

The Missing Intercept: A Demand Equivalence Approach*

Christian K. Wolf

November 18, 2019

JOB MARKET PAPER

[latest version]

Abstract: I prove that, in a broad class of structural macro models, shocks to private consumption demand elicit the same general equilibrium price responses as changes in aggregate public spending. This demand equivalence result implies that the aggregate effects of a large family of consumption demand shocks can be estimated in two steps. First, a researcher uses cross-sectional heterogeneity in shock exposure to recover the partial equilibrium effect on consumption demand. Second, indirect general equilibrium effects – the “missing intercept” of the micro regression – are, by demand equivalence, equal to the response of consumption to changes in aggregate public spending. I apply this method to deficit-financed tax rebates, and find (i) a large partial equilibrium response, but (ii) a fiscal multiplier of one and so little further general equilibrium feedback to private spending. Any calibrated structural model that instead implies a non-zero general equilibrium intercept either breaks demand equivalence or features fiscal multipliers far from one. Finally, I show that equivalence extends to generic investment demand shocks, and use it to estimate the effects of bonus depreciation stimulus.

*Email: ckwolf@princeton.edu. I am very grateful to my advisors Ben Moll, Mikkel Plagborg-Møller, Gianluca Violante and Mark Watson for their guidance and support. I also received helpful comments from Mark Aguiar, Adrien Auclert, Markus Brunnermeier, Thorsten Drautzburg, Oleg Itskhoki, Gregor Jarosch, Greg Kaplan, Nobuhiro Kiyotaki, Yann Koby, Moritz Lenel, Alisdair McKay, Emi Nakamura, Ezra Oberfield, Jonathan Payne, Diego Perez, Richard Rogerson, Matt Rognlie, Chris Sims, Jón Steinsson, Ludwig Straub, Chris Tonetti, Tom Winberry, and seminar participants at several venues. I acknowledge 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 is structural in that it builds on a property of models: I prove that, in a broad class of business-cycle models, identical changes in partial equilibrium private and public net excess demand also elicit identical general equilibrium price responses – intuitively, identical pressure on the economy’s resource constraint is accommodated in the same way. This “demand equivalence” justifies a two-step empirical strategy: First, a researcher leverages cross-sectional heterogeneity in shock exposure to recover the partial equilibrium response of private spending demand to the shifter. Second, she estimates the aggregate effects of an equally large shock to public spending using the existing toolkit for fiscal shocks (e.g. Ramey, 2018). By demand equivalence, summing (i) her micro estimates and (ii) the aggregate response of private spending to the fiscal shock will recover the *full* effect of the demand shifter on private spending.

I apply my method to study tax rebates. Micro data suggest a large but short-lived direct consumption response, while macro experiments for a similarly short-lived, deficit-financed change in public spending imply only little crowding-out of private consumption. By demand equivalence, it follows that full macro counterfactuals are close to the direct micro estimates. In order for a structural model to be inconsistent with such a near-zero “missing intercept,” it must either break demand equivalence or feature fiscal multipliers far from one. I find a similarly small missing intercept in a second application to investment tax stimulus.

¹Examples 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).

I first discuss the restrictions on household behavior and fiscal policy required for identical general equilibrium propagation of shocks to public spending and to private consumption demand. For a general class of quantitative 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. If 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. This invariance result, together with three further assumptions, allows me to prove the (first-order) equivalence of shocks to private consumption and to public spending. First, households and government need to consume the same final good. If so, identical changes in private or public spending will lead to identical partial equilibrium 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. The first two assumptions required for my exact equivalence result are satisfied in many popular structural general equilibrium models, including standard medium-scale New Keynesian models (Smets & Wouters, 2007), but also models with rich micro household and firm heterogeneity (e.g. Guerrieri & Lorenzoni, 2017; Khan & Thomas, 2013). While the third assumption is sometimes violated, I argue both empirically and theoretically that the associated error is robustly small.

I leverage the consumption demand equivalence result to formally justify my two-step empirical procedure. First I show that, if cross-sectional heterogeneity in exposure to a consumption demand shifter is independent of household characteristics, then the econometric estimands of micro difference-in-differences regressions are interpretable as the direct (partial equilibrium) response of consumption demand to the shifter. Second, by demand equivalence, the “missing intercept” can be recovered as the response of consumption to a public spending shock – or a combination of such shocks – that induces the same path of (partial equilibrium) net excess demand. For measurement, I link this response to the econometric estimands of the popular Structural Vector Autoregressive (VAR) or Local Projection (LP) approaches to fiscal shock transmission (Hall, 2009; Ramey, 2018). Under my identifying assumptions, the sum of the micro and macro estimates is then indeed interpretable as the desired semi-structural aggregate consumption counterfactual for the private demand shifter. Equivalently, the researcher could have written down *any* particular parametric model in

my 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 exact same macro counterfactual.

I demonstrate the feasibility and applicability of my methodology through the study of a popular consumption stimulus policy: a lump-sum, one-off income tax rebate. I first review previous empirical work (Parker et al., 2013; Jappelli & Pistaferri, 2014) and show that the direct partial equilibrium response of consumption to the stimulus is indeed either equal or at least tightly linked to the econometric estimands of those studies. Their different quasi-experiments consistently paint the picture of a large but short-lived expansion in consumption demand. Next, I construct a government spending news variable based on professional forecast errors, and treat this forecast error as a macro instrumental variable for government spending shocks (Stock & Watson, 2018). Following Plagborg-Møller & Wolf (2019), I project on this news variable using a recursive VAR. I find that the forecast error impulse leads to an uptick in government spending as short-lived as the private consumption demand increase, a persistent rise in government debt, and a fiscal multiplier of around one – output rises briefly, and consumption is 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 under my three key assumptions, I show that approximate demand equivalence is supported in the data and obtains in several quantitative model extensions. As discussed above, equivalence fails in an important benchmark class of models *only* because of short-term wealth effects in labor supply. Micro evidence is inconsistent with this labor supply channel (Cesarini et al., 2017), and general equilibrium model closures with even moderate degrees of nominal rigidity are well-known to largely neuter the aggregate effects of transitory labor supply shifts (Christiano, 2011a). I use 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, to illustrate the quantitative irrelevance of the labor supply channel. In the same model, saving and borrowing rates for households and the government are – in line with the evidence in Fagereng et al. (2018) – sufficiently similar to ensure that my second assumption is also nearly satisfied. Finally, to gauge the importance of the single common good assumption, I extend the model to a multi-sector economy, allowing for (i) a durable consumption good, (ii) productive benefits and consumption complementarities for public spending, and

(iii) imperfect factor mobility across sectors with heterogeneous production functions. In empirically disciplined variants of these models, the approximation remains accurate.

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), followed by a contraction in the future (excess supply). Importantly, these sufficient conditions impose no material restrictions on the production block of the economy; in particular, exact investment demand equivalence holds in most recent quantitative studies on 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 a large partial equilibrium increase in investment demand (Zwick & Mahon, 2017; Koby & Wolf, 2019), accommodated through a sharp rise in output, with little investment crowding-out and consumption relatively flat.

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 direct partial equilibrium effects. 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). I generalize my results to such cross-regional regressions in the companion note Wolf (2019). Second, the general principle underlying my approach – to leverage macro evidence on the general equilibrium propagation of plausibly equivalent shocks – is applicable to a rich family of consumption and investment shifters. It does not, however, solve the missing intercept problem for all possible shocks and policies. Third, while my demand 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 partial equilibrium impacts and general equilibrium accommodation that lies at the heart of my approach.

²For example, exact equivalence applies in the popular structural models of Khan & Thomas (2008), Khan & Thomas (2013), or Winberry (2018).

LITERATURE. This paper contributes to several strands of the 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. For example, Johnson et al. (2006), Agarwal et al. (2007), and Parker et al. (2013) leverage cross-sectional heterogeneity in policy exposure to study the response of household consumption expenditure to lump-sum payments. Estimates of household-level marginal propensities to consume have also been used to gauge the likely effects of income redistribution, either outright through policy changes or through gradual increases in inequality (Jappelli & Pistaferri, 2014). In Mian & Sufi (2009) and Mian & Sufi (2014), cross-*regional* heterogeneity in shock exposure is used to recover the direct consumption effects of changes in household balance sheets. Analogous micro causal effects can also be estimated for investment (e.g. Zwick & Mahon, 2017). As all of these studies control for macro fluctuations through time fixed effects, they are silent on any possible general equilibrium feedback. My key insight is that a second empirical literature – that on the aggregate effects of variations in government spending – can be informative about this “missing intercept.” Comprehensive literature summaries are Hall (2009) and Ramey (2018); overall, earlier empirical work quite consistently estimates output multipliers around 1, and zero (or slightly negative) responses of private spending. In connecting these two literatures, my two-step procedure is semi-structural in exactly the same way as conventional Structural Vector Autoregressive (SVAR) analysis (Sims, 1980): It relies on general identifying restrictions, rather than being tied to any particular parametric model of the macro-economy.

Second, the theoretical demand equivalence result itself builds on the burgeoning sufficient statistics literature in macroeconomics. Earlier contributions show that the estimands of micro difference-in-differences studies are linked to partial equilibrium spending elasticities (e.g. Kaplan & Violante, 2014; Berger et al., 2017). Very recently, several studies have tried to clarify the relationship between these partial equilibrium elasticities and the associated aggregate counterfactuals. Among those, my analysis relates most closely to Auclert & Rognlie (2018), Auclert et al. (2018), and Guren et al. (2019). Auclert et al. 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 – a special case of what I call “demand equivalence.” I extend the equivalence result to a larger family of models (and to investment demand), find support for approximate equivalence in micro and macro data, and measure the common general equilibrium effects through macro quasi-experiments. In ongoing related work, Guren et al. (2019) use estimates

of *local* government spending multipliers to cleanse regional consumption responses to house price changes from regional spending multipliers. Relative to their analysis, I use the reverse logic, aggregating partial equilibrium consumption effects *up* to macro outcomes, show that the equivalence logic applies to generic consumption and investment demand shifters, and study its accuracy in a larger model space, and with an emphasis on full impulse response dynamics.³ Overall, my results have a “sufficient statistics” interpretation similar to Chetty (2009) or Arkolakis et al. (2012): To deviate much from the conclusion of a near-0 missing intercept for private demand shifters, structural modelers must either leave the demand equivalence class or implicitly impose that fiscal multipliers are far from 1.

Third, my results connect to the large literature on estimation of quantitative business-cycle models. Dominant approaches are limited-information moment-matching, notably of impulse response functions, as well as full likelihood-based estimation (Christiano et al., 2005; Smets & Wouters, 2007; Nakamura & Steinsson, 2018). The equivalence result provides a novel justification for impulse response matching: By commonality of general equilibrium feedback, impulse responses to particular aggregate structural shocks can be informative for many different counterfactuals. This idea has a clear conceptual antecedent in the microeconomic program evaluation literature: Marschak’s Maxim suggests that economists should try to identify the *combinations* of structural parameters needed for policy analysis, rather than the hard-to-estimate parameters themselves (Marschak, 1974; Heckman, 2010). Demand equivalence suggests that fiscal multipliers are precisely such a useful combination. As such, using the language of Nakamura & Steinsson (2018), they can be a useful “identified moment” for the quantitative discipline of structural macro models.

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 under more general assumptions. 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.

³A common early antecedent to this line of work is Hall (2009), who argued that “the effects of higher consumer purchases [are likely to be] similar to the effects of higher government purchases” – an identical *output* multiplier. My additive decomposition relies on the same intuition and is equivalent to scaling micro estimates by one plus the *consumption* or *investment* multipliers associated with fiscal purchases.

2 Consumption demand equivalence

This section develops an exact equivalence result for the general equilibrium propagation of private consumption demand and public spending shocks. Section 2.1 outlines a benchmark quantitative business-cycle model. In Section 2.2 I proceed to formally define my notions of direct “partial equilibrium” responses and indirect “general equilibrium” feedback. To build intuition, Section 2.3 develops the equivalence result in a stylized special case with closed-form solution. In Section 2.4 I return to the rich benchmark class of models and give a simple set of sufficient conditions for exact demand equivalence.

2.1 The benchmark model

Time is discrete and runs forever, $t = 0, 1, \dots$. 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 mostly focus on two such innovations: first, a one-off transfer to households, and second, a transitory expansion in government spending. To nevertheless emphasize the generality of the demand equivalence result, I also consider a third shock: 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, ...).

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 three structural shocks indexed by $s \in \{\tau, g, v\}$ – tax rebates, government spending, and household impatience. I write individual shock paths as $\boldsymbol{\varepsilon}_s$, and use subscripts $\boldsymbol{\varepsilon}$ for transitions after a path $\boldsymbol{\varepsilon} \equiv (\boldsymbol{\varepsilon}'_{\tau}, \boldsymbol{\varepsilon}'_g, \boldsymbol{\varepsilon}'_v)'$. 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 $(u, \tau) \neq (s, 0)$.

⁴This 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).

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 in their baseline discount factor β_i . The discount factor of every household is further subject to an additional common shifter ζ_t , with $\zeta = \zeta(\varepsilon_v)$. Households can self-insure by investing in liquid nominal bonds b_{it}^h , with nominal returns i_t^b and subject to a borrowing constraint \underline{b} .⁵ Household 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

$$\max_{\{c_{it}, b_{it}^h\}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_i^t \zeta_t(\varepsilon_v) u(c_{it}, c_{it-1}, \ell_{it}) \right] \quad (1)$$

such that

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

and

$$b_{it}^h \geq \underline{b}$$

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,

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

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.⁶ Since

⁵I consider an extension with liquid and illiquid assets in Section 4.3.

⁶A 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

the wage-setting problem is standard, I relegate details to Appendix A.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.

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

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

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 $\lim_{t \rightarrow \infty} \left(\prod_{s=0}^t \frac{1 + \pi_s}{1 + i_{s-1}^b} \right) b_t = 0$.

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 A.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.⁷ To allow interpretation of perfect foresight transition paths as conventional first-order perturbation

labor supply for each individual household. I consider a model without unions in Appendix D.1.1.

⁷More 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.

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 A.1); I discuss an extension to transition paths with other starting points in Appendix D.5.

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.

The results in Section 2.4 will show what extra restrictions on this canonical model class are needed to attain an exact demand equivalence result and so justify my claims about model identification and empirical counterfactuals in Section 3.

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 partial equilibrium responses and indirect general equilibrium adjustment.

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

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

I decompose this aggregate impulse response into two parts: a direct “partial equilibrium” impulse and an indirect “general equilibrium” feedback part.⁸

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

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

Similarly, let the indirect (general equilibrium) feedback be

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

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

$$\hat{\mathbf{c}}_{\boldsymbol{\varepsilon}} = \hat{\mathbf{c}}_{\boldsymbol{\varepsilon}}^{PE} + \hat{\mathbf{c}}_{\boldsymbol{\varepsilon}}^{GE} \quad (4)$$

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 $\hat{\mathbf{c}}_d^{GE}$ of private demand shocks $d \in \{\tau, v\}$ – the demand equivalence result. In Section 3 I then argue that (i) cross-sectional regressions estimate the direct spending response $\hat{\mathbf{c}}_d^{PE}$ and (ii) it is in practice often possible to recover the aggregate effects of public spending shocks that can proxy for $\hat{\mathbf{c}}_d^{GE}$.

2.3 A simple example of demand equivalence

The intuition for the demand equivalence result is easily illustrated using a particular special case of my benchmark model – a simple spender-saver real business-cycle (RBC) model. In this 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

⁸My 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). 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., 2017; Fagereng et al., 2018). Nevertheless, in Appendix D.1.1, 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).

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.

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} . I assume that the tax increases (transfer cuts) $\hat{\tau}_e$ used to finance the two policies exclusively fall on savers. All model equations are stated in Appendix A.2.1.

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.

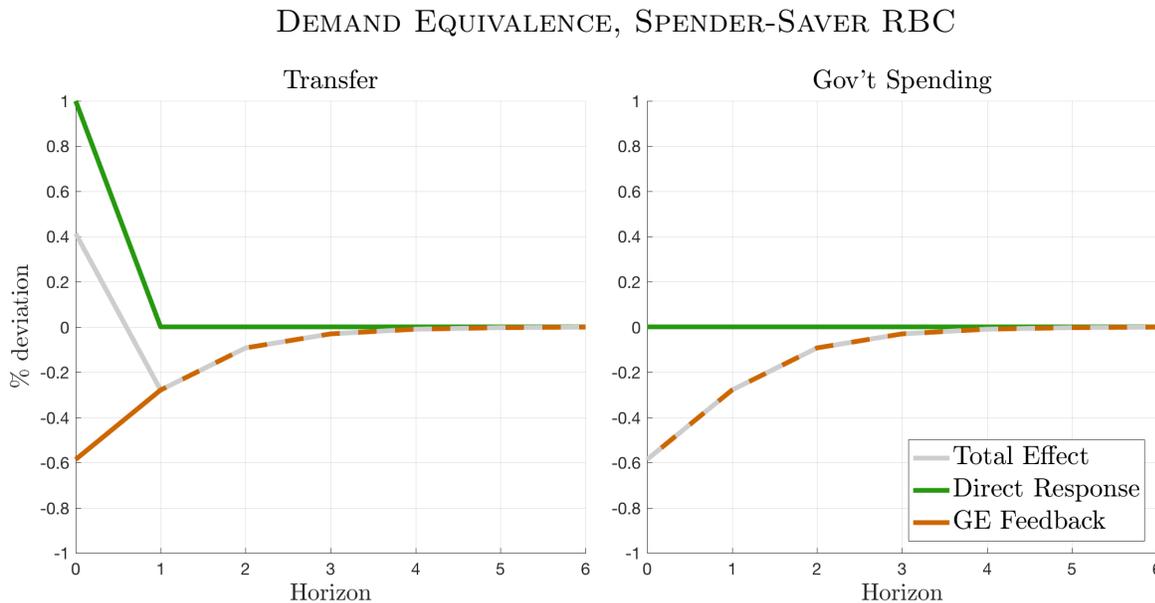


Figure 1: Impulse response decompositions after equally large, one-off tax rebate and government spending shocks in the simple spender-saver RBC model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

The left panel shows the consumption response to a one-off transfer, normalized to increase partial equilibrium consumption demand by one per cent. 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, future output drops and income declines.

The right panel then shows the consumption response to a one-period expansion in government spending, normalized to increase total fiscal consumption of the final good by one per cent of steady-state private consumption. By definition, household consumption does not respond directly to this second shock (the green line). In general equilibrium, consumption 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 appears to be 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 my simple spender-saver model.

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

$$\hat{c}_\tau^{GE} = \hat{c}_g^{GE} \quad (5)$$

and so

$$\hat{c}_\tau = \underbrace{\hat{c}_\tau^{PE}}_{PE \text{ response}} + \underbrace{\hat{c}_g}_{= GE \text{ feedback}} \quad (6)$$

Irrespective of the model parameterization, the *total* response of consumption to a government spending shock can proxy for the missing general equilibrium intercept of the private spending change.

PROOF. It is straightforward to establish the decomposition in Proposition 1 through the familiar closed-form solution of log-linearized RBC models. As I show in Appendix A.2.1, the response paths of (log-linearized) capital and consumption follow

$$\hat{k}_t = \alpha \times \hat{k}_{t-1} - \frac{1 - \alpha\beta}{1 - \lambda(1 - \alpha\beta)} \times (\varepsilon_{\tau t} + \varepsilon_{gt}) \quad (7)$$

$$\hat{c}_t = \alpha \times \hat{k}_{t-1} + \frac{\alpha\beta}{1 - \lambda(1 - \alpha\beta)} \times (\varepsilon_{\tau t} + \varepsilon_{gt}) - \varepsilon_{gt} \quad (8)$$

The key observation is that both shocks enter the law of motion for the capital stock (7) identically. In other words, consumption demand and government spending shocks have equal effects on capital accumulation, and so also output, interest rates, and wages. The sole difference between the two shocks lies in how the common amount of net output (output less investment) is split between household and government consumption, as evident from (8).

Unfortunately, this proof strategy is not particularly constructive – it relies on the *explicit* closed-form solution of the model, which of course will not be available for quantitatively relevant model variants. Instead, I find it convenient to write the equilibrium as a dynamic system of market-clearing equations (and prices adjusting to clear those markets).

Lemma 1. *Consider a shock path ε . Sequences of real rates \mathbf{r} and taxes on savers $-\tau_e$ are part of a perfect foresight equilibrium if and only if*

$$\mathbf{c}(\mathbf{r}, \mathbf{w}(\mathbf{r}), \mathbf{d}(\mathbf{r}), \tau_e; \varepsilon) + \mathbf{g}(\varepsilon) = \mathbf{y}(\mathbf{r}) - \mathbf{i}(\mathbf{r}) \quad (9)$$

$$\tau_e = \tau_e(\varepsilon) \quad (10)$$

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} .

Given a path of real interest rates \mathbf{r} , optimal firm behavior gives production \mathbf{y} , investment \mathbf{i} , and payments to households as dividends \mathbf{d} and wages \mathbf{w} . Similarly, given total household income and returns to saving, optimal household behavior implies a path for consumption demand \mathbf{c} ; finally, the path of government spending \mathbf{g} is exogenous. It is of course immediate that any possible equilibrium path of interest rates \mathbf{r} and saver transfers (taxes) τ_e must be such that the output market clears (equation (9)) and the government budget constraint holds (equation (10)). Lemma 1 then merely asserts that these conditions are also *sufficient*: Equilibria in the economy are fully characterized by adjustments of one intertemporal price – real interest rates – to clear one market – the output market.

Once this result is established, it is a small leap to go to Proposition 1: Totally differentiating both sides of (9) - (10) we find that, to first order,

$$\underbrace{\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \varepsilon} + \frac{\partial \mathbf{g}}{\partial \varepsilon} \\ \frac{\partial \tau_e}{\partial \varepsilon} \end{pmatrix}}_{\text{excess demand}} \times \varepsilon = \underbrace{\begin{pmatrix} \frac{\partial \mathbf{y}}{\partial \mathbf{r}} - \frac{\partial \mathbf{i}}{\partial \mathbf{r}} - \frac{\partial \mathbf{c}}{\partial \mathbf{r}} & -\frac{\partial \mathbf{c}}{\partial \tau_e} \\ \mathbf{0} & I \end{pmatrix}}_{\text{GE adjustment}} \times \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\tau}_e \end{pmatrix} \quad (11)$$

The initial disturbance ε 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 $\hat{\tau}_e$ solves

(11) for a tax rebate shock ϵ_τ . Then the same path $(\hat{\mathbf{r}}, \hat{\boldsymbol{\tau}}_e)$ also solves (11) for a government spending shock with the same intertemporal demand profile – that is, if $\hat{\mathbf{c}}_\tau^{PE} = \hat{\mathbf{g}}_g$. Intuitively, for general equilibrium feedback, it does not matter *why* there is a given amount of excess demand, or *why* there is a shortfall in the intertemporal government budget constraint – it just matters *how much*.⁹

INTERPRETATION. The decomposition (6) shows that, at least in the spender-saver model, government spending impulse responses are a useful *sufficient statistic* for the general equilibrium feedback effects associated with an income tax rebate. By the statement of Proposition 1, this result is not tied to any particular model parameterization; for example, it holds for arbitrary values of the saver discount rate or the capital share in production.

The proof strategy suggests that the demand equivalence result should in fact be quite a bit more general: Ultimately, the proof only relies on the sequential equilibrium characterization (9) - (10), and so should – for example – be invariant to largely arbitrary changes in production functions and preferences. Of course, the ability to characterize a model’s equilibrium through such a single market-clearing condition in a single intertemporal price is highly restrictive. However, as I show next, a variant of this proof strategy can be applied to justify the decomposition (6) in a very rich family of structural macro models.

2.4 A general equivalence result

This section establishes my most general consumption demand equivalence result. I will first state the result and its underlying assumptions, and then provide further intuition by linking the proof strategy back to the simple model.

Consumption demand equivalence relies on three key assumptions. The first assumption is implicitly embedded in the model of Section 2.1, but I explicitly state it here for emphasis.

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

The second assumption relates to the interest rates faced by households and government. The model already imposes that all agents borrow and lend at a common interest rate; As-

⁹Formally, the heuristic argument given here ensures only that a solution for the rebate transition path is also a solution for a particular public spending transition path; it is, however, silent on the existence and uniqueness of such transition paths. For the simple 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 (11) has a unique left-inverse. I provide further details in Appendix A.2.1.

sumption 2 again re-states this model property for emphasis, 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 ε_τ or ε_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.*

The third assumption restricts the economy's labor market. In response to the partial equilibrium increase in consumption demand \hat{c}_d^{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 $\hat{\ell}_d^{PE}$, defined formally in Appendix A.1. My third assumption provides two possible sufficient conditions to guarantee that $\hat{\ell}_d^{PE} = \mathbf{0}$.

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 $\{\tau, g, v\}$, the equilibrium transition path exists and is unique. Under Assumptions 1 and 2, the responses of consumption to a private demand shock d (either impatience v or tax rebate τ) and to a government spending shock g with $\hat{\mathbf{g}}_g = \hat{c}_d^{PE}$ satisfy, to first order,*

$$\hat{c}_d = \hat{c}_d^{PE} + \hat{c}_g + \text{error}\left(\hat{\ell}_d^{PE}\right) \quad (12)$$

where the error function is equal to $\mathbf{0}$ if $\hat{\ell}_d^{PE} = \mathbf{0}$. Under the additional Assumption 3,

$$\hat{c}_d = \underbrace{\hat{c}_d^{PE}}_{PE \text{ response}} + \underbrace{\hat{c}_g}_{= GE \text{ feedback}} \quad (13)$$

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 Lemma C.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 $\hat{\mathbf{g}}_g = \hat{\mathbf{c}}_d^{PE}$ – 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.¹⁰ 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 ϵ_d . Assumption 3 – a restriction on household behavior – is enough to rule this out: Following the shock ϵ_d , households either do not wish to or are not able to directly adjust their hours worked, i.e. $\hat{\ell}_d^{PE} = \mathbf{0}$. Together, Assumptions 1 to 3 ensure exact demand equivalence.

In the proof of Proposition 2, I establish the existence of a “demand multiplier” \mathbf{D} – a map transforming partial equilibrium net excess demand paths (such as $\hat{\mathbf{c}}_d^{PE}$ 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 \mathbf{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 2.1.¹¹ The intuition for such a common “demand multiplier” is particularly transparent in the standard static Keynesian cross, and was for example previously discussed in Hall (2009).

APPROXIMATION ERROR. The decomposition in (12) shows that, in the family of business-cycle models nested by the outline in Section 2.1, the *only* channel through which demand equivalence may fail is differential labor adjustment: Households that receive the rebate may decide to optimally work less. I assess the plausibility of this mechanism in two ways. First,

¹⁰Note 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 need not be (and I assume is not) financed through any change in taxes or transfers.

¹¹Auclert & 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.

in the remainder of this section, I analyze its strength in two particular numerical examples. Consistent with conventional wisdom in the recent business-cycle literature (e.g. Christiano, 2011a,b) I find that even moderate degrees of price and wage stickiness are sufficient to largely neuter the aggregate effects of transitory changes in labor supply. Second, in Section 4.2, I provide direct empirical discipline on the error term, and conclude that it is robustly small.

My first example is a heterogeneous-agent New Keynesian (HANK) model. The model falls into the benchmark class of Section 2.1 and features uninsurable income risk, moderate degrees of nominal price and wage stickiness, and several further frictions familiar from standard business-cycle models (e.g. investment adjustment costs, variable capital utilization, and a rich Taylor rule). Its parameterization is close to that of the estimated HANK model of Section 4; I relegate further details on model structure and parameterization to that section as well as Appendix A.2.2. Importantly, Assumption 1 holds and fiscal policy in the model is consistent with Assumption 2, but household preferences – of the typical separable kind – feature strong short-run wealth effects, and wages re-set every 2.5 quarters on average.

I compare impulse responses to a one-off income tax rebate and to a transitory increase in government spending. The two shocks give identical partial equilibrium net spending paths (normalized to 1% of steady-state consumption on impact) and are financed using identical delayed increases in taxes. Results are displayed in Figure 2.

APPROXIMATE DEMAND EQUIVALENCE, HANK MODEL

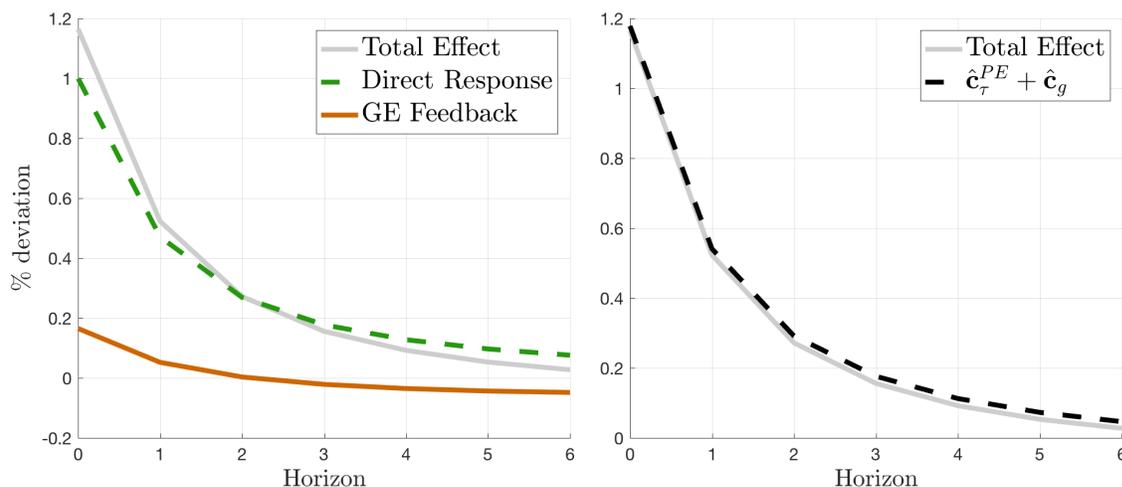


Figure 2: Impulse response decompositions and demand equivalence approximation in a simple HANK model, with details on the parameterization in Appendix A.2.2. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

The left panel decomposes the total consumption response to the rebate into direct partial

equilibrium effect and indirect general equilibrium feedback, in line with Definition 1. Since the model features a high average MPC, Keynesian multiplier effects dominate, and so general equilibrium effects further amplify the initial stimulus. The right panel approximates the consumption response by summing (i) the direct consumption response \hat{c}_τ^{PE} and (ii) the aggregate general equilibrium response of consumption to a similarly transitory and identically financed expansion in government spending, \hat{c}_g . Under Assumption 3, the decomposition would be exact, so the grey and black dotted lines would be indistinguishable. Instead, after receiving the rebate, households would like to work less, simultaneously pushing down consumption. The approximation $\hat{c}_\tau^{PE} + \hat{c}_g$ thus over-states the true consumption response; the associated error, however, is small, at just below 3 per cent of the true peak consumption response. Even with (moderately) sticky wages, labor is largely demand-determined, and so small and transitory shifts in labor supply are quantitatively irrelevant.

The second example is the popular quantitative business-cycle model of Justiniano et al. (2010). I solve the model at their estimated posterior mode, but – to allow for non-trivial effects of aggregate income tax rebates – add a small fringe λ of hand-to-mouth households.

APPROXIMATE DEMAND EQUIVALENCE, JUSTINIANO ET AL. (2010)

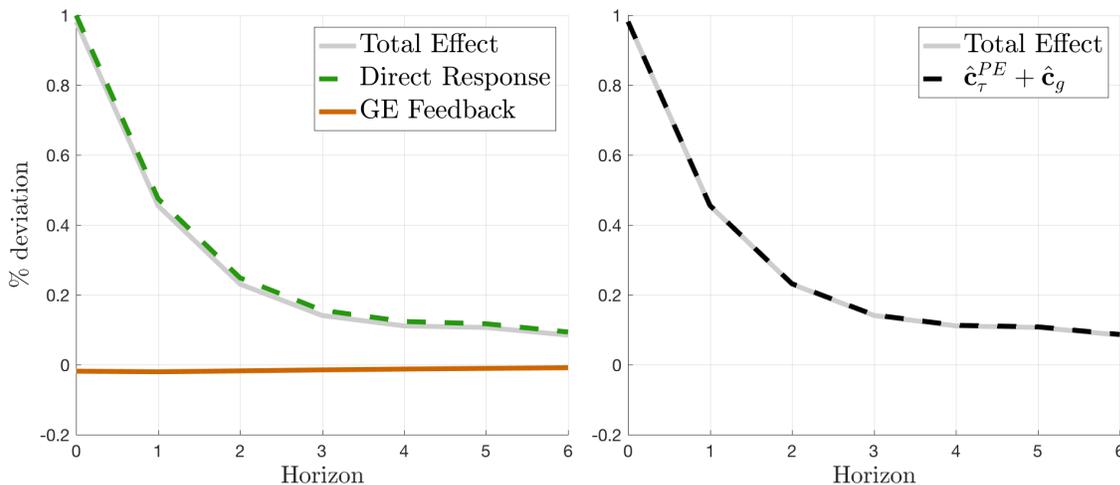


Figure 3: Impulse response decompositions and demand equivalence approximation in the model of Justiniano et al. (2010), solved at the posterior mode and with a fraction $\lambda \rightarrow 0$ of spenders. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

For the numerical experiment in Figure 3, I let $\lambda \rightarrow 0$, but keep the effective size of the rebate $\varepsilon_\tau \times \lambda$ fixed. Specifically, I consider a sequence of rebates given to spenders, inducing a spending response similar to the intertemporal demand profile in my estimated

HANK model.¹² Figure 3 plots the resulting aggregate impulse response decompositions. As before, general equilibrium feedback is the result of a complicated interaction of several model features. The flat orange line reveals that, at the model’s estimated mode parameterization, the crowding-in effects associated with higher household income just happen to be almost exactly offset by interest rate crowding-out. Crucially, however, in response to an equally large increase in government spending, consumption also barely moves, and so the additive approximation of Proposition 2 is again highly accurate, with a maximal error of only 0.4 per cent of the true peak consumption response.¹³ Relative to the HANK model, the approximation is even more accurate since wages are (much) stickier.

For additional insights on this documented near-equivalence, it is instructive to further dissect the approximation error in (12). As I show in the proof of Proposition 2, the error term can be re-expressed as the full general equilibrium response of consumption to a particular labor supply (wage cost-push) shock. The near-equivalence in Figures 2 and 3 is consistent with the discussion in Christiano (2011a): Even with moderate wage and price stickiness, the effects of transitory shifts in labor supply are largely neutralized in general equilibrium. However, my analytical results on the composition of the error term also justify a more direct measurement strategy: As I show in Section 4.2, micro and macro data can jointly provide direct empirical discipline on the size of the error, and robustly imply that it is small.

CONCLUSIONS. The analysis in this section has demonstrated that, in a large and empirically relevant *class* of structural models, private and public spending shocks share (either exactly or approximately) identical general equilibrium propagation.

This equivalence result is, however, completely silent on the *strength* of those common general equilibrium effects. In the simple spender-saver model of Section 2.3, partial equilibrium spending responses are crowded-out; in my two quantitative examples, feedback effects are instead relatively weak. In Appendix D.3, I show two extreme examples, one with full crowding-out, the other with strong amplification, yet both featuring exact demand equivalence. 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.

¹²It is straightforward to show that, in the limit $\lambda \rightarrow 0$ but with $\varepsilon_\tau \times \lambda = \text{constant}$, the aggregate dynamics of the model are identical to that of Justiniano et al. (2010), and the income tax rebate shock enters the household consumption-savings problem exactly like a standard impatience shock.

¹³Unsurprisingly, in both models, the approximation deteriorates for highly persistent shocks, as wages are not permanently sticky. Instead, the implied persistent shifts of labor supply materially affect aggregate quantities, and so the approximation error becomes larger. I provide an illustration in Appendix D.1.3.

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.

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.¹⁴

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 argue that (i) the econometric estimands of such cross-sectional regressions are often interpretable as direct (partial equilibrium) shock responses \hat{c}_d^{PE} and (ii) we can use estimates of the aggregate effects of changes in government spending to proxy for their missing general equilibrium intercept \hat{c}_d^{GE} .

MODEL. As before, I develop all arguments in the context of the structural model of Section 2.1. In my theoretical analysis of demand equivalence, the proof strategy dictated a focus on perfect foresight transition paths. For standard cross-sectional and macro regression estimands to be well-defined, however, I need a proper notion of aggregate risk. I thus now consider the linear vector moving-average representation induced by the first-order perturbation solution of the model, assuming that the shocks ε_{st} , $s \in \{\tau, g, v\}$ are mutually i.i.d. and $N(0, 1)$.¹⁵ I use s subscripts to indicate impulse response functions to such *one-time*

¹⁴More specifically, direct projection on measures of aggregate shocks (proxy variables) is ruled out by their endogeneity (Ramey, 2016). Other, more involved macro structural identification approaches are hindered by (i) likely non-invertibility due to the relative infrequency of shocks (Plagborg-Møller, 2019) and (ii) the dearth of plausible macro exclusion restrictions (e.g. zero restrictions in structural VAR representations).

¹⁵Previous studies that exploit the first-order equivalence of perturbation and perfect foresight transitions for estimation include Mankiw & Reis (2007) and Auclert et al. (2019). It is immediate from the properties

structural shocks; by certainty equivalence, these impulse responses are to first order identical to the transition paths for one-off structural shocks studied in Section 2, thus justifying the re-use of notation.

Finally, to introduce cross-sectional heterogeneity in shock exposure, I further assume that the rebate and impatience shocks faced by household i satisfy $\varepsilon_{sit} = \xi_{sit} \times \varepsilon_{st}$, where ξ_{sit} is i.i.d. across households and time (and uncorrelated with any household characteristics), with $\mathbb{E}(\xi_{sit}) = 1$ and $\text{Var}(\xi_{sit}) > 0$. In the proof of Proposition 3 I show that, under my assumptions on the exposure term ξ_{sit} , 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 demand shocks $d \in \{\tau, v\}$ takes the form

$$c_{it+h} = \alpha_i + \delta_t + \beta_{dh} \times \varepsilon_{dit} + u_{it+h}, \quad h = 0, 1, 2, \dots \quad (14)$$

where α_i and δ_t are individual and time fixed effects.¹⁶

It is straightforward to show that, under my assumptions, regressions such as (14) 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 $\{\varepsilon_{dit}\}$ generated from the linear vector moving average representation of the structural model of Section 2.1. Then the ordinary least-squares estimand of $\beta_d \equiv (\beta_{d0}, \beta_{d1}, \dots)'$ satisfies*

$$\beta_d = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \varepsilon_{d0}} di = \hat{\mathbf{c}}_d^{PE} \quad (15)$$

In words, regressions such as (14) do not estimate the true macro counterfactual $\hat{\mathbf{c}}_d$, but instead give a household-level average treatment effect that is interpretable as a *partial equilibrium* response, $\hat{\mathbf{c}}_d^{PE}$ – precisely the object defined in my decomposition in Definition 1. Obtaining such estimates from a sequence of cross-sectional micro regressions like (14) is the first step of my methodology.

of linear VMA representations that arbitrary further shocks could be added without affecting my results.

¹⁶The regression in (14) 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 my method to such regressions in Wolf (2019).

GENERAL EQUILIBRIUM EFFECTS. To map the micro estimates β_d 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 equivalence result in Proposition 2 suggests 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. Suppose that the econometrician can 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. Furthermore suppose that, for some linear combination of shocks with weights $\{\gamma_k\}$, it is the case that

$$\beta_d = \sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{g}}_{g_k} \quad (16)$$

In words, a linear combination of government spending shocks available from macro experiments gives similar partial equilibrium excess demand pressure as the private demand shock ε_d . 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 β_d and so shocks ε_d .^{17,18} It then remains to gauge the accuracy of the financing Assumption 2. For example, if the researcher is interested in counterfactuals for a deficit-financed rebate, then the composite public spending shock $\sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{g}}_{g_k}$ should also be deficit-financed. If so, then we can invoke Proposition 2 to conclude that

$$\hat{\mathbf{c}}_d^{GE} = \sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{c}}_{g_k} \quad (17)$$

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

$$\hat{\mathbf{c}}_d = \underbrace{\beta_d}_{\text{PE response}} + \underbrace{\sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{c}}_{g_k}}_{\text{GE feedback}} \quad (18)$$

¹⁷If researchers are willing to ignore anticipation effects of news shocks, then *any* path β_d can be replicated exactly with evidence on just a single government spending shock $\hat{\mathbf{g}}_{g_k}$. In my applications I do not need this additional assumption, but it is interesting to note that, at least for the government spending shocks studied in previous work, anticipation effects do appear rather limited (e.g. Ramey, 2011).

¹⁸In Appendix D.1.8 I document approximate equivalence when (16) is roughly satisfied.

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.¹⁹

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 E, I study the effects of a one-off (budget-neutral) income re-distribution from rich (low-MPC) to poor (high-MPC) households.²⁰ Finally, in Section 5, I establish a theoretical investment demand equivalence result, and use my two-step approach to estimate the aggregate 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 income tax rebate (i.e., lump-sum transfer). My main finding is that full general equilibrium counterfactuals are close to direct micro estimates: The partial equilibrium increase in consumption demand is accommodated one-for-one through an increase in output, with relatively limited general equilibrium crowding-in or -out.

DIRECT RESPONSE. I first require an estimate of the direct spending response \hat{c}_τ^{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.,

¹⁹Except 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 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 B.4).

²⁰Such budget-neutral re-distribution is also the topic of Auclert & Rognlie (2018). Due to the scarcity of evidence on heterogeneity in dynamic intertemporal MPCs across households, I in this application rely on a standard *partial equilibrium* consumption-savings problem to construct the direct response \hat{c}_d^{PE} . The mapping into macro counterfactuals is then again completed using the semi-structural second step from (17).

2006; Parker et al., 2013; Jappelli & Pistaferri, 2014; Fagereng et al., 2018).²¹ 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 to income tax rebates, estimate a differenced version of the micro regression (14); building on Proposition 3, Appendix B.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.²² In the left panel of Figure 4, 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 partial equilibrium spending response \hat{c}_{τ}^{PE} .

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 that government spending expansions – 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. In the context of my structural model, the IV relevance and exclusion

²¹By 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.

²²These estimates include the durables spending response. As I show in Appendix D.1.4, demand equivalence extends without change to a model with durable and non-durable consumption. To ensure consistency, my VAR analysis also throughout contains measures of total consumption. For completeness, however, I have repeated my analysis using evidence on non-durables consumption only. The direct partial equilibrium consumption response is smaller and more persistent, and the aggregate non-durables consumption response to a similarly persistent government spending is an even more tightly estimated 0.

MEASURING \hat{c}_τ^{PE} & \hat{c}_g

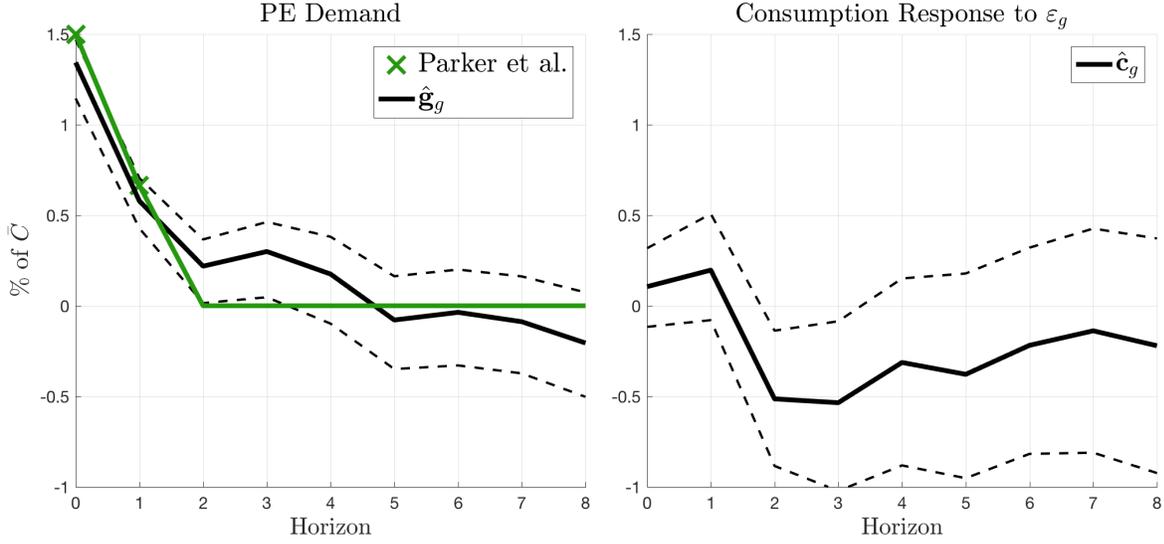


Figure 4: The left panel shows direct consumption responses to income tax rebate (green) vs. direct government spending response to identified spending shock (black). Estimated consumption responses from Parker et al. (2013) (Table 3). The dashed lines for the government spending shock correspond to 16th and 84th percentile confidence bands. The right panel shows the response of consumption to the same identified spending shock.

restrictions can be phrased follows:

Assumption 4. *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 $\tilde{z}_t \equiv z_t - \mathbb{E}(z_t | \{z_{t-l}, y_{t-l}\}_{l=1}^\infty)$ satisfies*

$$E(\tilde{z}_t \cdot \varepsilon_{gt}) \neq 0, \quad E(\tilde{z}_t \cdot \varepsilon_{ju}) = 0 \text{ for all } (j, u) \neq (g, t) \quad (19)$$

As shown in Plagborg-Møller & Wolf (2019), estimating a recursive vector autoregression (VAR) in instrument and macro aggregates $(z_t, y_t)'$, with the instrument ordered first, correctly identifies the impulse responses of all macro aggregates y_t to a structural innovation ε_{gt} in aggregate government spending, up to a scale parameter that is independent of horizon and response variable:

Proposition 4. *(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 4. Let θ_y denote the vector of impulse responses of y to the first shock in a*

recursive SVAR. Then the ordinary least-squares estimand of θ_y satisfies

$$\theta_y = \text{constant} \times \hat{y}_g \quad (20)$$

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

Relative impulse responses are thus identified – and since impulse responses will be re-scaled to match \hat{c}_τ^{PE} , relative responses are sufficient to recover \hat{c}_g .²³ Also note that the consistency proof requires no assumptions on invertibility of the shock ε_{gt} , mitigating concerns about timing and anticipation (Ramey, 2011; Leeper et al., 2013). I thus estimate a recursive VAR in forecast errors z_t and aggregates y_t , where the vector y_t includes measures of overall government spending, output, consumption, investment, hours worked, taxes, and government debt. To plausibly estimate \hat{c}_g in a stable macroeconomic regime, I restrict my sample to range from the third quarter of 1981 to the fourth quarter of 2007. Further details on exact variable definitions, data construction and the estimation procedure are relegated to Appendix B.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 4. 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.²⁴ The right panel shows the corresponding response of aggregate consumption – \hat{c}_g . Consumption appears to be somewhat crowded-in on impact, and mildly crowded-out in the following quarters. Overall, \hat{c}_g is close to 0 throughout, and in fact reasonably tightly estimated. Finally, in Appendix B.3, I show that the expansion in government spending leads to a delayed increase in taxes, as well as a persistent rise in total government debt. By Assumption 2, my counterfactuals for a transitory income tax rebate should thus be interpreted as pertaining to a particular, quite persistently *deficit-financed*

²³Strictly speaking, the estimated impulse response \hat{g}_g takes into account general equilibrium feedback to government spending; for the demand matching (16), any such feedback needs to be filtered out. In principle this can be done in two ways. First, researchers may simply assume that there is no such feedback. Second, if the *structural* government spending equation can be identified (e.g. as in Blanchard & Perotti, 2002), then feedback effects can be removed manually. For simplicity I choose the first path. Encouragingly, however, results using the second approach are almost identical, simply because I find limited feedback from macro aggregates to government expenditure, exactly as in previous work (e.g. Caldara & Kamps, 2017).

²⁴While the two demand paths are quite similar, they of course do not align perfectly. In Appendix D.1.8, I discuss the accuracy of my approximation under imperfect demand matching.

one-off transfer to households.²⁵

MACRO COUNTERFACTUALS. To construct a valid general equilibrium counterfactual, it now simply remains to sum the empirically estimated \hat{c}_τ^{PE} and \hat{c}_g . The results are displayed in the left panel of Figure 5. Note that, for construction of the plot, I take the point estimate of \hat{c}_τ^{PE} as given, and only account for macroeconomic estimation uncertainty.²⁶

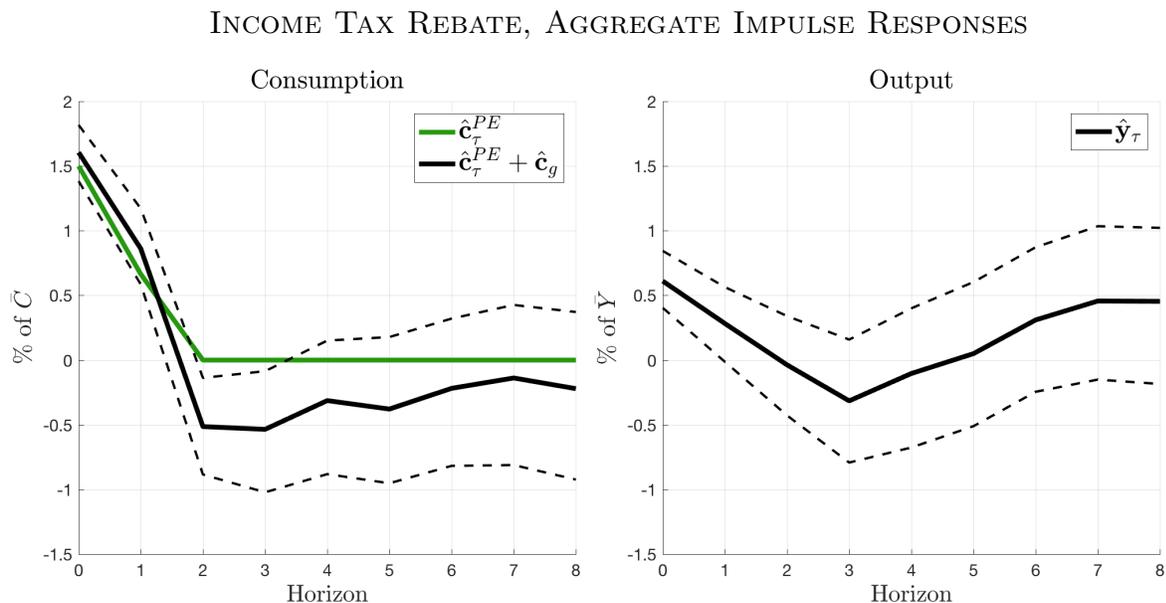


Figure 5: Consumption and output responses to an income tax rebate shock. The full consumption response is computed following the exact additive decomposition of Proposition 2, while the output response is simply equal to the response after a government spending shock. The dashed lines again correspond 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 $\hat{c}_\tau^{PE} + \hat{c}_g$ – appears to be quite close to the (large) micro-estimated direct spending response \hat{c}_τ^{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 then shows the corresponding impulse response of output which, by the

²⁵Other fiscal spending episodes could be used to construct counterfactuals for other financing schemes. For a detailed review of different spending episodes and their financing, see Ferriere & Navarro (2018).

²⁶This 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.

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.

My analysis suggests that, at least for income tax rebate stimulus, the “missing intercept” of general equilibrium feedback is a relatively tightly estimated 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. 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.²⁷

IMPLICATIONS FOR STRUCTURAL MODELING. My results have implications for structural modeling similar to those of the sufficient statistics characterizations in Chetty (2009) and Arkolakis et al. (2012): *Any* structural analysis that estimates a consumption response to tax rebates different from Figure 5 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.

The estimated HANK model of Section 4.1 is an example of a structural model that satisfies (i) and (ii): it features approximate demand equivalence (recall Figure 2) and large direct spending responses due to large household-level marginal propensities to consume (see Appendix A.2.2).²⁸ However, the overall model generates fiscal multipliers somewhat in excess of 1, and so – consistent with the equivalence result – it suggests that direct micro effects are a slight under-estimate of full counterfactuals. With somewhat less sticky prices and more aggressive monetary policy, the model almost perfectly matches my empirically estimated fiscal policy impulse responses (see Appendix D.4). This re-estimated HANK model is a promising laboratory for further *structural* analysis.²⁹

²⁷My pre-crisis VAR implicitly measures “normal-time” general equilibrium effects. As I discuss in Appendix D.5, 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).

²⁸Auclert et al. (2018) show that HANK models of the kind considered in my structural analysis can closely match empirically documented paths of intertemporal MPCs. An analogue of Figure 4 with a model-based direct spending response (instead of the estimated one) thus unsurprisingly looks very similar.

²⁹More conventional medium-scale DSGE models usually imply multipliers below 1 (Gechert, 2015; Ramey,

4 Approximation accuracy

The theoretical equivalence result of Section 2 relies on three main assumptions: (i) the existence of a common final good, (ii) identical borrowing and lending rates for households and government, and (iii) zero (short-run) wealth effects of labor supply, or fully rigid wages. All three assumptions are presumably incorrect. This section systematically studies the role of each in ensuring that the general equilibrium counterfactuals computed in Section 3 remain at least approximately valid.

I do so in two steps. First, in Section 4.1, I present a rich structural model, estimated to be consistent with evidence on both microeconomic household behavior and the time series properties of macroeconomic aggregates – in other words, a model suitable for structural fiscal policy counterfactuals. In this model, the approximation is accurate. Then, in Sections 4.2 to 4.4, I consider several departures from this benchmark, allowing me to study in isolation the importance of each individual assumption.

4.1 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. Arguably, the model is rich enough to serve as a quantitative laboratory for structural analysis of the aggregate effects of generic private and public spending shocks; as such, it is an example of a structural model that could plausibly be used to identify the missing general equilibrium intercept of, say, transitory transfer payments to households. I provide a brief outline of the model and my estimation strategy here, and relegate further details to Appendix A.2.2.

The household block is slightly more general than that of the benchmark model in Section 2.1. In particular, household borrowing in the liquid asset is now only possible at a penalty rate $i_t^b + \kappa$. As a result, households and government discount at different interest rates, and so equally large expansions in private and public spending cannot be financed using identical paths of taxes and transfers. 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 impatience and government spending shocks discussed in

2016). Hand-to-mouth behavior is thus central to the documented consistency between model and data.

Section 2.1, I also include shocks to total factor productivity and the marginal efficiency of investment, to price and 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.³⁰ 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.

RESULTS. I 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 6.

APPROXIMATE DEMAND EQUIVALENCE, ESTIMATED HANK MODEL

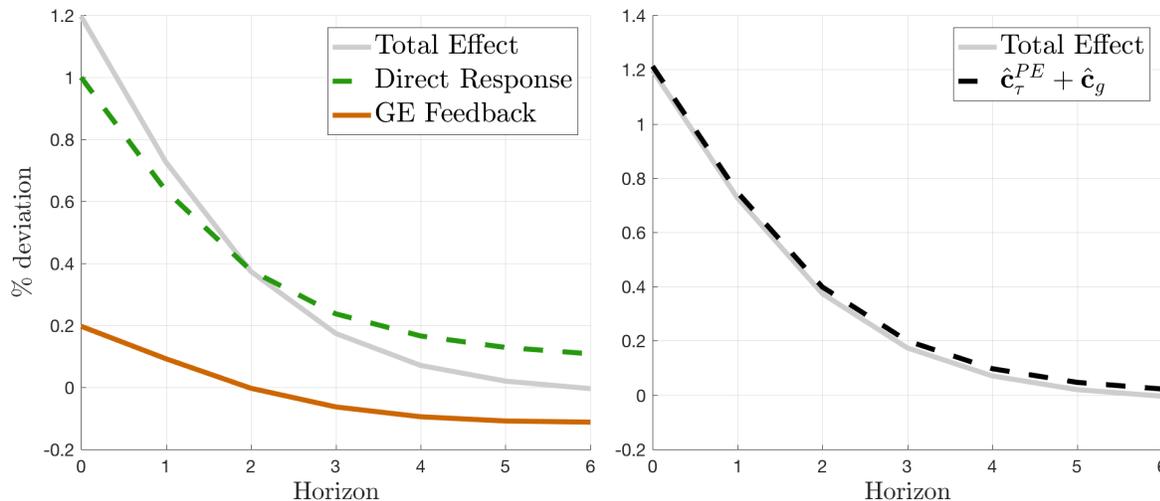


Figure 6: Impulse response decompositions and demand equivalence approximation in the estimated HANK model, with details on the parameterization in Appendix A.2.2. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

³⁰Specifically, I closely follow Justiniano et al. (2010) and include measures of output, inflation, a short-term interest rate, consumption, investment, hours worked, and a measure of aggregate wages.

The plot looks extremely similar to Figure 2. Of course this is unsurprising – the models are almost identical, differing only in the presence of a borrowing wedge. In terms of error metrics, the approximation here is in fact better than in the previous model: the maximal error (relative to the true peak response of aggregate consumption) is around 2.2 per cent, and around 3 per cent in the simpler model. The intuition is as follows: In the model with borrowing wedge, indebted households use the rebate to pay down (high-return) debt. As a result, the average return faced by households is higher than that faced by the government. Taxes thus have to rise by more to finance the government spending expansion $\hat{\mathbf{g}}_g$ compared to the consumption stimulus \hat{c}_τ^{PE} , and so the demand equivalence approximation tends to *under-state* the aggregate effects of a tax rebate, partially offsetting the labor supply error. I further elaborate on this intuition and on the size of the associated error in Section 4.3.

Figure 6 only reveals that my approximation is accurate at the estimated posterior mode of a particular structural model. In Appendix D.1.7 I go one step further and show that most of the estimated parameters governing model dynamics – including the monetary rule, the nature of investment adjustment costs, and the degree of variable capacity utilization in production – are in fact largely orthogonal to the accuracy of the approximation. Formally, I *randomly* draw model parameters from large supports, solve the implied model, and compute the approximation accuracy. I find that, of all estimated parameters, only the degree of price rigidity has a material impact on the accuracy of the approximation.³¹ However, even with near-flexible prices, and fixing the relatively moderate calibrated wage rigidity, the approximation error remains at only 9 per cent of the peak consumption response.

4.2 Labor supply

The demand equivalence approximation is accurate in all models studied so far because of sufficiently strong nominal rigidities. With fully flexible prices and wages, the strong wealth effects of my conventional separable preferences will invariably break the demand equivalence logic. Figure 7 provides an illustration, using the estimated structural HANK model of the previous section, but with flexible prices and wages.

The right panel shows that the quality of the demand equivalence approximation deteriorates sharply. Intuitively, following the tax rebate, households consume more and so – because the marginal utility of consumption is lower – optimally choose to work less. Labor

³¹With rigid prices, labor is demand-determined. Shifts in labor supply thus only affect relative wage and dividend pay-outs, and so the approximation is accurate (up to a redistributive effect).

FAILURE OF DEMAND EQUIVALENCE, FLEX-PRICE HA MODEL

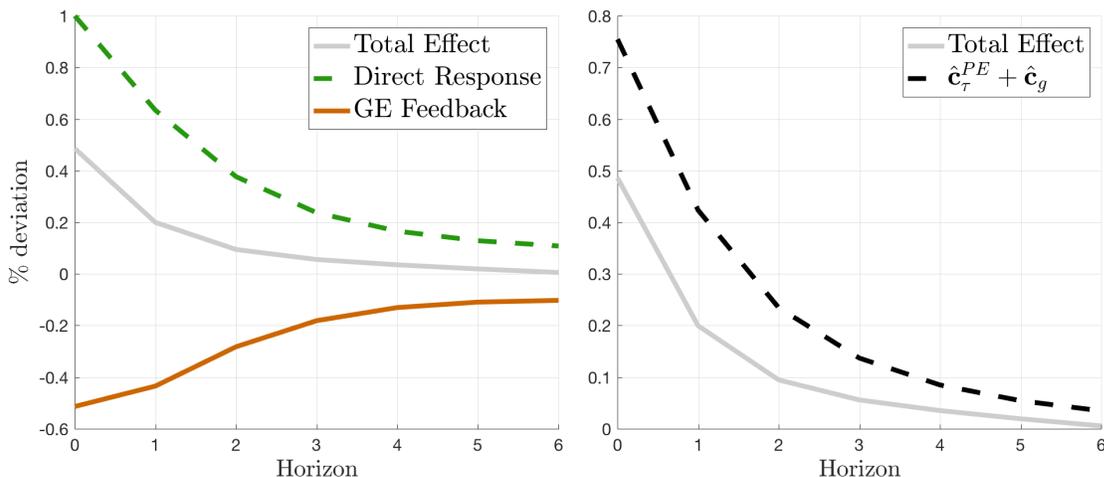


Figure 7: Impulse response decompositions and demand equivalence approximation in the estimated HANK model, but with flexible prices and wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

supply is not similarly reduced after expansions in government spending, so the demand equivalence approximation over-states the aggregate consumption response.

The mechanism underlying this inaccuracy is, however, sharply at odds with all kinds of micro and macro evidence. First, my estimation exercise based on standard aggregate time series data as usual calls for nominal rigidity. Without such rigidity, the model would feature large countercyclical (and counterfactual) swings in wages and inflation. Second, 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). As emphasized above, even very moderate degrees of nominal wage stickiness – with re-sets occurring every 2.5 quarters – are enough to make the demand equivalence approximation highly accurate. Third, the flexible-price model implies that, in response to the rebate, households would like to reduce their earnings by almost as much as they increase their spending. As emphasized by Auclert & Rognlie (2017), such large negative earnings responses are an inescapable feature of macro models with large wealth effects of labor supply, large average MPCs, and flexible wages. Micro data instead suggest an earnings response an order of magnitude smaller than the average MPC (e.g. Cesarini et al., 2017; Fagereng et al., 2018).³² In models with flexible

³²The strength of wealth effects in labor supply has also been estimated in macro data. Such studies routinely favor near-zero wealth effects (Schmitt-Grohé & Uribe, 2012; Born & Pfeifer, 2014).

wages but data-consistent small short-run wealth effects, the approximation is instead again highly accurate (see Appendix D.1.1).

Finally, as I show in Appendix D.2, it is actually possible to adapt my two-step methodology to directly account for the labor supply error term in the decomposition of Proposition 2. As discussed in Section 2.4, this 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 can thus recover a direct *empirical* correction for the error. Unsurprisingly, since empirical estimates of the desired labor supply contraction are small, the results of this augmented procedure are almost identical to the benchmark estimates of Section 3.2.

4.3 Interest rates

If households and government borrow and lend at different rates, then identical changes in private and public partial equilibrium 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.³³

In the estimated structural model of Section 4.1, some households pay down high-return liquid debt, so the implicit household discount rate is high. The associated bias, however, is negligible; in a model variant with fully rigid wages, the maximal error is equal to 0.7 per cent of the true peak consumption response. The intuition for this quantitative near-irrelevance 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

³³A very similar logic also applies to open economies: If home bias in private consumption demand is low (high) relative to home bias in government spending, then taxes need to increase by more (less) to finance private relative to public spending, and so the demand equivalence approximation will tend to over-state (under-state). The full formal argument is omitted in the interest of space but available upon request.

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.

Empirical evidence suggests that, in response to lump-sum transfer receipts, households mostly adjust their – arguably low-return – liquid deposits (Fagereng et al., 2018, Table 4). In Appendix D.1.2, I analyze the inaccuracy associated with such low household returns in a rich *two*-asset HANK model, similar to Kaplan et al. (2018). To threaten the quality of my approximation as much as possible, I assume that households earn the government interest rate on illiquid assets, and face a substantial return penalty of 1 per cent *per quarter* for transacting in liquid assets. As a result, household returns are weakly – and for most transactions strictly – below the interest rate paid on government debt. The separate biases associated with wealth effects in labor supply and return heterogeneity thus both push the demand equivalence approximation to over-state the aggregate consumption response. Even under these extreme assumptions, however, the approximation remains quite accurate, with a maximal error relative to the peak consumption response of around 7 per cent.³⁴

4.4 Beyond one-good economies

Exact equivalence requires households and government to consume a single, homogeneous final good. This section considers various deviations from this benchmark: (i) durable and non-durable consumption goods, (ii) valued and productive government spending and (iii) multiple goods with imperfect factor mobility and heterogeneous production functions.

DURABLES. All models considered so far abstract from the empirically relevant distinction between durable and non-durable consumption goods. As it turns out, even in a generalized model with separate durable and non-durable consumption, demand equivalence obtains under exactly the same assumptions as those discussed in Section 2.4. I relegate the formal argument to Appendix D.1.4, and only briefly discuss the intuition here.

The key assumption – routinely made in previous work featuring durable and non-durable consumption (e.g. Barsky et al., 2007; Berger & Vavra, 2015) – is that both the durable and the non-durable good can be produced costlessly from a common final good. If that is the case, then a generalized demand equivalence result applies to total household expenditure on non-durable *and* durable consumption.

³⁴If a researcher has a strong prior about return heterogeneity, then the implied difference in net present values can simply be returned to households as an additional rebate stimulus, perfectly analogous to my analysis of the effects of consumption complementarities in Appendix D.1.5.

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 D.1.5 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. I only briefly summarize my main conclusions here, with details and further intuition presented in the appendix.

I first show that, if private and public consumption are complements (substitutes), then the demand equivalence approximation is likely to over-state (under-state) the consumption response to the demand shifter. However, given a standard parametric form for the non-separability (as in Leeper et al., 2017), it is easy to correct for this bias. 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 caution against the use of public investment multipliers for my approximations. I illustrate the associated inaccuracy in an extension of my benchmark model where government expenditure directly shows up in aggregate production functions (following Leeper et al., 2010).

MULTI-GOODS MODELS. Private and public consumption baskets are different. Previous work has identified at least three channels through which such heterogeneity may break demand equivalence. First, if factors of production do not move freely between different sectors, then differences in consumption baskets will lead to heterogeneous relative price responses (Ramey & Shapiro, 1998). Second, if different goods have different production technologies, then the income generated by private and public demand shocks may flow to different factors of production, leading to heterogeneous general equilibrium propagation (Baqae, 2015; Alonso, 2017). And third, if firm investment demand features a higher intertemporal elasticity than private consumption demand, then government purchases of consumption and investment goods have different aggregate demand effects (Boehm, 2016).

The strength of these mechanisms is best tested directly with evidence on the aggregate effects of different *kinds* of government purchases. Reassuringly, with the notable exception of productive long-term investment, previous empirical work finds largely homogeneous multipliers by the type of spending (Gechert, 2015; Ramey, 2016). Since the resulting estimates are noisy, however, Appendix D.1.6 also presents indirect model-based evidence; specifically, I study the accuracy of the demand equivalence approximation in several multi-sector gen-

eralizations of the benchmark model, disciplined by empirical evidence on (i) the strength of relative price effects, (ii) heterogeneity in network-adjusted labor shares, (iii) the intertemporal elasticity of consumption and investment demand. I give intuition for the signs of the associated biases, but largely find that – in the empirically disciplined variants of these extended models – the asymmetry in multipliers is sufficiently small so as to not materially threaten the accuracy of the demand equivalence approximation, with maximal prediction errors always below 10 per cent.

5 Investment demand counterfactuals

This section extends my methodology to study the general equilibrium propagation of shocks to investment demand. Section 5.1 establishes 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

I again consider the benchmark model of Section 2.1. Relative to consumption demand shocks, the additional challenge of fluctuations in investment demand is that, through firms’ production technologies, investment will sooner or later translate into additional hiring and production. This adds two complications: First, firm behavior will induce a *net excess demand* path $i - y$, rather than a pure spending response. Second, tax financing of investment stimulus, changes in firm dividend pay-outs, as well as any potential expansion in labor hiring will have redistributive implications and directly feed back into consumption demand.

In this section, I will give sufficient conditions under which the net excess demand path $i - y$ fully determines general equilibrium feedback, thus allowing this aggregate feedback to be exactly replicated by a mix of expansionary and contractionary government spending shocks. Intuitively, the expansionary shock mirrors the impact excess (investment) demand, while the contractionary news shock synthesizes the implied future expansion in supply. As I will show, these sufficient conditions do not at all constrain the richness of the model’s investment block, but do impose some meaningful restrictions on household behavior.

THE EQUIVALENCE RESULT. The production block of the economy was sketched in Section 2.1, with a detailed outline in Appendix A.1. For purposes of the analysis here, it suffices to note that I allow for a rich set of real and financial frictions, including (convex

and non-convex) capital adjustment costs as well as a generic set of constraints on firm equity issuance and borrowing. The production block of the economy is thus general enough to nest essentially all recent contributions to the quantitative heterogeneous-firm investment literature (e.g. Khan & Thomas, 2008, 2013; Winberry, 2018; Koby & Wolf, 2019); importantly, my equivalence results do not require *any* restrictions on this rich firm side.

Anticipating the empirical application, I establish an exact equivalence result for investment tax credit shocks ϵ_q – shocks that reduce the cost of capital purchases at time t by an amount $\tau_{qt} = \tau_{qt}(\epsilon_q)$.³⁵ I obtain this investment demand equivalence result under four key restrictions on the non-production block of the economy. As before, I state the assumptions and the result first, and provide intuition as well as a proof sketch afterwards.

The first assumption – a single common final good – is again implicit in the model set-up.

Assumption 5. *A single final good is used for (government) consumption and investment.*

Implicitly, I 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 then rules out any *redistributional* effects associated with the firm subsidy.

Assumption 6. *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.*

This assumption effectively imposes a standard representative-household structure. The third assumption again concerns household labor supply decisions.

Assumption 7. *The Frisch elasticity of labor supply is either infinite (linear labor disutility), or wages are perfectly sticky.*

Linear labor disutility – clearly at odds with micro data on household labor supply – is sometimes justified at the aggregate level as a by-product of labor indivisibility (Hansen, 1985; Rogerson, 1988). Finally, the fourth assumption restricts monetary policy feedback.

Assumption 8. *The monetary authority's interest rate rule does not include an endogenous response to fluctuations in the level of aggregate output.*

³⁵More 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.

I define direct (partial equilibrium) responses and indirect (general equilibrium) feedback for firm investment and production exactly analogously to Definition 1, using the implied aggregate investment and production functions $\mathbf{i}(\bullet)$ and $\mathbf{y}(\bullet)$, respectively. Assumptions 5 to 8 are then enough for the following investment demand equivalence result.

Proposition 5. *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 5 to 8, the responses of investment and output to an investment tax credit shock q and to a government spending shock g with $\hat{\mathbf{g}}_g = \hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE}$ satisfy, to first order,*

$$\hat{\mathbf{i}}_q = \hat{\mathbf{i}}_q^{PE} + \hat{\mathbf{i}}_g \quad (21)$$

$$\hat{\mathbf{y}}_q = \hat{\mathbf{y}}_q^{PE} + \hat{\mathbf{y}}_g \quad (22)$$

The proof strategy is identical to that of the consumption demand equivalence result. First, Assumption 5 implies that I can consider a single aggregate output market-clearing condition. Second, Assumption 6 ensures the absence of any distributional effects that may lead to differential partial equilibrium *consumption* demand responses to the investment demand and public spending shocks, allowing me to restrict attention to the firm net demand path $\mathbf{i} - \mathbf{y}$. Third, Assumption 7 is sufficient to ignore labor market adjustments, either because households are willing to supply the additional demanded labor, or because they have no choice. And fourth, Assumption 8 is needed to ensure that only the *level* of net excess demand matters, not its composition. Without this assumption, the central bank would lean against any excess demand simultaneously associated with higher output *supply*, breaking demand equivalence.³⁶

ACCURACY. Assumptions 5 to 8 are routinely imposed in quantitative general equilibrium models of investment; in particular, they hold – and thus exact equivalence applies – in the well-known models of Khan & Thomas (2008), Khan & Thomas (2013), Winberry (2018), Ottonello & Winberry (2018) and Bloom et al. (2018). As such, the decomposition in (21) - (22) provides a useful exact identification result for a popular class of models.

To further gauge the accuracy of my approximation, Appendix D.1.9 studies shocks to investment demand in the estimated HANK model of Section 4.1. This model is a useful laboratory because it violates Assumptions 6 to 8: households are subject to non-trivial earnings risk and receive heterogeneous firm profit payments, the Frisch elasticity of labor

³⁶Alternatively, exact equivalence would obtain if the monetary authority responds to the output *gap*.

supply is relatively small (it is 1), wages are quite flexible, and the monetary authority responds to fluctuations in aggregate output. Even though each of these model ingredients individually biases my demand equivalence approximation upwards, I find that it remains accurate, in particular at short horizons.

5.2 Application: bonus depreciation

The investment equivalence result justifies a two-step procedure to study generic investment demand shocks, exactly analogous to my analysis of consumption shifters in Section 3. In this section I leverage the additive decomposition in Proposition 5 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).³⁷

DIRECT RESPONSE. My estimates of the direct response of investment to the shock rely heavily on Zwick & Mahon (2017) and Koby & Wolf (2019), who exploit cross-sectional firm-level heterogeneity in the exposure to bonus depreciation investment stimulus. In Koby & Wolf (2019), we estimate dynamic regressions akin to (14) 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 B.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 standard 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

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

Thus, given estimates of the capital depreciation rate δ , the capital share α , and the returns 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 $\nu = 1$.

³⁷As 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.

I take the regression estimates of $\hat{\mathbf{i}}_{qt}^{PE}$ for $t = 0, 1, 2, 3$ straight from Koby & Wolf (2019, Table 1). The green x 's in the investment panel of Figure 8 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 (but less persistent than) 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 Proposition 5, it remains to replicate the estimated net excess demand path through a suitable list of government spending shocks:

$$\hat{\mathbf{i}}_q^{PE} - \hat{\mathbf{y}}_q^{PE} = \sum_{k=1}^{n_k} \gamma_k \times \hat{\mathbf{g}}_{gk} \quad (23)$$

It is unlikely that any single estimated spending shock can replicate the reversal documented in Figure 8. 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 anticipation 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

INVESTMENT TAX CREDIT, IMPULSE RESPONSES

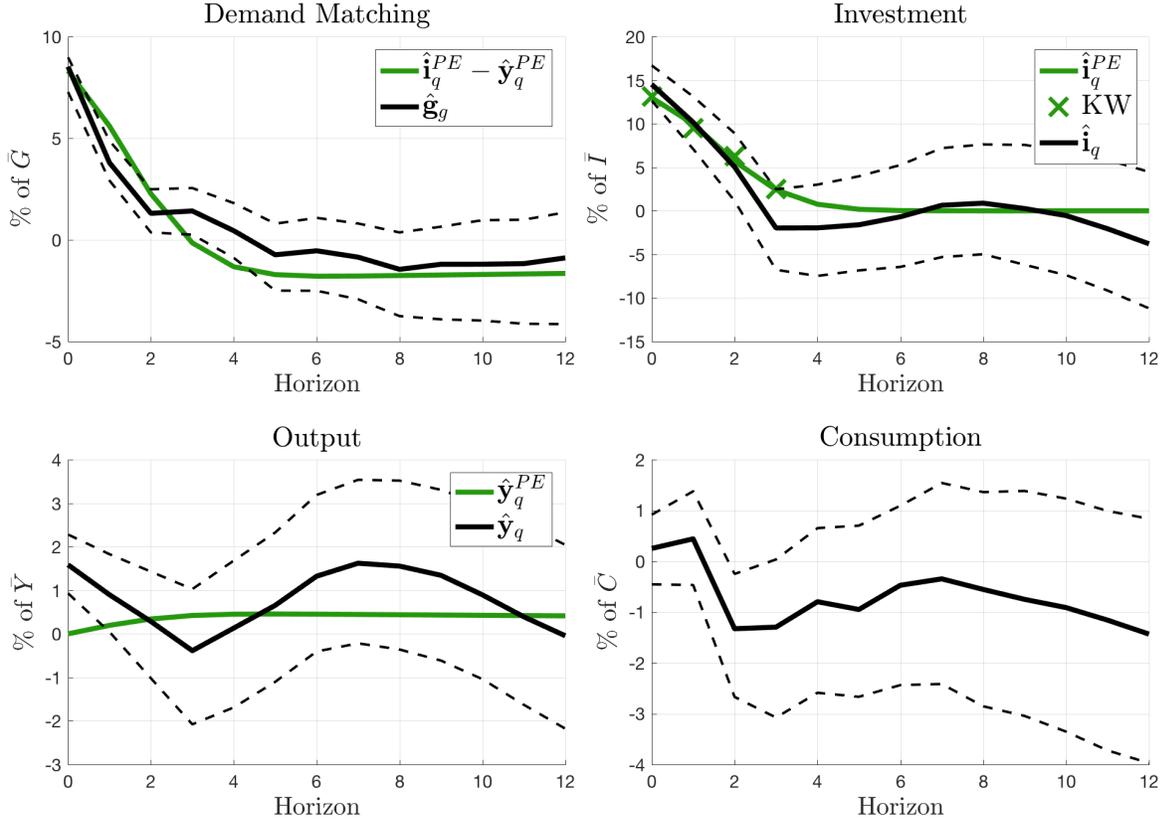


Figure 8: Investment, output and consumption responses to an investment tax incentive shock, with the partial equilibrium net output response path matched to a linear combination of government spending shocks. The investment and output responses are computed in line with Proposition 5, while the consumption response is simply equal to the response after the identified combination of government spending shocks. The dashed lines again correspond to 16th and 84th percentile confidence bands.

net excess demand path of the investment demand shock, as shown in the top left panel of Figure 8. Further details on the empirical implementation (in particular the construction of standard errors for general equilibrium feedback) are provided in Appendix B.3.

MACRO COUNTERFACTUALS. All results for general equilibrium counterfactuals are displayed in Figure 8. With the requirement that $\hat{g}_g = \hat{i}_q^{PE} - \hat{y}_q^{PE}$ satisfied, the investment and output panels implement the additive decompositions in (21) and (22), respectively. My main finding is that the substantial partial equilibrium investment demand responses

estimated in Zwick & Mahon (2017) and Koby & Wolf (2019) 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.

IMPLICATIONS FOR STRUCTURAL MODELING. My results contrast sharply with the predictions of the standard neoclassical model closure routinely entertained in quantitative models featuring rich investment micro-heterogeneity (e.g. Khan & Thomas, 2013; Bloom et al., 2018). As is well-known (e.g. Barro & King, 1984), investment demand shocks in such models are accommodated through large drops in consumption and only moderate impact increases in output. Real interest rates thus increase, leading to substantial general equilibrium crowding-out of rate-sensitive firm investment (Khan & Thomas, 2008). The results in Figure 8 are instead consistent with models featuring (i) relatively price-inelastic investment,³⁸ (ii) strong aggregate demand effects and (iii) little consumption crowding-out, for example due to hand-to-mouth spending or strong habit formation.

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

³⁸It may seem strange to claim that the large partial equilibrium investment responses documented in the cross-sectional regressions are consistent with price-inelastic investment. As we show in Koby & Wolf (2019), while these responses are indeed large in economic terms, they are orders of magnitude *smaller* than predicted by standard neoclassical models of investment.

not feature demand equivalence, or is inconsistent with conventional estimates on 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 my 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 extensions to future work.

References

- Agarwal, S., Liu, C., & Souleles, N. S. (2007). The reaction of consumer spending and debt to tax rebates—evidence from consumer credit data. *Journal of Political Economy*, *115*(6), 986–1019.
- 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.
- 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.
- Arkolakis, C., Costinot, A., & Rodriguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, *102*(1), 94–130.
- 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.
- Baqae, 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.
- Beraja, M., Hurst, E., & Ospina, J. (2019). The aggregate implications of regional business cycles. Technical report, National Bureau of Economic Research.
- 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., SaportaEksten, I., & Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, 86(3), 1031–1065.

- Boehm, C. (2016). Government spending and durable goods. Working Paper.
- 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.
- Canova, F. & Paustian, M. (2011). Business cycle measurement with some theory. *Journal of Monetary Economics*, 58(4), 345–361.
- 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.
- Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1), 451–488.
- 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.
- Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1), 1–45.
- Cocci, M. D. & Plagborg-Møller, M. (2019). Standard errors for calibrated parameters. Princeton University.
- 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.
- Fajgelbaum, P. D., Morales, E., Suárez Serrato, J. C., & Zidar, O. (2018). State taxes and spatial misallocation. *The Review of Economic Studies*, 86(1), 333–376.
- 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.
- Ferriere, A. & Navarro, G. (2018). The heterogeneous effects of government spending: It's all about taxes. *FRB International Finance Discussion Paper*, (1237).
- 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. (2019). 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.
- Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics*, 16(3), 309–327.
- Heckman, J. J. (2010). Building bridges between structural and program evaluation approaches to evaluating policy. *Journal of Economic Literature*, 48(2), 356–98.
- 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. (2019). Aggregation in Heterogeneous-Firm Economies: A Sufficient Statistics Approach. 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.
- Leeper, E. M., Walker, T. B., & Yang, S.-C. S. (2010). Government investment and fiscal stimulus. *Journal of Monetary Economics*, 57(8), 1000–1012.
- Leeper, E. M., Walker, T. B., & Yang, S.-C. S. (2013). Fiscal foresight and information flows. *Econometrica*, 81(3), 1115–1145.
- Mankiw, N. G. & Reis, R. (2007). Sticky information in general equilibrium. *Journal of the European Economic Association*, 5(2-3), 603–613.
- Marschak, J. (1974). Economic measurements for policy and prediction. In *Economic Information, Decision, and Prediction* (pp. 293–322). Springer.
- 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.
- Ottanello, 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. (2019). Bayesian inference on structural impulse response functions. *Quantitative Economics*, *10*(1), 145–184.
- 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.
- Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3), 649–665.
- Rogerson, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics*, 21(1), 3–16.
- 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.

A Model details

This appendix provides additional details on the rich class of structural models underlying my exact demand equivalence results. In Appendix A.1 I outline the full model and offer a formal definition of equilibrium transition paths. Appendix A.2 then discusses the particular parametric model variants used for illustration of approximate equivalence results.

A.1 The benchmark model

A.1.1 Full model outline

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. For the general estimated HANK model of Section 4, the only change is that I allow for a borrowing wedge; that is, the liquid interest rate satisfies

$$i^b(b^h) = \begin{cases} i^b & \text{if } b^h \geq 0 \\ i^b + \kappa_b & \text{if } b^h < 0 \end{cases}$$

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:³⁹

$$\begin{aligned} \pi_t^w(1 + \pi_t^w) &= \frac{\varepsilon_w}{\theta_w} \ell_t^h \left[\int_0^1 \left\{ -u_\ell(c_{it}, c_{it-1}, \ell_t^h) - \frac{\varepsilon_w - 1}{\varepsilon_w} (1 - \tau_\ell) w_t e_{it} \{u_c(c_{it}, c_{it-1}, \ell_t^h)\} \right. \right. \\ &\quad \left. \left. + \beta \mathbb{E}_t[u_{c-1}(c_{it+1}, c_{it}, \ell_{t+1}^h)] \right\} di \right] + \beta \pi_{t+1}^w (1 + \pi_{t+1}^w) \end{aligned} \quad (\text{A.1})$$

³⁹For 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.

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} , (A.1) provides a simple restriction on total labor supply $\boldsymbol{\ell}^h$.⁴⁰ Note that, without idiosyncratic labor productivity risk and so common consumption $c_{it} = c_t$, the derived wage-NKPC (A.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 (A.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 $\mathbf{s}^u = (\boldsymbol{\pi}, \mathbf{w}, \mathbf{c})$. 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 vector 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 $1 + r_t^b \equiv \frac{1 + i_{t-1}^b}{1 + \pi_t}$.⁴¹

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

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

such that

$$\begin{aligned} d_{jt}^I &= \underbrace{p_t^I y_{jt} - w_t \ell_{jt}}_{\pi_{jt}} - \xi_{jt} \times \mathbf{1}_{i_{jt} \neq 0} - (1 - \mathbf{1}_{i_{jt} < 0} \times \varphi) i_{jt} - \phi(k_{jt}, k_{jt-1}, i_{jt}, i_{jt-1}) \\ &\quad - b_{jt}^f + \frac{1 + i_{t-1}^b}{1 + \pi_t} b_{jt-1}^f \end{aligned}$$

⁴⁰In the special case $\theta_w \rightarrow \infty$, equation (A.1) is vacuous, so then I instead simply assume that $\boldsymbol{\ell}^h = \boldsymbol{\ell}^f$.

⁴¹Along 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. If firm saving and borrowing in the liquid asset is not constrained, then this choice of discount rate is needed to prevent arbitrarily large desired saving or borrowing.

$$\begin{aligned}
y_{jt} &= y(e_{jt}, u_{jt}k_{jt-1}, \ell_{jt}) \\
i_{jt} &= k_{jt} - [1 - \delta(u_{jt})]k_{jt-1} \\
-b_{jt}^f &\leq \Gamma(k_{jt-1}, k_{jt}, \pi_{jt}) \\
d_{jt}^I &\geq \underline{d}
\end{aligned}$$

Adjustment costs have a convex and continuously differentiable part ϕ , a firm-specific fixed adjustment cost ξ_{jt} (distributed with cdf $F(\xi)$ over support \mathbb{R}^+), and may feature partial irreversibility, with $\varphi \in [0, 1]$. 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 $\ell^f(\bullet)$, 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:

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

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

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

With some algebra, it is straightforward to show that in fact

$$d_t = y_t - w_t \ell_t - i_t$$

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, with an example of a concrete fiscal rule given in Appendix A.2.2. 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

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

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 condition⁴²

$$c_t + i_t + g_t = y_t$$

A.1.2 Equilibrium definition

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

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 $\boldsymbol{\ell}^h = \boldsymbol{\ell}^h(\mathbf{s}^h)$, 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,

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

⁴²So 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).

and the labor market clears,

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

for all $t = 0, 1, 2, \dots$

A.2 Parametric special cases

The quantitative illustrations and accuracy checks in Sections 2 and 4 largely rely on two particular structural models: the simple spender-saver RBC model and the estimated HANK model. This section provides further details on both models and discusses my preferred parameterizations.

A.2.1 The 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

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

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:

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

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

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

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

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

such that

$$d_t = y_t - i_t - w_t \ell_t \quad (\text{A.4})$$

$$y_t = k_{t-1}^{\alpha} \ell_t^{1-\alpha} \quad (\text{A.5})$$

$$i_t = k_t \quad (\text{A.6})$$

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

$$1 + r_t = \alpha k_t^{\alpha-1} \bar{\ell}^{1-\alpha} \quad (\text{A.7})$$

and wages satisfy

$$w_t = (1 - \alpha) k_{t-1}^{\alpha} \bar{\ell}^{-\alpha} \quad (\text{A.8})$$

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

$$g_t = \varepsilon_{gt} \quad (\text{A.9})$$

and sets rebates to spenders as

$$\tau_{ht} = \bar{\tau}_h + \frac{1}{\lambda} \varepsilon_{\tau t} \quad (\text{A.10})$$

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

$$\lambda \tau_{ht} + (1 - \lambda) \tau_{rt} + g_t = 0 \quad (\text{A.11})$$

$$b_t = 0 \quad (\text{A.12})$$

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

$$y_t = c_t + i_t + g_t \quad (\text{A.13})$$

where

$$c_t = \lambda c_{ht} + (1 - \lambda) c_{rt} \quad (\text{A.14})$$

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 (A.2) - (A.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 steady-state consumption 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 (A.13), and using (A.9), we get

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

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

$$\hat{c}_{rt+1} - \hat{c}_{rt} = (\alpha - 1)\hat{k}_t \quad (\text{A.16})$$

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

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

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

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

I solve the model exploiting the well-known equivalence between perfect foresight and first-order perturbation solutions. I thus treat (A.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

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

Plugging in and matching coefficients, we find that the guess is confirmed,⁴³ with

$$\begin{aligned} \theta_k &= \alpha \\ \omega_d &= -\frac{1 - \alpha\beta}{1 - \lambda(1 - \alpha\beta)} \end{aligned}$$

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

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

as claimed.

PARAMETERIZATION. For the graphical illustration in Figure 1 I use standard parameter values: $\beta = 0.99$, $\alpha = 1/3$ and $\lambda = 0.3$.

A.2.2 The estimated HANK model

The analysis in Section 4 builds on a rich 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 check in Section 2.4).

MODEL OUTLINE. The model is an extension of the rich baseline environment outlined in Section 2.1, violating Assumption 2 (households and government borrow at different interest rates) and Assumption 3 (strong wealth effects and imperfectly rigid wages).

⁴³It is straightforward to verify existence and uniqueness of the equilibrium following the arguments in Blanchard & Kahn (1980) or Sims (2000).

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

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

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

$$c_t^* \equiv \left[\int_0^1 e_{it} c_{it}^{-\gamma} di \right]^{-\frac{1}{\gamma}} \quad (\text{A.22})$$

Results are unchanged if I instead use the average marginal utility of aggregate consumption $-\gamma \hat{c}_t$ in the union wage target (as in Hagedorn et al., 2019). 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 model is populated by different illiquid wealth “types”, who each receive an exogenous (and time-invariant) endowment of illiquid shares.

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 then add structural shocks to output and investment productivity, monetary policy, price mark-ups and wage mark-ups to complement the already included impatience and government spending shocks. All shocks are modeled as in Justiniano et al. (2010), so I omit details. Finally, for purposes of the model estimation, I assume that

$$\hat{\tau}_{et} = -(1 - \rho_\tau) \times \hat{b}_{t-1} \quad (\text{A.23})$$

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 multiples of each other; without a borrowing wedge, they would be identical, at least in partial equilibrium. Since households spend most of the rebate

immediately, results are very similar if I instead simply use the rule (A.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 1.

STEADY-STATE PARAMETER VALUES, HANK MODEL

Parameter	Description	Value	Target	Model	Data
<i>Households</i>					
ρ_e, σ_e	Income Risk	-	Kaplan et al. (2018)	-	-
κ_d, p_d	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	-1	Kaplan et al. (2018)		
κ_b	Annual Borrowing Wedge	0.06	Fraction $b < 0$	0.15	0.15
<i>Firms</i>					
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
<i>Government</i>					
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 1: HANK model, steady-state calibration.

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 endowment, I split the illiquid wealth distribution from the 2016 SCF into four bins (< 15 , $15 - 50$, $50 - 85$, and > 85 percentiles), and then exactly match wealth in the four bins κ_w by allowing for four permanent illiquid wealth types with mass p_d . I set the average return on (liquid) assets in line 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. Households can borrow up to one time average quarterly labor earnings (which in turn are normalized to 1), and the borrowing wedge is set to discipline the fraction of households with negative liquid wealth. 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.⁴⁴ 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 30% 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 discrete-time variant of the methods developed in Ahn et al. (2017) 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.

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 estimation procedure then

⁴⁴More 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).

sticks as closely as possible to Justiniano et al. (2010): I consider the same set of macro observables (over the same time period), and impose identical priors whenever possible.⁴⁵ As a result, the estimation exercise does not really take advantage of the additional opportunities afforded by micro data; instead, it is merely a slightly more disciplined approach to arrive at a plausible parameterization for the non-household block of the model.

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 B.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 (A.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 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 D.1.7.⁴⁶

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

⁴⁵For the discussion of data construction I thus refer the interested reader to their appendix. I thank Brian Livingston for help in assembling the data.

⁴⁶The optimization routine is available at <http://sims.princeton.edu/yftp/optimize/>.

DYNAMICS PARAMETER VALUES, HANK MODEL

Parameter	Description	Density	Prior		Posterior
			Mean	Std	Mode
ϕ_p	Price Calvo Parameter	B	0.66	0.10	0.87
ζ	Capacity Utilization	G	5.00	1.00	3.71
κ	Investment Adjustment Cost	G	4.00	1.00	2.45
ρ_m	Taylor Rule Persistence	B	0.60	0.20	0.86
ϕ_π	Taylor Rule Inflation	N	1.70	0.30	2.08
ϕ_y	Taylor Rule Output	N	0.13	0.05	0.08
ϕ_{dy}	Taylor Rule Output Growth	N	0.13	0.03	0.06

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

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 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 accuracy check in Section 2.4 is identical to the estimated model except for one change: I set $\kappa_b = 0$ and $\underline{b} = 0$. As a result, households and government face identical interest rates, and any inaccuracy in my approximations is *exclusively* due to short-run wealth effects in labor supply.

B Empirical appendix

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

B.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), this appendix discusses how to interpret their regression estimands. Parker et al. estimate a differenced version of (14):

$$\Delta c_{it} = \text{time fixed effects} + \text{controls} + \beta_0 ESP_{it} + \beta_1 ESP_{it-1} + u_{it} \quad (\text{B.1})$$

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 $t = 0$ ($\varepsilon_{\tau i 0} = 1$), and that the remaining households receive the same rebate at $t = 1$ ($\varepsilon_{\tau i 1} = 1$). The model analogue of regression (B.1) is then

$$\Delta c_{it} = \delta_{\Delta t} + \beta_0 \varepsilon_{\tau it} + \beta_1 \varepsilon_{\tau it-1} + u_{it}, \quad t = 0, 1 \quad (\text{B.2})$$

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.⁴⁷ We can then

⁴⁷However, 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 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 4 is likely to be accurate.

My 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.

B.2 Direct response: micro investment elasticities

Koby & Wolf (2019) generalize the static analysis of Zwick & Mahon (2017) and estimate dynamic projection regressions of the form

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

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 from spanning the years 1993–2017; the sample period in particular 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. I briefly repeat the main insights here.

First, firms must not be subject to financial frictions. In the presence of financial frictions, the indicator $z_{n(j),t}$ does not remain a sufficient statistic summarizing the effects of a

given bonus depreciation policy. Second, *all* meaningful capital adjustment costs must be internal to the firm, while the aggregate supply of capital goods must be perfectly flexible.⁴⁸ Reassuringly, this assumption is consistent with the findings in House & Shapiro (2008), Edgerton (2010) and House et al. (2017), who all conclude that the supply of new capital goods is very elastic. Third, general equilibrium feedback associated with the investment demand stimulus should not co-vary with exposure to the stimulus itself. This assumption is satisfied if low- and high-depreciation firms do not systematically vary in their cyclicality (also see the discussion Zwick & Mahon, 2017).

Given the estimated partial 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{y}_q^{PE} , which requires parameter choices (α, ν, δ) . I have experimented with a wide range of parameter values, and found result to be robust; for example, with $\alpha = 0.3$ and $\nu = 0.8$ (a set-up closer to standard heterogeneous-firm model calibrations), the investment counterfactual barely changes.

B.3 The missing intercept: VAR estimation

My analysis of the transmission of transitory 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 forecast errors exactly as Ramey (2011). I then treat these forecast errors as a valid external instrument for structural government spending shocks, as formalized in Assumption 4. Following Plagborg-Møller & Wolf (2019), I study their transmission by ordering them first in a recursive VAR. In addition to the forecast error variable, the VAR contains measures of government spending, consumption, output, investment, taxes, and hours worked; in an expanded version, I also include total government debt.

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, a measure of the *federal* average marginal tax rate (Alexander & Seater, 2009),⁴⁹ and a measure of private business compensation. All variables are defined and measured as in Ramey (2011). To study the effects of the spending

⁴⁸This restriction is necessary because the regression (B.3) differences out aggregate capital price effects.

⁴⁹The 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.

shock on debt issuance, I construct a log per capita measure of total federal debt, deflating the nominal debt series in the St. Louis Fed’s FRED database (data series: GFDEBTN). In a further robustness check, I replace professional forecaster errors with Greenbook defense spending forecast errors, closely following the analysis in Drautzburg (2016).⁵⁰

ESTIMATION DETAILS. I specify both the benchmark and the extended VAR 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 B.1 shows the impulse responses of government spending, output, consumption, investment and taxes in the benchmark VAR, as well as the impulse response of total government debt in the expanded VAR.

As in most existing structural VAR work, I construct 16th and 84th percentile confidence bands; the output and debt 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 flat impulse response for 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. However, the tax response is somewhat sensitive to details of the model specification, and sometimes not significant.

ROBUSTNESS. My central results – the 1-1 increase in output, the limited crowding-out of private expenditure, and the persistent rise in debt – 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 spending, as in Ramey (2011). Depending on the set of included controls, the undershooting of consumption and investment is, in this earlier sample, usually

⁵⁰For 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.

BENCHMARK GOVERNMENT SPENDING SHOCK, VAR IRF'S

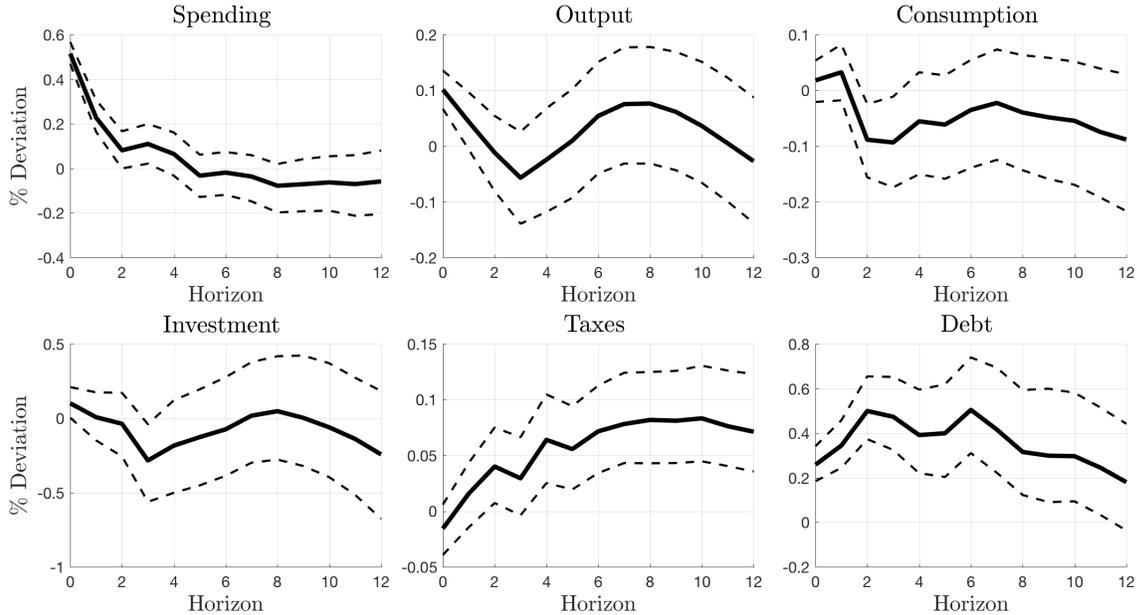


Figure B.1: Impulse responses after a one standard deviation innovation to the forecast error. The dashed lines correspond to 16th and 84th percentile confidence bands, constructed using 10,000 draws from the posterior distribution of the reduced-form VAR parameters.

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.⁵¹ 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).⁵² Second, replacing my benchmark measure of government spending forecast errors with Greenbook defense spending forecast errors leaves my results almost completely unchanged. This suggests that either (i) the benchmark VAR itself is largely picking up the response to military spending forecast errors

⁵¹Note, 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.

⁵²I 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.

or (ii) multipliers are invariant to the spending type (similar to Gechert (2015)). Third, removing individual controls in the benchmark specification does not materially impact the results. This is consistent with the intuition in Plagborg-Møller & Wolf (2019) – my estimands are projection coefficients which do not depend on the set of included controls, at least in terms of population estimands. Fourth, changes in the number of lags do not affect the overall flavor of my results. And fifth, frequentist inference (or a flat prior) gives almost identical impulse response estimates to those displayed here.

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.⁵³

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 (23). Detailed impulse response plots for this alternative identification scheme are available upon request.

B.4 Joint Uncertainty

I throughout ignore estimation uncertainty for microeconomic difference-in-differences estimators. 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, 2019). In principle, however, it is straightforward to account for joint estimation uncertainty: Under my

⁵³Also similar to Caldara & Kamps (2017), I find that additionally controlling for professional forecast errors has quite limited effects. Even in a pure Blanchard-Perotti recursive VAR (without the error control) I find an approximately unit fiscal spending multiplier and little response of private spending.

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.

I only provide a sketch of the argument here. To ease the notional burden, I consider a simple static model; the generalization to the dynamic case is conceptually straightforward, but more notationally involved. I assume that consumption of household i satisfies

$$c_{it} = \underbrace{\beta_\tau \varepsilon_{\tau t} + \beta_g \varepsilon_{gt} + \sigma_\nu \nu_t}_{\text{macro shocks}} + \underbrace{\beta_\tau^{PE} \xi_{it} + \sigma_\zeta \zeta_{it}}_{\text{micro shocks}}$$

where the macro shocks are distributed $(\varepsilon_{\tau t}, \varepsilon_{gt}, \nu_t)' \stackrel{iid}{\sim} N(0, I_3)$, and the micro shocks are distributed $(\xi_{it}, \zeta_{it})' \stackrel{iid}{\sim} N(0, I_2)$, independently of the macro shocks and any individual characteristics. It is then straightforward to see that a simple OLS estimator for (14) satisfies

$$\begin{aligned} \hat{\beta}_\tau^{PE} &= \frac{\frac{1}{N} \sum_{i=1}^N (c_{it} - \bar{c}_{it})(\varepsilon_{\tau it} - \bar{\varepsilon}_{\tau it})}{\frac{1}{N} \sum_{i=1}^N (\varepsilon_{\tau it} - \bar{\varepsilon}_{\tau it})^2} \\ &= \beta_\tau^{PE} + \frac{\frac{1}{N} \sum_{i=1}^N (\xi_{it} - \bar{\xi}_{it}) \sigma_\zeta (\zeta_{it} - \bar{\zeta}_{it})}{\frac{1}{N} \sum_{i=1}^N (\xi_{it} - \bar{\xi}_{it})^2} \end{aligned}$$

where the averages are taken over the N households i . Crucially, any macroeconomic uncertainty is differenced out. Similarly, by the projection arguments in Plagborg-Møller & Wolf (2019), a recursive VAR estimator gives

$$\begin{aligned} \hat{\beta}_g &= \frac{\frac{1}{T} \sum_{t=1}^T (c_t - \bar{c}_t)(\varepsilon_{gt} - \bar{\varepsilon}_{gt})}{\frac{1}{T} \sum_{t=1}^T (\varepsilon_{gt} - \bar{\varepsilon}_{gt})^2} \\ &= \beta_g + \frac{\frac{1}{T} \sum_{t=1}^T \beta_\tau (\varepsilon_{gt} - \bar{\varepsilon}_{gt})(\varepsilon_{\tau t} - \bar{\varepsilon}_{\tau t})}{\frac{1}{T} \sum_{t=1}^T (\varepsilon_{gt} - \bar{\varepsilon}_{gt})^2} + \frac{\frac{1}{T} \sum_{t=1}^T \sigma_\nu (\nu_t - \bar{\nu}_t)(\varepsilon_{\tau t} - \bar{\varepsilon}_{\tau t})}{\frac{1}{T} \sum_{t=1}^T (\varepsilon_{gt} - \bar{\varepsilon}_{gt})^2} \end{aligned}$$

where now all averages are taken over time t , in a total sample with length T . Importantly, by the proof of Proposition 3, micro shocks have no aggregate effects, so all estimation uncertainty for $\hat{\beta}_g$ is driven by macroeconomic uncertainty.

Since sampling uncertainty for the micro and macro estimators is independent, their covariance is zero, so construction of joint uncertainty bands is conceptually trivial. Under less stringent assumptions, it is always possible to construct conservative bounds using the methods developed in Cocci & Plagborg-Møller (2019).

C Proofs and auxiliary lemmas

C.1 Proof of Lemma 1

From the specification of the household and firm problems in Appendix A.2.1, 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 (A.8) and (A.4) we can also obtain $\mathbf{w} = \mathbf{w}(\mathbf{r})$, so the expression (9) is well-defined.

Next, since $\mathbf{g} = \mathbf{g}(\boldsymbol{\varepsilon})$ by (A.9), we can conclude that (9) is necessary for any perfect foresight transition equilibrium. Since $\boldsymbol{\tau}_e = \boldsymbol{\tau}_e(\boldsymbol{\varepsilon})$ by (A.11) it is similarly immediate that (10) is necessary. To show sufficiency, note that (A.2), (A.3) as well as (A.4), (A.5) and (A.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 (9) and (10) hold, then they are in fact part of a perfect foresight equilibrium. By existence and uniqueness of the perturbation solution (see Appendix A.2.1), 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. \square

C.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

$$\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{c}}{\partial \mathbf{r}} \times \hat{\mathbf{r}} + \frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}_e} \times \hat{\boldsymbol{\tau}}_e + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} = \left(\frac{\partial \mathbf{y}}{\partial \mathbf{r}} - \frac{\partial \mathbf{i}}{\partial \mathbf{r}} \right) \times \hat{\mathbf{r}} \quad (\text{C.1})$$

$$\hat{\boldsymbol{\tau}}_e = \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \quad (\text{C.2})$$

The existence of a unique perturbation solution (see the discussion in Appendix A.2.1) in conjunction with Lemma 1 implies that this equation also has a unique bounded solution for $(\hat{\mathbf{r}}, \hat{\boldsymbol{\tau}}_e)$. Thus there exists a unique linear map \mathcal{H} such that

$$\begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\tau}}_e \end{pmatrix} = \mathcal{H} \times \begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \end{pmatrix} \quad (\text{C.3})$$

where \mathcal{H} is the left inverse of

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

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

$$\hat{\mathbf{c}}_\varepsilon = \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon}}_{\hat{\mathbf{c}}_\varepsilon^{PE}} + \underbrace{\left(\frac{\partial \mathbf{c}}{\partial \mathbf{r}} \quad \frac{\partial \mathbf{c}}{\partial \tau_e} \right) \times \mathcal{H} \times \begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \\ \frac{\partial \tau_e}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \end{pmatrix}}_{\equiv \mathbf{D} \times \left(\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} + \frac{\partial \mathbf{g}}{\partial \boldsymbol{\varepsilon}} \times \boldsymbol{\varepsilon} \right)} \quad (\text{C.4})$$

The definition of the “demand multiplier” map \mathbf{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.⁵⁴ General equilibrium feedback is thus identical, giving (5), and (6) follows. \square

C.3 Auxiliary Lemma for Proposition 2

Lemma C.1. *Consider the structural model of Section 2.1. A perfect foresight equilibrium is a sequence of nominal interest rates $\{i_t^b\}_{t \geq 0}$, aggregate output $\{y_t\}_{t \geq 0}$, wages $\{w_t\}_{t \geq 0}$ and the endogenous part of tax rebates $\{\tau_{et}\}_{t \geq 0}$ such that*

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

where the consumption, production, investment, labor supply and labor demand 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, and $x_t = (i_t^b, y_t, w_t, \tau_{et})$.

To prove Lemma C.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.

⁵⁴Of 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).

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,⁵⁵ 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 $\tau_e \in \mathbf{x}$, 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.
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 C.1. Necessity is immediate, completing the argument. \square

C.4 Proof of Proposition 2

By Lemma C.1, a perfect foresight equilibrium is, to first order, a solution to the system of linear equations

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

⁵⁵Note that the path of the intermediate goods price p^I is obtained from the problem of retailers.

where J_i denotes the infinite-dimensional generalization of the selection matrix selecting the i th entry of a vector x_t . Assuming equilibrium existence and uniqueness,⁵⁶ there exists a unique linear 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 unique. Now consider consumption demand 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 the arguments in the proof of Lemma C.1, we know that $\frac{\partial \mathbf{i}}{\partial \boldsymbol{\varepsilon}_d} = \frac{\partial \mathbf{y}}{\partial \boldsymbol{\varepsilon}_d} = \frac{\partial \boldsymbol{\ell}^f}{\partial \boldsymbol{\varepsilon}_d} = \mathbf{0}$, 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 $\frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}_g} = \mathbf{0}$. I now distinguish two cases.

- (i) Suppose that Assumption 3 holds. Then we can also conclude that $\frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}_d} = \mathbf{0}$. The two direct shock responses are then

$$\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d} \\ \mathbf{0} \\ \mathbf{0} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}_d} \end{pmatrix} \times \boldsymbol{\varepsilon}_d = \begin{pmatrix} \hat{\mathbf{c}}_d^{PE} \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{ed}^{PE} \end{pmatrix}$$

⁵⁶Existence 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).

and

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

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

$$\hat{\mathbf{c}}_d = \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d}_{\hat{\mathbf{c}}_d^{PE}} + \underbrace{\frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{c}}_d^{PE} \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{ed}^{PE} \end{pmatrix}}_{\mathbf{D} \times \hat{\mathbf{c}}_d^{PE}}$$

and

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

respectively, where \mathbf{D} is a common *demand multiplier*. This establishes that

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

and so (13) follows from simple addition.

(ii) Without Assumption 3, the two direct shock responses are

$$\begin{pmatrix} \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d} \\ \frac{\partial \boldsymbol{\ell}^h}{\partial \boldsymbol{\varepsilon}_d} \\ \mathbf{0} \\ \frac{\partial \boldsymbol{\tau}_e}{\partial \boldsymbol{\varepsilon}_d} \end{pmatrix} \times \boldsymbol{\varepsilon}^d = \begin{pmatrix} \hat{\mathbf{c}}_d^{PE} \\ \hat{\boldsymbol{\ell}}^h \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{ed}^{PE} \end{pmatrix}$$

and

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

The response paths of consumption now satisfy

$$\hat{\mathbf{c}}_d = \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d}_{\hat{\mathbf{c}}_d^{PE}} + \frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{c}}_d^{PE} \\ \hat{\boldsymbol{\ell}}_d^{PE} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{ed}^{PE} \end{pmatrix}$$

and

$$\hat{\mathbf{c}}_g = \mathbf{0} + \frac{\partial \mathbf{c}}{\partial \mathbf{x}} \times \mathcal{H} \times \begin{pmatrix} \hat{\mathbf{g}}_g \\ \mathbf{0} \\ \mathbf{0} \\ \hat{\boldsymbol{\tau}}_{eg}^{PE} \end{pmatrix}$$

Combining the two:

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

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 $\hat{\boldsymbol{\ell}}_d^{PE}$, as claimed.

□

C.5 Proof of Proposition 3

The proof proceeds in three steps. First, I show that *aggregate* impulse responses to the heterogeneous shocks $\{\varepsilon_{di0}\}$ are identical to impulse responses to the common aggregate shock $\varepsilon_{d0} \equiv \int_0^1 \varepsilon_{di0}$. Second, I prove that $\hat{\mathbf{c}}_{di} - \hat{\mathbf{c}}_d = (\xi_{di0} - 1) \times \hat{\mathbf{c}}_d^{PE} + \boldsymbol{\zeta}_i$, where $\int_0^1 (\xi_{di0} -$

1) $\zeta_i di = \mathbf{0}$. And third, I exploit standard properties of fixed-effects regression 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 $\boldsymbol{\varepsilon}_d \equiv e_1$, where $e_1 = (1, 0, 0, \dots)'$. The direct partial equilibrium response of consumption to the shock is

$$\hat{\mathbf{c}}_d^{PE} \equiv \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_d} \times \xi_{di0} \times \boldsymbol{\varepsilon}_d di$$

where $\mathbf{c}_i(\bullet)$ is the consumption function of individual i , defined analogously to the aggregate consumption function $\mathbf{c}(\bullet)$. Since $\int_0^1 \xi_{di0} di = 1$ and since ξ_{di0} is assigned randomly across households (and so does not correlate with $\frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d$ at any t), we have that

$$\hat{\mathbf{c}}_d^{PE} = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_{d_0}} \times \boldsymbol{\varepsilon}_d di \times \left[1 + \int_0^1 (\xi_{di0} - 1) di \right] = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d di$$

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 $\boldsymbol{\varepsilon}_d$. The same argument applies to the desired partial equilibrium contraction in labor supply, $\hat{\boldsymbol{\ell}}_d^{PE}$. 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

$$\hat{\mathbf{c}}_{id} = \frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} \times \hat{\mathbf{x}} + \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_d} \times \xi_{di0} \times \boldsymbol{\varepsilon}_d$$

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

$$\hat{\mathbf{c}}_{id} - \hat{\mathbf{c}}_d = (\xi_{di0} - 1) \times \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d}_{\hat{\mathbf{c}}_d^{PE}} + \underbrace{\left(\frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} - \frac{\partial \mathbf{c}}{\partial \mathbf{x}} \right) \times \hat{\mathbf{x}} + \xi_{di0} \left(\frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_d} - \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d} \right) \times \boldsymbol{\varepsilon}_d}_{\equiv \zeta_i}$$

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 $\int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_d} di = \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_d}$, the residual term ζ_i must satisfy $\int_0^1 (\xi_{di0} - 1) \zeta_i di = 0$.

3. By the standard properties of fixed-effects regression, we can re-write regression (14) as

$$\hat{c}_{it+h} - \hat{c}_{t+h} = \beta_{dh} \times (\xi_{it} - 1) \varepsilon_{dt} + u_{it+h} - u_{t+h} \quad (\text{C.5})$$

By standard projection results, the estimand β_d satisfies

$$\begin{aligned}\beta_d &= \frac{\int_0^1 [(\xi_{di0} - 1)\hat{\mathbf{c}}_d^{PE} + \zeta_i] (\xi_{di0} - 1) di}{\int_0^1 (\xi_{di0} - 1)^2 di} \\ &= \hat{\mathbf{c}}_d^{PE}\end{aligned}$$

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

□

C.6 Proof of Proposition 4

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

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

$\hat{\mathbf{y}}_{gh}$ is thus the estimand of a local projection on ε_{gt} . (20) then follows immediately by Corollary 1 in Plagborg-Møller & Wolf (2019).⁵⁷ □

C.7 Auxiliary Lemma for Proposition 5

Lemma C.2. *Consider the structural model of Section 2.1. Under Assumptions 5 to 8, 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*

$$\mathbf{c} + \mathbf{i}(\mathbf{s}^f(\mathbf{c}); \boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) = \mathbf{y}(\mathbf{s}^f(\mathbf{c}); \boldsymbol{\varepsilon}) \quad (\text{C.6})$$

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

To prove Lemma C.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.

⁵⁷Strictly 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 .

1. By Assumption 6, the household block admits aggregation to a single representative household with period felicity function $u(c, c_{-1}, \ell)$. 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 8) 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 7, 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 (C.6) holds at all times, so the household lifetime budget constraint must hold.

Together, 1. - 2. establish sufficiency of the conditions in the statement of Lemma C.2. Necessity is immediate, completing the argument. \square

C.8 Proof of Proposition 5

By Lemma C.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

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

and

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

respectively. This establishes (21). The equations for output are exactly analogous. \square

C.9 Proof of Corollary D.1

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

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

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. \square

D Additional results

This appendix collects supplemental results. I discuss (i) robustness checks for accuracy of the demand equivalence approximation, (ii) a generalization of my methodology to correct for wealth effects in labor supply, (iii) a model variant with strong general equilibrium crowding-in, (iv) a re-parameterization of the estimated HANK model that more closely matches my empirically estimated output and consumption impulse responses to fiscal shocks, and (v) a generalization of the equivalence result away from the model’s deterministic steady state.

D.1 Approximation accuracy

This section provides details for the extensions and robustness checks referenced in Sections 2, 4 and 5. I discuss (i) a model variant without unions and with weak wealth effects in labor supply, (ii) a two-asset HANK model, (iii) inaccuracies for persistent demand shocks, (iv) a model with durables, (v) useful (valued or productive) government spending, (vi) multi-good economies, (vii) random draws for all parameters governing dynamics in the estimated HANK model, (viii) approximate equivalence under imperfect matching of private and public excess demand paths and (ix) approximate investment demand equivalence.

D.1.1 Weak wealth effects in labor supply

Empirical evidence suggests weak – but non-zero – short-term wealth effects associated with (small) unexpected income gains (Cesarini et al., 2017; Fagereng et al., 2018). My benchmark structural model – which features preferences with strong wealth effects, but sticky-wage unions – cannot directly speak to these weak short-term wealth effects, as micro-level difference-in-differences regressions invariably difference out the effects of direct labor adjustments (recall the proof of Proposition 3). In this section I thus instead consider a model with non-standard household preferences and without sticky-wage unions. Importantly, the model is designed to be consistent with the empirically documented weakness and cross-sectional homogeneity in short-run wealth effects in labor supply.

MODEL. I consider the estimated HANK model of Section 4, but with three changes. First, the model 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

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

where the preference shifter θ_{ift} satisfies⁵⁸

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

The variable x_{ft} is central. To jointly ensure arbitrarily weak short-run wealth effects in labor supply, homogeneous wealth effects in the cross section of households, and 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 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

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

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.⁵⁹ In particular, let $\ell^h = \ell^h(\mathbf{w}, \mathbf{c})$ denote the mapping from wages and family consumption into family labor supply induced by (D.1). The micro regression estimand in

⁵⁸Households do not internalize the effect of their consumption on the shifter.

⁵⁹I 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.

(14) then satisfies

$$\hat{\mathbf{c}}_{\tau}^{PE} = \left(I - \frac{\partial \mathbf{c}}{\partial \ell} \times \frac{\partial \ell^h}{\partial \mathbf{c}} \right)^{-1} \times \left(\frac{\partial \mathbf{c}}{\partial \tau} \cdot d\tau \right) \quad (\text{D.2})$$

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

PARAMETERIZATION. All parameters related to the sticky-wage block of the model are now irrelevant; the only new model parameter is ω . To ensure consistency with empirical evidence, I set $\omega = 0.05$. As in Cesarini et al. (2017), this specification results in a peak partial equilibrium labor supply response of around 4\$ for every 100\$ response in consumption.

RESULTS. Results are displayed in Figure D.1.

APPROXIMATE DEMAND EQUIVALENCE, WEAK WEALTH EFFECTS

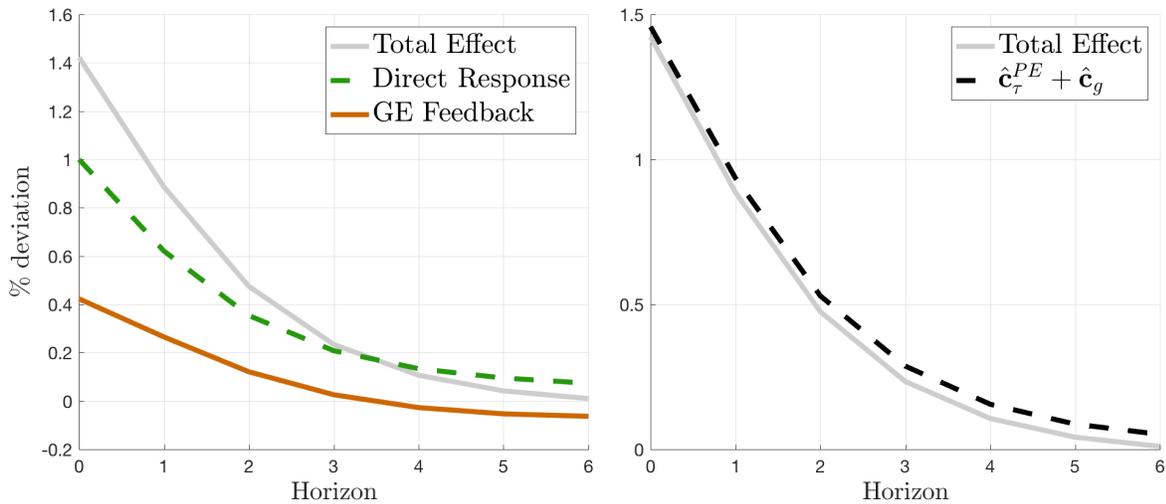


Figure D.1: Impulse response decompositions and demand equivalence approximation for the HANK model with weak wealth effects. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Estimated wealth effects are weak, so Assumption 3 is nearly satisfied, and the approximation is again highly accurate, with a maximal error of around 4 per cent. In addition to general equilibrium feedback associated with the labor supply contraction itself, the slight over-statement displayed in the right panel now also reflects a second, more subtle effect: Since the government spending expansion $\hat{\mathbf{g}}_{\varepsilon_g}$ only replicates the direct consumption response *net* of earnings changes, its present value is lower than that of the corresponding income tax

rebate, and so the associated tax burden is lower. Matching $\frac{\partial c}{\partial \tau} \cdot d\tau$ (rather than the true partial equilibrium response displayed in (D.2)) would instead make the approximation as accurate as in the benchmark model.

D.1.2 A two-asset HANK model

This section provides further details on the two-asset HANK model of Section 4.3. I sketch the expanded model (in particular the new consumption-savings problem), discuss the model parameterization, and finally display and interpret the results of my accuracy check.

MODEL SKETCH. Households invest in an illiquid asset with nominal return i^h and a liquid asset with return $i^h - \kappa_{1,b} + \mathbf{1}_{b^h < 0} \kappa_{2,b}$. The household consumption-savings problem then is

$$\max_{\{c_{it}, b_{it}^h, a_{it}^h\}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_i^t \zeta_t(\boldsymbol{\varepsilon}_v) u(c_{it}, c_{it-1}, \ell_{it}) \right]$$

such that

$$c_{it} + b_{it}^h + a_{it}^h = (1 - \tau_\ell) w_t e_{it} \ell_{it} + \frac{1 + i_{t-1}^h - \kappa_{1,b} + \mathbf{1}_{b_{it-1}^h < 0} \kappa_{2,b}}{1 + \pi_t} b_{it-1}^h + \frac{1 + i_{t-1}^h}{1 + \pi_t} a_{it-1}^h + \phi_a(a_{it}^h, a_{it-1}^h) + \tau_{it}$$

and

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

where $\phi_a(\bullet, \bullet)$ is the adjustment cost function for illiquid asset holdings. Similar to Kaplan et al. (2018), I assume that

$$\phi_a(a', a) = \frac{\chi_1}{\chi_2} \times \left(\frac{|a' - \frac{1+i^h}{1+\pi} a|}{\chi_0 + \frac{1+i^h}{1+\pi} a} \right)^{\chi_2} \times \left(\chi_0 + \frac{a' - \frac{1+i^h}{1+\pi} a}{a} \right)$$

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 $\omega_t \equiv b_t^h + a_t^h$ denote total funds managed by the mutual fund. Returns earned by the mutual fund i_t^m then satisfy

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

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 $\kappa_{1,b}$ to make assets liquid and $\kappa_{1,b} + \kappa_{2,b}$ to lend liquid assets. 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 the equilibrium 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 1-asset HANK model, and only re-calibrate the steady state of the model. Table 3 displays all parameters from the re-calibrated 2-asset model that are different from those displayed in Table 1 for the benchmark 1-asset model.

STEADY-STATE PARAMETER VALUES, 2-ASSET HANK MODEL

Parameter	Description	Value	Target	Model	Data
<i>Households</i>					
χ_0	Adjustment Cost Parameter	0.25	Kaplan et al. (2018)		
χ_1	Adjustment Cost Parameter	6.13	Fraction $b = 0$	0.29	0.30
χ_2	Adjustment Cost Parameter	2.03	A/Y	11.21	10.64
β	Discount Rate	0.98	B/Y	1.29	1.04
r^h	Return	0.0125	Kaplan et al. (2018)		
$\kappa_{1,b}$	Liquid Wedge	0.01	Upper Bound		
$\kappa_{2,b}$	Borrowing Wedge	0.03	Fraction $b < 0$	0.09	0.15
\underline{b}	Borrowing Constraint	-1	Kaplan et al. (2018)		
<i>Firms</i>					
δ	Depreciation	0.025	Firm Valuation		

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

I choose the parameters of the adjustment cost function to ensure a reasonable fit to the

liquid-illiquid wealth distribution in U.S. data (Kaplan et al., 2018). To provide a stringent test of the demand equivalence approximation, I set the wedge between returns on household deposits and government debt to be an (arguably implausible) 1 per cent per quarter. With smaller return gaps, the approximation would obviously improve.

As is typical for two-asset models, the average household MPC is lower than in a liquid-wealth calibration of a one-asset model, now at around 14 per cent. Intuitively, this is so because households have more vehicles to self-insure.

APPROXIMATION ACCURACY. Results are displayed in Figure D.2. Two features stand out. First, the model now features stronger general equilibrium crowding-out. Relative to the simpler one-asset HANK model, this model features (i) smaller average MPCs and (ii) no mechanical redistribution effects related to heterogeneous dividend exposure.⁶⁰ Both changes tend to dampen general equilibrium amplification. Second, even though both wealth effects in labor and (implausibly large) heterogeneity in borrowing and lending rates lead the demand equivalence to over-state the response of aggregate consumption, the approximation remains reasonable, with a maximal error around 7 per cent of the peak consumption response.

APPROXIMATE DEMAND EQUIVALENCE, 2-ASSET HANK MODEL

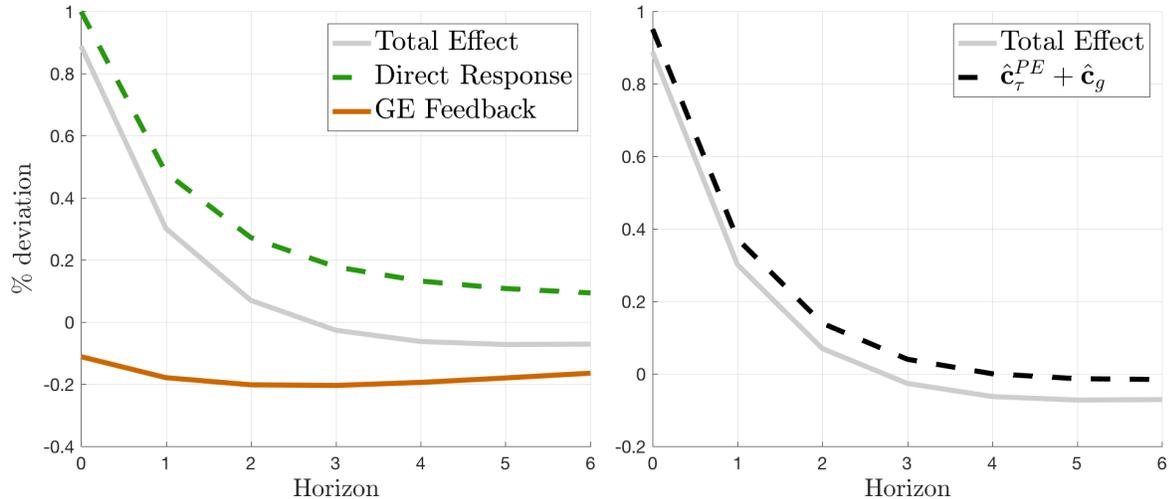


Figure D.2: Impulse response decompositions and demand equivalence approximation for the two-asset HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

⁶⁰All returns are received by the mutual fund and passed on to households, whereas before households were directly exposed to (mildly countercyclical) dividend payments.

D.1.3 Persistent demand shocks

Figure D.3 shows that, even in the rigid-wage model of Justiniano et al. (2010), the demand equivalence approximation deteriorates for very persistent demand shocks. This is not surprising: Persistent shocks induce – through pronounced long-term wealth effects – a persistent decline in desired labor supply. Since wages are not sticky forever, the decline in desired labor supply ultimately feeds into a decline in actual hours worked, so the approximation error is larger at long horizons.

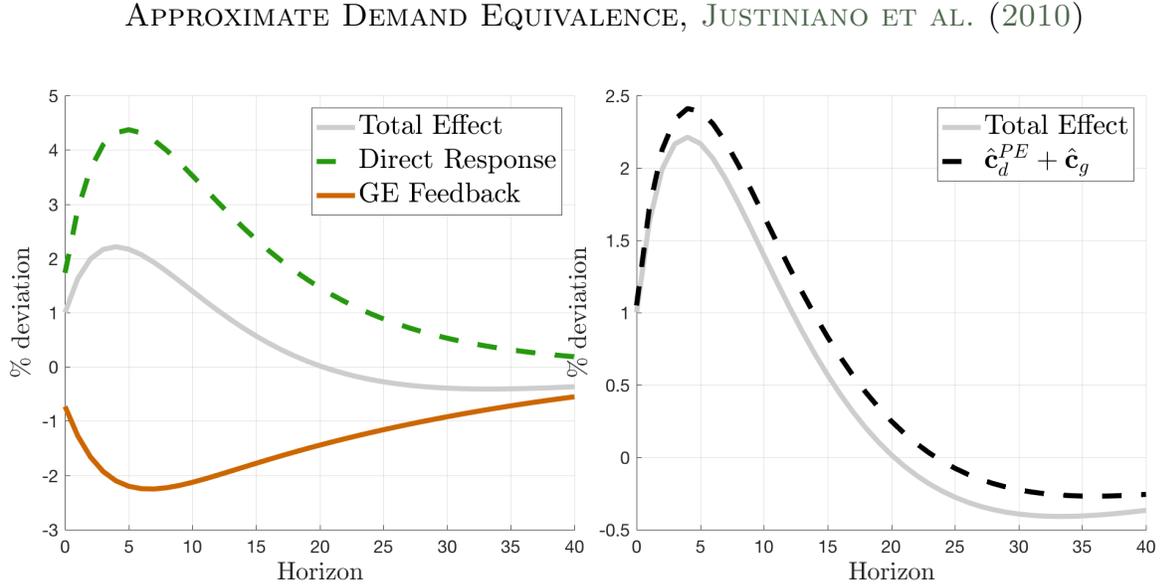


Figure D.3: Impulse response decompositions and demand equivalence approximation for persistent and hump-shaped demand shocks in the model of Justiniano et al. (2010), solved at the posterior mode and with a fraction $\lambda \rightarrow 0$ of spenders. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

D.1.4 Durables

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

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

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

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

where e_t is aggregate household expenditure. The equilibrium definition in Appendix A.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, the demand equivalence result then still applies, now for the aggregated household expenditure path \mathbf{e} :

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

$$\hat{\mathbf{e}}_d = \underbrace{\hat{\mathbf{e}}_d^{PE}}_{PE \text{ response}} + \underbrace{\hat{\mathbf{e}}_g}_{= GE \text{ feedback}} \quad (\text{D.4})$$

D.1.5 Useful government spending

In the benchmark model of Section 2.1 as well as the estimated HANK model in Section 4, government spending is useless – it is neither valued by households, nor does it have any productive benefits. Some previous work has instead allowed for such benefits of government spending. To gauge the extent to which such benefits threaten my approximations, I in

this section review related empirical evidence and provide model-based intuition as well as quantitative accuracy checks.

VALUED SPENDING. It is immediate that all equivalence results go through unchanged if government expenditure enters household utility in an additively separable fashion, i.e. the per-period felicity function can be represented as

$$\tilde{u}(c_{it}, c_{it-1}, \ell_{it}, g_t) = u(c_{it}, c_{it-1}, \ell_{it}) + v(g_t)$$

More interestingly, exact equivalence also holds with particular kinds of (popular) non-separable preferences. For example, suppose that preferences take the form

$$u(c_{it}, \ell_{it}, g_t) = \log(c_{it}^\nu g_t^{1-\nu}) - \chi \frac{\ell_{it}^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}}$$

Log preferences are popular in the business-cycle literature (and used in my own HANK model), while an inner Cobb-Douglas aggregator is popular in the trade literature (Fajgelbaum et al., 2018). It is straightforward to see that the marginal utility of consumption is then unaffected by changes in government expenditure, so consumption decisions are again unaffected, and demand equivalence survives under the same assumptions as before.

With other types of non-separabilities, exact equivalence does not survive. For example, Leeper et al. (2017) assume that households have conventional preferences over a synthetic consumption aggregate c_{it}^* , where

$$c_{it}^* = c_{it} + \alpha_G g_t$$

There is little direct empirical evidence on the magnitude or even sign of α_G . Since private and public consumption co-move in post-war aggregate data, standard likelihood-based estimation exercises with representative-agent models usually call for a negative coefficient (Leeper et al., 2017). Models with high average household MPCs instead endogenously tie private and public spending together, and so would likely require little or no complementarity (Galí et al., 2007). Nevertheless, and to threaten the accuracy of the demand equivalence approximation as much as possible, I consider a variant of my benchmark model with $\alpha_G = -0.24$, exactly as in Leeper et al. (2017). Results are displayed in Figure D.4.

Since private and public consumption are complements in household preferences, private consumption is directly stimulated by an increase in public spending, and so the demand equivalence approximation over-states (black line). However, given an estimate of α_G , it

APPROXIMATE DEMAND EQUIVALENCE, VALUED G

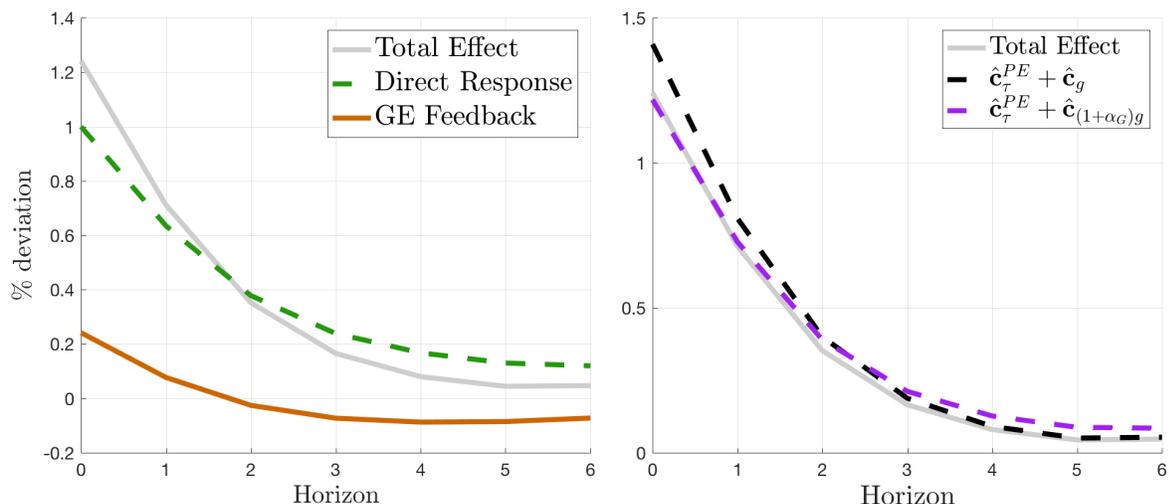


Figure D.4: Impulse response decompositions and demand equivalence approximation in the estimated HANK model, augmented to feature complementarities in private and public consumption, as in Leeper et al. (2017). The direct response and the indirect general equilibrium feedback are computed following Definition 1.

is straightforward to correct for this inaccuracy: Ignoring for simplicity the presence of potentially binding borrowing constraints, the wedge in household Euler equations associated with additional government spending is equal to $\alpha_G \times (g_t - g_{t+1})$. Thus, for every additional dollar of government spending, private partial equilibrium demand increases by $-\alpha_G$ dollars (ignoring the permanent level shift due to unchanged lifetime income). This suggests a simple fix: Instead of matching \hat{c}_τ^{PE} and \hat{g}_g , researchers should match \hat{c}_τ^{PE} and a scaled spending response $(1 + \alpha_G) \times \hat{g}_g$. The third purple line in Figure D.4 shows that, with this alternative approach, the demand equivalence approximation is again nearly exact.

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. Reassuringly, my empirical estimates are almost identical for overall and for pure military spending forecast errors, suggesting that my analysis is not picking up the effect of public investment spending.

For completeness, I also illustrate this conclusion through a fully structural analysis. In the models of Boehm (2016) and Leeper et al. (2010), government expenditure on investment

goods is productive in the sense that the aggregate stock of government “capital” k_t^g directly affects the production capabilities of intermediate goods producers:

$$y_{jt} = y(e_{jt}, k_t^g, u_{jt}k_{jt-1}, \ell_{jt})$$

where $k_t^g = (1 - \delta)k_{t-1}^g + g_t$. Analogously, I consider a variant of my estimated HANK model with a production function of the form

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

I set the output elasticity to $\alpha_g = 0.2$, large enough to generate substantial asymmetry in multipliers (and larger than usual estimates in this literature, e.g. $\alpha_g = 0.05$ in Leeper et al. (2010)). Results are displayed in Figure D.5.

APPROXIMATE DEMAND EQUIVALENCE, PRODUCTIVE G

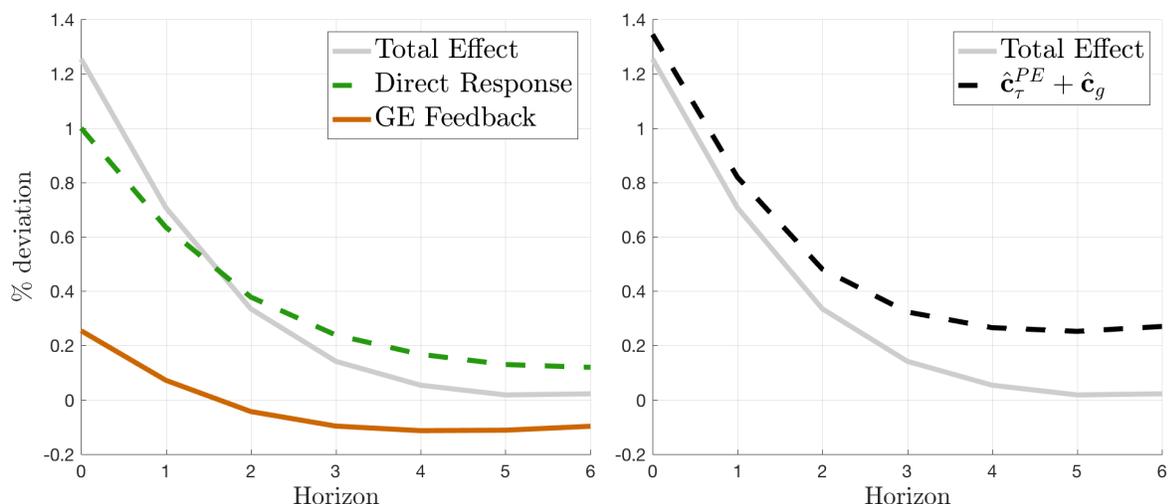


Figure D.5: Impulse response decompositions and demand equivalence approximation in the estimated HANK model, augmented to feature productive benefits of public spending, as in Leeper et al. (2010). The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Unsurprisingly, the approximation is reasonably accurate on impact (where the demand pressure of the shock dominates), but deteriorates over time, as higher government spending gradually expands the productive capacity of the economy. These results are entirely consistent with Leduc & Wilson (2013), who empirically document “an initial effect due nominal rigidities and a subsequent medium-term productivity effect.”

D.1.6 Multi-good economies

Heterogeneity in consumption baskets for private and public consumption can break the demand equivalence result, at least as long as factors of production are imperfectly mobile across sectors or production functions are sector-specific. In such a segmented economy, relative prices will respond to spending shocks (Ramey & Shapiro, 1998), and so the demand equivalence approximation will fail. Previous work has also emphasized that heterogeneity in the factor incidence of private and public demand shocks may matter for aggregate transmission (Alonso, 2017; Baqaee, 2015), and that the effects of government spending on consumption goods may differ from those on investment goods (Boehm, 2016).

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 a series of structural models, rich enough to allow for the mechanisms reviewed above and disciplined to be consistent with empirical evidence on their likely strength.

ENCOMPASSING MODEL. I consider a generalized variant of my benchmark model of Section 2. The model deviates from this benchmark framework 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 aggregate spending multipliers for each.

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; under my choice of household preferences, they take the form

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

for $m = 1, 2, I$, and where as before c_t^* is the virtual consumption aggregate defined in (A.22). Note that, with $\mu = 0$, 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} .

For all subsequent results, I build on the parameterization of the estimated HANK model of Section 4, but with one notable difference: I materially lower the 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.⁶¹ To be conservative, I choose a model calibration with $\phi_p = 0.6$, giving relative price responses consistent with this evidence.

In the remainder of this section I study the quality of the demand equivalence approximation for different types of government purchases. In particular, I consider two model variants: (i) a two-good model with a consumption good and an investment good, similar to Boehm (2016), and (ii) a model with two consumption goods and a separate investment good, and with heterogeneous factor shares. Both models implicitly allow for the relative price effects emphasized in Ramey & Shapiro (1998).

⁶¹For 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.

INVESTMENT GOODS. To study the effects of multiplier heterogeneity for investment and consumption goods, I consider a special case of the above economy with perfect factor mobility across the two consumption sectors (allowing aggregation to a single composite consumption good), but imperfect capital and labor mobility across the composite consumption and investment good sectors ($\mu = 1$). The government consumes both goods, and the overall size of the composite investment sector is calibrated to correspond to 20% of aggregate output in steady state. In keeping with Boehm (2016), I assume homogeneous production technologies across the two sectors, i.e. $\alpha_1 = \alpha_2 = \alpha_I \equiv \alpha$.

Results are displayed in Figure D.6. The right panel shows that the approximation is still accurate for government purchases of consumption goods. Perhaps more surprisingly, it remains reasonably accurate for government purchases of *investment* goods.⁶²

APPROXIMATE DEMAND EQUIVALENCE, INVESTMENT V CONSUMPTION

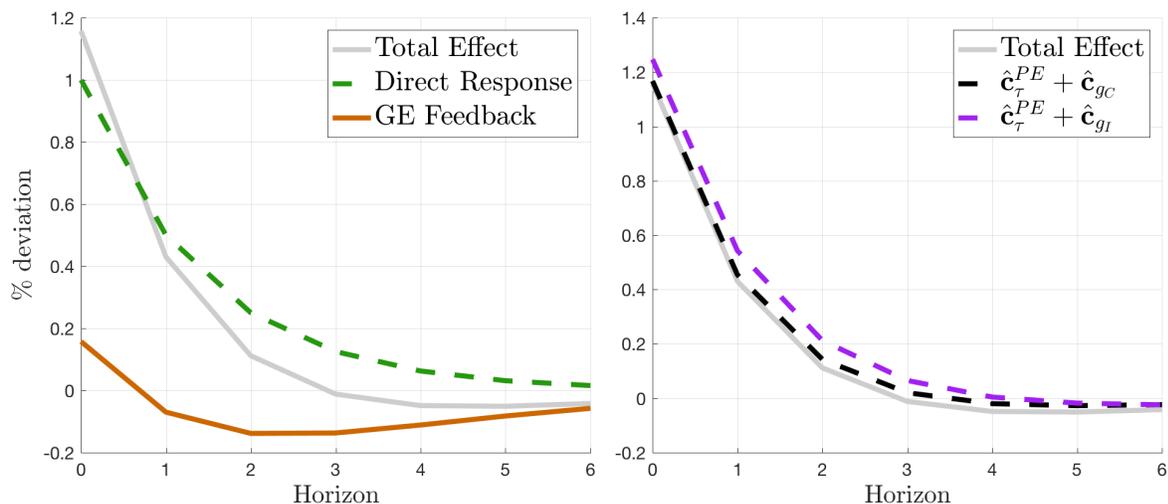


Figure D.6: Impulse response decompositions and demand equivalence approximation in the two-sector HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

My results are inconsistent with Boehm (2016). The differences between the two analyses can be traced back to three model features. First, Boehm’s model features flexible wages, while in my model wages are sticky. Wage flexibility turns out to be crucial to his mechanism: When the government buys the consumption good, its relative price rises, consumption is crowded out, and households work harder. When it buys the investment good, in contrast,

⁶²Correspondingly, I also find that the output responses for both types of government spending are almost identical.

relative prices barely change (because of the high intertemporal elasticity of investment demand), consumption is not crowded out, and labor supply does not respond. Nominal wage rigidity breaks this mechanism. Second, my shocks are much more transitory than his. As a result, in the presence of short-lived wage rigidity, the labor supply channel is particularly dampened. Third, his model features an extremely large intertemporal price elasticity of firm investment demand. Consistent with both my macro estimation as well as the micro evidence of Zwick & Mahon (2017), my model features much stronger adjustment costs, and so investment is not as easily crowded out. This increases the strength of aggregate demand effects after government purchases of investment goods.

CONSUMPTION BASKET HETEROGENEITY. I analyze relative price effects and heterogeneous factor incidence using a full three-sector version of my extended economy. I set $\mu = 1$, $\nu = 0.5$, labor shares for each of the three production sectors are exactly as in Alonso (2017, Table 3), and 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)). As before, I restrict factors to be imperfectly mobile across the three sectors. Results are displayed in Figure D.7.

APPROXIMATE DEMAND EQUIVALENCE, HETEROGENEOUS FACTOR INCIDENCE

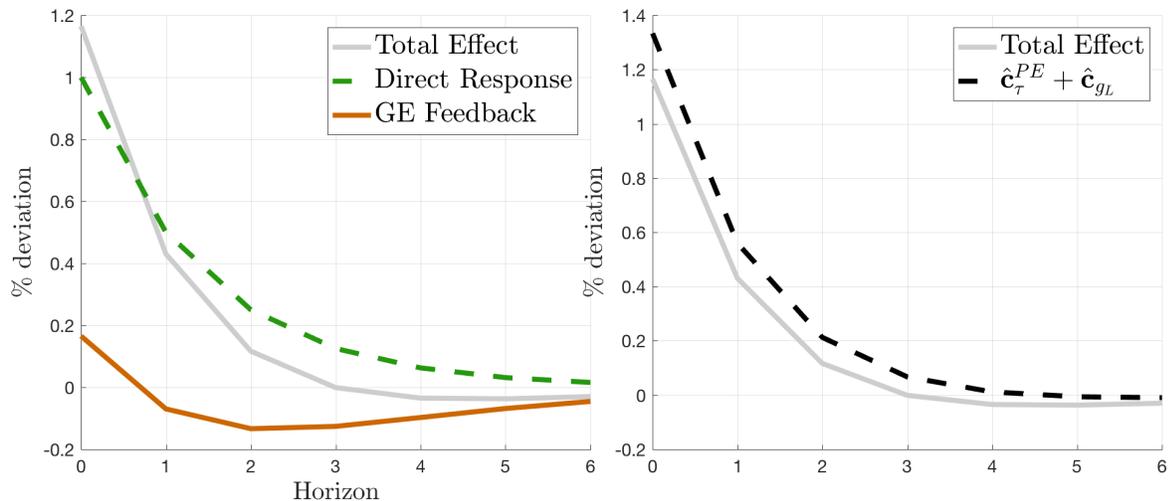


Figure D.7: Impulse response decompositions and demand equivalence approximation in the three-sector HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

In the data, the network-adjusted labor share of the average government consumption good exceeds that of the average consumption good. I thus in the right panel of Figure D.7

show the demand equivalence approximation for government purchases of the second (labor-intensive) consumption good. 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 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.⁶³

Finally, for completeness, I have re-computed the demand equivalence approximations in a model with imperfect factor mobility, but homogeneous production functions, leaving only the inaccuracies associated with relative price movements. I find that the demand equivalence approximation is then almost as accurate as in my benchmark model, suggesting that almost all of the inaccuracy in Figure D.7 is driven by the factor incidence mechanism.

D.1.7 Random parameter draws

The accuracy displayed in Figure 6 is not at all special to the mode parameterization of the estimated HANK 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 fixing the dynamics parameter values as in Table 2, I *randomly draw* their values from uninformative uniform distributions over wide supports, as displayed in Table 4.⁶⁴ 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 4.

I find that the approximation accuracy is largely orthogonal to all parameters except for the price stickiness ϕ_p . Fixing ϕ_p at the posterior mode and merely randomly drawing all other parameters, I find that 95 per cent of draws give a maximal prediction error below 3.3

⁶³Arguably, this is an upper bound for the likely size of the effect, since heterogeneity in MPCs by skill 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 D.7 shrinks.

⁶⁴This approach to documenting a generic property of a family of quantitative models closely follows the analysis in Canova & Paustian (2011).

DYNAMICS PARAMETER VALUES, RANDOM DRAWS

Parameter	Description	Lower Bound	Upper Bound
ϕ_p	Price Calvo Parameter	0.1	0.99
ζ	Capacity Utilization	0.1	10
κ	Investment Adjustment Cost	0.1	10
ρ_m	Taylor Rule Persistence	0.01	0.99
ϕ_π	Taylor Rule Inflation	1	5
ϕ_y	Taylor Rule Output	0	1
ϕ_{dy}	Taylor Rule Output Growth	0	1

Table 4: Supports for uniform parameter draws in the HANK model.

per cent. For $\phi_p = 0.1$, and fixing all other parameters at the posterior mode, the prediction error increases to almost 9 per cent; intuitively, output is now not demand-determined, but given my calibrated moderate degree of wage rigidity, shifts in desired household labor supply still have rather limited aggregate effects.⁶⁵

D.1.8 Imperfect demand matching

The excess demand paths in Figure 4 and Figure 8 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

$$\hat{g}_{gt} = (1 + \nu_t) \times \hat{c}_{\tau t}^{PE} \tag{D.5}$$

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 4; 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 8.7 per cent, so

⁶⁵Conversely, with rigid prices and flexible wages, labor is still demand-determined, so labor supply shifts only move wage relative to dividend income. In my model, these incidence effects turn out to be relatively small, so the demand equivalence approximation is also accurate with flexible wages and rigid prices.

the approximation remains accurate.⁶⁶ The intuition is simple: Since the model only features 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.⁶⁷ To illustrate, Figure D.8 shows the quality of the demand equivalence approximation for one particular draw of the error sequence ν .

APPROXIMATE DEMAND EQUIVALENCE, IMPERFECT MATCHING

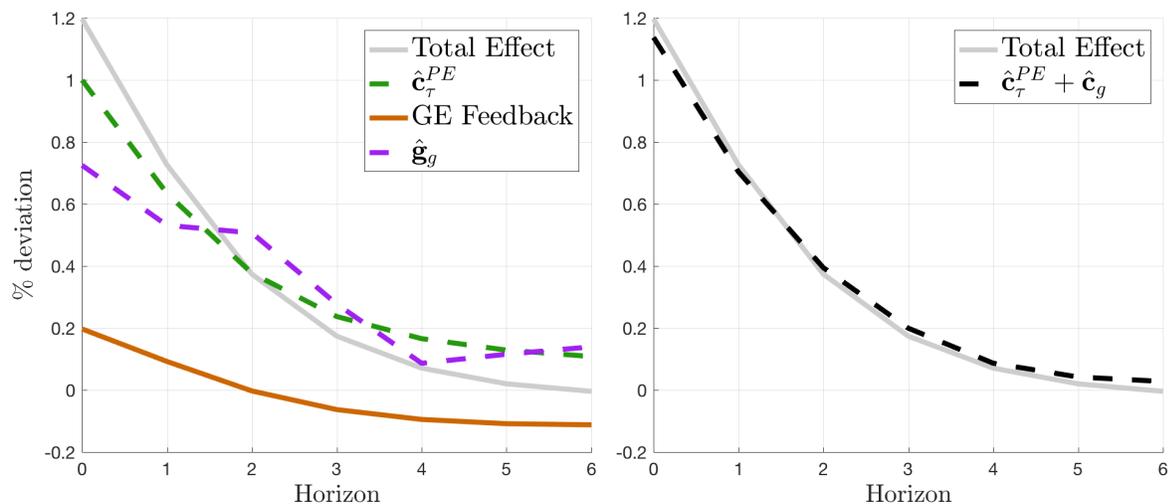


Figure D.8: Impulse response decompositions and demand equivalence approximation in the estimated HANK model, with imperfect demand matching, following (D.5). The direct response and the indirect general equilibrium feedback are computed following Definition 1.

D.1.9 Approximate investment demand equivalence

Consider again the estimated HANK model of Section 4. In keeping with my empirical investment application, I enrich the model to feature an investment tax credit shock that yields the same *partial equilibrium* investment response as that estimated in Section 5.2. I then construct an investment demand equivalence approximation for the aggregate investment response via the decomposition in (21). Results are displayed in Figure D.9.

⁶⁶Most of the large approximation errors come from draws in which the ν 's are so far from 0 that demand matching is clearly violated, so the results displayed here are actually an upper bound on likely inaccuracies.

⁶⁷Consistent with this intuition, the average error is even smaller when I repeat the same exercise in the model of Appendix D.4, which features general equilibrium amplification matched to my empirical estimates on government spending shock transmission.

APPROXIMATE INVESTMENT DEMAND EQUIVALENCE, HANK MODEL

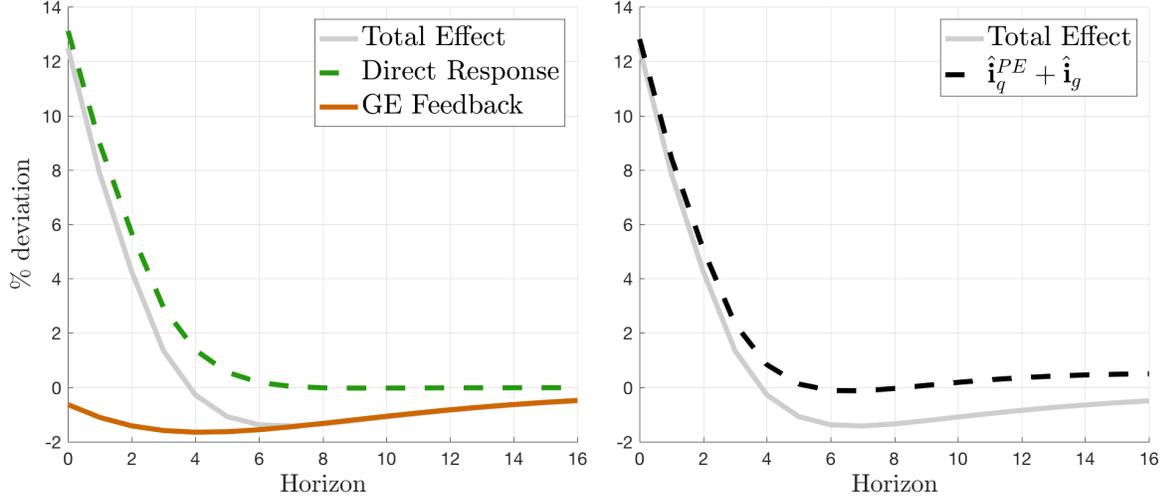


Figure D.9: Impulse response decompositions and investment demand equivalence approximation in the estimated HANK model, with details on the parameterization in Appendix A.2.2. The investment tax credit path is matched to replicate the direct investment response estimated in Section 5.2. The direct response and the indirect general equilibrium feedback are computed following the natural generalization of Definition 1.

The demand equivalence approximation remains reasonably accurate (in particular at short horizons), with the maximal error over all horizons equal to around 10 per cent of the true impact impulse response of investment. Intuitively, the equivalence approximation over-states the actual investment response since each of Assumptions 6 to 8 tends to weaken the extent of general equilibrium crowding-out. First, without Assumption 6, short-term increases in labor earnings (and cuts in dividend pay-outs) lead to excess consumption demand pressure. Second, without Assumption 7, wages need to increase more sharply to induce additional labor supply. And third, without Assumption 8, the monetary authority more aggressively leans against the investment demand stimulus. Nevertheless, all three effects are small, and so the additional degree of general equilibrium crowding-out in the estimated model is relatively modest. The investment demand equivalence approximation thus promises to be informative even if the underlying assumptions do not hold exactly.

D.2 Correcting for wealth effects in labor supply

Instead of ignoring the labor supply error term in Proposition 2, a simple alternative is to first estimate the direct partial equilibrium labor supply response $\hat{\ell}_d^{PE}$ from micro data and

then estimate the aggregate consumption effects of an equivalent household “leisure” shock.

GENERALIZED METHODOLOGY. Let $\hat{\mathbf{c}}_\psi$ denote the impulse response of aggregate consumption to a leisure shock – a labor wedge ε_ψ that changes desired household labor supply by $\hat{\ell}_d^{PE}$. Using the equilibrium construction of Lemma C.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 2 (but 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 $\hat{\ell}_d^{PE}$, 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 $\hat{\ell}_d^{PE}$ thus equals around 0.06% on impact, and little thereafter.

With a unit Frisch elasticity, 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 5 – a hardly relevant adjustment.

D.3 General equilibrium amplification

The equivalence result in Proposition 2 asserts that general equilibrium effects are tied together across shocks, but is silent on the strength of this common general equilibrium feedback. 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 rich benchmark model, 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 $\hat{\tau}$, consider an interest rate path \hat{r} such that, at $(\hat{\tau}, \hat{r})$ and facing steady-state wages forever, households are willing to consume steady-state consumption \bar{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 partial equilibrium consumption demand.

The second example is quantitative. I consider the benchmark estimated HANK model of Section 4, but set the household borrowing wedge to zero and further assume that preferences are as in Greenwood et al. (1988).

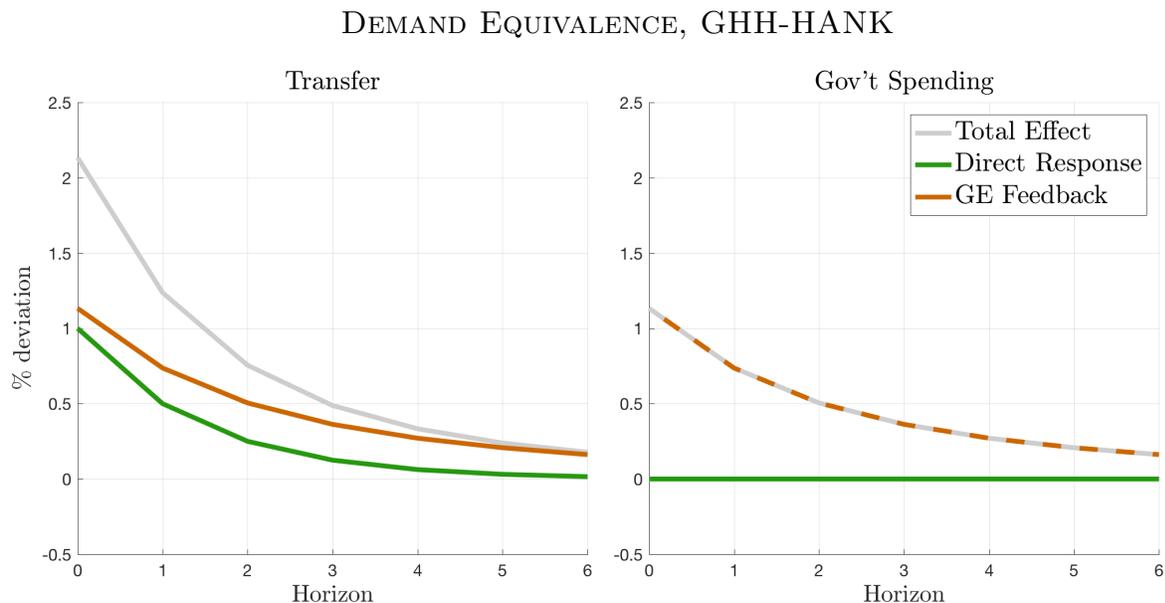


Figure D.10: Impulse response decompositions after equally large, one-off tax rebate and government spending shocks in the HANK model with GHH preferences. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

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 D.10.

D.4 Impulse response matching

The estimated model of Section 4 predicts some general equilibrium crowding-in of consumption following a transitory tax rebate. Minor parameter changes are however enough to ensure close agreement between model-implied and empirically estimated government spending impulse responses. In particular, it is enough to slightly lower the degree of nominal price rigidity (to $\phi_p = 0.65$) and to make monetary policy somewhat more aggressive (increasing the output response to $\phi_y = \phi_{dy} = 0.15$). Figure D.11 provides an illustration.

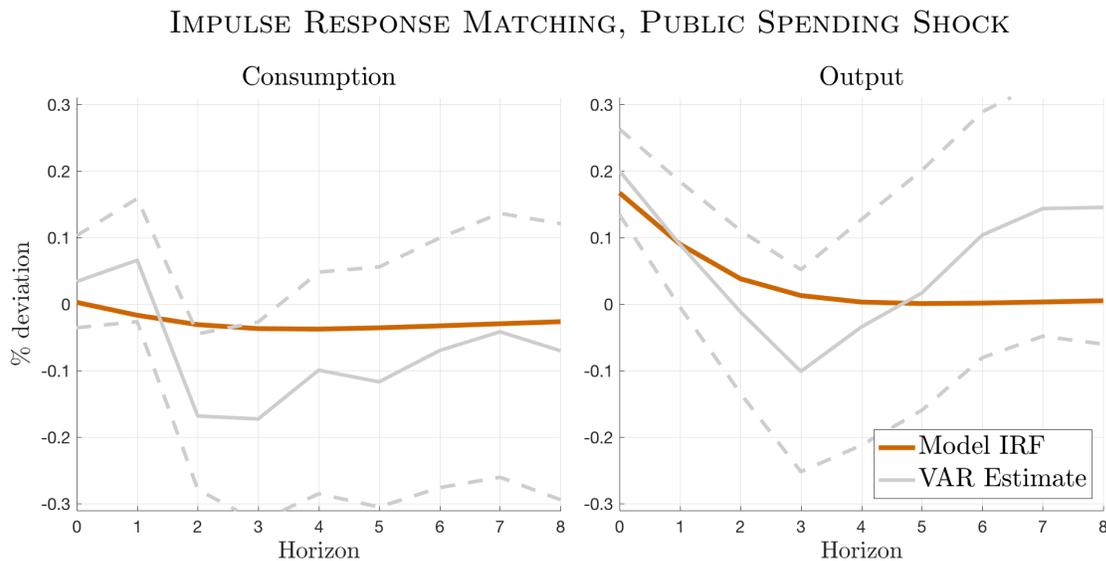


Figure D.11: Impulse responses to a transitory expansion in government spending. Empirical estimates (grey) exactly as in Section 3.2. Model-implied impulse responses (orange) in estimated HANK model, but with $\phi_p = 0.65$ and $\phi_y = \phi_{dy} = 0.15$. The government spending path exactly matches the model-implied household spending response to a transitory rebate, as in Figure 4.

The two panels show the consumption and output responses to a transitory increase in government spending. In both plots, the grey lines are the empirical estimates of Section 3.2, and the orange lines are model-implied analogues. Clearly, with the proposed parameter changes, the impulse responses align. The underlying model is thus a promising laboratory for the structural analysis of income tax rebates or other consumption demand shifters.

D.5 Demand equivalence along transition paths

All equivalence results presented in this paper were stated for transition paths starting at the deterministic steady state. However, it is immediate from the proofs of Propositions 2 and 5 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.

E Application: income redistribution

Both applications in the main text are semi-structural: 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 economy. Given this partial equilibrium block, my methodology can again be used to provide the mapping into full general equilibrium counterfactuals.

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 iMPCs 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 E.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

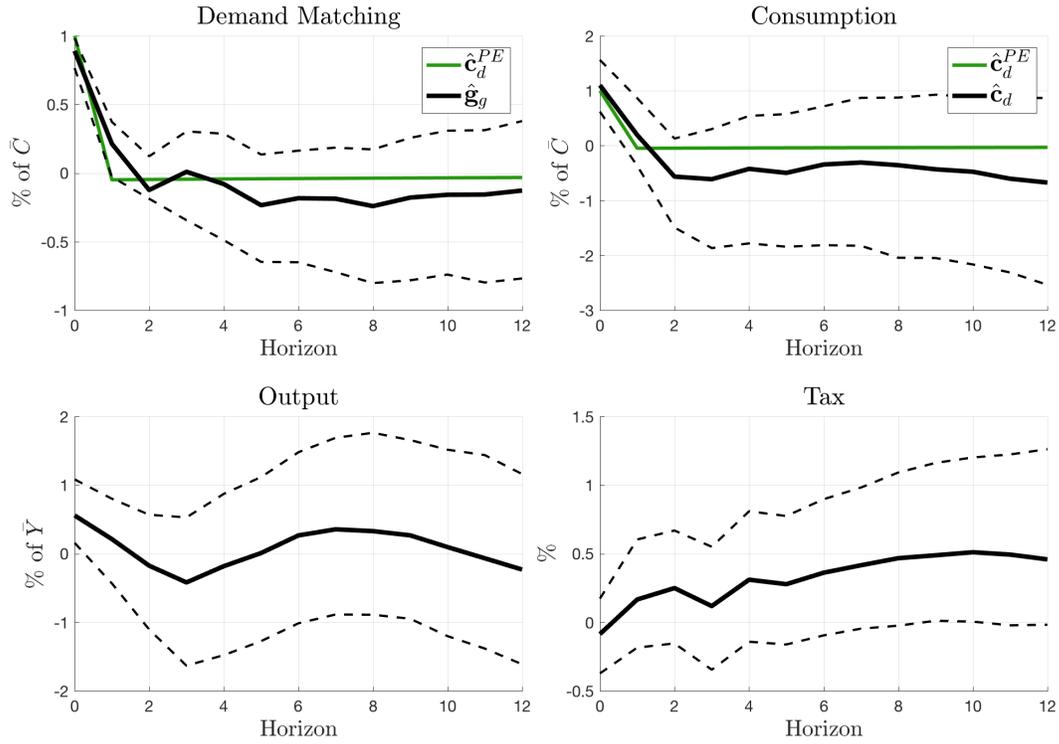


Figure E.1: Output and consumption responses to a redistribution shock, with the partial equilibrium net output response path matched to a linear combination of government spending shocks. The consumption response is computed in line with Proposition 2. The plot also shows the required demand matching 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 (13). 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.