UniCredit Foundation

The Missing Intercept: A Demand Equivalence Approach

Christian K. Wolf

2020 CopenhagenMacroDays Best Paper Award

Working Paper Series

n. 144 ■ September 2020

The Missing Intercept: A Demand Equivalence Approach

Christian K. Wolf[†]
University of Chicago

This version: October 9, 2020
First version: January 10, 2019

Abstract: I give conditions under which changes in private consumption demand are accommodated in general equilibrium exactly like changes in aggregate public spending. Under such demand equivalence, researchers can use the time series response of private consumption to fiscal purchases to recover the "missing general equilibrium intercept" of shocks to household spending identified in the cross section. I apply this method to deficit-financed tax rebates, and find (i) a large direct consumption demand response, and (ii) a fiscal multiplier of one and so a missing intercept close to zero. My theory and measurement also extend to shocks to investment demand.

[†]Becker Friedman Institute, University of Chicago, 5757 S. University Ave., Chicago, IL 60637 (email: ckwolf@uchicago.edu). This paper is a revised version of Chapter 1 of my PhD dissertation. I am very grateful to my advisors Ben Moll, Mikkel Plagborg-Møller, Gianluca Violante and Mark Watson for their continued guidance and support. I have particularly benefited from detailed comments by Mark Aguiar, Adrien Auclert, Markus Brunnermeier, Erik Hurst, Oleg Itskhoki, Greg Kaplan, Alisdair McKay, Emi Nakamura, Jonathan Payne, Monika Piazzesi, Richard Rogerson, Matt Rognlie, Martin Schneider, Jón Steinsson, Ludwig Straub, Chris Tonetti, Iván Werning, and Tom Winberry. I also thank seminar participants at several venues for their insights, and acknowledge generous support from the Alfred P. Sloan Foundation and the Macro Financial Modeling Project, as well as the Washington Center for Equitable Growth.

1 Introduction

A large literature in macroeconomics tries to estimate the aggregate effects of shocks to consumption and investment expenditure.[1] For most of these demand shifters, the experimental ideal – exogeneity at the macro level – is not attainable. In response, researchers increasingly leverage the cross-sectional variation available in micro data. Appealingly, because these estimates rely exclusively on cross-sectional information, they do not require macroeconomic identification restrictions. The well-known shortcoming is that such estimates are not interpretable as macro counterfactuals, simply because any potential general equilibrium effects – price changes, aggregate employment responses, tax financing, and so on – are differenced out. Previous work has tried to identify this "missing intercept" through fully specified structural models, with little systematic guidance on what model to choose, how to estimate it, and how to communicate uncertainty across the range of plausible models.

I develop an alternative *semi*-structural method, applicable to a general family of consumption and investment demand shifters. My method relies on an assumption that I refer to as "demand equivalence," positing that identical changes in private and public net excess demand also elicit identical general equilibrium feedback effects. I give a set of restrictions on economic primitives under which this equivalence holds, show that they are satisfied in a rich family of structural models, and document approximate equivalence in model extensions. The equivalence result is useful because it justifies a two-step empirical strategy: First, cross-sectional variation allows researchers to recover a shock's direct (or "partial equilibrium") effect on private spending. Second, the impulse response of private spending to a suitably chosen change in fiscal expenditure – estimable using time series methods – gives the missing general equilibrium effects. Summing (i) micro estimates and (ii) public expenditure impulse responses, researchers can recover the *full* effect of the demand shifter on private spending.

I illustrate my method with an application to income tax rebates. Micro data suggest a large but short-lived response of consumption, while time series evidence for a similarly short-lived, deficit-financed increase in public spending shows limited consumption crowding-out. Thus, under demand equivalence, the macro consumption counterfactual for a deficit-financed rebate is close to the micro estimate. For a structural model to be inconsistent with this near-zero intercept, it must either break equivalence or feature fiscal multipliers far from one. I find a similarly small intercept in a second application to investment tax cuts.

[1] Examples include tax rebates (Parker et al., 2013), redistribution (Jappelli & Pistaferri, 2014), credit tightening (Mian et al., 2013; Guerrieri & Lorenzoni, 2017), and bonus depreciation (Zwick & Mahon, 2017).

In the first part of the paper I give sufficient conditions for the equivalence of shocks to public spending and to private consumption. For a general class of business-cycle models, linearized impulse responses to macro shocks can be characterized implicitly as solutions to a linear infinite-horizon system of market-clearing conditions. The exclusion restriction required for shock equivalence is simply that two shocks perturb the same market-clearing conditions by the same amount; then, by the chain rule, the general equilibrium adjustment to these common perturbations must also be the same. I illustrate this abstract exclusion restriction through the lens of a rich heterogeneous-agent business-cycle model. Here, the equivalence of shocks to private consumption and to public spending maps into three main assumptions. First, households and government need to consume the same final good. If so, identical changes in private or public spending lead to identical excess demand for that common good. Second, households and government must borrow and lend at the same interest rate. The identical expansions in private and public demand can then be discounted at that common rate, and so can be financed using identical paths of taxes and transfers. Third, household labor supply must not respond differentially to the two shocks; sufficient conditions are either the absence of wealth effects in labor supply or fully demand-determined employment. Conditional on these restrictions, all other details of the model – preferences, technology, expectation formation, and so on – are irrelevant: shocks to consumption demand are invariably accommodated in general equilibrium exactly like changes in public spending.

I leverage the demand equivalence result to justify my two-step procedure. First, a researcher runs cross-sectional regressions of household consumption on a shifter of consumption demand (e.g., income tax rebate). If cross-sectional heterogeneity in shock exposure is independent of household characteristics, then the econometric estimand of these regressions is interpretable as the direct (or partial equilibrium) response path of consumption to the shifter. Second, invoking demand equivalence, we know that a suitably chosen public spending shock will induce the same general equilibrium effects. If such a shock can be identified using the macroeconometric toolkit for fiscal spending (e.g. Ramey, 2018), researchers can use the estimated time series response of private consumption to recover the path of missing intercepts for the cross-sectional regressions. The pointwise sum of cross-sectional and time series estimates is then interpretable as a valid semi-structural consumption counterfactual for the private demand shifter. Equivalently, the researcher could have written down *any* particular parametric model in the equivalence class, parameterized the model to be consistent with the estimated micro and macro moments, and solved it – the equivalence result guarantees that she would have recovered the same counterfactual as my simple sum.

I demonstrate the feasibility of the method through the study of a popular consumption stimulus policy: an income tax rebate. I first review previous empirical work (Parker et al., 2013; Jappelli & Pistaferri, 2014) and show that the direct consumption demand response to the stimulus is indeed either equal or at least tightly linked to the econometric estimands of those studies. Their different experiments consistently paint the picture of a large but short-lived expansion in spending. Next, by demand equivalence, I can recover the missing general equilibrium effects as the response of consumption to a similarly short-lived expansion in government spending. Thus, for the second step, I trace out the dynamic effects of a particular, highly transitory innovation to government spending: professional forecast errors. Following Plagborg-Møller & Wolf (2019), I project on these forecast errors using a recursive Vector Autoregression (VAR). I find that the short-lived spending expansion is persistently deficit-financed, so the results of my procedure are interpretable as pertaining to a similarly deficit-financed rebate. In response to the expansion, output rises one-to-one with spending, and consumption is largely flat. Summing the micro and macro estimates, I conclude that a one-off, deficit-financed transfer briefly but significantly stimulates aggregate consumption, with the overall response close to the direct effect estimated using micro data.

While the output of the two-step procedure is only exactly interpretable as a valid counterfactual if the underlying data-generating process satisfies my exclusion restrictions, I show that the method continues to perform well on artificial data generated from a variety of richer business-cycle models. My main laboratory is a heterogeneous-agent New Keynesian ("HANK") model, estimated to be consistent with salient features of cross-sectional earnings risk, the aggregate wealth distribution, and the time series distribution of macro aggregates. In this model, demand equivalence fails *only* because of short-term wealth effects in labor supply, resulting in the output of my two-step procedure to be biased upward. However, consistent with both micro evidence (Cesarini et al., 2017) and the results of much previous structural modeling (e.g. Christiano, 2011a), the inaccuracy resulting from this labor supply channel is quantitatively negligible, at least for transitory shocks. I also sign the bias and discuss the inaccuracy induced by several further model extensions. These include: a multi-asset structure, productive benefits and consumption complementarities for public spending, and imperfect factor mobility across sectors with heterogeneous production functions.

My methodology extends with little change to shifters of investment demand. In response to the shock, investment increases today (excess demand) while capital and so production build up gradually (excess supply). I give sufficient conditions under which the investment demand shifter is accommodated in general equilibrium exactly like an expansion in government expenditure today (excess demand), financed by a contraction in the future (excess supply). Importantly, these conditions impose no material restrictions on the production block of the economy; in particular, investment demand equivalence holds in most recent quantitative studies of the aggregate effects of firm-level investment frictions, including models with very rich firm heterogeneity.[2] Finally, I apply my results to study the aggregate effects of bonus depreciation stimulus: I find that the policy induces a large increase in investment demand (Zwick & Mahon, 2017; Koby & Wolf, 2020), accommodated in general equilibrium through a sharp rise in output, with little investment crowding-out.

Before proceeding further, I briefly comment on the scope and limitations of my analysis. First, my methodology requires first-stage micro regressions whose estimands are interpretable as free of any general equilibrium interactions. This is arguably the case for across-household or across-firm regressions, but not for cross-*regional* regressions (Mian et al., 2013; Mian & Sufi, 2014).[3] Second, I have demonstrated that my approach is feasible in practice for some much-studied shocks to consumption and investment demand. The general principle – to leverage macro evidence on the general equilibrium propagation of plausibly equivalent shocks – may allow researchers to construct counterfactuals for other shocks and policies, but those are beyond the scope of this paper. Third, while my equivalence results are only valid to first order, I impose no restriction on *where* the underlying Taylor series approximation is taken. Evidence on state dependence in the transmission of fiscal shocks thus applies without change to generic consumption and investment demand shifters. Finally, my two-step procedure relies sensitively on the assumption that all agents only interact through (a small set of) aggregate prices and quantities. Appealingly, I take little stand on the precise nature of that interaction, so my theory covers both conventional neoclassical as well as quite different Keynesian adjustment mechanisms. Less appealingly, strategic interaction between agents invariably breaks the neat separation into direct spending responses and general equilibrium accommodation that lies at the heart of my approach.

Literature. This paper relates and contributes to several strands of literature.

First, my methodology connects two empirical literatures. A fast-growing line of work uses variation at the individual or regional level to estimate spending responses to policy changes and other macro shocks (e.g. Mian & Sufi, 2009; Parker et al., 2013; Jappelli & Pistaferri, 2014; Zwick & Mahon, 2017). As all of these studies control for macro fluctuations

[2]For example, exact equivalence applies in the models of Khan & Thomas (2013) and Winberry (2018).

[3]I generalize my results to such cross-regional regressions in a companion note (Wolf, 2019).

through time fixed effects, they are silent on any possible general equilibrium feedback. I argue that a second literature – that on the aggregate effects of changes in government spending – can be informative about this "missing intercept." Comprehensive literature summaries are Hall (2009) and Ramey (2018); overall, earlier empirical work largely estimates output multipliers around 1, and zero (or slightly negative) responses of private spending.

Second, the demand equivalence result builds on the simple Keynesian cross intuition of a common "demand multiplier," prevalent in government policy evaluation and professional forecasting (e.g. Reichling & Whalen, 2012). Auclert et al. (2018) formalize this intuition to show that, in models with demand-determined labor, passive monetary policy, and without investment, consumption demand and government spending shocks have identical effects on aggregate output. I extend the equivalence result to a larger family of models (and to investment), find support for approximate equivalence in micro and macro data, and measure the common general equilibrium effects through macro quasi-experiments. In contemporaneous work, Guren et al. (2020) and Chodorow-Reich et al. (2019) use the reverse logic to strip out the *local* general equilibrium effects present in cross-regional regression estimates.

Third, my method naturally complements existing approaches to the estimation of counterfactuals in macroeconomics. In its reliance on general exclusion restrictions rather than parametric structural models, it is semi-structural in exactly the same way as conventional Structural Vector Autoregressive (SVAR) analysis (Sims, 1980). However, the equivalence result also provides a novel justification for impulse response matching as a limited-information approach to the estimation of fully structural models: By commonality of general equilibrium effects, impulse responses to particular aggregate structural shocks are robustly informative for many different counterfactuals. In my applications, the targeted moments are in fact *fully* informative, obviating the need for model solution (Andrews et al., 2018). This "sufficient statistics" approach is common in public finance (Chetty, 2009), and also increasingly widespread in macroeconomics (Nakamura & Steinsson, 2018).

Outline. Section 2 establishes the consumption demand equivalence result. In Section 3, I leverage commonality in general equilibrium propagation to propose a two-step procedure for estimation of consumption demand counterfactuals, with an application to income tax rebates. Section 4 then shows that the proposed approximation remains accurate in more general environments. The generalization to investment demand, including an application to investment tax stimulus, is discussed in Section 5. Section 6 concludes, and supplementary details, proofs and a third application are all relegated to several appendices.

2 Consumption demand equivalence

This section develops an exact equivalence result for the general equilibrium propagation of shocks to private consumption demand and to public spending. In Sections 2.1 to 2.4, I discuss the restrictions on economic primitives needed for demand equivalence in a general quantitative business-cycle model. Section 2.5 complements this model-based analysis with an abstract formulation of shock equivalence as a set of exclusion restrictions on (linearized) aggregate equilibrium conditions.

2.1 The benchmark model

Time is discrete and runs forever, $t = 0, 1, \ldots$. The model economy is populated by households, firms, and a government. There is no aggregate uncertainty, but households and firms are allowed to face idiosyncratic risk. I study perfect foresight transition paths back to steady state after one-time unexpected aggregate innovations at time 0; for vanishingly small innovations, these transition paths are mathematically equivalent to standard impulse response functions computed from the first-order perturbation solution to an otherwise identical model with aggregate risk.[4] Anticipating my main empirical application, I will focus on two such innovations: first, a one-off transfer to households, and second, a transitory expansion in government spending. Section 2.5 shows how the equivalence result extends to generic shifters of consumption demand (e.g., changes in borrowing constraints, redistribution, . . .).

Notation. The realization of a variable x at time t along the equilibrium perfect foresight transition path will be denoted x_t, while the entire time path will be denoted $\mathbf{x} = \{x_t\}_{t=0}^{\infty}$. Hats denote deviations from the deterministic steady state, bars denote steady-state values, and tildes denote logs. I study two structural shocks indexed by $s \in \{\tau, g\}$ – tax rebates and government spending. I write individual shock paths as $\boldsymbol{\varepsilon}_s$, and use subscripts $\boldsymbol{\varepsilon}$ for transitions after a generic path $\boldsymbol{\varepsilon} \equiv (\boldsymbol{\varepsilon}_\tau', \boldsymbol{\varepsilon}_g')'$. I reserve the simpler s subscripts for one-time single shocks – that is, shock paths with $\varepsilon_{s,0} = 1$ and $\varepsilon_{u,\tau} = 0$ for $(u, \tau) \neq (s, 0)$.

Households. A unit continuum of households $i \in [0, 1]$ has preferences over consumption c_{it} and labor ℓ_{it}. They are subject to idiosyncratic productivity risk e_{it} and potentially differ

[4] This result is an implication of certainty equivalence coupled with Taylor's theorem (Boppart et al., 2018). For ordinary business-cycle fluctuations, such first-order perturbations offer an accurate characterization of the model's global dynamics (e.g. Fernández-Villaverde et al., 2016; Ahn et al., 2017; Auclert et al., 2019).

in their baseline discount factor β_i. Households can self-insure by investing in liquid nominal bonds b^h_{it}, with nominal returns i^b_t and subject to a borrowing constraint $\underline{b}$. Income consists of labor earnings as well as (potentially type-specific) lump-sum rebates τ_{it} and dividend income d_{it}. Total hours worked ℓ_{it} are determined by demands of a unit continuum $k \in [0,1]$ of price-setting labor unions, as in Erceg et al. (2000); the problem of labor unions will be considered later. Given a path of prices, rebates, dividends, hours worked and inflation (π_t), the consumption-savings problem of household i is thus

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

such that

$$c_{it} + b^h_{it} = (1 - \tau_\ell) w_t e_{it} \ell_{it} + \frac{1 + i^b_{t-1}}{1 + \pi_t} b^h_{it-1} + \tau_{it} + d_{it}$$

and

$$b^h_{it} \geq \underline{b}$$

Labor productivity e_{it} follows a (stochastic) law of motion with $\int_i e_{it} di = 1$ at all times.

Labor unions behave as in conventional New Keynesian models (Erceg et al., 2000; Auclert et al., 2018). Worker i provides ℓ_{ikt} units of labor to union k, giving total hours worked for household i of $\ell_{it} \equiv \int_k \ell_{ikt} dk$. The total effective amount of labor intermediated by union k is $\ell_{kt} \equiv \int_i e_{it} \ell_{ikt} di$; each union then sells its labor services to a competitive labor packer at price w_{kt}. The labor packer aggregates union-specific labor to aggregate labor services,

$$\ell^h_t \equiv \left(\int_k \ell_{kt}^{\frac{\varepsilon_w - 1}{\varepsilon_w}} dk \right)^{\frac{\varepsilon_w}{\varepsilon_w - 1}}$$

sold at the aggregate wage index w_t, and where ε_w denotes the elasticity of substitution between different types of labor. Union k chooses its wage rate w_{kt} subject to wage-setting adjustment costs, and satisfies the corresponding demand for its labor services. I assume that it does so by demanding a common amount of hours worked from its members.[5] Since the wage-setting problem is standard, I relegate details to Appendix B.1. For the purposes of the analysis here, it suffices to note that union behavior can be summarized through a simple wage New Keynesian Phillips curve – effectively, an aggregate labor supply relation.

[5]A uniform hiring rule is the natural assumption in sticky-wage heterogeneous-household models, but is of course awkward in the flexible-wage limit, as it then does not nest the alternative natural case of flexible labor supply for each individual household. I consider a model without unions in Appendix E.3.

Fiscal policy. The fiscal authority consumes the same final good as households. Fiscal consumption g_t and total lump-sum transfers $\tau_t \equiv \int_0^1 \tau_{it} di$ are financed through debt issuance and taxes on labor income. The government budget flow constraint is

$$\frac{1+i^b_{t-1}}{1+\pi_t} b_{t-1} + g_t + \tau_t = \tau_\ell w_t \ell_t + b_t$$

I assume that total government spending $\mathbf{g} = \mathbf{g}(\boldsymbol{\varepsilon}_g)$ follows some exogenous process, and that the government freely sets a discretionary part of tax rebates $\boldsymbol{\tau}_x = \boldsymbol{\tau}_x(\boldsymbol{\varepsilon}_\tau)$. Given paths for spending targets $(\boldsymbol{\varepsilon}_g, \boldsymbol{\varepsilon}_\tau)$, initial nominal debt b_{-1} and a path of prices and quantities $(\mathbf{w}, \boldsymbol{\ell}, \mathbf{i}^b, \boldsymbol{\pi})$, a government debt financing rule is a path $\boldsymbol{\tau}_e$ such that $\boldsymbol{\tau} = \boldsymbol{\tau}_e + \boldsymbol{\tau}_x$, the flow government budget constraint holds at all periods t, and $\lim_{t\to\infty} \left(\prod_{s=0}^t \frac{1+\pi_s}{1+i^b_{s-1}}\right) b_t = 0$.

Rest of the economy. Since my focus is on the equivalence of private and public expansions in demand, I only sketch the rest of the model, with a detailed outline provided in Appendix B.1. The corporate sector is populated by three sets of firms: a unit continuum of heterogeneous, perfectly competitive intermediate goods producers j, a unit continuum of monopolistically competitive retailers with nominal price rigidities, and a final goods aggregator. Intermediate goods producers accumulate capital, hire labor, issue risk-free debt, and sell their composite intermediate good, possibly subject to (both convex and non-convex) capital adjustment costs as well as generic constraints on equity and debt issuance. Retailers purchase the intermediate good, costlessly differentiate, monopolistically set prices, and sell their differentiated good on to the competitive aggregator.

The last remaining entity in the model is the monetary authority. This monetary authority sets nominal rates on liquid bonds $\mathbf{i}^b$ in accordance with a conventional Taylor rule.

Equilibrium. I assume that there exists a unique deterministic steady state.[6] To allow interpretation of perfect foresight transition paths as conventional first-order perturbation solutions, I impose that the economy is indeed initially in steady state, and then study perfect foresight transition equilibria back to the initial deterministic steady state. The definition of equilibrium perfect foresight transition paths is then standard (see Appendix B.1); I discuss an extension to transition paths with other starting points in Appendix C.1.

[6] More precisely, I make implicit assumptions on functional forms and parameter values that guarantee that there is a unique deterministic steady state. In all numerical exercises, I have verified the uniqueness of the steady state and the (local) existence and uniqueness of transition paths.

Nested Models. My benchmark model is designed to nest several important earlier contributions to quantitative business-cycle analysis. In the absence of uninsurable household earnings risk and household borrowing limits, and without firm-level productivity differences and financial frictions, it becomes a standard New Keynesian model (e.g. Smets & Wouters, 2007). However, the environment is also rich enough to allow for non-trivial micro heterogeneity at the household and firm level. On the household side, income risk and limited self-insurance can endogenously generate hand-to-mouth behavior. With flexible prices, the model is identical to Aiyagari (1994) or Krusell & Smith (1998); with nominal rigidities, it is a HANK model in the mold of McKay et al. (2016) and Guerrieri & Lorenzoni (2017). On the firm side, I allow for a rich set of real and financial frictions to the capital accumulation process, as for example in Khan & Thomas (2008), Khan & Thomas (2013) and Winberry (2018). In other words, the benchmark model is as rich as most models that – in the absence of the identification results developed here – would be used to structurally pin down the missing general equilibrium intercept of, say, income tax rebate shocks.

2.2 Direct responses and general equilibrium feedback

The demand equivalence result will assert a commonality in the general equilibrium propagation of different shocks. A precise statement of such equivalence requires a formal definition of direct (or partial equilibrium) responses and indirect (or general equilibrium) effects.

I assume that the consumption-savings problem (1) has a unique solution for any path of prices, quantities and shocks faced by households. Aggregating the solutions across households, we obtain an aggregate consumption function $\mathbf{c} = \mathbf{c}(\mathbf{s}^h; \boldsymbol{\varepsilon})$, where $\mathbf{s}^h = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \boldsymbol{\ell}, \boldsymbol{\tau}_e, \mathbf{d})$ collects household income and saving returns – objects that adjust in general equilibrium. The total impulse response of consumption to the shock path $\boldsymbol{\varepsilon}$ is simply

$$\hat{\mathbf{c}}_\varepsilon \quad \equiv \quad \mathbf{c}(\mathbf{s}^h_\varepsilon; \boldsymbol{\varepsilon}) \; - \; \mathbf{c}(\bar{\mathbf{s}}^h; \mathbf{0})$$

I decompose this aggregate impulse response into two parts: a direct "partial equilibrium" impulse and an indirect "general equilibrium" feedback part.[7]

[7] My definition of the partial equilibrium consumption response abstracts from endogenous adjustments in earnings. I do so for three reasons. First, many empirical estimates of household spending responses to sudden income changes are actually interpretable as such netted spending elasticities (e.g. see Auclert, 2019, footnote 34). Second, in models with union-intermediated labor supply – like the one considered here –, replicating cross-sectional micro regressions invariably differences out labor responses (see Proposition 3). Third, microeconomic evidence suggests that short-run wealth effects are very weak anyway (Cesarini et al.,

Definition 1. *Let the direct (partial equilibrium) response of consumption to a shock path* $\boldsymbol{\varepsilon}$ *be defined as*

$$\hat{\mathbf{c}}_{\varepsilon}^{PE} \quad \equiv \quad \mathbf{c}(\bar{\mathbf{s}}^h; \boldsymbol{\varepsilon}) \;-\; \mathbf{c}(\bar{\mathbf{s}}^h; \mathbf{0}) \tag{2}$$

Similarly, let the indirect (general equilibrium) feedback be

$$\hat{\mathbf{c}}_{\varepsilon}^{GE} \quad \equiv \quad \mathbf{c}(\mathbf{s}_{\varepsilon}^h; \mathbf{0}) \;-\; \mathbf{c}(\bar{\mathbf{s}}^h; \mathbf{0}) \tag{3}$$

It is immediate that, to first order, the aggregate impulse response admits a simple additive decomposition into partial equilibrium response and general equilibrium feedback:

$$\hat{\mathbf{c}}_{\varepsilon} \quad = \quad \hat{\mathbf{c}}_{\varepsilon}^{PE} \;+\; \hat{\mathbf{c}}_{\varepsilon}^{GE} \tag{4}$$

The decomposition (4) is only interesting to the extent that its components can be tied to empirically measurable objects. The remainder of this section establishes conditions under which the consumption response to particular government spending shocks $\hat{\mathbf{c}}_g$ is informative about the general equilibrium feedback term $\hat{\mathbf{c}}_{\tau}^{GE}$ of lump-sum transfers – the demand equivalence result. In Section 3 I then argue that (i) cross-sectional regressions estimate the direct spending response $\hat{\mathbf{c}}_{\tau}^{PE}$ and (ii) it is in practice often possible to recover the aggregate effects of the particular public spending shocks that can proxy for $\hat{\mathbf{c}}_{\tau}^{GE}$.

2.3 A simple example of demand equivalence

I first illustrate the intuition for the demand equivalence result using a very simple special case of my benchmark model – a spender-saver real business-cycle (RBC) model. A mass λ of households are spenders (so $\beta_i = 0$), while the remaining households are savers ($\beta_i > 0$). Both types have log consumption utility and inelastically supply their labor endowment, and savers hold all risk-free real bonds and receive firm dividends. The firm sector admits aggregation to a representative firm which hires labor and accumulates capital; for simplicity I assume that capital depreciates fully within the period and that the production function is Cobb-Douglas, $y = k^{\alpha}\ell^{1-\alpha}$. The fiscal authority issues risk-free bonds, consumes the final good, and imposes (different) lump-sum taxes on savers and spenders. There are no nominal rigidities, so central bank behavior is irrelevant for all real quantities.

2017; Fagereng et al., 2018). Nevertheless, in Appendix E.3, I repeat all of my analysis in an alternative model without unions, but with a non-standard preference parameterization allowing for (data-consistent) weak short-run wealth effects (Jaimovich & Rebelo, 2009; Galí et al., 2012).

In this environment I compare the transmission of two structural shocks: (i) a one-off income tax rebate $\varepsilon_{\tau t}$ (to spenders) and (ii) a one-period expansion in aggregate government spending ε_{gt}, both worth one per cent of total steady-state consumption. I assume that the tax increases (transfer cuts) $\hat{\boldsymbol{\tau}}_e$ used to finance the two policies exclusively fall on savers. All model equations are stated in Appendix B.2.

Demand Equivalence. I begin with a concrete numerical example. I set the saver discount factor to $\beta = 0.99$, the capital share to $\alpha = 1/3$, and assume that a mass $\lambda = 0.3$ of households is hand-to-mouth. Figure 1 shows consumption impulse responses for one-period tax rebate and government spending shocks.

Demand Equivalence, Spender-Saver RBC Model

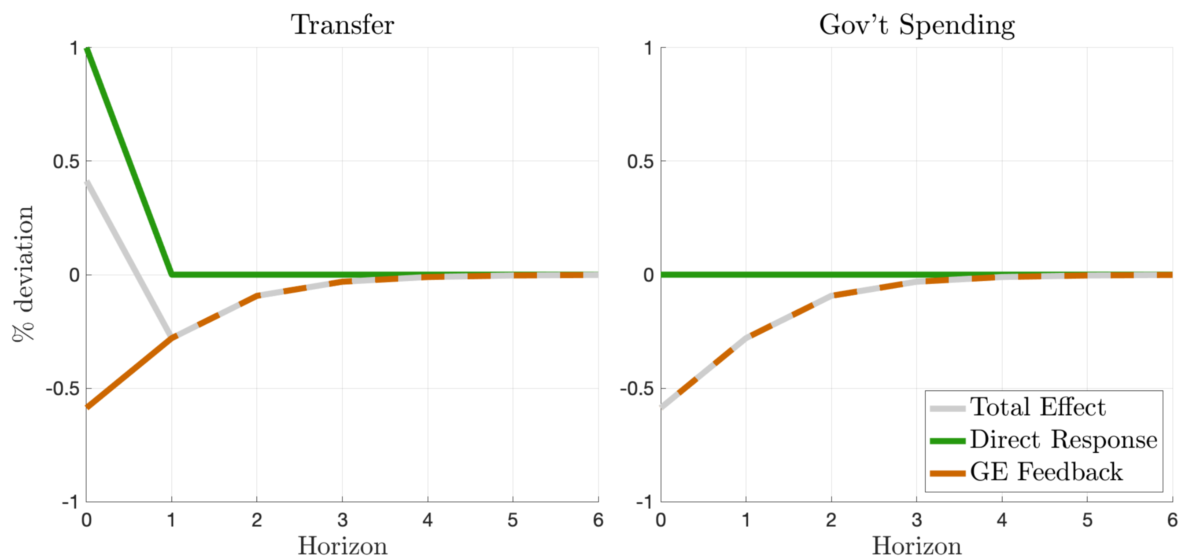

Figure 1: Consumption impulse response decompositions after equally large, one-off tax rebate and government spending shocks in the simple spender-saver RBC model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

The left panel shows the consumption response to the one-off transfer. In line with Definition 1, this aggregate impulse response is decomposed into direct partial equilibrium (green) and indirect general equilibrium (orange) responses. By assumption, spenders consume all of the rebate today. The grey line then shows that, after general equilibrium price adjustments, aggregate consumption only moderately rises on impact, then falls, and gradually returns to steady state. General equilibrium adjustment thus substantially crowds-out consumption. Intuitively, this is so because a rise in interest rates leads savers to postpone consumption; at the same time, investment is crowded out, so future output drops and income declines.

The right panel shows the consumption response to an increase in government spending. By definition, consumption does not respond directly to this second shock (green line). In general equilibrium, however, it drops substantially; exactly as for the tax rebate, this is largely due to higher rates crowding out both saver consumption and aggregate investment, and thus further pushing down future income. Crucially, the response of aggregate consumption to the public spending shock is identical to the general feedback associated with the tax rebate shock – a property of the model that I will refer to as "demand equivalence." As it turns out, demand equivalence is not an artifact of the particular parameterization chosen for Figure 1, but a general feature of the simple spender-saver model.

Proposition 1. *Suppose that* $\hat{\mathbf{c}}_\tau^{PE} = \hat{\mathbf{g}}_g$. *Then*

$$\hat{\mathbf{c}}_\tau^{GE} \quad = \quad \hat{\mathbf{c}}_g^{GE} \tag{5}$$

and so

$$\hat{\mathbf{c}}_\tau \quad = \quad \underbrace{\hat{\mathbf{c}}_\tau^{PE}}_{PE\ response} \quad + \quad \underbrace{\hat{\mathbf{c}}_g}_{=\ GE\ feedback} \tag{6}$$

Irrespective of the parameterization, the *total* response of consumption to a fiscal shock can proxy for the general equilibrium effects induced by the private spending change.

Intuition & proof sketch. The intuition for demand equivalence in the simple spender-saver model is entirely straightforward – it does not matter whether more final output is demanded for consumption by the government or by hand-to-mouth households. Either way, consumption of savers has to adjust to ensure market-clearing.

Proposition 1 follows immediately from inspection of the model's closed-form solution. To set the stage for the general equivalence results in Section 2.4 and Section 2.5, however, it will prove more instructive to follow an alternative, more readily generalizable proof strategy: writing the equilibrium as a system of market-clearing conditions and prices adjusting to clear those markets. For the spender-saver model, it is straightforward to show (see Lemma G.1) that a path of real interest rates $\mathbf{r}$ and lump-sum taxes on savers $-\boldsymbol{\tau}^e$ is part of a perfect foresight equilibrium if and only if

$$\mathbf{c}(\mathbf{r}, \mathbf{w}(\mathbf{r}), \mathbf{d}(\mathbf{r}), \boldsymbol{\tau}_e; \boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) \quad = \quad \mathbf{y}(\mathbf{r}) - \mathbf{i}(\mathbf{r}) \tag{7}$$

$$\boldsymbol{\tau}_e \quad = \quad \boldsymbol{\tau}_e(\boldsymbol{\varepsilon}) \tag{8}$$

where $\mathbf{y}(\bullet)$ and $\mathbf{i}(\bullet)$ are firm policy functions, optimal firm behavior implicitly pins down wages $\mathbf{w}(\bullet)$ and dividends $\mathbf{d}(\bullet)$ as functions of $\mathbf{r}$, and $\mathbf{c}(\bullet)$ aggregates the consumption policy functions of spenders and savers. Heuristically, we can totally differentiate both sides of (7) - (8) to find that, to first order, the equilibrium paths of interest rates and taxes must satisfy[8]

$$\underbrace{\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \end{pmatrix} \times \boldsymbol{\varepsilon}}_{\text{excess demand}} = \underbrace{\begin{pmatrix} \frac{\partial \mathbf{y}}{\partial \mathbf{r}} - \frac{\partial \mathbf{i}}{\partial \mathbf{r}} - \frac{\partial \mathbf{c}}{\partial \mathbf{r}} & -\frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}_e} \\ \mathbf{0} & I \end{pmatrix} \times \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\tau}}_e \end{pmatrix}}_{\text{GE adjustment}} \tag{9}$$

The initial disturbance $\boldsymbol{\varepsilon}$ leads to some time path of initial excess demand or supply, and some shortfall in the intertemporal government budget. Now suppose that a path $\hat{\mathbf{r}}$ and $\hat{\boldsymbol{\tau}}_e$ solves (9) for a tax rebate shock $\boldsymbol{\varepsilon}_\tau$. Then it is immediate that the same path $(\hat{\mathbf{r}}, \hat{\boldsymbol{\tau}}_e)$ also solves (9) for a fiscal spending shock with the same intertemporal demand profile – that is, if $\hat{\mathbf{c}}_\tau^{PE} = \hat{\mathbf{g}}_g$. Intuitively, it does not matter why there is excess demand for the final good or why is there a shortfall in the government budget constraint, it only matters *how much*.[9]

I will show next that the exact same proof strategy – exclusion restrictions on where different shocks can and cannot enter in a linearized representation of the equilibrium – continues to work in the much richer structural model of Section 2.1.

2.4 A general equivalence result

This section discusses the restrictions on the model of Section 2.1 necessary for exact demand equivalence. I discuss three assumptions: the first two are already embedded in the model set-up but explicitly re-stated here for emphasis, while the third one is a meaningful additional restriction. I will discuss the sensitivity of demand equivalence with respect to violations of each restriction in Section 4 and Appendix E of the Online Appendix.

The first assumption restricts goods bundles in the economy.

Assumption 1. *Households and government consume a single, homogeneous final good.*

[8]The functions in (7) - (8) map infinite-dimensional paths of inputs (e.g., $\mathbf{r}$) into infinite-dimensional paths of outputs (e.g., $\mathbf{y}$); as such, the derivatives in (9) should be understood as derivatives in infinite-dimensional vector spaces. For the spender-saver model existence of these derivatives is immediate from the characterization of the model solution in Appendix B.2; for the more general model, I simply assume differentiability (see Appendix B.1).

[9]This heuristic argument ensures only that a solution for the rebate policy is also a solution for a particular public spending shock; it is, however, silent on the existence and uniqueness of such paths. For the spender-saver application, existence and uniqueness are verified in the usual way for the recursive representation of the analogous linearized stochastic difference equation (Blanchard & Kahn, 1980), which implies that the infinite-dimensional general equilibrium feedback map in (9) has a unique left-inverse (see Appendix B.2).

The second assumption relates to the interest rates faced by households and government. In my model, all agents borrow and lend at a common interest rate; Assumption 2 first restates this model property, and then provides an additional restriction on the actual financing of government expenditure shocks.

Assumption 2. *Households and government borrow and lend at the same interest rate. The path of taxes and transfers used to finance a given public expenditure shock $\boldsymbol{\varepsilon}_\tau$ or $\boldsymbol{\varepsilon}_g$ depends only on the present value of the expenditure, not its time path.*

The third assumption restricts the economy's labor market. In response to the partial equilibrium increase in consumption demand $\hat{\mathbf{c}}_\tau^{PE}$, the average marginal utility of consumption declines, and so sticky-wage unions may try to bargain for higher wages. I denote the desired adjustment in aggregate hours worked at unchanged wages by $\hat{\boldsymbol{\ell}}_\tau^{PE}$, defined formally in Appendix B.1. My third assumption provides two possible sufficient conditions to guarantee that $\hat{\boldsymbol{\ell}}_\tau^{PE} = \mathbf{0}$.

Assumption 3. *There are either no wealth effects in labor supply, or wages are perfectly sticky (i.e., wage adjustment costs are infinitely large).*

These assumptions are sufficient for the following generalized equivalence result.

Proposition 2. *Consider the structural model of Section 2.1. Suppose that, for each one-time shock $\{\tau, g\}$, the equilibrium transition path exists and is unique. Under Assumptions 1 to 3, the responses of consumption to a tax rebate shock τ and to a government spending shock g with $\hat{\mathbf{g}}_g = \hat{\mathbf{c}}_\tau^{PE}$ satisfy, to first order,*

$$\hat{\mathbf{c}}_\tau \quad = \quad \underbrace{\hat{\mathbf{c}}_\tau^{PE}}_{PE\ response} \quad + \quad \underbrace{\hat{\mathbf{c}}_g}_{=\ GE\ feedback} \tag{10}$$

The proof strategy for Proposition 2 is almost identical to that of the spender-saver RBC model. Equilibria in the richer model can generally not be characterized as solutions to a single market-clearing condition in a single price; instead, as I show formally in Appendix A.1, they are solutions to a rich set of market-clearing conditions and other restrictions. Assumptions 1 to 3 are simply sufficient to ensure that private and public spending shocks perturb the same market-clearing conditions by the same amount, and thus elicit the same general equilibrium adjustment, exactly as in the proof of Proposition 1.

Assumption 1 – in conjunction with the requirement that $\hat{\mathbf{g}}_g = \hat{\mathbf{c}}_\tau^{PE}$ – ensures that the private and public demand shocks lead to the same excess demand pressure for the common final good. Since households and governments borrow and lend at identical rates, these identical net excess demand paths can *in principle* be financed using identical paths of taxes and transfers. Without Ricardian equivalence, however, the precise timing of the financing matters. Assumption 2 then simply ensures that, indeed, the two shocks are financed in exactly the same way. Under these restrictions alone, the general equilibrium propagation of private and public spending shocks may still differ, as households may also decide to *directly* adjust their desired labor supply following the shock $\boldsymbol{\varepsilon}_\tau$. Assumption 3 – a restriction on household behavior – is enough to rule this out: Following the shock $\boldsymbol{\varepsilon}_\tau$, households either do not wish to or are not able to directly adjust their hours worked, i.e. $\hat{\boldsymbol{\ell}}_\tau^{PE} = \mathbf{0}$. Together, Assumptions 1 to 3 ensure exact demand equivalence.

Relation to previous work. In the proof of Proposition 2, I establish the existence of a "demand multiplier" $\mathcal{D}$ – a map transforming net excess demand paths (such as $\hat{\mathbf{c}}_\tau^{PE}$ or $\hat{\mathbf{g}}_g$) into general equilibrium impulse responses. As such, it builds on results in Auclert & Rognlie (2018) and Auclert et al. (2018). In particular, in Auclert et al. (2018), the intertemporal Keynesian cross matrix $\mathcal{M}$ – a special case of the multiplier $\mathcal{D}$ – governs the transmission of private and public demand shocks, establishing demand equivalence. Their result applies in a model with passive monetary policy, demand-determined labor, and without investment; Proposition 2 provides the generalization to the model of Section 2.1.[10] The intuition for such a common "demand multiplier" is particularly transparent in the static Keynesian cross, and was for example previously discussed in Hall (2009) and Reichling & Whalen (2012).

Illustration. Figure 2 illustrates exact demand equivalence in a rich heterogeneous-agent New Keynesian (HANK) model. The parameterization of the model is identical to that of Section 4.1, with the sole difference that wages are fully – rather than just partially – sticky.

Together, Figures 1 and 2 reveal that demand equivalence is a general *semi*-structural feature of models. In the RBC model, excess demand is moderated through higher real interest rates, leading to significant general equilibrium crowding-out. In the HANK model, consumption is instead crowded in: interest rates increase relatively little, while the transitory increase in labor income pulls up household spending even further. Crucially, the nature of general equilibrium adjustment is in both cases invariant to the origin of the shock.

[10] Auclert & Rognlie (2018) is, to the best of my knowledge, the first paper to discuss general equilibrium multipliers for perfect foresight transition paths. In particular, they prove that different kinds of consumption demand shocks are propagated identically in general equilibrium; Proposition 2 shows under what conditions those same multipliers also apply to public demand shocks – that is, demand equivalence.

Demand Equivalence, Sticky-Wage HANK Model

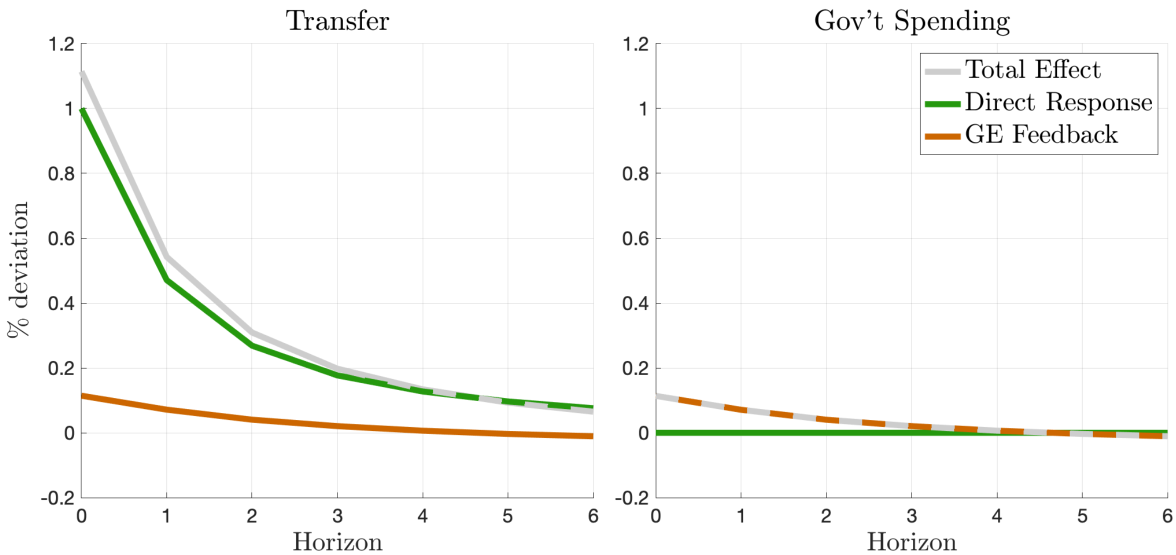

Figure 2: Consumption impulse response decompositions after tax rebate and government spending shocks in the estimated HANK model of Section 4.1, but with fully rigid wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

2.5 Extension to general exclusion restrictions

In the analysis so far I have used a particular shifter of consumer spending – lump-sum transfers – and a particular structural model – the general framework of Section 2.1 – to present demand equivalence as a set of restrictions on economic primitives. Proposition 2 covers a meaningful (and easily interpretable) space of models, but many of the restrictions implicit in this framework are in fact unnecessary. To make this point, this section complements the previous discussion with an abstract statement of shock equivalence in terms of exclusion restrictions in a linearized equilibrium representation. Throughout, I continue to use the same notational conventions as in my baseline structural model.

A general statement of consumption demand equivalence requires only two ingredients: an aggregate consumption function $\mathbf{c} = \mathbf{c}(\mathbf{s}^h; \boldsymbol{\varepsilon}_d)$ and a (differentiable) system of equations characterizing equilibrium aggregates $\mathcal{H}(\mathbf{x}; \boldsymbol{\varepsilon}_d, \boldsymbol{\varepsilon}_g) = \mathbf{0}$, where ε_{dt} and ε_{gt} are generic shocks to private and public spending, respectively, and the inputs to household consumption $\mathbf{s}^h$ are determined as part of the set of aggregates $\mathbf{x}$. Demand equivalence is then simply a set of exclusion restriction on derivatives of the equilibrium mapping $\mathcal{H}(\bullet)$: As long as

$$\frac{\partial \mathcal{H}}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d \quad = \quad \frac{\partial \mathcal{H}}{\partial \boldsymbol{\varepsilon}_g} \times \boldsymbol{\varepsilon}_g \tag{11}$$

it follows immediately that, to first order,

$$\hat{\mathbf{c}}_d = \underbrace{\hat{\mathbf{c}}_d^{PE}}_{\text{PE response}} + \underbrace{\hat{\mathbf{c}}_g}_{=\text{ GE feedback}} \tag{12}$$

(11) is a general exclusion restriction on the equilibrium system: both shocks must enter all equilibrium equations symmetrically. The proof sketch for the simple equivalence result in Section 2.3 illustrates this requirement transparently; similarly, the proof of Proposition 2 works because, under my imposed restrictions, the equilibrium of the general model can also be cast in a form consistent with (11).[11] In Appendices C.2 and C.3 I give examples of other shocks and models that can be written in this general form. First, I extend the baseline model to allow for time preference shocks as a generic consumption shifter. Equivalence obtains under the same restrictions as those discussed here. Second, I consider several interesting structural models beyond the familiar New Keynesian business-cycle tradition – including for example models with non-rational expectation formation of firms and households –, and show that they still all fit into the general semi-structure of (11).

Interpretation & Outlook. The results in this section assert that, for a quite general family of structural models, the general equilibrium effects of private and public spending shocks are invariably tied together. This equivalence, however, says nothing about the *strength* of those common general equilibrium effects: in the simple spender-saver model of Section 2.3, partial equilibrium spending responses are largely crowded-out; in my HANK example, feedback effects are relatively weak; and in Appendix C.4, I show two further extreme examples, one with full crowding-out, the other with very strong amplification. Ultimately, the strength of general equilibrium effects – and so the size of the missing intercept – is an empirical question. The next section presents my empirical strategy.

3 Estimating consumption demand counterfactuals

This section develops a two-step methodology to estimate semi-structural macro counterfactuals for generic consumption demand shifters. I describe the general approach in Section 3.1, and then in Section 3.2 apply it to study the effects of income tax rebates.

[11] Casting my results as exclusion restrictions on structural equilibrium representations suggests a connection to the identification of systems of simultaneous equations. This connection is explored further and formalized in Guren et al. (2020, Section 4).

3.1 The two-step methodology

Consider a researcher interested in the response of aggregate consumption to a generic "consumption demand" shifter – a shock that directly affects (incentives for) household spending. Examples of such shifters are plentiful in recent work; among the most notable are income tax rebate stimulus (Parker et al., 2013), household deleveraging due to tightened borrowing conditions (Mian et al., 2013; Berger et al., 2017), changes in household bankruptcy exemptions (Auclert et al., 2019) and redistribution across households through taxation (Jappelli & Pistaferri, 2014). As is well-known, estimation of the aggregate effects of such shocks is severely complicated by their likely endogeneity to wider macroeconomic conditions.

In response to these challenges, most recent work has tried to estimate shock propagation using household-level data, exploiting plausibly exogenous heterogeneity in shock *exposure*. In the remainder of this section I use the example of income tax rebate stimulus in the model of Section 2.1 to argue that: (i) the econometric estimands of such cross-sectional regressions are interpretable as direct (partial equilibrium) shock responses (here $\hat{\mathbf{c}}_\tau^{PE}$), and (ii) we can use existing estimates of the aggregate effects of changes in government spending to proxy for their missing general equilibrium intercept (here $\hat{\mathbf{c}}_\tau^{GE}$). By the discussion in Section 2.5, my results extend without change to generic consumption demand shifters ε_{dt}.

Data-generating process. I make two minor changes to the structural model of Section 2.1. The two changes ensure that the cross-sectional and time series regression estimands required for my two-step methodology are actually well-defined.

First, I now consider the linear vector moving-average representation induced by the first-order perturbation solution of the model, assuming that the shocks ε_{st}, $s \in \{\tau, g\}$ are mutually i.i.d. and $N(0,1)$. I use s subscripts to indicate impulse response functions to such *one-time* shocks; by certainty equivalence, these responses are to first order identical to the transition paths for one-off shocks studied in Section 2, justifying the re-use of notation. Such aggregate risk allows me to evaluate time-series regressions (i.e., VARs) run on model-generated data. Second, I assume that the rebate shock faced by household i satisfies $\varepsilon_{\tau it} = \xi_{\tau it} \times \varepsilon_{\tau t}$, where $\xi_{\tau it}$ is i.i.d. across households and time (and uncorrelated with any household characteristics), with $\mathbb{E}(\xi_{\tau it}) = 1$ and $\text{Var}(\xi_{\tau it}) > 0$.[12] Given this heterogeneity in shock exposure, I can study regressions run on the cross section of households.

[12] In the proof of Proposition 3 I show that, under my assumptions on the exposure term $\xi_{\tau it}$, all aggregates are – to first order – unaffected by this cross-sectional heterogeneity in shock exposure.

Micro regressions. A typical regression exploiting microeconomic heterogeneity in household exposure to the transfer shock $\varepsilon_{\tau t}$ takes the form

$$c_{it+h} = \alpha_i + \delta_t + \beta_{\tau h} \times \varepsilon_{\tau it} + u_{it+h}, \quad h = 0, 1, 2, \ldots \tag{13}$$

where α_i and δ_t are individual and time fixed effects.[13]

It is straightforward to show that, under my assumptions, regressions such as (13) estimate average household-level causal effects that are interpretable as direct partial equilibrium shock responses, consistent with Definition 1.

Proposition 3. *Suppose an econometrician observes a panel of household consumption $\{c_{it}\}$ and measures of shock exposure $\{\varepsilon_{\tau it}\}$ generated from the linear vector moving average representation of the structural model of Section 2.1. Then the ordinary least-squares estimand of $\boldsymbol{\beta}_\tau \equiv (\beta_{\tau 0}, \beta_{\tau 1}, \ldots)'$ satisfies*

$$\boldsymbol{\beta}_\tau \quad = \quad \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \varepsilon_{\tau 0}} \, di \quad = \quad \hat{\mathbf{c}}_\tau^{PE} \tag{14}$$

In words, regressions such as (13) do not estimate the true macro counterfactual $\hat{\mathbf{c}}_\tau$, but instead give a household-level average treatment effect that is interpretable as a *direct* (or partial equilibrium) response, $\hat{\mathbf{c}}_\tau^{PE}$ – precisely the object defined in my decomposition in Definition 1. Obtaining such estimates from a sequence of cross-sectional micro regressions like (13) is the first step of my methodology.

General equilibrium effects. To map the micro estimates $\boldsymbol{\beta}_\tau$ into full general equilibrium counterfactuals, researchers would typically use full structural models, calibrated to be consistent with the micro estimates themselves as well as various other formally or informally targeted macro moments (e.g. Kaplan & Violante, 2018). The results in Section 2 suggest that, for a large class of models, evidence on the aggregate effects of public spending shocks should be a highly informative macro moment – in fact informative enough to give some counterfactuals without ever having to solve any particular model.

The second step of my proposed methodology leverages this insight. Under demand equivalence, we know that – for any possible transfer policy – there exists *some* aggregate fiscal spending policy with the same general equilibrium effects. To make this observation

[13]The regression in (13) is at the individual level. My analysis here thus does *not* apply to cross-regional regressions, as for example in Mian et al. (2013). I generalize the method to such regressions in Wolf (2019).

practically useful, suppose that the econometrician uses the standard fiscal policy time series toolkit (Ramey, 2018) to jointly estimate the response of the macro-economy to a list of n_k different kinds of government spending shocks $\{\varepsilon_{g_k}\}_{k=1}^{n_k}$, where these shocks induce potentially different paths of aggregate government spending and tax financing. For my approach to work it must be the case that, for some linear combination of shocks with weights $\{\gamma_k\}$,

$$\boldsymbol{\beta}_\tau = \sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{g}}_{g_k} \tag{15}$$

In words, a linear combination of government spending shocks available from macro experiments gives similar partial equilibrium excess demand pressure as the transfer ε_τ. This is a restrictive requirement, but I will later demonstrate through several applications that such "demand matching" is possible in practice for many interesting partial equilibrium demand paths $\boldsymbol{\beta}_\tau$ and so shocks ε_τ.[14] The researcher would then estimate the missing general equilibrium intercept as

$$\hat{\mathbf{c}}_\tau^{GE} = \sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{c}}_{g_k} \tag{16}$$

By Assumption 2, (16) recovers the general equilibrium effects of a rebate financed exactly like the composite public spending shock $\sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{g}}_{g_k}$. For example, in my main application, I will consider a persistently deficit-financed spending expansion, so my counterfactuals will be interpretable as applying to a similarly deficit-financed transfer.

Macro counterfactuals. Putting all the pieces together, we finally get the full general equilibrium counterfactual

$$\hat{\mathbf{c}}_\tau = \underbrace{\boldsymbol{\beta}_\tau}_{\text{PE response}} + \underbrace{\sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{c}}_{g_k}}_{\text{GE feedback}} \tag{17}$$

Since by assumption the econometrician is able to jointly estimate the response of the macro-economy to the list $\{\varepsilon_{g_k}\}_{k=1}^{n_k}$ of fiscal shocks, she can straightforwardly construct frequentist standard errors or Bayesian confidence sets for the full general equilibrium term.[15]

[14] Alternatively, the researcher could fix the observed time series public spending paths $\{\hat{\mathbf{g}}_{g_k}\}$, and then interpret (16) as the *exact* missing intercept of any private spending path that can be written as (15).

[15] Except for a brief discussion in Section 3.2, I will largely ignore estimation uncertainty for the direct response. Under my assumptions, sampling uncertainty for the micro and macro parts is independent, so

In the remainder of this paper I illustrate my method with three examples. First, in Section 3.2, I use it to estimate the aggregate effects of a deficit-financed income tax rebate (Parker et al., 2013). Second, in Appendix F, I study the effects of a one-off (budget-neutral) income re-distribution from rich (low-MPC) to poor (high-MPC) households. Finally, in Section 5, I establish a theoretical investment demand equivalence result, and use the natural investment analogue of my approach to estimate the effects of bonus depreciation stimulus.

3.2 Application: income tax rebates

I combine micro and macro evidence to estimate the response of aggregate consumption and output to a one-off, deficit-financed income tax rebate (i.e., lump-sum transfer). My main finding is that full general equilibrium counterfactuals are close to direct micro estimates: The direct increase in consumption demand is accommodated almost one-for-one through an increase in output, with relatively limited general equilibrium crowding-out.

Direct Response. I first require an estimate of the direct spending response $\hat{\mathbf{c}}_\tau^{PE}$. For a one-off, one-quarter stimulus payment, this direct spending response is given as

$$\hat{c}_{\tau t}^{PE} \equiv MPC_{t,0} \times \hat{\tau}_0$$

where

$$MPC_{t,0} \equiv \int_0^1 \frac{\partial c_{it}}{\partial \tau_0} di$$

is the average marginal propensity to consume at time t out of an income gain at time 0.

Several recent studies have used rich household spending data to estimate objects that are either exactly or approximately interpretable as the desired average MPC (e.g. Johnson et al., 2006; Parker et al., 2013; Jappelli & Pistaferri, 2014; Fagereng et al., 2018).[16] A common finding in this literature is that households spend most of a (small) one-time income receipt on impact, and that the spending response decays back to zero relatively quickly. Johnson et al. (2006) and Parker et al. (2013), who specifically focus on the consumption response

construction of joint confidence sets is in principle straightforward. Intuitively, this is so because macro shocks are differenced out in micro regressions, and micro shocks have no aggregate effects (see Appendix D.4).

[16]By my definition of the consumption function $\mathbf{c}(\bullet)$ in Section 2, the MPC should be interpreted as an MPC after adjusting for any endogenous response of earnings. In the notation of Auclert (2019), it is the adjusted $\widehat{MPC}$. As discussed there, popular empirical studies arguably estimate this adjusted object. Furthermore, estimated earnings responses are usually small anyway, as discussed further in Section 4.2.

to income tax rebates, estimate a differenced version of the micro regression (13); building on Proposition 3, Appendix D.1 shows that – at least under some assumptions on household expectation formation – their regression estimates $MPC_{0,0}$ (and $MPC_{1,0}$).

The point estimates of Parker et al. (2013) suggest that, following the rebate stimulus of 2008, total consumption expenditures increased by about 50 to 90 per cent of payments in the quarter of the receipt. Given the overall size and (staggered) timing of the stimulus, this spending response corresponds to around 1.5 per cent of personal consumption expenditure on impact, and 0.7 per cent in the following quarter.[17] In the left panel of Figure 3, the green x's show the corresponding direct consumption responses $\hat{c}_{\tau 0}^{PE}$ and $\hat{c}_{\tau 1}^{PE}$; the solid green line shows what I take as the estimate of the full direct spending response $\hat{\mathbf{c}}_{\tau}^{PE}$.

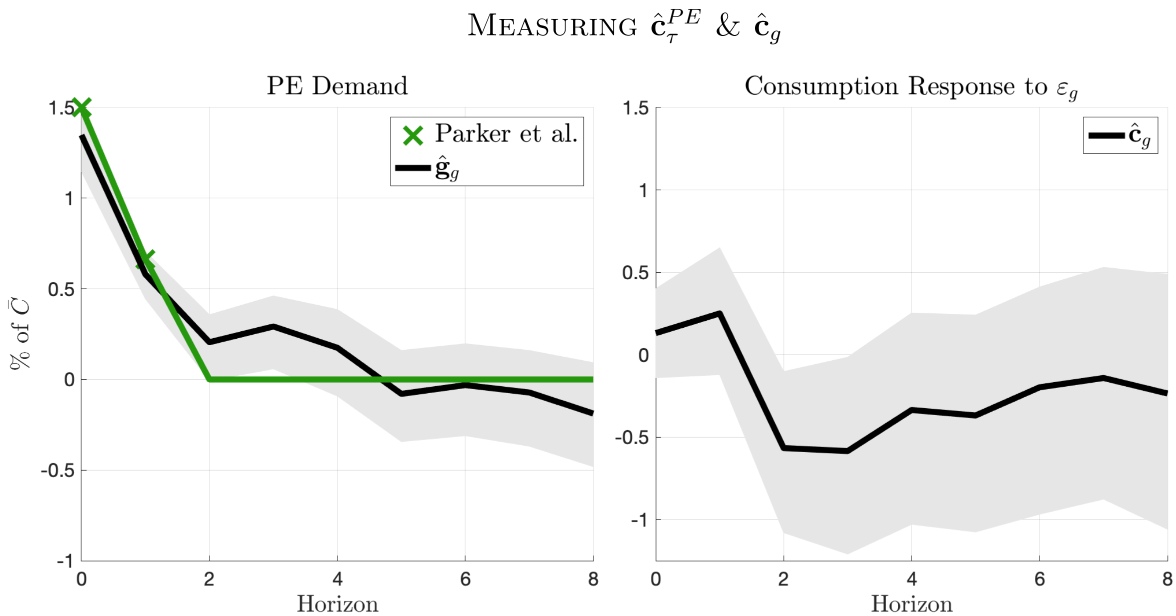

Figure 3: The left panel shows direct consumption responses to the tax rebate (green) vs. direct government spending response to identified spending shock (black), with 16th and 84th percentile confidence bands (grey), quarterly frequency. Estimated consumption responses from Parker et al. (2013) (Table 3). The right panel shows the response of consumption to the fiscal spending shock.

The Missing Intercept. It remains to estimate the aggregate effects of a similarly transitory and deficit-financed expansion in government spending. Previous studies often find

[17] These estimates include the durables spending response, consistent with the extended equivalence result in Appendix C.3. My VAR analysis thus also throughout contains measures of total consumption. For completeness, however, I have repeated my analysis using evidence on non-durables consumption only. Results are similar, with the missing intercept now an even more tightly estimated 0.

that increases in government consumption – both transitory and more persistent – are accommodated roughly one-for-one through increases in output, with relatively little feedback to private spending (Hall, 2009; Gechert, 2015; Caldara & Kamps, 2017; Ramey, 2018).

My identification of government spending shock propagation relies on professional forecast errors for real federal spending. Formally, I treat the forecast errors as a (noisy) measure of exogenous innovations to public expenditure; intuitively, this assumption can be justified by likely lags in the response of fiscal policy to any changes in macroeconomic fundamentals. In the language of macro identification, I assume that residualized forecast errors are valid external instruments (Stock & Watson, 2018). In Appendix D.3 I phrase the IV relevance and exclusion restrictions for the forecast errors z_t in terms of primitives in my data-generating process. For the purposes of the analysis here it suffices to note that, under those restrictions, a recursive vector autoregression (VAR) in instrument and macro aggregates $(z_t, y_t')'$ correctly identifies the impulse responses of all macro aggregates y_t to a structural innovation ε_{gt} in aggregate government spending. In my empirical implementation, the vector y_t includes measures of overall government spending, output, consumption, investment, hours worked, and federal marginal tax rates. Forecast errors are available from 1981:Q3 onwards; to plausibly estimate $\hat{\mathbf{c}}_g$ in a stable macroeconomic regime, I restrict my sample to end in 2008:Q4 (similar to Ramey, 2011). Since variable definitions, data construction and estimation procedure are all standard and in line with previous work, I relegate further details to Appendix D.3. The appendix also discusses several robustness checks, notably with respect to the vector of macro aggregates y_t, lag length selection, prior selection, and controls.

The results are also included in Figure 3. The left panel shows that, in response to the shock, government spending increases sharply, but returns to baseline quickly. Importantly, the time profile of the demand expansion quite closely mirrors the micro-estimated expansion in private consumption spending.[18] The right panel shows the corresponding response of aggregate consumption – $\hat{\mathbf{c}}_g$. Consumption appears to be somewhat crowded-in on impact, and mildly crowded-out in the following quarters. Finally, in Appendix D.3, I show that the expansion in government spending is deficit-financed, with a delayed increase in taxes. By Assumption 2, my counterfactuals for a transitory income tax rebate should thus be interpreted as pertaining to a particular, *deficit-financed* transfer to households.[19]

[18]While the two demand paths are quite similar, they of course do not align perfectly. In Appendix E.7, I discuss the accuracy of my approximation under imperfect demand matching.

[19]In the model of Section 2.1, tax financing is always non-distortionary. However, it is immediate that all results extend without change to distortionary tax financing – the only requirement is that the *same* financing is used for both policies.

Macro Counterfactuals. To construct a valid general equilibrium counterfactual, it now simply remains to sum the empirically estimated $\hat{\mathbf{c}}_\tau^{PE}$ and $\hat{\mathbf{c}}_g$. The results are displayed in the left panel of Figure 4. Note that, for construction of the plot, I take the point estimate of $\hat{\mathbf{c}}_\tau^{PE}$ as given, and only account for macroeconomic estimation uncertainty.[20]

Income Tax Rebate, Aggregate Impulse Responses

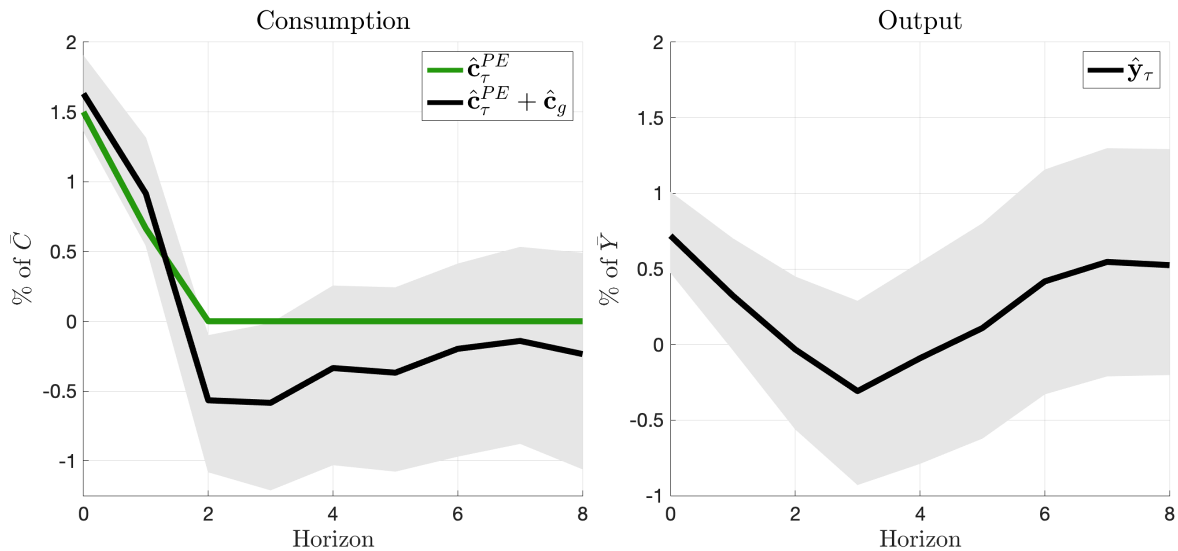

Figure 4: Consumption and output responses to an income tax rebate shock, quarterly frequency. The full consumption response is computed following the exact additive decomposition of Proposition 2, while the output response is simply equal to the response after a government spending shock. The grey areas again correspond to 16th and 84th percentile confidence bands.

The left panel shows the full general equilibrium counterfactual for consumption. The *aggregate* effect of the policy – according to my decomposition given as the simple sum $\hat{\mathbf{c}}_\tau^{PE} + \hat{\mathbf{c}}_g$ – appears to be quite close to the (large) micro-estimated direct spending response $\hat{\mathbf{c}}_\tau^{PE}$ documented in Parker et al. (2013). Thus, perhaps surprisingly, the various price and multiplier effects cited in previous empirical and theoretical work seem to roughly cancel. The right panel shows the corresponding impulse response of output which, by the demand equivalence result, is identical for tax rebate and government spending expansion. Here I find a significant (if short-lived) response, with output on impact rising by somewhat less than 1 per cent, and then returning to baseline. Overall, deficit-financed income tax rebates appear to provide meaningful stimulus to aggregate consumption and output.

[20]This is in keeping with my emphasis on the "missing intercept." However, since the direct spending response is only a function of the impact response coefficient of Parker et al. (2013), and since this coefficient is statistically significant, it is immediate that the full impact response is – by independence – also significant.

My analysis suggests that, at least for income tax rebate stimulus, the "missing intercept" of general equilibrium feedback is approximately zero. This conclusion is an immediate implication of the theoretical demand equivalence result in conjunction with a relatively standard piece of empirical evidence – deficit-financed government spending multipliers around 1, with limited feedback to private spending (Ramey, 2018). While direct micro estimates are thus actually a reliable guide to full general equilibrium counterfactuals, arriving at this conclusion nevertheless required important *macroeconomic* identifying assumptions, notably on demand equivalence and the identification of aggregate public spending shocks.[21]

Implications for Structural Modeling. The results presented here can also inform structural modeling. First, they provide a useful "sufficient statistics" summary of structural work: *Any* model with a consumption response to tax rebates different from Figure 4 either (i) breaks demand equivalence, (ii) is inconsistent with micro evidence on large direct spending responses, or (iii) is inconsistent with macro evidence that suggests around unit fiscal multipliers and limited feedback to private spending. Second, note that my semi-structural counterfactuals are only valid for the particular spending and financing paths of the underlying time-series experiments. Impulse responses for never-before implemented financing paths will in contrast invariably rely on fully specified structural models; for estimation of those, the cross-sectional and time-series impulse responses of my analysis promise to be highly informative "identified moments" in the sense of Nakamura & Steinsson (2018).

4 Approximation accuracy

Consistency of my proposed methodology requires that the exclusion restrictions of demand equivalence hold exactly. As such, my approach is similar in the spirit to the older structural VAR literature, which relied on exact contemporaneous zero restrictions for correct shock identification (e.g. Sims, 1980; Christiano et al., 1999). In this section I study the sensitivity of my approach to violations of these exclusion restrictions.

To do so, I consider quantitative structural business-cycle models as test laboratories. An econometrician observes cross-sectional and time-series data generated by the model and

[21] My pre-crisis VAR implicitly measures "normal-time" general equilibrium effects. As I discuss in Appendix C.1, however, the equivalence result applies to (small) shocks around *any* given current state of the economy. With an extended sample containing the recent period of low rates, I find suggestive evidence of slightly larger multipliers, consistent with the results in Ramey & Zubairy (2018) and Debortoli et al. (2019).

implements my methodology; I then ask how close the resulting output is to the true model-implied impulse responses. To abstract from estimation uncertainty, I suppose throughout that the econometrician in fact has access to infinitely large samples. The remainder of this section will consider two sets of laboratories: first, conventional business-cycle models estimated to fit standard macroeconomic time series aggregates; and second, further extensions of this baseline specifically designed to break demand equivalence.

4.1 Estimated business-cycle models

My model laboratories build on the structural framework of Section 2.1. In this benchmark model, demand equivalence fails *only* because of changes in household labor supply induced by the transfer.

Proposition 4. *Consider the structural model of Section 2.1. Suppose that, for each one-time shock $\{\tau, g\}$, the equilibrium transition path exists and is unique. Under Assumptions 1 and 2, the responses of consumption to a tax rebate shock τ and to a government spending shock g with $\hat{\mathbf{g}}_g = \hat{\mathbf{c}}_\tau^{PE}$ satisfy, to first order,*

$$\hat{\mathbf{c}}_\tau \quad = \quad \hat{\mathbf{c}}_\tau^{PE} \quad + \quad \hat{\mathbf{c}}_g \quad + \quad error\left(\hat{\boldsymbol{\ell}}_\tau^{PE}\right) \tag{18}$$

where the error function is characterized in Appendix G.5 and is equal to $\mathbf{0}$ if $\hat{\boldsymbol{\ell}}_\tau^{PE} = \mathbf{0}$.

The labor supply channel will result in the econometrician *over*-estimating the response of consumption to the transfer: Households want to work less upon receipt of the transfer, but do not similarly cut their labor supply after increases in government spending. The two model laboratories of this section – a novel estimated HANK model, as well as the canonical business-cycle model of Justiniano et al. (2010) – however show that this channel is largely absent in models with even moderate degrees of nominal rigidity. This finding is consistent with conventional wisdom in the business-cycle literature (e.g. Christiano, 2011a,b): at least for short-run fluctuations, hours worked are largely demand-determined.

The estimated HANK model. My main test for accuracy of the proposed approximation is an estimated HANK model, featuring a conventional consumption-savings problem under imperfect insurance embedded into an otherwise standard medium-scale DSGE model. I provide a brief outline of the model and my estimation strategy here, and relegate further details to Appendix B.3.

The household block is as described in Section 2.1. The rest of the economy is designed to be as close as possible to the medium-scale structural model of Justiniano et al. (2010). First, I allow for investment adjustment costs, variable capacity utilization, and a rich monetary policy rule. Second, in addition to the government spending shock discussed in Section 2.1, I also include shocks to total factor productivity and the marginal efficiency of investment, to household patience, to wage mark-ups, and to monetary policy. As I restrict attention to first-order transition paths, these additional shocks of course do not affect the propagation of private and public spending shocks; I only include them for estimation purposes.

I calibrate the model's steady state using targets familiar from the HANK literature (e.g. Kaplan et al., 2018). Importantly, because household self-insurance is severely limited, the average MPC is high, at around 30% quarterly out of a lump-sum 500$ income gain. Model parameters governing dynamics are then estimated using conventional likelihood methods (An & Schorfheide, 2007; Mongey & Williams, 2017) on a standard set of macroeconomic aggregates.[22] The key exception is the degree of wage stickiness which – in light of its centrality to my results – is directly calibrated to be consistent with recent micro evidence (Grigsby et al., 2019; Beraja et al., 2019), with wage re-sets every 2.5 quarters on average. I then solve the model at the estimated posterior mode, and implement the demand equivalence approximation for a one-off income tax rebate following the method of Section 3. Results are displayed in Figure 5.

As in the figures of Section 2, the left panel decomposes the response of aggregate consumption to a transfer into direct partial equilibrium (green) and indirect general equilibrium (orange) effects. The decomposition is almost identical to that in Figure 2 – of course unsurprising since the two models are essentially the same, with the sole exception being the degree of wage stickiness. The right panel then approximates this consumption response by summing (i) the direct response $\hat{\mathbf{c}}_\tau^{PE}$ and (ii) the full response of consumption to an expansion in government spending satisfying the demand matching in (15), $\hat{\mathbf{c}}_g$. Under Assumption 3, the decomposition would be exact, so the grey and black dotted lines would be indistinguishable. Now, since households want to optimally work less after receiving the transfer, my approximation overstates the full consumption response, as expected. The associated error, however, is small, at just below 3 per cent of the true peak consumption response. Intuitively, even with (moderately) sticky wages, labor is largely demand-determined, and so small and transitory shifts in labor supply are quantitatively irrelevant.

[22] Specifically, I include measures of output, inflation, a short-term interest rate, consumption, investment, and hours worked – six observables for six shocks.

Approximate Demand Equivalence, Estimated HANK Model

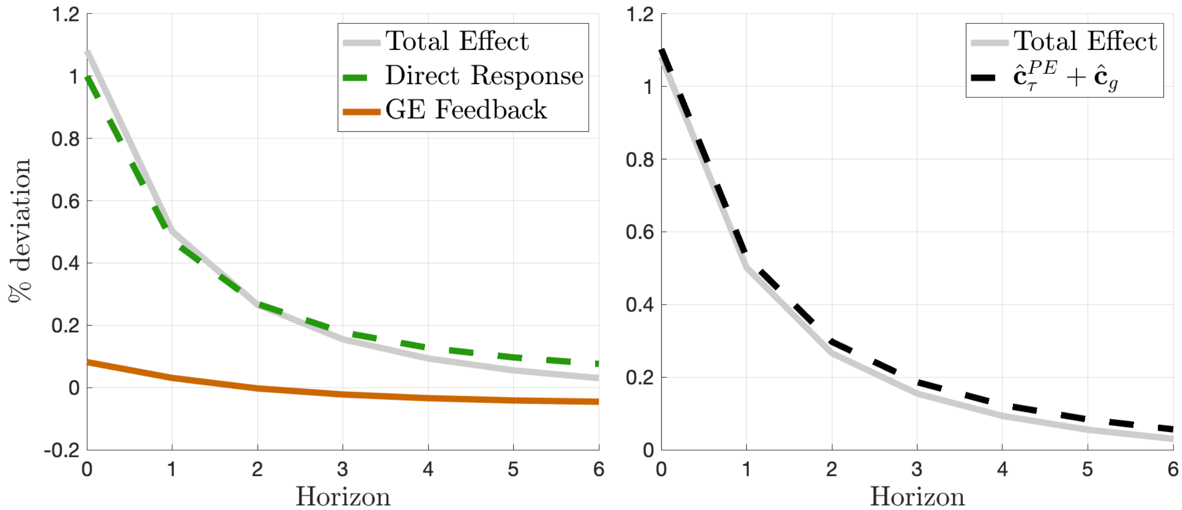

Figure 5: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, with details on the parameterization in Appendix B.3. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Finally, in Appendix E.1, I show that this conclusion is not at all special to the posterior mode of the estimated model. To make this point, I *randomly* draw model parameters from large supports, solve the implied model, and compute the approximation accuracy. The analysis reveals that, of all estimated parameters, only the degrees of price and wage rigidity have a material impact on the accuracy of the approximation.

Canonical business-cycle models. The approximate demand equivalence documented in Figure 5 is also not specific to my particular HANK model set-up, but in fact also a feature of various well-known estimated business-cycle models. In Appendix E.2, I use the demand equivalence approximation to construct counterfactuals for consumer impatience shocks in the popular structural model of Justiniano et al. (2010), solved at the posterior mode. Since wages are even stickier here, the approximation is even better than in the estimated HANK model, with the approximation error now barely visible.

4.2 Labor supply

In the estimated models considered so far, the labor supply channel is quantitatively irrelevant because transitory shifts in household labor supply have little effect on aggregate equilibrium outcomes. To illustrate a failure of demand equivalence, Figure 6 constructs the output of the two-step procedure in my estimated HANK model, re-parameterized to feature almost perfectly flexible prices and wages. Given strong wealth effects in labor supply, household earnings decline sharply in response to transfer receipt; indeed, by the argument in Auclert & Rognlie (2017), this decline in labor earnings can be as large as the spending increase associated with the transfer. The two-step approximation misses this substantial decline in hours worked, and so overstates the total consumption response.

Failure of Demand Equivalence, Flex-Price HA Model

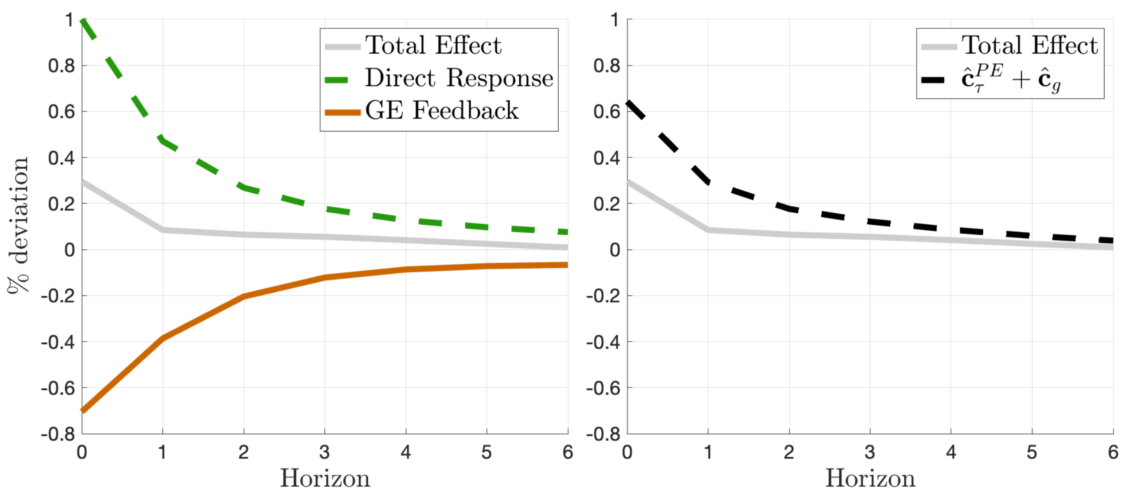

Figure 6: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, but with flexible prices and wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

There are two possible responses to this failure of my approach. The first, taken in this paper, is to argue that the mechanism underlying the inaccuracy is at odds with all kinds of micro and macro evidence, and thus ignore it. On the macro side, standard time series estimation exercises usually call for near-zero wealth effects in labor supply (Schmitt-Grohé & Uribe, 2012; Born & Pfeifer, 2014) as well as at least some degree of nominal rigidity. On the micro side, quasi-experimental evidence at the household level suggests that, in response to lump-sum transfer receipts, hours worked and earnings drop by an order of magnitude less than spending increases (e.g. Cesarini et al., 2017; Fagereng et al., 2018).[23] Finally, direct estimates of wage rigidity using microeconomic or cross-regional data suggest moderate, but non-trivial, amounts of stickiness (Grigsby et al., 2019; Beraja et al., 2019).

[23] Coibion et al. (2020) document similarly small earnings responses after stimulus payments in the COVID-19 recession.

The second response is to extend the two-step methodology to explicitly take into account the labor supply channel. The proof of Proposition 4 reveals that the error term is identical to the response of aggregate consumption to a particular kind of "labor wedge" shock – the sudden desire of households to work less. Combining (i) micro estimates of the size of this labor wedge shock and (ii) evidence on the aggregate effects of distortionary labor income taxes (Mertens & Ravn, 2013), we could recover a direct *empirical* correction for the error term. I briefly discuss this extension in Appendix E.3; unsurprisingly, since empirical estimates of the desired labor supply contraction are small, the results of this augmented procedure are almost identical to my benchmark estimates.

4.3 Other model extensions

In this section I use three extensions of my estimated HANK model to gauge the sensitivity of the demand equivalence approximation to violations of the required exclusion restrictions. These extensions are: (i) a richer household asset space, (ii) productive and direct utility benefits of government spending, and (iii) multiple production sectors. For all cases I discuss the likely sign of the bias based on a priori reasoning, and then evaluate it quantitatively using empirically disciplined extensions of my baseline estimated HANK model. The section here collects my main results, with details relegated to Appendix E in the Online Appendix.

Financial assets. If households and government borrow and lend at different rates, then identical changes in private and public net excess demand cannot be financed using identical paths of taxes and transfers, violating Assumption 2. In particular, if household returns are high (low) relative to government returns, then taxes need to increase by less (more) to finance private relative to public spending. These lower (higher) taxes will sooner or later feed back into consumption; if this happens immediately, then the simple demand equivalence approximation will tend to under-state (over-state) the true impulse response of aggregate consumption to the private demand shifter.[24]

Appendix E.4 considers a two-asset HANK model – similar to Kaplan et al. (2018) – in which households pay an intermediation fee on liquid deposits, biasing my equivalence approximation upwards.[25] This extension, however, has relatively little effect on the accuracy

[24]A similar logic applies to open economy settings. Assuming imperfect home bias, transfers to households will leak abroad, while government purchases will be fully concentrated on domestic output. The equivalence approximation will thus tend to over-state, with the bias increasing in the openness of the economy.

[25]I choose this asset structure to ensure that all sources of bias point in the same direction (upwards). If

of the demand equivalence approximation. The intuition is simple: Suppose that, in response to a shock, direct (partial equilibrium) household spending increases by 1$ for one year. My approximation compares the aggregate effects of this shock to those of an identical expansion in aggregate public spending. Crucially, even if the wedge between average household and government discount rates is an (arguably implausible) five per cent, the difference in present discounted values of the two spending expansions is just five cents – relatively small compared to the initial size of the stimulus. The implied difference in tax financing is thus also small, and so the approximation remains accurate.

Useful government spending. In the benchmark model, government spending is socially useless – it is neither valued by households, nor does it have productive benefits. In Appendix E.5 I study the extent to which my approximation is affected by non-separabilities in the private valuation of government spending (following Leeper et al., 2017) and productive benefits of government investment.

If private and public consumption are complements in household utility (as for example estimated in Leeper et al., 2017), then the demand equivalence approximation is likely to over-state the consumption response to the demand shifter. I show that, if the researcher is willing to commit to a particular functional form of the non-separability, then a straightforward modification of my methodology can correct for the bias. However, standard estimates of the complementarity are anyway small enough to not substantially threaten the accuracy of the equivalence approximation. Next, to gauge the importance of productive benefits of government spending, I review the empirical evidence on fiscal multipliers for public investment. The key take-away is that such multipliers are usually estimated to be larger than standard spending multipliers, with a cross-study average of around 1.5 (e.g. Gechert, 2015; Ramey, 2016). These findings suggest that the productive benefits of government investment are quantitatively significant, and thus caution against the use of public investment multipliers for my approximations.

Multi-goods models. While much of the business-cycle literature relies on the one-good abstraction in Assumption 1, private and public consumption baskets of course differ in reality. Most obviously, this heterogeneity in consumption goods will break equivalence because of relative price responses (Ramey & Shapiro, 1998). Somewhat more subtly, heterogeneity

government debt yields low returns, but household-held equity gives high returns, then the asset return bias would be negative and thus offset the labor supply bias.

can also matter if different goods have different production technologies. If so, the income generated by private and public demand shocks may flow to different factors of production, leading to heterogeneous general equilibrium propagation (Baqaee, 2015; Alonso, 2017).

To test the likely strength of those mechanisms, I in Appendix E.6 consider a multi-sector variant of my estimated HANK model, disciplined by empirical evidence on (i) the strength of relative price effects and (ii) heterogeneity in network-adjusted labor shares. With prices somewhat sticky in the short run, and given relatively moderate differences in network-adjusted labor shares across sectors, I find that demand equivalence continues to approximately hold. In particular, the relative weakness of the sectoral incidence channel documented here is consistent with the findings in Alonso (2017).

5 Investment demand counterfactuals

This section extends my methodology to study the general equilibrium propagation of shocks to investment demand. Section 5.1 sketches the theoretical equivalence result, and Section 5.2 leverages it to derive semi-structural aggregate counterfactuals for investment tax stimulus through accelerated bonus depreciation.

5.1 Investment demand equivalence

My formal analysis of investment demand equivalence again builds on the model of Section 2.1. Anticipating the empirical application, I augment this model to feature investment tax credit shocks $\boldsymbol{\varepsilon}_q$ – shocks that reduce the cost of capital purchases by intermediate goods producers at time t by an amount $\tau_{qt} = \tau_{qt}(\boldsymbol{\varepsilon}_q)$.[26]

THE EQUIVALENCE RESULT. I define direct (partial equilibrium) responses and indirect (general equilibrium) feedback for firm investment exactly analogously to Definition 1, using the implied aggregate investment function $\mathbf{i}(\bullet)$. As before, the question is – under what restrictions on primitives does the response of investment to suitably chosen fiscal spending experiments give the "missing intercept" $\hat{\mathbf{i}}_q^{GE}$?

My proof strategy is identical to that in Section 2: I characterize equilibrium response paths as a system of market-clearing conditions, and then impose enough restrictions on

[26] More generally, my results can be interpreted as applying to any kind of shock that appears as a reduced-form wedge in firm investment optimality conditions (Chari et al., 2007).

this system to ensure that the investment tax credit as well as a suitable fiscal experiment perturb the same equations by the same amount. In the baseline model, the investment tax credit policy has three main effects. First, investment responds; since investment invariably boosts the future productive capacity of the economy, production also increases, so the induced partial equilibrium net excess demand path for the final output and investment good is $\hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE}$.[27] Second, the policy may be redistributive: the cost of financing is borne by taxed households, but the benefits accrue to households receiving dividend payments. Third, more investment and so more capital will increase the marginal product of labor, and thus firm labor demand will increase.

Matching the first effect is straightforward: I consider a fiscal spending expansion with

$$\hat{\mathbf{g}}_g = \hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE} \tag{19}$$

For the other two effects I require additional exclusion restrictions. To rule out heterogeneous distributional implications of tax financing and dividend payments, I assume that household income risk is perfectly insurable, thus effectively imposing a representative-household structure. This restriction also implies that Ricardian equivalence holds, so the precise timing of the policy financing is irrelevant. Next, to ignore the labor demand response, I assume that labor supply is perfectly flexible, either because of a large Frisch elasticity of labor supply, or because labor is fully demand-determined. Under those two additional restrictions on primitives, a fiscal experiment satisfying (19) indeed gives the "missing intercept" of the investment response:

$$\hat{\mathbf{i}}_q \quad = \quad \hat{\mathbf{i}}_q^{PE} \quad + \quad \hat{\mathbf{i}}_g \tag{20}$$

Analogously we can also recover the output response as

$$\hat{\mathbf{y}}_q \quad = \quad \hat{\mathbf{y}}_q^{PE} \quad + \quad \hat{\mathbf{y}}_g \tag{21}$$

I formally state the equivalence result and its assumptions in Appendix A.2.

Scope. The formal investment demand equivalence proposition requires no additional restrictions on the production side of the economy: Firms can face a rich set of real and financial frictions, including (convex and non-convex) capital adjustment costs as well as a

[27]Strictly speaking, the output response may also appear as a wedge in the monetary policy rule. The restrictions required to rule out this policy effect are relatively standard and so not further discussed here; details are provided in Appendix A.2.

generic set of constraints on equity issuance and borrowing. The restrictions on the rest of the economy are meaningful, but – as discussed in Appendix C.6 – routinely imposed in quantitative general equilibrium models of investment; in particular, the assumptions hold in the well-known models of Khan & Thomas (2008), Khan & Thomas (2013), Winberry (2018), and Bloom et al. (2018). As such, the decomposition in (20) provides a useful identification result for a popular class of models.

5.2 Application: bonus depreciation

The investment equivalence result justifies a two-step procedure to study generic investment demand shifters, exactly analogous to my analysis of consumption shifters in Section 3. In this section I use the two-step procedure to construct a general equilibrium counterfactual for investment bonus depreciation stimulus – that is, the ability to tax-deduct investment expenditure at a faster rate, as implemented in the U.S. in the two most recent recessions (Zwick & Mahon, 2017). It is well-known that, in the absence of firm-level financial frictions, such accelerated bonus depreciation schedules are isomorphic to the investment tax credits covered by the investment equivalence result (Winberry, 2018).[28]

Direct Response. My estimates of the direct response of investment to the shock rely heavily on Zwick & Mahon (2017) and Koby & Wolf (2020), who exploit cross-sectional firm-level heterogeneity in the exposure to bonus depreciation investment stimulus. In Koby & Wolf (2020), we estimate dynamic regressions akin to (13) and give sufficient conditions under which the regression estimands are identical to or at least informative about the desired partial equilibrium investment spending responses $\hat{\mathbf{i}}_q^{PE}$. The discussion is largely analogous to that in Proposition 3, so I relegate further details to Appendix D.2.

Given a path for the direct investment spending response $\hat{\mathbf{i}}_q^{PE}$, I can recover the implied partial equilibrium production path using estimates of the capital elasticity of production. In particular, assuming a simple Cobb-Douglas production function $y = (k^{\alpha}\ell^{1-\alpha})^{\nu}$ as well as competitive spot labor markets, it is straightforward to show that

$$\hat{\hat{y}}_{qt}^{PE} = \frac{\alpha\nu}{1-(1-\alpha)\nu} \times \hat{\hat{k}}_{qt-1}^{PE}$$

Thus, given estimates of the capital depreciation rate δ, the capital share α, and the returns

[28] As in Section 3, the analysis in this section implicitly relies on the stochastic VMA representation of the model, and considers estimation of impulse responses to one-off structural shocks.

to scale parameter ν, it is possible to recover the implied partial equilibrium production path. Consistent with my estimated HANK model, I set $\delta = 0.016$, $\alpha = 0.2$ and $\nu = 1$.

I take the regression estimates of $\hat{\mathbf{i}}_{qt}^{PE}$ for $t = 0, 1, 2, 3$ straight from Koby & Wolf (2020, Table 1). The green x's in the investment panel of Figure 7 show the estimated path of direct investment spending responses to a one-quarter bonus depreciation shock worth around 8 cents, a shock similar in magnitude to the stimulus of 2008-2010. The solid green line extrapolates the empirical estimates to a full response path using a Gaussian basis function, similar to Barnichon & Matthes (2018). I take this extrapolated path to be the empirical estimate of the full spending response path $\hat{\mathbf{i}}_q^{PE}$.

Investment demand increases substantially and persistently in response to the stimulus. Since capital is pre-determined, and since all prices faced by firms (except for taxes and so effective capital goods prices) are fixed by the nature of the partial equilibrium exercise, output does not increase on impact, but instead only gradually increases over time. Together, the investment and output responses translate into a more complicated intertemporal net excess demand profile, displayed in the top left panel: Net excess demand is large and positive on impact (due to higher investment demand), but turns negative over time, as additional capital becomes productive and so expands the productive capacity of the economy.

THE MISSING INTERCEPT. Following (20), it remains to replicate the estimated net excess demand path through a suitable list of government spending shocks:

$$\hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE} \quad = \quad \sum_{k=1}^{n_k} \gamma_k \quad \times \quad \hat{\mathbf{g}}_{g_k} \tag{22}$$

It is unlikely that any single estimated spending shock can replicate the reversal documented in Figure 7. Encouragingly, much previous work on fiscal multipliers actually estimates the effects of *delayed* increases in government spending (Ramey, 2011; Caldara & Kamps, 2017) – that is, government spending news shocks. In principle, combining these delayed spending responses with the immediate spending effect estimated in Section 3.2 should allow me to replicate the net demand effects of the investment tax credit.

To operationalize this insight, I consider the same VAR as before, but now study the responses to residualized innovations in both the instrument equation as well as the equation for government expenditure itself. The first innovation is simply the shock studied in Section 3.2, while the second innovation is similar to the popular recursive identification scheme of Blanchard & Perotti (2002), augmented to include forecast errors as a control for anticipa-

Investment Tax Credit, Impulse Responses

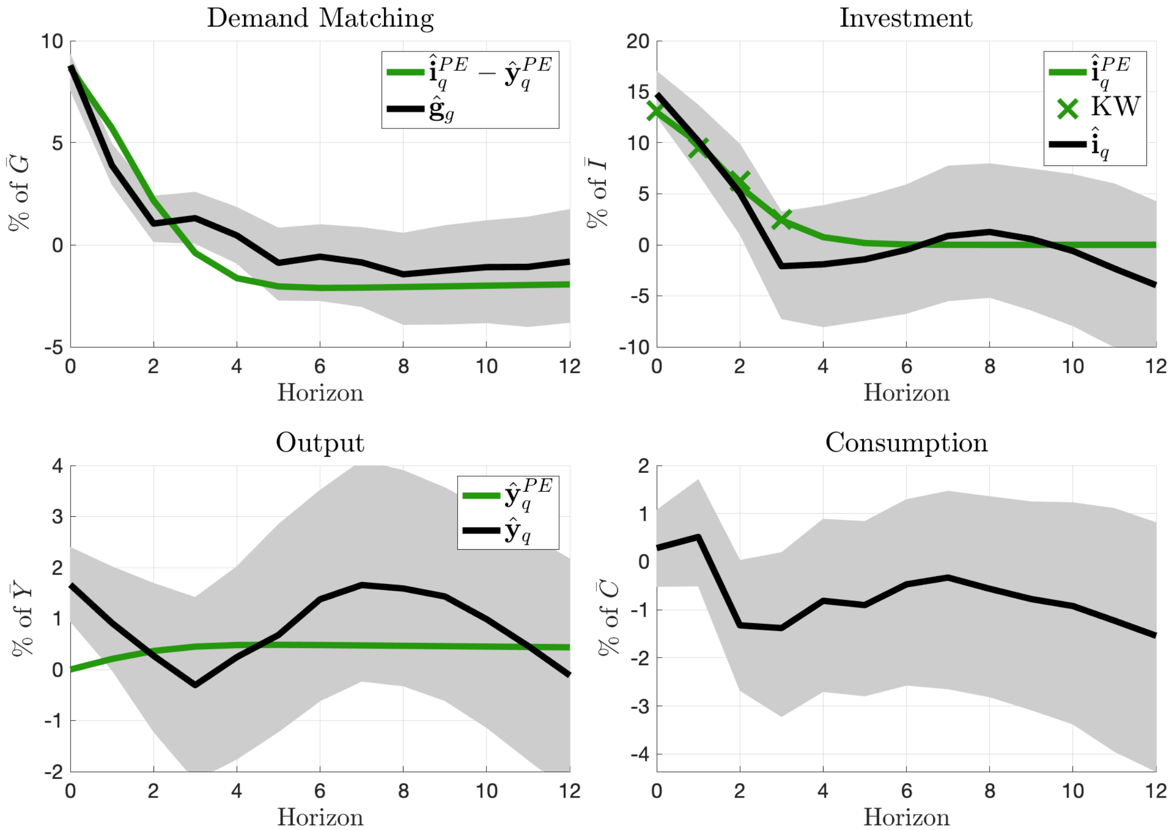

Figure 7: Investment, output and consumption responses to an investment tax incentive shock, quarterly frequency, with the partial equilibrium net output response path matched to a linear combination of government spending shocks. "KW" refers to Koby & Wolf (2020); details are given in Appendix D.2. The investment and output responses are computed in line with (20) - (21), while the consumption response is just the response after the identified combination of government spending shocks. The grey areas again correspond to 16th and 84th percentile confidence bands.

tion effects. Consistent with previous work, I find the effects of the Blanchard-Perotti shock to be delayed, and so a linear combination of the two shocks allows me to match the implied net excess demand path of the investment demand shock, as shown in the top left panel of Figure 7. Further details on the empirical implementation – in particular the construction of standard errors for general equilibrium feedback – are provided in Appendix D.3.

Macro Counterfactuals. All results for general equilibrium counterfactuals are displayed in Figure 7. With the requirement that $\hat{\mathbf{g}}_g = \hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE}$ satisfied, the investment and output panels implement the additive decompositions in (20) and (21), respectively. My main finding is that the substantial partial equilibrium investment demand responses estimated in Zwick & Mahon (2017) and Koby & Wolf (2020) also survive in general equilibrium. The increase in investment demand is accommodated through a sharp immediate increase in output as well as a smaller and somewhat delayed drop in consumption. Taken together, the large direct investment spending responses estimated in micro data as well as extant evidence on the transmission of aggregate government spending shocks suggest that bonus depreciation investment incentives provide a sizable macroeconomic stimulus.

6 Conclusion

I develop a new approach to the estimation of aggregate counterfactuals for a general family of consumption and investment demand shifters. Micro data can help us learn about the extent to which these demand shifters directly stimulate household and firm spending, and extant evidence on the transmission of public spending shocks to private expenditure contains valuable information about the "missing intercept" of general equilibrium accommodation. Applied to income tax rebates and investment bonus depreciation incentives, my methodology suggests that both policies substantively stimulate aggregate spending, and that this expansion in spending is accommodated in general equilibrium through a one-to-one increase in production, rather than being crowded out through price responses. Any calibrated structural model that implies large general equilibrium amplification or dampening either does not feature demand equivalence, or is inconsistent with conventional estimates of the size of the fiscal multiplier.

The methodology promises to be useful beyond the applications considered in this paper. In the companion paper Wolf (2019), I generalize my results to map cross-*regional* regression estimates into macro counterfactuals, with an application to household deleveraging due to tighter borrowing conditions (Mian et al., 2013; Guerrieri & Lorenzoni, 2017). Examples of other interesting macro shocks covered by the two-step method include firm uncertainty (Bloom, 2009; Bloom et al., 2018), shocks to firm credit conditions (Khan & Thomas, 2013) and household debt relief (Auclert et al., 2019). I leave those applications to future work.

A Appendix

A.1 Proof of consumption demand equivalence

Similar to the heuristic argument in Section 2.3, I begin by writing the equilibrium of the full baseline model as a dynamic system of market-clearing conditions.

Lemma A.1. *Consider the structural model of Section 2.1. A perfect foresight equilibrium is a sequence of nominal interest rates* $\mathbf{i}^b$*, aggregate output* $\mathbf{y}$*, wages* $\mathbf{w}$ *and the endogenous part of tax rebates* $\boldsymbol{\tau}_e$ *such that*

$$\begin{aligned}
\mathbf{c}(\mathbf{s}^h(\mathbf{x});\boldsymbol{\varepsilon}) + \mathbf{i}(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) &= \mathbf{y}(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) \\
\boldsymbol{\ell}^h(\mathbf{s}^u(\mathbf{x};\boldsymbol{\varepsilon})) &= \boldsymbol{\ell}^f(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) \\
\mathbf{y}(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) &= \mathbf{y} \\
\boldsymbol{\tau}_e(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) &= \boldsymbol{\tau}_e
\end{aligned}$$

where $\mathbf{x} = (\mathbf{i}^b, \mathbf{y}, \mathbf{w}, \boldsymbol{\tau}_e)$, $\mathbf{s}^h = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \boldsymbol{\ell}, \boldsymbol{\tau}_e, \mathbf{d})$, $\mathbf{s}^u = (\boldsymbol{\pi}, \mathbf{w}, \mathbf{c})$, $\mathbf{s}^f = (\mathbf{i}^b, \mathbf{w}, \boldsymbol{\pi})$*, and the consumption, production, investment, labor demand and labor supply functions* $\mathbf{c}(\bullet)$*,* $\mathbf{y}(\bullet)$*,* $\mathbf{i}(\bullet)$*,* $\boldsymbol{\ell}^h(\bullet)$ *and* $\boldsymbol{\ell}^f(\bullet)$ *are derived from optimal firm, household and union behavior.*

Proof. See Appendix G of the Online Appendix. □

A perfect foresight equilibrium is thus, to first order, a solution to the system of linear equations

$$\begin{aligned}
\left(\frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) + \left(\frac{\partial \mathbf{i}}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} &= \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) \\
\left(\frac{\partial \boldsymbol{\ell}^h}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) &= \left(\frac{\partial \boldsymbol{\ell}^f}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \boldsymbol{\ell}^f}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) \\
\left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) &= J_2 \times \hat{\mathbf{x}} \\
\left(\frac{\partial \boldsymbol{\tau}_e}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}\right) &= J_4 \times \hat{\mathbf{x}}
\end{aligned}$$

where J_i denotes the infinite-dimensional generalization of the selection matrix selecting the ith entry of a vector x_t. Assuming equilibrium existence and uniqueness,[29] there exists a unique linear

[29]Existence and uniqueness of a bounded transition path for representative-agent models can be shown as usual. For the heterogeneous-agent models, I have verified existence and uniqueness for particular numerical examples, using the conditions of Blanchard & Kahn (1980).

map $\mathcal{H}$ such that

$$\hat{\mathbf{x}} = \underbrace{\mathcal{H}}_{\text{GE adjustment}} \times \underbrace{\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} + \frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} - \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}} \\ \frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}} - \frac{\partial \boldsymbol{\ell}^f}{\partial \boldsymbol{\varepsilon}} \\ \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \end{pmatrix}}_{\text{direct shock response}} \times \boldsymbol{\varepsilon}$$

where $\mathcal{H}$ is a left inverse of

$$\begin{pmatrix} \frac{\partial \mathbf{y}}{\partial \mathbf{x}} - \frac{\partial \mathbf{c}}{\partial \mathbf{x}} - \frac{\partial \mathbf{i}}{\partial \mathbf{x}} \\ \frac{\partial \boldsymbol{\ell}^f}{\partial \mathbf{x}} - \frac{\partial \boldsymbol{\ell}^h}{\partial \mathbf{x}} \\ J_2 - \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \\ J_4 - \frac{\partial \boldsymbol{\tau}_e}{\partial \mathbf{x}} \end{pmatrix}$$

The assumed existence and uniqueness of the equilibrium ensures that this left inverse is in fact unique. Now consider transfer and government spending shocks. To reduce unnecessary clutter, I use the notation $\frac{\partial}{\partial \boldsymbol{\varepsilon}_s}$ (rather than the generic $\frac{\partial}{\partial \boldsymbol{\varepsilon}}$) to denote derivatives for a shock path where only entries of shock s are non-zero. By definition of the firm policy functions (see Appendix B.1), we know that $\frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}_\tau} = \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}_\tau} = \frac{\partial \boldsymbol{\ell}^f}{\partial \boldsymbol{\varepsilon}_\tau} = \mathbf{0}$, and similarly that $\frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}_g} = \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}_g} = \frac{\partial \boldsymbol{\ell}^f}{\partial \boldsymbol{\varepsilon}_g} = \mathbf{0}$. We also know that $\frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}_g} = \mathbf{0}$, and by Assumption 3 $\frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}_\tau} = \mathbf{0}$. The two direct shock responses are then

$$\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau} \\ \mathbf{0} \\ \mathbf{0} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}_\tau} \end{pmatrix} \times \boldsymbol{\varepsilon}_\tau = \begin{pmatrix} \hat{\mathbf{c}}_\tau^{PE} \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{e\tau}^{PE} \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}_g} \\ \mathbf{0} \\ \mathbf{0} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}_g} \end{pmatrix} \times \boldsymbol{\varepsilon}_g = \begin{pmatrix} \hat{\mathbf{g}}_g \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{eg}^{PE} \end{pmatrix}$$

By Assumption 2, we know that there exists a matrix $\mathcal{T}$ such that $\hat{\boldsymbol{\tau}}_{e\tau}^{PE} = \mathcal{T} \times \hat{\mathbf{c}}_\tau^{PE}$, $\hat{\boldsymbol{\tau}}_{eg}^{PE} = \mathcal{T} \times \hat{\mathbf{g}}_g$, and so $\hat{\boldsymbol{\tau}}_{e\tau}^{PE} = \hat{\boldsymbol{\tau}}_{eg}^{PE}$ if $\hat{\mathbf{c}}_\tau^{PE} = \hat{\mathbf{g}}_g$. Thus, in response to the lump-sum income tax rebate and government spending shocks, the response path of consumption satisfies

$$\hat{\mathbf{c}}_\tau = \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau} \times \boldsymbol{\varepsilon}_\tau}_{\hat{\mathbf{c}}_\tau^{PE}} + \underbrace{\frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{c}}_\tau^{PE} \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{e\tau}^{PE} \end{pmatrix}}_{\mathcal{D} \times \hat{\mathbf{c}}_\tau^{PE}}, \quad \text{and} \quad \hat{\mathbf{c}}_g = \mathbf{0} + \underbrace{\frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{g}}_g \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{eg}^{PE} \end{pmatrix}}_{\mathcal{D} \times \hat{\mathbf{g}}_g}$$

respectively, where $\mathcal{D}$ is a common *demand multiplier*. This establishes that $\hat{\mathbf{c}}_\tau^{GE} = \hat{\mathbf{c}}_g^{GE}$, and so (10) follows from simple addition. $\square$

A.2 Details on investment demand equivalence

I begin with the restrictions needed for an exact investment demand equivalence result. The first assumption – a single common final good – is again implicit in the model set-up.

Assumption A.1. *A single final good is used for (government) consumption and investment.*

In imposing this first restriction, I implicitly assume that all meaningful capital adjustment costs are internal to the firm, and that the aggregate supply of capital (out of the common final good) is perfectly elastic. This assumption is consistent with the empirical findings in House & Shapiro (2008), Edgerton (2010) and House et al. (2017).

The second assumption rules out any heterogeneous distributional implications associated with dividend and tax payments following the firm subsidy and the equivalent fiscal spending change.

Assumption A.2. *All households $i \in [0,1]$ have identical preferences, receive equal lump-sum government rebates τ_t and firm dividend income d_t, and face no idiosyncratic earnings risk.*

The third assumption allows me to ignore the labor demand response.

Assumption A.3. *Labor supply is perfectly elastic, either because the Frisch elasticity of labor supply is infinite (linear labor disutility), or because wages are perfectly sticky. Furthermore, the period household felicity function is separable in consumption and hours worked.*

Finally, I require an additional restriction on monetary policy feedback. If the monetary authority directly responds to the *level* of aggregate output, then the increase in production associated with the investment subsidy will induce a contractionary monetary response. I rule this out by assuming that the monetary authority targets the output gap (as for example in Justiniano et al., 2010), or does not respond at all to fluctuations in output.

Assumption A.4. *The monetary authority's interest rate rule does not include an endogenous response to fluctuations in the level of aggregate output.*

Under Assumptions A.1 to A.4 I can prove the following demand equivalence result.

Proposition A.1. *Consider the structural model of Section 2.1. Suppose that, for each one-time shock $\{q, g\}$, the equilibrium transition path exists and is unique. Then, under Assumptions A.1 to A.4, the response of investment to an investment tax credit shock q and to a government spending shock g with $\hat{\mathbf{g}}_g = \hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE}$ satisfies, to first order,*

$$\hat{\mathbf{i}}_q \quad = \quad \hat{\mathbf{i}}_q^{PE} \quad + \quad \hat{\mathbf{i}}_g$$

Proof. See Appendix G of the Online Appendix. □

It is immediate from the proof of Proposition A.1 that all results extend immediately to generic investment "wedges" (Chari et al., 2007).

References

Ahn, S., Kaplan, G., Moll, B., Winberry, T., & Wolf, C. K. (2017). When Inequality Matters for Macro and Macro Matters for Inequality. *NBER Macroeconomics Annual*, *32*.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, *109*(3), 659–684.

Alexander, E. & Seater, J. (2009). The Federal Income Tax Function. Working Paper.

Alonso, C. (2017). Cutting Back on Labor Intensive Goods? Implications for Fiscal Stimulus. Working Paper.

An, S. & Schorfheide, F. (2007). Bayesian analysis of DSGE models. *Econometric Reviews*, *26*(2-4), 113–172.

Andrews, I., Gentzkow, M., & Shapiro, J. M. (2018). On the informativeness of descriptive statistics for structural estimates. Technical report, National Bureau of Economic Research.

Arias, J. E., Rubio-Ramírez, J. F., & Waggoner, D. F. (2018). Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications. *Econometrica*, *86*(2), 685–720.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, *109*(6), 2333–67.

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2019). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. Technical report, National Bureau of Economic Research.

Auclert, A., Dobbie, W. S., & Goldsmith-Pinkham, P. (2019). Macroeconomic Effects of Debt Relief: Consumer Bankruptcy Protections in the Great Recession. Technical report, National Bureau of Economic Research.

Auclert, A. & Rognlie, M. (2017). A note on multipliers in NK models with GHH preferences. Technical report, Working Paper.

Auclert, A. & Rognlie, M. (2018). Inequality and aggregate demand. Technical report, National Bureau of Economic Research.

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

Auclert, A., Rognlie, M., & Straub, L. (2019). Micro Humps, Macro Jumps: Monetary Policy and Business Cycles in an Estimated HANK Model. Working Paper.

Baqaee, D. (2015). Targeted fiscal policy. Working Paper.

Barnichon, R. & Matthes, C. (2018). Functional approximation of impulse responses. *Journal of Monetary Economics*, *99*, 41–55.

Barro, R. J. & King, R. G. (1984). Time-separable preferences and intertemporal-substitution models of business cycles. *The Quarterly Journal of Economics*, *99*(4), 817–839.

Barro, R. J. & Redlick, C. J. (2011). Macroeconomic effects from government purchases and taxes. *The Quarterly Journal of Economics*, *126*(1), 51–102.

Barsky, R. B., House, C. L., & Kimball, M. S. (2007). Sticky-price models and durable goods. *American Economic Review*, *97*(3), 984–998.

Bayer, C. & Luetticke, R. (2020). Solving heterogeneous agent models in discrete time with many idiosyncratic states by perturbation methods. Working Paper.

Bayer, C., Lütticke, R., Pham-Dao, L., & Tjaden, V. (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, *87*(1), 255–290.

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

Beraja, M. & Wolf, C. K. (2020). Demand composition and the strength of recoveries. Working Paper.

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

Berger, D. & Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, *83*(1), 101–154.

Blanchard, O. & Perotti, R. (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *The Quarterly Journal of Economics*, *117*(4), 1329–1368.

Blanchard, O. J. & Kahn, C. M. (1980). The solution of linear difference models under rational expectations. *Econometrica*, *48*, 1305–1311.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, *77*(3), 623–685.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., & Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, *86*(3), 1031–1065.

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

Born, B. & Pfeifer, J. (2014). Policy risk and the business cycle. *Journal of Monetary Economics*, *68*, 68–85.

Caldara, D. & Kamps, C. (2017). The analytics of SVARs: a unified framework to measure fiscal multipliers. *The Review of Economic Studies*, *84*(3), 1015–1040.

Cesarini, D., Lindqvist, E., Notowidigdo, M. J., & Östling, R. (2017). The effect of wealth on individual and household labor supply: evidence from Swedish lotteries. *American Economic Review*, *107*(12), 3917–46.

Chari, V. V., Kehoe, P. J., & McGrattan, E. R. (2007). Business cycle accounting. *Econometrica*, *75*(3), 781–836.

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

Chodorow-Reich, G., Nenov, P. T., & Simsek, A. (2019). Stock market wealth and the real economy: A local labor market approach. Technical report, National Bureau of Economic Research.

Christiano, L., Eichenbaum, M., & Evans, C. (1999). Monetary Policy Shocks: What Have We Learned and to What End? In J. B. Taylor & M. Woodford (Eds.), *Handbook of Macroeconomics, Volume 1, Part A* chapter 2, (pp. 65–148). Elsevier.

Christiano, L. J. (2011a). Comment on" unemployment in an estimated new keynesian model". In *NBER Macroeconomics Annual 2011, Volume 26* (pp. 361–380). University of Chicago Press.

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

Coibion, O., Gorodnichenko, Y., & Weber, M. (2020). How did u.s. consumers use their stimulus payments? Working Paper 27693, National Bureau of Economic Research.

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

Drautzburg, T. (2016). A narrative approach to a fiscal DSGE model. FRB of Philadelphia Working Paper.

Edgerton, J. (2010). Estimating machinery supply elasticities using output price booms. Working Paper.

Erceg, C. J., Henderson, D. W., & Levin, A. T. (2000). Optimal Monetary Policy with Staggered Wage and Price Contract. *Journal of Monetary Economics*, *46*(2), 281–313.

Fagereng, A., Holm, M. B., & Natvik, G. J. J. (2018). MPC heterogeneity and household balance sheets. CESifo Working Paper Series.

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

Galí, J., López-Salido, J. D., & Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, *5*(1), 227–270.

Galí, J., Smets, F., & Wouters, R. (2012). Unemployment in an estimated New Keynesian model. *NBER Macroeconomics Annual*, *26*(1), 329–360.

Gechert, S. (2015). What fiscal policy is most effective? a meta-regression analysis. *Oxford Economic Papers*, *67*(3), 553–580.

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

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

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

Guren, A., McKay, A., Nakamura, E., & Steinsson, J. (2020). What Do We Learn From Cross-Sectional Empirical Estimates in Macroeconomics? Working Paper.

Hagedorn, M., Manovskii, I., & Mitman, K. (2019). The fiscal multiplier. Technical report, National Bureau of Economic Research.

Hall, R. E. (2009). By How Much Does GDP Rise If the Government Buys More Output? *Brookings Papers on Economic Activity*, *2*, 183–231.

House, C. L., Mocanu, A.-M., & Shapiro, M. D. (2017). Stimulus effects of investment tax incentives: Production versus purchases. Technical report, National Bureau of Economic Research.

House, C. L. & Shapiro, M. D. (2008). Temporary investment tax incentives: Theory with evidence from bonus depreciation. *American Economic Review*, *98*(3), 737–68.

Jaimovich, N. & Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, *99*(4), 1097–1118.

Jappelli, T. & Pistaferri, L. (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, *6*(4), 107–36.

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

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

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

Kaplan, G. & Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, *82*(4), 1199–1239.

Kaplan, G. & Violante, G. L. (2018). Microeconomic Heterogeneity and Macroeconomic Shocks. *Journal of Economic Perspectives, Special Issue on "The State of Macroeconomics a Decade After The Crisis*, *32*(3).

Khan, A. & Thomas, J. K. (2008). Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics. *Econometrica*, *76*(2), 395–436.

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

Koby, Y. & Wolf, C. K. (2020). Aggregation in Heterogeneous-Firm Economies: Theory and Measurement. Working Paper.

Krusell, P. & Smith, Jr, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, *106*(5), 867–896.

Leduc, S. & Wilson, D. (2013). Roads to prosperity or bridges to nowhere? theory and evidence on the impact of public infrastructure investment. *NBER Macroeconomics Annual*, *27*(1), 89–142.

Leeper, E. M., Traum, N., & Walker, T. B. (2017). Clearing up the fiscal multiplier morass. *American Economic Review*, *107*(8), 2409–54.

McKay, A., Nakamura, E., & Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, *106*(10), 3133–58.

Mertens, K. & Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, *103*(4), 1212–47.

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

Mian, A. & Sufi, A. (2009). The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis. *The Quarterly Journal of Economics*, *124*(4), 1449–1496.

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

Mongey, S. & Williams, J. (2017). Firm dispersion and business cycles: Estimating aggregate shocks using panel data. Manuscript, New York University.

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

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

Ottonello, P. & Winberry, T. (2018). Financial Heterogeneity and the Investment Channel of Monetary Policy. Unpublished working paper.

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*(6), 2530–53.

Perotti, R. (2007). In Search of the Transmission Mechanism of Fiscal Policy. *NBER Macroeconomics Annual*, *22*, 169–249.

Plagborg-Møller, M. & Wolf, C. K. (2019). Local Projections and VARs Estimate the Same Impulse Responses. Manuscript, Princeton University.

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *The Quarterly Journal of Economics*, *126*(1), 1–50.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. In J. B. Taylor & H. Uhlig (Eds.), *Handbook of Macroeconomics*, volume 2 chapter 2, (pp. 71–162). Elsevier.

Ramey, V. A. (2018). Ten Years After the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research? Working Paper.

Ramey, V. A. & Shapiro, M. D. (1998). Costly capital reallocation and the effects of government spending. In *Carnegie-Rochester Conference Series on Public Policy*, volume 48, (pp. 145–194). Elsevier.

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

Reichling, F. & Whalen, C. (2012). Assessing the short-term effects on output of changes in federal fiscal policies. *Congressional Budget Office, Washington, DC.*

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, *33*(3), 649–665.

Schmitt-Grohé, S. & Uribe, M. (2012). What's news in business cycles. *Econometrica*, *80*(6), 2733–2764.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, *48*(1), 1–48.

Sims, C. A. (2000). Solving Linear Rational Expectations Models. Unpublished working paper.

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

Stock, J. H. & Watson, M. W. (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal*, *128*(610), 917–948.

Werning, I. (2016). Incomplete Markets and Aggregate Demand. Unpublished working paper.

Winberry, T. (2018). Lumpy Investment, Business Cycles, and Stimulus Policy. Unpublished working paper.

Wolf, C. K. (2019). The Missing Intercept in Cross-Regional Regressions. Working Paper.

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

Online Appendix for:
The Missing Intercept:
A Demand Equivalence Approach

This online appendix contains supplemental material for the article "The Missing Intercept: A Demand Equivalence Approach". I provide (i) further details for the various structural models used in the paper, (ii) several additional results on exact demand equivalence, (iii) details on the empirical micro and macro experiments used to construct my semi-structural counterfactuals, (iv) results on approximation accuracy in models with approximate demand equivalence, and (v) a third application of my empirical methodology, to income redistribution. The end of this appendix contains further proofs and auxiliary lemmas.

Any references to equations, figures, tables, assumptions, propositions, lemmas, or sections that are not preceded "B." – "G." refer to the main article.

Contents

B Model details **53**

B.1 Rest of the economy and equilibrium definition 53
B.2 Spender-saver RBC model 57
B.3 Estimated HANK model 61
B.4 Two-asset model 66
B.5 Extension to multiple sectors 68
B.6 Extension to weak wealth effects 70

C Further results on demand equivalence **72**

C.1 General transition paths 72
C.2 Generic consumption demand shifters 72
C.3 Exact equivalence beyond the baseline model 73
C.4 Range of outcomes for the "missing intercept" 75
C.5 Correcting for wealth effects in labor supply 77
C.6 Nested models for investment demand equivalence 78

D Empirical appendix **79**

D.1 Direct response: micro consumption elasticities 79
D.2 Direct response: micro investment elasticities 80
D.3 The missing intercept: VAR estimation 81
D.4 Joint Uncertainty 84

E Approximation accuracy **85**

E.1 Random parameter draws 85
E.2 Other estimated business-cycle models 86
E.3 Labor supply 88
E.4 Interest rates 89
E.5 Useful government spending 90
E.6 Multi-sector economies 91
E.7 Imperfect demand matching 92

F Application: income redistribution **94**

G Further proofs and auxiliary lemmas 96
G.1 Auxiliary lemma for Proposition 1 96
G.2 Proof of Proposition 1 96
G.3 Proof of Lemma A.1 97
G.4 Proof of Proposition 3 98
G.5 Proof of Proposition 4 99
G.6 Auxiliary lemma for Proposition A.1 100
G.7 Proof of Proposition A.1 101
G.8 Proof of Corollary C.1 102
G.9 Proof of Proposition D.1 102

B Model details

This appendix provides additional details on the baseline structural model underlying my demand equivalence results. In Appendix B.1 I outline the full model and offer a formal definition of equilibrium transition paths. Appendices B.2 and B.3 then discuss (i) the simple spender-saver RBC model of Section 2.3 and (ii) the rich estimated HANK model. Finally, in Appendices B.4 to B.6, I sketch various model extensions that break the exclusion restrictions required for demand equivalence.

B.1 Rest of the economy and equilibrium definition

Recall that the model is populated by households, firms, and the government. Whenever there is no risk of confusion, I replace the full decision problems of agents by simple conditions characterizing their actual optimal behavior. I do so because many of the problems considered here (in particular those of price-setting entities) are notationally involved, but at the same time extremely well-known and so require no repetition.

Households. The household consumption-savings problem was described in detail in Section 2.1. It remains to specify the problem of a wage-setting union k. A union sets wages and labor to maximize weighted average utility of its members, taking as given optimal consumption-savings behavior of each individual member household, exactly as in Auclert et al. (2018). Following the same steps as those authors, it can be shown that optimal union behavior is summarized by a standard non-linear wage-NKPC:[30]

$$\begin{aligned}\pi_t^w(1+\pi_t^w) &= \frac{\varepsilon_w}{\theta_w}\ell_t^h\Big[\int_0^1 \Big\{-u_\ell(c_{it},\ell_t^h) - \frac{\varepsilon_w-1}{\varepsilon_w}(1-\tau_\ell)w_t e_{it} u_c(c_{it},\ell_t^h)\Big\}di\Big] \\ &\quad +\beta\pi_{t+1}^w(1+\pi_{t+1}^w) \end{aligned} \tag{B.1}$$

where $1+\pi_t^w = \frac{w_t}{w_{t-1}} \times \frac{1}{1+\pi_t}$, ε_w is the elasticity of substitution between different kinds of labor, and θ_w denotes the Rotemberg adjustment cost. Given prices $(\boldsymbol{\pi}, \mathbf{w})$ as well as a consumption path $\mathbf{c}$, (B.1) provides a simple restriction on total labor supply $\boldsymbol{\ell}^h$.[31] Note that, without idiosyncratic labor productivity risk and so common consumption $c_{it} = c_t$, the derived wage-

[30]For notational simplicity, in the derivation of this wage-NKPC assume that $\beta_i = \beta$ for all i. The generalization to heterogeneous β_i's is conceptually straightforward, but notationally cumbersome.

[31]In the special case $\theta_w \to \infty$, equation (B.1) is vacuous, so then I instead simply assume that $\boldsymbol{\ell}^h = \boldsymbol{\ell}^f$.

NKPC (B.1) is to first order identical to the standard specification in Erceg et al. (2000). An extension to partially indexed wages, as in Smets & Wouters (2007) or Justiniano et al. (2010), is straightforward and omitted in the interest of notational simplicity.

Together, the consumption-savings problem and (B.1) characterize optimal household and union behavior. I assume that the solutions to each problem exist and are unique, and summarize the solution in terms of aggregate consumption, saving and union labor supply functions $\mathbf{c}(\mathbf{s}^h, \boldsymbol{\varepsilon})$, $\mathbf{b}^h(\mathbf{s}^h, \boldsymbol{\varepsilon})$, and $\boldsymbol{\ell}^h(\mathbf{s}^u)$, where $\mathbf{s}^h = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \boldsymbol{\ell}, \boldsymbol{\tau}_e, \mathbf{d})$ and $\mathbf{s}^u = (\boldsymbol{\pi}, \mathbf{w}, \mathbf{c})$.[32] In particular, the union problem gives

$$\hat{\boldsymbol{\ell}}_{\varepsilon}^{PE} \equiv \boldsymbol{\ell}^h(\bar{\boldsymbol{\pi}}, \bar{\mathbf{w}}, \mathbf{c}(\bar{\mathbf{s}}^h; \boldsymbol{\varepsilon})) - \bar{\boldsymbol{\ell}}^h$$

For my theoretical equivalence results, I will impose the high-level assumption that all of those infinite-dimensional functions are at least once differentiable in their arguments.

Firms. I first study the problem of each of the three types of firms in isolation. I assume that all firms discount at the common rate $1 + r_t^b \equiv \frac{1+i_{t-1}^b}{1+\pi_t}$.[33]

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

$$\max_{\{d_{jt}^I, y_{jt}, \ell_{jt}, k_{jt}, i_{jt}, u_{jt}, b_{jt}^f\}} \quad \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \left(\prod_{q=0}^{t-1} \frac{1}{1+r_q^b} \right) d_{jt}^I \right]$$

such that

$$\begin{aligned}
d_{jt}^I &= \underbrace{p_t^I y_{jt} - w_t \ell_{jt}}_{\pi_{jt}} - \phi(k_{jt}, k_{jt-1}, i_{jt}, i_{jt-1}) - b_{jt}^f + \frac{1+i_{t-1}^b}{1+\pi_t} b_{jt-1}^f \\
y_{jt} &= y(e_{jt}, u_{jt} k_{jt-1}, \ell_{jt}) \\
i_{jt} &= k_{jt} - [1 - \delta(u_{jt})] k_{jt-1} \\
-b_{jt}^f &\leq \Gamma(k_{jt-1}, k_{jt}, \pi_{jt}) \\
d_{jt}^I &\geq \underline{d}
\end{aligned}$$

[32] Formally, the input to the union problem is a "virtual" consumption aggregate (see Auclert et al. (2018) and the discussion surrounding (B.22)). In a slight abuse of notation, the dependence on $\mathbf{c}$ in the equations here is a shorthand for dependence on overall household consumption decisions given $(\mathbf{s}^h, \boldsymbol{\varepsilon})$.

[33] Along a perfect foresight transition path, discounting at $1 + r_t^b$ is equivalent to discounting at the (common) stochastic discount factor of all households with strictly positive asset holdings.

The physical adjustment cost function $\phi(\bullet)$ is general: it may be convex and continuously differentiable, but it may also feature a fixed-cost component or partial irreversibility. Firms can vary capital utilization, with higher utilization leading to faster depreciation, i.e. $\delta'(\bullet) > 0$. The solution to the firm problem gives optimal production $\mathbf{y}(\bullet)$, labor demand $\boldsymbol{\ell}^f(\bullet)$, investment $\mathbf{i}(\bullet)$, intermediate goods producer dividends $\mathbf{d}^I(\bullet)$, capital utilization rates $\mathbf{u}(\bullet)$ and liquid corporate bond savings $\mathbf{b}^f(\bullet)$ as a function of nominal returns $\mathbf{i}^b$, inflation $\boldsymbol{\pi}$, wages $\mathbf{w}$, and the intermediate goods price $\mathbf{p}^I$.

2. *Retailers.* A unit continuum of retailers purchases the intermediate good at price p_t^I, costlessly differentiates it, and sells it on to a final goods aggregator. Price setting is subject to a Rotemberg adjustment cost. As usual, optimal retailer behavior gives rise to a standard NKPC as a joint restriction on the paths of inflation and the intermediate goods price. In log-linearized form:

$$\hat{\tilde{\pi}}_t = \underbrace{\frac{\varepsilon_p}{\theta_p}\frac{\varepsilon_p - 1}{\varepsilon_p}}_{\kappa_p} \times \hat{\tilde{p}}_t^I + \beta\hat{\tilde{\pi}}_{t+1}$$

where ε_p denotes the substitutability between different kinds of retail goods, and θ_p denotes the Rotemberg adjustment cost. In an equivalent (to first-order) Calvo formulation, the slope of the NKPC instead is given as

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

where $1 - \phi_p$ is the probability of a price re-set. A further extension to partially indexed prices, as in Smets & Wouters (2007) or Justiniano et al. (2010), is straightforward and omitted in the interest of notational simplicity. Total dividend payments of retailers are

$$d_t^R = (1 - p_t^I)y_t$$

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

Total dividend payments by the corporate sector are given as

$$d_t = d_t^I + d_t^R$$

Using the restriction on the intermediate goods price implied by optimal retailer behavior, aggregate dividends can thus be obtained solely as a function of $\mathbf{s}^f = (\mathbf{i}^b, \mathbf{w}, \boldsymbol{\pi})$.

We can now summarize the aggregate firm sector simply through a set of optimal production, labor hiring, investment, dividend pay-out and bond demand functions, $\mathbf{y} = \mathbf{y}(\mathbf{s}^f; \boldsymbol{\varepsilon})$, $\boldsymbol{\ell}^f = \boldsymbol{\ell}^f(\mathbf{s}^f; \boldsymbol{\varepsilon})$, $\mathbf{i} = \mathbf{i}(\mathbf{s}^f; \boldsymbol{\varepsilon})$, $\mathbf{d} = \mathbf{d}(\mathbf{s}^f; \boldsymbol{\varepsilon})$ and $\mathbf{b}^f = \mathbf{b}^f(\mathbf{s}^f; \boldsymbol{\varepsilon})$, as well as a restriction on the aggregate path of inflation, $\boldsymbol{\pi} = \boldsymbol{\pi}(\mathbf{s}^f; \boldsymbol{\varepsilon})$, where $\mathbf{s}^f = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w})$. As before, I will assume that these aggregate firm sector-level functions are at least once differentiable in their arguments.

Government. The fiscal authority was discussed in detail in Section 2.1; furthermore, in Appendix B.3, I give an example of a concrete fiscal rule. It remains to describe central bank behavior. In line with standard empirical practice I assume that the nominal rate on bonds i^b is set according to the conventional Taylor rule

$$\hat{\tilde{i}}_t^b = \rho_m \hat{\tilde{i}}_{t-1}^b + (1 - \rho_m) \left(\phi_\pi \hat{\tilde{\pi}}_t + \phi_y \hat{\tilde{y}}_t + \phi_{dy} \hat{\tilde{y}}_{t-1} \right)$$

A generalization to feature a notion of potential output, as in Justiniano et al. (2010), is straightforward.

Market-Clearing. Equating liquid asset demand from households and intermediate goods producers, as well as liquid asset supply from the government, we get

$$b_t^h + b_t^f = b_t$$

Equating labor demand and supply:

$$\ell_t^f = \ell_t^h$$

Finally, aggregating all household, firm and government budget constraints, we obtain the aggregate output market-clearing condition[34]

$$c_t + i_t + g_t = y_t$$

Equilibrium. All results in this paper rely on the following equilibrium definition.

[34] So as to not excessively clutter market-clearing conditions with various adjustment cost terms, I assume that adjustment costs are ex-post rebated lump-sum back to the agents facing the adjustment costs. Of course, all subsequent equivalence results are unaffected by this rebating. An alternative interpretation is that adjustment costs are instead just perceived utility costs, as in Auclert et al. (2018).

Definition 2. *Given initial distributions* $\mu_0^h = \bar{\mu}^h$ *and* $\mu_0^f = \bar{\mu}^f$ *of households and intermediate goods producers over their idiosyncratic state spaces, an initial real wage* $w_{-1} = \bar{w}$*, price level* p_{-1}*, and real government debt* $b_{-1} = \bar{b}$*, as well as exogenous shock paths* $\{\varepsilon_t\}_{t=0}^{\infty}$*, a recursive competitive equilibrium is a sequence of aggregate quantities* $\{c_t, \ell_t^h, \ell_t^f, b_t^h, b_t^f, b_t, y_t, i_t, d_t, k_t, g_t, \tau_t\}_{t=0}^{\infty}$ *and prices* $\{\pi_t, i_t^b, w_t\}_{t=0}^{\infty}$ *such that:*

1. Household Optimization. *Given prices and government rebates, the paths of aggregate consumption* $\mathbf{c} = \mathbf{c}(\mathbf{s}^h; \boldsymbol{\varepsilon})$*, labor supply* $\boldsymbol{\ell}^h = \boldsymbol{\ell}^h(\mathbf{s}^u)$*, and asset holdings* $\mathbf{b}^h = \mathbf{b}^h(\mathbf{s}^h; \boldsymbol{\varepsilon})$ *are consistent with optimal household and wage union behavior.*

2. Firm Optimization. *Given prices, the paths of aggregate production* $\mathbf{y} = \mathbf{y}(\mathbf{s}^f; \boldsymbol{\varepsilon})$*, investment* $\mathbf{i} = \mathbf{i}(\mathbf{s}^f; \boldsymbol{\varepsilon})$*, capital* $\mathbf{k}$*, labor demand* $\boldsymbol{\ell}^f = \boldsymbol{\ell}^f(\mathbf{s}^f; \boldsymbol{\varepsilon})$*, dividends* $\mathbf{d} = \mathbf{d}(\mathbf{s}^f; \boldsymbol{\varepsilon})$ *and asset holdings* $\mathbf{b}^f = \mathbf{b}^f(\mathbf{s}^f; \boldsymbol{\varepsilon})$ *are consistent with optimal firm behavior. Furthermore, the path of inflation is consistent with optimal retailer behavior.*

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

4. Market Clearing. *The goods market clears,*

$$c_t + i_t + g_t = y_t$$

the bond market clears,

$$b_t^h + b_t^f = b_t$$

and the labor market clears,

$$\ell_t^h = \ell_t^f$$

for all $t = 0, 1, 2, \ldots$.

B.2 Spender-saver RBC model

The simple model is a special case of the rich benchmark model of Section 2.1. For convenience, I here explicitly state the equations characterizing the model equilibrium.

Model Sketch. A mass $\lambda \in (0,1)$ of households are spenders, indexed by h. They inelastically supply labor and receive lump-sum transfers τ_{ht}; since $\beta_h = 0$, their consumption satisfies

$$c_{ht} = w_t \bar{\ell} + \tau_{ht} \tag{B.2}$$

The residual fraction of households are savers, indexed by r. Since there are no adjustment costs or portfolio restrictions, I can characterize the consumption-savings problem of savers as a simple one-asset problem with exogenous dividend receipts:

$$\max_{\{c_{rt}, b_t\}} \quad \sum_{t=0}^{\infty} \beta^t \log(c_{rt})$$

such that

$$c_{rt} + b_t = (1 + r_t) b_{t-1} + d_t + w_t \bar{\ell} + \tau_{rt}$$

Optimal saver behavior is characterized by the Euler equation

$$c_{rt}^{-1} = \beta(1 + r_{t+1}) c_{rt+1}^{-1} \tag{B.3}$$

A single representative firm chooses investment to maximize the present value of dividend payments to savers, discounted at the real rate faced by savers (and so their stochastic discount factor, to first order). Its problem is

$$\max_{\{d_t, y_t, k_t, \ell_t, i_t\}} \quad \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t-1} \frac{1}{1 + r_s} \right) d_t$$

such that

$$d_t = y_t - i_t - w_t \ell_t \tag{B.4}$$

$$y_t = k_{t-1}^{\alpha} \ell_t^{1-\alpha} \tag{B.5}$$

$$i_t = k_t \tag{B.6}$$

where the final relation uses full depreciation of the capital good. From now on I use that, in equilibrium, $\ell_t = \bar{\ell}$. Optimal firm investment is characterized by the relation

$$1 + r_t = \alpha k_t^{\alpha - 1} \bar{\ell}^{1-\alpha} \tag{B.7}$$

and wages satisfy

$$w_t = (1-\alpha)k_{t-1}^{\alpha}\bar{\ell}^{-\alpha} \tag{B.8}$$

The government consumes an exogenously determined amount of the final good,

$$g_t = \varepsilon_{gt} \tag{B.9}$$

and sets rebates to spenders as

$$\tau_{ht} = \bar{\tau}_h + \frac{1}{\lambda}\varepsilon_{\tau t} \tag{B.10}$$

The scaling factor $\frac{1}{\lambda}$ is chosen to ensure that the amount of stimulus rebate given to spenders overall is independent of the mass of spenders in the economy. Expenditure is financed fully through contemporaneous lump-sum taxation on savers, so

$$\begin{aligned} \lambda\tau_{ht} + (1-\lambda)\tau_{rt} + g_t &= 0 && \text{(B.11)} \\ b_t &= 0 && \text{(B.12)} \end{aligned}$$

Note that, in the notation of Section 2.1, $\hat{\tau}_{ht} = \frac{1}{\lambda}\hat{\tau}_{xt}$ and $\hat{\tau}_{rt} = -\frac{1}{1-\lambda}\hat{\tau}_{et}$. Finally, aggregate market-clearing dictates that

$$y_t = c_t + i_t + g_t \tag{B.13}$$

where

$$c_t = \lambda c_{ht} + (1-\lambda)c_{rt} \tag{B.14}$$

A recursive competitive equilibrium for aggregate prices and quantities $\{c_t, c_{ht}, c_{rt}, r_t, w_t, d_t, y_t, k_t, i_t, g_t, \tau_{ht}, \tau_{rt}, b_t\}$ is fully characterized by the relations (B.2) - (B.14).

Without loss of generality, and to simplify the algebra, I normalize

$$\bar{\ell} = \left[(\alpha\beta)^{\frac{\alpha}{1-\alpha}} - (\alpha\beta)^{\frac{1}{1-\alpha}}\right]^{-1}$$

which ensures that $\bar{c} = 1$ and $\bar{y} = 1/(1-\alpha\beta)$. I furthermore, and also for notational simplicity, assume that steady-state rebates $\bar{\tau}_h$ are such that the steady-state consumption levels of spenders and savers are equalized.

Log-linear Solution. It is straightforward to characterize the (log-linear) solution of this model in closed form. Log-linearizing (B.13), and using (B.9), we get

$$(1-\alpha\beta)\hat{\hat{c}}_t + \alpha\beta\hat{\hat{k}}_t + \frac{1}{\bar{y}}\varepsilon_{gt} = \alpha\hat{\hat{k}}_{t-1} \tag{B.15}$$

Log-linearizing (B.3) and plugging into (B.7), we get

$$\hat{\hat{c}}_{rt+1} - \hat{\hat{c}}_{rt} = (\alpha-1)\hat{\hat{k}}_t \tag{B.16}$$

Expressing saver consumption in terms of aggregate and spender consumption using (B.14), solving for spender consumption in terms of the rebate shock and capital using (B.2) and (B.8), and plugging into (B.16), we get

$$\frac{1}{1-\lambda}\hat{\hat{c}}_{t+1} - \frac{\lambda}{1-\lambda}(\alpha\hat{\hat{k}}_t + \frac{1}{\lambda}\varepsilon_{\tau t+1}) - \frac{1}{1-\lambda}\hat{\hat{c}}_t + \frac{\lambda}{1-\lambda}(\alpha\hat{\hat{k}}_{t-1} + \frac{1}{\lambda}\varepsilon_{\tau t}) = (\alpha-1)\hat{\hat{k}}_t \tag{B.17}$$

All other equilibrium objects are immediately determined from the remaining equilibrium relations, so the equilibrium is fully characterized by (B.15) and (B.17). Plugging (B.15) into (B.17) to eliminate consumption, we get the single equation

$$\begin{aligned} &\frac{1}{1-\lambda}\left[\frac{\alpha}{1-\alpha\beta}\hat{\hat{k}}_t - \frac{\alpha\beta}{1-\alpha\beta}\hat{\hat{k}}_{t+1} - \varepsilon_{gt+1}\right] - \frac{\lambda}{1-\lambda}\left(\alpha\hat{\hat{k}}_t + \frac{1}{\lambda}\varepsilon_{\tau t+1}\right) \\ &\quad - \frac{1}{1-\lambda}\left[\frac{\alpha}{1-\alpha\beta}\hat{\hat{k}}_{t-1} - \frac{\alpha\beta}{1-\alpha\beta}\hat{\hat{k}}_t - \varepsilon_{gt}\right] + \frac{\lambda}{1-\lambda}\left(\alpha\hat{\hat{k}}_{t-1} + \frac{1}{\lambda}\varepsilon_{\tau t}\right) = (\alpha-1)\hat{\hat{k}}_t \end{aligned} \tag{B.18}$$

I solve the model exploiting the well-known equivalence between perfect foresight and first-order perturbation solutions. I thus treat (B.18) as a second-order expectational difference equation (replacing all variables dated $t+1$ by their expectation), and find its unique stable solution. To this end conjecture that

$$\hat{\hat{k}}_t = \theta_k\hat{\hat{k}}_{t-1} + \omega_d(\varepsilon_{gt} + \varepsilon_{\tau t}) \tag{B.19}$$

Plugging in and matching coefficients, we find that the guess is confirmed,[35] with

$$\theta_k = \alpha$$

[35]It is straightforward to verify existence and uniqueness of the equilibrium following the arguments in Blanchard & Kahn (1980) or Sims (2000).

$$\omega_d = -\frac{1-\alpha\beta}{1-\lambda(1-\alpha\beta)}$$

Plugging this back into (B.15), we get

$$\hat{\hat{c}}_t = \alpha\hat{\hat{k}}_{t-1} + \frac{\alpha\beta}{1-\lambda(1-\alpha\beta)} \times (\varepsilon_{gt} + \varepsilon_{\tau t}) - \varepsilon_{gt} \tag{B.20}$$

Demand equivalence follows immediately.

B.3 Estimated HANK model

Much of my analysis builds on an estimated one-asset HANK model, featuring a consumption-savings problem under imperfect insurance embedded into an otherwise standard medium-scale DSGE environment. This section provides details on the model, the solution algorithm, my approach to likelihood-based estimation, and the final parameterization used to generate the results in Section 4 (as well as the simpler illustration in Section 2.4).

Model Outline. The model is a particular variant of the rich baseline environment outlined in Section 2.1, consistent with Assumptions 1 and 2 but violating Assumption 3 (strong wealth effects and imperfectly rigid wages).

Households have separable preferences over consumption and labor,

$$u(c,\ell) = \frac{c^{1-\gamma}-1}{1-\gamma} - \chi\frac{\ell^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}},$$

and discount the future at rate β. The log-linearized wage-NKPC then takes the form

$$\hat{\hat{\pi}}_t^w = \kappa_w \times \left[\frac{1}{\varphi}\hat{\hat{\ell}}_t - (\hat{\hat{w}}_t - \gamma\hat{\hat{c}}_t^*)\right] + \beta\hat{\hat{\pi}}_{t+1}^w \tag{B.21}$$

where κ_w is a function of model parameters and c_t^* satisfies

$$c_t^* \equiv \left[\int_0^1 e_{it}c_{it}^{-\gamma}di\right]^{-\frac{1}{\gamma}} \tag{B.22}$$

To facilitate comparison with the standard New Keynesian business-cycle literature, I will throughout replace this virtual consumption aggregate with aggregate consumption c_t, thus giving me an entirely standard wage-NKPC (as in Hagedorn et al., 2019); results are, however, almost unchanged if I use c_t^* instead.[36] I furthermore slightly generalize the model of Section 2.1 to allow for stochastic death with probability ξ. All households receive identical lump-sum transfers τ_t but are heterogeneous in dividend payment receipts. In particular, I assume that the most productive households receive larger dividend payments, so that stock wealth is effectively concentrated among a small share of households.

The intermediate goods production block – in particular the production function $y(\bullet)$, the investment adjustment cost function $\phi(\bullet)$, and the capacity utilization depreciation rate $\delta(\bullet)$ – is set up exactly as in Justiniano et al. (2010). Relative to the model outlined in Section 2.1, I omit the transfer shocks, but then add structural shocks to output and investment productivity, monetary policy, household impatience, and wage mark-ups to complement the already included government spending shocks. All shocks are assumed to follow simple AR(1) processes. Finally, for purposes of the model estimation, I assume that

$$\hat{\tau}_{et} = -(1-\rho_\tau) \times \hat{b}_{t-1} \tag{B.23}$$

The endogenous part of transfers is cut in response to increases in $\hat{b}_t$. For plots of approximate equivalence results, I let transfer shocks be financed using this rule, and then assume that government spending shocks are financed using the same (potentially scaled) intertemporal tax *profile*, consistent with Assumption 2. The partial equilibrium financing paths of the two shocks will thus always be the same. Since households spend most of the rebate immediately, results are very similar if I instead simply use the rule (B.23) for all shocks.

Steady-State Calibration. Solving for the deterministic steady-state of the model requires specification of several parameters. On the household side, I need to set income risk and share endowment processes, specify preferences, and choose liquid borrowing limits as well as the substitutability between different kinds of labor. On the firm side, I need to specify production and investment technologies, as well as the substitutability between different kinds of goods. Finally, on the government side, I need to set taxes, transfers, and total bond supply. Government spending is then backed out residually. My preferred parameter values and associated calibration targets are displayed in Table B.1.

The first block shows parameter choices on the household side. For income risk, I adopt the 33-state specification of Kaplan et al. (2018), ported to discrete time. For share endow-

[36] Using c_t has the advantage that union wage-setting is not affected by the *distributional* implications of the shock. However, since labor is largely demand-determined in the short run, even those distributional considerations have little effect on equilibrium hours worked.

Parameter	Description	Value	Target	Model	Data
Households					
ρ_e, σ_e	Income Risk	-	Kaplan et al. (2018)	-	-
$\underline{e}^p$, χ_0, χ_1	Dividend Endowment	-	Illiquid Wealth Shares	-	-
β	Discount Rate	0.97	B/Y	1.04	1.04
r^b	Average Return	0.01	Annual Rate	0.04	0.04
ξ	Death Rate	1/180	Average Age	45	45
γ	Preference Curvature	1	Standard		
φ	Labor Supply Elasticity	1	Standard		
ε_w	Labor Substitutability	10	Standard		
$\underline{b}$	Borrowing Limit	0	McKay et al. (2016)		
Firms					
α	Capital Share	0.2	Justiniano et al. (2010)		
δ	Depreciation	0.016	Total Wealth/Y	10.64	10.64
ε_p	Goods Substitutability	16.67	Profit Share	0.06	0.06
Government					
τ_ℓ	Labor Tax	0.3	Average Labor Tax	0.30	0.30
τ/Y	Transfer Share	0.05	Transfer Share	0.05	0.05
B/Y	Liquid Wealth Supply	1.04	Government Debt/Y	1.04	1.04

Table B.1: HANK model, steady-state calibration.

ments, I assume that

$$d_{it} = \begin{cases} 0 & \text{if } e_{it}^p \leq \underline{e}^p \\ \chi_0 (e_{it}^p - \underline{e}^p)^{\chi_1} \times d_t & \text{otherwise} \end{cases}$$

where e_{it}^p is the permanent component of household i's labor productivity. I set the cutoff $\underline{e}^p$ so that the bottom half of households receive no dividends (consistent with the illiquid wealth distribution in the 2016 SCF), χ_1 so that the top 10 per cent of households receive the same share of total dividends (and so total illiquid wealth) as in Kaplan et al. (2018), and then back out χ_0 to ensure that $\int_0^1 d_{it} di = d_t$.[37] Next, I set the average return on (liquid) assets in line

[37] A natural alternative assumption would be to set $d_{it} = d_t$, as in McKay et al. (2016) or Auclert et al. (2018). This alternative choice of course changes impulse responses, but has little effect on the accuracy of the demand equivalence approximation.

with standard calibrations of business-cycle models. The discount and death rates are then disciplined through targets on the total amount of liquid wealth as well as average household age. For my baseline model, I further assume that households cannot borrow. All remaining parameters are set in line with conventional practice. The second block shows parameter choices on the firm side. I discipline the Cobb-Douglas production function $y = k^{\alpha}\ell^{1-\alpha}$ by setting α in line with Justiniano et al. (2010), identify goods substitutability by targeting the profit share, and finally back out the depreciation rate from my target of total wealth (and so corporate sector valuation) in the economy as a whole.[38] The third block informs the fiscal side of the model. The average government tax take, transfers, and debt issuance are all set in line with direct empirical evidence.

Importantly, because household self-insurance is severely limited, the average MPC in the economy is high, around 28% out of an unexpected 500$ income gain. As a result, the model can replicate the large (yet gradual) empirically observed consumption response to income tax rebates, as argued previously in Auclert et al. (2018).

Dynamics: Computational Details. I solve the model using a variant of the popular Reiter method (Reiter, 2009). In particular, I use a mixture of the methods developed in Ahn et al. (2017) and Bayer & Luetticke (2020) to reduce the dimensionality of the state space. Without dimensionality reduction, the number of idiosyncratic household-level states is too large to allow likelihood-based estimation. With dimensionality reduction, the number of states is reduced to around 300, making estimation feasible.

While the estimation relies on standard state-space methods, my displays of exact and approximate demand equivalence are instead computed in sequence-space. For the perfect foresight sequence-space solution I proceed similarly to Boppart et al. (2018) and Auclert et al. (2019); throughout, I verify that the results are not sensitive to the choice of transition endpoint T or the step size used to approximate the Jacobians of the linearized system.

Dynamics: Estimation. With two exceptions, I estimate the remaining model parameters (which exclusively govern dynamics around the deterministic steady state) using standard likelihood methods, as in An & Schorfheide (2007). The set of observables is: aggregate output (y), consumption (c), investment (i), inflation (π), the short-term nominal interest rate (r_t^n), and total hours worked (ℓ). The construction of all series follows Justiniano et al.

[38]More conventional higher values of α change impulse responses, but do not break demand equivalence. Similarly, the results also remain accurate with the low value of α entertained in Auclert & Rognlie (2018).

(2010), and my sample period is 1960:Q1 – 2006:Q4.[39] Priors are reported in Table B.2.

The first exception is the transfer adjustment parameter ρ_τ; since I do not include data on government debt, this parameter would likely be poorly identified. I thus simply set $\rho_\tau = 0.85$, in line with the VAR evidence documented in Galí et al. (2007) and Appendix D.3. Second, as it is central to my approximate equivalence results, I directly discipline the degree of wage stickiness from micro data. Exploiting the standard first-order equivalence of Calvo price re-sets and Rotemberg adjustment costs, it is easy to show that the slope parameter of the wage-NKPC (B.21) can be equivalently written as

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

where $1 - \phi_w$ is the probability of wage adjustment in the quarter. I set the wage stickiness parameter consistent with the micro evidence in Grigsby et al. (2019) and Beraja et al. (2019), giving $\phi_w = 0.6$ – price re-sets every 2.5 quarters. Direct estimation of this parameter would instead suggest a much larger value, consistent with the findings of Justiniano et al. (2010) and other estimated New Keynesian models.

The results of the estimation are displayed in Table B.2. Since they are not relevant for my purposes here, I omit estimates of shock persistence and volatility; some brief remarks on those follow at the end. I find the posterior mode using the `csminwel` routine provided by Chris Sims; for accuracy of the demand equivalence approximation beyond the mode parameterization of the model, see the discussion in Appendix E.1.[40]

On the whole, the results are quite consistent with the parameter estimates in Justiniano et al. (2010). Relative to their rich framework, the two central changes in my model are, first, the introduction of uninsurable income risk, and second, the absence of habit formation. The first change ties consumption and income more closely together, while the second leads to less endogenous persistence and worsens the Barro-King puzzle (Barro & King, 1984). Jointly, these changes dampen the importance of impatience shocks as a driving force of consumption fluctuations, but also give a somewhat smaller role for investment efficiency shocks as a source of cyclical fluctuations. These findings are consistent with the intuition in Werning (2016) and the estimation results on the no-habit model in Justiniano et al. (2010). Ultimately, given the similarity in model environment and data sources, the similarity of

[39] I thank Brian Livingston for help in assembling the data.

[40] The optimization routine is available at `http://sims.princeton.edu/yftp/optimize/`.

Parameter	Description	Density	*Prior* Mean	Std	*Posterior* Mode
ϕ_p	Price Calvo Parameter	B	0.7	0.10	0.80
ζ	Capacity Utilization	N	5.00	1.00	4.54
κ	Investment Adjustment Cost	N	4.00	1.00	2.40
ρ_m	Taylor Rule Persistence	B	0.80	0.20	0.67
ϕ_π	Taylor Rule Inflation	N	2.00	0.10	1.95
ϕ_y	Taylor Rule Output	N	0.15	0.05	0.05
ϕ_{dy}	Taylor Rule Output Growth	N	0.15	0.05	0.30

Table B.2: HANK model, parameters governing dynamics, estimated using conventional likelihood-based methods. For the priors, N stands for Normal and B for Beta.

the resulting parameter estimates should not come as a surprise. A more serious estimation exercise on the effects of micro heterogeneity on macro fluctuations would also leverage the advantages afforded by time series of richer micro data, and is left for future work.

Simplified model. The simplified HANK model considered for the illustration in Section 2.4 is identical to the estimated model except for one change: I set $\phi_w = 1$. As a result, demand equivalence holds exactly.

B.4 Two-asset model

In this section I describe the two-asset HANK model discussed in Section 4.3. The results from the approximate demand equivalence check are reported in Appendix E.4.

Model sketch. Households invest in an illiquid asset with real return r^h and a liquid asset with real return $r^h - \kappa_b$, where $1 + r_t^h = \frac{1+i_{t-1}^h}{1+\pi_t}$. The household consumption-savings problem then is

$$\max_{\{c_{it}, b_{it}^h, a_{it}^h\}} \quad \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_i^t u(c_{it}, \ell_{it}) \right]$$

such that

$$c_{it} + b_{it}^h + a_{it}^h = (1 - \tau_\ell) w_t e_{it} \ell_{it} + \left[\frac{1 + i_{t-1}^h}{1 + \pi_t} - \kappa_b \right] b_{it-1}^h + \frac{1 + i_{t-1}^h}{1 + \pi_t} a_{it-1}^h + \phi_a(a_{it}^h, a_{it-1}^h; \zeta_{it}) + \tau_{it}$$

and

$$b_{it}^h \geq \underline{b}, \quad a_{it}^h \geq \underline{a}$$

where $\phi_a(\bullet, \bullet; \zeta)$ is the adjustment cost function for illiquid asset holdings. Similar to Bayer et al. (2019), I assume that a randomly chosen fraction η of households can freely adjust their illiquid wealth holdings ($\zeta = 1$), while the remaining households cannot adjust ($\zeta = 0$). The adjustment cost function can then be written as

$$\phi_a(a', a) = \begin{cases} 0 & \text{if } \zeta = 1 \\ \infty & \text{if } \zeta = 0 \end{cases}$$

Returns in the economy are determined as follows. Both liquid and illiquid assets are issued by a mutual fund, which in turn owns all government debt and all claims to corporate profits in the economy. Let $\omega_t \equiv b_t^h + b_t^f + a_t^h$ denote total funds managed by the mutual fund. Returns earned by the mutual fund i_t^m then satisfy

$$\omega_{t-1} \times \frac{1 + i_{t-1}^m}{1 + \pi_t} = b_{t-1} \frac{1 + i_{t-1}^b}{1 + \pi_t} + (d_t + v_t)$$

where v_t denotes the value of the corporate sector, which by arbitrage satisfies

$$\frac{1 + i_{t-1}^b}{1 + \pi_t} = \frac{v_t + d_t}{v_{t-1}}$$

except possibly at $t = 0$. I assume that the mutual fund is competitive, and faces intermediation costs κ_b to make assets liquid. It follows immediately that we must have $i_t^h = i_t^m$.

The rest of the economy is unchanged; in particular, firms still discount at $\frac{1+i_{t-1}^b}{1+\pi_t}$, which in the absence of aggregate risk is equivalent to discounting at $\frac{1+i_{t-1}^m}{1+\pi_t} = \frac{1+i_{t-1}^h}{1+\pi_t}$. The only change to Definition 2 is the new asset market-clearing condition:

$$b_t^h + b_t^f + a_t^h = b_t + v_t$$

Parameterization. For simplicity, I keep all parameters governing dynamics identical to the estimated one-asset HANK model, and only re-calibrate the steady state. Table B.3 displays all parameters from the re-calibrated 2-asset model that are different from those displayed in Table B.1 for the benchmark one-asset model.

To provide a stringent test of the demand equivalence approximation, I set the wedge

Parameter	Description	Value	Target	Model	Data
Households					
η	Probability of Adjustment	0.15	A/Y	9.70	10.64
β	Discount Rate	0.98	B/Y	1.28	1.04
r^h	Return	0.015	Upper Bound		
κ_b	Liquid Wedge	0.0125	Upper Bound		
Firms					
δ	Depreciation	0.025	Firm Valuation		

Table B.3: 2-asset HANK model, steady-state calibration.

between returns on household deposits and government debt to be an (arguably implausible) 1.25 per cent *per quarter*. Given this large difference, I then choose the adjustment probability η to ensure a reasonable fit to total liquid and illiquid wealth in the U.S. economy.

B.5 Extension to multiple sectors

In this section I extend the baseline estimated HANK to feature multiple sectors, producing different consumption goods as well as a separate investment good. The results from the corresponding approximate demand equivalence check are reported in Appendix E.6.

Model sketch. The extended model departs from my one-sector baseline in the following ways. First, it features three goods – two consumption goods and an investment good. Households have preferences over a consumption basket c_{it}, which is given as a mix of the two individual consumption goods:

$$c_{it} = c_{i1t}^{\nu} c_{i2t}^{1-\nu}$$

I let the ideal price index of the consumption bundle be the numeraire of my economy, and denote the relative prices of two consumption goods by q_{1t} and q_{2t}. Investment is only possible using the economy's investment good, whose real relative price is denoted q_{It}. The government purchases each of the three goods, with potentially different spending multipliers for each, and the monetary authority responds to changes in consumer price inflation.

Second, total household labor supply ℓ_t^h is an aggregator of labor supply for each of the three goods in the economy:

$$\ell_t^h \equiv \left[\ell_{1t}^{\frac{\varphi+\mu}{\varphi}} + \ell_{2t}^{\frac{\varphi+\mu}{\varphi}} + \ell_{It}^{\frac{\varphi+\mu}{\varphi}} \right]^{\frac{\varphi}{\varphi+\mu}}$$

$\mu = 0$ corresponds to perfect labor mobility across the sectors, while $\mu = 1$ corresponds to perfect immobility, with all labor types entering separately into my particular choice of household utility functions. For each type of labor, labor supply is intermediated by a unit continuum of sticky-wage unions. Optimal union behavior then gives the three log-linearized wage-NKPCs; proceeding as in Appendix B.3, we find

$$\begin{aligned}\hat{\hat{w}}_{mt} = \frac{\beta}{1+\beta}\hat{\hat{w}}_{mt+1} - \kappa_w \left[\hat{\hat{w}}_{mt} - \left(\frac{1-\mu}{\varphi}\hat{\hat{\ell}}_t^h + \frac{\mu}{\varphi}\hat{\hat{\ell}}_{mt} \right) - \gamma \hat{\hat{c}}_t \right] \\ - \frac{1}{1+\beta}\hat{\hat{\pi}}_t + \frac{\beta}{1+\beta}\hat{\hat{\pi}}_{t+1} + \frac{1}{1+\beta}\hat{\hat{w}}_{mt-1}\end{aligned}$$

for $m = 1, 2, I$. Note that, with $\mu = 0$ (i.e., perfect labor mobility), wages in all sectors are at all times equalized. Overall, household i then receives $e_{it} w_t \ell_t$ worth of labor earnings, where w_t is the aggregated wage index.

Third, there are separate production sectors for each of the three goods. Briefly, I simply repeat the production sector of the estimated 1-asset HANK model three times, but with good-specific final prices q_{mt} and potentially heterogeneous capital shares α_m. All three sectors then purchase capital goods at price q_{It}, hire labor at cost w_{mt}, and sell their own good at real price q_{mt}.

PARAMETERIZATION. I build on the parameterization of the estimated HANK model of Section 4, with one notable difference: a smaller degree of nominal price rigidities. In the model, the probability of price re-sets governs relative price movements after a demand shock for a specific good. I have included measures of relative prices in my VARs and find little response, similar to Nakamura & Steinsson (2014); however, Ramey & Shapiro (1998) show that, after large government spending shocks that move output by almost 4 per cent, relative prices move by 2.5 per cent.[41] To be conservative, I choose a model calibration with $\phi_p = 0.6$, giving relative price responses consistent with this evidence.

Next, I set $\mu = 1$ (i.e., fully sector-specific labor). I set the average capital share $\bar{\alpha} \equiv$

[41] For my VAR analysis, I follow Ramey & Shapiro (1998) and – in a VAR with military spending forecast errors – include a measure of the relative price of manufacturing goods.

$\alpha_1 \frac{\bar{y}_1}{\bar{y}} + \alpha_2 \frac{\bar{y}_2}{\bar{y}} + \alpha_I \frac{\bar{y}_I}{\bar{y}} = 0.2$ (as in my baseline model), and then set relative labor shares as in Alonso (2017, Table 3.3), giving $\alpha_1 = 0.48$, $\alpha_2 = 0.04$, $\alpha_I = 0.17$. Finally, the fraction of labor in each of the three sectors is set so that their relative sizes are also data-consistent; again following Alonso (2017), this gives $\bar{y}_1/\bar{y} = 0.29$, $\bar{y}_2/\bar{y} = 0.48$ $\bar{y}_3/\bar{y} = 0.23$.

B.6 Extension to weak wealth effects

In this section I describe a variant of the estimated HANK model without unions, but with weak wealth effects in labor supply. The results from the corresponding approximate demand equivalence check are reported in Appendix E.3.

Model sketch. I make three changes to the baseline model. First, the economy is now populated by a *double* unit continuum of households – a unit continuum of *families* $f \in [0, 1]$, and a unit continuum of households $i \in [0, 1]$ for each f. Each family is a replica of the unit continuum of households in the benchmark model, but shock exposures may be heterogeneous across families. I will explain the purpose of this artificial construction momentarily. Second, there are no unions – each household decides on its own labor supply. Third, I change household preferences. Similar to Jaimovich & Rebelo (2009) and Galí et al. (2012), I assume that

$$u_{ft}(c_{ift}, \ell_{ift}) = \frac{c_{ift}^{1-\gamma} - 1}{1-\gamma} - \chi \theta_{ift} \frac{\ell_{ift}^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}}$$

where the preference shifter θ_{ift} satisfies[42]

$$\theta_{ift} = x_{ft}^{\gamma} \times c_{ift}^{-\gamma}$$

The variable x_{ft} is central. To jointly ensure (i) arbitrarily weak short-run wealth effects in labor supply, (ii) homogeneous wealth effects in the cross section of households (both consistent with the estimates in Cesarini et al. (2017)), and (iii) direct earnings responses showing up in cross-sectional regressions, I assume that

$$x_{ft} = x_{ft-1}^{1-\omega} \times c_{ft}^{\omega}$$

This preference specification is the simplest design with all three desired properties. First, by varying the parameter ω, I can control the strength of short-term wealth effects, exactly

[42] Households do not internalize the effect of their consumption on the shifter.

as in Galí et al. (2012). With $\omega = 0$ wealth effects are 0, and so Assumption 3 is satisfied. Second, solving for optimal household labor supply decisions, we get

$$\chi \ell_{ift}^{\frac{1}{\varphi}} = w_t x_{ft}^{-\gamma} \tag{B.24}$$

If all "families" are equally affected by the shock, then everyone's labor supply is identical, giving the desired homogeneity. Thus, for the first two requirements, the family construction is not necessary – we could simply replace c_{ft} by c_t, giving the natural heterogeneous-agent analogue of the preferences in Galí et al. (2012). But third, with heterogeneous family-level shock exposures, cross-sectional regressions as in Proposition 3 will pick up direct earnings responses.[43] In particular, let $\boldsymbol{\ell}^h = \boldsymbol{\ell}^h(\mathbf{w}, \mathbf{c})$ denote the mapping from wages and family consumption into family labor supply induced by (B.24). The micro regression estimand in (13) then satisfies

$$\hat{\mathbf{c}}_\tau^{PE} = \left(I - \frac{\partial \mathbf{c}}{\partial \boldsymbol{\ell}} \times \frac{\partial \boldsymbol{\ell}^h}{\partial \mathbf{c}}\right)^{-1} \times \left(\frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}} \cdot d\boldsymbol{\tau}\right) \tag{B.25}$$

For my accuracy checks, I simply match this regression estimand with an identical expansion in aggregate government spending.

Parameterization. The parameters related to the sticky-wage block of the baseline model are irrelevant for this model variant; all other parameters are set exactly as before. Finally, the single new model parameter is ω. To ensure consistency with empirical evidence, I set $\omega = 0.043$. As in Cesarini et al. (2017), this specification results in a peak measured cross-sectional (partial equilibrium) labor supply response of around 4$ for every 100$ response in consumption.

[43] I could have used a similar family construction for the union model. Without changes in preferences, however, this model would be inconsistent with empirical evidence on the weakness of wealth effects.

C Further results on demand equivalence

This appendix collects several supplementary results to my discussions of exact consumption and investment demand equivalence. In Appendix C.1 I show that my arguments apply without change to perturbations around arbitrary transition paths. Appendices C.2 and C.3 then emphasize the generality of consumption demand equivalence by considering a larger family of shocks and models, and in Appendix C.4 I illustrate the range of general equilibrium outcomes consistent with exact equivalence. Appendix C.5 then shows that, instead of ignoring the labor supply channel, researchers can use a generalized variant of my methodology to explicitly account for it. Finally, in Appendix C.6, I argue that many popular structural models of investment are nested by the investment demand equivalence result.

C.1 General transition paths

All equivalence results in this paper are stated for transition paths starting at the deterministic steady state. However, it is immediate from the proof of Proposition 2 (and similarly that of Proposition A.1) that nothing in my logic hinges on the starting point. Intuitively, the crucial restriction in my arguments is that they are valid to first order, but not that they only apply to particular expansion points. All results can thus equivalently be interpreted as applying to first-order perturbation solutions around a given (deterministic) transition path.

For example, initial states μ_0^h, μ_0^f, w_{-1} and p_{-1} could be such that the economy is in a deep recession or brisk expansion. My equivalence results would then apply to deviations from the unshocked transition path of the economy back to steady state. These deviations need not agree with impulse responses at steady state, but they remain tied together across different kinds of demand shocks.

C.2 Generic consumption demand shifters

Consumption demand equivalence extends without change to generic shifters of consumption demand. To establish this claim, I augment the baseline model to feature fluctuations in household patience as a simple reduced-form stand-in for various more plausibly structural shocks to household spending (e.g. changes in borrowing constraints, redistribution, ...).

The discount factor of every household is now subject to an additional common shifter

ζ_t, with $\boldsymbol{\zeta} = \boldsymbol{\zeta}(\boldsymbol{\varepsilon}_v)$, giving preferences as

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_i^t \zeta_t(\boldsymbol{\varepsilon}_v) u(c_{it}, \ell_{it}) \right] \tag{C.1}$$

For equivalence of private demand shocks ν and public spending shocks g I need to slightly modify my assumption on fiscal financing:

Assumption 2'. *Households and government borrow and lend at the same interest rate. The path of taxes and transfers used to finance a given public expenditure shock $\boldsymbol{\varepsilon}_\tau$ or $\boldsymbol{\varepsilon}_g$ depends only on the present value of the expenditure, not its time path. A spending path with zero net present value is purely deficit-financed, and so elicits no* direct *tax response.*

Note that impatience shocks – shocks that just shift the intertemporal profile of private consumption spending – have zero net present value. As a result, the analogous government spending change also has zero net present value, and so need not be financed through any change in taxes or transfers. Assumption 2' formally states the required restriction. The proof of Proposition 2 then applies without change to give

$$\hat{\mathbf{c}}_\nu \quad = \quad \underbrace{\hat{\mathbf{c}}_\nu^{PE}}_{\text{PE response}} \quad + \quad \underbrace{\hat{\mathbf{c}}_g}_{= \text{GE feedback}} \tag{C.2}$$

C.3 Exact equivalence beyond the baseline model

My proof of consumption demand equivalence applies to any model that satisfies the set of semi-structural exclusion restrictions characterized in Section 2.5. This section briefly discusses several prominent examples of such models.

Durables. I extend the household consumption-savings problem to feature durable and non-durable consumption:

$$\max_{\{c_{it}, d_{it}^h, b_{it}^h\}} \quad \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_i^t u(c_{it}, d_{it}^h, \ell_{it}) \right] \tag{C.3}$$

such that

$$c_{it} + d_{it}^h + b_{it}^h = (1 - \tau_\ell) w_t e_{it} \ell_{it} + \frac{1 + i_{t-1}^b(b_{it-1}^h)}{1 + \pi_t} b_{it-1}^h + (1 - \delta) d_{it-1}^h + \tau_{it} + d_{it} + \phi_d(d_{it-1}^h, d_{it}^h)$$

and

$$b_{it} \geq \underline{b} - (1-\theta)d_{it}$$

where $\phi_d(\bullet)$ is the durables adjustment cost function, $1-\theta$ is the share of durable holdings that can be collateralized, and – in a slight abuse of notation – I only use the superscript h to distinguish between household durables consumption d_{it}^h and dividend receipts d_{it}. Note that this specification allows for all of the bells and whistles considered in quantitative studies of durable and non-durable consumption (Barsky et al., 2007; Berger & Vavra, 2015): Households have potentially non-separable preferences over c and d^h, adjustments in durables may incur additional costs, and households can borrow against their durable goods holdings.

Crucially, I assume that the common final good y_t can be costlessly turned into the durable good d_t^h, so that the aggregate resource constraint becomes

$$y_t = \underbrace{c_t + d_t^h - (1-\delta)d_t^h}_{e_t} + i_t + g_t$$

where e_t is aggregate household expenditure. The equilibrium definition in Appendix B.1 thus generalizes straightforwardly, with aggregate household expenditure now replacing pure (non-durable) consumption expenditure. Defining a PE-GE decomposition for total household expenditure as in Definition 1, we can easily show that the demand equivalence result still applies, now for the aggregated household expenditure path $\mathbf{e}$:

Corollary C.1. *Consider the structural model of Section 2.1, extended to feature durable goods, as in Problem (C.3). Suppose that, for each one-time shock $\{\tau, g\}$, the equilibrium transition path exists and is unique. Then, under Assumptions 1 to 3, the response of consumption to a tax rebate shock τ and to a government spending shock g with $\hat{\mathbf{g}}_g = \hat{\mathbf{e}}_\tau^{PE}$ satisfy, to first order,*

$$\hat{\mathbf{e}}_\tau = \underbrace{\hat{\mathbf{e}}_\tau^{PE}}_{\text{PE response}} + \underbrace{\hat{\mathbf{e}}_g}_{= \text{ GE feedback}} \tag{C.4}$$

As argued in Beraja & Wolf (2020), consumption dynamics in models with durables generally look very different from those in models with only non-durable consumption. Corollary C.1 reveals, however, that this change in aggregate outcomes is completely orthogonal to demand equivalence.

Preferences. My baseline structural model assumes time separability in household preferences. It is, however, immediate that general forms of time non-separability are similarly nested: As long as the consumption block of the model admits aggregation to some aggregate consumption function $\mathbf{c}(\bullet)$, the equivalence proof goes through unchanged. My approximate equivalence results for the model of Justiniano et al. (2010) – with habit formation as a very simple form of non-separability – illustrate this claim.

Valued government spending. In my baseline model, households do not value government expenditure. However, it is immediate from my general characterization of consumption demand equivalence that this assumption is stronger than necessary – the key restriction is that the aggregate consumption function $\mathbf{c}(\bullet)$ does not directly depend on government consumption. Possible sufficient conditions are (i) government spending entering the per-period felicity function in an additively separable fashion,

$$\tilde{u}(c, \ell, g) = u(c, \ell) + v(g)$$

or (ii) a CES preference specification

$$u(c, \ell, g) = \frac{\left[\phi^{\rho} c^{1-\rho} + (1-\phi)^{\rho} g^{1-\rho}\right]^{\frac{1-\gamma}{1-\rho}} - 1}{1-\gamma} - \chi \frac{\ell^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}},$$

with $\rho = \gamma$.

Expectation formation. All of the models considered in this paper impose rational expectation formation for households and firms. A simple alternative is the sticky information structure in Auclert et al. (2019). It follows from Figure D.2 of that paper – the directed acylic graph representation of the model – that consumption and investment blocks under sticky information still admit aggregation to consumption and investment functions $\mathbf{c}(\bullet)$ and $\mathbf{i}(\bullet)$, respectively, and so all equivalence results apply unchanged.[44]

C.4 Range of outcomes for the "missing intercept"

Proposition 2 asserts that private and public spending shocks induce the same general equilibrium effects, but is silent on the strength of this common general equilibrium feedback.

[44]Specifically, aggregate consumption can be obtained as a function of the paths of post-tax income as well as (liquid and illiquid) returns (also see their Equation (25)), while investment can be obtained as a function of returns, wages, and aggregate output, similar to the characterization in Lemma A.1.

In this section I give two extreme examples, one with full general equilibrium crowding-out, and one with strong general equilibrium amplification.

The first example is a variant of the baseline model of Section 2.1, restricted to feature flexible prices and wages, labor-only production, and household preferences as in Greenwood et al. (1988). In this model, an income tax rebate does not move aggregate output, consumption, or labor. The argument is well-known and straightforward: Given a rebate path $\hat{\boldsymbol{\tau}}$, consider an interest rate path $\hat{\mathbf{r}}$ such that, at $(\hat{\boldsymbol{\tau}}, \hat{\mathbf{r}})$ and facing steady-state wages forever, households are willing to consume steady-state consumption $\bar{\mathbf{c}}$ forever. But then the output and labor markets clear by construction, and so we have indeed found an equilibrium. Thus, in this model, interest rate feedback fully crowds out any partial equilibrium perturbations to consumption demand.

The second example is quantitative. I consider the estimated New Keynesian business-cycle model of Justiniano et al. (2010), but now assume that preferences are as in Greenwood et al. (1988). It is immediate that this model satisfies all assumptions in Section 2.4, and so exact demand equivalence holds.

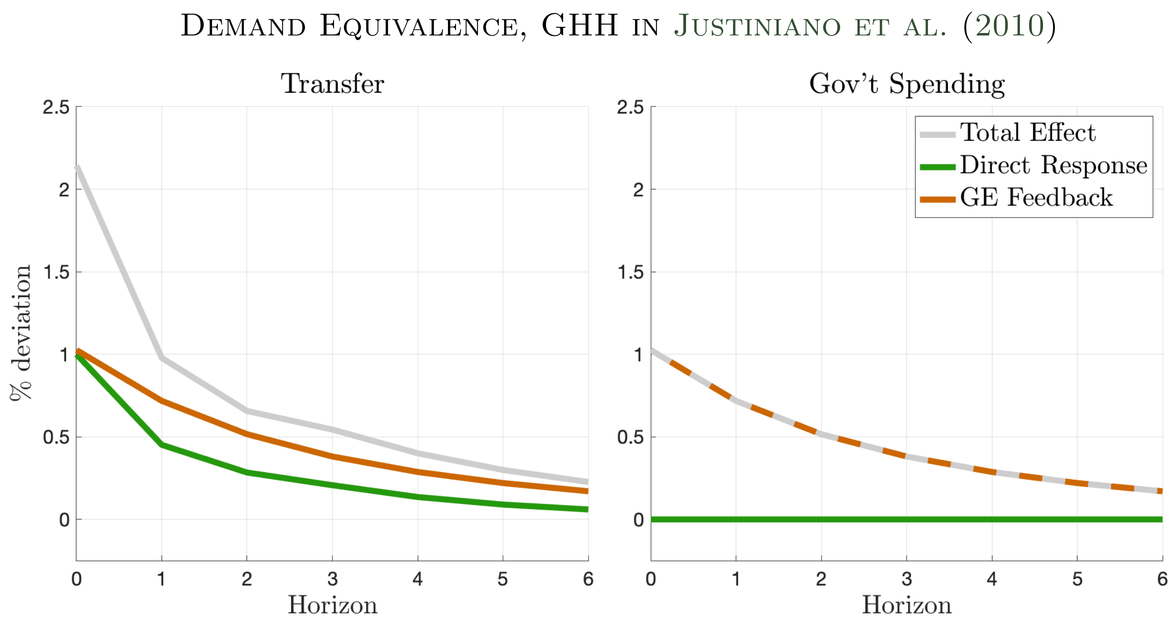

Figure C.1: Consumption impulse response decompositions after equally large, one-off impatience and government spending shocks in the model of Justiniano et al. (2010) with GHH preferences. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Given strong complementarities in consumption and labor supply, the extra production induced by the demand shock will lead to yet more consumption demand, setting in motion a strong general equilibrium feedback cycle (see Auclert & Rognlie, 2017, for an analytical characterization). Results are displayed in Figure C.1.

C.5 Correcting for wealth effects in labor supply

Instead of ignoring the labor supply error term in Proposition 2, a possible alternative is to first estimate the direct partial equilibrium labor supply response $\hat{\boldsymbol{\ell}}_d^{PE}$ from micro data and then estimate the aggregate consumption effects of an equivalent household "leisure" shock, consistent with the decomposition in Proposition 4.

Generalized methodology. Let $\hat{\mathbf{c}}_\psi$ denote the impulse response of aggregate consumption to a leisure shock – a labor wedge $\boldsymbol{\varepsilon}_\psi$ that changes desired household labor supply by $\hat{\boldsymbol{\ell}}_d^{PE}$. Using the equilibrium construction of Lemma A.1, it is straightforward to see that such a shock has no other *direct* partial equilibrium effect. It is thus immediate that, under the assumptions of Proposition 4 (i.e., without imposing Assumption 3), we have

$$\hat{\mathbf{c}}_d = \hat{\mathbf{c}}_d^{PE} + \hat{\mathbf{c}}_g + \hat{\mathbf{c}}_\psi$$

In practice, $\hat{\mathbf{c}}_\psi$ is presumably not available, since there is no good evidence (to the best of my knowledge) on the aggregate effects of pure shocks to the labor wedge. Instead, the best related evidence is on changes in labor income taxes (Mertens & Ravn, 2013). Estimates of the consumption response to labor income tax changes are likely to be informative about $\hat{\mathbf{c}}_\psi$, but have two problems. First, to translate the size of the tax change into a partial equilibrium change in labor supply $\hat{\boldsymbol{\ell}}_d^{PE}$, we need an estimate of the Frisch elasticity of labor supply. Second, the tax may generate revenue, which could be used to finance greater government spending or reduce future tax burdens.

Results. Direct micro estimates (Cesarini et al., 2017; Fagereng et al., 2018) suggest that, for every 100\$ consumption spending response to a one-off unexpected income receipt, total labor income very briefly dips by around 4\$. For the income tax rebate studied by Parker et al. (2013), partial equilibrium consumption spending increased by around 1.5%. Assuming that consumption spending roughly equals labor income, the direct labor supply response $\hat{\boldsymbol{\ell}}_d^{PE}$ thus equals around 0.06% on impact, and little thereafter.

For example, assuming a unit aggregate Frisch elasticity of labor supply, a labor supply drop of this magnitude would correspond to a transitory labor income tax increase of 0.06

percentage points. According to the point estimates of Mertens & Ravn, such a transitory tax hike in turn induces a general equilibrium drop of consumption of around 0.07%. Abstracting from the effects of future tax adjustments associated with the tax hike today, we would thus subtract around 0.07% from the benchmark estimates of the impact consumption response in Figure 4 – a hardly relevant adjustment.

C.6 Nested models for investment demand equivalence

Exact investment demand equivalence holds in the popular structural models of Khan & Thomas (2008), Khan & Thomas (2013), Winberry (2018), and Bloom et al. (2018). I verify this claim by checking that each of the assumptions necessary for the result is in fact satisfied.

First, in all of those models, capital adjustment costs are internal to the firm, so Assumption A.1 holds. Second, each model is closed with a simple representative household with linear labor disutility, so Assumptions A.2 and A.3 hold. Finally, since none of these models feature nominal rigidities, Assumption A.4 is irrelevant.[45]

[45] Well-known heterogeneous-firm models with nominal rigidities include Ottonello & Winberry (2018) and Koby & Wolf (2020). In both cases Assumption A.4 is satisfied.

D Empirical appendix

This appendix provides additional details on the empirical results needed to implement my two-step methodology. Appendix D.1 discusses estimates of the direct (partial equilibrium) consumption response to income tax rebates, Appendix D.2 does the same for investment tax credits, and Appendix D.3 offers supplemental information on my VAR-based identification of government spending shock transmission. Finally, in Appendix D.4, I briefly discuss how to account for joint estimation uncertainty in micro and macro estimators.

D.1 Direct response: micro consumption elasticities

Proposition 3 shows that, with truly exogenous cross-sectional heterogeneity in shock exposure, micro difference-in-differences regressions estimate direct partial equilibrium responses. In the empirical analysis of Johnson et al. (2006) and Parker et al. (2013), matters are slightly more subtle – all households are exposed to the shock, but exposure differs *over time* for exogenous reasons. Building on Kaplan & Violante (2014), I here discuss how to interpret their regression estimands. Parker et al. estimate a differenced version of (13):

$$\Delta c_{it} = \text{time fixed effects} + \text{controls} + \beta_0 ESP_{it} + \beta_1 ESP_{it-1} + u_{it} \tag{D.1}$$

where ESP_{it} is the dollar amount of the rebate receipt at time t. To establish that the regression estimands are interpretable as $MPC_{0,0}$ and $MPC_{1,0} - MPC_{0,0}$, respectively, consider again the structural model of Section 2.1, and suppose – roughly in line with the actual policy experiment (see Kaplan & Violante, 2014) – that a randomly selected fraction ω of households receive a lump-sum rebate at $t = 0$ ($\varepsilon_{\tau i0} = 1$), and that the remaining households receive the same rebate at $t = 1$ ($\varepsilon_{\tau i1} = 1$). The model analogue of regression (D.1) is then

$$\Delta c_{it} = \delta_{\Delta t} + \beta_0 \varepsilon_{\tau it} + \beta_1 \varepsilon_{\tau it-1} + u_{it}, \quad t = 0, 1 \tag{D.2}$$

Now suppose additionally that *receipt* of the rebate is a surprise for all households; in particular, it is a surprise at $t = 1$ for households who receive the delayed check.[46] We can then

[46] However, note that I still assume that the aggregate perfect foresight transition path is perfectly anticipated by all households; in that case, aggregate general equilibrium feedback is still differenced out. I discuss below what happens if the transition path *and* all individual rebates are anticipated by households.

follow exactly the same steps as in the proof of Proposition 3 to show that, to first order,

$$\begin{aligned}\beta_0 &= MPC_{0,0} \\ \beta_1 &= MPC_{1,0} - MPC_{0,0}\end{aligned}$$

Of course, as emphasized by Kaplan & Violante (2014), it may be dubious to assume that the delayed check was a surprise to all households. If instead the delayed check was perfectly anticipated, then the regression estimands are $\beta_0 = MPC_{0,0} - MPC_{0,1}$ and $\beta_1 = MPC_{1,0} - MPC_{1,1}$, where $MPC_{t,1} \equiv \int_0^1 \frac{\partial c_{it}}{\partial \tau_1} di$ is the response of consumption at t to a rebate received at $t = 1$, but anticipated at $t = 0$. Encouragingly, at least in my estimated HANK model, $MPC_{0,0}$ and $MPC_{1,1}$ are quite similar, and $MPC_{0,1}$ is relatively small (similar to Auclert et al. (2018)). Thus, even if the rebate was partially anticipated, the approximation underlying my estimate of the direct response in Figure 3 is likely to be accurate.

The analysis in Section 3.2 relies on the estimates of Parker et al. (2013). Since their lagged spending estimates are not significant, I base my direct spending path on the significant impact spending response in their Table 3, consistent with the headline presentation of their results in the introduction. My conclusions are, however, quite similar if the impact and delayed spending responses are evaluated at the point estimates in their Table 5.

D.2 Direct response: micro investment elasticities

Koby & Wolf (2020) generalize the static analysis of Zwick & Mahon (2017) and estimate dynamic projection regressions of the form

$$\hat{i}_{jt+h} = \alpha_j + \delta_t + \beta_{qh} \times z_{n(j),t} + u_{jt} \tag{D.3}$$

where $z_{n(j),t}$ is the size of the bonus depreciation investment stimulus for industry $n(j)$ of firm j. We estimate this regression on a quarterly Compustat sample spanning the years 1993–2017; this sample period features the two bonus depreciation episodes of 2001-2004 and 2008-2010, exactly as in Zwick & Mahon (2017). We then give sufficient conditions under which the estimands $\{\beta_{qs}\}$ are interpretable as the direct partial equilibrium response of investment to a one-time bonus depreciation stimulus.

Given the estimated partial equilibrium path $\{\hat{i}_{qt}^{PE}\}_{t=0}^{3}$, I recover the full partial equilibrium investment response by fitting a single Gaussian basis function, exactly as in Barnichon & Matthes (2018). It then remains to construct the corresponding output path $\hat{\mathbf{y}}_q^{PE}$, which requires parameter choices (α, ν, δ). To construct Figure 7 I set $\alpha = 0.2$, $\nu = 1$ and $\delta = 0.016$, in line with standard modeling practice in general and my estimated HANK model in particular.

D.3 The missing intercept: VAR estimation

My analysis of government spending shocks closely follows the important contributions of Perotti (2007) and Ramey (2011), both in terms of data and in terms of model specification. I construct the government spending forecast errors exactly as Ramey (2011). I then treat these forecast errors as a valid external instrument for structural government spending shocks; formally I assume the following:

Assumption D.1. *Suppose that an econometrician observes time series of macroeconomic aggregates y_t and professional forecast errors of real federal spending z_t, where the residualized forecast error $\tilde{z}_t \equiv z_t - \mathbb{E}\left(z_t \mid \{z_{t-\ell}, y_{t-\ell}\}_{\ell=1}^{\infty}\right)$ satisfies*

$$E(\tilde{z}_t \cdot \varepsilon_{gt}) \neq 0, \quad E(\tilde{z}_t \cdot \varepsilon_{ju}) = 0 \textit{ for all } (j, u) \neq (g, t) \tag{D.4}$$

Then, following Plagborg-Møller & Wolf (2019), I study the transmission from forecast error to macro aggregates by ordering the error first in a recursive VAR. Under Assumption D.1, this strategy will consistently estimate the desired structural impulse responses.

Proposition D.1. *(Plagborg-Møller & Wolf, 2019) Suppose that the researcher estimates a VAR in $(z_t, y_t)'$, where y_t is a vector of observed macroeconomic aggregates and z_t satisfies Assumption D.1. Let $\boldsymbol{\theta}_y$ denote the vector of impulse responses of y to the first shock in a recursive SVAR. Then the ordinary least-squares estimand of $\boldsymbol{\theta}_y$ satisfies*

$$\boldsymbol{\theta}_y \quad = \quad \textit{constant} \; \times \; \hat{\mathbf{y}}_g \tag{D.5}$$

where the constant term is a scalar*, independent of the individual response variable in y or the impulse response horizon.*

Data. My benchmark VAR consists of the log real per capita quantities of total government spending, total output (GDP), total (non-durable, durable and services) consumption, private fixed investment, total hours worked, and a measure of the *federal* average marginal tax rate (Alexander & Seater, 2009).[47]. All variables are defined and measured as in Ramey (2011). As further robustness checks, I also consider VARs with (i) total tax revenue in lieu of the marginal tax rate (following Caldara & Kamps, 2017), (ii) Greenbook defense spending forecast errors in lieu of professional forecaster errors (following Drautzburg, 2016),[48] and (iii) a log per capita measure of total federal debt.[49]

Estimation Details. I estimate all VARs in levels, with a quadratic time trend and four lags. The lag length selection is informed by standard information criteria, and is also consistent with the recommendation of Ramey (2016) in the postscript to her handbook chapter. For estimation of the model, I use a uniform-normal-inverse-Wishart distribution over the orthogonal reduced-form parameterization (Arias et al., 2018). Throughout, I display confidence bands constructed through 10,000 draws from the model's posterior.

Benchmark Results. Figure D.1 shows the impulse responses of government spending, output, consumption, investment, the marginal tax rate and total federal debt.

As in most existing structural VAR work, I construct 16th and 84th percentile confidence bands; the output and tax responses, however, remain significant at the more conventional 95 per cent level. In line with most of the previous literature I find a significant positive output response (corresponding to around a unit multiplier), and a largely flat reaction of consumption. Total debt rises immediately and significantly, suggesting that the government spending expansion is debt-financed. In fact, I also find a delayed and persistent increase in labor income taxes, as well as a similarly timed increase in total tax revenues (not shown).

Robustness. My central results – the 1-1 increase in output and the limited crowding-out of private expenditure – are robust to various changes in model specification. First, I have experimented with different sub-samples. Starting earlier (1971Q1) means that I need to link forecasts on real federal spending (available after 1981) to earlier forecasts of military

[47]The tax measure of Barro & Redlick (2011) includes state income taxes; given my focus on federal expenditure, I regard the Alexander & Seater series as more suitable for my purposes.

[48]For demand matching I need to re-scale public and private demand shocks to be in common dollar (and not percentage) terms. This is easily done using information on the GDP shares of consumption, investment, and government consumption plus investment. I take those data from FRED, and then simply compute averages for the different shares across the VAR sample period.

[49]I construct this series by scaling the total federal debt-to-output ratio (GFDEGDQ188S in the St. Louis Fed's FRED database) by my per capital real output measure. Given the short sample, the VAR with debt drops hours worked as an additional control.

Benchmark Government Spending Shock, VAR IRFs

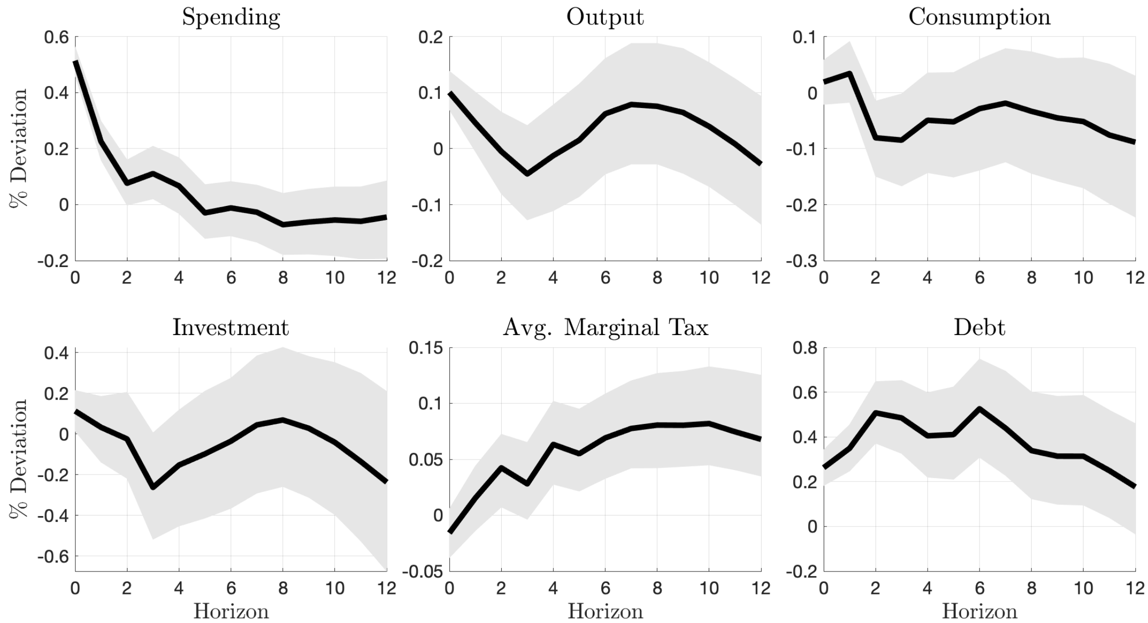

Figure D.1: Impulse responses after a one standard deviation innovation to the forecast error, quarterly frequency. The grey areas correspond to 16th and 84th percentile confidence bands, constructed using 10,000 draws from the posterior distribution of the reduced-form VAR parameters.

spending, as in Ramey (2011). Depending on the set of included controls, the undershooting of consumption and investment is, in this earlier sample, usually more pronounced (similar to Ramey, 2011). However, the undershooting then goes hand-in-hand with a similar undershooting of spending itself, invalidating the required demand matching.[50] Continuing the sample to 2016Q4 means that I need to stop controlling for taxes, as my available measures only continue until 2009. Results in this expanded sample suggest that crowding-in is slightly stronger, consistent with standard intuition on zero lower bound constraints. The results are, however, not particularly robust, similar to the findings in Ramey & Zubairy (2018) and Debortoli et al. (2019).[51] Second, replacing my benchmark measure of government spending forecast errors with Greenbook defense spending forecast errors leaves my results almost

[50]Note, however, that – unlike the impact co-movement of fiscal spending and output – the dynamic under-shooting of consumption and output is not statistically significant at the 95 per cent level. It is also somewhat dependent on the set of controls; for example, with most controls dropped, I instead find (again largely insignificant) over-shooting.

[51]I have also allowed the aggregate effects of spending shocks to be heterogeneous across expansions and recessions, through a local projection implementation identical to Ramey & Zubairy (2018). Similar to those authors, I find no evidence of such state dependence.

completely unchanged. This suggests that either (i) the benchmark VAR itself is largely picking up the response to military spending forecast errors or (ii) multipliers are invariant to the spending type (similar to Gechert (2015)). Third, changes in the number of lags or included controls do not affect the overall flavor of my results. And fourth, frequentist inference (or a flat prior) gives almost identical impulse response estimates to those displayed in Figure D.1 above.

Alternative identification. Following Blanchard & Perotti (2002), I consider a second approach to the analysis of government spending shock propagation. I estimate the same benchmark VAR as before, but now consider the dynamic propagation of an innovation to the equation for government spending g_t itself, rather than for its forecast error. This identification scheme is identical to the original approach of Blanchard & Perotti (2002), except for the fact that I now control implicitly for past government spending forecast errors.

Similar to Caldara & Kamps (2017), I find that this alternative identification scheme identifies a government spending shock with a more persistent response of government spending itself. Qualitatively, the responses of other macroeconomic aggregates – in particular output, consumption and investment – look similar to those for my benchmark identification. Importantly, because both sets of impulse responses are identified in the same reduced-form VAR, I can easily account for joint uncertainty by drawing from the posterior of that reduced-form VAR, rotating forecast residuals in line with either my benchmark or the Blanchard-Perotti identification scheme, and then finding the best fit to net demand paths following (22).

D.4 Joint Uncertainty

I throughout ignore estimation uncertainty for the first-step cross-sectional estimates. This approach is in line with standard empirical practice, which largely takes microeconomic point estimates of household MPCs and investment price elasticities at face value (e.g. Kaplan & Violante, 2014; Auclert et al., 2018; Koby & Wolf, 2020). In principle, however, it is straightforward to account for joint estimation uncertainty: Under my identifying assumptions, microeconomic and macroeconomic estimation uncertainty are independent, so sampling uncertainty for the micro and macro estimators is independent. Joint standard errors can thus be straightforwardly constructed from the individual standard errors of the micro and macro estimators. The detailed argument is available upon request.

E Approximation accuracy

In this section I gauge the robustness of my empirical methodology by applying it to artificial data generated from structural models that violate exact demand equivalence. Appendix E.1 begins with the baseline HANK model, but *randomly* draws parameters (rather than estimating them). Appendix E.2 then considers other canonical estimated business-cycle models, while Appendices E.3 to E.6 extend the baseline model to accommodate different kinds of plausible deviations from exact equivalence. Finally, in Appendix E.7, I construct my approximation under the assumption that the demand matching in (15) is not exact. Throughout I construct the population estimands of my methodology, effectively assuming that the econometrician has access to infinitely large cross-sectional and time-series samples.

E.1 Random parameter draws

The accuracy displayed in Figure 5 is not special to the particular (mode) parameterization of my estimated model, but a generic feature of standard business-cycle models with at least moderate wage and price stickiness. To illustrate this point, I proceed as follows: Rather than setting the parameter values governing dynamics as in Table B.2, I *randomly draw* their values from uninformative uniform distributions over wide supports, as displayed in Table E.1. For each parameter draw, I compute the maximal demand equivalence error relative to the true model-implied peak consumption response. This procedure is repeated for 10,000 draws from the joint uniform distributions in Table E.1.

Parameter	**Description**	**Lower Bound**	**Upper Bound**
ϕ_p	Price Calvo Parameter	0.15	0.95
ζ	Capacity Utilization	0.5	5
κ	Investment Adjustment Cost	0.5	5
ρ_m	Taylor Rule Persistence	0.15	0.95
ϕ_π	Taylor Rule Inflation	1.1	2.5
ϕ_y	Taylor Rule Output	0	0.5
ϕ_{dy}	Taylor Rule Output Growth	0	0.5

Table E.1: Supports for uniform parameter draws in the HANK model.

I find that the approximation accuracy is largely orthogonal to all parameters except for the price stickiness ϕ_p. Figure E.1 provides a graphical illustration. The grey line shows a kernel density estimate of the error distribution when all parameters except for ϕ_p are drawn randomly. It is clear that the estimated parameters have little effect on approximation accuracy – most mass of the error distribution is concentrated around the error estimate at the posterior mode. If ϕ_p is also drawn randomly, then larger errors are more likely; however, given my calibrated moderate degree of wage rigidity, shifts in household labor supply still have limited aggregate effects, and so the maximal error remains relatively small.

Demand Equivalence Error Distribution, Random Draws

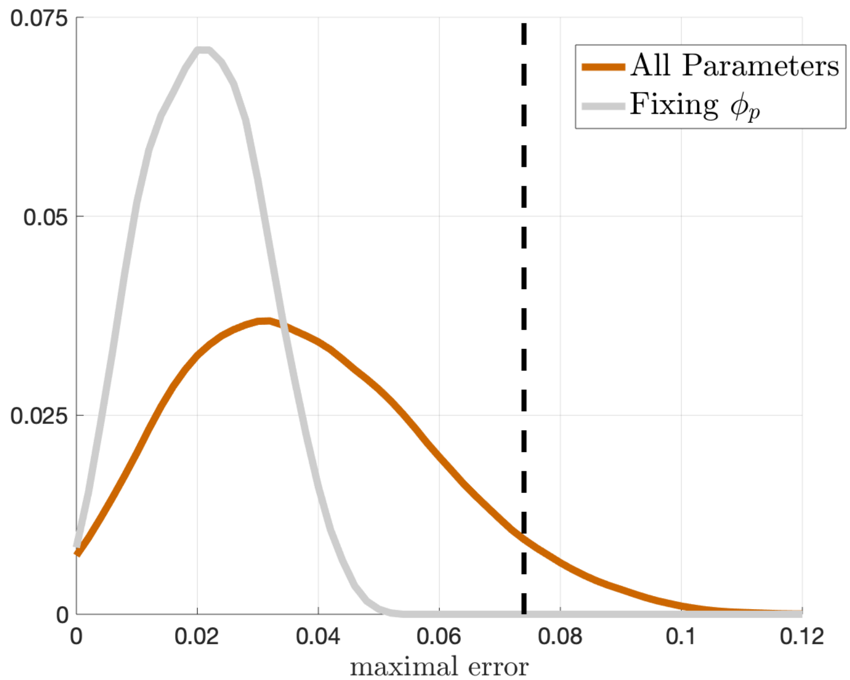

Figure E.1: Kernel estimate of maximal error distribution, with parameters drawn randomly according to Table E.1 (orange). The grey lines show the same kernel density estimate when ϕ_p is fixed at its estimated posterior mode.

E.2 Other estimated business-cycle models

Approximate consumption demand equivalence is not just a feature of my estimated HANK model, but similarly holds in many canonical models of the previous business-cycle literature. In this section I illustrate this claim with two examples: (i) Justiniano et al. (2010) as an example of an estimated New Keynesian model, and (ii) Schmitt-Grohé & Uribe (2012) as an example of an estimated neoclassical business-cycle model.

JUSTINIANO ET AL. (2010). In the estimated model of Justiniano et al. (2010), consumption demand equivalence fails only because Assumption 3 is not satisfied: wealth effects in labor supply are not zero, and hours worked are not fully demand-determined. However, prices and wages are estimated to be very sticky, and so – consistent with Christiano (2011a) – hours worked are still largely demand-determined, at least in the short run. Figure E.2 shows that this is indeed the case: For a consumption demand (impatience) shock with persistence $\rho_b = 0.1$ (and thus short-lived, similar to the tax rebate in my baseline HANK model), the error associated with the demand equivalence approximation is barely visible.

APPROXIMATE DEMAND EQUIVALENCE, JUSTINIANO ET AL. (2010)

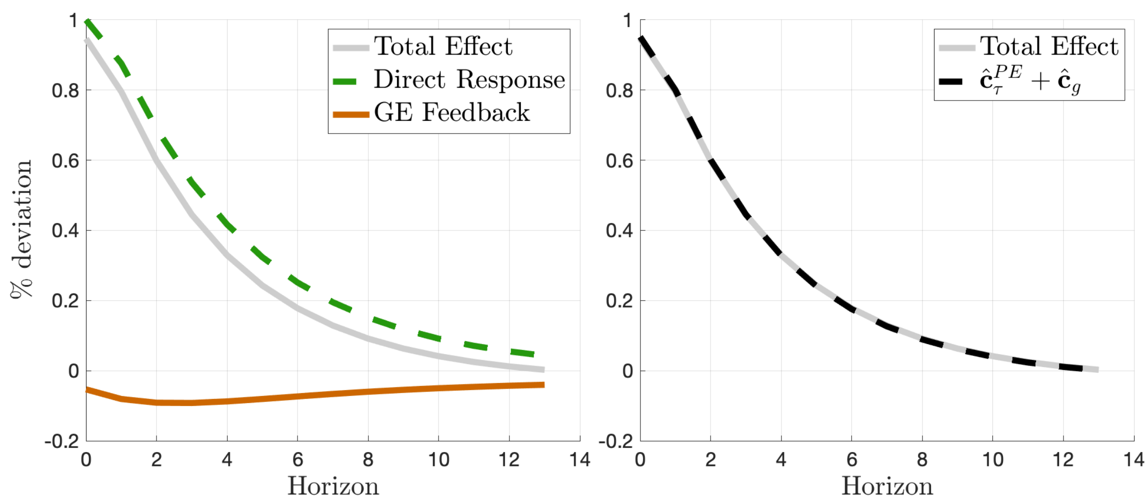

Figure E.2: Consumption impulse response decompositions and demand equivalence approximation in the model of Justiniano et al. (2010), solved at the posterior mode and for an impatience shock with persistence $\rho_b = 0.1$. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

SCHMITT-GROHÉ & URIBE (2012). The model of Schmitt-Grohé & Uribe (2012) similarly breaks consumption demand equivalence only through violation of Assumption 3. Wages and prices are now flexible, so labor is never demand-determined; however, near-exact demand equivalence still obtains because wealth effects in labor supply are essentially absent. Adapted to the notation of this paper, household preferences are given as

$$u(v) = \frac{v^{1-\sigma} - 1}{1 - \sigma}$$

where

$$v_t = c_t - bc_{t-1} - \psi \ell_t^\theta s_t$$

and

$$s_t = (c_t - bc_{t-1})^\gamma s_{t-1}^{1-\gamma}$$

As $\gamma \to 0$, there are no wealth effects in labor supply. Since both the Bayesian and frequentist estimation exercises in the paper give a very precise point estimate of $\gamma = 0$ (see Table II), I conclude that Assumption 3 holds almost exactly.

E.3 Labor supply

In the model of Appendix B.6, households are always on their labor supply curve, but preferences are non-standard, modified to feature data-consistent small short-term wealth effects in labor supply. Figure E.3 constructs my demand equivalence approximation in this model, with the direct (partial equilibrium) consumption spending response now defined as in (B.25). Estimated wealth effects are very weak, so Assumption 3 is nearly satisfied, and the approximation is again highly accurate, with a maximal error of around 4 per cent.

Approximate Demand Equivalence, Weak Wealth Effects

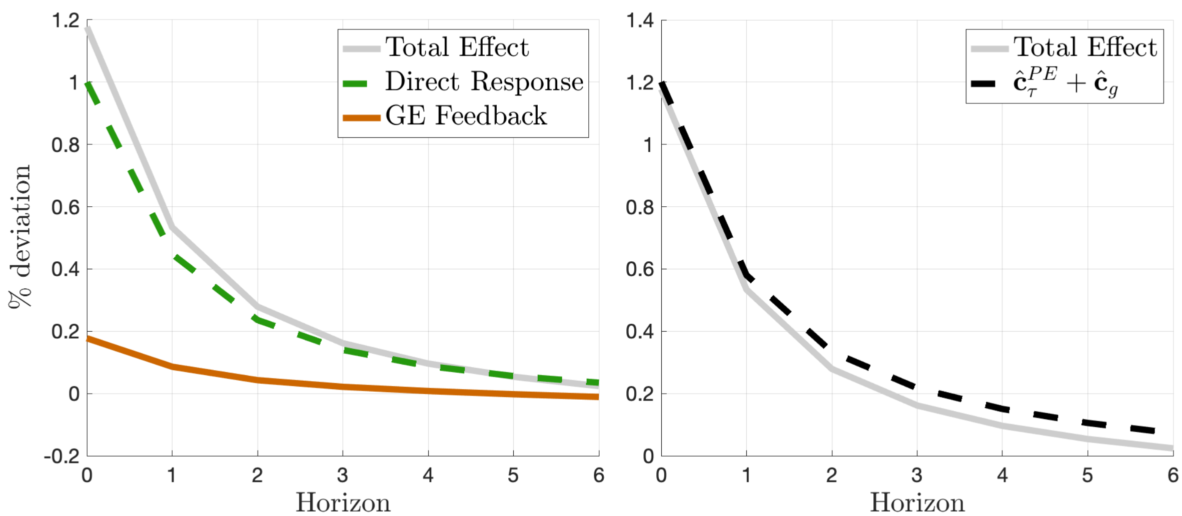

Figure E.3: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with weak wealth effects. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

E.4 Interest rates

The two-asset model sketched in Appendix B.4 is well-suited to test the role of the interest rate channel in breaking the demand equivalence approximation. With household deposit rates (significantly) below government borrowing rates, my approximation will tend to overstate the true consumption response, thus reinforcing the labor supply channel operative in the baseline model. However, as shown in Figure E.4, the approximation remains reasonable – I find a maximal error of around 7 per cent of the peak consumption response, even for a *quarterly* return gap of 1.25 per cent.[52]

Approximate Demand Equivalence, 2-Asset HANK Model

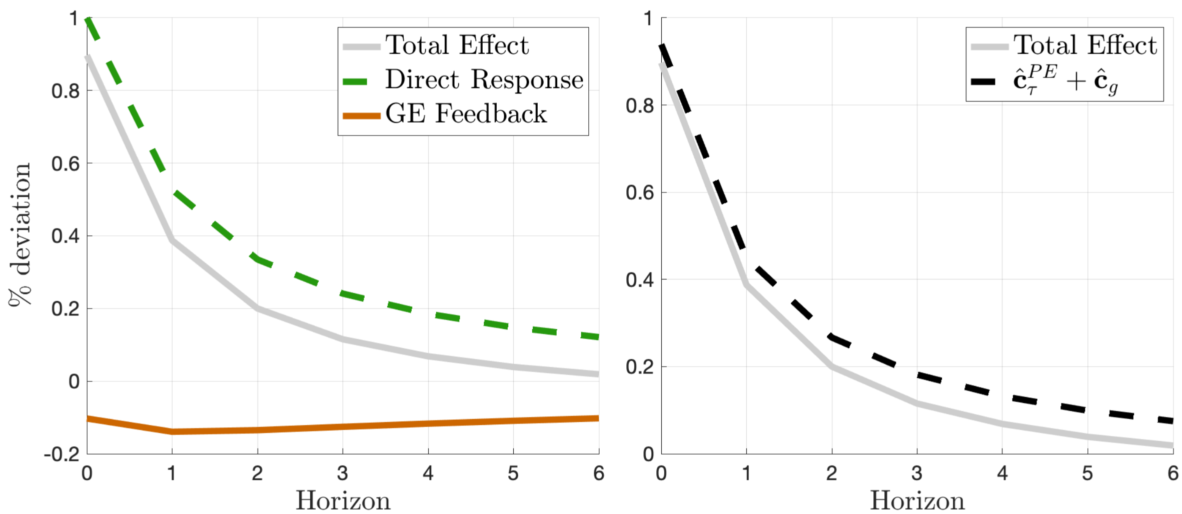

Figure E.4: Consumption impulse response decompositions and demand equivalence approximation for the two-asset HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

I have also considered model variants with large borrowing wedges (i.e., credit card debt). With indebted households facing large effective rates of return, the interest rate channel now imparts a (small) negative bias, largely offsetting the small positive bias of the labor supply channel. Results for these model variants are available upon request.

[52]The impulse responses reveal that the two-asset model features stronger general equilibrium crowding-out. This is so because of (i) smaller average MPCs and (ii) the lack of mechanical redistribution effects related to heterogeneous dividend exposure (which stimulate spending in the baseline model).

E.5 Useful government spending

My equivalence approximation generally fails if government expenditure is *useful*, in the sense that it is either (non-separably) valued by households and/or productive. In this section I argue that: (i) if the form of the non-separability is known, then a straightforward correction of my baseline methodology is available, and (ii) approximations based on productive government investment are likely to be poor, in particular at long horizons.

Valued spending. In Appendix C.3 I discuss examples of household preferences over government expenditure that are consistent with exact demand equivalence. For a generic preference specification $u(c, \ell; g)$, any change in government spending $\hat{\mathbf{g}}$ will induce a (zero net present value) consumption response path $\hat{\mathbf{c}}_g^{PE}$, defined exactly as in Definition 1. Following the same logic as in the proof of Proposition 2, it follows immediately that, if

$$\hat{\mathbf{c}}_g^{PE} \quad + \quad \hat{\mathbf{g}}_g \quad = \quad \hat{\mathbf{c}}_\tau^{PE}$$

then we have

$$\hat{\mathbf{c}}_\tau \quad = \quad \hat{\mathbf{c}}_\tau^{PE} \quad - \quad \hat{\mathbf{c}}_g^{PE} \quad + \quad \hat{\mathbf{c}}_g \tag{E.1}$$

Leeper et al. (2017) estimate that private and public consumption are (weak) complements. Given their preference specification, researchers could in principle back out $\hat{\mathbf{c}}_g^{PE}$ and thus implement the extended equivalence approximation of (E.1). However, since the estimated complementarity is small, I have – in unreported simulations – found a naive approximation (i.e., ignoring $\hat{\mathbf{c}}_g^{PE}$) to actually not do much worse. These results are available upon request.

Productive spending. If government spending has productive benefits, then the aggregate effects of private and public spending should differ. Consistent with this intuition, empirical estimates of public investment multipliers are usually larger than those of public spending (Leduc & Wilson, 2013; Gechert, 2015). These results caution against the use of public investment multipliers for the demand equivalence approximation. I have confirmed this prediction in a version of my baseline estimated HANK model in which government spending is productive: I assume that production functions take the form, with

$$y_{jt} = (k_t^g)^{\alpha_g}(u_{jt}k_{jt-1})^{\alpha}\ell_{jt}^{1-\alpha}$$

where k_t^g denotes the aggregate stock of government "capital" k_t^g. In this extended model, the equivalence approximation substantially overstates the true consumption response, in particular at long horizons.

To safeguard against this investment channel, I have re-run my empirical VAR analysis for pure military spending forecast errors (in lieu of the unconditional forecast errors considered in the main experiment). Reassuringly my empirical estimates are almost unchanged, suggesting that my analysis is not picking up the effect of public investment spending. This conclusion is consistent with results reported in Drautzburg (2016).

E.6 Multi-sector economies

Heterogeneity in consumption baskets for private and public consumption can break the demand equivalence result. Previous work has emphasized two channels. First, relative prices will move in response to spending shocks (Ramey & Shapiro, 1998). Second, if goods differ in their factor incidence (e.g., capital vs. labor income), and if factor income covaries with household characteristics (e.g., households with little non-labor income have high MPCs), then general equilibrium effects will be shock-specific (Alonso, 2017; Baqaee, 2015).

With the notable exception of productive long-lived investments, evidence on asymmetry in government spending multipliers by the type of spending is relatively scarce (Gechert, 2015; Ramey, 2016). I complement this evidence with a less direct, model-based approach: I study the accuracy of the demand equivalence approximation in the model of Appendix B.5. Consistent with the findings in Alonso (2017), I implement my demand equivalence approximation using government expenditure on the relatively more labor-intensive good. Results are displayed in Figure E.5. The approximation error is clearly visible, and goes in the expected direction: Since the MPC out of labor income is higher than that out of capital income, the approximation using the second consumption good over-states; similarly, for the first good, it under-states (not shown). However, and consistent with the conclusions in Alonso (2017) and Baqaee (2015), these incidence effects are not particularly strong. The intuition is simple: In the data, the average consumption good has a labor share of around of 0.4, while the network-adjusted labor share of government consumption is around 0.65. Assuming an (extreme) average quarterly MPC out of labor income of around 0.5, and an MPC out of any residual income of 0.05, the resulting second-round demand difference from spending on the two goods would be around 11 cents for every dollar of spending.[53]

[53]Arguably, this is an upper bound for the likely size of the effect, since heterogeneity in MPCs by skill

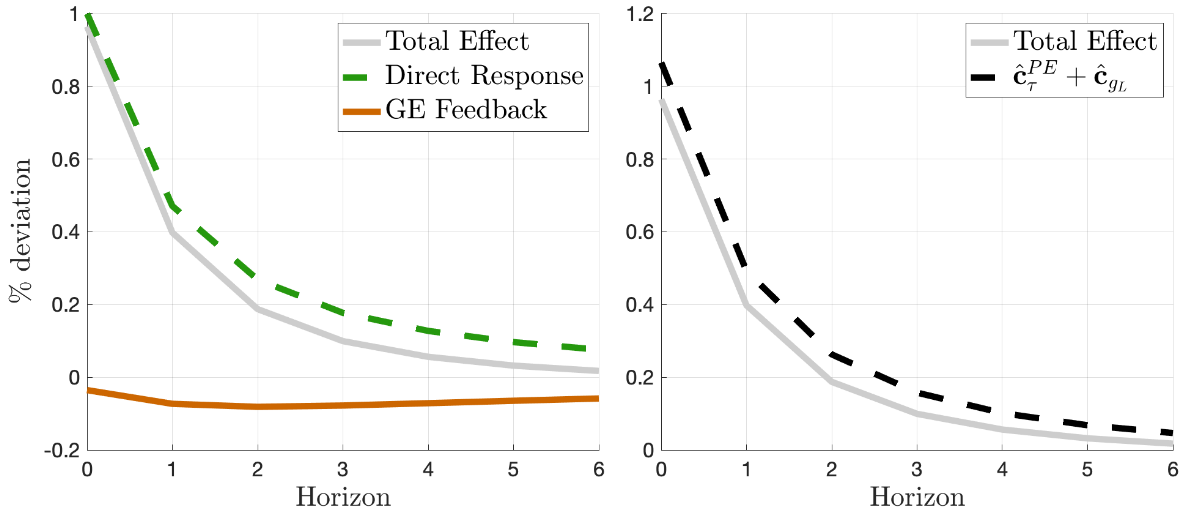

Figure E.5: Consumption impulse response decompositions and demand equivalence approximation in the three-sector HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

E.7 Imperfect demand matching

The excess demand paths in Figure 3 and Figure 7 are matched well, but of course not perfectly. To gauge the distortions associated with moderate mis-matching, I again consider the estimated HANK model of Section 4.1, but now do not assume perfectly matched excess demand paths; instead, I construct the demand equivalence approximation for an inaccurately matched government spending path $\hat{\mathbf{g}}_g$ with

$$\hat{g}_{gt} \quad = \quad (1 \, + \, \nu_t) \quad \times \quad \hat{c}_{\tau t}^{PE} \tag{E.2}$$

where $\nu_t \sim N(0, \sigma_\nu^2)$. I set σ_ν^2 to get average errors identical in size to those displayed in Figure 3; this gives $\sigma_\nu^2 = 0.123$.

I then construct the demand equivalence approximation for 10,000 draws of the error sequence $\boldsymbol{\nu}$, and for each compute the maximal prediction error relative to the peak true consumption response. I find that 95 per cent of prediction errors lie below 9.8 per cent, so the approximation remains accurate.[54] The intuition is simple: Since the model only features

implies the opposite conclusion: Government expenditure is concentrated on relatively high-skilled labor (Baqaee, 2015); if MPCs out of skilled labor are smaller, then the gap displayed in Figure E.5 shrinks.

[54] Most of the large approximation errors come from draws in which the $\boldsymbol{\nu}$'s are so far from 0 that demand

relatively moderate general equilibrium amplification, prediction errors for consumption can only be large if the demand path perturbation itself is substantial. The errors in demand matching, however, are by construction small, and thus so are the overall approximation errors. To illustrate, Figure E.6 shows the quality of the demand equivalence approximation for one particular draw of the error sequence $\boldsymbol{\nu}$.

Approximate Demand Equivalence, Imperfect Matching

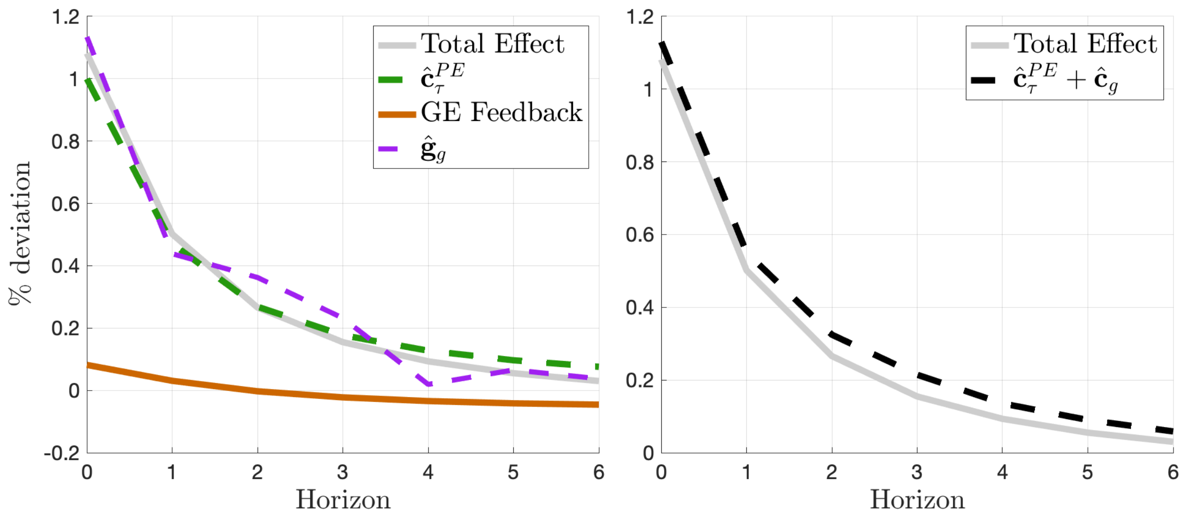

Figure E.6: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, with imperfect demand matching, following (E.2). The direct response and the indirect general equilibrium feedback are computed following Definition 1.

matching is clearly violated, so the results displayed here are actually an upper bound on likely inaccuracies.

F Application: income redistribution

The two applications in the main text are fully semi-structural: For both I recover spending impulses from micro data, and then use demand equivalence to map micro estimates into general equilibrium counterfactuals. For some interesting shocks, however, micro data are not rich enough to estimate the required direct spending responses. Appealingly, however, construction of direct spending responses only requires researchers to specify one *block* of the model economy. Given this partial model block, my methodology can again be used to provide the mapping into full general equilibrium counterfactuals, without requiring the researcher to specify all other parts of the model.

I illustrate this insight with an application to a simple redistributive, budget-neutral stimulus policy: The government imposes a lump-sum tax on the richest 10 per cent (in terms of liquid wealth holdings) of households, and uses the proceeds to finance a lump-sum rebate to the poorest 10 per cent.

Direct response. Jappelli & Pistaferri (2014) document that, because poor households on average have higher MPCs than rich households, a redistributive policy of this sort should stimulate short-term demand. However, as pointed out in Auclert & Rognlie (2018), all households spend their income at *some* point in time, so the demand stimulus today is necessarily offset by a demand contraction in the future. Since estimates of heterogeneity in *dynamic* MPCs (and so the desired spending paths) across the household wealth distribution are hard to obtain, I instead use the partial equilibrium consumption-savings problem (1) – parameterized exactly as in my estimated HANK model – to construct the partial equilibrium consumption demand path associated with the budget-neutral redistributive policy.

The solid green in the top right panel of Figure F.1 shows the estimated direct consumption response. Consistent with the empirical estimates of Jappelli & Pistaferri (2014), consumption sharply increases on impact. Since the taxed rich households behave almost exactly in line with the permanent income hypothesis, their consumption decreases slightly but persistently, so overall consumption demand decreases slightly but persistently over time.

The missing intercept. I match the implied partial equilibrium excess demand path through a combination of expansionary and contractionary government spending shocks, similar to the bonus depreciation application in Section 5.2. The top left panel shows that the partial equilibrium excess demand path is matched reasonably well, if with substantial uncertainty at higher horizons. The bottom right panel shows that taxes – which in theory

Redistribution Shock, Impulse Responses

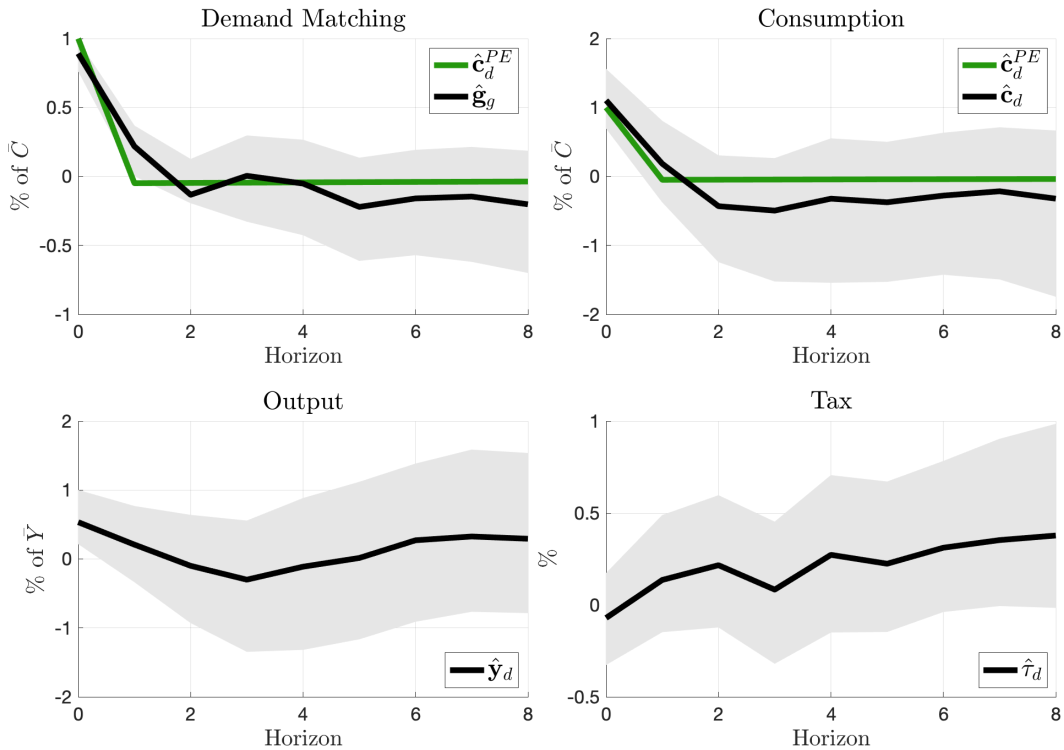

Figure F.1: Output and consumption responses to a redistribution shock, with the partial equilibrium net output response path matched to a linear combination of government spending shocks. The consumption response is computed in line with Proposition 2. The plot also shows the required demand matching as well as the implied labor tax response (cf. Assumption 2). The dashed lines again correspond to 16th and 84th percentile confidence bands.

need not respond, since the implied partial equilibrium excess demand path has zero net present value – only respond very little, so Assumption 2 is reasonable.

Macro counterfactuals. The top right panel computes the general equilibrium consumption counterfactual implied by the demand equivalence decomposition (10). Importantly, while the direct consumption response was derived from my partial equilibrium consumption-savings block, all general equilibrium feedback is estimated semi-structurally. Consistent with the results in the rest of this paper, I find limited general equilibrium feedback, so consumption rises significantly (if briefly) following the redistributive shock. The bottom left panel shows that this general equilibrium increase in consumption is accommodated through an (imprecisely estimated) increase in aggregate output.

G Further proofs and auxiliary lemmas

G.1 Auxiliary lemma for Proposition 1

Lemma G.1. *Consider a shock path $\boldsymbol{\varepsilon}$ in the spender-saver model. Sequences of real rates $\mathbf{r}$ and taxes on savers $-\boldsymbol{\tau}_e$ are part of a perfect foresight equilibrium if and only if*

$$\mathbf{c}(\mathbf{r}, \mathbf{w}(\mathbf{r}), \mathbf{d}(\mathbf{r}), \boldsymbol{\tau}_e; \boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) = \mathbf{y}(\mathbf{r}) - \mathbf{i}(\mathbf{r}) \tag{G.1}$$

$$\boldsymbol{\tau}_e = \boldsymbol{\tau}_e(\boldsymbol{\varepsilon}) \tag{G.2}$$

where $\mathbf{y}(\bullet)$ and $\mathbf{i}(\bullet)$ are firm policy functions, and optimal firm behavior implicitly pins down wages $\mathbf{w}(\bullet)$ and dividends $\mathbf{d}(\bullet)$ as functions of $\mathbf{r}$.

From the specification of the household and firm problems in Appendix B.2, it is immediate that there exist differentiable functions $\mathbf{c}(\mathbf{r}, \mathbf{w}, \mathbf{d}, \boldsymbol{\tau}_e; \boldsymbol{\varepsilon})$, $\mathbf{y}(\mathbf{r})$, $\mathbf{i}(\mathbf{r})$ and $\mathbf{d}(\mathbf{r})$ that fully characterize optimal firm and household behavior. But by (B.8) and (B.4) we can also obtain $\mathbf{w} = \mathbf{w}(\mathbf{r})$, so the expression (G.1) is well-defined.

Next, since $\mathbf{g} = \mathbf{g}(\boldsymbol{\varepsilon})$ by (B.9), we can conclude that (G.1) is necessary for any perfect foresight transition equilibrium. Since $\boldsymbol{\tau}_e = \boldsymbol{\tau}_e(\boldsymbol{\varepsilon})$ by (B.11) it is similarly immediate that (G.2) is necessary. To show sufficiency, note that (B.2), (B.3) as well as (B.4), (B.5) and (B.7) hold by optimal household and firm behavior, respectively, and that all other equations simply residually determine remaining model variables. Thus, if an interest rate path $\mathbf{r}$ and a saver transfer path $\boldsymbol{\tau}_e$ are such that (G.1) and (G.2) hold, then they are in fact part of a perfect foresight equilibrium. By existence and uniqueness of the perturbation solution (see Appendix B.2), and by equivalence of perfect foresight transition paths and perturbation solutions (Fernández-Villaverde et al., 2016; Boppart et al., 2018), we know that this transition path exists and is unique. □

G.2 Proof of Proposition 1

By differentiability of the consumption, investment and output supply functions, a perfect foresight equilibrium is, to first order, a solution to the *linear* system of equations

$$\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{c}}{\partial \mathbf{r}} \times \hat{\mathbf{r}} + \frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}_e} \times \hat{\boldsymbol{\tau}}_e + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} = \left(\frac{\partial \mathbf{y}}{\partial \mathbf{r}} - \frac{\partial \mathbf{i}}{\partial \mathbf{r}}\right) \times \hat{\mathbf{r}} \tag{G.3}$$

$$\hat{\boldsymbol{\tau}}_e = \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \tag{G.4}$$

The existence of a unique perturbation solution (see the discussion in Appendix B.2) in conjunction with Lemma G.1 implies that this equation also has a unique bounded solution for $(\hat{\mathbf{r}}, \hat{\boldsymbol{\tau}}_r)$. Thus there exists a unique linear map $\mathcal{H}$ such that

$$\begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\tau}}_e \end{pmatrix} = \mathcal{H} \times \begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \end{pmatrix} \tag{G.5}$$

where $\mathcal{H}$ is the left inverse of

$$\begin{pmatrix} \frac{\partial \mathbf{y}}{\partial \mathbf{r}} - \frac{\partial \mathbf{i}}{\partial \mathbf{r}} - \frac{\partial \mathbf{c}}{\partial \mathbf{r}} & -\frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}_r} \\ \mathbf{0} & \frac{\partial \boldsymbol{\tau}_r}{\partial \boldsymbol{\varepsilon}} \end{pmatrix}$$

Since there exists a unique bounded solution, this left inverse is unique. Thus, in response to a generic shock $\boldsymbol{\varepsilon}_s$, the response path of consumption satisfies

$$\hat{\mathbf{c}}_\varepsilon = \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}}_{\hat{\mathbf{c}}_\varepsilon^{PE}} + \underbrace{\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \mathbf{r}} & \frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}_e} \end{pmatrix} \times \mathcal{H} \times \begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \end{pmatrix}}_{\equiv \mathcal{D} \times \left(\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \right)} \tag{G.6}$$

The definition of the "demand multiplier" map $\mathcal{D}$ uses my assumptions on the government financing rule – both policy experiments can be (and in fact are) financed using identical paths of lump-sum saver taxes.[55] General equilibrium feedback is thus identical, giving (5), and (6) follows. □

G.3 Proof of Lemma A.1

To prove Lemma A.1 I proceed in two steps. First, I show that all relevant inputs to the household and firm problems can be obtained as functions only of $\mathbf{x}$ and $\boldsymbol{\varepsilon}$. Second, I show sufficiency of the four equations in the statement of the result.

1. Given $(\mathbf{i}^b, \mathbf{y})$, the Taylor rule of the monetary authority allows us to back out the path of inflation $\boldsymbol{\pi}$. Thus all inputs to the firm problem are known,[56] so indeed $\mathbf{s}^f = \mathbf{s}^f(\mathbf{x})$. We thus obtain $\mathbf{y}$, $\mathbf{i}$ and $\boldsymbol{\ell}^f$. Setting $\boldsymbol{\ell} = \boldsymbol{\ell}^f$ and since $\boldsymbol{\tau}_e \in \mathbf{x}$, all inputs to the household problem are known, so indeed $\mathbf{s}^h = \mathbf{s}^h(\mathbf{x})$. We can thus also solve for the path of consumption, so that indeed $\mathbf{s}^u = \mathbf{s}^u(\mathbf{x}; \boldsymbol{\varepsilon})$, and we finally recover union labor supply.

[55] Of course, given Ricardian equivalence, this does not really matter. It only matters that the present value of the implied tax burdens on savers is the same, which is ensured by the intertemporal government budget constraint (and since government and savers borrow and save at a common rate).

[56] Note that the path of the intermediate goods price p^I is obtained from the problem of retailers.

2. Optimal household, firm and government behavior is assured by assumption. It thus remains to check that (i) all markets clear (ii) that the input path of output is consistent with firm production, and (iii) that the rebate path is consistent with the government budget constraint. Output and labor market-clearing are ensured by the first two equations in the statement of the lemma, and asset market-clearing then follows from Walras' law. The third set of equations in the lemma statement then ensures consistency in aggregate production, while the fourth set – which uses that the only relevant quantities for the government budget constraint are $(\mathbf{r}, \mathbf{w}, \boldsymbol{\ell})$ – ensures that the government budget constraint holds period-by-period.

Together, 1. - 2. establish sufficiency of the conditions in the statement of Lemma A.1. Necessity is immediate, completing the argument. □

G.4 Proof of Proposition 3

The proof proceeds in three steps. First, I show that *aggregate* impulse responses to the heterogeneous shocks $\{\varepsilon_{\tau i0}\}$ are identical to impulse responses to the common aggregate shock $\varepsilon_{\tau 0} \equiv \int_0^1 \varepsilon_{\tau i0}$. Second, I prove that $\hat{\mathbf{c}}_{\tau i} - \hat{\mathbf{c}}_{\tau} = (\xi_{\tau i0} - 1) \times \hat{\mathbf{c}}_{\tau}^{PE} + \boldsymbol{\zeta}_i$, where $\int_0^1 (\xi_{\tau i0} - 1)\boldsymbol{\zeta}_i di = \mathbf{0}$. And third, I exploit standard properties of fixed-effects regressions to complete the argument. As in the proof of Proposition 2, I use the notation $\frac{\partial}{\partial \boldsymbol{\varepsilon}_s}$ to denote derivatives for a shock path where only entries of shock s are non-zero.

1. We study impulse responses to the shock path $\boldsymbol{\varepsilon}_{\tau} \equiv e_1$, where $e_1 = (1, 0, 0, \ldots)'$. The direct partial equilibrium response of consumption to the shock is

$$\hat{\mathbf{c}}_{\tau}^{PE} \equiv \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_{\tau}} \times \xi_{\tau i0} \times \boldsymbol{\varepsilon}_{\tau} di$$

where $\mathbf{c}_i(\bullet)$ is the consumption function of individual i, defined analogously to the aggregate consumption function $\mathbf{c}(\bullet)$. Since $\int_0^1 \xi_{\tau i0} di = 1$ and since $\xi_{\tau i0}$ is assigned randomly across households (and so does not correlate with $\frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_{\tau}} \times \boldsymbol{\varepsilon}_{\tau}$ at any t), we have that

$$\hat{\mathbf{c}}_{\tau}^{PE} = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_{\tau_0}} \times \boldsymbol{\varepsilon}_{\tau} di \times \left[1 + \int_0^1 (\xi_{\tau i0} - 1) di\right] = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_{\tau}} \times \boldsymbol{\varepsilon}_{\tau} di$$

The direct partial equilibrium response of aggregate consumption is thus identical to the response in an economy where all individuals i face the common shock $\boldsymbol{\varepsilon}_{\tau}$. The same argument applies to the desired partial equilibrium contraction in labor supply, $\hat{\boldsymbol{\ell}}_{\tau}^{PE}$. But if direct partial equilibrium responses are the same, then general equilibrium adjustment is the same, and so all aggregates are the same.

2. Consumption of household i along the transition path satisfies

$$\hat{\mathbf{c}}_{i\tau} = \frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} \times \xi_{\tau i0} \times \boldsymbol{\varepsilon}_\tau$$

where $\mathbf{x}$ was defined in Lemma A.1. We thus get

$$\hat{\mathbf{c}}_{\tau i} - \hat{\mathbf{c}}_\tau = (\xi_{\tau i0} - 1) \times \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau} \times \boldsymbol{\varepsilon}_\tau}_{\hat{\mathbf{c}}_\tau^{PE}} + \underbrace{\left(\frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} - \frac{\partial \mathbf{c}}{\partial \mathbf{x}}\right) \times \hat{\mathbf{x}} + \xi_{\tau i0}\left(\frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} - \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau}\right) \times \boldsymbol{\varepsilon}_\tau}_{\equiv \boldsymbol{\zeta}_i}$$

Note that, since by definition we have $\int_0^1 \frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} di = \frac{\partial \mathbf{c}}{\partial \mathbf{x}}$ and $\int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} di = \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau}$, the residual term $\boldsymbol{\zeta}_i$ must satisfy $\int_0^1 (\xi_{\tau i0} - 1)\boldsymbol{\zeta}_i di = 0$.

3. By the standard properties of fixed-effects regression, we can re-write regression (13) as

$$\hat{c}_{it+h} - \hat{c}_{t+h} = \beta_{\tau h} \times (\xi_{it} - 1)\varepsilon_{dt} + u_{it+h} - u_{t+h} \tag{G.7}$$

By standard projection results, the estimand $\boldsymbol{\beta}_\tau$ satisfies

$$\begin{aligned} \boldsymbol{\beta}_\tau &= \frac{\int_0^1 \left[(\xi_{\tau i0} - 1)\hat{\mathbf{c}}_\tau^{PE} + \boldsymbol{\zeta}_i\right] (\xi_{\tau i0} - 1) di}{\int_0^1 (\xi_{\tau i0} - 1)^2 di} \\ &= \hat{\mathbf{c}}_\tau^{PE} \end{aligned}$$

where I have used the fact that $\text{Var}(\xi_{\tau it}) > 0$.

□

G.5 Proof of Proposition 4

Following the same steps as in the proof of Proposition 2, but without imposing Assumption 3, we get the two direct shock responses as

$$\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau} \\ \frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}_\tau} \\ \mathbf{0} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}_\tau} \end{pmatrix} \times \boldsymbol{\varepsilon}_\tau = \begin{pmatrix} \hat{\mathbf{c}}_\tau^{PE} \\ \hat{\boldsymbol{\ell}}_\tau^{PE} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{e\tau}^{PE} \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}_g} \\ \mathbf{0} \\ \mathbf{0} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}_g} \end{pmatrix} \times \boldsymbol{\varepsilon}_g = \begin{pmatrix} \hat{\mathbf{g}}_g \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{eg}^{PE} \end{pmatrix}$$

The response paths of consumption now satisfy

$$\hat{\mathbf{c}}_\tau = \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau} \times \boldsymbol{\varepsilon}_\tau}_{\hat{\mathbf{c}}_\tau^{PE}} + \frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{c}}_\tau^{PE} \\ \hat{\boldsymbol{\ell}}_\tau^{PE} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{e\tau}^{PE} \end{pmatrix}, \quad \text{and} \quad \hat{\mathbf{c}}_g = \mathbf{0} + \frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{g}}_g \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{eg}^{PE} \end{pmatrix}$$

Combining the two:

$$\hat{\mathbf{c}}_\tau \quad = \quad \hat{\mathbf{c}}_\tau^{PE} \quad + \quad \hat{\mathbf{c}}_g \quad + \quad \underbrace{\frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \mathbf{0} \\ \hat{\boldsymbol{\ell}}_\tau^{PE} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}}_{\text{error}\left(\hat{\boldsymbol{\ell}}_\tau^{PE}\right)}$$

In particular, the third term is immediately seen to be the general equilibrium response of consumption to a leisure shock leading to a desired union labor supply adjustment of $\hat{\boldsymbol{\ell}}_\tau^{PE}$, as claimed. $\square$

G.6 Auxiliary lemma for Proposition A.1

Lemma G.2. *Consider the structural model of Section 2.1. Under Assumptions A.1 to A.4, all firm sector price inputs $\mathbf{s}^f$ can be derived as functions only of the path of aggregate consumption $\mathbf{c}$. Sequences of consumption $\mathbf{c}$ and shocks $\boldsymbol{\varepsilon}$ are part of a perfect foresight equilibrium if and only if*

$$\mathbf{c} + \mathbf{i}(\mathbf{s}^f(\mathbf{c}); \boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) = \mathbf{y}(\mathbf{s}^f(\mathbf{c}); \boldsymbol{\varepsilon}) \tag{G.8}$$

where the production and investment functions $\mathbf{y}(\bullet)$, $\mathbf{i}(\bullet)$ are derived from optimal firm behavior.

To prove Lemma G.2 I as before proceed in two steps. First, I show that all relevant inputs to the firm problem can be obtained as functions only of $\mathbf{c}$ and $\boldsymbol{\varepsilon}$. Second, I show sufficiency of the aggregate market-clearing equation.

1. By Assumptions A.2 and A.3, the household block admits aggregation to a single representative household with period felicity function $u(c) - v(\ell)$. Given $\mathbf{c}$, the Euler equation

of the representative household allows us to back out the path of real interest rates $\mathbf{r}$. Given $\mathbf{r}$, the Fisher equation and the Taylor rule of the monetary authority (by Assumption A.4) allow us to recover the path of aggregate inflation $\boldsymbol{\pi}$, and so by the NKPC of retailers we recover $\mathbf{p}^I$. Next, given Assumption A.3, the wage-NKPC allows us to recover the path of real wages $\mathbf{w}$. Together with $\boldsymbol{\varepsilon}$ we thus have all inputs to the firm problem, and in particular indeed $\mathbf{s}^f = \mathbf{s}^f(\mathbf{c})$, as claimed.

2. Optimal firm and government behavior is assured by construction. Next, since the Euler equation and wage-NKPC hold, the only missing condition for household optimality is the lifetime budget constraint. But by assumption the aggregate market-clearing condition (G.8) holds at all times, so the household lifetime budget constraint must hold. Finally, the labor market clears by Assumption A.3.

Together, 1. - 2. establish sufficiency of the conditions in the statement of Lemma G.2. Necessity is immediate, completing the argument. □

G.7 Proof of Proposition A.1

By Lemma G.2, a perfect foresight equilibrium is, to first order, a solution to the system of linear equations

$$\hat{\mathbf{c}} + \frac{\partial \mathbf{i}}{\partial \hat{\mathbf{c}}} \times \hat{\mathbf{c}} + \frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} = \frac{\partial \mathbf{y}}{\partial \hat{\mathbf{c}}} \times \hat{\mathbf{c}} + \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}$$

As before, we thus in general have

$$\hat{\mathbf{c}} = \mathcal{H} \times \left(\frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} - \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \right)$$

for a unique linear map $\mathcal{H}$. Now again use the notation $\frac{\partial}{\partial \boldsymbol{\varepsilon}_s}$ to denote derivatives for a shock path where only entries of shock s are non-zero. In response to investment tax and government spending shocks, the response path of investment satisfies

$$\hat{\mathbf{i}}_q = \underbrace{\frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}_q} \times \boldsymbol{\varepsilon}_q}_{\hat{\mathbf{i}}_q^{PE}} + \underbrace{\frac{\partial \mathbf{i}}{\partial \mathbf{c}} \times \mathcal{H} \times \left(\frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}_q} \times \boldsymbol{\varepsilon}_q - \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}_q} \times \boldsymbol{\varepsilon}_q \right)}_{\mathcal{D}_i \times \left(\hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE} \right)}$$

and

$$\hat{\mathbf{i}}_g \quad = \quad \mathbf{0} \quad + \quad \underbrace{\frac{\partial \mathbf{i}}{\partial \mathbf{c}} \times \mathcal{H} \times \left(\frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}_g \right)}_{\mathcal{D}_i \times \hat{\mathbf{g}}_g}$$

respectively. This establishes (20). The equations for output are exactly analogous. □

G.8 Proof of Corollary C.1

It is straightforward to show that a generalization of Lemma A.1 now holds for the system

$$\begin{aligned} \mathbf{e}(\mathbf{s}^h(\mathbf{x});\boldsymbol{\varepsilon}) + \mathbf{i}(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) &= \mathbf{y}(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) \\ \boldsymbol{\ell}^h(\mathbf{s}^u(\mathbf{x};\boldsymbol{\varepsilon})) &= \boldsymbol{\ell}^f(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) \\ \mathbf{y}(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) &= \mathbf{y} \\ \boldsymbol{\tau}_e(\mathbf{s}^f(\mathbf{x});\boldsymbol{\varepsilon}) &= \boldsymbol{\tau}_e \end{aligned}$$

where $\mathbf{e}$ is now the aggregated optimal household expenditure function for durable and non-durable consumption. Applying the same steps as in the proof of Proposition 2 to this new system, the result follows. □

G.9 Proof of Proposition D.1

By definition of $\hat{\mathbf{y}}_g$, we know that

$$\hat{y}_{gh} = \text{Cov}(y_{t+h}, \varepsilon_{gt})$$

$\hat{y}_{gh}$ is thus the estimand of a local projection on ε_{gt}. (D.5) then follows immediately by Corollary 1 in Plagborg-Møller & Wolf (2019).[57] □

[57] Strictly speaking, it remains to verify their Assumption 1, ensuring that the process $(z_t, y_t')'$ is not stochastically singular. It is straightforward to augment the model of Section 2.1 with more structural shocks or measurement errors to ensure that this is the case for any vector of macroeconomic observables y_t.

UniCredit Foundation

Piazza Gae Aulenti, 3

UniCredit Tower A

20154 Milan

Italy

Giannantonio De Roni – *Secretary General*

e-mail: giannantonio.deroni@unicredit.eu

Annalisa Aleati - *Scientific Director*

e-mail: annalisa.aleati@unicredit.eu

Info at:

www.unicreditfoundation.org

UniCredit Foundation