber, M., 2024. Big g. Journal of Political Economy 132, 3260–3297.

Dupor, B., Karabarbounis, M., Kudlyak, M., Mehkari, M., 2018. Regional consumption responses and the aggregate fiscal multiplier. Available at SSRN 3131248 .

Elliott, M., Golub, B., Leduc, M. V., 2020. Supply network formation and fragility. arXiv preprint arXiv:2001.03853 .

Fagereng, A., Holm, M. B., Natvik, G. J. J., 2019. Mpc heterogeneity and household balance sheets. Tech. rep.

Farhi, E., Werning, I., 2017. Fiscal unions. American Economic Review 107, 3788–3834.

Feyrer, J., Mansur, E. T., Sacerdote, B., 2017. Geographic dispersion of economic shocks: Evidence from the fracking revolution. American Economic Review 107, 1313–34.

Gabaix, X., 2011. The granular origins of aggregate fluctuations. Econometrica 79, 733–772.

Greenwood, J., Hercowitz, Z., Huffman, G. W., 1988. Investment, capacity utilization, and the real business cycle. The American Economic Review pp. 402–417.

Gruber, J., 1997. The consumption smoothing benefits of unemployment insurance. The American Economic Review 87, 192.

Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I., 2020. Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? NBER Working Papers 26918, National Bureau of Economic Research, Inc.

Guvenen, F., Smith, A. A., 2014. Inferring labor income risk and partial insurance from economic choices. Econometrica 82, 2085–2129.

Hastings, J., Shapiro, J. M., 2018. How are snap benefits spent? evidence from a retail panel. American Economic Review 108, 3493–3540.

Hubmer, J., 2019. The race between preferences and technology. Tech. rep., University of Pennsylvania.

Internal Revenue Service, Statistics of Income Division, 2012. Soi tax stats: Individual income tax statistics, 2012 zip code data. Accessed May 13, 2026.

Jappelli, T., Pistaferri, L., 2014. Fiscal policy and mpc heterogeneity. American Economic Journal: Macroeconomics 6, 107–36.

Kaplan, G., Moll, B., Violante, G. L., 2018. Monetary policy according to hank. American Economic Review 108, 697–743.

Lewis, D. J., Melcangi, D., Pilossoph, L., 2019. Latent heterogeneity in the marginal propensity to consume. FRB of New York Staff Report .

Long, J. B., Plosser, C. I., 1987. Sectoral vs. aggregate shocks in the business cycle. The American Economic Review 77, 333–336.

Miyazawa, K., 1976. Input-output analysis and interrelational income multiplier as a matrix. In: *Input-Output Analysis and the Structure of Income Distribution*, Springer, pp. 22–42.

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

Parker, J. A., Souleles, N. S., Johnson, D. S., McClelland, R., 2013. Consumer spending and the economic stimulus payments of 2008. American Economic Review 103, 2530–53.

Patinkin, D., 1949. Involuntary unemployment and the keynesian supply function. The Economic Journal pp. 360–383.

Patterson, C., 2023. The matching multiplier and the amplification of recessions. American Economic Review 113, 982–1012.

Proebsting, C., 2022. Market segmentation and spending multipliers. Journal of Monetary Economics 128, 1–19.

Ramey, V. A., 2011. Can government purchases stimulate the economy? Journal of Economic Literature 49, 673–85.

Rubbo, E., 2019. Networks, phillips curves, and monetary policy. Tech. rep., Harvard University.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Rivera Drew, J. A., Richards, S., Rogers, R., Schroeder, J., Williams, K. C. W., 2025. Ipums usa: Version 16.0. Extract used in this paper includes ACS samples; accessed May 7, 2026.

Stiglitz, J. E., 1970. Non-substitution theorems with durable capital goods. The Review of Economic Studies 37, 543–553.

U.S. Bureau of Economic Analysis, 2012. Input-output accounts data: 2012 use, make, and imports tables. BEA Input-Output Interactive Data Application, tables used: use table, make table, and imports table for 2012; accessed April 30, 2026.

U.S. Census Bureau, 2015. 2012 commodity flow survey datasets: 2012 cfs public use file. Accessed March 11, 2026.

Werning, I., 2015. Incomplete markets and aggregate demand. Tech. rep., National Bureau of Economic Research.

Woodford, M., 2020. Effective demand failures and the limits of monetary stabilization policy. Tech. rep.

Supplemental Appendix to *Fiscal Policy in a Networked Economy* by Becko, Flynn, and Patterson

B. Rationing Equilibrium Microfoundation

In this appendix, we set up a generalization of our baseline model to allow for fully general rationing on either the supply or the demand side of the market when interest rates do not adjust to clear labor markets. We show that this model reduces to our baseline model under the following four economic conditions: 1) agents cannot be forced to work more or employ more workers than they would like, 2) rationing is minimally efficient in the sense that it meets either demand or supply, 3) aggregate labor demand is smaller than aggregate labor supply, and 4) there are no income effects in household labor supply. Finally, we sketch an extension of the model that accommodates multiple types of labor.

B.1. A general model of labor rationing

As in the main text, a finite number of competitive household types $n \in N$ with mass μ_n and firms $i \in I$ respectively supply and demand a homogeneous labor factor in order to produce goods over two periods $t \in \{1,2\}$. Wages in both periods are sticky and so are normalized to one, i.e., $w^t = 1$. Due to a binding ZLB on nominal interest rates, the real interest rate r is exogenous. We denote by $p^t = \{p_i^t\}_{i \in I}$ the vector of prices in period t.

Given prices, each household n forms a hypothetical first-period labor supply ℓ_n^{*1} consistent with optimization subject to only (a) a borrowing constraint in the form of minimum savings $\underline{s}_n$ and (b) a budget constraint incorporating lump-sum taxes $\tau = (\tau_n^1, \tau_n^2)$.

$$
\begin{aligned}
\ell_n^{*1} \in \arg\max_{\ell^1} \max_{\ell^2, c^1, c^2} & \sum_{t=1,2} \beta_n^{t-1} u_n^t(c^t, \ell^t) \\
\text{s.t. } & p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell^1 + \frac{\ell^2}{1+r} \\
& \ell^1 - p^1 c^1 - \tau_n^1 \geqslant \underline{s}_n
\end{aligned}
\tag{A26}
$$

Similarly, each firm i forms a hypothetical first-period labor demand L_n^{*1} consistent with profit maximization, given a CRS production function $F_i^{t=1}$ that incorporates the single labor factor as well as a vector of inputs (any goods).

$$
L_i^{*1} \in \arg\max_{L^1} \max_{X^1} \; p_i^1 F_i^1(X^1, L^1) - p^1 X^1 - L^1 \tag{A27}
$$

Next, a non-price mechanism that we refer to as the *rationing function* assigns to each household and firm its realized first-period labor supply and demand, respectively, as a function of the vectors $\ell^{*1} = \{\ell_n^{*1}\}_{n \in N}$ and $L^{*1} = \{L_i^{*1}\}_{i \in I}$ of all hypothetical, preferred labor supplies and demands, respectively.

$$\ell_n^1 = R_n^S\left(\ell^{*1}, L^{*1}\right) \qquad\qquad L_i^1 = R_i^D\left(\ell^{*1}, L^{*1}\right) \tag{A28}$$

Although we remain agnostic to the details of labor assignment, we assume the rationing function assigns equal amounts of labor supply and demand, i.e., $\sum_{n \in N} \mu_n R_n^S\left(\ell^{*1}, L^{*1}\right) = \sum_{i \in I} R_i^D\left(\ell^{*1}, L^{*1}\right)$.

Taking its rationed labor supply as given, each household n chooses second-period labor supply ℓ_n^2 and vectors of first- and second-period consumption $c_n^t = \{c_{ni}^t\}_{i \in I}$.

$$\begin{aligned}
(\ell_n^1, \ell_n^2, c_n^1, c_n^2) \in \underset{\ell^1, \ell^2, c^1, c^2}{\arg\max} & \sum_{t=1,2} \beta_n^{t-1} u_n^t(c^t, \ell^t) \\
\text{s.t. } & p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell^1 + \frac{\ell^2}{1+r} \\
& \ell^1 - p^1 c^1 - \tau_n^1 \geqslant \underline{s}_n \\
& \ell^1 = \ell_n^1
\end{aligned} \tag{A29}$$

Taking its rationed labor demand as given, each firm i chooses a vector of first-period inputs $X_i^1 = \{X_{ij}^1\}_{j \in I}$ to maximize profits; in the second period, i chooses both labor L^2 and inputs $X_i^2 = \{X_{ij}^2\}_{j \in I}$ to maximize profits given a CRS production function $F_i^{t=2}$.

$$\begin{aligned}
(X_i^t, L_i^t) \in \underset{X^t, L^t}{\arg\max} \; & p_i^t F_i^t(X^t, L^t) - p^t X^t - L^t \\
\text{s.t. } & L^1 = L_i^1
\end{aligned} \tag{A30}$$

In addition to levying lump-sum taxes τ_n^t on households, the government purchases G_i^t units of good $i \in I$ subject to running a balanced budget over the two periods.

$$\sum_{n \in N} \mu_n \left(\tau_n^1 + \frac{1}{1+r}\tau_n^2\right) = p^1 G^1 + \frac{1}{1+r} p^2 G^2 \tag{A31}$$

Finally, goods markets and the second period labor market clear.

$$F_i^t(X_i^t, L_i^t) = \sum_{j \in I} X_{ji}^t + \sum_{n \in N} \mu_n c_{ni}^t + G_i^t, \qquad\qquad \sum_{i \in I} L_i^2 = \sum_{n \in N} \mu_n l_n^2 \tag{A32}$$

A *general* rationing equilibrium is therefore defined as follows:

Definition 2. *Given a profile of government spending and transfers* $\{\tau_n^t, G_i^t\}_{t \in \{0,1\}, n \in N, i \in I}$ *that satisfies (A31), a general rationing equilibrium is a profile of prices and quantities* $\{p_i^t, \ell_n^{*1}, L_i^{*1}, l_n^t, c_n^t, L_i^t, X_{ij}^t\}_{t \in \{0,1\}, n \in N, i, j \in I}$ *that satisfy conditions (A26) – (A32).*

B.2. Microfounding the Reduced-Form Model

In this section we show how the more reduced-form model of rationing in the main text can be obtained as a special case of the more general rationing model above. This occurs under three conditions: That the rationing function satisfies feasibility and allocative efficiency properties, demand is deficient, and that there are no income effects in labor supply.

We begin by imposing two intuitive assumptions on the rationing function.

Assumption 1. *The rationing function (R^S, R^D) satisfies the following conditions:*

1. *Feasibility: For all ℓ^{*1} and L^{*1}, for all $i \in I$ and $n \in N$,*

$$R_n^S\left(\ell^{*1}, L^{*1}\right) \leqslant \ell_n^{*1} \qquad \text{and} \qquad R_i^D\left(\ell^{*1}, L^{*1}\right) \leqslant L_i^{*1}. \tag{A33}$$

2. *Allocative efficiency: For all ℓ^{*1} and L^{*1}, there do not exist $i \in I$ and $n \in N$ such that*

$$R_n^S\left(\ell^{*1}, L^{*1}\right) < \ell_n^{*1} \qquad \text{and} \qquad R_i^D\left(\ell^{*1}, L^{*1}\right) < L_i^{*1}. \tag{A34}$$

Feasibility captures the idea that firms and households may not be forced to hire more workers or work more than they wish, respectively. Allocative efficiency posits that the rationing mechanism ensures there are not firms and workers who would like to be matched, but are not. We view this as a reasonable assumption when wage stickiness—and not for example search frictions—is the main friction.

We now impose that we are in a region where aggregate ideal labor demand is less than aggregate ideal labor supply.

Assumption 2. *There is strictly deficient demand in any equilibrium $\sum_{i\in I} L_i^{*1} < \sum_{n\in N} \mu_n \ell_n^{*1}$.*

Finally, we assume that households have the intratemporal preferences of Greenwood, Hercowitz, and Huffman (1988), so that there are no income effects in labor supply:

Assumption 3. *Household preferences are GHH, i.e., for all $n \in N, t \in \{1, 2\}$, there exist concave functions U_n^t, strictly increasing and homothetic functions ν_n^t, and strictly convex functions Λ_n^t, all real-valued and increasing, such that for all consumption vectors c and labor supplies ℓ,*

$$u_n^t(c, \ell) = U_n^t\left(\nu_n^t(c) - \Lambda_n^t(\ell)\right). \tag{A35}$$

The following result shows that rationing equilibria and general rationing equilibria coincide under these assumptions:

Proposition 3. *Under Assumptions 1, 2, and 3, each household n has a unique preferred level of labor supply $\ell_n^{*1} = \bar{\ell}_n^{*1}$, independent of government spending and transfer policies. Moreover, under* *any government spending and transfer policies, we have that* $\mu_n \ell_n^1 = R_n^1(L^{*1}) \equiv R_n^S(\bar{\ell}^{*1}, L^{*1})$ *and* $L_i^1 = L_i^{*1}$*. Thus, the set of rationing equilibria and general rationing equilibria are equal.*

Proof. Since, by Assumption 3, ν_n^t is strictly increasing and homothetic, and since p^t is independent of policy, each household consumes a fixed basket of goods α_n^t with non-negative weights on each good. We can therefore write $\nu_n^t(c^t) = \nu_n^t(\{\alpha_{ni}^1 \tilde{c}_n\}_{i\in\mathcal{I}}) = \omega_n^t \tilde{c}_n^t$ for some constant $\omega_n^t > 0$. Each household's choice of consumption can then be reduced to choosing $\tilde{c}_n^t$. Recall we can express the household's budget constraint as:

$$p^1 \alpha_n^1 \tilde{c}^1 + \frac{p^2 \alpha_n^2 \tilde{c}^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell^1 + \frac{\ell^2}{1+r} \tag{A36}$$

By the intratemporal Euler equation of the household, by strict convexity of Λ_n^t, we can express its optimal labor supply as:

$$\ell_n^{*t} = \bar{\ell}_n^{*t} \equiv \left(\Lambda_n^{t\,\prime}\right)^{-1} \left(\omega_n^t\right) \tag{A37}$$

if $\left(\Lambda_n^{t\,\prime}\right)^{-1} \left(\omega_n^t\right)$ lies in the positive reals and zero otherwise. The definition of generalized rationing equilibrium therefore implies $\mu_n \ell_n^1 = R_n^S(\ell^{*1}, L^{*1}) = R_n^S(\bar{\ell}^{*1}, L^{*1}) \equiv R_n^1(L^{*1})$, as desired.

By Assumption 2, we have that demand is strictly deficient $\sum_{i\in I} L_i^{*1} < \sum_{n\in N} \ell_n^{*1}$. Next, we claim Assumption 1 implies that if there is strictly deficient demand, then $L_i^1 = L_i^{*1}$. To see this, first note that the feasibility property and deficient demand together imply there exists some n such that $\ell_n^1 < \ell_n^{*1}$. But then if there is some i such that $L_i^1 < L_i^{*1}$, allocative efficiency cannot hold, a contradiction. □

B.3. Extension to multiple labor types and/or labor markets

The microfoundation above can be extended to accommodate many labor types (firm preferences over workers) and many labor markets (worker preferences over firms) under a two key additional assumptions. We sketch the arguments here, as they are essentially identical to the ones above. First, *all* first period wages are rigid, as are the expectations of all second period wages (or their absolute and relative inflation rates). This prevents any adjustment through prices, as in the one-factor model. Second, there is deficient demand not in just one labor market, but in all of them. This guarantees that firms are rationed the labor they demand. Again, households' GHH preferences ensure that fiscal policies do not affect their preferred levels of labor supply locally, so that we obtain the same formulation as before. Since the rationing function in the reduced-form model in the main text conditions on the identities of all workers and all firms, we may interpret it as matching the appropriate labor types to firms in the appropriate labor markets, under these conditions of extreme rigidity and demand deficiency.

C. Additional Results

In this appendix, we present results on properties (including existence) of rationing equilibrium (C.1), provide comparative statics for the multiplier (C.2), and analyze benchmark cases in which various network adjustments to the Keynesian multiplier are zero (C.3).

C.1. Equilibrium Properties

In this appendix, we ensure our analysis of the multiplier is well-posed and eliminate any nuisance terms that unnecessarily complicate the analysis. To this end, we first provide a non-substitution theorem that ensures prices are technologically determined—and thus independent of fiscal policies—and, second, prove the existence of a rationing equilibrium.

The following technical conditions on production technologies and household preferences are sufficient for the non-substitution theorem. Assumption 4 provides basic technical conditions on production and Assumption 5 imposes a simple positivity condition on demand such that there is demand for all goods.

Assumption 4. *For all $i \in N$ and $t \in \{1, 2\}$, the production functions F_i^t are continuous, weakly increasing, strictly quasi-concave, and homogeneous of degree one. Further, labor is essential in production, i.e., $F_i^t(X_i^t, 0) \equiv 0$, and production is strictly increasing in labor. Finally, for $t \in \{1, 2\}$, there exists some $\overline{p}^t \in \mathbb{R}^I_{>0}$ and $\{\overline{X}_i^t, \overline{L}_i^t\}_{i\in I}$ s.t. for all i, $F_i^t(\overline{X}_i^t, \overline{L}_i^t, z_i^t) \geqslant 1$ and $\overline{p}^t\overline{X}_i^t + \overline{L}_i^t \leqslant \overline{p}_i^t$.*[42]

Assumption 5. *For any p, ℓ^1, τ, θ: for each good i and time t, some household type n has $c_{ni}^t > 0$.*

Under these two rather weak assumptions, we can show that:

Proposition 4. *Fix a profile of government spending and transfers $\mathcal{P} = \{\tau_n^t, G_i^t\}_{t\in\{0,1\},n\in N,i\in I}$ that satisfies (4) and suppose that Assumptions 4 and 5 hold. There exists a unique $p^t(\mathcal{P}) \in \mathbb{R}^I_{\geqslant 0}$ consistent with rationing equilibrium and, moreover, $p^t(\mathcal{P}) \gg 0$. Furthermore, $p^t(\mathcal{P})$ is invariant to $\mathcal{P}$.*

Proof. We prove the existence of demand-independent prices with lattice-theoretic argument similar to that of Acemoglu and Azar (2020) and their uniqueness with an argument similar to that of Stiglitz (1970).

Preliminaries. Fix a time period $t \in \{1, 2\}$ and a firm $i \in I$. We define i's unit cost function (at this time and technology) $\kappa_i : \mathbb{R}^I_{>0} \to \mathbb{R}$ as the function that maps any strictly positive vector of prices to i's least cost of production:

$$\kappa_i^t(p) \equiv \min_{X_i^t, L_i^t \geqslant 0,\ F_i^t(X_i^t, L_i^t)\geqslant 1} pX_i^t + L_i^t. \tag{A38}$$

[42] A sufficient but not necessary condition is that every good can be produced using only labor.

This function is well-defined because, by Assumption 4, i can produce a unit of output with the input bundle $\overline{X}_i^t, \overline{L}_i^t$. Thus, we may restrict the domain over which the firm optimizes in (A38) to the set

$$\left\{ X_i^t \in \mathbb{R}_{\geqslant 0}^I, L_i^t \in \mathbb{R}_{\geqslant 0} \;\middle|\; \forall j \in I, X_{ij}^t \leqslant \frac{\overline{\kappa}_i^t(p)}{p_j},\ L_i^t \leqslant \overline{\kappa}_i^t(p),\ F_i^t(X_i^t, L_i^t) \geqslant 1 \right\} \tag{A39}$$

where $\overline{\kappa}_i^t(p) \equiv p\overline{X}_i^t + \overline{L}_i^t$. This set is compact by continuity of F_i^t (from Assumption 4). Thus, since the objective of (A38) is continuous, a minimum exists by Weierstrass' Theorem.

Now note that by Assumption 5 and market clearing, each firm $i \in I$ has strictly positive output in any equilibrium. Firm optimization and CRS (homogeneity of degree one of F_i^t from Assumption 4) thus imply that each firm's price is equal to its marginal cost, which also equals its average cost. Since labor is essential (from Assumption 4) and has a wage normalized to one, this cost must moreover be strictly positive in any equilibrium. We conclude that any equilibrium price vector at time $t \in \{1, 2\}$ $p^t \in \mathbb{R}_{\geqslant 0}^I$ must be strictly positive and satisfy:

$$\forall i \in I, \quad p_i^t = \kappa_i^t(p^t). \tag{A40}$$

The remainder of the proof shows that such a price vector exists, is unique, and invariant to $\mathcal{P}$.

Existence. Recall the price vector $\overline{p}^t$ from Assumption 4. By that assumption and the definition of κ_i^t, we have

$$\forall i \in I, \quad \kappa_i^t(\overline{p}^t) \leqslant \overline{p}_i^t. \tag{A41}$$

Moreover, we claim that there exists an $\alpha > 0$ such that:

$$\forall i \in I, \quad \kappa_i^t(\alpha\overline{p}^t) \geqslant \alpha\overline{p}_i^t. \tag{A42}$$

To see this, suppose it does not hold. Then there exists a sequence $\{\alpha^k\}$ such that $\alpha^k \to 0$, a firm i, and a sequence of input bundles $\{X_i^k, L_i^k\}$ such that $X_i^k \in \mathbb{R}_{\geqslant 0}^I, L_i^k \in \mathbb{R}_{\geqslant 0}$ and for all $k \in \mathbb{N}$,

$$F_i^t(X_i^k, L_i^k) \geqslant 1 \qquad \text{and} \qquad \alpha\overline{p}^t X_i^k + L_i^k < \alpha\overline{p}_i^t. \tag{A43}$$

Since by assumption all components of $\overline{p}$ are strictly positive, this implies $L_i^k \to 0$ and each $X_i^k \leqslant \frac{\overline{p}_i^t}{\overline{p}^t}$. But then F_i^t's continuity and the Bolzano-Weierstrass Theorem imply that $F_i^t(X_i^\infty, 0) \geqslant 1$ at some limit point X_i^∞ of $(X_i^k)_{k\in\mathbb{N}}$. This contradicts that labor is essential (Assumption 4).

Since—as is immediate from the definition (A38)—each cost function κ_i is monotonically increasing, the observations above imply that the function $\kappa = (\kappa_1, ..., \kappa_{|I|})$ maps the set

$$\mathcal{D} \equiv \prod_{i \in I} [\alpha\overline{p}_i, \overline{p}_i] \tag{A44}$$

to itself for some $\alpha > 0$, namely one sufficiently small that the previous argument carries. Since $\mathcal{D}$ endowed with the standard component-wise partial order is a complete lattice, Tarski's fixed point theorem implies that the set of fixed points of κ on $\mathcal{D}$ is a complete lattice. In particular, it is non-empty. This implies the existence of a strictly positive solution to the fixed-point problem (A40).

Uniqueness. Suppose that p^t and $p^{t\prime}$ are two strictly positive solutions (A40). We will show $p^t = p^{t\prime}$. Let $(X_i^t)_{i\in I}$ and $(L_i^t)_{i\in I}$ be any vectors of firm-specific cost-minimizing unit input demands at p^t, and let $\widehat{\boldsymbol{X}}^{\boldsymbol{t}}$ be the $I \times I$ matrix with (i,j) entry X_{ji}^t and $\widehat{\boldsymbol{L}}^{\boldsymbol{t}}$ be the $I \times I$ diagonal matrix with ii^{th} entry L_i^t. The price-cost fixed point equation (A40) at p^t can therefore be written:

$$p^t = (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T p^t + \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1} \tag{A45}$$

We moreover claim that $\boldsymbol{I} - (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T$ is invertible, so that we have

$$p^t = \left(\boldsymbol{I} - (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T\right)^{-1} \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1}. \tag{A46}$$

To see why, recall that since labor is essential for each good, $p_i^t > ((\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T p^t)_i$ for all $i \in I$. This implies $\boldsymbol{I} - (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T$ is invertible, because if for any vector v, $v = (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T v$, then for all $\alpha \in \mathbb{R}$ such that $p^t + \alpha v \geqslant 0$, the fact that $(\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T$ has all positive entries implies then $p^t + \alpha v > (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T (p^t + \alpha v) \geqslant 0$. Since p^t is strictly positive, this is only possible when $v = 0$.

By cost-minimization, we have the component-wise inequality

$$p^{t\prime} \leqslant (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T p^{t\prime} + \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1}, \qquad \text{which implies} \qquad p^{t\prime} \leqslant \left(\boldsymbol{I} - (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T\right)^{-1} \widehat{\boldsymbol{L}}^{\boldsymbol{t}} = p^t. \tag{A47}$$

The same argument implies $p^{t\prime} \geqslant p^t$, and so $p^t = p^{t\prime}$.

Invariance. We observe that the previous step constructed the unique equilibrium price vector and this construction was invariant to $\mathcal{P}$, completing the proof. □

The existence of unique, positive prices $p^t \in \mathbb{R}^I_>$ consistent with rationing equilibrium and that are invariant to policy is what allows us to set $p_i^t = 1$ for all $i \in I$ and $t \in \{1, 2\}$ in our analysis. Moreover, combining Proposition 4 with constant returns to scale technology implies a simple form for aggregate input and labor demands. To see this, first define the unit input and labor demands

$$\left(\widehat{X}_i^t, \widehat{L}_i^t\right) \equiv \arg \min_{(X_i^t, L_i^t) \text{ s.t. } F_i^t(X_i^t, L_i^t) \geqslant 1} p^t X_i^t + L_i^t \tag{A48}$$

where by Proposition 4, p^t is the unique price vector consistent with rationing equilibrium. These demands are well-defined because (a) a minimum cost exists by the argument in the first step of the proof of Proposition 4 and (b) production is strictly quasi-concave by Assumption 4. We now claim:

Corollary 1. *Suppose that Assumptions 4 and 5 hold. In any rationing equilibrium, input demands X^t and labor demands L^t are given by:*

$$X^t = \widehat{\boldsymbol{X}}^{\boldsymbol{t}} Q^t \quad L^t = \widehat{\boldsymbol{L}}^{\boldsymbol{t}} Q^t \tag{A49}$$

where $\widehat{\boldsymbol{X}}^{\boldsymbol{t}}$ is the matrix with i^{th} column $\widehat{X}_i^t$ and $\widehat{\boldsymbol{L}}^{\boldsymbol{t}}$ is the diagonal matrix with (i,i) entry $\widehat{L}_i^t$.

Proof. CRS implies that for a firm producing Q_i^t units in equilibrium,

$$X_i^t = Q_i^t \widehat{X}_i^t \quad L_i^t = Q_i^t \widehat{L}_i^t \tag{A50}$$

Stacking these equations over I yields the result. □

Proposition 4 implies two additional, useful results. First, the Leontief-inverse matrix always exists. Second, one can use the Leontief-inverse to obtain a useful closed-form expression for the policy-independent prices. This is stated formally in the following corollary:

Corollary 2. *Suppose that Assumptions 4 and 5 hold. For any $t \in \{1,2\}$, the Leontief-inverse matrix $(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{t}})^{-1}$ exists. Moreover, prices are given uniquely by the following expression:*

$$p^t = \left(\boldsymbol{I} - (\widehat{\boldsymbol{X}}^{\boldsymbol{t}})^T\right)^{-1} \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1} \tag{A51}$$

Proof. This follows from the argument in the "Uniqueness" step of the proof of Proposition 4 that any unit-inputs matrix corresponding to non-substitution prices is invertible. □

We now proceed to establish that the analysis of equilibrium is well posed by providing regularity conditions under which equilibria exist. To this end, we assume basic continuity properties of demand and that household consumption in the first period is bounded away from fully consuming first period income as income grows large.

Because it does not complicate the analysis, we prove the result when consumption and second-period labor supply take the general functional forms $c_n^t = c_n^t(p, \ell_n^1, \tau_n)$ and $\ell_n^2 = \ell_n^2(p, \ell_n^1, \tau_n)$.

Assumption 6. *The primitives satisfy the following properties:*

1. *The consumption and labor functions c_n^t and ℓ_n^2 are continuous in ℓ_n^1.*
2. *For any p, τ: there exists $\overline{y} \in \mathbb{R}_+$ and $\overline{c} < 1$ such that for all $n \in N$ and $\ell_n^1 > \overline{y}$, we have that $p^1 c_n^1(p, \ell_n^1, \tau_n) \leqslant \overline{c}\ell_n^1$.*

With this additional structure, we are now able to prove the existence of rationing equilibria for the economy under consideration.

Proposition 5. *Fix a profile of government spending and transfers* $\mathcal{P} = \{\tau_n^t, G_i^t\}_{t\in\{0,1\},n\in N,i\in I}$ *that satisfies (4) and suppose that Assumptions 4, 5, and 6 hold. There exists a rationing equilibrium.*

Proof. In each period $t \in \{1, 2\}$, prices p^t are given by Corollary 2. Moreover, by Assumption 6 we have the following fact: for any p^1, p^2, τ: there exists some $\overline{y} \in \mathbb{R}_+$ and some $\overline{c} < 1$ such that $p^1 c_n^1(p, \ell_n^1, \tau_n) \leqslant \overline{c}\ell_n^1$ for all n, $\ell_n^1 > \overline{y}$.

Thus, first period aggregate spending $(p^1)^T C^1 + (p^1)^T G^1$ satisifes:

$$(p^1)^T C^1 + (p^1)^T G^1 \leqslant \overline{c}\,\overline{y} + \overline{c}\mu^T \ell^1 + (p^1)^T G^1 \tag{A52}$$

Since $\overline{c} < 1$, it follows that there exists $\overline{Y}$ such that if $\ell^1 \in \mathcal{Y}^1 \equiv \{\ell^1 \in \mathbb{R}_+^N \mid \mu^T \ell^1 \leqslant \overline{Y}\}$, then aggregate spending is weakly less than $\overline{Y}$; since all spending flows to wages, i.e., $\mu^T \ell^1 = \mathbb{1}^T L^1 = (p^1)^T (C^1 + G^1)$, aggregate income is also then less than $\overline{Y}$. Formally:

$$\forall \ell^1 \in \mathcal{Y}^1 : R^1 \left(\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(1 - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1} \left(C^1(p, \ell^1, \tau) + G^1\right)\right) \in \mathcal{Y}^1 \tag{A53}$$

This observation allows us to define a function $\Psi : \mathcal{Y}^1 \to \mathcal{Y}^1$ given by:

$$\Psi(\ell^1) = R^1 \left(\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(1 - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1} \left(C^1(p, \ell^1, \tau) + G^1\right)\right) \tag{A54}$$

where the previous argument establishes that $\Psi(\ell^1)$ is indeed contained in $\mathcal{Y}^1$. Moreover, continuity of $C^1(\cdot)$ establishes that Ψ is a continuous function.

Since Ψ is a continuous function on a compact, convex domain, it has a fixed point ℓ^1 by Brouwer's theorem.

Given this fixed point ℓ^1 of Ψ, we can construct a rationing equilibrium as follows: Let p^t be the non-substitution-theorem prices. Let c_n^t and ℓ_n^2 by the relevant functions taking in prices p^t and incomes ℓ^1. Let production in each period be:

$$Q^t = (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{t}})^{-1}(G^t + C^t) \tag{A55}$$

The definitions of the consumption and labor supply functions ensure that household budget constraints hold. The construction of Q^t ensures that each goods market clears. Because ℓ^1 is a fixed point, first period income is consistent with the rationing function and the first period labor market clears. Finally, the second period labor market clears by Walras' law. $\square$

As we have established natural conditions under which an equilibrium exists, our analysis of equilibria is well-posed.

C.2. Comparative Statics for the Multiplier

We use the structure of the multiplier from Proposition 1 to provide comparative statics of the multiplier in the various objects that contribute towards it. To this end, define the matrix:

$$\mathcal{M} = \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^1}\boldsymbol{l}^{\mathbf{1}}_{\boldsymbol{L}^1}\widehat{\boldsymbol{L}}^{\mathbf{1}}\left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{\mathbf{-1}} \tag{A56}$$

and assume that all entries of this matrix are non-negative.[43] $\mathcal{M}$ serves the role of a generalized MPC in our multiplier expression. We first consider the effect of arbitrary changes in this object on the response of final output to an arbitary shock.

Proposition 6. *Consider a change in the economy such that $\mathcal{M}$ is replaced with $\mathcal{M}' = \mathcal{M} + \varepsilon\mathcal{E}$. The effect on dY^1 of this change is given to first order in ε by:*

$$\frac{d}{d\varepsilon}dY^1|_{\varepsilon=0} = (\boldsymbol{I} - \mathcal{M})^{-1}\mathcal{E}(\boldsymbol{I} - \mathcal{M})^{-1}\partial Y^1 \tag{A57}$$

Proof. We start from the multiplier derived in Proposition 1 and then differentiate with respect to ε:

$$\frac{d}{d\varepsilon}dY^1|_{\varepsilon=0} = \frac{d}{d\varepsilon}\left(\boldsymbol{I} - \mathcal{M}\right)^{-1}|_{\varepsilon=0}\partial Y^1 = \left(\boldsymbol{I} - \mathcal{M}\right)^{-1}\mathcal{E}\left(\boldsymbol{I} - \mathcal{M}\right)^{-1}\partial Y^1 \tag{A58}$$

This uses the standard formula from matrix calculus that $\partial\left(\boldsymbol{A}^{\mathbf{-1}}\right) = -\boldsymbol{A}^{\mathbf{-1}}(\partial\boldsymbol{A})\boldsymbol{A}^{\mathbf{-1}}$, taking $\boldsymbol{A} = \boldsymbol{I} - \mathcal{M}$, and noting that $\partial\boldsymbol{A} = -\mathcal{E}$. □

Thus, for changes in network structure given by any $\mathcal{E}$, we can compute how the multiplier from any shock changes.

Corollary 3. *Suppose that either (i) all household MPCs increase or (ii) at any firm, the share of income rationed to some zero-MPC household decreases and the share rationed to all other households increases. For any $\partial Y^1 \geqslant 0$, dY^1 increases in all dimensions.*

Proof. See in both cases that $\frac{d}{d\varepsilon}dY^1|_{\varepsilon=0} = \left(\boldsymbol{I} - \mathcal{M}\right)^{-1}\mathcal{E}\left(\boldsymbol{I} - \mathcal{M}\right)^{-1}\partial Y^1$, where $\mathcal{E} \geqslant 0$. The result follows immediately. □

While the general formulae above permit exact computation of the effects on the full vector of final output, given the potentially unrestricted network structures that we allow, it is hard to draw qualitative conclusions. For the remainder of this analysis, we report comparative statics of total final output in the empirically relevant case where there exists some reference incidence h^*

[43] Recall we have also assumed in the main text that the modulus of $\mathcal{M}$ is strictly below 1.

around which the bias and homophily effects are zero for all possible ∂h^1. In this case, the analysis of Appendix A.3 shows that the total effect of a purchases shock on output follows

$$\mathbb{1}^T dY^1 \propto \boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\partial G^1 \tag{A59}$$

It is then immediate that a rationing matrix that places higher entries on higher MPC households increases the output effect of any uniformly positive or negative shock. A more interesting result is shown in the following proposition, which establishes that adding IO linkages to an economy without them leads to a contraction in the range of output multipliers across sectors:

Proposition 7. *Consider two economies in which the bias and homophily effects are zero for all possible shocks and that are identical except that one features no input-output linkages* $\widehat{\boldsymbol{X}}^{\boldsymbol{1}} = 0$ *and the other has some arbitrary input output matrix* $\widehat{\boldsymbol{X}}^{\boldsymbol{1}}$*. The maximum purchases multiplier in the first economy is larger than the maximum purchases multiplier in the second economy and the minimum purchases multiplier is smaller in the first economy than the minimum purchases multiplier in the second economy.*

Proof. The maximum purchases multiplier in the first economy is given (up to constants irrelevant for this comparison) by the following:

$$\max_{i\in\mathcal{I}} \mathbb{1}^T\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1}e_i = \max_{i\in\mathcal{I}}(\mathbb{1}^T\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1})_i \tag{A60}$$

where e_i is the vector with one in dimension i and zeros elsewhere. In the second economy, see that the maximum purchases multiplier is given by:

$$\max_{i\in\mathcal{I}} \mathbb{1}^T\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}e_i = \max_{i\in\mathcal{I}}\left(\mathbb{1}^T\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1}\right)\widehat{\boldsymbol{L}}^{\boldsymbol{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}_{\cdot,i} \tag{A61}$$

As production is CRS, recall by Corollary 2 that the columns of $\widehat{\boldsymbol{L}}^{\boldsymbol{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}$ sum to one. Thus, the maximum multiplier is some weighted average of the elements of $\mathbb{1}^T\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1}$. The value of which is necessarily bounded above by the maximum element of $\mathbb{1}^T\boldsymbol{m}\boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^1}$, which is the maximum purchases multiplier in the first economy. The proof that the minimum purchases multiplier is smaller in the first economy than the minimum purchases multiplier in the second economy follows the same argument, where we instead note that the weighted average is necessarily bounded below by the minimum. □

C.3. Special Cases Where Incidence, Bias, and Homophily Effects Vanish

In the main text, we briefly discussed a simple case where all of the incidence, bias, and homophily in shock propagation vanish. Here, we state and more formally discuss this case.

Proposition 8. *Suppose that all industries have a common rationing-weighted average MPC,* $m = \sum_{n \in N} (\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}})_{ni} m_n$. *The incidence, bias, and homophily effects are zero, and the change in GDP corresponding to any change in government purchases is:*

$$dGDP^1 = \frac{1}{1-m} \mathbb{1}^T dG^1 \tag{A62}$$

Proof. To begin, recall from Proposition 1 and the definition $\widehat{\boldsymbol{C}}^{\mathbf{1}} \boldsymbol{m} = \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^{\mathbf{1}}}$ that for any shock to first-period government purchases we have

$$\begin{aligned} \mathbb{1}^T dY^1 &= \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^{\mathbf{1}}} \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \hat{\boldsymbol{L}}^{\mathbf{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1} \right)^{-1} dG^1 = \sum_{k=0}^{\infty} \mathbb{1}^T \left(\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^{\mathbf{1}}} \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \hat{\boldsymbol{L}}^{\mathbf{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1} \right)^k dG^1 \\ &= \sum_{k=0}^{\infty} m^k \mathbb{1}^T dG^1 = \frac{1}{1-m} \mathbb{1}^T dG^1 \end{aligned} \tag{A63}$$

where the last line follows from the fact that $\mathbb{1}^T \left(\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^{\mathbf{1}}} \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \hat{\boldsymbol{L}}^{\mathbf{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1} \right) = m \mathbb{1}^T$. It remains to show that this fact follows from the conditions provided in the statement of the Proposition. Namely, we must show that

$$\mathbb{1}^T \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^{\mathbf{1}}} \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} = m \mathbb{1}^T \quad \implies \quad \mathbb{1}^T \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{h}^{\mathbf{1}}} \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \hat{\boldsymbol{L}}^{\mathbf{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1} = m \mathbb{1}^T \tag{A64}$$

This is immediate from the transpose of the no-profit condition

$$p^1 = (\boldsymbol{I} - (\widehat{\boldsymbol{X}}^{\mathbf{1}})^T)^{-1} \hat{\boldsymbol{L}}^{\mathbf{1}} \mathbb{1}, \tag{A65}$$

and our normalization $p = \mathbb{1}$. □

The proposition imposes that each firm's marginal employees have the same average MPC as one another. This eliminates the incidence, bias, and homophily effects, leaving only the classical Keynesian multiplier. That is, wherever in the economy a purchases shock is directed, and however it spreads through directed consumption and the IO network, the change in aggregate consumption generated by the reduction in firm revenue is the same. Of course, a particular special case that satisfies these conditions is when there is a single good and a single household (in which case $\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} = 1$). Note that even when the traditional Keynesian multiplier obtains, the aggregate MPC need not equal either the average MPC or the income-weighted average MPC of the population; this is the case only when each firm's marginal employees have that average MPC.

D. Model Extensions and Results

In this appendix, we extend the baseline model to allow for imperfect competition with fixed markups (D.1) and many periods (D.2).

D.1. Imperfect Competition

In this appendix, we show how to incorporate imperfect competition in the form of fixed markups on marginal costs. Now, instead of each sector being populated by a continuum of perfectly competitive firms, we suppose that for all $i \in I$ there is a single monopolist producing each good, charging a fixed markup of $\frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t}$ over their marginal cost and making (and distributing) profits $\Pi_i^t = \frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t} p_i^t Q_i^t$.[44] Despite this, we argue that a non-substitution theorem still holds and we can obtain analogous multiplier formulae once we augment labor income rationing with profit rationing. To do this, we have to slightly modify Assumption 4:

Assumption 7. *For each $t \in \{1,2\}$ there exists some $\overline{p}^t \in \mathbb{R}_+^I$ and $\{X_i^t, L_i^t\}_{i \in I}$ such that for all $i \in I$: $F_i^t(X_i^t, L_i^t) \geqslant 1$ and $(1 + \frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t})(\overline{p}^t X_i^t + L_i^t) \leqslant \overline{p}_i^t$*

Under this modified assumption, we can state and prove the modified non-substitution theorem with markups:

Proposition 9. *Fix a profile of government spending and transfers $\mathcal{P} = \{\tau_n^t, G_i^t\}_{t \in \{0,1\}, n \in N, i \in I}$ that satisfies (4) and suppose that Assumptions 4, 5, and 7 hold. There exists a unique $p^t(\mathcal{P}) \in \mathbb{R}_{\geqslant 0}^I$ consistent with rationing equilibrium and, moreover, $p_t(\mathcal{P}) \gg 0$. Furthermore, $p^t(\mathcal{P})$ is invariant to $\mathcal{P}$.*

Proof. We modify the proof of proposition 4. Each firm i in period t now sets a price $p_i^t = (1 + \frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t})\kappa_i(p^t) = \kappa_i(p^t)/(1 - \widehat{\Pi}_i^t)$, where κ_i^t is i's unit cost function in period t. That is, i prices goods as though it were a competitive firm with production function $(1 - \widehat{\Pi}^t) F_i^t(X_i^t, L_i^t)$. Consider now a modified economy without markups and production functions given by the previously-stated markup-adjusted production functions. Assumption 7 implies that Assumption 4 holds in this modified economy. The result then follows by application of Proposition 4. □

Since, by Proposition 9, prices are invariant to policy, we may normalize them to one without loss of generality. This implies that firms earn profits per unit sold of $\frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t} / \left(1 + \frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t}\right) = \widehat{\Pi}_i^t$. Each firm i and time t therefore pays $\widehat{\Pi}_i^t Q_i^t$ to its shareholders, in close analogy to its payments $\widehat{L}_i^t Q_i^t$ to employees and its expenditures $\widehat{X}_{ij}^t Q_i^t$ on inputs from each other firm j.

We assume that profits from each firm are distributed to households in each period t according to an exogenous dividend (a.k.a. profit rationing) function $D^t : \mathbb{R}^I \to \mathbb{R}^N$ satisfying $\sum_{i \in \mathcal{I}} \Pi_i^t = \sum_{n \in N} D^t(\Pi^t)_n$ for all $\Pi^t \in \mathbb{R}^{\mathcal{I}}$. With profits, household income is comprised of rationed first-period labor income, chosen second-period labor income, and (not chosen) dividend income in both periods. For simplicity, we continue to assume that (by the additive separability of labor supply

[44] One microfoundation for constant markups is that industries are comprised of a continuum of firms, with each other firm's and household's demands having the same CES aggregator for these firms' varieties.

and consumption preferences) each household n's consumption in each period depends on its labor and profit incomes only through its total income $\{h_n^t\}_{t=1,2}$.

We can now state a profit-inclusive Keynesian cross. The main difference to Proposition 1 comes from the need to account for changes in profits and how these are distributed to households as dividends. Additionally, the presence of profits links the first period to the second—because changes in future profits affect lifetime incomes, which affect consumption today.

Proposition 10. *For any shock inducing a partial equilibrium effect $\partial Y = \boldsymbol{C_\tau} d\tau + dG$, the general equilibrium response in production satisfies:*

$$dQ = \widehat{\boldsymbol{X}} dQ + \boldsymbol{C_{h^1}} \boldsymbol{R^1_{L^1}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} dQ^1 + \boldsymbol{C_h} \boldsymbol{D_\Pi} \widehat{\boldsymbol{\Pi}} dQ + \partial Y \tag{A66}$$

where $\boldsymbol{C_h}$ is the matrix of household directed MPCs, where $\boldsymbol{D_\Pi}$ is the block diagonal matrix composed of $\boldsymbol{D^1_{\Pi^1}}$ and $\boldsymbol{D^2_{\Pi^2}}$—which are each $N \times I$ matrices with entries $D^t_{\Pi^t_i}(\Pi^t)_n$—and where $\widehat{\boldsymbol{\Pi}}$ is the block diagonal matrix composed of $\widehat{\boldsymbol{\Pi}}^{\boldsymbol{1}}$ and $\widehat{\boldsymbol{\Pi}}^{\boldsymbol{2}}$—themselves each diagonal matrices with entries $\widehat{\Pi}^t_i$. All quantities are evaluated at the initial equilibrium.

Proof. Stacking the vectors that represent periods 1 and 2, we perturb the goods market equilibrium conditions:

$$dQ \ = \widehat{\boldsymbol{X}} dQ + \boldsymbol{C_{h^1}} \boldsymbol{\mu} d\ell^1 + dG + \boldsymbol{C_h} \boldsymbol{D_\Pi} \widehat{\boldsymbol{\Pi}} dQ \tag{A67}$$

Plugging in for $\boldsymbol{\mu} d\ell^1 = \boldsymbol{R^1_{L^1}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} dQ^1$, we obtain the result. □

In the special case where there are no profits in the second period, i.e., $\widehat{\Pi}^2_i = 0$ for all i, restricting (A66) to the first period and using $dY_1 = (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}) dQ_1$ implies

$$dY^1 = \left[\boldsymbol{I} - \boldsymbol{C^1_{h^1}} (\boldsymbol{R^1_{L^1}} \hat{\boldsymbol{L}}^{\boldsymbol{1}} + \boldsymbol{D^1_\Pi} \widehat{\boldsymbol{\Pi}}^{\boldsymbol{1}})(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}\right]^{-1} \partial Y_1. \tag{A68}$$

This is the same expression as our baseline multiplier, except for that both labor income and profit income is rationed to households. This expression microfounds our quantitative calibration with firm profits in Section 4.

D.2. *Multiple Time Periods*

Consider the benchmark model from Section 2, but with many periods. In particular, suppose that real interest rates are constrained not just for a single period but for all T time periods, so that we may take them as exogenously fixed at rates r^t. Labor is rationed in periods $1, ..., T-1$, whereas in period T, labor is supplied competitively. Household consumption c^t_n and final period labor supply ℓ_n^T is chosen as a function of preferences, taxes, and total income in each period in a way that satisfies the dynamic budget constraint:

$$\sum_{t\leqslant T}\frac{\ell_n^t}{\prod_{t'\leqslant t}1+r^{t'}}=\sum_{t\leqslant T}\frac{p^tc_n^t+\tau_n^t}{\prod_{t'\leqslant t}1+r^{t'}} \tag{A69}$$

Similarly, lump-sum taxes and spending $\{\{\tau_n^t\}_{n\in N},\{G_i^t\}_{i\in I}\}_{t\leqslant T}$ satisfy a lifetime budget constraint for the government:

$$\sum_{n\in N}\mu_n\Big(\sum_{t\in T}\frac{1}{\prod_{t'\leqslant t}1+r^{t'}}\tau_n^t\Big)=\sum_{t\in T}\frac{1}{\prod_{t'\leqslant t}1+r^{t'}}p^tG^t \tag{A70}$$

At all periods $t\leqslant T-1$—those in which labor is rationed—and for all $n\in N$, we have $\mu_n\ell_n^t=R_n^t(L^t)$ for some period-specific rationing function R^t satisfying $\sum_{n\in N}\mu_nR_n^t(L^t)=\sum_{n\in N}L_n^t$.

Definition 3. *(Dynamic rationing equilibrium) Given a profile of government spending and transfers* $\{\tau_n^t,G_i^t\}_{t\in\{0,\ldots,T\},n\in N,i\in I}$ *that satisfies (A70), a dynamic rationing equilibrium is a profile* $\{p_i^t,\ell_n^t,c_n^t,L_i^t,X_{ij}^t\}_{t\in\{0,\ldots,T\},n\in N,i,j\in I}$ *that satisfies firm optimization in every period (as in* (1)*), consumption and final period labor supply consistent with household optimization, labor rationing in periods* $t\leqslant T-1$*, and goods and labor market clearing.*

Under this dynamic equilibrium concept, we obtain an analogous fiscal multiplier to that of our two period model. In order to present this result, let us introduce a bit of notation: Since prices are exogenous, normalize the units of all goods so that $p_i^t=1$ for all $t\leqslant T$, $i\in I$. Below, for any T-length vector v, we will use the notation v^{-T} to denote the $(T-1)$-length $(v_1,\ldots,v_{T-1})$. In a similar fashion, we let $\widehat{\boldsymbol{L}}^{-\boldsymbol{T}}$ and $\widehat{\boldsymbol{X}}^{-\boldsymbol{T}}$ denote the block-diagonal $\big((T-1)\times I\big)\times\big((T-1)\times I\big)$ matrices comprised of the within-period-t unit labor demand and IO matrices on the t^{th} block. $\boldsymbol{R}^{-\boldsymbol{T}}_{\boldsymbol{L}^{-\boldsymbol{T}}}$ is the analogous block-diagonal $\big((T-1)\times N\big)\times\big((T-1)\times I\big)$ matrix corresponding to marginal labor rationing stacked across $T-1$ periods. Somewhat more interestingly, let $\boldsymbol{m}^{-\boldsymbol{T}}$ denote the $\big((T-1)\times N\big)\times\big((T-1)\times N\big)$ matrix whose (t,n,t',n') entry is zero if $n\neq n'$ and—if $n=n'$—is equal to n's MPC into period t consumption spending out of period t' income. Finally, let $\widehat{\boldsymbol{C}}^{-\boldsymbol{T}}$ be the $\big((T-1)\times N\big)\times\big((T-1)\times I\big)$ matrix that captures the direction of this marginal spending, i.e., the matrix with (t,i,t',n) entry equal to the fraction of n's period-t'-income-induced spending into period t that is directed toward good i.

Proposition 11. *Given any rationing equilibrium, the local change in equilibrium final output* dY^{-T} *in periods* $1,\ldots,T-1$ *following a fiscal shock with direct effect* ∂Y^{-T} *on final output in periods* $1,\ldots,T-1$ *is given by:*

$$dY^{-T}=\Big(\boldsymbol{I}\ -\ \widehat{\boldsymbol{C}}^{-\boldsymbol{T}}\ \boldsymbol{m}^{-\boldsymbol{T}}\ \boldsymbol{R}^{-\boldsymbol{T}}_{\boldsymbol{L}^{-\boldsymbol{T}}}\ \widehat{\boldsymbol{L}}^{-\boldsymbol{T}}\ \Big(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{-\boldsymbol{T}}\Big)^{-1}\Big)^{-1}\partial Y^{-T} \tag{A71}$$

Proof. The proof is directly analogous to that of Proposition 1. □

The presence of many periods with rationing introduces one key qualitative difference: Shocks can spill over across periods with labor rationing. As a result, it is no longer sufficient to consider the directed MPC of households; one must work with the directed *intertemporal* MPC of households that represents marginal changes in consumption across goods and time. These intertemporal MPCs are precisely the object of study of Auclert et al. (2018). Indeed, (A71) coincides with their intertemporal Keynesian cross when there is a single good.

E. Additional Tables and Figures

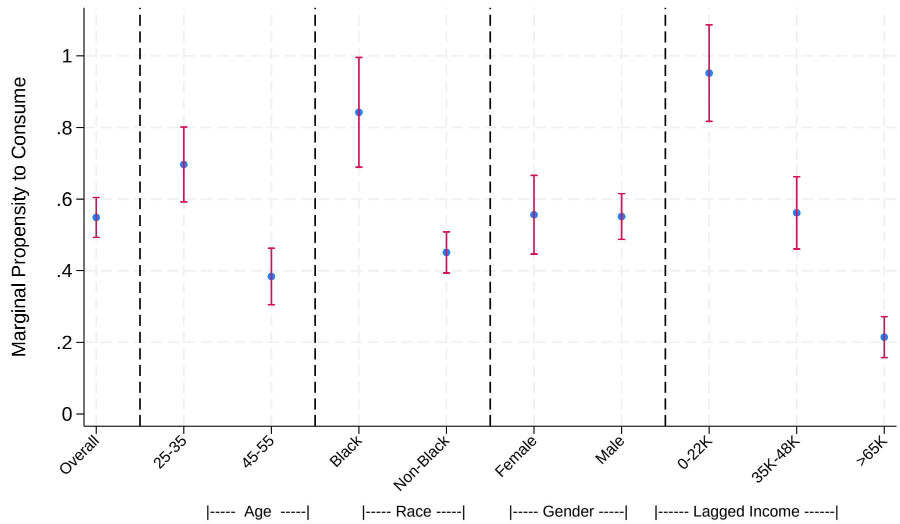

Fig. A1. Heterogeneity in estimated MPCs for total consumption across demographic groups.

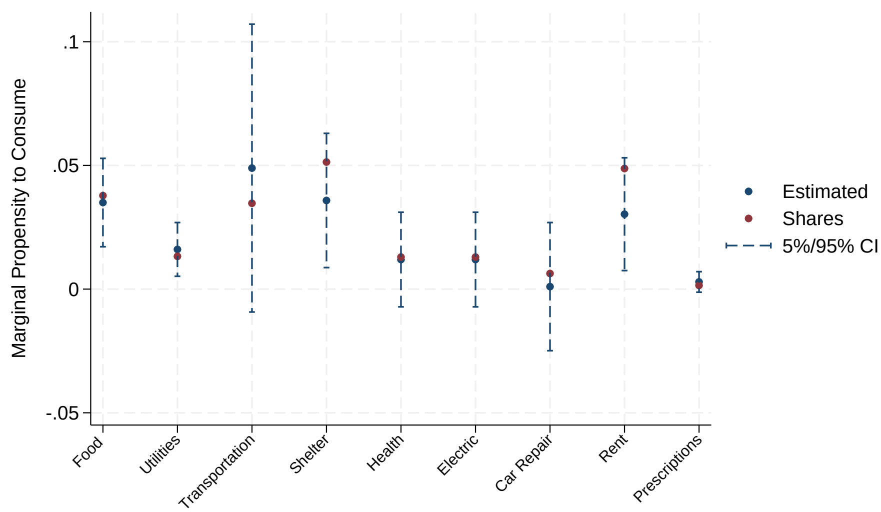

Fig. A2. Estimated Directed MPCs Vs. CEX basket-weighted MPCs

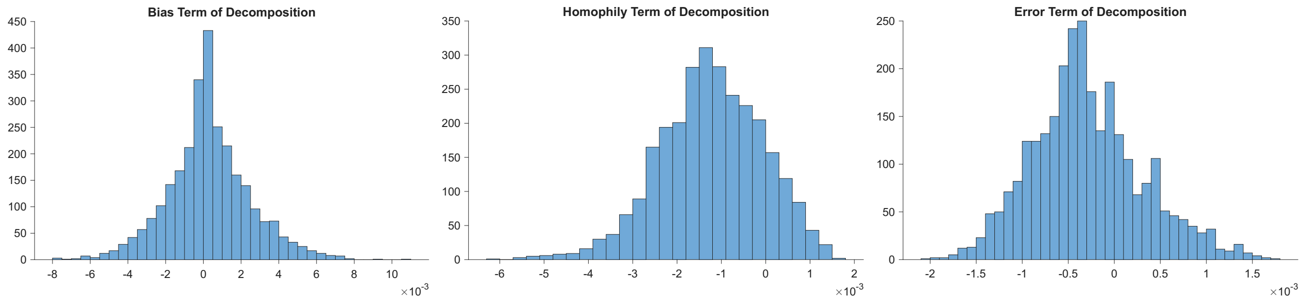

Fig. A3. Histograms of the bias term (left) and homophily term (middle) and overall error terms (right) from the decomposition in Proposition 2. For all subfigures, the distribution reflects a unit demand shock to each of the 2805 sector-region pairs, with baseline h^* given by the income incidence of a shock to demand proportional to 2012 state-industry GDP.

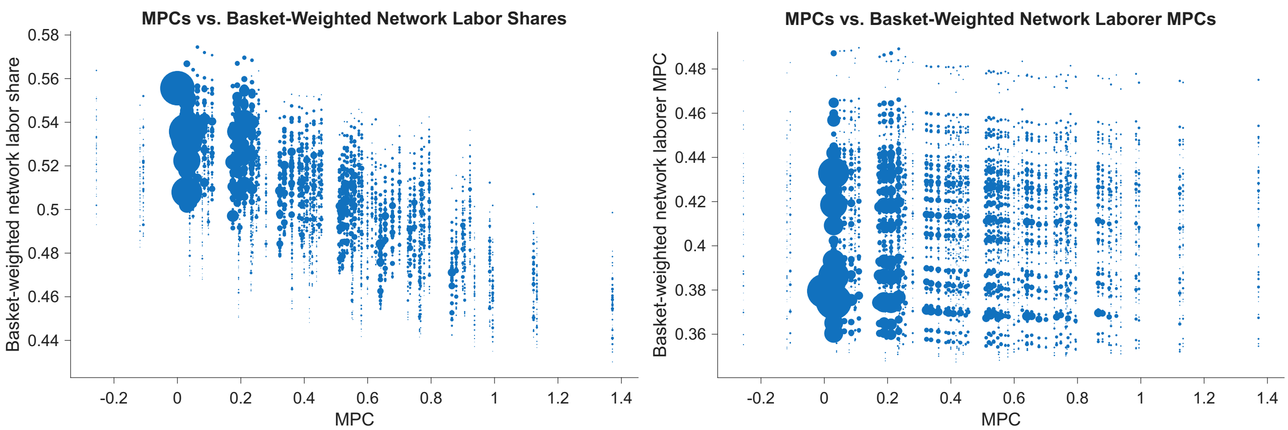

Fig. A4. The left panel shows the scatter plot of worker MPCs against the basket-weighted labor share of the sectors on which they consume. The right panel shows a scatter plot of worker MPCs against the basket-weighted MPCs of the labor employed in the sectors producing the goods they ultimately consume. In both plots, each dot corresponds to a region × demographic group level and its size corresponds to the incidence of a GDP-proportional shock onto the region-demographic's income.

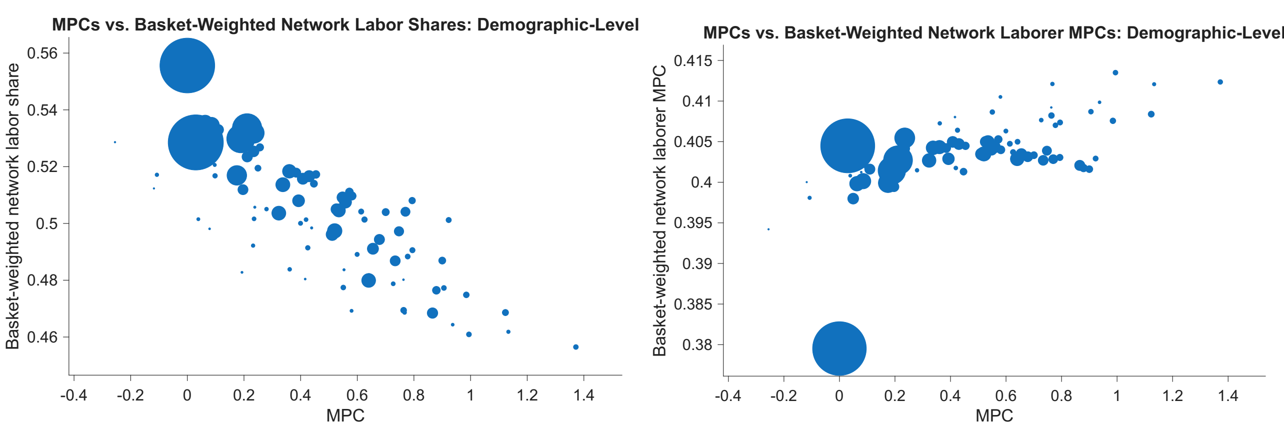

Fig. A5. The left panel shows the scatter plot of worker MPCs against the basket-weighted labor share of the sectors on which they consume. The right panel shows a scatter plot of worker MPCs against the basket-weighted MPCs of the labor employed in the sectors producing the goods they ultimately consume. In both plots, each dot corresponds to a demographic group level—we add group members across regions, weighting by the incidence of a GDP-proportional shock—and its size corresponds to the incidence of a GDP-proportional shock onto the demographic's income.

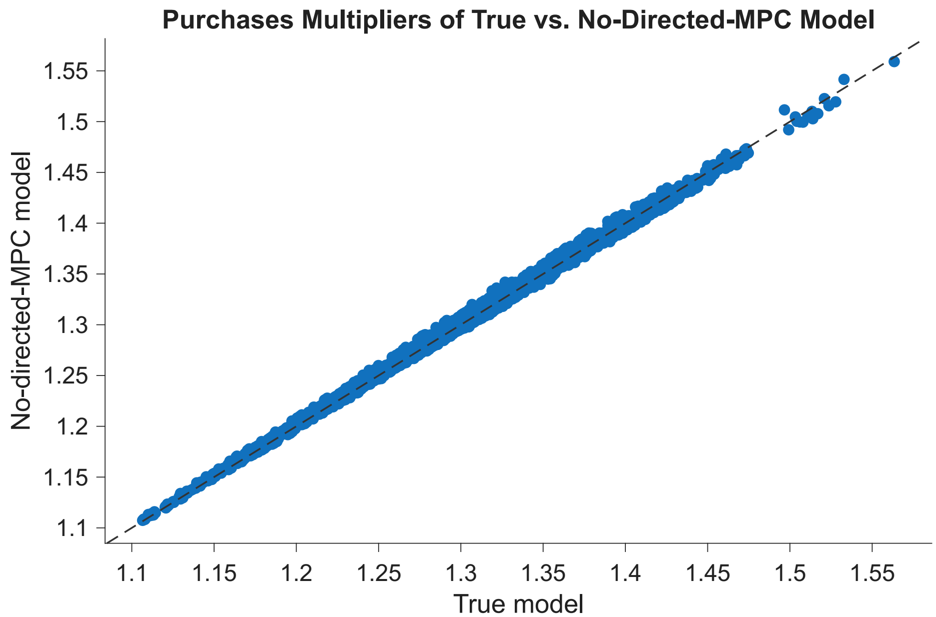

Fig. A6. Scatter plot of purchases multipliers for each of the 2805 industry-region pairs in the baseline model (x-axis) and the model in which all households have homogeneous consumption baskets in proportion to GDP (y-axis).

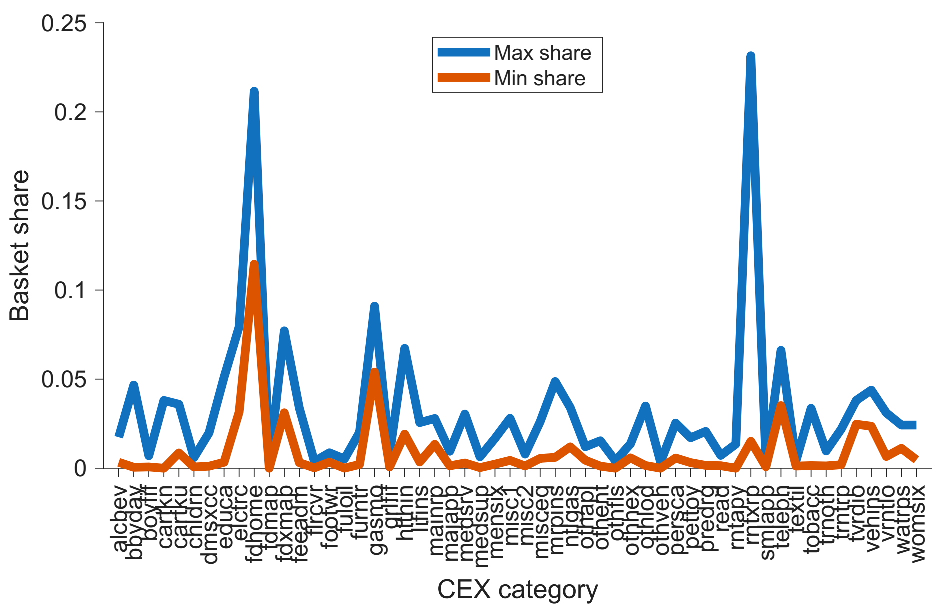

Fig. A7. Maximum and minimum consumption basket weights across all demographic groups in each CEX consumption category.

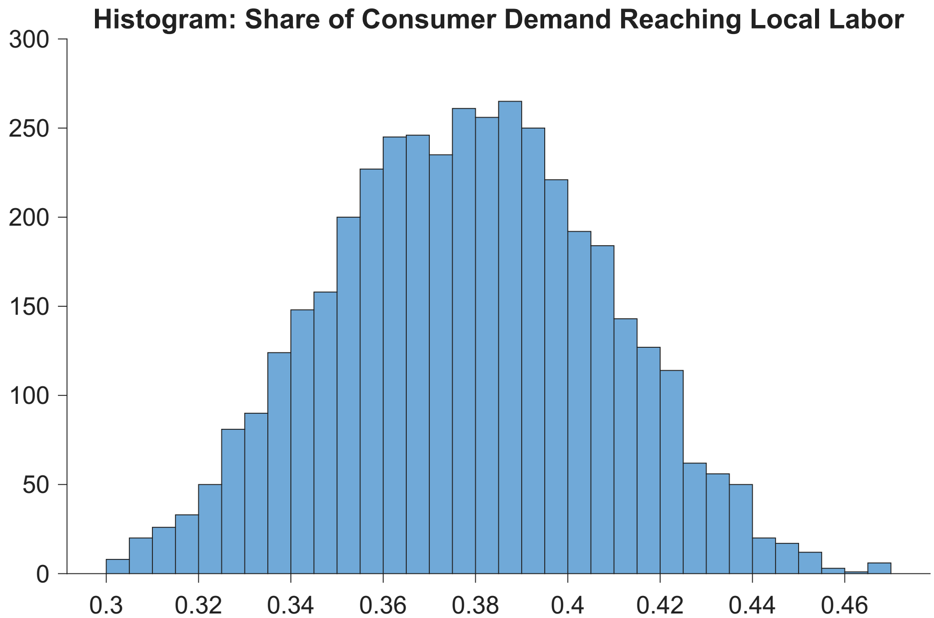

Fig. A8. Histogram of the fraction of consumer demand resulting in income for labor within the same state for each state-demographic pair, excluding "foreigners" demographic group.

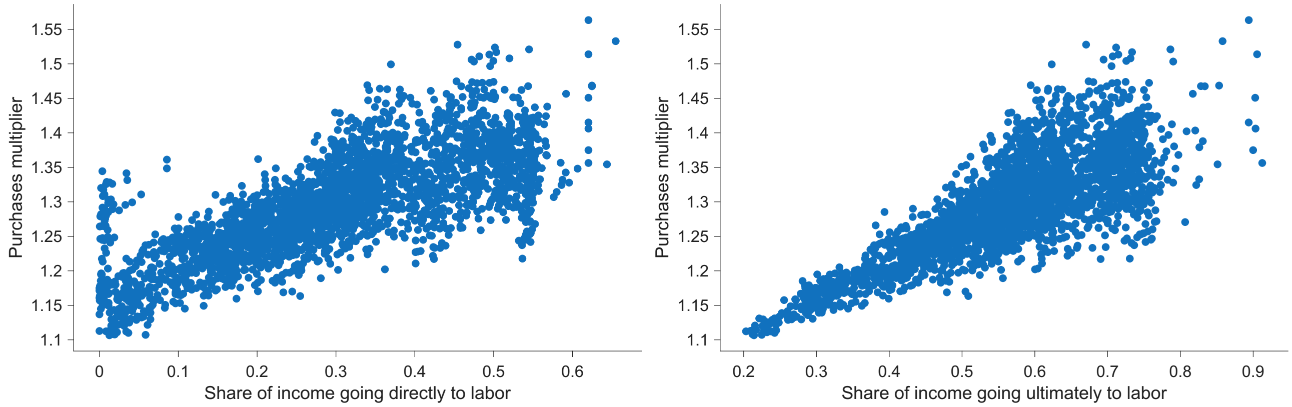

Fig. A9. Left Panel: Scatter plot comparing, for a $1 purchases shock to each industry-region, the change in aggregate consumption (y-axis) and the share of income from production that goes directly to labor (as opposed to capital, foreigners, or inputs) (x-axis). Right Panel: Scatter plot comparing, for a $1 purchases shock to each industry-region, the change in aggregate consumption (y-axis) and the ultimate labor share accounting for labor employed in the production of intermediates (x-axis).

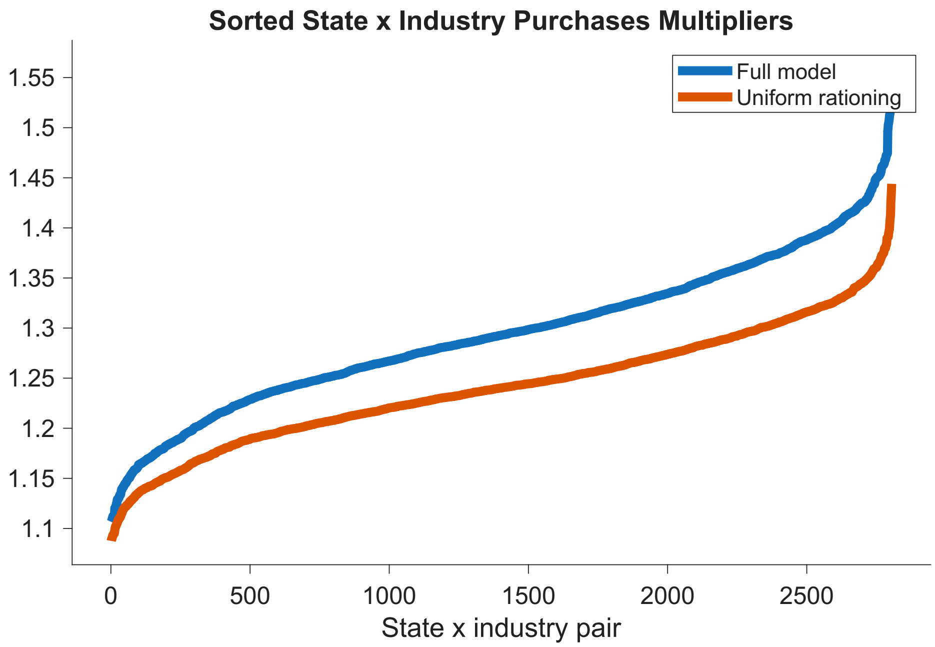

Fig. A10. Sorted change in aggregate consumption resulting from a one dollar demand shock to each industry-region pair. Full model is the baseline. Uniform rationing corresponds to all households' labor income being scaled in proportion to their labor income.

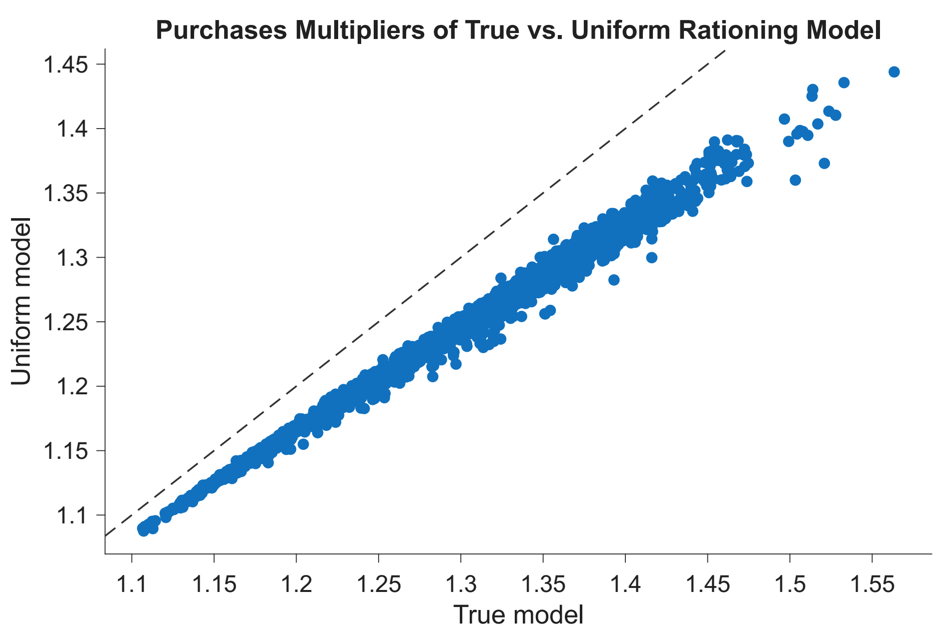

Fig. A11. Scatter plot of the change in aggregate consumption resulting from a one dollar demand shock to each industry-region pair. Full model is the baseline. Uniform rationing corresponds to all households' labor income being scaled in proportion to their income.

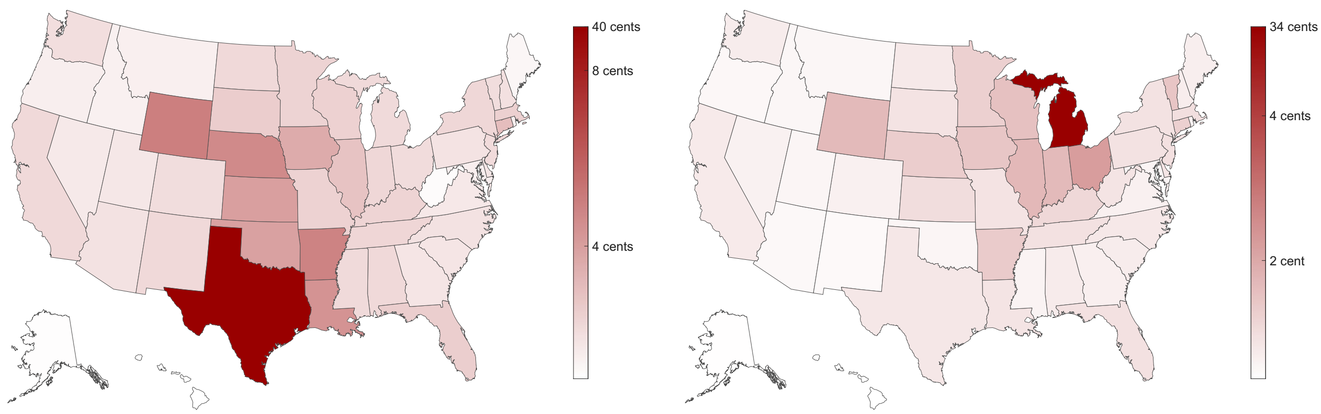

Fig. A12. Changes in state-level value added per worker following a uniform $1-per-capita transfer shock to households in Texas (left panel) and Michigan (right panel).

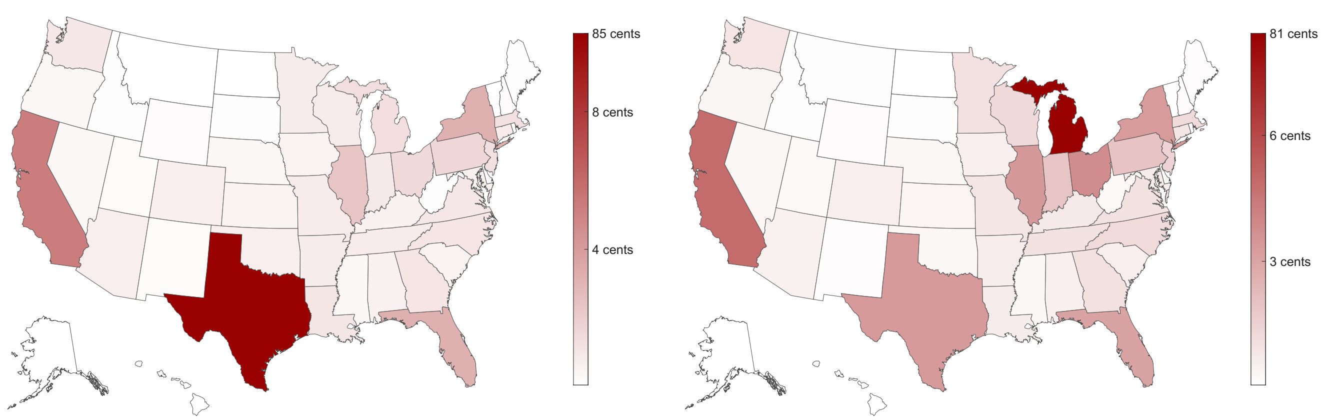

Fig. A13. Changes in state-level value added following a GDP-proportional $1 purchases shock to Texas (left panel) and Michigan (right panel).

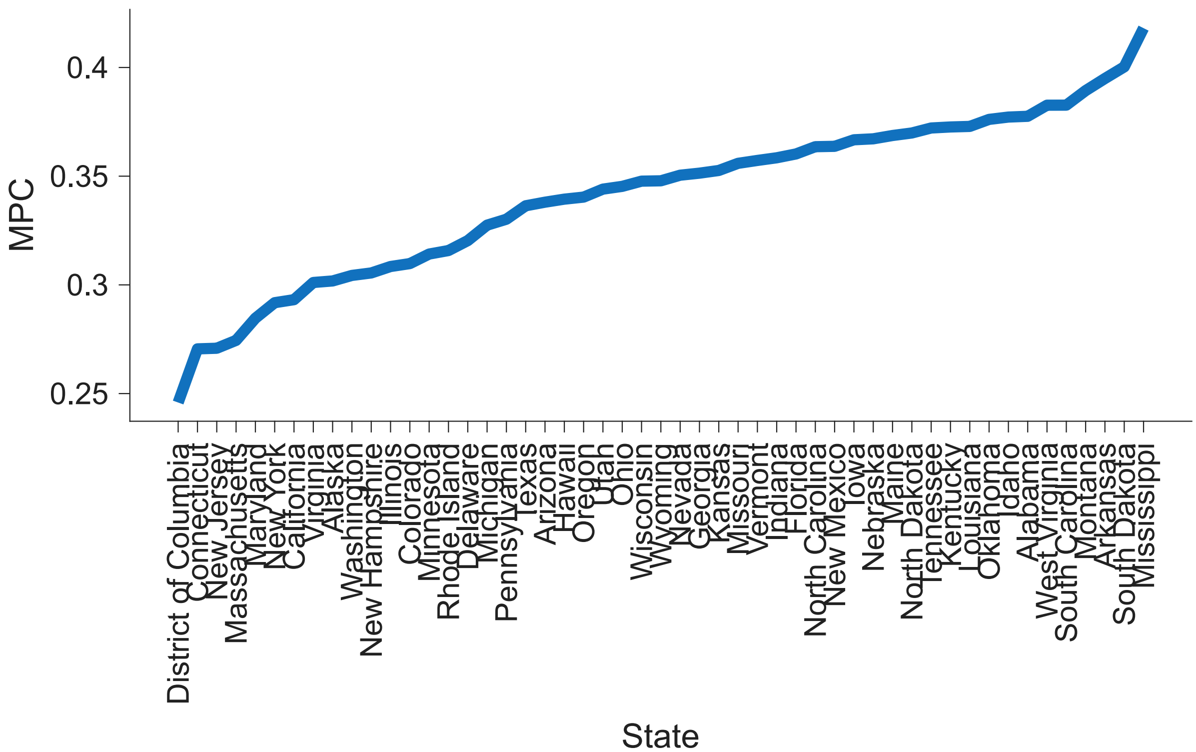

Fig. A14. Income-weighted average worker MPC by state.

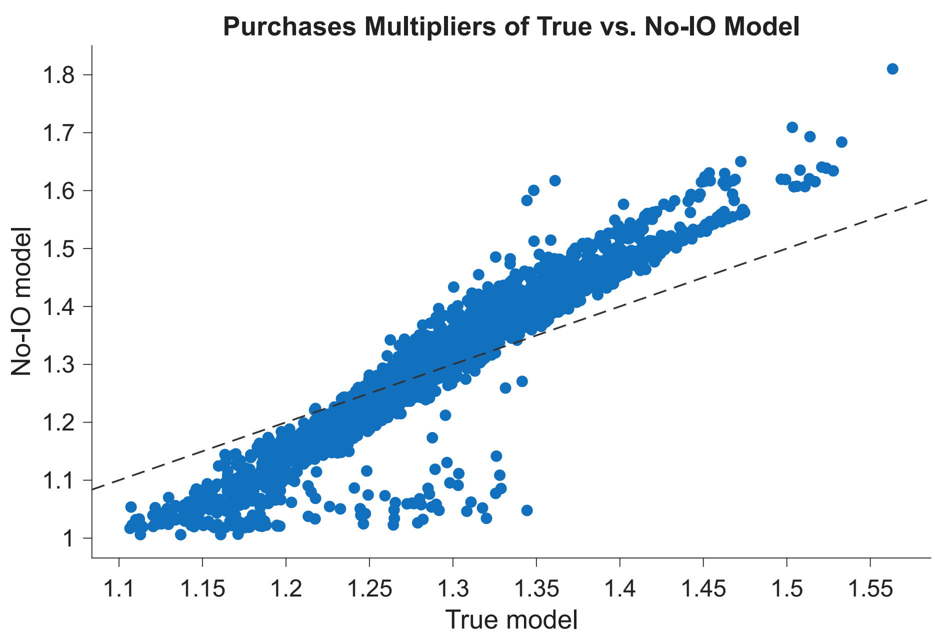

Fig. A15. Scatter plot of the change in aggregate consumption resulting from a one dollar demand shock to each industry-region pair. Full model is the baseline. No-IO corresponds to the case where firms do not purchase intermediate inputs and instead scale up their factor payments proportionally.

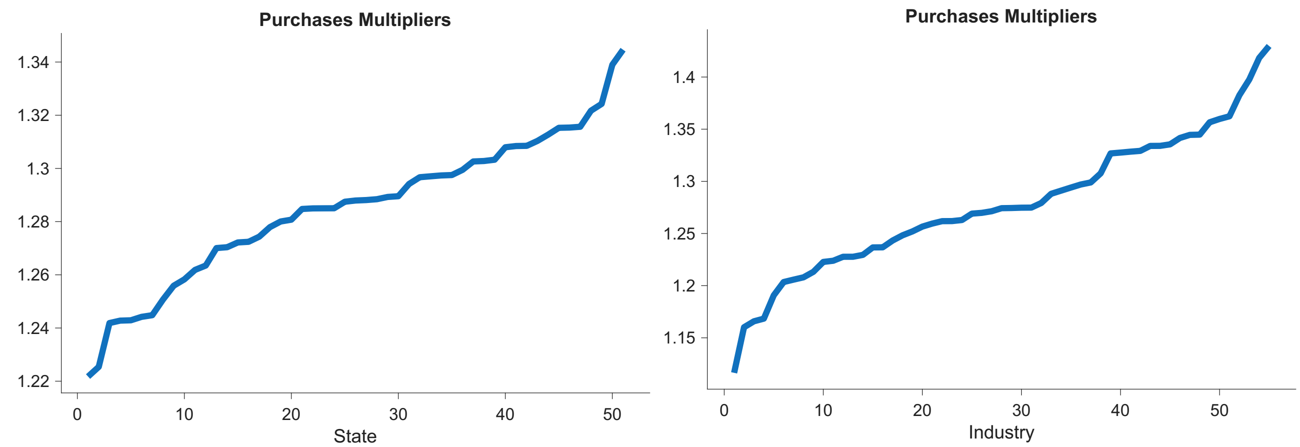

Fig. A16. Multipliers for state-level and industry level purchases shocks. Formally, for the state-level shock, we purchase from each state one dollar of output, in proportion to the sectoral composition of that state. For the industry-level shock, we purchase from each demographic group one dollar of output, in proportion to the state composition of that sector.

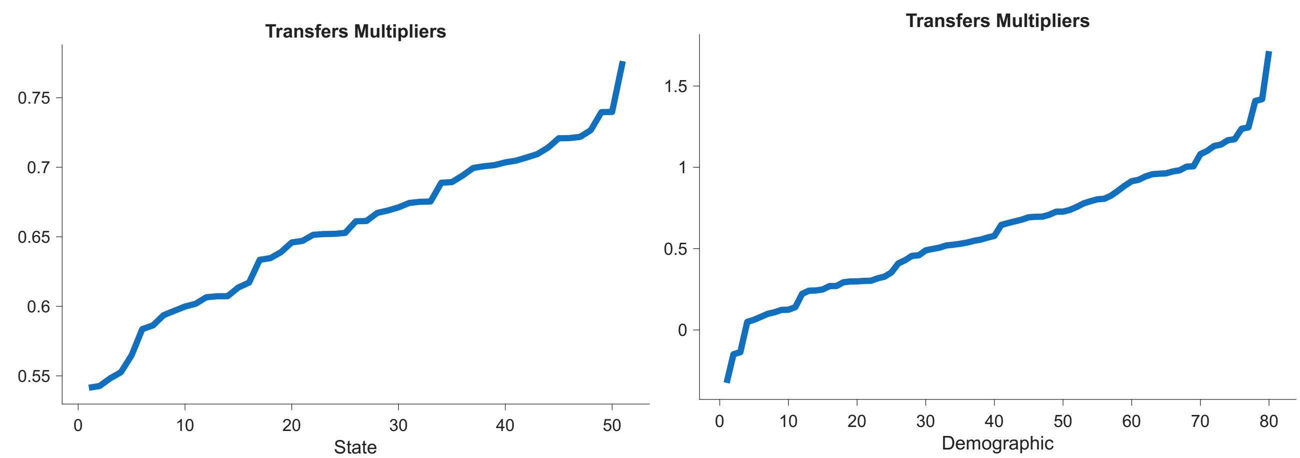

Fig. A17. Multipliers for state-level and demographic level transfers shocks. Formally, for the state-level shock, we transfer each state one dollar, in proportion to the demographic composition of that state. For the demographic-level shock, we transfer each demographic group one dollar, in proportion to the distribution of that demographic across states.