Fiscal Policy in a Networked Economy

John Sturm Becko* Princeton
Joel P. Flynn† Yale
Christina Patterson‡ University of Chicago and NBER

December 13, 2023

Abstract

Advanced economies feature complicated networks that connect households, firms, and regions. How do these structures affect the impact of fiscal policy and its optimal targeting? We study these questions in a model with input-output linkages, regional structure, and household heterogeneity in MPCs, consumption baskets, and shock exposures. Theoretically, we derive estimable formulae for the effects of fiscal policies on aggregate GDP, or fiscal multipliers, and show how network structures determine their size. Empirically, we find that multipliers vary substantially across policies, so targeting is important. Beneath these aggregate effects are large spatial and sectoral spillovers from policies directed to any one firm or household. However, virtually all variation in multipliers stems from differences in policies' direct incidence onto households' MPCs. Thus, while the distributional effects of fiscal policy depend on the detailed structure of the economy, 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, Ricardo Caballero, Arnaud Costinot, Dave Donaldson, 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 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

Economic shocks present policymakers with the challenge of designing stimulus programs to alleviate economic distress. Most recently, the United States Congress has adopted three common forms of stimulus: undirected transfers (stimulus checks), targeted transfers (expanded unemployment insurance benefits), and targeted spending (industry programs, such as for the airlines). This wide range of alternative policies draws attention to questions that policymakers face during every major recession. Which forms of fiscal stimulus are the most effective, whom do they help, and how should they be targeted?

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 determining fiscal multipliers at the macroeconomic level. 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. While these channels interact in complex ways, we show how to decompose all of their interactions into three distinct effects, on top of a baseline representative firm and agent Keynesian multiplier.

The second part of the paper takes this decomposition to the data and documents three key empirical findings. First, fiscal multipliers vary widely across different government purchase and transfer policies, so targeting is important. Second, there are large 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 heterogeneity in fiscal multipliers can be explained by the heterogeneous incidence of fiscal shocks onto households with different MPCs, whose particular patterns of consumption are, by contrast, irrelevant. 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, general equilibrium 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. Within this rich setting, we study a rationing equilibrium where wages are sticky and demand is deficient. Labor is therefore rationed, meaning that households can lie off their labor supply curves and be involuntarily un(der)employed. This assumption, as well as a focus on the case where an effective lower bound binds, makes our model applicable to severe recessions.

The various microeconomic interconnections between households complicate the transmission of fiscal shocks into output. Our first result tracks these linkages in a generalized Keynesian income multiplier that expresses the change in the vector of first period output across industries and regions in response to a partial equilibrium shock to final demand. First, the input-output network translates shocks to final demand into changes in the production of each sector. Second, a rationing function captures how those 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.

Despite this complexity, we show that the total effect of any fiscal policy on aggregate GDP – 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 GDP 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, 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 at 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. 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* output, all dimensions of heterogeneity shape the *distribution* of this output across households, industry, and space. 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 the targeting of fiscal policy. First, consider a planner concerned only with aggregate underemployment. 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 transfer policy in the model: Our estimates suggest that government transfers of one thousand dollars to each employed worker would increase GDP by 69 cents per dollar spent, whereas transfers of two thousand dollars to each worker with above-median MPC would increase GDP by 96 cents per dollar spent. Second, consider a planner addressing a recession where underemployment is concentrated on certain demographic groups and regions. While this heterogeneity com-

[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 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.

plicates policy design, our detailed estimates of demand spillovers still allow us to compare various forms of targeting. Using estimates from the Great Recession, we show that targeting government purchases according to the combination of MPCs and labor wedges achieves the bulk of the available gains.

Finally, we use the structure of our exercise to provide theoretical results and counterfactual simulations that explore how the distribution of multipliers may vary across economies and across time within the US. We first construct a counterfactual economy with no input-output linkages, finding that while this does not affect the amplification of a GDP-proportional shock, rich IO connections do tighten the distribution of multipliers across industry-regions. Second, we consider the general decrease in sectoral labor shares between 2000 and 2012, finding that this decreases the multiplier of shocks to most, but not all, industry-region pairs. Third, we show that a "hollowing-out" of the labor income distribution within occupations has negligible effects on fiscal multipliers, but that rising income inequality could affect the distribution of multipliers if it were to change average MPC levels or shift workers across states and industries.

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 solely in understanding the response of *aggregate* GDP to fiscal policy. Our results suggest that accounting for heterogeneous incidence is the only important margin to consider. Thus, our results echo recent work that stresses amplification of shocks that load more heaving onto households with higher-than-average MPCs (Werning, 2015; Kaplan, Moll, and Violante, 2018; Auclert, 2019; Patterson, 2019; 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).

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 in a quantitatively meaningful way to the aggregate GDP impact of shocks. 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, while we study heterogeneity within a single period. This emphasis on cross-sectional het-

erogeneity 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) develop rich macroeconomic models featuring these channels as well as endogenous prices and markups. By abstracting away from price movements, 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 1.30, 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, fiscal multipliers range from 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 approximately 44% of the total change in GDP 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.

Outline The rest of the paper proceeds as follows. Section 2 introduces the model and defines the rationing equilibrium. Section 3 derives and analyzes the generalized multiplier. Section 4 introduces the data and methodology we use to estimate the generalized multiplier. Section 5 studies the estimated distribution of fiscal multipliers. Section 6 explores the implications of these findings for the design of fiscal policy. Section 7 explores the stability of these conclusions across economies and over time. Section 8 concludes.

[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.

2. The Model and Rationing Equilibrium

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 very flexible labor rationing. Our model is rich enough to capture many dimensions of household, industrial, and regional heterogeneity, but sufficiently tractable to deliver equations that we later bring directly to the data.

In Appendices D.1 and D.3, 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.

2.1. 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 $t \in \{1, 2\}$.

First period nominal wages w^1 are fixed due to binding downward wage rigidity and nominal interest rates ι are fixed due to a binding lower bound. In addition, we assume nominal wage inflation is fixed.[4] As a result, 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 function," 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

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

[4] It suffices to assume that—while one remains in a recessionary environment with deficient demand—fiscal policy does not affect households beliefs about inflation.

supply – which is determined completely by firms' labor demands – as given. We offer a detailed microfoundation of this rationing process in Appendix B.

As first and second period wages are fixed, we normalize each to one, i.e. $w^t \equiv 1$.[5] We denote the real interest rate by $r = \frac{w^1}{w^2}\iota$ and denote by $p^t = \{p_i^t\}_{i \in I}$ the vector of prices in period t. Note that r is exogenously fixed because wages and the nominal interest rate are.

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 CRS production function F_i^t.

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

Each household n's labor supply in the first period is then determined as a function of the vector of firm-specific labor demands $L^1 = \{L_i^1\}_{i \in I}$. This rationing occurs without regard for household labor supply preferences.

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

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. $\sum_{n \in N} R_n^1(L^1) = \sum_{i \in I} L_i^1$.

Taking not only prices but also first-period labor supply as given, each household n chooses a second-period labor supply ℓ_n^2 and vectors of first- and second-period consumption $c_n^t = \{c_{ni}^t\}_{i \in I}$. For most of the paper, we remain agnostic to the nature of household behavior – simply specifying Marshallian demands and supplies – except for that we impose that those demands satisfy a life-time budget constraint that incorporates lump-sum taxes $\tau_n = (\tau_n^1, \tau_n^2)$ and that households consume out of net first period income $h_n^1 = \ell_n^1 - \tau_n^1$.[6]

$$\begin{aligned} c_n^t = c_n^t(h_n^1, \tau_n^2) \qquad & \qquad \ell_n^2 = \ell_n^2(h_n^1, \tau_n^2) \\ p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} &= \ell_n^1 + \frac{\ell^2}{1+r} \end{aligned} \tag{3}$$

For the policy analysis in Section 6, we further assume that these Marshallian demands are derived from standard utility maximization with additively separable preferences over con-

[5] The former is a price normalization; the latter is a normalization on the size of a unit of labor supply.

[6] While we prove versions of all results in the general case where household behavior depends independently on ℓ_n^1 and τ_n^1, the assumption that they enter additively simplifies the exposition of the main text by equating the MPC out of first-period labor and transfer income. This rules out complementarities between labor supply and consumption, but is consistent with any utility function that is additively separable in consumption and labor supply. It also rules out certain behavioral interpretations, such as mental accounting of the sort studied by Lian (2019).

sumption and labor supply subject to the budget constraint as well as a borrowing constraint in the form of minimum savings $\underline{s}_n$.

$$\begin{aligned}(\ell_n^2, c_n^1, c_n^2) \in \underset{\ell^2, c^1, c^2}{\arg\max}\ & u_n^1(c^1) - v_n^1(\ell_n^1) + \beta_n \left(u_n^2(c^2) - v_n^2(\ell^2)\right) \\ \text{s.t } & p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell_n^1 + \frac{\ell^2}{1+r} \\ & \ell_n^1 - p^1 c^1 - \tau_n^1 \geqslant \underline{s}_n\end{aligned} \tag{3'}$$

for some time-and-household-specific consumption utility functions u_n^t, labor disutility functions v_n^t, and subjective discount factors β_n.

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.

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

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

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

A rationing equilibrium is therefore defined as follows:

Definition 1. *A rationing equilibrium is a profile of prices and quantities* $\{p_i^t, \ell_n^t, c_n^t, L_i^t, X_{ij}^t\}_{t \in \{0,1\}, n \in N, i,j \in I}$ *that satisfy conditions (1) – (5).*

In Appendix C.1, we provide mild technical assumptions that guarantee that an equilibrium exists and that prices are determined independently from demand, a non-substitution theorem.[7] As we do not consider supply shocks in the main text, we now normalize the units of consumption goods so that $p_i^t = 1$ for all $t \in \{1, 2\}$, $i \in I$. This does not stem from an assumption that prices are sticky, but rather follows from that wages are sticky and prices are technologically determined as a function of wages, by the non-substitution theorem.[8]

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

[7] The latter is a consequence of the fact that we have a single labor factor.

[8] Indeed, in Appendix A we consider the case of supply shocks, under which prices do move, underscoring that we have not assumed that prices are sticky but rather that they are technologically determined.

level final output,[9]

$$Y_i^t \equiv Q_i^t - \sum_{j \in I^1} X_{ji}^t \tag{7}$$

where aggregate output is simply given by:

$$\text{GDP}^t = \mathbb{1}^T Y_i^t \tag{8}$$

2.2. Interpreting the Model

The main assumptions that justify our focus on rationing – a binding a zero lower bound on nominal interest rates and downward nominal wage rigidity – have strong empirical support in the context of economic downturns in the United States. Indeed, a zero lower bound on nominal rates has constrained monetary policy in the United States and Europe for much of the last decade, particularly during recessions. A wide empirical literature documents that wages in the United States – both for incumbent workers and new hires – exhibit downward nominal wage rigidity, responding extremely little to both aggregate output and unemployment (Solon, Barsky, and Parker, 1994; Grigsby, Hurst, and Yildirmaz, 2019; Hazell and Taska, 2019).

At the same time, it is important to note that fixed intertemporal prices are inessential for our results on optimal targeting of fiscal policy. In particular, if monetary policy and wages respond to aggregate output (and not more disaggregated output), our results still recover the optimal (least cost) way of obtaining a given level of stimulus to aggregate output. Concretely, this is because the effects on real interest rates are common to all policies that cause identical aggregate stimulus, making these effects irrelevant for comparing the relative fiscal cost of such policies.

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.[10] While our analysis relies on the absence of relative wage changes, our microfoundation of this feature is not essential. In fact, 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

[9] A related but distinct notion of sector-level production is value added:

$$V_i^t \equiv Q_i^t - \sum_{j \in I} X_{ij}^t. \tag{6}$$

which also sums to aggregate output. In equilibrium, value added is always equal to factor demand, i.e. $V_i^t = L_i^t$.

[10] 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.

downward pressure on all wages, all wages are downwardly rigid, and relative wage-inflation expectations are sticky. Downward pressure on all wages is an appropriate assumption given our focus on severe economic downturns. Hazell and Taska (2019) document that downward wage rigidity is not merely an aggregate phenomenon, but exists across occupations, wage quartiles, and firm sizes.

Given that neither intratemporal nor intertemporal prices can adjust, markets must clear 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).[11]

3. The Generalized Multiplier

Within this setting, we explore the general equilibrium impact of shocks to government purchases and transfers. Our goal is to derive an expression for the generalized Keynesian income multiplier that maps the effect of shocks in partial equilibrium to their general equilibrium impact. 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 network structure in the macroeconomy, and that we later take to the data to study implications for the design of fiscal policy.

3.1. The Generalized Fiscal Multiplier in a Networked Economy

Our main results express the economy's general equilibrium responses to fiscal shocks as a function of their partial equilibrium effect on goods demand before incomes have been allowed to adjust. The first-period partial equilibrium effect of a shock to government purchases or

[11] Among the advantages of our reduced 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 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.

transfers is given by:

$$\partial Y^1 = dG^1 - \boldsymbol{C^1_{h^1}} \mu d\tau^1 + \boldsymbol{C^1_{\tau^2}} \mu d\tau^2 \tag{9}$$

where here ∂Y^1 and dG^1 are length-I vectors 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 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.

Our first proposition captures the way that an *income multiplier* amplifies a partial equilibrium shock ∂Y^1 to generate the general equilibrium change in the vector of first-period values added dY^1. To state this result, we let $\widehat{\boldsymbol{X}}^{\boldsymbol{1}}$ and $\widehat{\boldsymbol{L}}^{\boldsymbol{1}}$ denote the $I \times I$-dimensional unit-production input-output and labor demand matrices, respectively, in the first-period.[12] Similarly, we denote by $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}$ the 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 $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m} = \boldsymbol{C^1_{h^1}}$.

Throughout the rest of the paper, we assume that the Leontief-inverse matrix $(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}$ exists.[13] Moreover, in analogy to the assumption that the aggregate MPC is less than one in the simple Keynesian multiplier, we assume the moduli of $\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}\boldsymbol{R^1_{L^1}}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}$ and $\boldsymbol{R^1_{L^1}}\widehat{\boldsymbol{L}}^{\boldsymbol{1}}(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1}\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}$ are less than one, which guarantees that the generalized multiplier is well-defined.[14]

Proposition 1. *Given any rationing equilibrium, the local change in equilibrium first period final output dY^1 following a fiscal shock with partial equilibrium effect on first-period final output ∂Y^1 is given by:*[15]

$$dY^1 = \left(\boldsymbol{I} \; - \; \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \; \boldsymbol{m} \; \boldsymbol{R^1_{L^1}} \; \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \; \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \right)^{-1} \partial Y^1 \tag{10}$$

Proof. See Appendix A.1. □

[12] 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 7 and Corollary 1 in the Appendix. $\widehat{\boldsymbol{X}}^{\boldsymbol{1}}$ is the $I \times I$-dimensional matrix with i^{th} column equal to $\widehat{X}^1_i$; $\widehat{\boldsymbol{L}}^{\boldsymbol{1}}$ is the diagonal $I \times I$-dimensional matrix with (i,i) entry $\widehat{L}^1_i$.

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

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

[15] 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.

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

$$\underbrace{\widehat{\boldsymbol{C}}^{\mathbf{1}}}_{I\times N}\ \underbrace{\boldsymbol{m}}_{N\times N}\ \underbrace{\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}}_{N\times I}\ \underbrace{\widehat{\boldsymbol{L}}^{\mathbf{1}}}_{I\times I}\ \left(\boldsymbol{I} - \underbrace{\widehat{\boldsymbol{X}}^{\mathbf{1}}}_{I\times I}\right)^{-1} \tag{11}$$

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. Following a demand shock to firms, the term $\left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{-1}$ maps changes in final demand to changes in production via the input-output network. Having pinned down the change in required production, $\widehat{\boldsymbol{L}}^{\mathbf{1}}$ maps these to changes in firms' demand for labor. Next, the marginal rationing matrix $\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}$ 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 $\widehat{\boldsymbol{C}}^{\mathbf{1}}$ 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*.[16]

In terms of terminology, we refer to the matrix of Equation 10 that maps vectors of partial equilibrium changes in final output to vectors of general equilibrium 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 (2019). Moreover, the *interaction* between the input-output network and the directed consumption network matters: multipliers are largest when it is not only partial equilibrium 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

[16]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.

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 9). Specifically, we show that the multipliers are higher for any partial equilibrium 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 10).[17] 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. Our results may therefore be interpreted as either (a) the effect of changes in the targeting of a first-period fiscal program of a fixed size or (b) the *partial* effect of a general change in fiscal policy *occurring through changes in first-period policies* rather than through changes in second-period policy. We consider the general case in Appendix A.2.

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

$$\boldsymbol{\mathcal{G}} \equiv \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \widehat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} \widehat{\boldsymbol{C}}^{\mathbf{1}} \tag{12}$$

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

$$\partial h^1 \equiv \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \widehat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} dG^1 - \boldsymbol{\mu} d\tau^1 \tag{13}$$

as the partial equilibrium incidence of a shock to policy in the first period on household first

[17] 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.

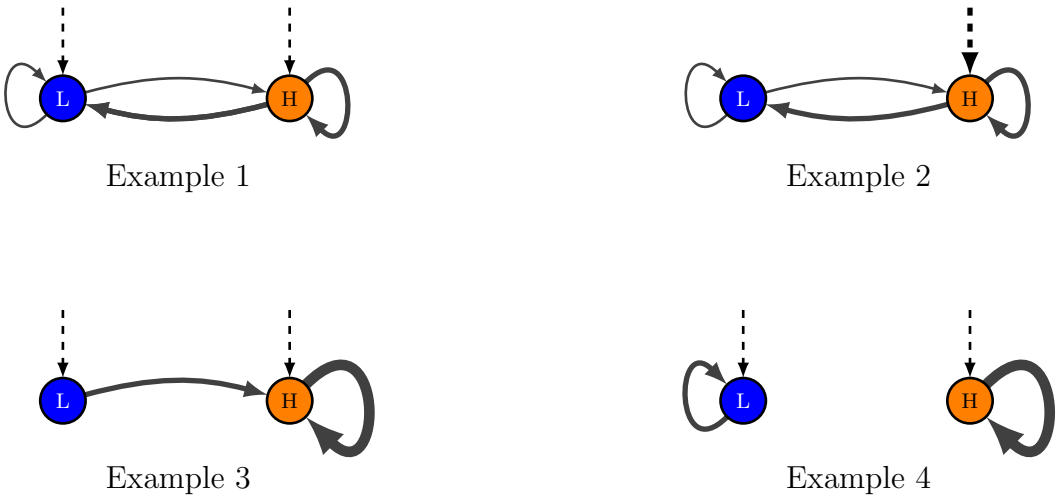

Fig. 1. H and L are nodes corresponding to the households with MPCs m_H and m_L, respectively. The thickness of solid arrows captures the magnitude of households' marginal spending. The thickness of dashed arrows capture the magnitude of the direct incidence of policies onto household income. Example 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").

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 partial equilibrium shock is "as if" the economy had a single household with MPC $\overline{m} = \frac{m_L+m_H}{2}$. In the example, the partial equilibrium 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 partial equilibrium 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 partial equilibrium 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 partial-equilibrium 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 the the higher-MPC households' MPC. In particular the multiplier is given by $1 + \frac{\overline{m}}{1-m_H} = 1.60$, generating 43% 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 33% 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.[18]

Proposition 2. *The total change in first-period GDP due to a change in first-period fiscal*

[18]We provide an exact decomposition in terms of Bonacich centralities of $\boldsymbol{\mathcal{G}}$ in Appendix A.2. The approximation in Proposition 14 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.

policy with income incidence ∂h^1 that sums to one can be approximated as:

$$\begin{aligned} dGDP^1 =& \mathbb{1}^T dG^1 + \frac{1}{1-\mathbb{E}_{h*}[m_n]}\Bigg(\underbrace{\mathbb{E}_{h*}[m_n]}_{\text{RA Keynesian effect}} + \underbrace{\mathbb{E}_{\partial h^1}[m_n] - \mathbb{E}_{h*}[m_n]}_{\text{Incidence effect}} \\ &+ \underbrace{\mathbb{E}_{\partial h^1}[m_n]\left(\mathbb{E}_{\partial h^1}[m_n^{next}] - \mathbb{E}_{h*}[m_n]\right)}_{\text{Biased spending direction effect}} + \underbrace{\mathbb{C}ov_{\partial h^1}[m_n, m_n^{next}]}_{\text{Homophily effect}} \Bigg) + O^3(|m|) \end{aligned} \tag{14}$$

where h^* is any reference income weighting that sums to one and $m_n^{next} = \left(m^T \boldsymbol{\mathcal{G}}\right)_n$ is the average MPC of households who receive as income i's marginal dollar of spending.[19]

Proof. See Appendix A.2. □

The above proposition holds for all reference partial equilibrium changes in net income h^* of unit size, but naturally the choice of this h^* affects the accuracy of the approximation. In this sense, each of these effects are relative, meaning that they capture the amplification relative to some reference shock to earnings. In our later empirical analysis, we take h^* as the change in income induced by a GDP proportional demand shock. In this case, we show that the error term always accounts for less than 0.3% of the multiplier, so that this approximation is very tight.

In Appendix C.3 we discuss how Proposition 2 applies to several benchmark economies, highlighting cases in which the various alterations to the Keynesian multiplier are zero. One important benchmark is a "homothetic economy" where both consumption and labor rationing functions are homothetic. In this case, a final-output-proportional shock has no incidence or bias effect, but heterogeneity in household consumption baskets and sectoral employment can still generate additional amplification through the homophily effect (Proposition 11, part 1.). An even more "neutral" case occurs when all firms in the economy employ workers at the margin who have the same average MPC as one another. In this case, the bias and homophily effects are zero and the income multiplier for any shock is simply the Keynesian multiplier evaluated at this average MPC (Proposition 11, part 2.). Note that even when the traditional Keynesian multiplier obtains, the aggregate MPC need not equal either the average MPC or the income-weighted MPC of the population; this is the case only when each firm's marginal employees have the population average 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.

[19] 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.

Moreover, one might reasonably expect that each effect formalized by Propostion 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 (2019) 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 heterogenous 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.

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 general equilibrium response of demand dY^1 to any partial equilibrium shock ∂Y^1 is given by:

$$dY^1 = \left(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1} \right)^{-1} \partial Y^1 \tag{15}$$

To estimate the generalized multiplier, we therefore 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

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 $\widehat{\boldsymbol{X}}^{\mathbf{1}}$ is an $(S \times I) \times (S \times I)$ matrix where S is the number of regions and I is the number of industries. 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. 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.[20] We also make an adjustment to account for linkages across industries in 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 compute its industrial composition with the same procedure – using the use, make, and import

[20]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.

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. 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.[21] For all nontradable industries, we assume that the commodity is sourced entirely within the state.[22]

Finally, we construct the regional input-output matrix by combining the national industry-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 $(S \times I) \times (S \times N)$ matrix where N is the number of demographic groups and recall S and I 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.[23]

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

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

[22] 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. We assume that all other sectors are domestically non-tradable. 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.

[23] 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.

Price Index (CPI) and Consumer Expenditure Survey (CEX) following the methodology in Patterson (2019). Specifically, we follow the procedure of Gruber (1997), using the panel structure of the PSID to estimate the equation:

$$\Delta C_{Ht} = \sum_x (\beta_x \Delta E_{Ht} \times x_{Ht} + \alpha_x \times x_{Ht}) + \delta_{s(H)t} + \varepsilon_{Ht} \tag{16}$$

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 $\delta_{s(i)t}$ is a state by time fixed effect. Guided by the frequency at which PSID data is available from 1997 to 2015, we consider two-year changes for both consumption and income. The first time period in our model should therefore be interpreted as lasting for two years. Conveniently, this is similar to the length of a typical recession. Estimating Equation 16, we obtain the following estimate of the MPC for household H at time t:

$$\widehat{MPC}_{Ht} = \sum_x \hat{\beta}_x x_{Ht} \tag{17}$$

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, 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.[24]

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

[24]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. Moreover, Patterson (2019) documents that both

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 G). For discussion of the robustness of these estimates, see Patterson (2019).

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:

$$\widehat{MPC}_{n(Ht)i} = \alpha_{n(Ht)i}\widehat{MPC}_{n(Ht)} \tag{18}$$

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 G), suggesting that we are 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:

$$\widehat{MPC}_{risn} = \lambda_{irs}\widehat{MPC}_{ni} \tag{19}$$

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.[25] We assume all nontradable goods are consumed within the state.

The procedure above provides the directed MPC entries for the 80 demographic groups.

[25]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 modelling 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 the results.

It remains to estimate the directed MPCs for shareholders and foreigners. For foreigners, we simply set all entries to zero. 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. For shareholders, we take the MPC out of stock market wealth as estimated by Chodorow-Reich, Nenov, and Simsek (2019) at 0.028. 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 will be assumed in Section 6. 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]

4.3. The Rationing Matrix

The rationing matrix $\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}$ is a $(S \times N) \times (S \times I)$ matrix where recall S, N, and I 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 the rationing matrix in three steps. We first use the American Community Survey (ACS) to compute, within each state-industry pair, the total labor earnings by each demographic group in 2012. We also use state-level data from the BEA on compensation and gross output by industry to compute labor shares of final output for each state-industry pair.

Second, we use these two components, along with the estimated demographic group MPCs, to construct the the labor rationing entries for workers. Concretely, we employ the following formula:

$$\left(\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}\right)_{rnsi} = \mathbb{I}[r=s]\frac{y_{inr}}{\sum_n y_{inr}}\alpha_{ir}\beta_i\left(1+\xi\left(MPC_n - \overline{MPC}_{ir}\right)\right) \tag{20}$$

[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.

where y_{inr} is total earnings of demographic n in industry i in region r, α_{ir} is state-by-industry labor share of value added, β_i is the national value added to gross output ratio in industry i, ξ is the correlation between MPCs and earnings elasticities, and $\overline{MPC}_{ir}$ is the earnings-weighted MPC of all workers in industry i in region r. The indicator function imposes the condition that all labor earnings are received within the state where production occurs. This is the unique functional form that both preserves a constant correlation between MPC and earnings elasticity, of which there is strong evidence from Patterson (2019) and preserves total income received across all demographic groups in each industry-region pair. We set $\xi = 1.332$, the correlation of MPC with earnings elasticity to aggregate shocks measured in Patterson (2019).[27] While our model can in principle incorporate regional migration in response to shocks, we – by assuming that employment at each firm only depends on its own labor demand – only partially allow for this possibility. In particular, our calibration rules out the possibility that the share of labor each firm rations from each demographic group may depend on changes in the group's share of the population due to migration.

Finally, it remains to allocate factor payments that are not received by labor. These take two forms: payments made to the 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, from the BEA use table. We distribute a share of these profits domestically, according to the domestic share of US equity ownership in 2012, around 86.5%. We allocate the profit of each domestic firm nationally, according to each demographic by state's share of dividend income in the Internal Revenue Service's Statistics of Income (IRS SOI) data. 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.[28]

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 GDP and demand spillovers across regions. We begin by

[27] See Patterson (2019) for more details and discussion.

[28] 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.

quantifying how the variation in the fiscal multiplier for government purchases or transfers – i.e. the total change in GDP 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.

5.1. Extent of Heterogeneity in Multipliers

We estimate that the response of GDP to a demand shock which is GDP-proportional across industries and regions, or *aggregate purchases multiplier*, is equal to 1.30, 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 GDP 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 six. Indeed, our range of multipliers, from about 1.1 to 1.6, provides one rationalization for the variation in purchases multipliers estimated in the literature. Transfers multipliers, i.e. the effect on aggregate GDP 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 GDP of transferring a dollar to a household ranges from slightly below zero for some households (some types have negative MPCs) to nearly 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 a factor of 1.5 across states (see Figure A18 in Appendix G); transfers multipliers differ by a factor of 1.3 across states and by nearly as much across demographic groups as across demographic-region pairs (See Figure A19 in Appendix G).

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*

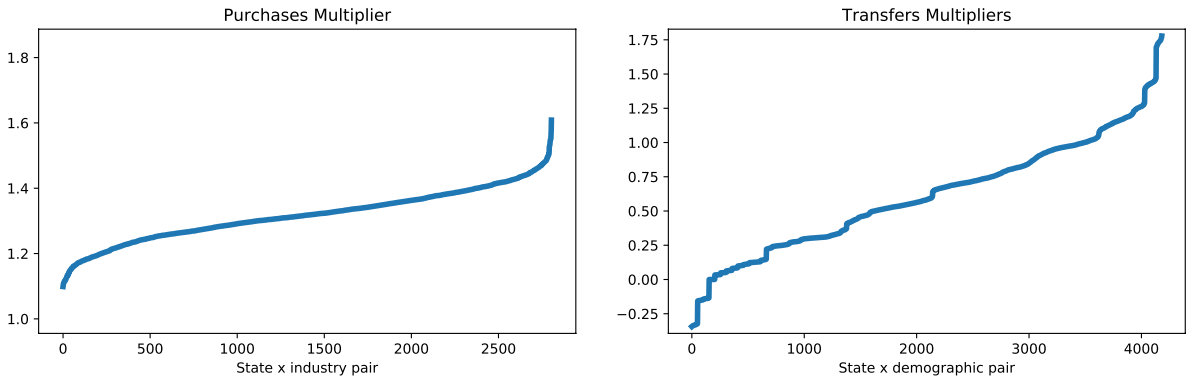

Fig. 2. Left: Change in aggregate output resulting from a $1 government purchase shock to each state-by-industry pair, i.e. distribution of government purchase multipliers. The x-axis sorts state-by-industry pairs by their purchase multipliers and the y-axis plots the estimated multiplier. Right: Change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

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 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 and 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

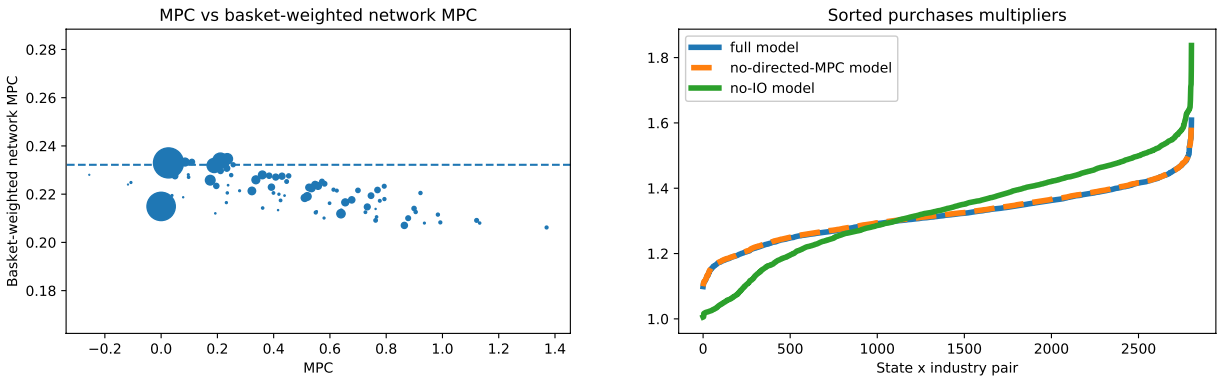

Fig. 3. The left panel shows a scatter of MPCs m_n against basket-weighted MPCs m_n^{next}. The dashed line gives the average MPC $\mathbb{E}_{h*}[m_n]$ for h^* given by the income incidence of a shock to demand proportional to 2012 state-industry GDP. The right panel shows the change in GDP for each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

more than 99 percent of the multiplier.[29] To drive this point home, the orange 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.[30]

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 A9 in Appendix G), 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 Figure A5 in

[29]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 A4 in Appendix G, respectively). Across the full distribution of shocks, the contributions of the bias and homophily terms range between zero to four tenths of a percent increase in 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 A4 in Appendix G) and find that they are uniformly an order of magnitude smaller than the bias and homophily terms. Our approximation is therefore very tight for any feasible shock.

[30]In Figure A8 of Appendix G we show a scatter plot of the multipliers from these two models. The correlation in multipliers across the two models is nearly perfect.

Appendix G), directing more spending toward capital, the owners of which have low MPCs.[31] Second, our estimates feature substantial within-region non-tradeables demand, with around a third of total labor demand remaining within the state from which consumption originates (see Figure A10 in Appendix G). Moreover, there is spatial heterogeneity in MPCs, with income-weighted MPCs differing by a factor of 1.5 across states (See Figure A16 in Appendix G). 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.

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.

We consider alternatives to our baseline estimates along four dimensions. 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. To bound the importance of this assumption, we consider two extreme alternative calibrations for interstate trade, one in which all buyers of a good 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-

[31] Conditional on reaching labor, the average MPC of workers producing consumption baskets is very slightly increasing across the MPC distribution (see Figure A5 in Appendix G), so labor share differences account for the bulk of differences in basket-weighted MPCs stemming from heterogeneous consumption baskets. This finding is also consistent with the empirical patterns in Hubmer (2019).

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 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.

Fourth, our baseline estimates assume, following Patterson (2019), 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.

Appendix Figure A7 shows the size of the bias and homophily effects in the $(3+1)\times 2\times 2 = 16$ versions of the model described above. In all cases, the bias and homophily effects contribute less than 0.01 to the multiplier of a GDP-proportional shock. 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 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 causes 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 lots 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 A11 in Appendix G, 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 A12 and A13 in Appendix G 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 green line corresponds to the model without input-output linkages, which features a much more dispersed distribution of multipliers.[32] 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.[33] 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.

5.3. Regional Demand Spillovers

Finally, we turn our focus away from the effects of fiscal shocks on *aggregate* GDP 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.[34] 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-

[32]See Figure A17 in Appendix G for a scatter plot of the multipliers across both the full model and that without input-output linkages.

[33]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.

[34]We consider the problem of such a planner in Section 6.3.

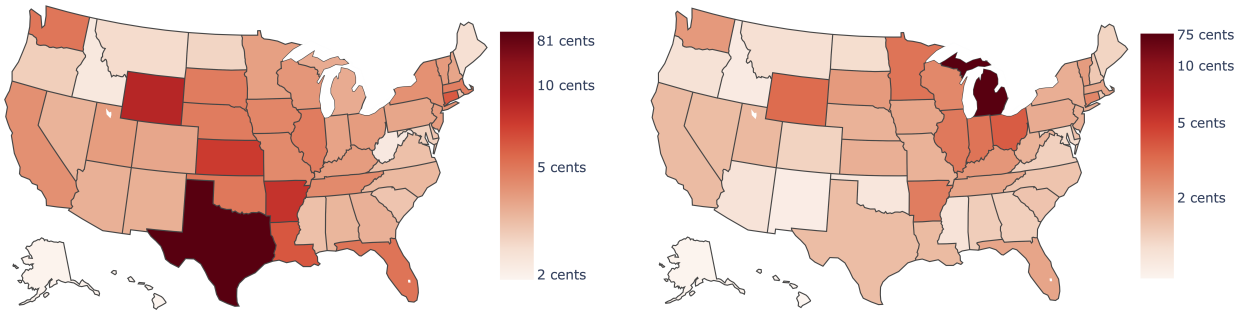

Fig. 4. Changes in state GDP per capita following a GDP-proportional $1-per-capita purchases shock to Texas (left panel) and Michigan (right panel).

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 GDP increases by $1.31 in response to $1 of additional spending. However, only an average of 74 cents of this increase in GDP occur within the state to which purchases are directed, with the remaining 57 cents diffusing to other states. These interstate spillovers occur 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.[35]

Figure 4 shows these spillovers cartographically, plotting the per-capita changes in state GDP that result from $1-per-capita shock to output in Texas and Michigan, respectively.[36] Consistent with a regional bias in trade, we observe slightly greater increases in GDP-per-capita in states that neighbor those shocked directly.

[35] Appendix Figure A14 shows the results of an analogous exercise for state-specific uniform transfers rather government purchases. In this case, stimulus reaches household incomes directly rather than spreading across states through the input-output network. We estimate that, out of the 70-cent increase in output resulting from a $1 uniform transfer to households in an average state, only 29 cents—less than half—occurs within that state. Intuitively, the spending-to-income network crosses more state lines than does the IO network, since it directs money across states not only through IO network, but also through household purchases of goods from other states.

[36] Appendix Figure A15 shows versions of these maps that measure spillovers in absolute, rather than per-capita terms.

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 consistent with this, as our estimated spillovers are largest for the geographically closer states.[37] 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 GDP and income 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 begin by characterizing the motives of the social planner and clarifying how the estimated multipliers from Section 5 directly inform the design of policy. We break our subsequent analysis into two parts: First, we consider a setting with widespread underemployment. Here, we assume the social planner seeks solely to maximize aggregate income and show that – due to the bias and homophily effects' empirically small role in shaping multipliers – they can achieve this objective with simple MPC targeting. Second, we consider the more general case when underemployment is more severe in some regions than others. We illustrate the effectiveness of targeting policy on the combination of MPCs and labor wedges with an application to the Great Recession.

6.1. Welfare and Fiscal Policy

In order to compute the welfare effects of fiscal policy, we adopt the micro-foundation of households' Marshallian demands stated in (3'), in which households maximize utility taking first-period labor supply as given, subject to budget and borrowing constraints. Recognizing that each household n's labor income is a function of fiscal policy (G, τ), we denote its indirect utility by $W_n(\ell_n^1(G, \tau), \tau_n)$.

[37] In Appendix F, we more formally explore the extent to which our model predicts the cross-state spillovers in response to identified demand shocks. We find that we are empirically underpowered to assess this claim. Indeed, the high dimensionality of the spillovers and the lack of statistical power motivates our semi-structural approach to uncovering spillovers.

We assume that the planner is utilitarian, placing some welfare weight λ_n on households of type n. Social welfare is then given by:

$$W(G,\tau) \equiv \sum_{n \in N} \mu_n \lambda_n W_n(\ell_n^1(G,\tau),\tau_n) \tag{21}$$

Below, we denote by $\widetilde{\lambda}_n \equiv \lambda_n u_{ni}^1(c_n^1)$ the planner's marginal value of transferring a dollar to n in the first period.[38] We denote by $\Delta_n \equiv \frac{v_n^{1\prime}(\ell_n^1)}{u_{ni}^1(c_n)} - 1$ each household's labor wedge, which measures how far households are off their intratemporal labor supply curves.

Proposition 3 characterizes the welfare impact of a change in first-period fiscal policy in terms of these welfare statistics. In particular, the Proposition fixes second-period policy – as well as *total* government spending in the first period – and considers how first period policies should be targeted for a given level of spending.[39]

Proposition 3. *The change in welfare dW due to a small change in taxes and government purchases in the first period can be expressed as:*

$$dW = \sum_{n \in N} \mu_n \widetilde{\lambda}_n \left[\underbrace{-\Delta_n d\ell_n^1}_{\text{Address under-emp.}} - \underbrace{d\tau_n^1}_{\text{Make transfers}} \right] \tag{22}$$

Proof. See Appendix A.3. □

Equation 22 clarifies that, in general, the social planner aims to do more than simply maximize aggregate GDP. First, she seeks to alleviate involuntary un(der)employment by changing the labor allocation so as to provide more employment to households with large negative labor wedges (i.e. the underemployed). Second, she may make transfers between households, in the name of pure redistribution. Finally, different households have differing marginal utility and potentially different social welfare weights, determining the relative weight placed on these effects.[40]

[38] Here, i is any good of which n consumes a positive amount at t.

[39] For a more general result that allows for changes in second-period policies, see the proof in Appendix A.3.

[40] In Appendix D.2, we show that Proposition 3 carries over directly to environments with non-zero markups in the first period. Intuitively, profit owners can be thought of as providing capital services with completely elastic supply. This allows us to treat capital owners "as if" they simply supply labor and are rationed to in proportion to firms' markups. The only modification required to accommodate this broader interpretation is that the generalized income multiplier must be extended to include capital income. This interpretation contrasts sharply with Baqaee (2015), who proposes that a labor-wedge-reducing planner should target the industry with the highest network-adjusted labor share. The difference comes from the fact that Baqaee's model features competitive firms (hence no markups) and efficiently-allocated capital (no capital wedge).

This result moreover allows us to derive necessary conditions for optimal fiscal policy. In Proposition 12 in Appendix C.4, we do this and show that locally optimal fiscal stimulus balances five possible motives: a Keynesian stimulus channel to alleviate wedges from involuntary unemployment; a generalized opportunistic channel of Werning (2011) capturing that valuable fiscal stimulus is cheaper in welfare terms when there is underemployment; redistributing income between agents; relaxing agents binding borrowing constraints; and a short-termism channel when agents have binding borrowing constraints and a higher value for stimulus today. We moreover provide simple tests that do not require detailed structure on the network structure of the economy for the optimality of stimulus in Proposition 13 in Appendix C.5.

Of course, the planner cannot directly reallocate labor income, but rather must induce a reallocation of income through the multiplier effects on spending and transfer policies. Nevertheless, as we have empirically estimated the equilibrium mapping between fiscal policy and income (as given in Proposition 1), we can calculate how any policy affects $d\ell^1$ and therefore dW. We can now apply our findings from Section 5 to understand the welfare effects of targeting fiscal stimulus.

6.2. Simple MPC-Targeting

We first focus on the case where the social planner solely seeks to maximize aggregate income. This corresponds to the case where the direct social value of transfers $\widetilde{\lambda}_n$ is equal across households and all labor is rationed on the margin to un(der)employed households, who have no marginal disutility of labor.[41] We view this second condition as sensible in the context of a severe depression, where underemployment is widespread and not concentrated in particular demographic groups or regions. In the empirical case demonstrated above, where the bias and homophily effects are zero for all possible policies, Propositions 2 and 3 imply that, when MPCs out of labor and transfer income are the same, the change in welfare from first period fiscal policies is given by:[42]

$$dW \propto \sum_{n \in N} m_n \partial h_n^1 \tag{23}$$

[41] $\widetilde{\lambda}_n$ is equal across households when the social planner is, on the margin, indifferent to transferring a dollar from any one household to any other. While this formulation is consistent with borrowing constraints, a neutral utilitarian planner (i.e. λ_n constant) would, in many models, prefer to transfer money to households who are more borrowing constrained, since they have higher marginal utilities of income today.

[42] For a formal statement and proof, see the proof of Proposition 3.

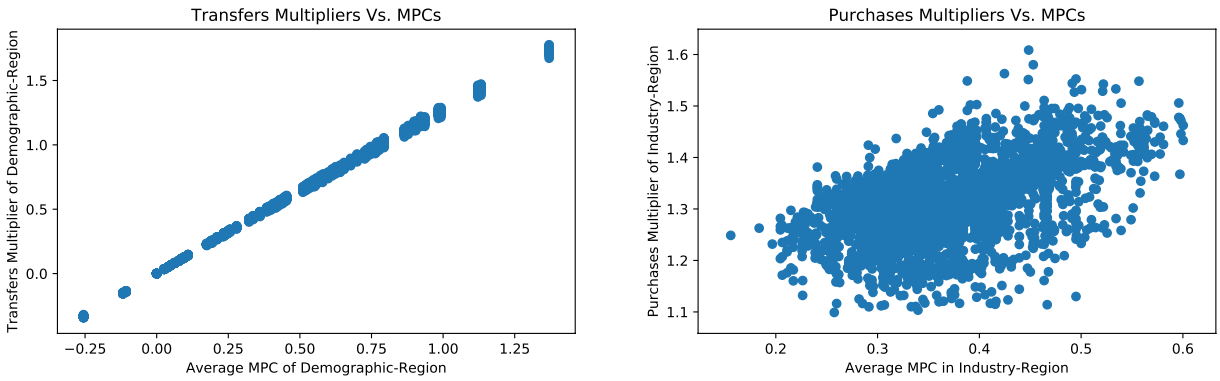

Fig. 5. Left: the effectiveness of targeting transfer stimulus by household MPCs. Right: the effectiveness of targeting government purchases stimulus by the average MPC of workers in each industry-region pair.

and recall that $\partial h^1 = \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \widehat{\boldsymbol{L}}^{\mathbf{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1} dG^1 - \boldsymbol{\mu} d\tau^1$ is the partial equilibrium change in household incomes induced by the fiscal policy. Thus, the change in welfare 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 income 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 an income-maximizing planner can do is simply target households with the highest MPCs.

Figure 5 empirically demonstrates the effectiveness of simple MPC targeting for maximizing aggregate output. The left panel scatters household MPCs against the resulting transfers multiplier from giving them a dollar, revealing an effectively perfect relationship between the two.[43] 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 GDP.

By contrast, the right panel shows that 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

[43]Note that the MPC that we use in Figure 5 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).

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 5 shows how naively targeting sectors employing the highest 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 have to be sacrificed to obtain that direct value, enabling a policymaker with knowledge of the value of direct government purchases to determine which policy to optimally pursue.

6.2.1. 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 GDP due to a shock that targets each industry-region pair proportionally to its value added is 1.30. While this number is sizeable, 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 1.61 in the oil and gas extraction industry in Georgia – the additional GDP induced by the same amount of spending policy would be twice as large. For transfer spending, we find that uniformly distributing a dollar to all household would lead to an increase in GDP of 77 cents. In contrast, if the government instead gave that dollar to the group with the highest multiplier – which we estimate is black men in South Carolina aged between 25-35 who earn less than \$22,000 – it would generate 1.78 additional units of final output, a 130% 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.[44] Putting these stylized transfers into our model,[45] we find that GDP increases by 79 cents for each dollar spent. Figure 6 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 MPCs. For example, the value at $2,000 shows the multiplier that the model predicts if the government gives $2,000 dollars to each worker in order of their MPCs until they exhaust their budget. This calculation shows that, with a maximum transfer size of $1,200, the multiplier on the income-targeted transfer (0.79) is very close to the multiplier with maximal MPC targeting (0.8), 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-MPC workers. Indeed, increasing the transfer to $2,500 produces a multiplier of 1.02, almost 30% higher than the benchmark policy, with the same budget.[46]

6.3. Fiscal Policy with Localized Shocks

While our theoretical and empirical results imply very simple fiscal rules in the case where the planner seeks to simply maximize aggregate GDP, our framework also provides a flexible toolkit for evaluating the welfare effects of fiscal policy targeting more localized economic downturns, where not all marginal labor is supplied by underemployed households with zero marginal disutility from labor. For example, suppose the initial shock to which the policymaker is responding is very concentrated in some areas and underemployment is less widespread. In this case, if the planner continues to have no redistributional preferences, the welfare effect of government purchases equals:

[44]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.

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

[46]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.

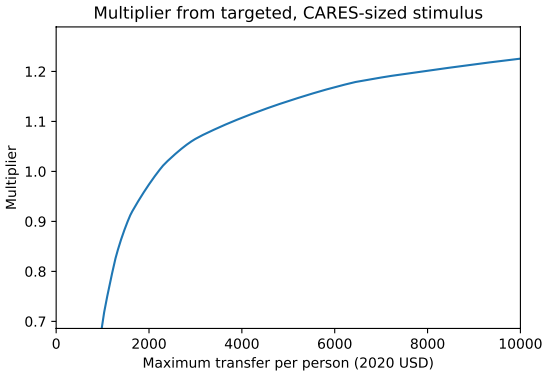

Fig. 6. Transfer amounts are inflated to 2020 dollars.

$$dW = -\sum_{n \in N} \mu_n \Delta_n d\ell_n^1 \tag{24}$$

Equation 24 demonstrates that the planner does not simply wish to maximize aggregate income, but also wants to direct stimulus to those households who are most severely underemployed. Accordingly, the planner should target not only households with high MPCs but also those with high labor wedges, those who buy goods produced by households with high labor wedges, etc.

To illustrate how our framework can be applied in this setting, we study the optimal targeting of government purchases during the Great Recession, which – although widespread – had much more severe impacts on certain regions and demographic groups. To do so, we augment our existing estimates with estimates of regional-demographic-specific rationing wedges. In Appendix E, we provide two microfoundations in which the rationing wedge for each demographic group in each state is given by the percentage change in labor hours worked by that group in the Recession relative to the preceding period.[47] To compute the welfare effects of fiscal policy, we can then simply combine changes in hours worked at the state-demographic level in the ACS from 2005-6 and 2009-10, and take the product with the induced spending-to-labor-income map that we have already estimated. This delivers the welfare gain from the stimulus benefit associated with spending one dollar in a specific industry in a specific state in the middle of the Great Recession.

This analysis illustrates that in the presence of localized shocks, targeting industry-regions simply based on their multipliers or their labor wedges alone is somewhat effective but leaves significant gains on the table. To the first point, we find that – across all sector-

[47]In particular, this is true if either (i) all households within a group are homogeneously employed, have quadratic labor disutility and apply a zero utility discount rate to the future or (ii) all households within a group are probabilistically totally unemployed or fully employed.

region pairs – the multiplier on a dollar of government purchases is somewhat, but not perfectly predictive of its welfare effect, with an estimated R^2 of 69% (see the left panel of Figure A6 in Appendix G). The average level of labor wedges of workers in a given region and industry is similarly predictive of the welfare effect of stimulus targeting that industry and region, with an R^2 of 72% (See the right panel of Figure A6 in Appendix G). Combined, the average level of labor wedges and multipliers have an R^2 of 78% in predicting welfare effects. Thus, while the planner can achieve large welfare gains by considering either sectoral-regional multipliers or underemployment in isolation, there are also welfare gains from considering them jointly and incorporating more detailed information.[48]

7. Historical and Counterfactual Exercises

In order to shed light on how our analysis might differ in other countries or time periods within the US, we perform three counterfactual exercises. Each exercise varies one the three key blocks of the generalized multiplier in Proposition 1 – the IO network, employment linkages, or consumption patterns – while keeping the other two fixed at our estimates for the 2012 US economy. We alternately consider the effects of weakened input-output linkages, historical (i.e. higher) sectoral labor shares, and increased income inequality.

7.1. The Role of IO Linkages

A salient feature of our calibration is that firms have large input shares (the average is around 45%). As a result, demand shocks to one industry-region are spread across other, upstream sectors as well as other regions – even before reaching households. In this section we explore the effects of the rich IO linkages in our economy by considering a stark counterfactual in which firms have zero input shares, instead directing revenues only to primary factors.

More formally, we modify our calibration in the following way: First, we set $\widehat{X}^1 = 0$, so that no firms use inputs. Second, so as to hold fixed the total cost of production, we rescale each firm's marginal demand for each labor type, capital, and foreign factors proportionally to the reduction in input expenditure. Interpreting this alternative calibration as a "no-inputs" economy, we ask two questions. First, how does input usage shape the fiscal multiplier of a GDP-proportional purchases shock? Second, how does input usage shape the degree of heterogeneity in purchases and transfers multipliers across different shocks?

[48] As the Great Recession featured widespread underemployment, it is perhaps unsurprising that around two thirds of the welfare gains from fiscal stimulus can be explained by the size of fiscal multipliers. In the case of a more localized shock or recessionary episode, heterogeneity in wedges would play a greater role; our framework still facilitates such an evaluation.

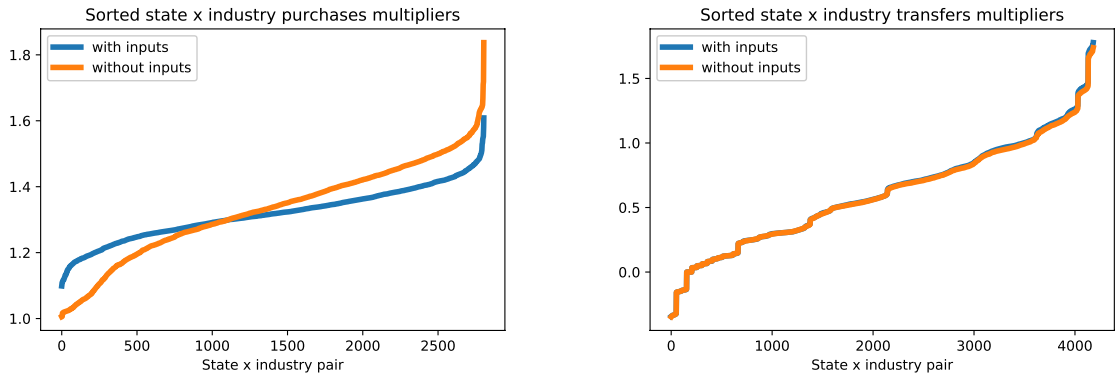

Fig. 7. Multipliers for state-industry-level purchases shocks and state-demographic-level transfers shocks. Differences in input linkages are more relevant for purchases shocks.

We begin with the aggregate question. Recall we have already shown in Section 5.2.1 that our finding of negligible bias and homophily effects holds in a range of alternative calibrations (including this one). How IO linkages affect the fiscal multiplier of a GDP-proportional purchases shock, then, only depends on how they determine the shock's incidence onto high- or low-MPC households. As an empirical matter, we find that the presence of IO linkages has no effect on incidence: the aggregate purchases multiplier is 1.30 in either case.

However, IO linkages do affect the *distribution* of multipliers. In fact, starting from an economy with zero bias and homophily effects, weakening IO linkages always increases the difference between the largest and smallest industry-region purchases multipliers, as we show theoretically in Appendix C.2. As discussed in Section 5.2.3, this is because IO linkages serve to dilute the incidence of a shock by spreading shocks to firms with the highest-MPC (or lowest-MPC) workers across others in their supply chain whose workers have more moderate MPCs. Figure 7 shows that, in practice, removing inputs spreads out not only the highest and lowest industry-region multipliers, but also the whole distribution. At the same time, the distribution of transfers multipliers is unaffected, since transfers reach households directly, rather than through the IO network. Thus, targeting government purchases (but not transfers) is relatively more important in economies with weaker input linkages.

7.2. The Decline of the Labor Share

Our model allows us to consider not only hypothetical counterfactuals but also actual, historical changes in the structure of the economy. One salient change over the past several years is the well-documented decline in the labor share in the US (Karabarbounis and Neiman, 2014; Dorn, Katz, Patterson, and Van Reenen, 2017). Indeed, Hazell (2019) provides empirical evidence that this reduction in the labor share has dampened unemployment

fluctuations. In this section, we perform a similar exercise in our model, comparing the purchases and transfers multipliers as industry-specific labor shares change from their 2000 to 2012 levels. Intuitively, if spending is directed away from high-MPC workers and toward low-MPC shareholders, aggregate amplification should fall.

Our methodology is as follows. We assume that, within each year and each industry, the shares of employee compensation in revenue is constant across states. We obtain these shares from the BEA use tables in 2000 and 2012. The left panel of Figure A20 shows the distribution of labor shares of revenue by industry in each year. The aggregate labor share of value added fell from 59.2% in 2000 to 54.9% in 2012; the aggregate labor share of revenue fell from 32.1% to 30.0%. We maintain our earlier, 2012-based, estimates of demographic-specific consumption baskets and MPCs, demographic employment by region, and input-output network. We allocate the difference in labor income between 2000 and 2012 to a factor with MPC zero; this can be understood as a foreign factor or as profits accruing to MPC-zero shareholders.[49]

Unsurprisingly, in line with our theoretical results in Appendix C.2, the reduction in the labor share leads to a smaller fiscal multiplier, as revenues are directed to lower-MPC households. We estimate an aggregate purchases multiplier of 1.338 in 2000 and 1.300 in 2012. Figure 8 shows the sorted distributions of purchases and transfers multipliers across all shocks, for 2000 and 2012. The distribution of purchases multipliers shifts down, as less of the income from a given change in demand flows to workers and more flows to low-MPC factors. Still, the multiplier does not fall for every state-industry pair. The right panel of Figure A20 shows that a few industries – namely those with sufficiently increased labor shares – have higher multipliers in 2012 than in 2000.

For transfers multipliers, the response to changing labor shares is almost zero. This is because transfers target households of each MPC directly, so that differences in the labor share only affect the multiplier to second order in MPCs.

7.3. Rising Labor Income Inequality

While the previous exercise speaks to changes in the distribution of aggregate income between labor and capital, labor incomes themselves have also become more unequally divided between US workers over the last several decades, seen both in an increase in the college wage premium and a steep rise in the labor incomes at the top of the distribution (Smith, Yagan, Zidar, and Zwick (2019), Autor (2014)). We consider the effect of a "hollowing

[49]We obtain qualitatively and quantitatively similar results if we allocate this income to capitalists with MPC 0.028, following Chodorow-Reich et al. (2019).

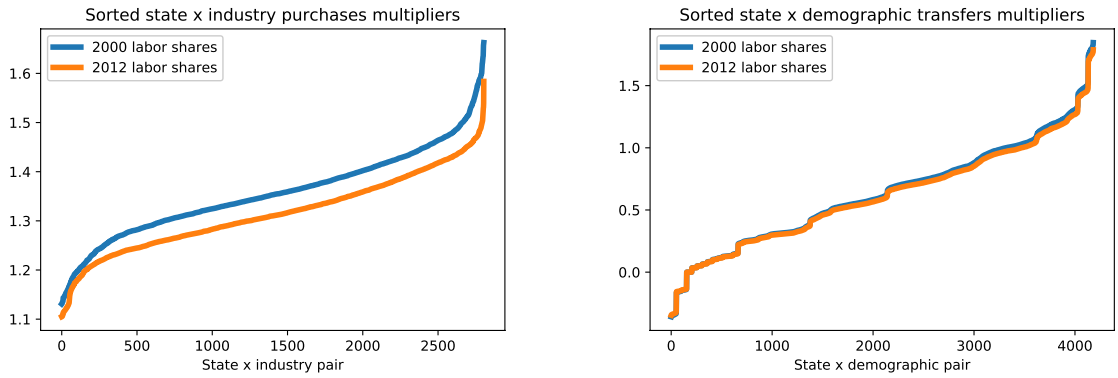

Fig. 8. Multipliers for state-industry-level purchases shocks and state-demographic-level transfers shocks. Differences in labor shares are more relevant for purchases shocks.

out" of the income distribution within occupations, for example due to a within-industry decline in cognitive routine tasks, on the distribution of multipliers. Concretely, within in each industry, within each race-sex-age group, we reallocate labor income initially earned by the middle income group evenly between the lowest and highest income earners of the same demographic and industry of employment. This has an almost zero effect on the distributions of multipliers. This is because, within race-age-sex groups, the dependence of MPC on income is approximately linear (see Figure A21). Therefore, while labor income inequality may in principle affect fiscal multipliers, it must do so by changing MPCs conditional on income or re-sorting workers across industries or regions.

8. 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 transfers among the highest MPC households can increase the effect of the policy on GDP by up to 130%. Despite the empirical unimportance of the details of the economy for fiscal multipliers, large spillovers make fine-grained modelling of the network structure of the economy necessary for understanding the distributional 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 responding to the Great Recession and COVID-19 may have left substantial gains on the table – on the order of several hundred billions of dollars.

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.

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.

Autor, D. H., 2014. Skills, education, and the rise of earnings inequality among the "other 99 percent". Science 344, 843–851.

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.

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

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.

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., 2019. Stock market wealth and the real economy: A local labor market approach. Tech. rep., National Bureau of Economic Research.

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.

Dorn, D., Katz, L. F., Patterson, C., Van Reenen, J., 2017. Concentrating on the fall of the labor share. American Economic Review 107, 180–85.

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.

Grigsby, J., Hurst, E., Yildirmaz, A., 2019. Aggregate nominal wage adjustments: New evidence from administrative payroll data. Tech. rep., National Bureau of Economic Research.

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.

Hazell, J., 2019. The falling labor share dampens unemployment fluctuations. Tech. rep.

Hazell, J., Taska, B., 2019. Downward rigidity in the wage for new hires. Tech. rep.

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

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.

Karabarbounis, L., Neiman, B., 2014. The global decline of the labor share. The Quarterly journal of economics 129, 61–103.

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

Lian, C., 2019. A theory of narrow thinking. The Review of Economic Studies .

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., 2019. The matching multiplier and the amplification of recessions. In: *2019 Meeting Papers*, Society for Economic Dynamics, no. 95.

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.

Smith, M., Yagan, D., Zidar, O., Zwick, E., 2019. Capitalists in the Twenty-First Century. The Quarterly Journal of Economics 134, 1675–1745.

Solon, G., Barsky, R., Parker, J. A., 1994. Measuring the cyclicality of real wages: how important is composition bias? The quarterly journal of economics 109, 1–25.

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

Werning, I., 2011. Managing a liquidity trap: Monetary and fiscal policy. Tech. rep., National Bureau of Economic Research.

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.

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

A. Omitted Proofs

The formal results in Sections 2, 3, and 6 are stated for the model presented in the text, but also hold in a generalization of the model in which households experience preference shocks (such as discount rate shocks) and—except where noted otherwise—firms experience technology shocks. They also allow household consumption to respond differently to first period labor supply and first-period transfers.

We model technology shocks in a reduced form way by allowing the production function of each firm j in each period t to depend on a technology parameter z_j^t, so that its output is given by $F_j^t(X_j^t, L_j^t, z_j^t)$. Since, in the presence of technology shocks, the vector of all prices $p = \{p_i^t\}_{t\in\{1,2\},i\in I}$ is endogenous, we assume that households' Marshallian demands depend on prices in addition to taxes and first-period labor supply.

As supply shocks can affect input demands directly, partial equilibrium supply shocks directly affect gross output – i.e. the vector of firm production $Q_i^t = F_i^t(X_i^t, L_i^t, z_i^t)$ – rather than final output. As it is strictly more general, we will in the context of supply shocks consider partial equilibrium shocks ∂Q^1 to first-period gross output rather than to first-period final output (∂Y^1 of the main text).

We model household preference shocks in a similarly reduced form way by allowing each household type n's first- and second-period household consumption and second-period labor supply to depend additionally on a preference shock θ_n. Incorporating preference shocks and prices, consumption and second-period labor supply are therefore given by $c_j^t(p, \ell_n^1, \tau_n, \theta_n)$ and $l_j^2(p, \ell_n^1, \tau_n, \theta_n)$. Throughout our proofs we let $\theta \equiv (\theta_1, ..., \theta_n)$ and denote by $C_i^t(p, \ell^1, \tau, \theta)$ aggregate consumption demands as a function of, among other things, preference shocks:

$$C_j^t(p, \ell^1, \tau, \theta) = \sum_{n\in N} \mu_n \; c_{nj}^t(p, \ell_n^1, \tau_n, \theta_n). \tag{A1}$$

A.1. Proof of Proposition 1

We prove the proposition in the case of more general partial equilibrium shocks to gross output, thereby nesting supply and preference shocks in addition to fiscal shocks.

Proposition 4. *Given any rationing equilibrium, the local change in equilibrium first period final output dY^1 following any small shock to parameters contained in Span$\{dG, d\theta, d\tau, dz\}$*

with a first-period partial equilibrium effect ∂Q^1 is given by:[50]

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

Proof. We simply apply the implicit function theorem to the goods market clearing condition. Namely, our differentiability assumptions allow us to express

$$\begin{aligned} dQ^1 \; &= \widehat{\boldsymbol{X}}^{\boldsymbol{1}} dQ^1 + \widehat{\boldsymbol{X}}^{\boldsymbol{1}}_{\boldsymbol{z}} dz Q^1 + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{p}} \boldsymbol{p}_{\boldsymbol{z}} dz + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\ell^1}} d\ell^1 + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\tau}} d\tau + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\theta}} d\theta + dG \\ d\ell^1 \; &= \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} dY^1. \end{aligned} \tag{A3}$$

Defining $\partial Q^1 = \widehat{\boldsymbol{X}}^{\boldsymbol{1}}_{\boldsymbol{z}} dz Q^1 + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{p}} \boldsymbol{p}_{\boldsymbol{z}} dz + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\tau}} d\tau + \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\theta}} d\theta + dG$, recalling the definition $\widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} = \boldsymbol{C}^{\boldsymbol{1}}_{\boldsymbol{\ell^1}}$, and substituting for $d\ell^1$, we have

$$\left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right) dQ^1 \; = \; \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} dY^1 + \partial Q^1 \tag{A4}$$

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

In this Appendix, we provide a general version of Proposition 2 in order to accommodate (a) changes to fiscal policy in both periods, (b) preference shocks, and (c) more general household preferences.

We consider the general case in which household behavior may depend flexibly on the variables outside of their control. Namely, consumption c^t_n and second-period labor supply ℓ^2_n are given by Marshallian demands

$$c^t_n = c^t_n(\ell^1_n, \tau_n, \theta_n) \qquad\qquad \ell^2_n = \ell^2_n(\ell^1_n, \tau_n, \theta_n) \tag{A5}$$

where here we have excluded price dependence since we will not consider supply shocks.

Next, for any change in fiscal policies and preferences $\Delta \equiv (\Delta G, \Delta\tau, \Delta\theta)$, define by $\partial_\Delta c^1_n$ each household of type n's change in first-period consumption due to the direct effects Δ on labor supply, income, and preferences, to first-order in Δ

$$\partial_\Delta c^1_n \; \equiv \; c^1_{n\ell^1_n}(\ell^1_n, \tau_n, \theta_n) \frac{1}{\mu_n} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1} \Delta G^1 \; + \; c^1_{n\tau_n}(\ell^1_n, \tau_n, \theta_n) \Delta\tau_n \; + \; c^1_{n\theta_n}(\ell^1_n, \tau_n, \theta_n) \Delta\theta_n. \tag{A6}$$

[50] In the proof, we derive an explicit expression for ∂Q^1

Defining $m_{n\Delta} \equiv \mathbb{1}^T \partial_\Delta c_n^1$, we express $\partial_\Delta c_n^1$ as the product of n's MPC out of Δ with its MPC direction out of Δ , i.e. $\partial_\Delta c_n^1 \equiv \widehat{c}_{n\Delta}^1 m_{\Delta n}$. Finally, we let $\boldsymbol{m_\Delta}$ be the $N \times N$ diagonal matrix with entries $m_{\Delta n}$, we let $\widehat{\boldsymbol{C}}_{\boldsymbol{\Delta}}^{\boldsymbol{1}}$ be the $I \times N$ matrix of consumption directions $\widehat{c}_{n\Delta}^1$ across households, and we let

$$\boldsymbol{\mathcal{G}_\Delta} \equiv \boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \widehat{\boldsymbol{C}}_{\boldsymbol{\Delta}}^{\boldsymbol{1}} \tag{A7}$$

denote the spending-to-income network out of Δ.

Our generalization of Proposition 2 decomposes the change in first-period aggregate final output due to a small change in fiscal policy along the direction Δ:

Proposition 5. *The change in first-period aggregate final output with respect to a small change in fiscal policy and preferences in direction Δ indexed by ϵ can be approximated as:*

$$\begin{aligned} \mathbb{1}^T \frac{dY^1}{d\epsilon} =& \mathbb{1}^T \Delta G^1 + \frac{1}{1 - \mathbb{E}_{h^*}[m_n]} \Bigg(\underbrace{\mathbb{E}_{h^*}[m_n]}_{\substack{RA \\ Keynesian\ effect}} + \underbrace{\mathbb{E}_\mu[m_{\Delta n}] - \mathbb{E}_{h^*}[m_n]}_{Incidence\ effect} \\ &+ \underbrace{\mathbb{E}_\mu[m_{\Delta n}] \left(\mathbb{E}_\mu[m_{\Delta n}^{next}] - \mathbb{E}_{h^*}[m_n] \right)}_{Biased\ spending\ direction\ effect} + \underbrace{\mathbb{C}ov_\mu[m_{\Delta n}, m_{\Delta n}^{next}]}_{Homophily\ effect} \Bigg) + O^3(|m|) \end{aligned} \tag{A8}$$

where h^* is any reference income weighting of unit-magnitude and $m_{\Delta n}^{next} = \left(m^T \boldsymbol{\mathcal{G}_\Delta} \right)_n$ is the average MPC-out-of-labor-income of households who receive as income i's marginal dollar of spending induced by Δ.

Proof. From the general version of Proposition 1 in Appendix A.1, and using that $\mathbb{1}^T \widehat{C}^1 = \mathbb{1}^T$ we have

$$\begin{aligned} \mathbb{1}^T \frac{dY^1}{d\epsilon} &= \mathbb{1}^T \left(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \right)^{-1} \left(\Delta G_1 + \boldsymbol{C}_{\boldsymbol{\tau}}^{\boldsymbol{1}} \Delta\tau + \boldsymbol{C}_{\boldsymbol{\theta}}^{\boldsymbol{1}} \Delta\theta \right) \\ &= \mathbb{1}^T \Delta G_1 + \mathbb{1}^T \left(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \right)^{-1} \underbrace{\left(\widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \Delta G^1 + \boldsymbol{C}_{\boldsymbol{\tau}}^{\boldsymbol{1}} \Delta\tau + \boldsymbol{C}_{\boldsymbol{\theta}}^{\boldsymbol{1}} \Delta\theta \right)}_{\equiv \widehat{C}_\Delta^1 \boldsymbol{m_\Delta} \mu} \\ &= \mathbb{1}^T \Delta G_1 + \mathbb{1}^T \widehat{C}_\Delta^1 \boldsymbol{m_\Delta} \mu + \underbrace{\mathbb{1}^T \left(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \right)^{-1} \widehat{\boldsymbol{C}}^{\boldsymbol{1}}}_{b^T \equiv \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{m} \boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \right)^{-1}} \boldsymbol{m} \underbrace{\boldsymbol{R}_{\boldsymbol{L}^1}^{\boldsymbol{1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}} \right)^{-1} \widehat{C}_\Delta^1}_{\boldsymbol{\mathcal{G}_\Delta}} \boldsymbol{m_\Delta} \mu \\ &= \mathbb{1}^T \Delta G_1 + \left(\mathbb{1}^T + b^T \boldsymbol{m} \boldsymbol{\mathcal{G}_\Delta} \right) \boldsymbol{m_\Delta} \mu \end{aligned} \tag{A9}$$

Next, letting $b_\Delta^T \equiv \mathbb{1}^T + b^T \boldsymbol{m}_{\boldsymbol{\Delta}}^{\boldsymbol{next}}$ for $\boldsymbol{m}_{\boldsymbol{\Delta}}^{\boldsymbol{next}} \equiv \boldsymbol{m}\boldsymbol{\mathcal{G}}_{\boldsymbol{\Delta}}$, we rewrite this as

$$\begin{aligned}
\mathbb{1}^T \frac{dY^1}{d\epsilon} &= \mathbb{1}^T \Delta G_1 + \mathbb{E}_\mu[m_{\Delta n} b_{\Delta n}] \\
&= \mathbb{1}^T \Delta G_1 + \mathbb{E}_\mu[m_{\Delta n}]\mathbb{E}_\mu[b_{\Delta n}] + \mathbb{C}\text{ov}_\mu[m_{\Delta n}, b_{\Delta n}] \\
&= \mathbb{1}^T \Delta G_1 + \mathbb{E}_\mu[m_{\Delta n}]\mathbb{E}_\mu[1 + b_n m_{\Delta n}] + \mathbb{C}\text{ov}_\mu[m_{\Delta n}, b_{\Delta n}] \\
&= \mathbb{1}^T \Delta G_1 + \mathbb{E}_\mu[m_{\Delta n}]\mathbb{E}_{h*}[1 + b_n m_n] + \mathbb{E}_\mu[m_{\Delta n}]\left(\mathbb{E}_\mu[b_n m_{\Delta n}] - \mathbb{E}_{h*}[b_n m_n]\right) + \mathbb{C}\text{ov}_\mu[m_{\Delta n}, b_n m_{\Delta n}^{next}]
\end{aligned} \tag{A10}$$

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

$$\begin{aligned}
\mathbb{1}^T \frac{dY^1}{d\epsilon} &= \mathbb{1}^T \Delta G_1 + \mathbb{E}_\mu[m_{\Delta n}] \frac{1}{1 - \mathbb{E}_{h*}[m_n]} \\
&\qquad + \mathbb{E}_\mu[m_{\Delta n}] \frac{\mathbb{E}_\mu[m_{\Delta n}] - \mathbb{E}_{h*}[m_n]}{1 - \mathbb{E}_{h*}[m_n]} + \frac{\mathbb{C}\text{ov}_\mu[m_{\Delta n}, b_n m_{\Delta n}^{next}]}{1 - \mathbb{E}_{h*}[m_n]} + O^3(|m|) \\
&= \mathbb{1}^T \Delta G_1 + \frac{1}{1 - \mathbb{E}_{h*}[m_n]} \Big(\mathbb{E}_{h*}[m_n] + \mathbb{E}_\mu[m_{\Delta n}] - \mathbb{E}_{h*}[m_n] \\
&\qquad + \mathbb{E}_\mu[m_{\Delta n}] \left(\mathbb{E}_\mu[m_{\Delta n}] - \mathbb{E}_{h*}[m_n]\right) + \mathbb{C}\text{ov}_\mu[m_{\Delta n}, b_n m_{\Delta n}^{next}] \Big) + O^3(|m|)
\end{aligned} \tag{A11}$$

□

To see how Proposition 2 follows from this result, observe that $\mu_n m_{\Delta n}$ collapses to $m_n \partial h_n^1$ when we restrict to changes in spending and transfers, and restrict consumption to be a function of net household first period income. It follows immediately that we can write Equation A8 in the form claimed in the main text:

$$\begin{aligned}
\mathbb{1}^T dY^1 = \mathbb{1}^T dG^1 + \frac{1}{1 - \mathbb{E}_{h*}[m_n]} \Bigg(& \underbrace{\mathbb{E}_{h*}[m_n]}_{\text{RA Keynesian effect}} + \underbrace{\mathbb{E}_{\partial h^1}[m_n] - \mathbb{E}_{h*}[m_n]}_{\text{Incidence effect}} \\
& + \underbrace{\mathbb{E}_{\partial h^1}[m_n]\left(\mathbb{E}_{\partial h^1}[m_n^{\text{next}}] - \mathbb{E}_{h*}[m_n]\right)}_{\text{Biased spending direction effect}} + \underbrace{\mathbb{C}\text{ov}_{\partial h^1}[m_n, m_n^{\text{next}}]}_{\text{Homophily effect}} \Bigg) + O^3(|m|)
\end{aligned} \tag{A12}$$

A.3. Proof of Proposition 3

Proof. We first provide a lemma characterizing the marginal change in welfare for a more general welfare function in which households have instrumental value of government purchases given by $w_n^t(G^t)$ and that considers changes in policy in both periods.

Lemma 1. *The change in welfare dW due to a small change in taxes and government*

purchases in the first period—at a constant interest rate—can be expressed as:

$$dW = \sum_{n\in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 - \left(d\tau_n^1 + (1-\phi_n)\frac{d\tau_n^2}{1+r} \right) + \left(WTP_n^1 dG^1 + (1-\phi_n)\frac{WTP_n^2}{1+r} dG^2 \right) \right] \tag{A13}$$

where $WTP_n^t = \frac{w_{nG}^t}{\kappa_n^t}$ *is household* n*'s marginal willingness to pay for government expenditure in period* t *and* ϕ_n *is the household's borrowing wedge, defined implicitly by:*

$$\kappa_n^1 = \beta_n \frac{1+r}{1-\phi_n} \kappa_n^2 \tag{A14}$$

Proof. To begin, we define κ_n^t to be n's marginal value of additional expenditure in period t, i.e. for all i, $u_{nc_i}^t = \kappa_n^t$ (recall prices are normalized to one). Therefore,

$$\begin{aligned} dW &= \sum_{n\in N} \lambda_n \mu_n \sum_{t=1,2} \beta_n^{t-1} \left(u_{nc}^t dc_n^t - v_n^{t\prime} d\ell_n^t + w_{nG}^t dG^t \right) \\ &= \sum_{n\in N} \lambda_n \mu_n \sum_{t=1,2} \beta_n^{t-1} \left[\kappa_n^t \left(\mathbb{1}^T dc_n^t - \frac{v_n^{t\prime}}{\kappa_n^t} d\ell_n^t \right) + w_{nG}^t dG^t \right] \end{aligned} \tag{A15}$$

Next note that in the second period, free labor supply implies $v_n^{2\prime} = \kappa_n^2$. In the first, there may be some wedge Δ_n such that $v_n^{1\prime} = \kappa_n^1(1+\Delta_n)$; 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. In these terms, we have

$$dW = \sum_{n\in N} \lambda_n \kappa_n^1 \mu_n \left[-\Delta_n d\ell_n^1 + \sum_{t=1,2} \frac{\kappa_n^t}{\kappa_n^1} \beta_n^{t-1} \left(\mathbb{1}^T dc_n^t - d\ell_n^t \right) + \left(\frac{w_{nG}^1}{\kappa_n^1} dG^1 + \frac{\beta_n w_{nG}^2}{\kappa_n^1} dG^2 \right) \right] \tag{A16}$$

Next, define $\widetilde{\lambda}_n = \lambda_n \kappa_n^1$. Also note that $\frac{\kappa_n^t}{\kappa_n^1}\beta_n^{t-1} = 1$ for $t = 1$. For $t = 2$, we use the modified Euler equation:

$$\kappa_n^1 = \beta_n \frac{1+r}{1-\phi_n} \kappa_n^2 \tag{A17}$$

where ϕ_n is a borrowing wedge. $\phi_n \geqslant 0$ is positive when households behave as if interest rates are higher than in reality, i.e. consume more in the future than they would like; this corresponds to borrowing constraints. This gives us

$$dW = \sum_{n\in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 + \left(\mathbb{1}^T dc_n^1 - d\ell_n^1 \right) + \frac{1-\phi_n}{1+r} \left(\mathbb{1}^T dc_n^2 - d\ell_n^2 \right) + \left(\frac{w_{nG}^1}{\kappa_n^1} dG^1 + \left(\frac{1-\phi_n}{1+r} \right) \frac{w_{nG}^2}{\kappa_n^2} dG^2 \right) \right] \tag{A18}$$

Differentiating the household's lifetime budget constraint (at constant r):

$$\mathbb{1}^T dc_n^1 - d\ell_n^1 + \frac{\mathbb{1}^T dc_n^2 - d\ell_n^2}{1+r} = -d\tau_n^1 - \frac{d\tau_n^2}{1+r} \tag{A19}$$

Plugging this in, we have:

$$dW = \sum_{n \in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 + \phi_n \left(\mathbb{1}^T dc_n^1 - d\ell_n^1 \right) - (1-\phi_n) \left(d\tau_n^1 + \frac{d\tau_n^2}{1+r} \right) + \left(\frac{w_{nG}^1}{\kappa_n^1} dG^1 + \left(\frac{1-\phi_n}{1+r} \right) \frac{w_{nG}^2}{\kappa_n^2} dG^2 \right) \right] \tag{A20}$$

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

$$\underline{s}_n = l_n^1 - \tau_n^1 - \mathbb{1}^T c_n^1 \implies \mathbb{1}^T dc_n^1 - d\ell_n^1 = -d\tau_n^1 \tag{A21}$$

Defining the within-period willingness to pay for government purchases as $WTP_n^t = \frac{w_{nG}^t}{\kappa_n^t}$, we arrive at the final expression:

$$dW = \sum_{n \in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n d\ell_n^1 - \left(d\tau_n^1 + (1-\phi_n) \frac{d\tau_n^2}{1+r} \right) + \left(WTP_n^1 dG^1 + (1-\phi_n) \frac{WTP_n^2}{1+r} dG^2 \right) \right] \tag{A22}$$

Completing the proof of the lemma. □

The claimed result then follows immediately by zeroing all second period policy variations and willingness to pay for government spending.

We now prove the additional welfare formulae claimed in the text (Equations 23 and 24). Equation 24 follows immediately from the results above. To derive equation 23, recall that we assumed: the modified welfare weights are uniform across households ($\widetilde{\lambda}_n = 1$ for all $n \in N$), there is no willingness to pay for government purchases ($WTP_n^1 = 0$ for all $n \in N$), and all underemployed households have zero marginal disutility of labor (if $\Delta_n < 0$ then $\Delta_n = -1$). Thus, Equation 22 reduces to:

$$dW = \mu^T d\ell^1 - \mu^T d\tau^1 \tag{A23}$$

Moreover, by the formula for the multiplier, we have that:

$$\boldsymbol{\mu} d\ell^1 = \boldsymbol{\Gamma^1} \left(\boldsymbol{I} - \boldsymbol{C_{\ell^1}^1 \Gamma^1} \right)^{-1} \left(dG^1 - \boldsymbol{C_{\ell^1}^1 \mu} d\tau^1 \right) \tag{A24}$$

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

$$\begin{aligned} dW &= \mathbb{1}^T \mathbf{\Gamma^1} \left(\boldsymbol{I} - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \mathbf{\Gamma^1}\right)^{-1} \left(dG^1 - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \boldsymbol{\mu} d\tau^1\right) - \mu^T d\tau^1 \\ &= \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \mathbf{\Gamma^1}\right)^{-1} dG^1 - \mathbb{1}^T \left[\left(\boldsymbol{I} - \mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\right)^{-1} \mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} + I\right] \boldsymbol{\mu} d\tau^1 \\ &= \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \mathbf{\Gamma^1}\right)^{-1} dG^1 - \mathbb{1}^T \left(\boldsymbol{I} - \mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\right)^{-1} \boldsymbol{\mu} d\tau^1 \end{aligned} \tag{A25}$$

We now use the final assumption made that the bias and homophily effects are zero for all possible policies. To this end, we first show that, if the bias and homophily effects are zero for all purchases and transfers shocks relative to some baseline income incidence h^*, then either $m_n = 0$ or $m_n^{\text{next}} = \mathbb{E}_{h^*}[m_{n'}]$.

To start, fixing a single type $n \in N$, consider the bias term corresponding to a transfers shock with direct incidence $\partial h^1 = \hat{e}_n$ (i.e. only transfering to n).

$$\text{bias}^{h^*}_{\partial h^1} = \mathbb{E}_{\partial h^1}[m_n] \left(\mathbb{E}_{\partial h^1}[m_n^{\text{next}}] - \mathbb{E}_{h^*}[m_{n'}]\right) = m_n \left(m_n^{\text{next}} - \mathbb{E}_{h^*}[m_{n'}]\right) \tag{A26}$$

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

Applying this fact to the multipliers in Equation A25, we have

$$\begin{aligned} \mathbb{1}^T \left(\boldsymbol{I} - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \mathbf{\Gamma^1}\right)^{-1} &= \sum_{k=0}^{\infty} \mathbb{1}^T \left(\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \mathbf{\Gamma^1}\right)^k = \mathbb{1}^T + \mathbb{1}^T \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \mathbf{\Gamma^1} + \sum_{k=1}^{\infty} \mathbb{1}^T \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \left(\mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\right)^k \mathbf{\Gamma^1} \\ &= \mathbb{1}^T + m^T \mathbf{\Gamma^1} + \sum_{k=1}^{\infty} \mathbb{E}_{h^*}[m_n]^k m^T \mathbf{\Gamma^1} = \left(\mathbb{1} + \frac{1}{1 - \mathbb{E}_{h^*}[m_n]} m\right)^T \mathbf{\Gamma^1} \end{aligned} \tag{A27}$$

Similarly, we have that:

$$\begin{aligned} \mathbb{1}^T \left(\boldsymbol{I} - \mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\right)^{-1} &= \mathbb{1}^T + \mathbb{1}^T \mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} + \sum_{k=1}^{\infty} \mathbb{1}^T \mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \left(\mathbf{\Gamma^1} \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\right)^k \\ &= \mathbb{1}^T + \frac{1}{1 - \mathbb{E}_{h^*}[m_n]} m^T = \left(\mathbb{1} + \frac{1}{1 - \mathbb{E}_{h^*}[m_n]} m\right)^T \end{aligned} \tag{A28}$$

Substituting (A27) and (A28) into Equation A25 shows that:

$$dW = \left(\mathbb{1} + \frac{1}{1 - \mathbb{E}_{h^*}[m_n]} m\right)^T \left(\mathbf{\Gamma^1} dG^1 - \boldsymbol{\mu} d\tau^1\right) \tag{A29}$$

By budget balance, we moreover have that:

$$\mathbb{1}^T dG^1 - \mathbb{1}^T \boldsymbol{\mu} d\tau^1 = 0 \tag{A30}$$

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

$$\mathbb{1}^T \mathbf{\Gamma^1} dG^1 - \mathbb{1}^T \boldsymbol{\mu} d\tau^1 = 0 \tag{A31}$$

Thus, the change in welfare is simply given by:

$$dW = \frac{1}{1 - \mathbb{E}_{h*}[m_n]} m^T \left(\mathbf{\Gamma^1} dG^1 - \boldsymbol{\mu} d\tau^1 \right) \tag{A32}$$

yielding the claim given in the text. □

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 / wage expectations 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 $\tau = (\tau_n^1, \tau_n^2)$.

$$\begin{aligned} \ell_n^{*1} \in \arg\max_{\ell^1} \max_{\ell^2, c^1, c^2} & \sum_{t=1,2} \beta_n^{t-1} u_n^t(c^t, \ell^t) \\ & \text{s.t } p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell^1 + \frac{\ell^2}{1+r} \\ & \ell^1 - p^1 c^1 - \tau_n^1 \geqslant \underline{s}_n \end{aligned} \tag{A33}$$

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).

$$L_i^{*1} \in \arg\max_{L^1} \max_{X^1} \; p_i^1 F_i^1(X^1, L^1) - p^1 X^1 - L^1 \tag{A34}$$

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.

$$\ell_n^1 = R_n^S\left(\ell^{*1}, L^{*1}\right) \qquad\qquad L_i^1 = R_i^D\left(\ell^{*1}, L^{*1}\right) \tag{A35}$$

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 $c_n^t = \{c_{ni}^t\}_{i\in I}$.

$$\begin{aligned}(\ell_n^1, \ell_n^2, c_n^1, c_n^2) \in \underset{\ell^1,\ell^2,c^1,c^2}{\arg\max} & \sum_{t=1,2} \beta_n^{t-1} u_n^t(c^t, \ell^t) \\ \text{s.t } & p^1 c^1 + \frac{p^2 c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell^1 + \frac{\ell^2}{1+r} \\ & \ell^1 - p^1 c^1 - \tau_n^1 \geqslant \underline{s}_n \\ & \ell^1 = \ell_n^1\end{aligned} \tag{A36}$$

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 $F_i^{t=2}$.

$$\begin{aligned}(X_i^t, L_i^t) \in \underset{X^t, L^t}{\arg\max}\ & p_i^t F_i^t(X^t, L^t) - p^t X^t - L^t \\ \text{s.t } & L^1 = L_i^1\end{aligned} \tag{A37}$$

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

$$\sum_{n\in N} \mu_n \left(\tau_n^1 + \frac{1}{1+r}\tau_n^2\right) = p^1 G^1 + \frac{1}{1+r} p^2 G^2 \tag{A38}$$

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

$$F_i^t(X_i^t, L_i^t) = \sum_{j\in I} X_{ji}^t + \sum_{n\in N} \mu_n c_{ni}^t + G_i^t, \qquad\qquad \sum_{i\in I} L_i^2 = \sum_{n\in N} \mu_n l_n^2 \tag{A39}$$

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

Definition 2. *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 (A33) – (A39).*

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

$$R_n^S\left(\ell^{*1}, L^{*1}\right) \leqslant \ell_n^{*1} \qquad \text{and} \qquad R_i^D\left(\ell^{*1}, L^{*1}\right) \leqslant L_i^{*1}. \tag{A40}$$

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

$$R_n^S\left(\ell^{*1}, L^{*1}\right) < \ell_n^{*1} \qquad \text{and} \qquad R_i^D\left(\ell^{*1}, L^{*1}\right) < L_i^{*1}. \tag{A41}$$

Free disposal captures the idea that firms and households may always work any amount less than they are allocated by the rationing function, by simply hiring fewer workers or by choosing to work fewer jobs or hours, respectively.[51] 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. *Parameters are such that there is strictly deficient demand in any equilibrium* $\sum_{i\in I} L_i^{1*} < \sum_{n\in N} \ell_n^{1*}$.

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:

[51] Of course in practice, some workers are not able to simply accept a fraction of a job. Still, we view this as a reasonable assumption in the context of deficient labor demand, particularly for lower-income workers, who are the most exposed to underemployment.

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 v_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,*

$$u_n^t(c,\ell) = U_n^t\left(v_n^t(c) - \Lambda_n^t(\ell)\right). \tag{A42}$$

That our general rationing model reduces under these three assumptions to the one studied in the main text is shown by the following result:

Proposition 6. *Under assumptions 1, 2, and 3, we have that $\mu_n \ell_n^1 = R_n^1(L^{*1}) \equiv R_n^S(\ell^{*1}, L^{*1})$ and $L_i^1 = L_i^{*1}$.*[52] *Thus, the set of rationing equilibria and general rationing equilibria are equal.*

Proof. 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 ℓ_n^{1*}, then $L_i^1 = L_i^{1*}$. To see this, first note that the free disposal property and deficient demand together imply there exists some n such that $\ell_n^1 < \ell_n^*$. But then if there is some i such that $L_i^1 < L_i^{1*}$, allocative efficiency cannot hold, a contradiction.

To complete the proof, it remains only to show that $R_n^S(\ell^{1*}, L^1) = R_n(L^1)$. Observe by strict monotonicity and homotheticity of v_n^t and the fact that p^t is fixed by Proposition 7, that each household consumes a fixed basket of goods α_n^t with non-negative weights on each good. Moreover, we can write $v_n^t(c^t) = v_n^t(\{\alpha_{ni}^1 \tilde{c}_n\}_{i\in\mathcal{I}}) = \omega_n^t \tilde{c}_n^t$ 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:

$$p^1 \alpha_n^1 \tilde{c}^1 + \frac{p^2 \alpha_n^2 \tilde{c}^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} = \ell^1 + \frac{\ell^2}{1+r} \tag{A43}$$

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

$$\ell_n^{t*} = \left(\Lambda_n^{t\,\prime}\right)^{-1}\left(\omega_n^t\right) \tag{A44}$$

if $\left(\Lambda_n^{t\,\prime}\right)^{-1}\left(\omega_n^t\right)$ lies in the positive reals and zero otherwise. Thus, ℓ_n^{t*} is constant. Hence, we may write $R_n^S(\ell^{1*}, L^1) = R_n(L^1)$, as claimed. □

[52] Implicit in the statement that $R_n^1(L^{*1}) \equiv R_n^S(\ell^{*1}, L^{*1})$ is the claim that $R_n^S(\ell^{*1}, L^{*1})$ does not depend on ℓ^{*1}.

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. 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

Here we present a provide results on properties (including existence) of rationing equilibrium (C.1), provide comparative statics for the multiplier (C.2), analyze benchmark cases in which various network adjustments to the Keynesian multiplier are zero (C.3), provide first order conditions for policy changes (C.4), and analyze policy at the optimum (C.5).

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 demand – 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, t and z_i^t, production $F_i^t(X_i^t, L_i^t, z_i^t)$ is continuous, weakly increasing, strictly quasi-concave, and homogeneous of degree one in (X_i^t, L_i^t). Further, labor is essential in production, i.e. $F_i^t(X_i^t, 0, z_i^t) = 0$, and production is strictly increasing in labor. Finally, for all t and z^t, there exists some $\overline{p}^t(z^t) \in \mathbb{R}^I_{>0}$ 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 $\overline{p}^t \overline{X}_i^t + \overline{L}_i^t \leqslant \overline{p}_i^t$.*[53]

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 7. *Under Assumptions 4 and 5, for each $t \in \{1, 2\}$ and for any given z^t, there exists a unique $p^t \in \mathbb{R}^I_{\geqslant 0}$ consistent with rationing equilibrium, independent of demand. Also, all components of p^t are strictly positive.*

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

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

Fix a time period $t \in \{1, 2\}$, a vector of productivity parameters z^t, 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 firm prices to i's least cost of production:

$$\kappa_i(p) \equiv \min_{X_i^t, L_i^t \geqslant 0,\ F_i^t(X_i^t, L_i^t, z^t) \geqslant 1} pX_i^t + L_i^t. \tag{A45}$$

This function is well-defined since—because i can produce a unit of output with the input bundle $\overline{X}_i^t, \overline{L}_i^t$—we may WLOG restrict the domain of (A45) to the set

$$\left\{ X_i^t \in \mathbb{R}^I_{\geqslant 0}, L_i^t \in \mathbb{R}_{\geqslant 0} \;\middle|\; \forall j \in I, X_{ij}^t \leqslant \frac{\overline{\kappa}_i(p)}{p_j},\ L_i^t \leqslant \overline{\kappa}_i(p),\ F_i^t(X_i^t, L_i^t, z^t) \geqslant 1 \right\} \tag{A46}$$

where $\overline{\kappa}_i(p) \equiv p\overline{X}_i^t + \overline{L}_i^t$ This set is compact by $F_i^t(\cdot, \cdot, z^t)$'s continuity, from Assumption 4, so since the objective of (A45) is continuous, a minimum exists.

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 therefore imply that each firm's price is equal to its cost; otherwise it could profit by scaling up or down. 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 $p \in \mathbb{R}^I_{\geqslant 0}$ must be strictly positive and satisfy:

$$\forall i \in I, \quad p_i = \kappa_i(p). \tag{A47}$$

The remainder of the proof shows that such a price vector exists and is unique.

Existence

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

$$\forall i \in I, \quad \kappa_i(\overline{p}) \leqslant \overline{p}_i. \tag{A48}$$

Moreover, we claim that for small enough $\alpha > 0$

$$\forall i \in I, \quad \kappa_i(\alpha\overline{p}) \geqslant \alpha\overline{p}_i. \tag{A49}$$

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

$$F_i^t(X_i^k, L_i^k, z_i^t) \geqslant 1 \qquad \text{and} \qquad \alpha\overline{p}X_i^k + L_i^k < \alpha\overline{p}_i. \tag{A50}$$

Since by assumption all components of $\overline{p}$ are strictly positive, this implies $L_i^k \to 0$ and each

$X_i^k \leqslant \frac{\overline{p}_i}{\overline{p}}$. But then $F_i^t(\cdot,\cdot,z_i^t)$'s continuity (from Assumption 4) and the Bolzano-Weierstrass Theorem imply that $F_i^t(X_i^\infty, 0, z_i^t) \geqslant 1$ at some (finite) 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 (A45)—each cost function κ_i is monotonically increasing, the observations above imply that the function $\kappa = (\kappa_1, ..., \kappa_{|I|})$ maps the set

$$\mathcal{P} \equiv \prod_{i\in I}[\alpha\overline{p}_i, \overline{p}_i] \tag{A51}$$

to itself for some $\alpha > 0$, namely one sufficiently small that the previous argument carries. Since $\mathcal{P}$ 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 $\mathcal{P}$ is a complete lattice. In particular, it is non-empty. This implies the existence of a strictly positive solution to the fixed-point problem (A47).

Uniqueness

Suppose that p^t and $p^{t\prime}$ are two strictly positive solutions (A47). 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 i^{th} entry L_i^t. The price-cost fixed point equation (A47) at p^t can therefore be written:

$$p^t = \widehat{\boldsymbol{X}}^{\boldsymbol{t}} p^t + \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1} \tag{A52}$$

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

$$p^t = \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{t}}\right)^{-1} \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1}. \tag{A53}$$

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

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

$$p^{t\prime} \leqslant \widehat{\boldsymbol{X}}^{\boldsymbol{t}} p^{t\prime} + \widehat{\boldsymbol{L}}^{\boldsymbol{t}} \mathbb{1}, \qquad \text{which implies} \qquad p^{t\prime} \leqslant \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{t}}\right)^{-1} \widehat{\boldsymbol{L}}^{\boldsymbol{t}} = p^t. \tag{A54}$$

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

□

The existence of unique, positive prices $p^1(z^1), p^2(z^2) \in \mathbb{R}^I_>$ consistent with equilibrium

allows us to reduce the number of endogenous price variables in considering comparative statics that keep z^1 and z^2 fixed. Implicit in this non-substitution economy is the assumption that good prices respond instantaneously to changes in technology, which is irrelevant in the case of demand shocks.

Moreover, combining Proposition 7 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

$$\left(\widehat{X}_i^t(z^t), \widehat{L}_i^t(z^t)\right) \equiv \arg \min_{(X_i^t, L_i^t) \text{ s.t. } F_i^t(X_i^t, L_i^t, z_i^t) \geqslant 1} p^t(z^t) X_i^t + L_i^t \tag{A55}$$

where by Proposition 7, $p^t(z^t)$ 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 7 and (b) production is strictly quasi-concave by Assumption 4. We now claim:

Corollary 1. *In any rationing equilibrium, aggregate demands X^t and the vector of firm-specific labor demands L^t are given by:*

$$X^t = \widehat{\boldsymbol{X}}^{\boldsymbol{t}}(z^t) Q^t \quad L^t = \widehat{\boldsymbol{L}}^{\boldsymbol{t}}(z^t) Q^t \tag{A56}$$

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

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

$$X_i^t = Q_i^t \widehat{X}_i^t(z^t) \quad L_i^t = Q_i^t \widehat{L}_i^t(z^t) \tag{A57}$$

Stacking these equations over $\mathcal{I}^t$ gives

$$X^t = \widehat{\boldsymbol{X}}^{\boldsymbol{t}}(z^t) Q^t \quad L^t = \widehat{\boldsymbol{L}}^{\boldsymbol{t}}(z^t) Q^t \tag{A58}$$

□

Proposition 7 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 demand-independent prices. This is stated formally in the following corollary:

Corollary 2. *For any t, z^t the Leontief-inverse matrix $(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{t}}(z^t))^{-1}$ exists. Moreover,*

prices are given uniquely by the following expression:

$$p^t(z^t) = \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{t}}(z^t)^T\right)^{-1} \hat{\boldsymbol{L}}^{\boldsymbol{t}}(z^t)\mathbb{1} \tag{A59}$$

Proof. This follows from the argument in the "Uniqueness" step of the proof of Proposition 7 that any unit-inputs matrix corresponding to non-substitution prices is invertible. □

Throughout the paper we will write $\widehat{\boldsymbol{X}}^{\boldsymbol{t}}, \hat{\boldsymbol{L}}^{\boldsymbol{t}}$ for $\widehat{\boldsymbol{X}}^{\boldsymbol{t}}(z^t), \hat{\boldsymbol{L}}^{\boldsymbol{t}}(z^t)$ when z^t is fixed. We write $\widehat{\boldsymbol{X}}$ and $\hat{\boldsymbol{L}}$ for the block-diagonal matrices composed of $\widehat{\boldsymbol{X}}^{\boldsymbol{1}}$ and $\widehat{\boldsymbol{X}}^{\boldsymbol{2}}$, and $\hat{\boldsymbol{L}}^{\boldsymbol{1}}$ and $\hat{\boldsymbol{L}}^{\boldsymbol{2}}$ respectively.

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 $c_n^t = c_n^t(p, \ell_n^1, \tau_n, \theta_n)$ and $\ell_n^2 = \ell_n^2(p, \ell_n^1, \tau_n, \theta_n)$.

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, θ: 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 $p^1 c_n^1(p, \ell_n^1, \tau_n, \theta_n) \leqslant \overline{c}\ell_n^1$.*

This assumption is extremely mild and satisfied by virtually all standard household problems of which we are aware.[54] With this additional structure we are now able to prove the existence of rationing equilibria for the economy under consideration.

Proposition 8. *Under Assumptions 4, 5, and 6, there exists a rationing equilibrium.*

Proof. Fix all exogenous parameters. Note that by Proposition 7, prices p^1 and p^2 are pinned down by technology and so can be taken as given as well.

First, by Assumption 6 we have the following fact: For any p^1, p^2, τ, θ: there exists some $\overline{y} \in \mathbb{R}_+$ and some $\overline{c} < 1$ such that $p^1 c_n^1(p, \ell_n^1, \tau_n, \theta_n) \leqslant \overline{c}\ell_n^1$ for all n, $\ell_n^1 > \overline{y}$.

[54] It is easy to see how Assumption 6 holds if households are utility maximizers whose utility functions satisfy various standard assumptions. Existence and continuity of the consumption and labor functions follow from continuity and quasiconcavity of utility, and from Berge's theorem. Satisfying the lifetime budget constraint follows from non-satiation. Consumption being asymptotically bounded away from first-period income follows from sufficiently decreasing marginal utility.

Thus, for any vector of incomes ℓ^1, first period aggregate spending $(p^1)^T C^1 + (p^1)^T G^1$ is bounded above:

$$(p^1)^T C^1 + (p^1)^T G^1 \leqslant \overline{cy} + \overline{c}\mu^T \ell^1 + (p^1)^T G^1 \tag{A60}$$

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:

$$\forall \ell^1 \in \mathcal{Y}^1 : R^1 \left(\widehat{\boldsymbol{L}^1} (1 - \widehat{\boldsymbol{X}^1})^{-1} \left(C^1(p, \ell^1, \tau, \theta) + G^1 \right) \right) \in \mathcal{Y}^1 \tag{A61}$$

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

$$\Psi(\ell^1) = R^1 \left(\widehat{\boldsymbol{L}^1} (1 - \widehat{\boldsymbol{X}^1})^{-1} \left(C^1(p, \ell^1, \tau, \theta) + G^1 \right) \right) \tag{A62}$$

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 implied by z^t. Let c^t_n and ℓ^2_n by the relevant functions taking in prices p^t and incomes ℓ^1. Let production in each period be:

$$Q^t = (\boldsymbol{I} - \widehat{\boldsymbol{X}^t})^{-1} (G^t + C^t) \tag{A63}$$

The definition 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. □

As we have established 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:

$$\mathcal{M} = \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}} \boldsymbol{l}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \widehat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}^{\mathbf{1}}} \right)^{-\mathbf{1}} \tag{A64}$$

and assume that all entries of this matrix are non-negative. $\mathcal{M}$ 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 9. *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:*

$$\frac{d}{d\varepsilon} dY^1|_{\varepsilon=0} = (\boldsymbol{I} - \mathcal{M})^{-1} \mathcal{E} (\boldsymbol{I} - \mathcal{M})^{-1} \partial Q^1 \tag{A65}$$

where recall ∂Q^1 generalizes ∂Y^1 to the case with supply shocks (see the beginning of Appendix A).

Proof. We start from the multiplier derived in Proposition 1:

$$dY^1 = (\boldsymbol{I} - \mathcal{M})^{-1} \partial Q^1 \tag{A66}$$

We then have that:[55]

$$\begin{aligned} \frac{d}{d\varepsilon} dY^1|_{\varepsilon=0} &= \frac{d}{d\varepsilon} (\boldsymbol{I} - \mathcal{M})^{-1}|_{\varepsilon=0} \partial Q^1 \\ &= (\boldsymbol{I} - \mathcal{M})^{-1} \mathcal{E} (\boldsymbol{I} - \mathcal{M})^{-1} \partial Q^1 \end{aligned} \tag{A67}$$

□

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 $\partial Q^1 \geqslant 0$, dY^1 increases in all dimensions.*

Proof. See in both cases that $\frac{d}{d\varepsilon} dY^1|_{\varepsilon=0} = (\boldsymbol{I} - \mathcal{M})^{-1} \mathcal{E} (\boldsymbol{I} - \mathcal{M})^{-1} \partial Q^1$, 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

[55] 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 $\partial \boldsymbol{A} = -\mathcal{E}$.

comparative statics of total final output in the empirically relevant case where there exists some reference incidence h^* around which the bias and homophily effects are zero for all possible ∂h^1. In this case, the proof of Proposition 3 shows that the total effect of a purchases shock on output is

$$\mathbb{1}^T dY^1 \propto \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{-1}\partial G^1 \tag{A68}$$

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 10. *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* $\widehat{\boldsymbol{X}}^{\mathbf{1}} = 0$ *and the other has some arbitrary input output matrix* $\widehat{\boldsymbol{X}}^{\mathbf{1}\prime}$*. 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:

$$\max_{i\in\mathcal{I}} \mathbb{1}^T \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} e_i = \max_{i\in\mathcal{I}} (\mathbb{1}^T \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}})_i \tag{A69}$$

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:

$$\max_{i\in\mathcal{I}} \mathbb{1}^T \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{-1} e_i = \max_{i\in\mathcal{I}} \left(\mathbb{1}^T \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\right)\widehat{\boldsymbol{L}}^{\mathbf{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{-1}_{.,i} \tag{A70}$$

As production is CRS, recall by Corollary 2 that the columns of $\widehat{\boldsymbol{L}}^{\mathbf{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{-1}$ sum to one. Thus, the maximum multiplier is some weighted average of the elements of $\mathbb{1}^T \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}$. The value of which is necessarily bounded above by the maximum element of $\mathbb{1}^T \boldsymbol{m}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}$, which is the maximum purchases multiplier in the first economy. The proof the 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/or Homophily Effects in Propagation Vanish

In the main text, we briefly discussed two important cases where some or all of the incidence, bias, and homophily effects in shock propagation vanish. Here, we state and more formally discuss these results.

Proposition 11. *The following statements are true:*

1. *(No incidence or bias effects) Suppose that consumption preferences and labor rationing are homothetic, that no households are net borrowers in period* 1*, and that there are no government purchases.*[56] *Then, for either (a) unit-magnitude final-output-proportional, first-period government purchase or (b) a unit-magnitude, labor-supply-proportional, first-period transfer the incidence and bias effects are zero, so that we have:*

$$\mathbb{1}^T dY^1 = \mathbb{1}^T dG^1 + \frac{1}{1 - \mathbb{E}_{\boldsymbol{\mu}\ell^1}[m_n]} \left(\mathbb{E}_{\boldsymbol{\mu}\ell^1}[m_n] \; + \; \underbrace{\mathbb{C}ov_{\boldsymbol{\mu}\ell^1}[m_n, m_n^{next}]}_{\textit{Homophily effect}} \right) + \; O^3(|m|) \tag{A71}$$

   where $\boldsymbol{\mu}$ is the diagonal matrix of type weights and $\boldsymbol{\mu}\ell^1$ is the vector of total first-period labor supply by each demographic group.
2. *(No incidence, bias, or homophily effects) Suppose that all industries have a common rationing-weighted average MPC, m.*[57] *Then the incidence, bias, and homophily effects are zero, the change in GDP corresponding to any change in government purchases is:*

$$\mathbb{1}^T dY^1 = \frac{1}{1-m} \mathbb{1}^T dG^1 \tag{A72}$$

Proof. We prove the two claims separately:

1. Since $\partial h^1 \propto \ell^1$, there is no incidence effect. It remains to show there is no bias effect, and so suffices to show that $\mathbb{E}_{\boldsymbol{\mu}\ell_1}[m_n^{\text{next}} - m_n] = m^T \boldsymbol{\mathcal{G}} \boldsymbol{\mu}\ell^1 - m^T \boldsymbol{\mu}\ell^1 = 0$. So it suffices to show $\boldsymbol{\mathcal{G}}\boldsymbol{\mu}\ell^1 = \boldsymbol{\mu}\ell^1$.

   Plugging in the definition of $\boldsymbol{\mathcal{G}}$, we have $\boldsymbol{\mathcal{G}}\boldsymbol{\mu}\ell^1 = \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L^1}} \widehat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} \widehat{\boldsymbol{C}}^{\mathbf{1}} \boldsymbol{\mu}\ell^1$. Since each household saves zero on net, ℓ^1 is equal to total spending. Homotheticity of

[56] By homothetic labor rationing, we mean that marginal and average rationing of income are equal, i.e. $\boldsymbol{\mu}\ell^1 = \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L^1}} L^1$.

[57] Formally, $\sum_{n \in N} (\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L^1}})_{ni} m_n = m$ for all i.

consumption implies that $\widehat{\boldsymbol{C}}^{\mathbf{1}}\boldsymbol{\mu}\ell^1$, then, is the vector of total consumption of goods; since there are no government purchases, this equals aggregate output net of inputs, i.e. Y^1. Finally, homotheticity of rationing implies that $\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}}\right)^{-1}Y^1=\boldsymbol{\mu}\ell^1$.

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

$$\begin{aligned}\mathbb{1}^T dY^1 &= \mathbb{1}^T\left(\boldsymbol{I}-\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1}\right)^{-1}dG^1=\sum_{k=0}^{\infty}\mathbb{1}^T\left(\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1}\right)^{k}dG^1\\ &=\sum_{k=0}^{\infty}m^k\mathbb{1}^T dG^1=\frac{1}{1-m}\mathbb{1}^T dG^1\end{aligned} \tag{A73}$$

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

$$\mathbb{1}^T\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}=m\mathbb{1}^T\quad\Longrightarrow\quad\mathbb{1}^T\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}}\widehat{\boldsymbol{L}}^{\mathbf{1}}(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{\mathbf{1}})^{-1}=m\mathbb{1}^T \tag{A74}$$

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

$$p^1=(\boldsymbol{I}-(\widehat{\boldsymbol{X}}^{\mathbf{1}})^T)^{-1}\widehat{\boldsymbol{L}}^{\mathbf{1}}\mathbb{1}, \tag{A75}$$

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

□

The first part of the proposition shows how, even in a "homothetic economy," heterogeneity in household consumption baskets and sectoral employment can generate additional amplification through homophily. This happens even at the same time as – for a GDP-proportional shock – homotheticity eliminates the incidence and bias effects by ensuring that shock incidence is proportional to household income and that each household's marginal consumption is proportional to its initial consumption, so that the income-weighted average of marginal consumption is proportional to the GDP vector. Still, when households with different MPCs direct their spending toward different goods, the households employed to produce the goods consumed by higher-MPC households experience a greater change in income – not from the initial, uniform shock, but from the economy's response to it. Insofar as these households have different MPCs from the average, homophily is still possible. This mechanism generates non-neutrality for the multiplier, even if the economy and the shock considered are "neutral" in all other aspects. Aggregate neutrality requires (to third order in MPCs) that the economy feature exactly zero correlation between households' MPCs and

the MPCs of the households they spend on.

The second part of 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 MPC of the population; this is the case only when each firm's marginal employees have the population average MPC.

C.4. Optimal Policy at a Global Optimum

In the main text, we focused primarily on small changes in welfare corresponding to small changes in policy and without intrinsic value of government purchases. In this section, we specialize to the case of small changes in policy *at an optimum* where households may also value government purchases. To do so, we consider the following planner's problem: The full version of the planner's problem, Equation 21, is

$$\begin{aligned}\max_{\{c^t_{ni},\ell^t_n,Q^t_i,G^t_i,\tau^t_n\}_{t\in\{1,2\},n\in N,i\in I}} \quad & W \equiv \sum_{n\in N}\mu_n\lambda_n\sum_{t=1,2}\beta_n^{t-1}\left[u^t_n(\tilde{c}^1)-v^t_n(\tilde{\ell}^t)+w^t_n(G^t)\right] \\ \text{s.t. } & (c^1_n,c^2_n,\ell^2_n) \text{ solves Equation 3' given } \ell^1_n \\ & Q^t = \mu^T c^t + \widehat{\boldsymbol{X}}^{\boldsymbol{t}} Q^t + G^t \\ & \boldsymbol{\mu}\ell^1 = R^1(\widehat{\boldsymbol{L}}^{\mathbf{1}}Q^1),\ \mu^T\ell^2 = \mathbb{1}^T\widehat{\boldsymbol{L}}^{\mathbf{2}}Q^2 \\ & \mathbb{1}^T G^1 + \frac{\mathbb{1}^T G^2}{1+r} = \mu^T\tau^1 + \frac{\mu^T\tau^2}{1+r}\end{aligned} \tag{A76}$$

Our first result decomposes the first-order condition for optimal government purchases and transfers into five distinct mechanisms. This is closely related to Proposition 3 in the main text, which considers the change in welfare away from the global optimum.

Proposition 12. *Suppose taxes τ^{1*},τ^{2*} and purchases G^{1*},G^{2*} solve the planner's problem. Now consider a change in policy $\tau^t = \tau^{t*}+\varepsilon\tau^t_\varepsilon, G^t = G^{t*}+\varepsilon G^t_\varepsilon$, indexed by ε. The following*

first-order condition holds:

$$
\begin{aligned}
0 = & \underbrace{\left(\widetilde{\lambda}^T \boldsymbol{\mu}\boldsymbol{WTP^1} - (\gamma \mathbb{1}^T + \widetilde{\lambda}^T \boldsymbol{\Delta}\boldsymbol{\Gamma^1})\right) G_\varepsilon^1}_{\textit{Opportunistic government purchases}} \quad \underbrace{+ \frac{\left(\widetilde{\lambda}^T \boldsymbol{\mu}(\boldsymbol{I}-\boldsymbol{\phi})\boldsymbol{WTP^2} - \gamma \mathbb{1}^T\right) G_\varepsilon^2}{1+r}}_{\textit{Short-termist government purchases}} \\
& \underbrace{- (\widetilde{\lambda} - \gamma \mathbb{1})^T \boldsymbol{\mu} \left(\tau_\varepsilon^1 + \frac{\tau_\varepsilon^2}{1+r}\right)}_{\textit{Pure redistribution}} \quad \underbrace{+ \widetilde{\lambda}^T \frac{\boldsymbol{\phi}\boldsymbol{\mu}\tau_\varepsilon^2}{1+r}}_{\textit{Relaxation of borrowing constraints}} \\
& \underbrace{- \widetilde{\lambda}^T \boldsymbol{\Delta}\boldsymbol{\Gamma^1} \left(\boldsymbol{I} - \boldsymbol{C_{\ell^1}^1}\boldsymbol{\Gamma^1}\right)^{-1} \boldsymbol{C_{\ell^1}^1} \left(\boldsymbol{\Gamma^1} G_\varepsilon^1 - \boldsymbol{\mu}\tau_\varepsilon^1 - \frac{1_{\phi_n=0}\boldsymbol{\mu}\tau_\varepsilon^2}{1+r}\right)}_{\textit{Keynesian stimulus (alleviation of involuntary unemployment)}}
\end{aligned}
\tag{A77}
$$

where γ is the marginal value of public funds, $\boldsymbol{\Gamma^1} \equiv \boldsymbol{R_{L^1}^1}\widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}$, $\boldsymbol{\mu}$, $\boldsymbol{\phi}$, and $\boldsymbol{\Delta}$ are the diagonal matrices of type weights, borrowing wedges, and labor wedges, respectively.

Proof. The planner takes prices and the interest rate as given. Goods and labor market clearing and first-period rationing determine the change in first-period employment as a function of G_ε^1 and τ_ε^1. We are left with the following first-order condition:

$$
0 = dW + \gamma \left[\mu^T \tau_\varepsilon^1 + \frac{\mu^T \tau_\varepsilon^2}{1+r} - \mathbb{1}^T G_\varepsilon^1 - \frac{\mathbb{1}^T G_\varepsilon^2}{1+r}\right] \tag{A78}
$$

where dW is as in Equation A13. This gives an expression for the change in welfare in terms of τ_ε, G_ε, and ℓ_ε^1, the change in first-period employment. By Equation 10, $\boldsymbol{\mu}\ell_\varepsilon^1 = \boldsymbol{\Gamma^1}(\boldsymbol{I} - \boldsymbol{C_{\ell^1}^1}\boldsymbol{\Gamma^1})^{-1}\partial Y^1$, where $\boldsymbol{\Gamma^1} \equiv \boldsymbol{R_{L^1}^1}\widehat{\boldsymbol{L}}^{\boldsymbol{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}$ and $\partial Y^1 = G_\varepsilon^1 - \boldsymbol{C_{\ell^1}^1}\boldsymbol{\mu}\tau_\varepsilon^1 + \boldsymbol{C_{\tau^2}^1}\boldsymbol{\mu}\tau_\varepsilon^2$.[58] For borrowing-constrained households, $\boldsymbol{C_{\tau^2}^1} = 0$; they would already like to substitute additional consumption toward the first period but are constrained not to do so. Other households are Ricardian, implying (since preferences are additively spearable in consumption and labor) that $\boldsymbol{C_{\tau^2}^1} = -\frac{\boldsymbol{C_{\ell^1}^1}}{1+r}$. Plugging in for dW, and using matrix notation, we have

$$
\begin{aligned}
0 = \widetilde{\lambda}^T \Bigg[& - \boldsymbol{\Delta}\boldsymbol{\Gamma^1}(\boldsymbol{I} - \boldsymbol{C_{\ell^1}^1}\boldsymbol{\Gamma^1})^{-1} \left(G_\varepsilon^1 - \boldsymbol{C_{\ell^1}^1}\boldsymbol{\mu}\left(\tau_\varepsilon^1 + \frac{1_{\phi_n=0}\,\tau_\varepsilon^2}{1+r}\right)\right) - \left(\boldsymbol{\mu}\tau_\varepsilon^1 + \frac{\boldsymbol{\mu}(\boldsymbol{I}-\boldsymbol{\phi})\tau_\varepsilon^2}{1+r}\right) \\
& + \left(\boldsymbol{\mu}\boldsymbol{WTP^1} G_\varepsilon^1 + \boldsymbol{\mu}(\boldsymbol{I}-\boldsymbol{\phi})\frac{\boldsymbol{WTP^2}}{1+r} G_\varepsilon^2\right) \Bigg] + \gamma \left(\mu^T \tau_\varepsilon^1 + \frac{\mu^T \tau_\varepsilon^2}{1+r} - \mathbb{1}^T G_\varepsilon^1 - \frac{\mathbb{1}^T G_\varepsilon^2}{1+r}\right)
\end{aligned}
\tag{A79}
$$

[58] Here we have used that since preferences are additively separable in consumption and labor, each household's MPCs out of labor and transfer income are the same.

Now, observe that the term on the first line can be rewritten:

$$\begin{aligned}
&\boldsymbol{\Gamma^1}(\boldsymbol{I}-\boldsymbol{C^1_{\ell^1}\Gamma^1})^{-1}\left(G^1_\varepsilon-\boldsymbol{C^1_{\ell^1}\mu}\left(\tau^1_\varepsilon+\frac{1_{\phi_n=0}\ \tau^2_\varepsilon}{1+r}\right)\right)=\boldsymbol{\Gamma^1}\left(\sum_{k=0}^{\infty}(\boldsymbol{C^1_{\ell^1}\Gamma^1})^k\right)\left(G^1_\varepsilon-\boldsymbol{C^1_{\ell^1}\mu}\left(\tau^1_\varepsilon+\frac{1_{\phi_n=0}\ \tau^2_\varepsilon}{1+r}\right)\right)\\
=&\boldsymbol{\Gamma^1}\left(G^1_\varepsilon+\left(\sum_{k=0}^{\infty}(\boldsymbol{C^1_{\ell^1}\Gamma^1})^k\right)\boldsymbol{C^1_{\ell^1}\Gamma^1}G^1_\varepsilon\right)-\boldsymbol{\Gamma^1}\left(\sum_{k=0}^{\infty}(\boldsymbol{C^1_{\ell^1}\Gamma^1})^k\right)\boldsymbol{C^1_{\ell^1}\mu}\left(\tau^1_\varepsilon+\frac{1_{\phi_n=0}\ \tau^2_\varepsilon}{1+r}\right)\\
=&\boldsymbol{\Gamma^1}G^1_\varepsilon+\boldsymbol{\Gamma^1}\left(\boldsymbol{I}-\boldsymbol{C^1_{\ell^1}\Gamma^1}\right)^{-1}\boldsymbol{C^1_{\ell^1}}\left(\boldsymbol{\Gamma^1}G^1_\varepsilon-\boldsymbol{\mu}\tau^1_\varepsilon-\boldsymbol{\mu}\frac{1_{\phi_n=0}\ \tau^2_\varepsilon}{1+r}\right)
\end{aligned}\tag{A80}$$

Substituting this back in and rearranging, we obtain Equation A77. □

The opportunistic government purchases term is as in Werning (2011) and Baqaee (2015). It augments the standard first-order condition for government purchases with a labor-wedge term, reflecting the fact that the social cost of additional government purchases is lower than the market cost when they are produced using underemployed labor. The second term is also an augmented version of the standard expression for government purchases—this time in the second period. The borrowing wedge reflects the fact that households with binding borrowing constraints implicitly discount the future at a higher-than-market rate; the planner must account for this when deciding whether to make purchases on their behalf.

The third term of Equation A77 is a standard, pure redistribution term, weighing the private benefits of transfers against the social cost (the MVPF). The fourth term augments this, when there are borrowing constraints. In particular, taxes in the second period are less costly to borrowing-constrained households since they discount the future more heavily than the market rate indicates.

Finally, the last line captures the value of stimulus brought on by changes in income—those corresponding to pure income transfers via taxes and labor market income earned by employees producing government purchases.[59]

C.5. Evaluating Optimality of Fiscal Policy

We now use Proposition 12 to more fully characterize optimal fiscal policies in two benchmark cases where the planner's indifference between transfers to each household and/or purchases in each sector leads to optimality conditions that can be evaluated without knowledge of the rich interconnections between households.

Proposition 13. *The following two statements are true:*

[59] If second period purchases are held constant, then the net income transfer is zero, i.e. this term operates solely through redistribution to different households (who may spend differently).

1. *(Optimal transfer policy) Suppose that the marginal social dis-utility of labor supply is constant across all households rationed to on the margin at the optimum.*[60] *Then* $dW = 0$ *with respect to marginal changes in first-period transfers if and only if, for all* $n \in N$*,*

$$\gamma = \widetilde{\lambda}_n \left(1 + \frac{m_n}{1 - m_n}(-\Delta_n)\right) \tag{A81}$$

where γ *is the marginal value of public funds.*

2. *(Optimal purchases policy) Suppose that the social gains from first-period government purchases are equal to some* $\tilde{v}$ *across all goods and constraints bounding purchases above zero do not bind. Then* $dW = 0$ *with respect to marginal changes in first-period purchases if and only if, for all* $i \in I$*,*

$$\gamma = \tilde{v} + \frac{1}{1 - \widetilde{m}_i}\left(-\widetilde{\lambda\Delta}_i\right) \tag{A82}$$

where $\widetilde{m}_i$ *is the rationing-weighted average MPC in the production of good* i *and* $\widetilde{\lambda\Delta}_i$ *is the rationing-and-welfare-weighted average rationing wedge in the production of good* i*.*[61]

Proof. We first prove the result for first-period transfers. At any optimum, we know that Equation A77 must hold for all policy variations $\tau_\varepsilon^1 \in \mathbb{R}^N$ that only vary first-period transfers, keeping other instruments fixed. Taking $\tau_\varepsilon^1 = e_n$, the nth basis vector, we see that:

$$\left(\widetilde{\lambda}^T - \gamma \mathbb{1}\right)_n^T = \left(\widetilde{\lambda}^T \boldsymbol{\Delta\Gamma^1}\left(\boldsymbol{I} - \boldsymbol{C^1_{\ell^1}\Gamma^1}\right)^{-1}\boldsymbol{C^1_{\ell^1}}\right)_n \tag{A83}$$

Stacking these equations over n, we obtain:

$$\left(\widetilde{\lambda} - \gamma \mathbb{1}\right)^T = \widetilde{\lambda}^T \boldsymbol{\Delta\Gamma^1}\left(\boldsymbol{I} - \boldsymbol{C^1_{\ell^1}\Gamma^1}\right)^{-1}\boldsymbol{C^1_{\ell^1}} \tag{A84}$$

Since $\{e_n\}$ is a basis and Equation A77 is linear, this equation fully encompasses the optimality condition of Proposition 12 with respect to first period transfers.

We can simplify this system of equations. First, see that:

$$\boldsymbol{\Gamma^1}(\boldsymbol{I} - \boldsymbol{C^1_{\ell^1}\Gamma^1})^{-1}\boldsymbol{C^1_{\ell^1}} = \sum_{k=0}^{\infty} \boldsymbol{\Gamma^1}(\boldsymbol{C^1_{\ell^1}\Gamma^1})^k \boldsymbol{C^1_{\ell^1}} = \sum_{k=1}^{\infty} \boldsymbol{\Gamma^1 C^1_{\ell^1}} \tag{A85}$$

[60] Formally, if $\left[\boldsymbol{\Gamma^1 C^1_{\ell^1}}\right]_{n,0} \neq \vec{0}$ then $\widetilde{\lambda}_n(1 + \Delta_n) = const$, where $\boldsymbol{\Gamma^1} \equiv l_L^1 \widehat{\boldsymbol{L}}^{\boldsymbol{1}}\left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}$.

[61] Formally, $\widetilde{m}_i \equiv \left(m^T \boldsymbol{\Gamma^1}\right)_i$ and $\widetilde{\lambda\Delta}_i \equiv \left(\widetilde{\lambda}^T \boldsymbol{\Delta\Gamma^1}\right)_i$.

Adding $\widetilde{\lambda}^T\boldsymbol{\Delta}$ to both sides of Equation A84, we therefore obtain:

$$\left(\widetilde{\lambda}^T(\boldsymbol{I}+\boldsymbol{\Delta})-\gamma\mathbb{1}^T\right)=\widetilde{\lambda}^T\boldsymbol{\Delta}\left(\boldsymbol{I}-\boldsymbol{\Gamma^1}\boldsymbol{C^1_{\ell^1}}\right)^{-1}\implies\left(\widetilde{\lambda}^T(\boldsymbol{I}+\boldsymbol{\Delta})-\gamma\mathbb{1}^T\right)\left(\boldsymbol{I}-\boldsymbol{\Gamma^1}\boldsymbol{C^1_{\ell^1}}\right)=\widetilde{\lambda}^T\boldsymbol{\Delta} \tag{A86}$$

Now, express $\boldsymbol{\Gamma^1}\boldsymbol{C^1_{\ell^1}}=\boldsymbol{\Gamma^1}\widehat{\boldsymbol{C}}^{\boldsymbol{1}}\boldsymbol{m}$. Recognizing that all columns of the spending-to-income matrix $\boldsymbol{\Gamma^1}\widehat{\boldsymbol{C}}^{\boldsymbol{1}}$ sum to one, as total spending is equal to total factor income, and—by assumption—that $\widetilde{\lambda}_n(1+\Delta_n)$ is constant across all households n except for those for which the n^{th} row of $\boldsymbol{\Gamma^1}\boldsymbol{C^1_{\ell^1}}$ is zero, (A86) can be rewritten as:

$$\left(\widetilde{\lambda}^T(\boldsymbol{I}+\boldsymbol{\Delta})-\gamma\mathbb{1}^T\right)(\boldsymbol{I}-\boldsymbol{m})=\widetilde{\lambda}^T\boldsymbol{\Delta} \tag{A87}$$

We therefore have all, for all n, that

$$\widetilde{\lambda}_n(1+\Delta_n)-\gamma=\frac{1}{1-m_n}\widetilde{\lambda}_n\quad\implies\quad\gamma=\widetilde{\lambda}_n\left(1+\frac{m_n}{1+m_n}(-\Delta_n)\right) \tag{A88}$$

We prove the result for first-period government purchases in an analogous way. To begin, consider Equation A77 for policy variations $G^1_\varepsilon\in\mathbb{R}^I$ that only vary first period purchases. Again considering each basis vector of $\mathbb{R}^I$ and stacking we obtain:

$$0=\widetilde{\lambda}^T\boldsymbol{W}\boldsymbol{T}\boldsymbol{P^1}-(\gamma\mathbb{1}^T+\widetilde{\lambda}^T\boldsymbol{\Delta}\boldsymbol{\Gamma^1})-\widetilde{\lambda}^T\boldsymbol{\Delta}\boldsymbol{\Gamma^1}(\boldsymbol{I}-\boldsymbol{C^1_{\ell^1}}\boldsymbol{\Gamma^1})^{-1}\boldsymbol{C^1_{\ell^1}}\boldsymbol{\Gamma^1} \tag{A89}$$

This can be rewritten as:

$$\widetilde{\lambda}^T\boldsymbol{W}\boldsymbol{T}\boldsymbol{P^1}-\gamma\mathbb{1}^T=\widetilde{\lambda}^T\boldsymbol{\Delta}\boldsymbol{\Gamma^1}(\boldsymbol{I}-\boldsymbol{C^1_{\ell^1}}\boldsymbol{\Gamma^1})^{-1} \tag{A90}$$

From the assumption that the social gains from government purchases equal $\tilde{v}$, we have that $\widetilde{\lambda}^T\boldsymbol{W}\boldsymbol{T}\boldsymbol{P^1}=\tilde{v}\mathbb{1}^T$. Moreover, by definition $\widetilde{\lambda\Delta}^T=\widetilde{\lambda}^T\boldsymbol{\Delta}\boldsymbol{\Gamma^1}$. Hence (A90) can be rewritten as

$$\tilde{v}\mathbb{1}^T-\gamma\mathbb{1}^T=\widetilde{\lambda\Delta}^T\left(\boldsymbol{I}-\boldsymbol{C^1_{\ell^1}}\boldsymbol{\Gamma^1}\right)^{-1} \tag{A91}$$

Next, define $\widetilde{m}_i\equiv\left(m^T\boldsymbol{\Gamma^1}\right)_i$ to be the rationing-weighted average MPC in the production of good i and let $\widetilde{\boldsymbol{m}}$ be the corresponding matrix with $\widetilde{m}$ on the diagonal. Moreover, define $\underline{C}_{ji}\equiv(\boldsymbol{C^1_{\ell^1}}\boldsymbol{\Gamma^1})_{ji}/\widetilde{m}_i$ to be the average direction of consumption of workers producing i, weighted by their MPC and marginal rationing in i's production.[62] Crucially, note that

[62]For any i with $\tilde{m}_i=0$, define $\underline{C}_{ji}$ in any way satisfying $\sum_j\underline{C}_{ji}=1$.

$\underline{\boldsymbol{C}}\widetilde{\boldsymbol{m}} = \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{\Gamma}^{\mathbf{1}}$ by construction and that $\mathbb{1}^T\underline{\boldsymbol{C}}\widetilde{\boldsymbol{m}} = \mathbb{1}^T\widetilde{\boldsymbol{m}}$:

$$\mathbb{1}^T\underline{\boldsymbol{C}}\widetilde{\boldsymbol{m}} = \mathbb{1}^T\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{\Gamma}^{\mathbf{1}} = m^T\boldsymbol{\Gamma}^{\mathbf{1}} = \widetilde{m}^T \tag{A92}$$

The first order condition for purchases (A91) is therefore equivalent to:

$$(\tilde{v}-\gamma)\mathbb{1}^T\left(\boldsymbol{I}-\boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{\ell}^{\mathbf{1}}}\boldsymbol{\Gamma}^{\mathbf{1}}\right) = (\tilde{v}-\gamma)\mathbb{1}^T\left(\boldsymbol{I}-\widetilde{\boldsymbol{m}}\right) = \widetilde{\lambda\Delta}^T \iff \gamma = \tilde{v} + \frac{1}{1-\widetilde{m}_i}(-\widetilde{\lambda\Delta}_i) \quad \forall i. \tag{A93}$$

□

Proposition 13 says that the planner may verify whether the current policy is optimal despite having very partial knowledge about the economy. In the transfer case, the planner only needs information on household level welfare weights ($\widetilde{\lambda}_n$), rationing wedges (Δ_n), and MPCs (m_n) – not the network of marginal spending flows between households. In the purchases case, the planner needs to know the average MPC and welfare-weighted rationing wedge by industry; these require knowledge of the rationing function linking output to incomes, but not the directed consumption matrix.

The main idea underlying Proposition 13 is that—at an optimum—the social value of additional spending by any household is independent of how that spending is directed. This is clearest in the case of transfers: For any household employed in order to produce marginally-demanded goods, the social value of their employment is equal to the value of a transfer to that household, less the dis-utility of labor. Since (by assumption) the dis-utility of labor is constant across households, and since—at an optimum—the value of transfers must also be constant across households, it follows that the social value of additional employment is constant across households. Since the planner is indifferent over the direction of household spending, she targets solely based on the magnitude of that spending—i.e. household MPCs—as well as household welfare weights. A similar argument applies in the case of government purchases.

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 (including an infinite horizon) (D.3). We also provide a version of Proposition 3 on optimal policy with fixed markups (D.2).

D.1. Imperfect Competition

In this section 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 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 $\Pi_i^t = \frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t} p_i^t Q_i^t$.[63] 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 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, $F_i^t(X_i^t, L_i^t, z_i^t) \geqslant 1$ 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 14. *Under Assumptions 4, 5, and 7, for a given z^t and $\widehat{\Pi}^t$, there exists a unique p^t consistent with rationing equilibrium, independent of demand.*

Proof. We modify the proof of proposition 7 to accommodate markups. 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 $(1 - \widehat{\Pi}^t)F_i^t(X_i^t, L_i^t, z_i^t)$. 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 7. □

Since, by Proposition 14, prices are exogenous, we may without loss of generality normalize them to one. This implies that firms earn profits per unit sold of $\frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t} / \left(1 + \frac{\widehat{\Pi}_i^t}{1-\widehat{\Pi}_i^t}\right) = \widehat{\Pi}_i^t$.

[63] 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.

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 $\pi^t \in \mathbb{R}^{\mathcal{I}}$. We let $\pi_n^t = D^t(\Pi^t)_n / \mu_n$ represent each household of type n's total dividend income in period t. 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. We therefore allow household consumption and labor supply functions to depend on π_n^t directly.

We can now state a profit-inclusive Keynesian cross. The main difference to Proposition 4 comes from the need to account for changes in profits, how these are distributed to households as dividends, and their directed MPCs out of dividends. A second difference is that the presence of profits inextricably links the first period to the second – because changes in future profits affect lifetime incomes, which affect consumption today.

Proposition 15. *For any shock inducing a first-period partial equilibrium effect ∂Q, the general equilibrium response in production satisfies:*

$$dQ = \widehat{\boldsymbol{X}} dQ + \boldsymbol{C_{\ell^1}} \boldsymbol{R^1_{L^1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} dQ^1 + \boldsymbol{C_\pi} \boldsymbol{D_\Pi} \widehat{\boldsymbol{\Pi}} dQ + \partial Q \tag{A94}$$

where $\boldsymbol{C_\pi}$ is the matrix of household directed MPCs out of profit income, 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 $D^t_{\Pi_i^t}(\Pi^t)_n$ – 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 $\widehat{\Pi}_i^t$. 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:

$$dQ \;= \widehat{\boldsymbol{X}} dQ + \widehat{\boldsymbol{X}}_{\boldsymbol{z}} dz Q + \boldsymbol{C_p p_z} dz + \boldsymbol{C_{\ell^1}} d\ell^1 + \boldsymbol{C_\theta} d\theta + dG + \boldsymbol{C_\pi} \boldsymbol{D_\Pi} \widehat{\boldsymbol{\Pi}} dQ \tag{A95}$$

Plugging in for $d\ell^1 = \boldsymbol{R^1_{L^1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} dQ^1 + \boldsymbol{R^1_{L^1}} d\widehat{\boldsymbol{L}}^{\boldsymbol{1}} Q^1$, we have

$$dQ = \widehat{\boldsymbol{X}} dQ + \boldsymbol{C_{\ell^1}} \boldsymbol{R^1_{L^1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} dQ^1 + \boldsymbol{C_\pi} \boldsymbol{D_\Pi} \widehat{\boldsymbol{\Pi}} dQ + \partial Q \tag{A96}$$

where here $\partial Q = \widehat{\boldsymbol{X}}_{\boldsymbol{z}} dz Q + \boldsymbol{C_p p_z} dz + \boldsymbol{C_{\ell^1}} \boldsymbol{R^1_{L^1}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}}_z dz Q^1 + \boldsymbol{C_\tau} d\tau + \boldsymbol{C_\theta} d\theta + dG$. □

D.2. Optimal Policy with Imperfect Competition

In this section, we extend the optimal policy results of Section 6 to the more general environment with constant, non-zero markups. We continue to normalize prices p_i^t to one without loss of generality.

To highlight as clearly as possible the parallels to the case without profits, we make two important assumptions. First—although in the first period, profit-creation is uninternalized by households—we assume that the government incentivizes second-period profit-creation with Pigouvian subsidies funded lump-sum by shareholders.

Assumption 8. *There is an ad-valorem subsidy s_i^2 on the purchase of i in the second period (for consumption or production), set equal to the profit rate $\widehat{\Pi}_i^2$. It is funded directly by an additional lump-sum, second-period tax $\hat{\tau}_n^2$ defined by $\mu_n \hat{\tau}_n^2 = \mu^n \pi_n^2 = D^2(\Pi^2)_n = D^2(\{s_i^2 Q_i^2\}_{i \in I})_n$.*

Second, we assume that the MPC out of future profits is zero. This is a rather weak assumption, as the MPC out of even *current* capital income is small empirically.

Assumption 9. $C_{\pi^2}^1 = 0$.

With these assumptions in mind, we begin by defining the household's problem. It is the same as in Equation 3', except that households now also receive profit income, which affects their borrowing constraints and budget constraints in the same manner as lump-sum taxes (but with the opposite sign). Note that this microfoundation implies $\boldsymbol{C_{\ell^1}} = -\boldsymbol{C_{\tau^1}} = \boldsymbol{C_{\pi^1}}$. That is, additional income from rationed, first-period labor has the same effects on consumption as additional income from first-period transfers and profits.

As in Section 6, we study the policy problem of a fiscal planner. Formally, the planner's problem is the same as in Equation A76 except that household behavior solves Equation 3' with the profit-inclusive budget constraint, and aggregate variables evolve according to Equation A94.

Our main result considers the change in welfare induced by changes in transfers and government expenditure, analogously to Proposition 3.

Lemma 2. *Under Assumptions 8 and 9, the change in welfare dW due to a small change in taxes and government expenditure—at a constant interest rate—can be expressed as:*

$$dW = \sum_{n \in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n dl_n^1 + d\pi_n^1 - \left(d\tau_n^1 + (1 - \phi_n) \frac{d\tau_n^2}{1 + r^1} \right) + \left(WTP_n^1 dG^1 + (1 - \phi_n) \frac{WTP_n^2}{1 + r^1} dG^2 \right) \right] \tag{A97}$$

where $\widetilde{\lambda}_n$ is the value the planner places on the marginal transfer of first-period wealth to a household of type n, Δ_n and ϕ_n are n's implicit first-period labor wedge and borrowing wedge, and WTP_n^t is the vector of n's marginal willingness to pay for period t government expenditures on each good, in period t dollars. The changes in first-period employment and profits are in turn given by

$$\begin{aligned} \boldsymbol{\mu} dl^1 &= \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \widehat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} dY^1, \qquad \boldsymbol{\mu} d\pi^1 = \boldsymbol{D}^{\mathbf{1}}_{\boldsymbol{\Pi}^{\mathbf{1}}} \widehat{\boldsymbol{\Pi}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}} \right)^{-1} dY^1, \\ dY^1 &= \left(\boldsymbol{I} - \boldsymbol{C}^{\mathbf{1}}_{\boldsymbol{y}^{\mathbf{1}}} \left(\boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L}^{\mathbf{1}}} \widehat{\boldsymbol{L}}^{\mathbf{1}} + \boldsymbol{D}^{\mathbf{1}}_{\boldsymbol{\Pi}^{\mathbf{1}}} \widehat{\boldsymbol{\Pi}}^{\mathbf{1}} \right) \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} \right)^{-1} \partial Q^1 \end{aligned} \tag{A98}$$

Proof. We follow the same steps as the proof of Proposition 3 (see Appendix A.3) up to the substitution of the budget constraint, which now includes profits. With profits, differentiating the household's lifetime budget constraint (at constant r^1) gives:

$$p^1 dc_n^1 - dl_n^1 - d\pi_n^1 + \frac{p^1 dc_n^2 - dl_n^2}{1+r^1} = -d\tau_n^1 + \frac{d\pi_n^2 - d\hat{\tau}_n^2 - d\tau_n^2}{1+r^1} \tag{A99}$$

Note that by construction, $d\hat{\tau}_n^2 = d\pi_n^2$. Cancelling these terms and substituting in the differentiated budget constraint, we have:

$$dW = \sum_{n \in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n dl_n^1 + \phi_n \left(p^1 dc_n^1 - dl_n^1 \right) + (1-\phi_n) \left(d\pi_n^1 - d\tau_n^1 - \frac{d\tau_n^2}{1+r^1} \right) + \left(\frac{w_{nG}^1}{\kappa_n^1} dG^1 + \left(\frac{1-\phi_n}{1+r^1} \right) \frac{w_{nG}^2}{\kappa_n^2} dG^2 \right) \right] \tag{A100}$$

For households with non-strictly-binding borrowing constraints, $\phi_n = 0$. For households with $\phi_n > 0$, the borrowing constraint $\underline{s}_n^1 = l_n^1 + \pi_n^1 - \tau_n^1 - p^1 c_n^1$ implies $p^1 dc_n^1 + d\tau_n^1 = dl_n^1 + d\pi_n^1$. Defining the within-period willingnesses to pay $WTP_n^t = \frac{w_{nG}^t}{\kappa_n^t}$, we arrive at the final expression:

$$dW = \sum_{n \in N} \widetilde{\lambda}_n \mu_n \left[-\Delta_n dl_n^1 + \left(d\pi_n^1 - d\tau_n^1 - (1-\phi_n) \frac{d\tau_n^2}{1+r^1} \right) + \left(WTP_n^1 dG^1 + (1-\phi_n) \frac{WTP_n^2}{1+r^1} dG^2 \right) \right] \tag{A101}$$

Finally, the expressions for $dl, d\pi, dY$ come from rearranging Equation A94 under Assumption 9 and using $dY^1 = (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}}) dQ^1$. □

Studying Equation A97 reveals a key insight: Under Assumptions 8 and 9, the change in welfare due to a change in taxes and expenditures is the same as in an *as-if* economy without profits but where share-holders supply labor with a wedge -1. This labor supply wedge corresponds to complete under-employment; share-holders—who experience no marginal disutility of holding shares—would continue to be willing to hold shares until profits-per-revenue reached zero. Just like labor suppliers, share-holders do not choose their income but rather take it as given.

One application of this observation is that in the simple case where (i) there is no intrinsic value of government purchases, (ii) λ_n is constant across households, (iii) all marginally-employed households have no disutility of labor ($\Delta_n = -1$), and (iv) there are no bias or homophily effects, we obtain a tight generalization of the policy result presented in Section 6.2. Formally,

$$dW \propto \sum_{n \in N} m_n \partial h_n^1 \tag{A102}$$

where $\partial h^1 = \boldsymbol{R}^{\mathbf{1}}_{\boldsymbol{L^1}} \widehat{\boldsymbol{L}}^{\mathbf{1}} \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\mathbf{1}} \right)^{-1} dG^1 - \boldsymbol{\mu} d\tau^1$. In particular, MPC targeting continues to be optimal in widespread recessions, despite the inclusion of profits.

Inutitively, the assumption that all marginal labor supplies have a labor supply wedge of -1 matches with the shareholders' implicit labor supply wedge of -1: both are indifferent to supplying more of their factor. Thus, there is zero social cost to any marginal employment, so the optimal policy maximizes output. As without markups, the output-maximizing policy targets MPC when bias and homophily effects are absent.

D.3. 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_n^t 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:

$$\sum_{t \leqslant T} \frac{\ell_n^t}{\prod_{t' \leqslant t} 1 + r^{t'}} = \sum_{t \leqslant T} \frac{p^t c_n^t + \tau_n^t}{\prod_{t' \leqslant t} 1 + r^{t'}} \tag{A103}$$

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

$$\sum_{n \in N} \mu_n \Big(\sum_{t \in T} \frac{1}{\prod_{t' \leqslant t} 1 + r^{t'}} \tau_n^t \Big) = \sum_{t \in T} \frac{1}{\prod_{t' \leqslant t} 1 + r^{t'}} p^t G^t \tag{A104}$$

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 $\sum_{n \in N} \mu_n R_n^t(L^t) = \sum_{n \in N} L_n^t$.

Definition 3. *(Dynamic rationing equilibrium) A dynamic rationing equilibrium is a profile*

of prices and quantities $\{p_i^t, \ell_n^t, c_n^t, L_i^t, X_{ij}^t\}_{t\in\{0,...,T\},n\in N,i,j\in I}$ *that satisfies firm optimization (as in* (1)*), labor rationing in periods* $t \leqslant T-1$*, consumption and final period labor supply consistent with the Marshallian demands described above, government budget balance, and goods and labor market clearing.*

Under this dynamic equilibrium concept, we obtain an analogous fiscal multiplier to that of our two period model. This reflects that the non-substitution theorem continues to hold and all inter-temporal prices are assumed to be constant.

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 $(v_1, ..., v_{T-1})$. In a similar fashion, we let $\widehat{\boldsymbol{L}}^{-\boldsymbol{T}}$ and $\widehat{\boldsymbol{X}}^{-\boldsymbol{T}}$ denote the block-diagonal $\big((T-1)\times I\big)\times\big((T-1)\times I\big)$ matrix 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 $\big((T-1)\times N\big)\times\big((T-1)\times I\big)$ 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 16. *For any small shock to fiscal policy inducing a partial equilibrium effect* ∂Y^{-T} *in periods* $1, ..., T-1$*, there exists a selection from the equilibrium set such that the general equilibrium response of* $1, ..., T-1$ *period values added* dY^{-T} *is given by:*

$$dY^{-T} = \left(\boldsymbol{I} \;-\; \widehat{\boldsymbol{C}}^{-\boldsymbol{T}}\,\boldsymbol{m}^{-\boldsymbol{T}}\,\boldsymbol{R}^{-\boldsymbol{T}}_{\boldsymbol{L}^{-\boldsymbol{T}}}\,\widehat{\boldsymbol{L}}^{-\boldsymbol{T}}\,\left(\boldsymbol{I}-\widehat{\boldsymbol{X}}^{-\boldsymbol{T}}\right)^{-1}\right)^{-1}\partial Y^{-T} \tag{A105}$$

Proof. The proof is directly analogous to that of Proposition 1. One must simply expand the goods and labor market clearing conditions and rearrange terms. □

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, (A105) coincides with their intertemporal Keynesian cross when there is a single good.

E. Measuring Rationing Wedges

In this Appendix, we describe how to recover rationing wedges in the data and how we estimate the counterfactual welfare effects of fiscal stimulus in the Great Recession. We present two microfoundations for the same, particularly simple form of the rationing wedge in terms of the demographic level percentage change in unemployment from before the Great Recession to during the Great Recession. Namely:

$$\Delta_n = \underbrace{\frac{\ell_n^2 - \ell_n^1}{\ell_n^2}}_{\text{Gap in Labor Income}} \tag{A106}$$

E.1. Intensive margin microfoundation

Our first microfoundation assumes that all households within each demographic group can be treated as having the same quantity of labor supply. This is equivalent to the assumption that all labor supply adjustment happens on the intensive margin (hours worked), and that workers within any demographic group experience the same change in hours.

We now assume that (i) all households have slack borrowing constraints $\phi_n = 0$ and that (ii) for all households n, $\beta_n(1 + r) = 1$. The latter is empirically justifiable with standard estimates for discount factors and the real interest rate in the US during the Great Recession.[64] By household optimization, these assumptions imply that second period labor disutility and first-period consumption utility are related, for any goods i_t consumed positively by n at t

$$v_n^{2\prime}(l_n^2) = u_{ni_2}^2(c_n^2) = u_{ni_1}^1(c_n^1) \tag{A107}$$

Each household n's first-period labor wedge Δ_n therefore satisfies:

$$\Delta_n = \frac{v_n^{1\prime}(\ell_n^1)}{v_n^{2\prime}(\ell_n^2)} - 1 \tag{A108}$$

Finally, we follow a standard calibration that assumes labor disutility is quadratic. In this case, the rationing wedge is the percentage gap in labor supply from the steady state:

$$\Delta_n = \frac{l_n^1 - l_n^2}{l_n^2}. \tag{A109}$$

[64] Alternatively, it follows exactly from (i) and the inter-temporal Euler equation when in the special case where consumption utility is linear.

E.2. Extensive margin microfoundation

Our second microfoundation focuses on the polar case in which, before a change in government purchases, all households within a demographic group are either fully employed or fully unemployed. Each household of type n supplies labor inelastically up to some level $\tilde{\ell}_n$ after which their marginal dis-utility of labor supply sharply—but continuously—increases, so that they always supply close to $\tilde{\ell}_n$ in equilibrium.

At the initial equilibrium, a total mass $\zeta_n \leqslant \mu_n$ of type-n households are employed at the efficient level in the initial equilibrium; the remainder are unemployed. This implies that initially employed households supply $\approx \tilde{\ell}_n$ in both periods, whereas initially unemployed households supply 0 in the first period and $\approx \tilde{\ell}_n$ in the second period. Total first- and second-period labor supplies by group n are therefore approximately $\ell_n^1 \approx \zeta_n \mu_n \tilde{\ell}_n$ and $\ell_n^2 \approx \mu_n \tilde{\ell}_n$, respectively.[65]

Now consider a change in labor demand. We assume that – within each demographic group – each firm rations the same expected amount of marginal labor to each worker. However, in the case of employed workers, they do this by rationing infinitesimally more labor to a continuum of workers, whereas in the case of unemployed workers, they do this by hiring workers at their efficient level of labor supply. The former has only second-order welfare consequences. The latter increases welfare by

$$\begin{aligned}\widetilde{\Delta}_n \equiv & \max_{\substack{\mathbb{1}^T c^1 + \frac{\mathbb{1}^T c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} \leqslant l_n^1 + \frac{l_n^2}{1+r} \\ l_n^1 - \mathbb{1}^T c_n^1 - \tau_n^1 \geqslant \underline{s}_n}} u_n^1(c^1) - v_n^1(l_n^1) + \beta_n \left[u_n^2(c^2) + v_n^2(l_n^2)\right] \\ & - \max_{\substack{\mathbb{1}^T c^1 + \frac{\mathbb{1}^T c^2}{1+r} + \tau_n^1 + \frac{\tau_n^2}{1+r} \leqslant \frac{l_n^2}{1+r} \\ -\mathbb{1}^T c_n^1 - \tau_n^1 \geqslant \underline{s}_n}} u_n^1(c^1) - v_n^1(0) + \beta_n \left[u_n^2(c^2) + v_n^2(l_n^2)\right]\end{aligned} \tag{A110}$$

per newly employed worker. We assume $\widetilde{\Delta}_n$ is constant across n and normalize it to 1. Alternatively, one may think of differences in $\widetilde{\Delta}_n$ as embedded through differences in $\widetilde{\lambda}_n$.

Finally, note that the fraction of the f_n-sized subgroup which is initially unemployed – and therefore experiences the welfare gain $\widetilde{\Delta}_n$ if employed – is equal to

$$\frac{\mu_n - \zeta_n}{\mu_n} = \frac{(\mu_n - \zeta_n) l_n^*}{\mu_n l_n^*} \approx \frac{l_n^2 - l_n^1}{l_n^2} \tag{A111}$$

The average rationing wedge among type ns, and – since labor is rationed to each member

[65] These approximations are exact in the limit where labor disutility is kinked at $\tilde{\ell}_n$.

equally – the appropriate aggregate rationing wedge for type ns is therefore

$$\Delta_n \approx \frac{l_n^2 - l_n^1}{l_n^2} \tag{A112}$$

E.3. Estimation

For our Great Recession analysis, we compute labor wedges in the ACS by taking the percentage change in labor hours worked from 2005-06 to 2009-10 in each of our state-by-demographic bins. When there are no observations in any given bin, we assume that the change in labor hours is given by the state-level average.

For robustness, we compute a version of the state-by-demographic level rationing wedge by imposing that each demographic group's rationing wedge is the change in hours for that demographic group nationwide compared to the average multiplied by the average change in hours across demographics at the state level. The results are very similar, with the R^2 of multipliers and labor wedges in explaining welfare changes dropping slightly to 54% from 78%.

F. Validating the Model

The model that we develop and estimate in this paper makes stark predictions about the propagation of industry- and region-specific shocks. In this section, we attempt to empirically validate those quantitative predictions. Letting $\boldsymbol{T}$ be the $S \times (S \times I)$ matrix that adds all changes in final output within states, note that Proposition 1 provides an expression for changes in state-level total final output dY^{S1}:

$$dY^{S1} = \boldsymbol{T}\left(\boldsymbol{I} \; - \; \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \; \boldsymbol{m} \; \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \; \widehat{\boldsymbol{L}}^{\boldsymbol{1}} \; \left(\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}}\right)^{-1}\right)^{-1} \partial Q^1 \equiv \boldsymbol{M}\partial Q^1 \tag{A113}$$

where recall ∂Q^1 generalizes ∂Y^1 to the case with supply shocks.

Any identified partial equilibrium shock V will be some component of the many partial equilibrium shocks hitting the economy, which we can express as $\partial Q^1 = V^1 + U^1$, where U is the partial equilibrium effect on demand of the unobserved shocks hitting the economy. Plugging this in, we arrive at the foundation for our estimating equation:

$$dY^S_t = \boldsymbol{M}(V_t + U_t) = \beta \boldsymbol{M} V_t + \epsilon_{i,t} \tag{A114}$$

where V_t is the vector of identified industry-by-region shocks and $\boldsymbol{M}$ is our estimated generalized multiplier. The strict prediction of our model is that $\beta = 1$, meaning that we have perfectly predicted the heterogeneous effects of the shocks on final output growth. Note that the matrix $\boldsymbol{M}$ includes not only heterogeneity in the response to a shock in one's own market, but also how each market will respond to other markets through spillovers arising from spending network effects. Therefore, in addition to testing $\beta = 1$, we also test separately for the existence of spillovers of the nature predicted by the model. More specifically, we consider the following regression:

$$dY^S_t = \boldsymbol{M}(V_t + U_t) = \alpha_0 \; (\boldsymbol{M}_{\boldsymbol{diag}}) \; V_t + \alpha_1 \; (\boldsymbol{M}_{\boldsymbol{offdiag}}) \; V_t + \epsilon_{i,t} \tag{A115}$$

where $\boldsymbol{M}_{\boldsymbol{diag}}$ is the matrix obtained by using only the diagonal elements of the income multiplier $(\boldsymbol{I} - \widehat{\boldsymbol{C}}^{\boldsymbol{1}} \boldsymbol{m} \boldsymbol{R}^{\boldsymbol{1}}_{\boldsymbol{L}^{\boldsymbol{1}}} \widehat{\boldsymbol{L}}^{\boldsymbol{1}} (\boldsymbol{I} - \widehat{\boldsymbol{X}}^{\boldsymbol{1}})^{-1})^{-1}$ in (A113) and $\boldsymbol{M}_{\boldsymbol{offdiag}}$ is is the matrix obtained by using only its off-diagonal elements. α_0 measures the degree to which $\boldsymbol{M}$ accurately captures the direct effect of a shock and the within-state multiplier and α_1 measures the degree to which the model accurately captures the nature of the spillovers across regions and industries.

In the following section, we use identified shocks for the demand shock V_t from state-level military spending shocks from Nakamura and Steinsson (2014). Of course, bringing this to

the data presents several identification challenges particular to the shock in question. We address the challenges particular to the shock below by slightly modifying Equation A114 to fit its particular setting.

F.1. Government Purchase Shocks from Nakamura and Steinsson (2014)

We now consider the local government purchases shock developed by Nakamura and Steinsson (2014) to estimate the local purchases multiplier. We refer the reader to that paper for the details on the construction of the shock. We closely follow their original specification, using data on US states from 1966-2006. We restrict our attention to variation across states and our dependent variable is the 2-year change in state GDP per capita, divided by the level of state GDP lagged 2 periods. The state spending shock is the 2-year change in predicted GDP effects of military spending per capita V_t, also divided by the level of state GDP $y_{s,t}$ lagged 2 periods. Specifically, we run the following regression

$$\frac{y_{s,t} - y_{s,t-2}}{y_{s,t-2}} = \beta \frac{(\boldsymbol{M}V_t)_s - (\boldsymbol{M}V_{t-2})_s}{y_{s,t-2}} + \gamma_s + \gamma_t + e_{s,t} \tag{A116}$$

where γ_s and γ_y are state and year fixed effects, respectively. The central concern is that military spending is not random and may be directed towards states based on their economic performance. Therefore, we follow Nakamura and Steinsson (2014) and instrument the state changes in spending with state dummies interacted with national changes in military spending. Table A1 shows the results. First, Column 1 shows the replication of the result in Nakamura and Steinsson (2014), which is the equivalent of imposing that $\boldsymbol{M} = \boldsymbol{T}$ (call this $\boldsymbol{M_1}$). Column 2 shows the estimate of Equation A116. The estimates are noisy, but two small pieces of evidence suggest that including the multiplier provides a better fit for the data than the simple specification. First, while we cannot reject that the estimated coefficient β in either the $\boldsymbol{M_1}V$ or $\boldsymbol{M}V$ specification is 1, $\beta_{\boldsymbol{M}V}$ is closer to 1 than is $\beta_{\boldsymbol{M_1}V}$, suggesting that the heterogeneity embedded in $\boldsymbol{M}$ is getting us closer to capturing all of the variation in the data. Second, the R-squared in Column 2 is slightly higher than that in Column 1. However, the estimates are noisy and largely inconclusive.

The remaining columns of Table A1 show the estimates separating the own and spillover effects as in Equation A115. A finding that the coefficient on the spillover term was positive and close to 1 would suggest that our measure was accurately picking up the experienced spillovers. However, the estimates are also too noisy to be conclusive.

	Baseline			Robustness		
				No State FE	post-1980	post-1990
State Spending (M_1V)	1.474***					
	(0.373)					
Model Prediction (MV)		1.189***				
		(0.299)				
Model Prediction ($M_{diag}V$)			1.251***	1.166***	1.569**	0.657
			(0.355)	(0.309)	(0.611)	(0.908)
Model Prediction ($M_{offdiag}V$)			-0.145	0.496	-7.112	-8.899
			(3.367)	(3.242)	(5.443)	(9.385)
Constant						
Observations	1989	1989	1989	1989	1377	867
R-Squared	0.316	0.319	0.316	0.309	0.305	0.308

Table A1: Reduced Form Validation: Government Spending from Nakamura and Steinsson (2014)

G. Additional Tables and Figures

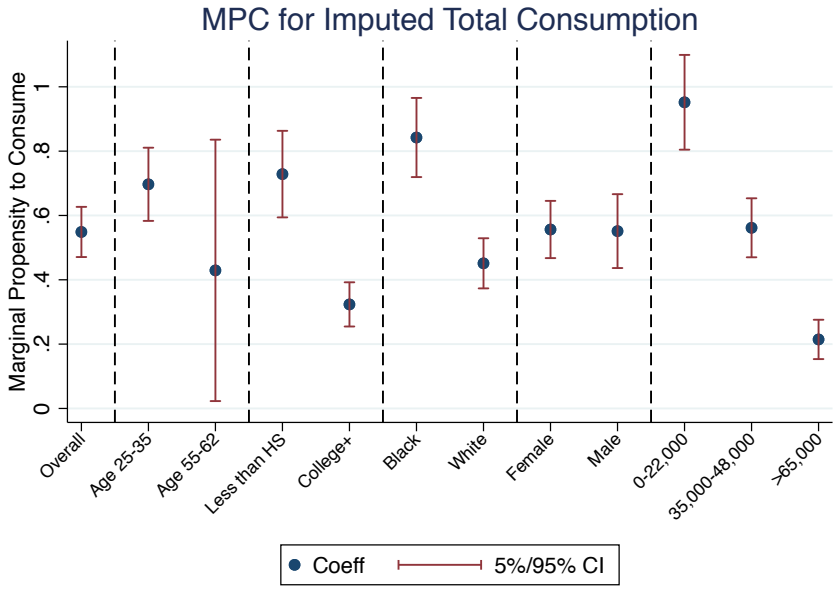

Fig. A1. Heterogeneity in estimated MPCs for total consumption across demographic groups.

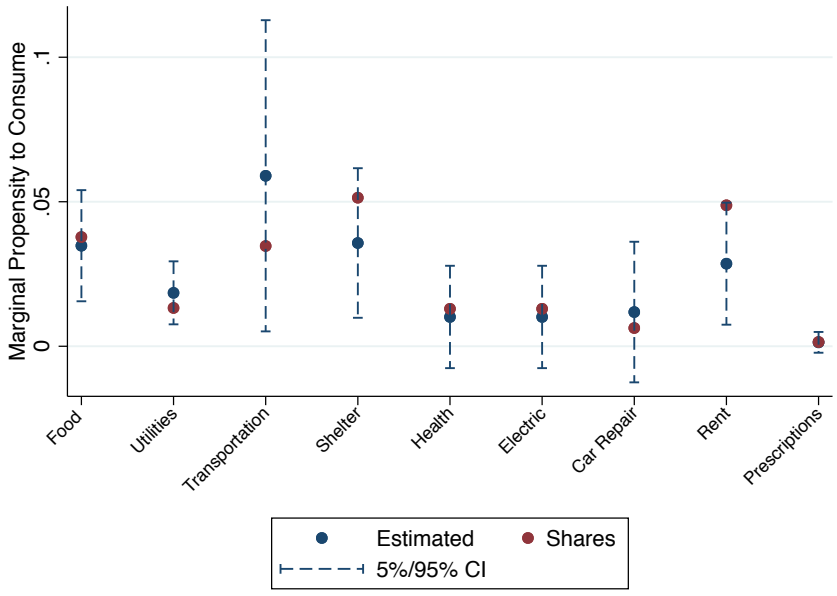

Fig. A2. Estimated Directed MPCs Vs. CEX basket-weighted MPCs

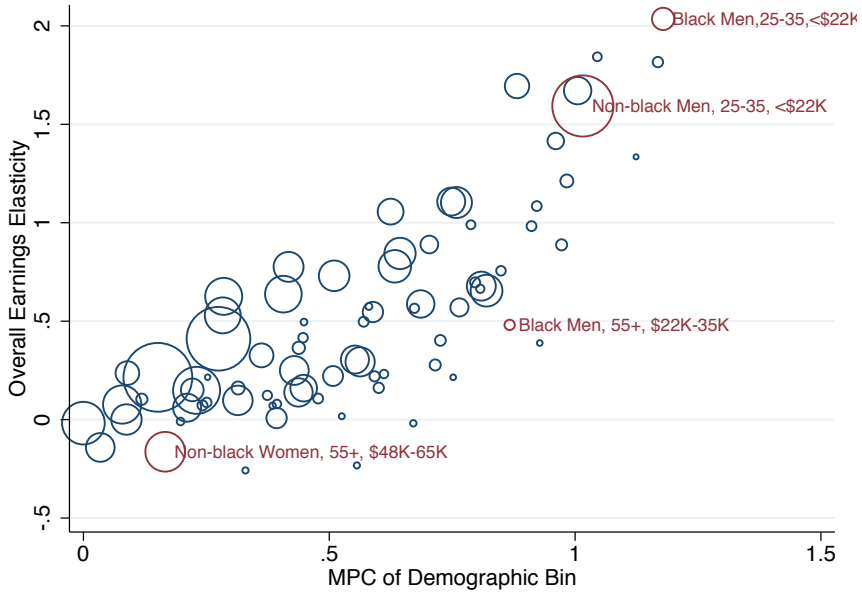

Fig. A3. Earnings elasticity to GDP shocks scattered against estimated MPC. See Patterson (2019) for more details.

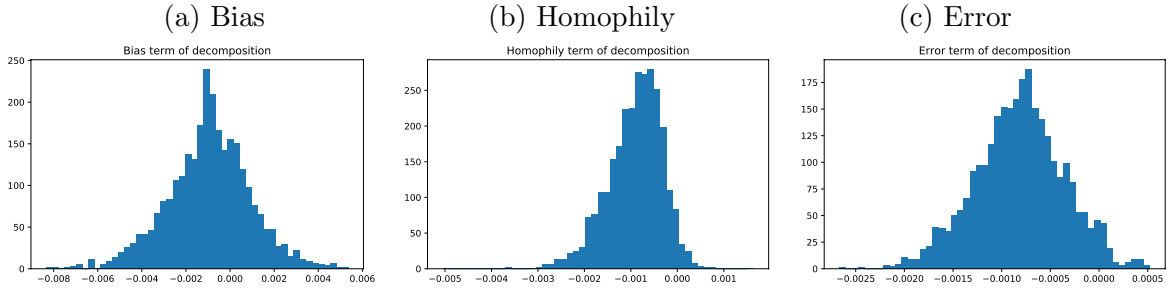

Fig. A4. Histograms of the bias term (left) and homophily term (middle) and overall error terms (right) from the decomposition in Proposition 2. For all subfigures, the distribution reflects a unit demand shock to each of the 2805 sector-region pairs, with baseline h^* given by the income incidence of a shock to demand proportional to 2012 state-industry GDP.

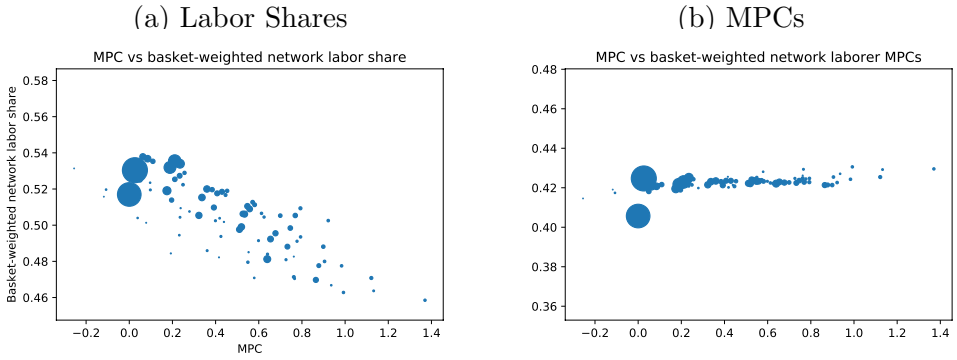

Fig. A5. The left panel shows the scatter plot of worker MPCs against the basket-weighted labor share of the sectors on which they consume. The right panel shows a scatter plot of worker MPCs against the basket-weighted MPCs of the labor employed in the sectors producing the goods they ultimately consume.

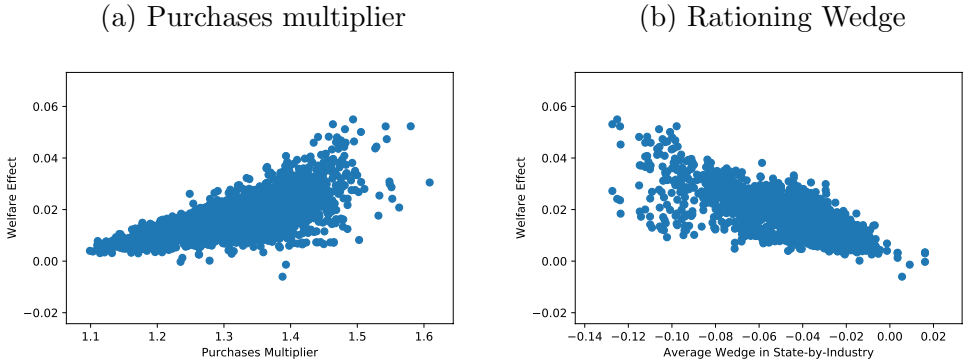

Fig. A6. Left panel: The x-axis gives the purchases multiplier for a dollar of government purchases targeting each of the 2805 state-industry pairs. The y-axis gives the estimated welfare effect of a dollar of government purchases targeting each of the 2805 state-industry pairs using rationing wedges from the Great Recession. Right panel: The x-axis gives the population-weighted Great Recession rationing wedge of employees in each of the 2805 state-industry pairs. The y-axis gives the estimated welfare effect of a dollar of government purchases targeting each of the 2805 state-industry pairs using rationing wedges from the Great Recession.

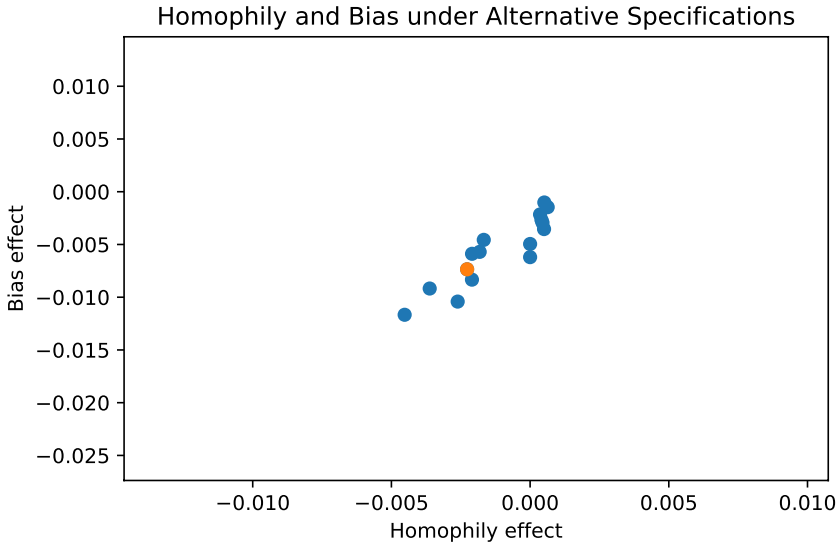

Fig. A7. Scatter plot of bias and homophily for a GDP-proportional shock in 16 alternative models where we compare all combinations of the modifications described in Section 5.2.2. These modifications adopt extreme alternative assumptions on regional trade, the input-output network, consumption heterogeneity, and income rationing. The orange dot corresponds to the baseline model. Here, the reference incidence h^* is that induced by a GDP-proportional shock.

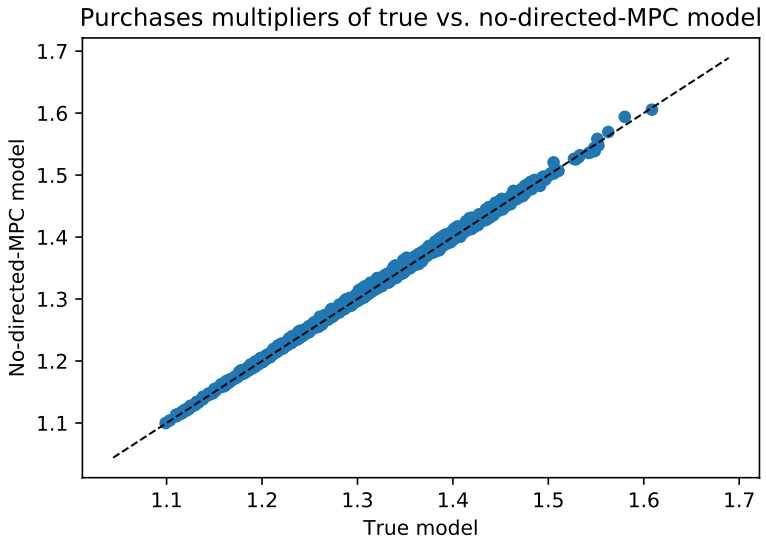

Fig. A8. Scatter plot of purchases multipliers for each of the 2805 industry-region pairs in the baseline model (x-axis) and the model in which all households have homogeneous consumption baskets in proportion to aggregate consumption (y-axis).

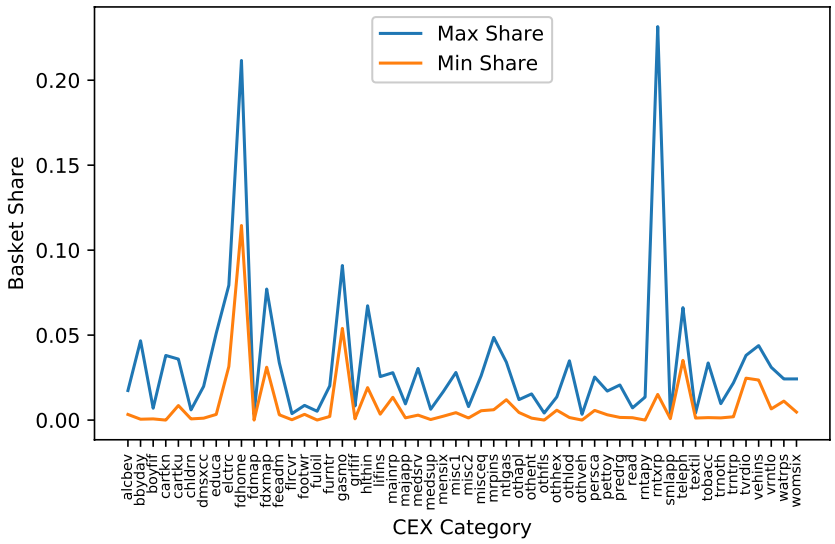

Fig. A9. Maximum and minimum consumption basket weights across all demographic groups in each CEX consumption category.

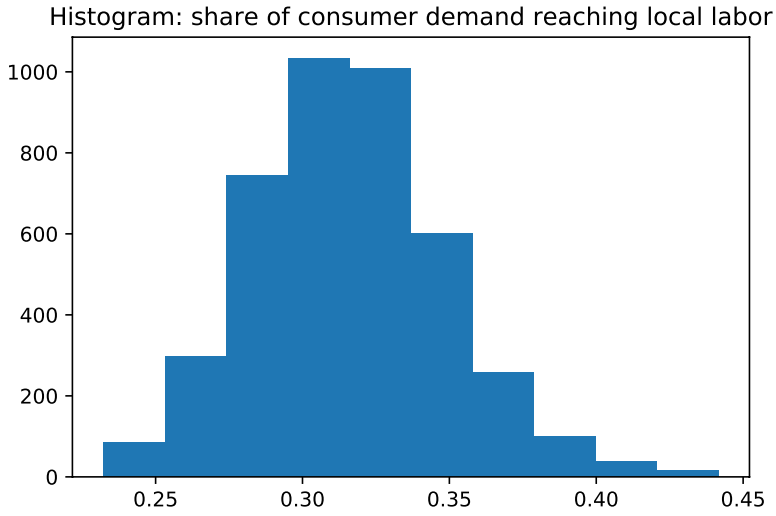

Fig. A10. Histogram of the fraction of consumer demand resulting in income for labor within the same state for each state-demographic pair.

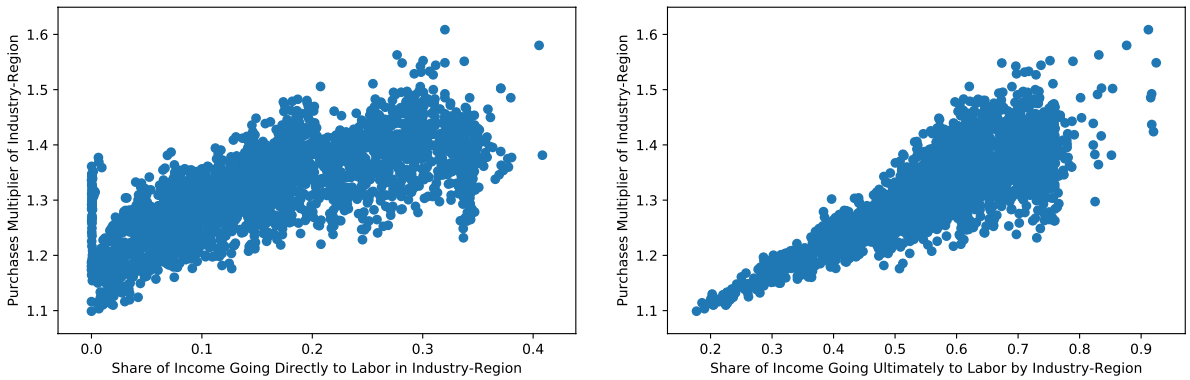

Fig. A11. Left Panel: Scatter plot of the change in GDP for each industry-region pair according to a one dollar demand shock in each pair against the share of income from production that goes directly to labor (as opposed to capital, foreigners, or inputs). Right Pannel: Scatter plot of the change in GDP for each industry-region pair according to a one dollar demand shock in each pair against the ultimate labor share accounting for labor employed in the production of intermediates.

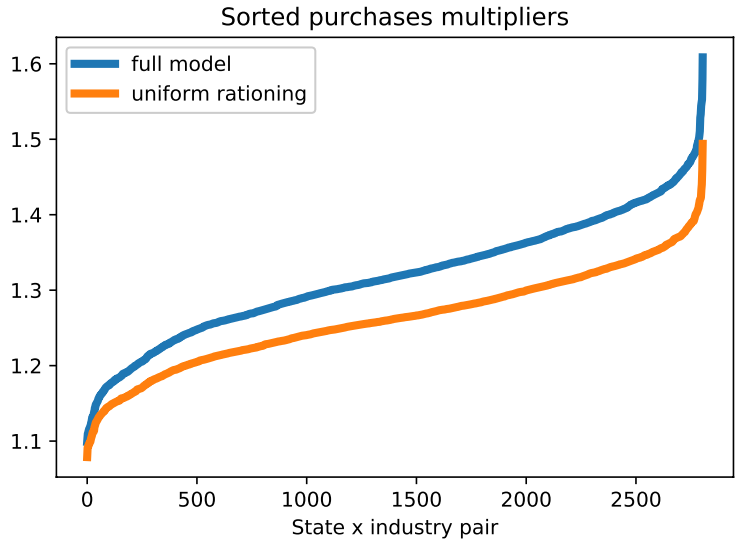

Fig. A12. Sorted change in GDP for each industry-region pair according to a one dollar demand shock in each pair. Full model is the baseline. Uniform rationing corresponds to all households' labor income being scaled in proportion to their labor income.

Fig. A13. Scatter plot of the change in GDP for each industry-region pair according to a one dollar demand shock in each pair. Full model is the baseline. Uniform rationing corresponds to all households' labor income being scaled in proportion to their income.

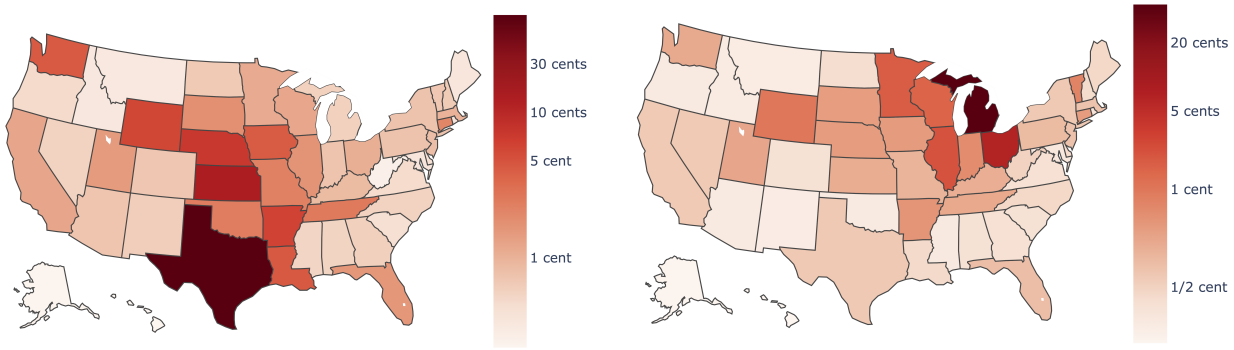

Fig. A14. Changes in state GDP per capita following a uniform $1-per-capita transfer shock to households in Texas (left panel) and Michigan (right panel).

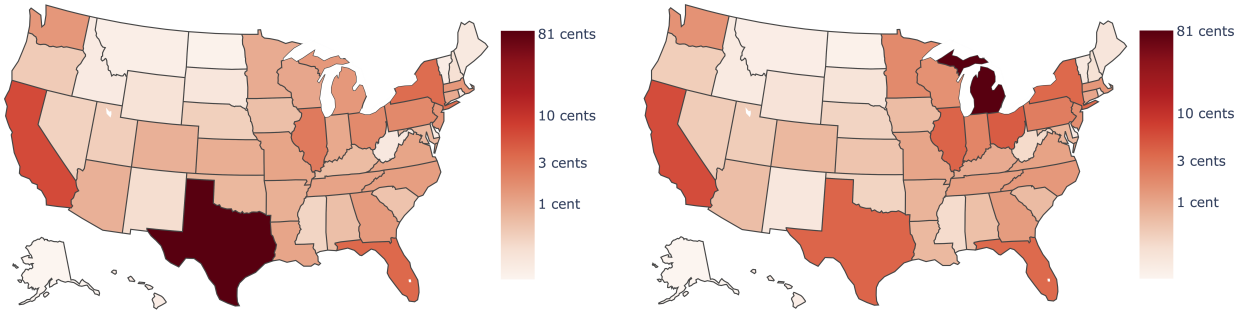

Fig. A15. Changes in state GDP following a GDP-proportional $1 purchases shock to Texas (left panel) and Michigan (right panel).

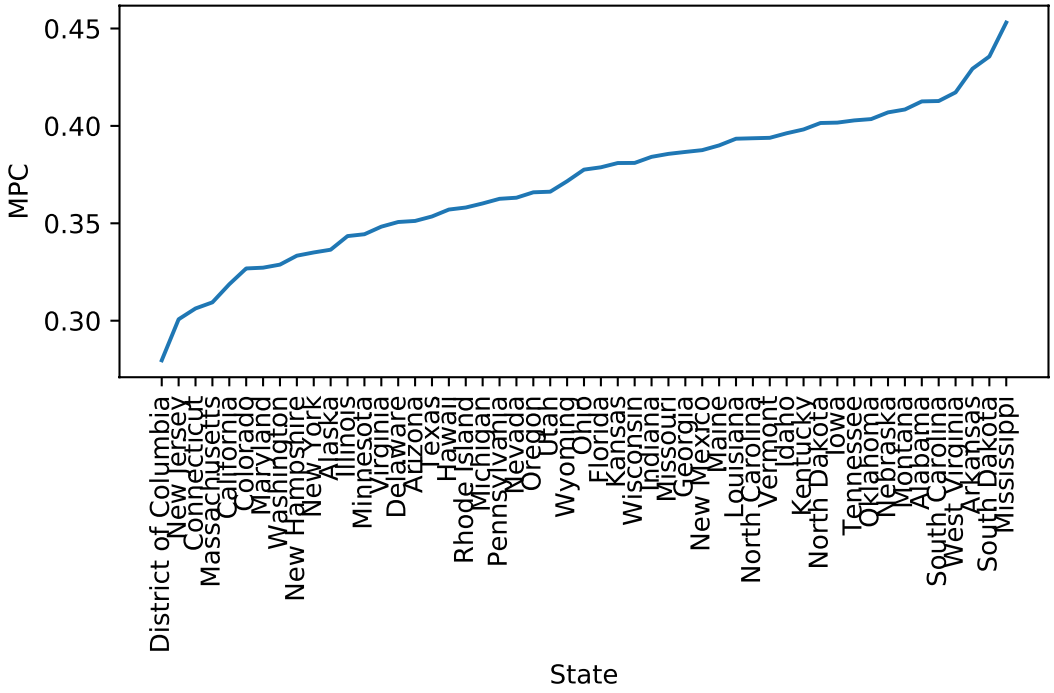

Fig. A16. Income-weighted average MPC by state.

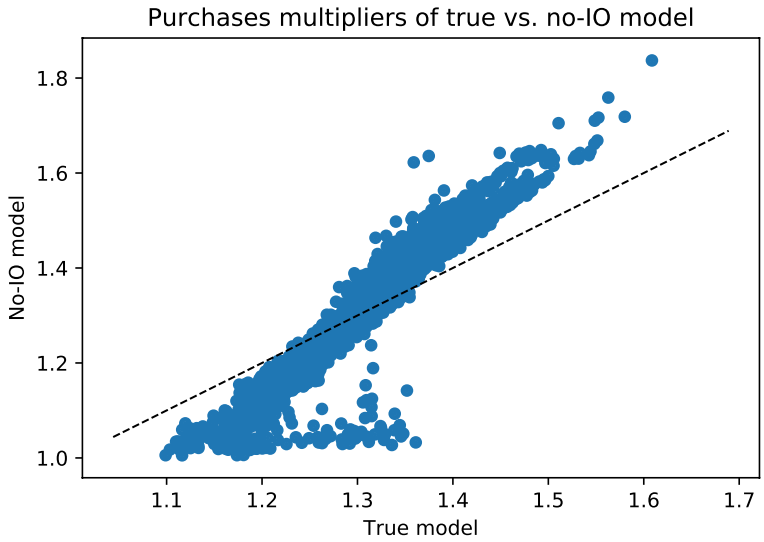

Fig. A17. Scatter plot of purchases multipliers for each of the 2805 industry-region pairs in the baseline model (x-axis) and the model in which there is no intermediate goods use by firms (y-axis).

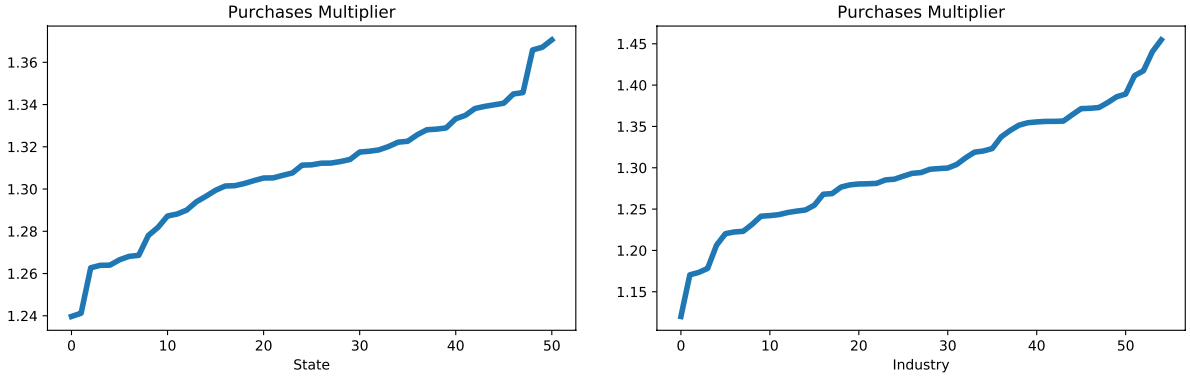

Fig. A18. Multipliers for state-level and industry level shocks. Formally, we take the shock for each state r as $\partial Q_r = \left(\mathbb{I}[s=r]\frac{y_{sj}}{\sum_k y_{rk}}\right)_{sj}$, where y_{rj} is BEA gross output for sector j in state r and each industry j as $\partial Q_j = \left(\mathbb{I}[k=j]\frac{y_{rk}}{\sum_s y_{sj}}\right)_{rk}$.

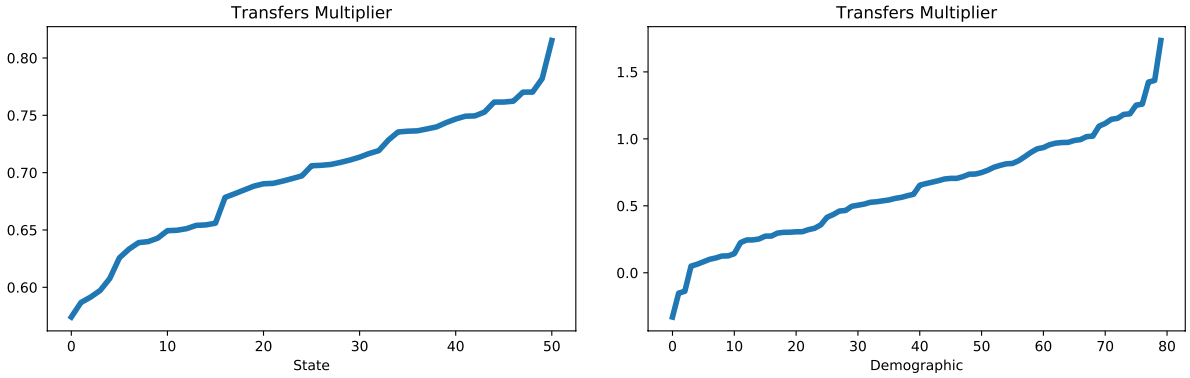

Fig. A19. Multipliers for state-level and demographic level transfers shocks. Formally, for the state-level shock, we transfer each state one dollar, in proportion to the demographic composition of that state. For the demographic-level shock, we transfer each demographic group one dollar, in proportion to the distribution of that demographic across states.

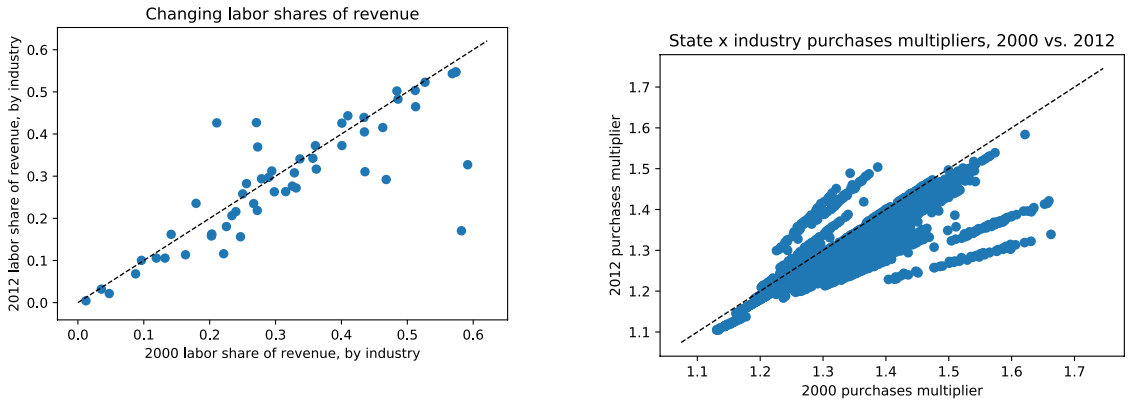

Fig. A20. Left Panel: Labor shares of revenue, by industry, in 2000 vs. 2012. Most industries experience a modest decline in labor share. The most dramatic decline is in the sector labelled "data processing, internet publishing, and other information services." The most dramatic increase is in the sector labelled "apparel and leather and allied products." Right Panel: Scatter plot of purchases multipliers in 2000 vs. 2012, by state-industry pair

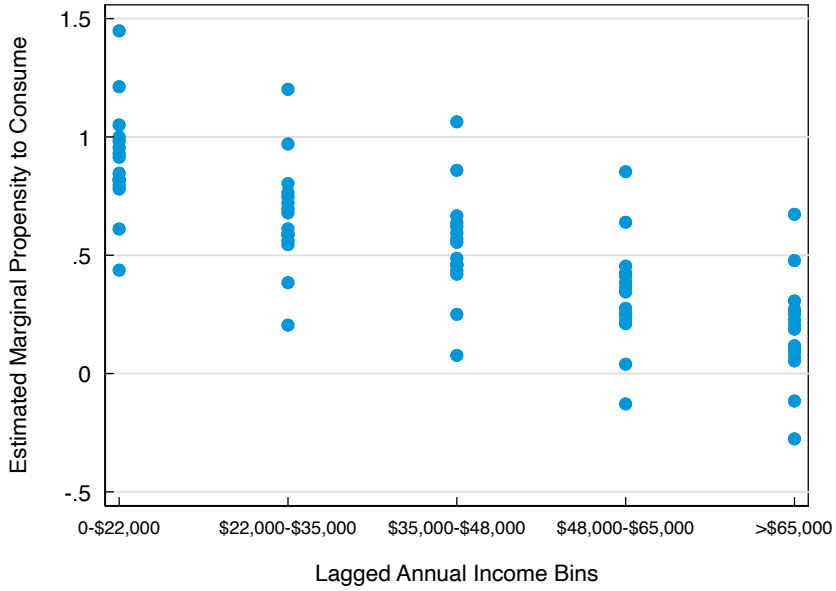

Fig. A21. Scatter plot of MPC vs. income for each age-sex-race demographic group.