The Missing Intercept: A Demand Equivalence Approach

Christian K. Wolf[†]
MIT & NBER

This version: March 29, 2023
First version: January 10, 2019

Abstract: I give conditions under which changes in private spending are accommodated in general equilibrium exactly like changes in aggregate fiscal expenditure. Under such demand equivalence, researchers can use time series evidence on fiscal multipliers to recover the general equilibrium "missing intercept" of shocks to private spending identified in the cross section. Through the lens of this theory, time series estimates of a fiscal multiplier around one suggest a missing intercept close to zero—an observation that I illustrate with an application to the 2008 tax rebates. I also discuss the robustness of this aggregation approach to plausible violations of demand equivalence.

[†]Email: ckwolf@mit.edu. This paper is a subtantially revised version of Chapter 1 of my PhD dissertation. I am grateful to my advisors Ben Moll, Mikkel Plagborg-Møller, Gianluca Violante and Mark Watson for their continued guidance and support. I have particularly benefited from detailed comments by the Editor, Emi Nakamura, as well as four anonymous referees. I also received helpful comments from Mark Aguiar, Adrien Auclert, Markus Brunnermeier, Dave Donaldson, Thorsten Drautzburg, Peter Ganong, Erik Hurst, Oleg Itskhoki, Greg Kaplan, Alisdair McKay, Jonathan Payne, Monika Piazzesi, Richard Rogerson, Matt Rognlie, Martin Schneider, Jón Steinsson, Ludwig Straub, Chris Tonetti, Iván Werning, Tom Winberry, and seminar participants at several venues. I acknowledge generous support from the Alfred P. Sloan Foundation and the Macro Financial Modeling Project as well as the Washington Center for Equitable Growth.

1 Introduction

A large literature in macroeconomics tries to estimate the aggregate effects of shocks to consumption and investment expenditure.[1] For most of these demand shifters, the experimental ideal—exogeneity at the macro level—is not attainable. In response, researchers increasingly leverage the cross-sectional variation available in micro data. Appealingly, because micro estimates rely exclusively on cross-sectional information, they do not require macroeconomic identification restrictions. The well-known shortcoming is that micro estimates miss general equilibrium effects that affect all micro units (price changes, labor demand, ...), and thus do not give macro counterfactuals—a so-called "missing intercept" problem.

So how can researchers aggregate their micro estimates into macro counterfactuals? The familiar Keynesian cross suggests one line of attack: if changes in private demand propagate similarly to changes in public expenditure, then simply scaling cross-sectionally identified private spending impulses by estimates of aggregate fiscal multipliers may approximately give a true macro effect. Back-of-the-envelope aggregation along these lines is already popular in policy practice (see Reichling & Whalen, 2012) and some academic work (e.g. Hausman, 2016). A second, largely separate literature instead turns micro estimates into macro counterfactuals through rich structural models (e.g. Kaplan & Violante, 2018).

This paper offers a hybrid perspective: I use structural models to gauge the informativeness of fiscal spending shocks for the general equilibrium effects of private demand shifters, clarify the conditions under which those fiscal shocks solve the missing intercept problem, and finally assess the plausibility of these conditions. The analysis proceeds in three steps. First, in the context of a relatively general structural macro model, I give a set of restrictions on economic primitives ensuring "demand equivalence"—that is, identical changes in private and public demand eliciting identical general equilibrium feedback. Second, I leverage this theoretical equivalence result to propose a measurement strategy. Researchers begin with a cross-sectional analysis, estimating the *direct* response of spending to some private demand shifter. Under demand equivalence, *aggregation* can then be achieved through time series evidence on the effects of a very particular change in fiscal expenditure: one that (i) induces the same path of aggregate demand; (ii) is financed using the same path of future taxes (if any); and (iii) occurs in the same macro environment. Requirements (i)–(iii) are obviously stringent, so I discuss how applied researchers can try to at least approximately satisfy them.

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

I illustrate with an application to the aggregate effects of stimulus checks. Here I argue that much prior empirical time series work points to fiscal multipliers around one (e.g., see Ramey, 2018) and thus suggests a "missing intercept" close to zero—i.e., a macro counterfactual close to the cross-sectionally estimated direct spending effects of such checks. Third, I discuss the robustness of this approach to aggregation to violations of demand equivalence. I find that most plausible violations lead to an upward bias: aggregation via fiscal multipliers understates the degree of general equilibrium crowding-out, and so overstates the actual macro causal effect. Using model simulations, I present general conditions under which the bias is likely to be small, and find them to be satisfied for stimulus checks.

Consumption demand equivalence. I begin by identifying conditions under which shocks to household spending—e.g., stimulus checks or credit tightenings—propagate in general equilibrium just like changes in public expenditure. The main building block result is that, in standard business-cycle models, linearized impulse responses to aggregate shocks can be characterized implicitly as solutions to a linear infinite-horizon system of market-clearing conditions. Private and public spending shocks thus induce identical general equilibrium effects as long as they perturb the same market-clearing conditions by the same amount.

Viewed through the lens of a relatively standard medium-scale business-cycle model, this abstract exclusion restriction maps into three substantive economic assumptions. First, households and government consume the same final good. If so, identical changes in private or public spending lead to identical excess demand for that common good. Second, households and government borrow and lend at the same interest rate. The identical expansions in private and public demand then induce the same fiscal deficit (in net present value terms), and so can be financed using identical paths of future taxes. In particular, if the excess demand path has zero net present value (as is the case for non-policy demand shocks like credit tightenings), then the identical change in public spending can be purely deficit-financed, with no direct tax response. Third, labor supply does not respond differentially to the two shocks; sufficient conditions for this are either the absence of wealth effects in labor supply or fully demand-determined employment. Under these three restrictions, for any given shock to private spending, I prove that *there exists* a public spending shock that solves the missing intercept problem—that is, the response of aggregate consumption to the private demand shifter is equal to the sum of a) the shifter's *direct* effect on consumer spending and b) the *total* response of consumption to the public spending shock. The constructive proof reveals the properties of this shock: it must (i) induce the same path of excess demand and (ii) be

associated with the same path of future taxes as the private demand shifter. My focus on linearized equilibria furthermore automatically imposes a third condition (iii): both shocks occur in the same macro environment (e.g., in a recession, or at the zero lower bound).

Measurement strategy. Next I show how to combine this theoretical demand equivalence insight with empirical time series and cross-sectional evidence on public and private spending shock transmission. Cross-sectional regressions of private spending on a demand shifter (e.g., check receipt) recover that shifter's *direct* effect on private demand. Time series methods on the other hand have been applied widely to estimate the aggregate effects of fiscal purchases. Demand equivalence tells us how to put the two together: time series fiscal multiplier estimates pin down the general equilibrium effects of *any* shifter of private demand that satisfies the conditions (i) - (iii)—same net excess demand, same financing, and same macro environment. The key challenge to leveraging demand equivalence in practice is thus the need to find cross-sectionally and time series identified shocks that align in this particular way. I propose the following solution. First, a researcher uses cross-sectional variation to estimate the direct demand response. Second, she turns to the time series fiscal policy literature to find a fiscal spending experiment—or a linear combination of such experiments—that induces as similar an aggregate demand path as possible.[2] With condition (i) ensured in this way, the researcher goes ahead and computes the sum of the a) cross-sectional and b) time series estimated consumption impulse responses. Conditions (ii) and (iii) then tell us how to interpret this sum: under demand equivalence, it is a valid counterfactual for a particular private demand shifter that (ii) is associated with the same tax response as the time series identified fiscal shock(s) and (iii) occurred in the same macro environment.

The chief appeal of this methodology is that it promises to allow researchers to empirically estimate the aggregate effects of shocks and policies for which no credible macro experiment is available. Its limitations are, first, that it requires the combination of very particular pieces of cross-sectional and time series variation; and second, that demand equivalence itself rests on restrictive assumptions. The remainder of the paper discusses these challenges.

Application. I showcase the method's applicability through a study of the 2008 "stimulus check" policy. Household-level evidence suggests that stimulus checks lead to a meaningful

[2]I elaborate in Section 3 on why combining different pieces of time series variation in this way is credible. The key insight is that different structural identification schemes pick up different structural shocks—transitory spending innovations for some, more persistent changes for others. McKay & Wolf (2022) discuss this point further, though mostly with a focus on monetary policy shocks.

but short-lived expansion in spending (e.g., see Parker et al., 2013). By the discussion above, under demand equivalence, a researcher wishing to aggregate these estimates would need to find a (i) similarly transitory fiscal spending expansion; furthermore, for her counterfactual to be informative about the 2008 experiment, this expansion should be (ii) deficit-financed and (iii) accommodated by the monetary authority. I argue that empirical evidence for fiscal spending experiments with those features suggests a multiplier of around one, with output moving approximately one-to-one and limited crowding-out of private spending. Putting the pieces together, I conclude that the stimulus check policy briefly but significantly stimulated consumption, with the overall response close to the direct effect estimated using micro data alone. Importantly, *any* structural model that is consistent with the assumptions of demand equivalence as well as my two pieces of evidence will invariably agree with this conclusion.

Assessment. Next I discuss whether this "demand equivalence" solution to the missing intercept problem is robust to empirically plausible violations of its strong assumptions. The nature of the exercise is as follows: I consider structural models violating demand equivalence, simulate data, implement my method, and report the error. My main laboratory is a rich heterogeneous-agent New Keynesian ("HANK") model. In this setting, demand equivalence fails *only* because of short-term wealth effects in labor supply. However, consistent with both micro-level evidence (Cesarini et al., 2017) and the results of much previous modeling (Christiano, 2011a), I find the inaccuracy resulting from this channel to be negligible. Several further model extensions—including relative price movements between public and private consumption bundles, productive benefits of public spending, and openness of the economy—unsurprisingly all tend to increase the approximation error. Interestingly, I find that the *sign* of the error is common for almost all of those extensions: fiscal multiplier estimates now miss some of the general equilibrium crowding-out relevant for private demand shocks, and so my procedure tends to *overstate* the aggregate response of private spending. I conclude that the output of the demand equivalence approximation should generally be interpreted as giving an upper bound for actual counterfactuals. This bound will be tight if: the spending shock is transitory, muting wealth and relative price effects; there is little leakage of spending abroad (e.g., because economy is quite closed); and the researcher uses time series fiscal multiplier estimates that capture changes in public consumption (and not productive investment). I verify all of these conditions in my application to stimulus checks.

Extensions. I conclude with a brief discussion of the scope of the demand equivalence approach. First, while my main application is to uniform stimulus checks, the method applies

without change to many other shifters of consumer demand, including targeted transfers, household deleveraging, or increases in inequality. To illustrate, I show that a fiscal contraction today offset by an expansion in the future identifies the missing intercept of a temporary increase in earnings inequality. Second, I extend the equivalence theory to shifters of investment demand and provide an application to bonus depreciation stimulus.

Literature. My analysis relates and contributes to several strands of literature.

First, the demand equivalence approach to the missing intercept problem 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 shocks (e.g. Mian & Sufi, 2009; Parker et al., 2013; 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. I give formal conditions under which a second literature—that on the aggregate effects of changes in government spending—can be informative about this "missing intercept." Comprehensive literature summaries are Hall (2009) and Ramey (2018).

Second, the theoretical demand equivalence result elaborates on the familiar Keynesian cross intuition of a common "demand multiplier" (e.g., Reichling & Whalen, 2012; Hausman, 2016). Building on the important work of Auclert & Rognlie (2018) and Auclert et al. (2018), I give sufficient conditions under which *there exist* aggregate public spending shocks that solve the missing intercept problem for private demand shifters, and in particular characterize the properties required of these shocks. In contemporaneous and independent work, Guren et al. (2020) and Chodorow-Reich et al. (2019) use the reverse logic to strip out the *local* general equilibrium effects present in cross-regional regressions.

Third, the proposed methodology naturally complements existing strategies for the estimation of (policy) counterfactuals in macroeconomics. In its reliance on general exclusion restrictions rather than parametric models, it is semi-structural in exactly the same way as Structural Vector Autoregressive (SVAR) analysis (Sims, 1980). My analysis also connects with the "sufficient statistics" approach that is common in public finance (Chetty, 2009) and by now also increasingly widespread in macroeconomics (e.g., Auclert et al., 2018; Nakamura & Steinsson, 2018; McKay & Wolf, 2022). I show that, across a particular family of structural models, certain moments—the cross-sectional and time series estimates required by my methodology—fully pin down the desired macro counterfactual.

Outline. Section 2 establishes the consumption demand equivalence result. In Section 3, I show how to connect this theoretical result with cross-sectional and time series empirical

evidence, and discuss an application to stimulus checks. Section 4 then critically assesses the output of the proposed procedure. Applications to other shifters of consumer demand as well as the extension to investment are presented in Section 5, and Section 6 concludes.

2 Consumption demand equivalence

This section builds on the familiar Keynesian cross logic to develop an equivalence result for shocks to private consumption and to public spending. In Sections 2.1 and 2.2 I first consider a standard business-cycle model and discuss what restrictions on primitives are needed to ensure such equivalence. Section 2.3 complements this analysis with a general formulation of shock equivalence as a set of exclusion restrictions on equilibrium conditions.

2.1 Model

I begin by presenting a particular model environment, general enough to nest many seminal contributions to quantitative business-cycle analysis.[3] The purpose of the model is twofold. First, it allows me to present a set of economically interpretable sufficient conditions for demand equivalence in a familiar, canonical environment. Second, the model will form the backbone of my critical assessment of the demand equivalence methodology in Section 4.

Time is discrete and runs forever, $t = 0, 1, \ldots$. The 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 the model's steady state after one-time unexpected aggregate innovations at time 0; for vanishingly small innovations, these transition paths are equivalent to standard impulse response functions computed from the first-order perturbation solution of an otherwise identical model with aggregate risk.[4] Anticipating my main application, I will focus on two such innovations: first, stimulus checks sent to households, and second, a transitory expansion in government spending. Section 2.3 shows how the demand equivalence result extends to generic policy and non-policy shifters of consumption demand (e.g., changes in borrowing constraints).

[3] The environment nests conventional estimated New Keynesian models (e.g. Smets & Wouters, 2007), models with uninsurable household income risk (Aiyagari, 1994; McKay et al., 2016), and models with rich real and financial firm-level investment frictions (Khan & Thomas, 2013; Winberry, 2018). A thorough assessment of its generality and limitations is relegated to Section 4.

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

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

HOUSEHOLDS. A unit continuum of households $i \in [0, 1]$ has preferences over consumption c_{it} and labor ℓ_{it}. The real relative price of the consumption bundle in terms of the economy's numeraire is denoted p_t^c. Households are subject to idiosyncratic productivity risk e_{it}, and can self-insure by investing in liquid nominal bonds b_{it}^h, with nominal returns i_t^b and subject to a fixed borrowing constraint $\underline{b}$. Borrowing incurs an additional intermediation cost $\kappa_b \geq 0$. Income consists of labor earnings as well as (potentially type-specific) lump-sum transfers 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, transfers, dividends, hours worked and inflation (π_t), the consumption-savings problem of household i is thus

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

such that

$$p_t^c c_{it} + b_{it}^h = (1 - \tau_\ell) w_t e_{it} \ell_{it} + \left(\frac{1 + i_{t-1}^b}{1 + \pi_t} + \kappa_b \mathbf{1}_{b_{it-1}^h < 0} \right) b_{it-1}^h + \tau_{it} + d_{it}, \quad b_{it}^h \geq \underline{b}$$

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

Because of frictions in the labor market, household hours worked are determined by labor unions, as in Erceg et al. (2000) and 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,

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

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

Fiscal policy. The fiscal authority consumes a bundle with real relative price p_t^g. 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 flow budget constraint is

$$\frac{1+i_{t-1}^b}{1+\pi_t} b_{t-1} + p_t^g g_t + \tau_t = \tau_\ell w_t \ell_t + b_t \tag{2}$$

I assume that total government spending $\mathbf{g} = \mathbf{g}(\boldsymbol{\varepsilon}_g)$ and the discretionary part of stimulus checks $\boldsymbol{\tau}^x = \boldsymbol{\tau}^x(\boldsymbol{\varepsilon}_\tau)$ follow exogenous processes. A financing rule is a mapping from spending targets $(\boldsymbol{g}, \boldsymbol{\tau}^x)$, initial nominal debt b_{-1} and prices and quantities $(\mathbf{w}, \boldsymbol{\ell}, \mathbf{i}^b, \boldsymbol{\pi}, \boldsymbol{p}^g)$ into the endogenous part of transfers $\boldsymbol{\tau}^e$ such that $\boldsymbol{\tau} = \boldsymbol{\tau}^e + \boldsymbol{\tau}^x$, the flow government budget constraint holds at all periods t, and $\lim_{t\to\infty} \widehat{b}_t = 0$. That is, lump-sum taxes adjust in response to fiscal outlays—both outright expenditure and stimulus checks—to ultimately return government debt to its steady-state level. I emphasize that all results below extend without change to the alternative assumption of outlays financed with time-varying distortionary taxes τ_ℓ; the key restriction for demand equivalence will be that stimulus checks and spending increases are financed using *identical* paths of taxes, distortionary or not.

Rest of the economy. Since my focus is on the equivalence of private and public expansions in demand, I only sketch the rest of the model, with a detailed outline provided in Appendix B.1. The corporate sector is populated by three sets of firms: a unit continuum of heterogeneous, perfectly competitive intermediate goods producers; a unit continuum of monopolistically competitive retailers with nominal price rigidities; and aggregators for final (private and public) consumption and investment goods. Intermediate goods producers ac-

[5] A uniform hiring rule is the natural assumption in sticky-wage heterogeneous-household models, but is of course awkward in the flexible-wage limit, as it then does not nest the alternative natural case of flexible labor supply for each individual household. I consider a model without unions in Appendix B.3.4.

cumulate capital, hire labor, issue risk-free debt, and sell their composite intermediate good, possibly subject to capital adjustment costs as well as constraints on their equity and debt issuance. Retailers purchase the intermediate good, costlessly differentiate, monopolistically set prices, and sell their differentiated good on to the competitive aggregators. Dividends from the corporate sector as a whole are paid out to households.

The last remaining entity in the model is the monetary authority. This monetary authority sets nominal rates on liquid bonds i_t^b following a conventional Taylor rule.

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

2.2 The equivalence result

I now formalize the Keynesian cross intuition of different "demand" shocks propagating similarly in general equilibrium. A precise statement of such equivalence requires a definition of direct (or partial equilibrium) responses and indirect (or general equilibrium) effects.

I assume that the consumption-savings problem (1) has a unique solution for any path of prices, quantities and shocks faced by households. Aggregating the solutions across households, we obtain a consumption function $\mathbf{c} = \mathbf{c}(\mathbf{s}^h; \boldsymbol{\varepsilon})$, where $\boldsymbol{\varepsilon}$ pins down stimulus check income $\boldsymbol{\tau}^x$ and $\mathbf{s}^h = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \boldsymbol{\ell}, \boldsymbol{\tau}^e, \mathbf{d}, \boldsymbol{p}^c)$ collects the residual household income as well as saving returns and prices—i.e., objects that adjust in general equilibrium. The total impulse response of consumption to the generic shock path $\boldsymbol{\varepsilon}$ is then given as

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

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

[6] More precisely, I make implicit assumptions on functional forms and parameter values that guarantee that there is a unique deterministic steady state. In all numerical exercises, I have verified the uniqueness of the steady state and the (local) existence and uniqueness of transition paths (see Appendix B.2).

[7] My definition of the partial equilibrium consumption response abstracts from endogenous partial equilibrium adjustments in earnings. I do so for three reasons. First, many empirical estimates of household

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

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

Similarly, let the indirect (general equilibrium) feedback be

$$\widehat{\mathbf{c}}_{\varepsilon}^{GE} \quad \equiv \quad \mathbf{c}(\mathbf{s}_{\varepsilon}^h; \mathbf{0}) \quad - \quad \mathbf{c}(\bar{\mathbf{s}}^h; \mathbf{0}) \tag{4}$$

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

$$\widehat{\mathbf{c}}_{\varepsilon} \quad = \quad \widehat{\mathbf{c}}_{\varepsilon}^{PE} \quad + \quad \widehat{\mathbf{c}}_{\varepsilon}^{GE} \tag{5}$$

For example, for a stimulus check policy, the direct response part captures the effect of the stimulus check $\boldsymbol{\tau}^x(\boldsymbol{\varepsilon}_\tau)$ on spending in isolation, while the indirect effect contains both the tax financing as well as all other general equilibrium effects (e.g., labor demand, prices, ...).

The remainder of this section establishes properties of the decomposition (5)—a general "demand equivalence" result. Section 3 will then connect theory and measurement, linking the components of (5) to measurable objects and so in particular to the "missing intercept" (or aggregation) problem of cross-sectional regressions.

Demand equivalence & its implications. To state a demand equivalence result in the model of Section 2.1, I require three additional restrictions on model primitives.

The first assumption restricts goods bundles in the economy.

Assumption 1. *Households and government consume a single, homogeneous final good. It follows that* $p_t^c = p_t^g = 1$ *for all* t*.*

The second assumption relates to the interest rates faced by households and government: all agents must borrow and lend at a common interest rate.

Assumption 2. *There is no borrowing wedge (*$\kappa_b = 0$*), so households and government borrow and lend at the same interest rate* i_t^b*.*

---

spending responses to sudden income changes are actually interpretable as such netted spending elasticities (e.g. see Auclert, 2019, footnote 34). Second, in models with union-intermediated labor supply, replicating cross-sectional regressions differences out labor responses (see Proposition 2). Third, microeconomic evidence suggests that short-run wealth effects are weak (Cesarini et al., 2017). Nevertheless, in Appendix D.3, I repeat my analysis in an alternative model without unions, but with a non-standard preference parameterization allowing for weak short-run wealth effects (as in Jaimovich & Rebelo, 2009; Galí et al., 2012).

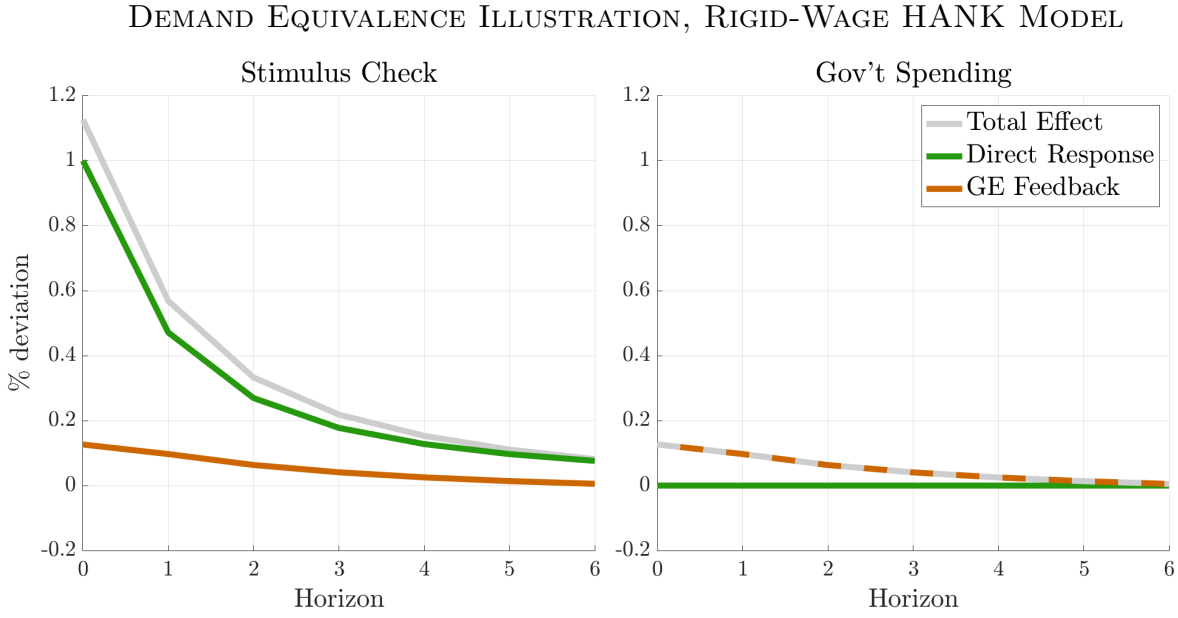

Figure 1: Consumption impulse response decompositions after stimulus check and government spending shocks in the estimated HANK model of Section 4.1, but with fully rigid wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

The third assumption restricts the economy's labor market. In response to the partial equilibrium increase in consumption demand $\widehat{\mathbf{c}}_\tau^{PE}$ induced by a given stimulus check policy, 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 $\widehat{\boldsymbol{\ell}}_\tau^{PE}$, defined formally in Appendix B.1. My third assumption provides two possible sufficient conditions to guarantee that $\widehat{\boldsymbol{\ell}}_\tau^{PE} = \mathbf{0}$.

Assumption 3. *Either household preferences are such that there are no wealth effects in labor supply, or wages are perfectly sticky (i.e., wage adjustment costs are infinitely large).*

These three assumptions are sufficient for the following "demand equivalence" result.

Proposition 1. *Consider a stimulus check policy $\boldsymbol{\varepsilon}_\tau$, and suppose that Assumptions 1 to 3 hold. Then, for a fiscal spending policy $\boldsymbol{\varepsilon}_g$ such that (i) $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}_\tau^{PE}$ (identical net excess demand) and (ii) $\widehat{\boldsymbol{\tau}}_g^e = \widehat{\boldsymbol{\tau}}_\tau^e$ (identical tax response), we have that, to first order,*

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

Under Assumptions 1 to 3, shocks to private and public net excess demand induce the exact same general equilibrium feedback effects. Proposition 1 presents the key implication

of such demand equivalence that is relevant for the analysis in this paper: the response of aggregate consumption to a fiscal spending shock with the particular properties (i) and (ii) at the same time gives the general equilibrium feedback effects associated with the stimulus check policy $\boldsymbol{\varepsilon}_\tau$. Figure 1 provides a graphical illustration: in a quantitative HANK model satisfying Assumptions 1 to 3, the general equilibrium feedback effects on consumption after a stimulus check policy (orange line, left panel) and a fiscal spending expansion (orange line, right panel) are exactly the same. In the chosen model parameterization, interest rates and tax financing induce some general equilibrium crowding-out, while Keynesian employment effects lead to crowding-in, with the latter effect dominating slightly.

Finally, I emphasize that my focus on linearized equilibria also implicitly imposes a further condition (iii) on the two shocks: they need to occur in the same macro environment (e.g., in a recession, or when nominal interest rates are constrained by a binding effective lower bound). This assumption will be important in communicating the results of my empirical measurement strategy in Section 3.

Proof sketch. The proof of the demand equivalence decomposition in (6) leverages the "sequence-space" approach to equilibrium characterization and in particular builds on the important contributions of Auclert & Rognlie (2018) and Auclert et al. (2018).

The argument proceeds as follows. Equilibria even in the rich model of Section 2.1 can be characterized as a system of several aggregate prices and quantities adjusting to clear several markets. Assumptions 1 to 3 turn out to be sufficient to ensure that, for any given shock to private spending, a public expenditure shock with properties (i) and (ii) will perturb the same market-clearing conditions by the same amount, thus eliciting the same general equilibrium adjustment and so implying (6). First, Assumption 1—in conjunction with requirement (i) on the fiscal shock, $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}_\tau^{PE}$—ensures that the private and public demand shocks lead to the same excess demand pressure for the common final good. Second, since by Assumption 2 households and governments borrow and lend at identical rates, this common excess demand path can *in principle* be financed using identical paths of future taxes. Property (ii) of the fiscal spending shock ensures that this is indeed the case. Third, Assumption 3—a restriction on household behavior—ensures that the consumption increase induced by stimulus checks does not lead to any *direct* adjustment in hours worked.[8]

Overall, my statement of demand equivalence in Proposition 1 offers two key insights relative to previous work. First, it explicitly characterizes the properties of the fiscal shock

[8] In general equilibrium, however, hours worked can and generally will respond to both shocks.

required for the decomposition in (6) to hold. These properties will take center stage in the connection of theory to measurement in Section 3. Second, it reveals that not all assumptions necessary to arrive at a Keynesian cross-type equilibrium characterization (as e.g. in Auclert et al., 2018) are also necessary for a demand equivalence result. Notably, neither the absence of investment nor the assumption of a fixed real rate of interest are required here.

2.3 Extension to general exclusion restrictions

In the analysis so far I have used a particular shifter of consumer spending—stimulus checks—and a particular structural model—the general framework of Section 2.1—to present demand equivalence as a set of restrictions on economic primitives. As it turns out, however, many of the restrictions implicit in this framework are in fact unnecessary. To make this point, I now complement the previous model-based discussion with an abstract statement of shock equivalence as a set of exclusion restrictions imposed on a general linearized system. Throughout, I continue to use the same notational conventions as in my baseline model.

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

$$\frac{\partial \mathcal{H}}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d \quad = \quad \frac{\partial \mathcal{H}}{\partial \boldsymbol{\varepsilon}_g} \times \boldsymbol{\varepsilon}_g \tag{7}$$

it follows immediately that, to first order,

$$\widehat{\mathbf{c}}_d \quad = \quad \underbrace{\widehat{\mathbf{c}}_d^{PE}}_{\text{PE response}} \quad + \quad \underbrace{\widehat{\mathbf{c}}_g}_{=\text{ GE feedback}} \tag{8}$$

exactly as in Proposition 1.[9] Condition (7) is a general exclusion restriction on the equilibrium system: both shocks must enter all equilibrium equations symmetrically. The proof of Proposition 1 works because, under my imposed restrictions, the equilibrium can be cast in

[9] An equilibrium is a solution of $\frac{\partial \mathcal{H}}{\partial \boldsymbol{x}} \times \boldsymbol{x} + \frac{\partial \mathcal{H}}{\partial \boldsymbol{\varepsilon}_d} \times \boldsymbol{\varepsilon}_d + \frac{\partial \mathcal{H}}{\partial \boldsymbol{\varepsilon}_g} \times \boldsymbol{\varepsilon}_g = \mathbf{0}$. In stating (8) I am assuming that this system has a solution for $\boldsymbol{\varepsilon}_d$; it then follows from my assumptions that the same path of $\boldsymbol{x}$ also solves the system for $\boldsymbol{\varepsilon}_g$. Equilibrium *uniqueness* would require further assumptions on $\frac{\partial \mathcal{H}}{\partial \boldsymbol{x}}$ (e.g., invertibility).

a form consistent with (7).[10] In Appendices C.2 and C.3 I give examples of other shocks and models that can be written in this form. First, I extend the model of Section 2.1 to allow for time preference shocks as a generic non-policy consumption shifter, capturing in reduced-form primitive shocks like a tightening of household borrowing constraints. Equivalence obtains under the same restrictions as in Section 2.2; in particular, since non-policy private demand shifters $\boldsymbol{\varepsilon}_d$ necessarily induce partial equilibrium spending paths $\widehat{\boldsymbol{c}}_d^{PE}$ with zero net present value, requirement (ii) now dictates that the equivalent public spending shock $\boldsymbol{\varepsilon}_g$ is *purely deficit-financed*, with taxes moving only because of general equilibrium feedback into the government budget, and not directly because of the shock $\boldsymbol{\varepsilon}_g$. Second, I consider several examples of popular models beyond the usual New Keynesian tradition—including for example models with non-rational expectation formation of firms and households—, and show that they still all fit into the general semi-structure of (7).

2.4 Summary & Outlook

This section has presented conditions ensuring the commonality of general equilibrium effects of private and public spending shocks. I emphasize, however, that this equivalence result by itself says nothing about the *strength* of that common general equilibrium feedback: in the illustrative HANK model example of Figure 1, general equilibrium effects are relatively weak; in Appendix C.4, I instead show two other examples, one with full crowding-out, the other with very strong amplification. The chief appeal of the demand equivalence result is instead as a measurement device, as I discuss next.

3 Solving the missing intercept problem

This section shows how the demand equivalence result together with time series evidence on the dynamic causal effects of fiscal purchases can be used to solve the "missing intercept" aggregation problem for cross-sectionally identified consumption demand shifters. Section 3.1 begins by tying the theoretical decomposition in (8) to empirically measurable objects. Section 3.2 then uses this mapping between theory and empirics to propose a general empirical measurement strategy. Finally in Section 3.3 I discuss how to use the method to study an important policy instrument: deficit-financed stimulus checks.

[10] Casting my results as exclusion restrictions on equilibrium representations suggests a connection to the identification of systems of simultaneous equations. This connection is explored in Guren et al. (2020).

3.1 From theory to measurement

The general demand equivalence decomposition (8) has two parts: a) the direct (or partial equilibrium) response of consumer spending to some private demand shifter $\widehat{\mathbf{c}}_d^{PE}$ and b) the effect of a change in fiscal purchases on total household consumption $\widehat{\mathbf{c}}_g$, which under the conditions of the proposition equals the private shifter's general equilibrium term $\widehat{\mathbf{c}}_d^{GE}$. Each of those two components can be tied to objects estimated in previous empirical work.

Micro regressions. Cross-sectional regressions of household-level consumption on idiosyncratic shock exposure promise to identify part a): the direct consumption response.

To make this claim precise, I return to the model of Section 2.1, but now assume that the transfer stimulus received by household i is $\varepsilon_{\tau it} = \xi_{\tau it} \times \varepsilon_{\tau t}$, where $\xi_{\tau it}$ is i.i.d. across households and time (and uncorrelated with any household characteristics), with $\mathbb{E}(\xi_{\tau it}) = 1$ and $\text{Var}(\xi_{\tau it}) > 0$. Given this heterogeneity in shock exposure, I can study regressions run on the cross-section of households. A typical cross-sectional regression takes the form

$$c_{it+h} = \alpha_i + \delta_t + \beta_{\tau h} \times \varepsilon_{\tau it} + u_{it+h}, \quad h = 0, 1, 2, \ldots \tag{9}$$

where α_i and δ_t are individual and time fixed effects. It is straightforward to show that, under my assumptions, regressions such as (9) estimate average household-level causal effects that are interpretable as direct partial equilibrium shock responses, as claimed.[11,12]

Proposition 2. *Suppose an econometrician observes a panel of consumption $\{c_{it}\}$ and shock exposure $\{\varepsilon_{\tau it}\}$. Then the ordinary least-squares estimand of $\boldsymbol{\beta}_\tau \equiv (\beta_{\tau 0}, \beta_{\tau 1}, \ldots)'$ satisfies*

$$\boldsymbol{\beta}_\tau \quad = \quad \widehat{\mathbf{c}}_\tau^{PE} \tag{10}$$

Note that common general equilibrium effects are absorbed by the time fixed effect. It is in light of this result that $\widehat{\mathbf{c}}_\tau^{GE}$ is often referred to as the "missing intercept."

Fiscal multipliers. The dynamic causal effects of changes in fiscal purchases on macroeconomic outcomes can be estimated using the conventional semi-structural macroeconometric

[11] Formally, for Proposition 2, I consider the first-order perturbation solution of the model in Section 2.1 with aggregate shocks ε_{st}, $s \in \{\tau, g\}$. This ensures that all regression estimands are well-defined.

[12] Note that the regression (9) is run at the *household* level. This is important: cross-*regional* regressions (e.g. as in Mian et al., 2013) contain local general equilibrium effects and so do not identify my notion of direct effects. I extend my approach to such cross-regional regressions in Wolf (2019).

toolkit (see Ramey, 2016, for a review). Under demand equivalence, such time series analysis promises to identify part b): the "missing intercept" of general equilibrium effects.

Most previous empirical work on fiscal multipliers has relied on one of three possible sets of identifying assumptions. First, researchers may have access to direct, "narrative" measures of fiscal shocks (Ramey, 2011). Second, outside information—either in the form of zero or sign restrictions (Blanchard & Perotti, 2002; Mountford & Uhlig, 2009) or via measures of other structural shocks (Caldara & Kamps, 2017)—can help researchers pin down fiscal rules and so identify spending shocks. Third, professional forecast errors of government spending may be interpreted as orthogonal to any (known) rules-based spending, and so may identify fiscal shocks when coupled with additional exogeneity or timing assumptions (Ramey, 2011; Drautzburg, 2020). In all of these cases, the desired causal effects can then be estimated using Vector Autoregressions (VARs) or Local Projections (LPs). Across this range of identifying assumptions, estimated fiscal multipliers tend to lie around 1 (Ramey, 2018).

The key implication of demand equivalence is that these fiscal multiplier estimates are actually informative about the propagation of a larger *menu* of structural shocks. Any identified aggregate fiscal shock is associated with an (estimable) implied path of fiscal purchases $\widehat{\mathbf{g}}_g$ and underlying financing (taxes and/or deficits). Under demand equivalence, this shock in general equilibrium propagates exactly like *any* shock to private spending—say, stimulus checks—with (i) the same excess demand path $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}_d^{PE}$ and (ii) the same tax response, and furthermore (iii) occurring in the same macroeconomic environment, which in particular means the same response of the monetary authority. It follows that conventional time series estimates of fiscal multipliers actually contain much more information than commonly believed: for a suitable private demand shifter $\boldsymbol{\varepsilon}_d$, they for free give that shifter's general equilibrium effects on consumption $\widehat{\mathbf{c}}_d^{GE}$ as well as the total causal effects on all other macro aggregates (employment, investment, inflation, ...).

Combining cross section & time series. The previous discussion reveals that, through the lens of the theory of demand equivalence, cross-sectional evidence on private demand shifters and time series estimates of public spending shock propagation are useful complements. When put together in accordance with Proposition 1, they fully characterize aggregate counterfactuals for the private demand shifters, allowing researchers to estimate the causal effects of these shifters even in the absence of exogenous macro variation.

To make this potentially powerful insight operational, however, we must confront an important challenge: the demand equivalence result only allows us to combine *particular*

cross-sectional and time series identified shocks—those that satisfy the conditions (i) - (iii). This is challenging; for example, nothing guarantees that any given time series and cross-sectional causal effect estimates give demand paths that are aligned as required by (i). The next subsection presents a concrete empirical methodology that tackles these challenges.

3.2 The empirical methodology

I consider a researcher that wishes to predict the aggregate effects of some shifter of private demand ε_d. As a first step towards doing so she has successfully leveraged cross-sectional identifying information to estimate the direct spending response $\widehat{\mathbf{c}}_d^{PE}$. By the previous discussion, she now needs to find a time series aggregate fiscal spending shock that (i) induces a similar path of excess demand, (ii) is financed in the same way, and (iii) occurs in the same macroeconomic environment. I begin with the first requirement.

Aligning excess demand paths. To ensure requirement (i), I suggest that the researcher begins by leveraging one or several of the semi-structural time series identification approaches discussed in Section 3.1 to estimate the aggregate effects of a *menu* of n_k different kinds of government spending shocks $\{\varepsilon_{g_k}\}_{k=1}^{n_k}$ with implied spending paths $\widehat{\mathbf{g}}_{g_k}$. To see why it actually makes sense to expect different identification schemes to yield different spending profiles $\widehat{\mathbf{g}}_{g_k}$, it will be instructive to consider a general policy rule for fiscal purchases:

$$g_t = \underbrace{f(\Omega_t)}_{\text{systematic component}} + \underbrace{\varepsilon_{g,t,0}}_{\text{contemp. shock}} + \underbrace{\sum_{\ell=1}^{\infty} \varepsilon_{g,t-\ell,\ell}}_{\text{news shocks}}$$

Ω_t denotes the policymaker's information set at date t, and $f(\Omega_t)$ is the systematic component of policy—government purchases responding to aggregate economic conditions, through both automatic as well as discretionary fiscal stabilization. The second and third terms then give deviations from that systematic rule; that is, policy shocks. $\varepsilon_{g,t,0}$ in this set-up is a classical contemporaneous policy shock: a deviation from the policy rule at date t that was announced at date t. The third term instead collects news shocks: deviations from the rule at date t that were already announced in prior periods, at $t-\ell$. The various empirical identification schemes reviewed above aim to isolate exogenous variation in government purchases—in other words, they aim to isolate some linear combination of the policy shocks $\{\varepsilon_{g,t,\ell}\}$. Some experiments will capture transitory spending impulses (i.e., mostly correlating with $\varepsilon_{g,t,0}$), while others will reflect more persistent spending dynamics and thus more gradual deviations from the rule

(i.e., mostly correlating with $\varepsilon_{g,t,\ell}$ for $\ell > 0$). Different empirical strategies yielding different time profiles $\widehat{\boldsymbol{g}}_g$ of fiscal purchases is thus not necessarily a sign of mis-specification; rather, it is what we should *expect*, with each identified exogenous fiscal intervention consisting of a different mix of dynamic policy treatments. This observation is developed in more detail in McKay & Wolf (2022), though there largely with a focus on monetary policy shocks.

Given a list of identified fiscal expenditure paths $\widehat{\mathbf{g}}_{g_k}$, I then propose that the researcher projects her cross-sectionally identified private spending path $\widehat{\mathbf{c}}_d^{PE}$ on the space spanned by those fiscal shocks. This gives

$$\widehat{\mathbf{c}}_d^{PE} \quad = \quad \sum_{k=1}^{n_k} \gamma_k \quad \times \quad \widehat{\mathbf{g}}_{g_k} \quad + \quad \text{error} \tag{11}$$

If the error term in (11) is sufficiently small, then we may consider the weighted average

$$\sum_{k=1}^{n_k} \gamma_k \quad \times \quad \widehat{\mathbf{c}}_{g_k} \tag{12}$$

as a promising candidate to learn about $\widehat{\mathbf{c}}_d^{GE}$. In Section 3.3 I will provide a concrete example of a case in which this matching error is very small, implying that the proposed approach to aggregation can indeed be operationalized.[13]

Constructing the general equilibrium counterfactual. Having addressed the demand path matching problem in this way, I suggest that researchers construct the demand equivalence approximation as

$$\widehat{\mathbf{c}}_d \quad = \quad \underbrace{\widehat{\mathbf{c}}_d^{PE}}_{\text{PE response}} \quad + \quad \underbrace{\sum_{k=1}^{n_k} \gamma_k \quad \times \quad \widehat{\mathbf{c}}_{g_k}}_{\text{GE feedback}} \tag{13}$$

with the weights γ_k as in (11).

Conditions (ii) and (iii) then affect the *interpretation* of the result of (13). Under demand equivalence, this sum is a valid general equilibrium counterfactual for a *particular* demand shock $\boldsymbol{\varepsilon}_d$ that is associated with the same endogenous response of taxes as the identified public spending shock(s) *and* occurred in the same macro environment, in particular including the

[13]Note that my construction in (11) leverages a linear combination of distinct *date*-0 fiscal policy shocks. Alternatively, researchers could consider following the strategy of Sims & Zha (1995) and using a *sequence* of date $t = 0, 1, 2, \ldots$ fiscal policy shocks to *exactly* align net excess demand paths. As discussed in McKay & Wolf (2022), this approach is immune to the Lucas critique only in the absence of forward-looking expectational effects. Since I wish to allow for such expectational effects I do not consider this strategy further.

same response of interest rates. For example, if the private demand shifter is lump-sum income, and the aligned fiscal shock is persistently deficit-financed and largely accommodated by the monetary authority, then the sum will give us a counterfactual for a deficit-financed, accommodated stimulus check policy.[14] Concrete applications of this proposed methodology are provided in the next subsection (for stimulus checks) as well as in Section 5.

3.3 Applying the method to stimulus checks

In this section I discuss how to leverage the demand equivalence approach to learn about the aggregate effects of a popular fiscal policy tool: stimulus checks. The study of such stimulus checks is well-suited to illustrate the proposed approach, for two main reasons. First, even though such checks are an increasingly popular fiscal policy tool, there are few estimates of their aggregate effects. Intuitively, the core estimation challenge is that there is little-to-no *exogenous* time series variation in those stimulus payments. Second, a wealth of micro data has allowed researchers to credibly estimate the direct spending response of households to the receipt of (small) lump-sum gains, giving the required micro identification.

I begin by reviewing empirical evidence on the direct spending effects of stimulus checks. Then, leveraging the theory of demand equivalence, I ask what a researcher would need to believe about the aggregate effects of fiscal purchases to draw conclusions about the missing general equilibrium effects of stimulus checks. Appendix F.1 complements this high-level discussion with a concrete worked-out application in which I leverage my own empirical estimates of fiscal spending multipliers to construct the desired stimulus check counterfactual. Throughout my main focus is on the fiscal stimulus check episode of 2008.

Direct effect. The direct response of consumer spending to a one-off stimulus check policy is given as

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

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

[14] It follows from the arguments of McKay & Wolf (2022) that, to construct counterfactuals for arbitrary alternative monetary policy reactions, it suffices to add suitable date-0 monetary policy shocks—again something identified by the prior time series literature. I leave this extension of my method to future work.

are either exactly or approximately interpretable as the desired average MPC (e.g. Johnson et al., 2006; Parker et al., 2013; Jappelli & Pistaferri, 2014; Fagereng et al., 2018). 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 then decays back to zero quickly. In particular, the point estimates of Parker et al. (2013) and Broda & Parker (2014) suggest that, following a one-off stimulus check, total consumption expenditures increase by around 50 cents on the dollar on impact and 20 cents in the subsequent quarter. Translated to the size of the 2008 stimulus check policy, this corresponds to around 1.5 per cent of total personal consumption expenditure on impact, and then around 0.6 per cent in the following quarter.[15] I provide a further discussion of these empirical estimates—and in particular challenges to interpreting them as the required direct spending response $\widehat{c}^{PE}_{\tau t}$—in Appendix E.1.

General equilibrium aggregation. What would a researcher need to know about the propagation of fiscal spending shocks to aggregate those cross-sectional estimates? By the theory of demand equivalence, she would need to know the aggregate effects of a similarly transitory increase in fiscal purchases, ensuring requirement (i)—the alignment of net excess demand paths. Furthermore, for her to be able to construct a counterfactual that is in fact plausibly informative about the 2008 stimulus check policy, that transitory spending increase should furthermore be (ii) quite persistently deficit-financed and (iii) accommodated by the monetary authority, with nominal rates remaining largely unchanged.[16]

In Appendix F.1 I discuss one suitable candidate to satisfy these particular requirements: fiscal shocks identified through professional forecast errors of government purchases (following Ramey, 2011). I there show that such shocks indeed induce transitory and deficit-financed fiscal spending expansions, with those expansions furthermore largely accommodated by the central bank.

Aggregate counterfactual. By demand equivalence, if a researcher has been able to estimate the effects of a fiscal spending expansion with the required properties, then the last remaining step would be to simply sum a) the cross-sectionally estimated direct spending

[15] Recent work (Orchard et al., 2022) has argued that these MPC estimates are likely to be biased upward. I repeat my analysis with their favored alternative MPC estimates in Appendix F.2. My headline conclusion—general equilibrium counterfactuals close to direct spending responses—still applies. The main change is that the smaller impact direct spending response invariably maps into a smaller total general equilibrium effect.

[16] In principle the researcher should also ensure that the spending shock is taken from a time period in which there was slack in the aggregate economy. By the results of Ramey & Zubairy (2018), however, such slack is likely to be of limited importance compared to the response of interest rates.

response and b) the effect of the fiscal spending shock on private consumption. As argued in Section 3.1, most empirical evidence on fiscal spending shocks suggests multipliers around 1 and only limited crowding-out of private spending. The particular fiscal expenditure shock that I estimate in Appendix F.1 is no exception: output increases approximately one-to-one with the increase in fiscal purchases, while total consumption responds very little, with some slight crowding-out over time. These observations suggest an *aggregate* consumption effect of stimulus checks close to the micro-estimated effect—that is, a general equilibrium transfer multiplier that is similar to the household-level average MPC, or a "missing intercept" close to zero. Thus, perhaps surprisingly, the various price and multiplier effects cited in previous empirical and theoretical work seem to roughly cancel.

How should we interpret this finding? By the theoretical analysis in Section 2, we know immediately that *any* structural model satisfying demand equivalence and estimated to match the cross-sectional and time series empirical evidence reviewed here will invariably arrive at that same conclusion. It thus remains to discuss the plausibility of the demand equivalence assumption itself. I do so in the next section.

4 How plausible is demand equivalence?

The methodology and results presented in Section 3 are *exactly* valid only under the strong conditions required for demand equivalence. These conditions, however, are rather unlikely to hold in practice. I thus now apply the proposed methodology to artificial data generated from a large number of models violating demand equivalence, and ask whether the estimated counterfactuals at least *approximately* equal the model's true causal effects.[17] In keeping with my theoretical analysis in Section 2 and the empirical application in Section 3.3, I throughout consider stimulus checks as my example of a private demand shifter. In the interest of space I here mostly just provide intuition for the size and sign of the implied biases, with the interested reader referred to Appendix D for detailed results.

I proceed in three steps. First, in Section 4.1, I consider a quantitative HANK model, enriched to feature many of the bells and whistles of the business-cycle literature. Second, in Section 4.2, I further extend this baseline environment with various additional frictions specifically designed to break demand equivalence. Finally in Section 4.3 I summarize the results from my model laboratories in the form of recommendations for applied practice.

[17] Formally, I consider an econometrician with access to infinitely large samples of cross-sectional and time-series model-generated data. Using the data, the econometrician implements the method of Section 3.

4.1 Quantitative business-cycle models

My first model laboratory is a structural HANK model, rich enough to feature many of the frictions popular in the quantitative business-cycle literature (e.g., sticky prices and wages, variable capacity utilization, and investment adjustment costs, as in Smets & Wouters, 2007). Models of this sort are routinely used for policy evaluation, and so in particular are a natural candidate for a fully structural solution to the aggregation (missing intercept) problem.

I build on the general framework of Section 2.1 and continue to impose Assumptions 1 and 2, but now relax Assumption 3. Demand equivalence in this generalized environment thus fails *only* because of the labor supply channel.[18]

Proposition 3. *Consider a stimulus check policy $\boldsymbol{\varepsilon}_\tau$, and suppose that Assumptions 1 and 2 hold. Then, for a fiscal spending policy $\boldsymbol{\varepsilon}_g$ such that (i) $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}_\tau^{PE}$ (identical net excess demand) and (ii) $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$ (identical* direct *tax response), we have that, to first order,*

$$\widehat{\mathbf{c}}_\tau \quad = \quad \widehat{\mathbf{c}}_\tau^{PE} \quad + \quad \widehat{\mathbf{c}}_g \quad + \quad \textit{error}\left(\widehat{\boldsymbol{\ell}}_\tau^{PE}\right) \tag{14}$$

where the error function is characterized in Appendix G.3 and is equal to $\mathbf{0}$ if $\widehat{\boldsymbol{\ell}}_\tau^{PE} = \mathbf{0}$.

My choice to only relax Assumption 3 is motivated by previous work: contributions to the quantitative business-cycle literature rarely depart from the common-goods assumption and feature households borrowing and lending in government bonds, but usually do not impose Assumption 3 (for canonical examples see Christiano et al., 2005; Smets & Wouters, 2007).

Model estimation. I provide only a brief outline of the model and my estimation strategy here, and relegate further details to Appendix B.2.

The household block is as described in Section 2.1, with preferences specialized to be of a standard separable form. The rest of the economy is designed to be as close as possible to the canonical model of Justiniano et al. (2010). First, I allow for investment adjustment costs, variable capacity utilization, and a rich monetary policy rule. Second, I extend the economy to be subject to a standard menu of aggregate shocks: to total factor productivity and the marginal efficiency of investment, to household patience, to wage mark-ups, to government

[18] In stating Proposition 3, I have relaxed the equal financing assumption to one of equal *direct* financing, where the direct tax response is defined analogously to Definition 1. With Assumption 2 and $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}_\tau^{PE}$, such equal direct financing is still feasible. Identical *overall* financing—i.e., $\widehat{\boldsymbol{\tau}}_g^e = \widehat{\boldsymbol{\tau}}_\tau^e$—however is generally not feasible without Assumption 3. This is because differences in general equilibrium feedback imply that the *other* inputs to the fiscal budget may not respond identically to the two shocks.

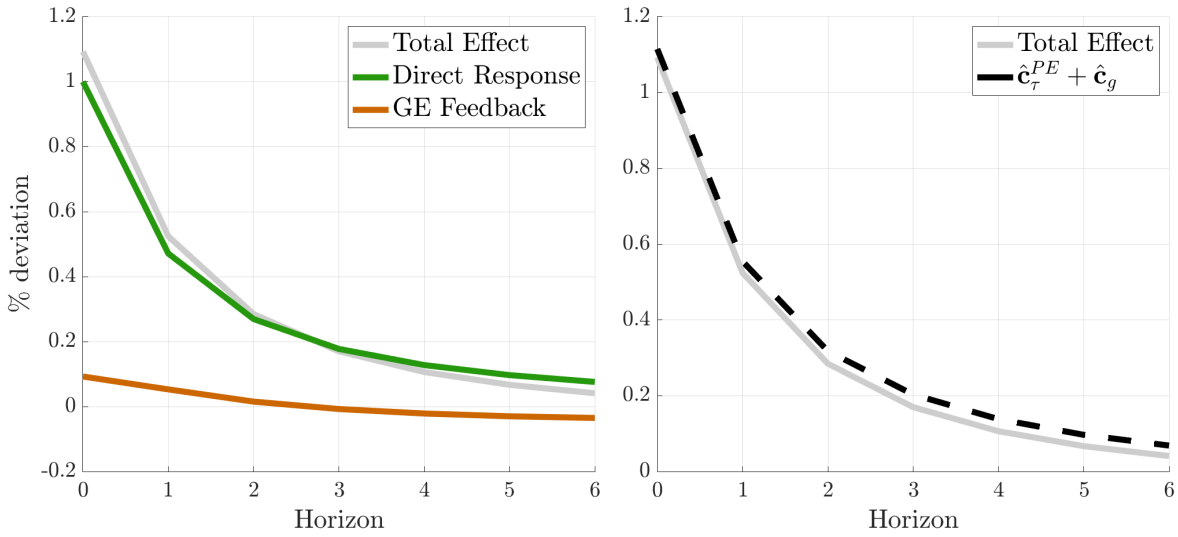

Figure 2: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, with details on the parameterization in Appendix B.2. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

spending, and to monetary policy. The only purpose of these additional shocks is to allow the model to fit aggregate U.S. business-cycle dynamics reasonably well, opening the door for a conventional likelihood-based estimation approach (An & Schorfheide, 2007). 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 likelihood methods on a standard set of macroeconomic aggregates.[19] 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. Most of the results in the remainder of this section refer to the model's posterior mode parameterization.

Main results. I subject the economy to a stimulus check policy shock, and then consider a researcher that tries to use the methodology of Section 3 to estimate that policy's aggregate

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

causal effects.[20] Results are displayed in Figure 2: the left panel decomposes the response of aggregate consumption to the stimulus check into direct partial equilibrium (green) and indirect general equilibrium (orange) effects, while the right panel compares the actual model-implied causal effect (grey) with the output of my procedure (dashed black),

The main take-away from the results in Figure 2 is that the demand equivalence approximation remains excellent, with the grey and black dashed lines in the right panel close to each other throughout.[21] The left panel first of all reveals that general equilibrium effects in the estimated model are rather small, reflecting largely offsetting interest rate, tax financing, and Keynesian amplification effects. Following a similarly short-lived and deficit-financed fiscal spending expansion, the same forces imply that aggregate consumption barely moves, giving the small approximation error displayed in the right panel. The intuition for the sign and magnitude of that approximation error is simple. Following receipt of the stimulus check, households consume more. Given their lower marginal utility of consumption, they would optimally like to work less, thus in general equilibrium depressing aggregate output and consumption. This labor supply channel is absent after an increase in (unvalued) fiscal purchases, so the demand equivalence approximation *overstates* the response of consumption to the transfer stimulus. However, even with quite strong wealth effects, this channel is largely irrelevant quantitatively: as long as wages are at least moderately sticky, labor is mostly demand-determined in the short run, so transitory shifts in labor supply do not matter much. This finding is consistent with conventional wisdom in the business-cycle literature (e.g. Christiano, 2011a,b): at least for relatively transitory fluctuations, hours worked in conventional (New Keynesian) business-cycle models are largely demand-determined.

Extensions & other models. The results in Figure 2 are neither special to the posterior mode of my model, nor to the particular HANK setting considered here.

First, in Appendix D.1, I *randomly* draw model parameters from large supports, solve the implied model, and compute the approximation accuracy. The analysis reveals that only the degrees of price and wage rigidity—and so the extent to which labor is demand-determined—have a material impact on the accuracy of the approximation, exactly as expected. Second, in Appendix D.2, I use the demand equivalence approximation to construct counterfactuals

[20] I compute the response of the endogenous component of taxes to the stimulus check policy, $\widehat{\boldsymbol{\tau}}^e_\tau$, using the particular rule (B.7). I then set $\widehat{\boldsymbol{\tau}}^e_g \propto \widehat{\boldsymbol{\tau}}^e_\tau$, with the factor of proportionality chosen so that $\lim_{t\to\infty} \widehat{b}_t \to 0$ after the fiscal spending shock $\boldsymbol{\varepsilon}_g$. With Assumption 3 this specification would ensure identical *overall* tax financing, exactly as in Proposition 1. I maintain this specification of fiscal rules for all of Section 4.

[21] At its largest, the associated error equals just above three per cent of the true peak consumption response.

for private demand shocks in the popular model of Justiniano et al. (2010). Since wages are even stickier there, the approximation is in fact better than in my estimated HANK model, with the approximation error now barely visible.

Discussion. The analysis in this section has demonstrated that demand equivalence is, at least approximately, a feature of standard quantitative models of business-cycle fluctuations. Intuitively, the features typically added to such models to ensure agreement with time series aggregates—e.g., investment adjustment costs, variable capacity utilization, or price- and wage-indexing—are all entirely consistent with demand equivalence. In particular, the common goods and financing assumptions are imposed regularly, and the labor supply channel is generally found to be quantitatively insignificant. We can thus materially strengthen the conclusions of Section 3: we now know that even popular, quantitatively relevant business-cycle models—though they break *exact* demand equivalence—will still robustly yield private demand shock counterfactuals that are almost perfectly governed by a) the shock's direct effect on private spending and b) aggregate fiscal multipliers.

While promising, this result is however only a first step to gauging the empirical relevance of the demand equivalence approximation. All three assumptions required for Proposition 1—and not just the labor supply restriction—are likely to be violated in practice, so I now extend the baseline HANK model in several directions to understand better why and how the approximation can fail.

4.2 Breaking equivalence

In this section I consider a large number of model extensions, each designed to challenge the quality of the demand equivalence approximation by breaking one or several of Assumptions 1 to 3. For each model variant I begin with the baseline HANK model studied in Section 4.1, and then add a further friction that is inconsistent with demand equivalence. Model details, the calibration strategy, and full results are presented in Appendices B.3 and D; here I just focus on sign and size of the induced bias as well as the core economic intuition.

My main results are reported in Figure 3, which plots the demand equivalence error for each of my experiments, with this error defined as

$$\text{error} \quad = \quad \frac{(\widehat{\mathbf{c}}_\tau^{PE} + \widehat{\mathbf{c}}_g) - \widehat{\mathbf{c}}_\tau}{\widehat{c}_{\tau 0}} \tag{15}$$

Note that (15) does *not* normalize the approximation errors to be positive. The fact that the

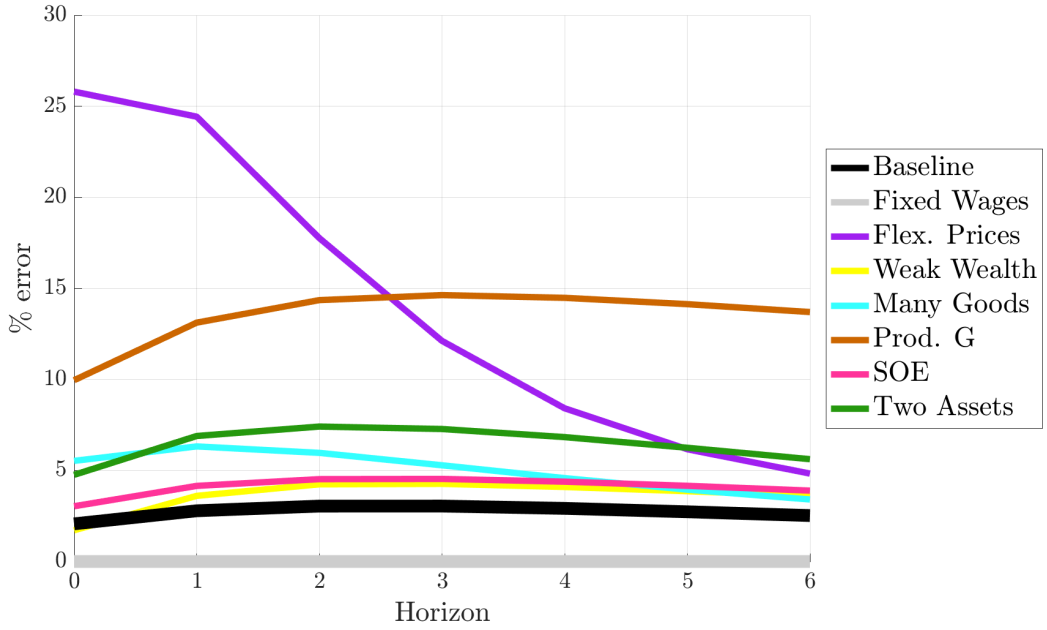

Figure 3: Errors (relative to the true impact consumption response) of the demand equivalence approximation in several model extensions. Details for all extensions and their parameterizations are relegated to Appendix B.

errors displayed in Figure 3 all turn out to be positive is thus not an artifact of normalization, but in fact a key result.

Labor supply & wealth effects. Four experiments—the baseline HANK model, fixed wages, a model with flexible prices and wages, and a model with household preferences that imply weak wealth effects—illustrate the role of Assumption 3 on labor supply in breaking demand equivalence. As shown previously in Figures 1 and 2, in the estimated HANK model, the demand equivalence approximation is very accurate even with moderately sticky wages (black), and exact in the case of fully rigid wages (grey). The purple line in Figure 3 shows that, with (nearly) flexible prices and wages, the quality of the approximation deteriorates sharply: because of quite strong wealth effects in labor supply, households cut hours worked following transfer receipt, and so the demand equivalence approximation—which misses these wealth effects—significantly overstates the response of aggregate consumption.

How material is this particular threat to the demand equivalence approach? I have already emphasized that, for relatively transitory shocks (such as one-off stimulus checks),

even moderately sticky wages are enough to mute the labor supply channel. Other pieces of macro and micro evidence suggest the same conclusion. First, on the macro side, standard time series estimation usually calls for near-zero wealth effects in labor supply (e.g., Schmitt-Grohé & Uribe, 2012; Born & Pfeifer, 2014; Bayer et al., 2022). Second, on the micro side, quasi-experimental evidence at the household level suggests that, at least in response to moderately sized lump-sum transfers, hours worked and earnings drop by an order of magnitude less than spending increases (e.g., Cesarini et al., 2017).[22] The yellow line in Figure 3 shows that, with fully flexible wages but household preferences taking a non-standard form to feature such weak wealth effects, the approximation error is again small.

Many goods. Heterogeneity in public and private consumption baskets is a further obvious threat to demand equivalence: without the restriction of a common final good (Assumption 1), changes in public and private purchases may set in motion very different general equilibrium effects. First, relative prices will move in response to sectoral spending shocks (Ramey & Shapiro, 1998). Second, if goods differ in their factor incidence (e.g., capital vs. labor income), and if factor income covaries with household characteristics (e.g., households with little non-labor income have high MPCs), then general equilibrium effects will necessarily be shock-specific (Alonso, 2017; Baqaee, 2015).

To gauge the importance of these channels, I construct the demand equivalence approximation in a multi-good model in which: goods differ in their labor intensity; real relative prices fluctuate in response to (sectoral) shocks; and government expenditure is concentrated on the relatively more labor-intensive good. The cyan line in Figure 3 shows that these model extensions further reinforce the (still-present) wealth effect baseline error, with the positive bias now even more pronounced. The logic is as follows. First, the real relative price of the private consumption bundle naturally increases by more after stimulus checks than after an increase in fiscal purchases. Demand equivalence thus misses one channel of general equilibrium crowding-out. Second, since in my model MPCs out of labor income exceed those out of capital income, fiscal purchases have larger general equilibrium multipliers.

While positive throughout, the error is again moderate, at around double of the baseline

[22] Coibion et al. (2020) document similarly small earnings responses after stimulus payments in the COVID-19 recession. Mogstad et al. (2021) estimate marginal propensities to earn that are roughly twice as large, at around 3 per cent (see the discussion in their Appendix J); relative to MPCs of the size considered here ($\approx$ 30 per cent), wealth effects of this magnitude are still immaterial. However, for larger income gains (and thus smaller MPCs), the bias coming from household labor supply responses is likely to be more significant. Consistent with this observation, I in Section 4.3 recommend that researchers should only apply the demand equivalence approximation to moderately sized spending shocks.

model. First, even with prices adjusting every 2.5 quarters on average, transitory spending shocks induce only moderate relative price fluctuations. Since the price elasticity of consumer demand in my model is furthermore relatively small, it follows that the price channel is almost completely irrelevant.[23] Second, and consistent with Alonso (2017), Baqaee (2015), as well as Flynn et al. (2022), I find that plausible differences in MPCs and factor incidence are not enough to yield sizable differences in total multipliers. In the data, the average consumption good has a labor share of around of 0.4, while the network-adjusted labor share of government consumption is around 0.65. Even assuming a quarterly MPC out of labor income of around 0.4, and an MPC out of any residual income of only 0.1, the resulting second-round demand difference would only be around 7.5 cents for a dollar of spending.[24]

Productive government purchases. As a second violation of the common-good assumption, I extend my baseline HANK model to allow for productive benefits of government spending, with the stock of government "capital" $k_t^g \equiv (1-\delta)k_{t-1}^g + g_t$ directly entering the production function of firms. I calibrate the model to match empirical evidence on public investment multipliers (Leduc & Wilson, 2013; Gechert, 2015).

The orange line in Figure 3 reveals that productive benefits of government purchases can quite materially undermine the quality of the demand equivalence approximation. The approximation error is positive throughout, reflecting the fact that government purchases increase the economy's productive capacity and so crowd-in consumption—an amplification channel missing for transfer stimulus. In fact, since the productive benefits are long-lived, the error remains quite persistently elevated even for transitory shocks.

Open economy. As a third violation of the common-good assumption, I consider an open-economy version of my HANK model. In this environment, private consumption purchases in response to stimulus checks partially leak abroad, while government purchases are assumed to fall exclusively on domestic goods. By equating government purchases $\widehat{\mathbf{g}}_g$ and private expenditure on the *domestic* good $\widehat{\mathbf{c}}_\tau^{H,PE}$, the demand equivalence approach can still ensure identical net excess demand for the domestic good; however, because of demand leakage, the

[23] Orchard et al. (2022) present a model with large relative price responses as well as large price elasticities of consumer demand, opening the door for stronger general equilibrium crowding-out. I further discuss the relationship between my results and theirs in Appendix F.2.

[24] Arguably, 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 3 shrinks.

government purchases $\widehat{\mathbf{g}}_g$ induce a strictly smaller deficit (in net present value terms) than the stimulus check policy $\boldsymbol{\varepsilon}_\tau$, thus breaking demand equivalence.

The pink line in Figure 3 shows the approximation error for an open economy with a home bias of 0.89 (matching the U.S.). As expected, openness increases the approximation error relative to the baseline economy: because of the lack of demand leakage, government spending is cheaper than the equivalent stimulus check, so taxes rise by less, leading to less general equilibrium crowding-out. However, given the substantial degree of home bias, it is not surprising that the error remains close to the baseline model throughout.

Interest rates. The third assumption required for demand equivalence is that of identical interest rates for household and government (Assumption 2). This assumption is necessarily violated in models with multiple savings vehicles, such as the two-asset model of Kaplan & Violante (2014) and Kaplan et al. (2018). In contrast to the various other sources of bias, it is not clear ex-ante in which direction a violation of this assumption will bias the approximation: if household returns are high relative to government returns (e.g., due to credit card debt or savings in equity), then taxes need to increase by less to finance private relative to public spending, and so the approximation will be biased downward. Conversely, if returns are low (e.g., due to bank intermediation), then the bias is positive.

To get a sense of the likely magnitude of the implied approximation error, the green line in Figure 3 shows results for a two-asset model in which households pay an intermediation fee on liquid deposits, giving a positive bias and thus reinforcing the always-present labor supply error. I find that, even for a *quarterly* return gap of 1.25 per cent, the error remains rather moderate, peaking at around 7.5 per cent of the true impact consumption response. To see why, suppose that, in response to a transfer stimulus policy, 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 public spending. Even if annual household and government discount rates differ by $4 \times 1.25 = 5$ per cent, the difference in present discounted values of the two spending expansions is just five cents—small compared to the initial size of the stimulus. The implied difference in tax financing is thus also small, and so the approximation remains quite accurate. Thus, even for large return gaps (in either direction), the bias is comparable in magnitude to the (small) labor supply-related bias.

4.3 Summary and recommendations for practice

The results in Sections 4.1 and 4.2 shed light on the appeals and limitations of the proposed demand equivalence approach. Key necessary conditions for its accuracy include: transitory and relatively small shocks, ensuring that wealth effects in labor supply as well as real relative price movements are indeed negligible; a fiscal time series experiment that does not pick up productive government investment; and a relatively closed economy, or more generally a private spending shock that verifiably fell largely onto domestically produced goods. Interest rate effects or sectoral heterogeneity in spending multipliers on the other hand appear somewhat less likely to materially threaten the accuracy of the approximation. Thus, if these necessary conditions are satisfied, then the output of the demand equivalence approximation can be interpreted as a rather tight *upper bound* to the actual general equilibrium response of private spending to the private demand shifter.

To summarize, while substantial care is necessary in applying the demand equivalence approach, I have also argued that it *can* be highly informative under the right circumstances. In particular, stimulus checks—the main application of Section 3, and a topic of much policy interest—appear rather well-suited: the stimulus is relatively short-lived, and wealth effects are estimated to be quite weak; the U.S. economy is relatively closed; and in Appendix F.1 I made sure to use fiscal spending experiments that do not capture government investment.

Finally, I emphasize that my conclusions in this section are also informative for researchers who wish to use structural models to solve the "missing intercept" aggregation problem. To the extent that a structural modeler finds a missing intercept path far from zero, we know from the results in this paper that this finding must be tied either to fiscal multipliers far from one—if the model is close to standard business-cycle modeling practice—or to departures from demand equivalence, with Section 4.2 providing a list of the most important ones. I hope that these insights will prove useful in relating and interpreting the results from model-based aggregation exercises in various existing studies (e.g. Kaplan & Violante, 2018; Auclert & Rognlie, 2018; Auclert et al., 2019; Orchard et al., 2022).

5 Extensions

My theoretical and empirical analysis so far has been largely restricted to the stimulus check application. This section discusses the wider scope of the demand equivalence approach. First, in Section 5.1, I give examples of other shocks and policies covered by the consumption equivalence result. Second, in Section 5.2, I present the conditions required for an analogous

investment demand equivalence result, and then use the result to discuss the likely aggregate effects of another well-known fiscal stimulus policy: investment bonus depreciation.

5.1 Other consumer spending shocks

As argued in Section 2.3, the consumption demand equivalence result—and so the measurement strategy of Section 3—applies to *any* shifter of private household spending, not just uniform stimulus checks. In this section I discuss two examples. First, as another instance of a policy-induced shifter, I consider stimulus checks targeted at certain sub-populations, consistent with recent U.S. policy design. Second, I study the effects of a short-lived increase in income inequality as an example of a non-policy shifter of household spending.

Targeted Transfers. Consider a one-off stimulus check policy ε_τ targeted at some sub-population $\mathcal{T} \subseteq [0,1]$ of households. Proceeding analogously to the discussion in Section 3.3, we get the direct consumption response as

$$\widehat{C}^{PE}_{\tau t} \equiv \underbrace{|\mathcal{T}|}_{\text{\# of recipients}} \times \underbrace{MPC^{\mathcal{T}}_{t,0}}_{\text{MPC of recipients}} \times \underbrace{\widehat{\tau}_0}_{\text{check size}}$$

where $|\mathcal{T}| \equiv \int_{i\in\mathcal{T}} di$ and

$$MPC^{\mathcal{T}}_{t,0} \equiv \frac{1}{|\mathcal{T}|}\int_{i\in\mathcal{T}} \frac{\partial c_{it}}{\partial \tau_0} di$$

The direct response is thus estimable using information on household MPCs in the targeted sub-group, and so general equilibrium counterfactuals can be estimated as in Section 3.3. The demand equivalence approach can thus be used to gauge the effects of stimulus check policies even if the desired targeting has no historical precedent—we only need the corresponding demand path to have been studied before in a time series experiment.

General private demand shifters. Examples of much-studied non-policy shifters of private consumer demand include a tightening of financial constraints (e.g. Guerrieri & Lorenzoni, 2017), changes in income inequality (e.g. Auclert & Rognlie, 2018) and changing tastes, e.g. due to infection risk associated with the consumption of certain goods (Beraja & Wolf, 2020). Such shocks are covered by the more general decomposition in (8): they induce some *zero net present value* change $\widehat{\mathbf{c}}^{PE}_d$ of consumption demand, and so full general equilibrium counterfactuals can be constructed with knowledge of $\widehat{\mathbf{c}}_g$ for a *purely deficit-financed* fiscal spending experiment with $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}^{PE}_d$—i.e., a change in spending today that is financed

with a spending reversal in the future. I present an application to a transitory increase in labor income inequality in Appendix F.4, following Auclert & Rognlie (2018).

5.2 Investment

The demand equivalence logic can also be leveraged to estimate general equilibrium counterfactuals for shifters of *investment* demand. In this section I first sketch the conditions required for investment demand equivalence and then discuss an application to bonus depreciation stimulus. Details for theory and application are provided in Appendices A.2 and F.3.

THEORY: INVESTMENT DEMAND EQUIVALENCE. I again use the model of Section 2.1. Anticipating the empirical application, I augment the model to feature investment tax credit shocks $\boldsymbol{\varepsilon}_q$—shocks that reduce the cost of capital purchases by intermediate goods producers at time t by an amount $\tau_{qt} = \tau_{qt}(\boldsymbol{\varepsilon}_q)$.[25] I define direct (partial equilibrium) responses and indirect (general equilibrium) feedback for firm investment exactly analogously to Definition 1, using the implied aggregate investment function $\mathbf{i}(\bullet)$. As before, the question is: under what restrictions on primitives does the response of investment to a suitably chosen fiscal spending experiment give the "missing intercept" $\widehat{\mathbf{i}}_q^{GE}$?

The proof strategy is identical to that in Section 2: I characterize equilibrium response paths as a system of market-clearing conditions, and then impose enough restrictions on this system to ensure that the investment tax credit as well as a suitable fiscal experiment perturb the same equations by the same amount. In my model, the investment tax credit policy has three direct effects. First, investment responds; since investment invariably boosts the future productive capacity of the economy, production also increases, so the induced partial equilibrium net excess demand path for the final output and investment good is $\widehat{\mathbf{i}}_q^{PE} - \widehat{\mathbf{y}}_q^{PE}$. Second, the policy may be redistributive: the cost of financing is borne by taxed households, but the benefits accrue to households receiving dividend payments. The two groups need not be the same. Third, more investment and so more capital will increase the marginal product of labor, pushing up firm labor demand.

Matching the first effect through a fiscal shock is conceptually straightforward: under a common-good assumption, we simply need to consider a fiscal spending expansion with

$$\widehat{\mathbf{g}}_g = \widehat{\mathbf{i}}_q^{PE} - \widehat{\mathbf{y}}_q^{PE} \tag{16}$$

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

For the other two effects I require additional exclusion restrictions. To rule out heterogeneous distributional implications of tax financing and dividend payments, I assume that household income risk is perfectly insurable, thus effectively imposing a representative-household structure. This restriction also implies that Ricardian equivalence holds, so the precise timing of the policy financing is irrelevant. Next, to ignore the labor demand response, I assume that labor supply is perfectly flexible, either because of a large Frisch elasticity of labor supply, or again because labor is fully demand-determined. Under those two additional restrictions on primitives, a fiscal experiment satisfying (16) indeed gives the "missing intercept" of the investment response, in the precise sense that

$$\widehat{\mathbf{i}}_q \quad = \quad \widehat{\mathbf{i}}_q^{PE} \quad + \quad \widehat{\mathbf{i}}_g \tag{17}$$

Analogously we can also recover the output response as

$$\widehat{\mathbf{y}}_q \quad = \quad \widehat{\mathbf{y}}_q^{PE} \quad + \quad \widehat{\mathbf{y}}_g \tag{18}$$

I formally state the equivalence result and its assumptions in Appendix A.2. Crucially, the proposition requires no additional restrictions on the production side of the economy: firms can face a rich set of real and financial frictions, including (convex and non-convex) capital adjustment costs as well as a generic set of constraints on equity issuance and borrowing. A detailed discussion of nested (heterogeneous-firm) models is provided in Appendix C.5.

Application: bonus depreciation. I apply the investment demand equivalence result to estimate general equilibrium counterfactuals for bonus depreciation stimulus—that is, the ability to tax-deduct investment expenditure at a faster rate.[26] I focus on bonus depreciation for three reasons. First, it is popular: when conventional monetary policy is constrained by an effective lower bound, bonus depreciation has arguably become the go-to tool for investment stimulus. Second, previous empirical work has leveraged heterogeneity in firm exposure to the stimulus to estimate its direct effect on investment, $\widehat{\mathbf{i}}_q^{PE}$ (Zwick & Mahon, 2017; Koby & Wolf, 2020)—the key empirical input needed for my approach. Third, given the endogeneity of bonus depreciation policies to wider macroeconomic conditions, I am not aware of any studies that credibly identify the aggregate causal effects of such policies.

I only provide a brief summary of the results here, with details provided in Appendix F.3.

[26] 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 (see Winberry, 2018).

Overall, my findings closely echo those of Section 3.3. Micro data indicate a large and persistent response of investment to bonus depreciation stimulus of the size seen in 2008. By the investment demand equivalence result, solving the aggregation problem thus requires the researcher to take a stand on the aggregate effects of a rather persistent and deficit-financed increase in government purchases, again with little monetary offset. My own empirical implementation again suggests a fiscal multiplier around 1, implying that 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.

6 Conclusion

How can researchers learn about the "missing intercept" of cross-sectionally identified shifters of private spending? In this paper I ask whether—consistent with a simple Keynesian cross intuition—time series estimates of aggregate fiscal expenditure shocks can solve this aggregation problem. I give a set of restrictions on economic primitives under which aggregation via such "demand equivalence" is indeed possible, show how to operationalize this result, and discuss its empirical plausibility. In an application to deficit-financed stimulus checks, I find that cross-sectionally identified spending estimates are likely to be close to full general equilibrium counterfactuals, corresponding to a missing intercept close to zero.

I leave several avenues for future research. First, to be widely applicable, the demand equivalence approach requires time series estimates for a wide *menu* of fiscal spending experiments with different persistence and financing. More research on fiscal multipliers is thus welcome—it promises to not only tell us narrowly about those fiscal experiments, but also about the propagation of many other demand-type shocks and policies. Second, other interesting macro shocks covered by the demand equivalence approach 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). The methodology developed here could be applied to all of them. Third, I emphasize that the general conceptual idea of this paper—to leverage equivalence in the general equilibrium propagation of different shocks and policies—is not necessarily limited to fiscal spending and demand amplification, and so may help to solve the "missing intercept" problem in other contexts as well.

A Appendix

A.1 Proof of consumption demand equivalence

I begin by writing the equilibrium of the baseline model (but with Assumption 1 imposed) as a dynamic system of market-clearing conditions.

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

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

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

Proof. See Appendix G of the Online Appendix. □

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

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

where J_i denotes the infinite-dimensional generalization of the selection matrix selecting the ith entry of a vector x_t. Assuming equilibrium existence and uniqueness,[27] there exists a unique linear map $\mathcal{H}$ such that

[27] Existence and uniqueness of a bounded transition path for representative-agent models can be shown as usual. For the heterogeneous-agent models, I have verified existence and uniqueness for particular numerical examples, using either the conditions of Blanchard & Kahn (1980) or invertibility of the perfect-foresight general equilibrium adjustment matrix (see Auclert et al., 2019).

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

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

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

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

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

The impulse response paths of consumption thus satisfy

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

respectively. By Assumption 2 we know that, if $\widehat{\mathbf{c}}_\tau^{PE} = \widehat{\mathbf{g}}_g$, then setting $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$ is consistent with $\lim_{t\to\infty} \widehat{b}_t = 0$, since $\widehat{\boldsymbol{\tau}}_\tau^x$ and $\widehat{\mathbf{g}}_g$ by construction have the same net present value.[28] This establishes that, if $\widehat{\mathbf{c}}_\tau^{PE} = \widehat{\mathbf{g}}_g$ and $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$, then $\widehat{\mathbf{c}}_\tau^{GE} = \widehat{\mathbf{c}}_g$. (6) then follows. Finally note that, given the assumed fiscal financing rule $\boldsymbol{\tau}^e(\bullet)$, $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$ also implies $\widehat{\boldsymbol{\tau}}_g^e = \widehat{\boldsymbol{\tau}}_\tau^e$—the stated property (ii) of the fiscal spending shock. □

[28] By assumption $\widehat{\mathbf{c}}_\tau^{PE} = \widehat{\mathbf{g}}_g$. Furthermore it follows from the household budget constraint that we must have $\sum_{t=0}^\infty \left(\frac{1}{1+\bar{r}^b}\right)^t \widehat{c}_{\tau,t}^{PE} = \sum_{t=0}^\infty \left(\frac{1}{1+\bar{r}^b}\right)^t \widehat{\tau}_{\tau,t}^x$. Combining the two: $\sum_{t=0}^\infty \left(\frac{1}{1+\bar{r}^b}\right)^t \widehat{g}_{g,t} = \sum_{t=0}^\infty \left(\frac{1}{1+\bar{r}^b}\right)^t \widehat{\tau}_{\tau,t}^x$.

A.2 Details on investment demand equivalence

I begin with the restrictions needed for an exact investment demand equivalence result. The first assumption is again that of a single common final good.

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

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

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

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

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

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

These are the three material restrictions discussed in Section 5.2. Finally, I require an additional restriction on monetary policy. If the monetary authority directly responds to the *level* of output, then the increase in production associated with the investment subsidy will induce a contractionary monetary response. One simple sufficient condition for ruling this out is that the monetary authority does not respond to fluctuations in output.[29]

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

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

Proposition A.1. *Consider an investment stimulus $\boldsymbol{\varepsilon}_q$, and suppose that Assumptions A.1 to A.4 hold. Then, for a fiscal spending policy $\boldsymbol{\varepsilon}_g$ such that $\widehat{\mathbf{g}}_g = \widehat{\mathbf{i}}_q^{PE} - \widehat{\mathbf{y}}_q^{PE}$, we have that, to first order,*

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

Proof. See Appendix G of the Online Appendix. □

The proof reveals that all results extend to generic investment "wedges" (Chari et al., 2007).

[29] More generally, a researcher leveraging my methodology simply needs to communicate that her counterfactual assumes that same interest rate response as that observed after the fiscal shock used for aggregation.

References

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

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

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

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

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

Angeletos, G.-M. & Lian, C. (2022). Determinacy without the taylor principle. Technical report, National Bureau of Economic Research.

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

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

Auclert, A., Bardóczy, B., & Rognlie, M. (2020). MPCs, MPEs and multipliers: A trilemma for new keynesian models. Technical report, National Bureau of Economic Research.

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. (2018). Inequality and aggregate demand. Technical report, National Bureau of Economic Research.

Auclert, A., Rognlie, M., Souchier, M., & Straub, L. (2021). Exchange rates and monetary policy with heterogeneous agents: Sizing up the real income channel. Technical report, National Bureau of Economic Research.

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

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

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

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

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

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

Baugh, B., Ben-David, I., Park, H., & Parker, J. A. (2021). Asymmetric consumption smoothing. *American Economic Review*, *111*(1), 192–230.

Bayer, C., Born, B., & Luetticke, R. (2022). Shocks, frictions, and inequality in us business cycles.

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

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

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

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

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

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

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

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

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

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

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

Broda, C. & Parker, J. A. (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics*, *68*, S20–S36.

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

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

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

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

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

Christiano, L. 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.

Cochrane, J. H. (2011). Determinacy and identification with taylor rules. *Journal of Political economy*, *119*(3), 565–615.

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

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

Drautzburg, T. (2020). A narrative approach to a fiscal dsge model. *Quantitative Economics*, *11*(2), 801–837.

Dupaigne, M. & Fève, P. (2016). Persistent government spending and fiscal multipliers: The investment-channel. *European Economic Review*, *89*, 425–453.

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

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

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

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

Flynn, J. P., Patterson, C., & Sturm, J. (2022). Fiscal policy in a networked economy. NBER Working Paper.

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.

Ganong, P. & Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American economic review*, *109*(7), 2383–2424.

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

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

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

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

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

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

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

Hausman, J. K. (2016). Fiscal policy and economic recovery: The case of the 1936 veterans' bonus. *American Economic Review*, *106*(4), 1100–1143.

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

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

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

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

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

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

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

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

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

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

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

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

Kueng, L. (2018). Excess sensitivity of high-income consumers. *The Quarterly Journal of Economics*, *133*(4), 1693–1751.

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., Walker, T. B., & Yang, S.-C. S. (2013). Fiscal foresight and information flows. *Econometrica*, *81*(3), 1115–1145.

Li, D., Plagborg-Møller, M., & Wolf, C. K. (2021). Local projections vs. vars: Lessons from thousands of dgps. arXiv preprint arXiv:2104.00655.

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

McKay, A. & Wolf, C. K. (2022). What can time-series regressions tell us about policy counterfactuals? Technical report, MIT Working Paper.

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.

Mogstad, M., Golosov, M., Graber, M., & Novgorodsky, D. (2021). How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (76).

Mountford, A. & Uhlig, H. (2009). What are the effects of fiscal policy shocks? *Journal of applied econometrics*, *24*(6), 960–992.

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.

Orchard, J., Ramey, V. A., & Wieland, J. (2022). Micro mpcs and macro counterfactuals: The case of the 2008 rebates. Technical report, Working Paper.

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

Parker, J. A., Souleles, N. S., Johnson, D. S., & McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, *103*(6), 2530–53.

Plagborg-Møller, M. & Wolf, C. K. (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica*, *89*(2), 955–980.

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

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

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

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

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

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

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

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

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

Sims, C. A. & Zha, T. (1995). Does monetary policy generate recessions? 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.

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.

Wolf, C. K. (2021). Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result. Working Paper.

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

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

This online appendix contains supplemental material for the article "The Missing Intercept: A Demand Equivalence Approach". I provide: (i) further details for the various structural models used in the paper; (ii) several additional results on exact demand equivalence, supplementing the theoretical analysis in Sections 2 and 5; (iii) results on approximation accuracy in structural models where demand equivalence fails, elaborating on my discussion in Section 4; and (iv) details on the cross-sectional and time series evidence that is then used for (v) my empirical applications. The end of this appendix contains several further proofs as well as auxiliary lemmas.

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

Contents

B Model details

This appendix provides additional details on the structural models considered in the main body of the paper. In Appendix B.1 I begin by outlining the full baseline model and offering a formal definition of equilibrium transition paths. Appendix B.2 then discusses the baseline HANK model of Section 4.1. Finally, in Appendix B.3, I give modeling details for the various extensions considered in Section 4.2.

B.1 Rest of the baseline economy and equilibrium definition

Recall that the model is populated by households, firms, and the government. Whenever there is no risk of confusion, I replace the full decision problems of agents by simple conditions characterizing their optimal behavior.[30] Since for much of the paper I impose the one-good restriction of Assumption 1, I here present the equilibrium for this baseline case, and relegate a discussion of the notationally involved multi-good extension to Appendix B.3.1.

Households & unions. The household consumption-savings problem was described in Section 2.1. Since I for now consider a simpler one-good economy, we have that $p_t^c = 1 \ \forall t$. For all simulations I specialize household preferences to be of a standard separable form:

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

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

$$\begin{aligned} \pi_t^w(1+\pi_t^w) \ &= \ \frac{\epsilon_w}{\theta_w}\ell_t^h\Big[\int_0^1 \Big\{ -u_\ell(c_{it}, \ell_t^h) - \frac{\epsilon_w - 1}{\epsilon_w}(1-\tau_\ell) w_t e_{it} u_c(c_{it}, \ell_t^h)\Big\} di\Big] \\ &\quad + \beta \pi_{t+1}^w (1+\pi_{t+1}^w) \end{aligned} \tag{B.2}$$

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

[30] 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.

of labor, and θ_w denotes the Rotemberg adjustment cost. Given prices $(\boldsymbol{\pi}, \mathbf{w})$ as well as household consumption, (B.2) provides a simple restriction on total labor supply $\boldsymbol{\ell}^h$.[31] Note that, without idiosyncratic labor productivity risk and so common consumption $c_{it} = c_t$, the derived wage-NKPC (B.2) 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. Note that, in the special case of preferences as in (B.1), (B.2) simplifies to

$$\pi_t^w(1+\pi_t^w) = \frac{\epsilon_w}{\theta_w}\ell_t^h\left[\chi(\ell_t^h)^{\frac{1}{\varphi}} - \frac{\epsilon_w - 1}{\epsilon_w}(1-\tau_\ell)w_t\int_0^1 e_{it}c_{it}^{-\gamma}di\right] + \beta\pi_{t+1}^w(1+\pi_{t+1}^w) \tag{B.3}$$

and so, in log deviations,

$$\widehat{\tilde{\pi}}_t^w = \kappa_w \times \left[\frac{1}{\varphi}\widehat{\ell}_t^h - (\widehat{\tilde{w}}_t - \gamma\widehat{\tilde{c}}_t^*)\right] + \beta\widehat{\tilde{\pi}}_{t+1}^w \tag{B.4}$$

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

$$c_t^* \equiv \left[\int_0^1 e_{it}c_{it}^{-\gamma}di\right]^{-\frac{1}{\gamma}} \tag{B.5}$$

Together, the consumption-savings problem and the general wage-NKPC (B.2) characterize optimal household and union behavior. I assume that the solutions to each problem exist and are unique, and summarize the solution in terms of aggregate consumption, saving and union labor supply functions $\mathbf{c}(\mathbf{s}^h, \boldsymbol{\varepsilon})$, $\mathbf{b}^h(\mathbf{s}^h, \boldsymbol{\varepsilon})$, and $\boldsymbol{\ell}^h(\mathbf{s}^u)$, where $\mathbf{s}^h = (\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \boldsymbol{\ell}, \boldsymbol{\tau}^e, \mathbf{d})$ and $\mathbf{s}^u = (\boldsymbol{\pi}, \mathbf{w}, \mathbf{c})$.[32] In particular, the union problem gives

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

To state and prove the equivalence results, I will impose the high-level assumption that all of those functions are at least once differentiable in their arguments.

Firms. Recall that there are three types of firms: intermediate goods producers who make investment and labor hiring decisions (and who are subject to a very rich menu of real and

[31] In the special case $\theta_w \to \infty$, equation (B.2) is vacuous, so then I instead simply assume that $\boldsymbol{\ell}^h = \boldsymbol{\ell}^f$.

[32] Formally, the input to the union problem is the "virtual" consumption aggregate in (B.5). In a slight abuse of notation, the dependence on $\mathbf{c}$ in the equations here is a shorthand for dependence on overall household consumption decisions given $(\mathbf{s}^h, \boldsymbol{\varepsilon})$.

financial frictions), and retailers and aggregators whose sole purpose is to introduce nominal rigidities. As discussed throughout the paper, this structure is just rich enough to allow me to nest many canonical quantitative business-cycle models. I furthermore assume that all of those firms discount at the common rate $1 + r_t^b \equiv \frac{1+i_{t-1}^b}{1+\pi_t}$.

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

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

such that

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

The physical adjustment cost function $\phi(\bullet)$ is general: it may be convex and continuously differentiable, but it may also feature a fixed-cost component or partial irreversibility. Firms can also vary capital utilization at additional cost $a(u)k_{-1}$, as in Justiniano et al. (2010). The solution to the firm problem gives optimal production $\mathbf{y}(\bullet)$, labor demand $\boldsymbol{\ell}^f(\bullet)$, investment $\mathbf{i}(\bullet)$, intermediate goods producer dividends $\mathbf{d}^I(\bullet)$, capital utilization rates $\mathbf{u}(\bullet)$ and liquid corporate bond savings $\mathbf{b}^f(\bullet)$ as a function of nominal returns $\mathbf{i}^b$, inflation $\boldsymbol{\pi}$, wages $\mathbf{w}$, and the intermediate goods price $\mathbf{p}^I$.

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

$$\widehat{\tilde{\pi}}_t = \underbrace{\frac{\epsilon_p}{\theta_p} \frac{\epsilon_p - 1}{\epsilon_p}}_{\kappa_p} \times \widehat{\tilde{p}}_t^I + \frac{1}{1+\bar{r}^b} \widehat{\tilde{\pi}}_{t+1}$$

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

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

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

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

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

Total dividend payments by the corporate sector are given as

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

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

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

GOVERNMENT. I denote the fiscal financing rule by $\boldsymbol{\tau}^e = \boldsymbol{\tau}^e(\boldsymbol{w}, \boldsymbol{\ell}, \boldsymbol{i}^b, \boldsymbol{\pi}, \boldsymbol{\tau}^x, \boldsymbol{g})$. This fiscal rule must imply that, with debt evolving in accordance with the government budget constraint (2), we have $\lim_{t\to\infty} \widehat{b}_t = 0$. This in particular implies that the path $\boldsymbol{\tau}^e$ is such that the following log-linearized *lifetime* government budget constraint holds:

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\bar{r}^b}\right)^t \bar{b}\left(\widehat{\tilde{i}}_{t-1}^b - \widehat{\tilde{\pi}}_t\right) + \sum_{t=0}^{\infty}\left(\frac{1}{1+\bar{r}^b}\right)^t \bar{g}\widehat{\tilde{g}}_t + \sum_{t=0}^{\infty}\left(\frac{1}{1+\bar{r}^b}\right)^t \bar{\tau}\left(\frac{\bar{\tau}^e}{\bar{\tau}}\widehat{\tilde{\tau}}_t^e + \frac{\bar{\tau}^x}{\bar{\tau}}\widehat{\tilde{\tau}}_t^x\right)$$
$$= \sum_{t=0}^{\infty}\left(\frac{1}{1+\bar{r}^b}\right)^t \tau_\ell \bar{w}\bar{\ell}\left(\widehat{\tilde{w}}_t + \widehat{\tilde{\ell}}_t\right) \quad \text{(B.6)}$$

It remains to describe central bank behavior. In line with standard modeling practice I assume that the nominal rate on bonds i^b is set according to the conventional Taylor rule

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

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

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

Equating labor demand and supply:

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

Finally, aggregating all household, firm and government budget constraints, we obtain the aggregate output market-clearing condition for the single final good:[33]

$$c_t + i_t + g_t - \kappa_b \int_0^1 \mathbf{1}_{b_{it-1}^h<0} b_{it-1}^h di = y_t$$

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

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

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

3. Government. *The liquid nominal rate is set in accordance with the monetary authority's Taylor rule. The government spending, lump-sum tax, and debt issuance paths are jointly consistent with the government's budget constraint, its exogenous laws of motion for spending and discretionary transfers, and its financing rule* $\boldsymbol{\tau}^e(\bullet)$.

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

$$c_t + i_t + g_t - \kappa_b \int_0^1 \mathbf{1}_{b^h_{it-1}<0} b^h_{it-1} di = y_t$$

the bond market clears,

$$b^h_t + b^f_t = b_t$$

and the labor market clears,

$$\ell^h_t = \ell^f_t$$

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

B.2 Quantitative HANK model

Much of my analysis builds on a particular one-asset HANK model. 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.

MODEL OUTLINE. The model is a particular variant of the rich baseline environment outlined in Section 2.1, consistent with Assumptions 1 and 2 but violating Assumption 3.

Households have preferences as in (B.1). To facilitate comparison with the standard New Keynesian business-cycle literature, I will throughout replace the virtual consumption aggregate (B.5) in the wage-NKPC with aggregate consumption c_t, thus giving me an entirely standard wage-NKPC (as in Hagedorn et al., 2019); results are, however, almost unchanged if I use c^*_t instead.[34] 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 (so $\tau_{it} = \tau_t \ \forall i$) but are heterogeneous in dividend payment receipts. In particular, I assume that the most productive households receive larger dividend payments, so that stock market

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

wealth is effectively concentrated among a small share of households.

The intermediate goods production block—in particular the production function $y(\bullet)$, the investment adjustment cost function $\phi(\bullet)$, and the capacity utilization resource cost $a(\bullet)$—is set up following Justiniano et al. (2010). I furthermore assume that there are no firm-level financial frictions. For model estimation, I allow for structural shocks to output and investment productivity, monetary policy, government spending, household impatience, and wage mark-ups. All shocks are assumed to follow simple AR(1) processes. Finally I assume a fiscal financing rule of the form

$$\widehat{\tau}_t^e = -(1-\rho_\tau) \times \widehat{b}_{t-1} \tag{B.7}$$

The endogenous part of transfers is cut in response to increases in $\widehat{b}_t$. 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 (i.e., $\widehat{\boldsymbol{\tau}}_g^e \propto \widehat{\boldsymbol{\tau}}_\tau^e$). In particular, for all models in which Assumption 2 is imposed, this fiscal rule ensures that the partial equilibrium financing paths of the two shocks will always be the same (since $\widehat{\boldsymbol{g}}_g = \widehat{\boldsymbol{c}}_\tau^{PE}$).

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

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

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

where e_{it}^p is the permanent component of household i's labor productivity. I set the cutoff $\underline{e}^p$ so that the bottom half of households receive no dividends (consistent with the illiquid wealth distribution in the 2016 SCF), χ_1 so that the top 10 per cent of households receive the same

Parameter	Description	Value	Target	Model	Data
Households					
ρ_e, σ_e	Income Risk	-	Kaplan et al. (2018)	-	-
$\underline{e}^p$, χ_0, χ_1	Dividend Endowment	-	Illiquid Wealth Shares	-	-
β	Discount Rate	0.97	B/Y	1.04	1.04
$\bar{r}^b$	Average Return	0.01	Annual Rate	0.04	0.04
ξ	Death Rate	1/180	Average Age	45	45
γ	Preference Curvature	1	Standard		
φ	Labor Supply Elasticity	1	Standard		
ϵ_w	Labor Substitutability	10	Standard		
$\underline{b}$	Borrowing Limit	0	McKay et al. (2016)		
Firms					
α	Capital Share	0.2	Justiniano et al. (2010)		
δ	Depreciation	0.024	Total Wealth/Y	10.64	10.64
ϵ_p	Goods Substitutability	20	Profit Share	0.06	0.06
Government					
τ_ℓ	Labor Tax	0.3	Average Labor Tax	0.30	0.30
τ/Y	Transfer Share	0.05	Transfer Share	0.05	0.05
B/Y	Liquid Wealth Supply	1.04	Government Debt/Y	1.04	1.04

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

share of total dividends (and so total illiquid wealth) as in Kaplan et al. (2018), and then back out χ_0 to ensure that $\int_0^1 d_{it} di = d_t$.[35] Next, 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. For my baseline model, I further assume that households cannot borrow. All remaining parameters are set in line with conventional practice. The second block shows parameter choices on the firm side. I discipline the Cobb-Douglas production function $y = k^\alpha \ell^{1-\alpha}$ by setting α in line with Justiniano et al. (2010), identify goods substitutability by targeting the profit share, and finally back out the depreciation rate from my target of total wealth

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

(and so corporate sector valuation) in the economy as a whole.[36] The third block informs the fiscal side of the model. The average government tax take, transfers, and debt issuance are all set in line with direct empirical evidence.

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

As already mentioned in Footnote 6 I verify numerically the existence and uniqueness of the model's steady state. To see how this is done note that my calibration strategy *directly* pins down all parameters relevant for the steady state except for β. I then choose β to ensure asset market-clearing, given my target for total liquid wealth. I search across a large range of β's and always find a unique solution.

Dynamics: computational details. For my likelihood-based estimation I solve the model using a variant of the popular Reiter method (Reiter, 2009). In particular, I use a mixture of the methods developed in Ahn et al. (2017) and Bayer & Luetticke (2020) to reduce the dimensionality of the state space. Without dimensionality reduction, the number of idiosyncratic household-level states is too large to allow likelihood-based estimation. With dimensionality reduction, the number of states is reduced to around 300, making estimation feasible. My displays of exact and approximate demand equivalence are instead computed in sequence-space, as in Boppart et al. (2018) and Auclert et al. (2019). I construct all sequence-space Jacobians using simple finite-difference approximations.

As mentioned in Footnote 6 I verify numerically the existence and uniqueness of linearized transition paths. To do so I numerically check the Blanchard-Kahn conditions (when using first-order perturbation methods) or verify the invertibility of the general equilibrium adjustment matrix (see Auclert et al., 2019).

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

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

(2010), and my sample period is 1960:Q1—2006:Q2.[37] Priors are reported in Table B.2.

The first exception is the transfer adjustment parameter ρ_τ; since I do not include data on government debt, this parameter would likely be poorly identified. I thus simply set $\rho_\tau = 0.9$, in line with the VAR evidence documented in Galí et al. (2007) and Appendix E.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 can be shown that the slope parameter of the wage-NKPC (B.4) can be equivalently written as

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

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.[38]

The results of the estimation are displayed in Table B.2. Since they are not relevant for my purposes here, I omit estimates of shock persistence and volatility. I find the posterior mode using the `csminwel` routine provided by Chris Sims; accuracy of the demand equivalence approximation beyond the mode parameterization is discussed in Appendix D.1.[39] Overall the results are consistent with the estimates in Justiniano et al.. 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.[40]

Simplified model. The simplified HANK model considered for the illustration in Figure 1 is identical to the estimated model except for one change: I set $\phi_w = 1$ (and so $\kappa_w = 0$). As a result, demand equivalence holds exactly.

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

[38] Direct estimation of κ_w in my set-up tends to yield a similar (if slightly steeper) wage-NKPC. Most previous work that has used time series data to estimate the degree of wage stickiness instead finds much larger numbers for ϕ_w and so smaller numbers for κ_w (e.g. Justiniano et al., 2010).

[39] The routine only ensures that the estimated mode is a (local) peak of the posterior. In Appendix D.1 I thus also check accuracy very far away from this local maximum.

[40] Relative to the framework of Justiniano et al., 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). Auclert et al. (2019) discuss the effects of these changes on the decomposition of business cycles into structural shocks.

			Prior		*Posterior*
Parameter	**Description**	**Density**	**Mean**	**Std**	**Mode**
ϕ_p	Price Calvo Parameter	B	0.7	0.10	0.80
ζ	Capacity Utilization	N	5.00	1.00	5.21
κ	Investment Adjustment Cost	N	4.00	1.00	3.92
ρ_m	Taylor Rule Persistence	B	0.80	0.20	0.69
ϕ_π	Taylor Rule Inflation	N	2.00	0.10	1.74
ϕ_y	Taylor Rule Output	N	0.15	0.05	0.03
ϕ_{dy}	Taylor Rule Output Growth	N	0.15	0.05	0.30

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

B.3 Model extensions of Section 4.2

For all model extension I first describe the change in the environment and then discuss my approach to parameterization, as needed for the quantitative analysis of Section 4.

B.3.1 Multiple goods

The full model with multiple goods departs from the one-sector baseline in three ways.

First, it features three goods—two consumption goods and an investment good. The household consumption basket c_{it} now satisfies

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

I let the ideal price index of the consumption bundle be the numeraire of my economy (so that we again have $p_t^c = 1 \ \forall t$), and I denote the real relative prices of the two consumption goods by x_t^1 and x_t^2. Investment is only possible using the economy's investment good, whose real relative price is denoted x_t^I. The government purchases each of the three goods, with potentially different spending multipliers for each, and the monetary authority responds to changes in consumer price inflation.

Second, household disutility over labor supply takes the same functional form as before, with ℓ_t^h now given as an aggregate of labor supply for each of the three goods:

$$\ell_t^h \equiv \left[\chi_1(\ell_{1t}^h)^{\frac{\varphi+\mu}{\varphi}} + \chi_2(\ell_{2t}^h)^{\frac{\varphi+\mu}{\varphi}} + \chi_I(\ell_{It}^h)^{\frac{\varphi+\mu}{\varphi}}\right]^{\frac{\varphi}{\varphi+\mu}}$$

where $\{\chi_1, \chi_2, \chi_I\}$ govern disutility from work in each of the sectors. $\mu = 0$ corresponds to perfect labor mobility across the sectors, while $\mu = 1$ corresponds to perfect immobility, with all labor types entering separately into household utility. 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:

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

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

Third, there are separate production sectors for each of the three goods. Briefly, I simply repeat the production sector of the baseline model described in Appendix B.1 three times, but with good-specific final prices x_t^m and potentially heterogeneous capital shares α_m. All three sectors then purchase capital goods at price x_t^I, hire labor at cost w_t^m, and sell their own good at real price x_t^m.

Parameterization. I build on the parameterization of the estimated HANK model, with one notable difference: a smaller degree of nominal price rigidities. In the model, the probability of price re-sets governs relative price movements after a demand shock for a specific good. I have included measures of relative prices in my VARs and find little response (Figure E.1), similar to Nakamura & Steinsson (2014); however, Ramey & Shapiro (1998) show that, after large and persistent government spending shocks that move output by almost 4 per cent, relative prices move by 2.5 per cent. To be conservative, I set $\phi_p = 0.6$.

Next, I set $\mu = 1$ (i.e., fully sector-specific labor). I set the average capital share $\bar{\alpha} \equiv \alpha_1\frac{\bar{y}_1}{\bar{y}} + \alpha_2\frac{\bar{y}_2}{\bar{y}} + \alpha_I\frac{\bar{y}_I}{\bar{y}} = 0.2$ (as in my baseline model), and then pin down relative α's as in Alonso (2017, Table 3.3), giving $\alpha_1 = 0.30$, $\alpha_2 = 0.15$, $\alpha_I = 0.19$. The fraction of labor in each of the three sectors is set so that their relative sizes are also data-consistent; again following Alonso (2017), this gives $\bar{y}_1/\bar{y} = 0.29$, $\bar{y}_2/\bar{y} = 0.48$ $\bar{y}_I/\bar{y} = 0.23$. I recover $\bar{g}_I$ residually from the market-clearing condition for good I, and then set $\bar{g}_1 = \bar{g}_2 = \frac{1}{2}(\bar{g} - \bar{g}_I)$, with $\bar{g}$ set as before. Finally I then recover the weight ν in household preferences as $\nu = \frac{\bar{c}_1}{\bar{c}}$, and set the labor preference weights $\{\chi_1, \chi_2, \chi_I\}$ to clear the labor market given $\{\bar{w}_1, \bar{w}_2, \bar{w}_I\}$.

B.3.2 Open economy

To study the role of demand leakage abroad I consider a small open economy version of my baseline HANK model, following Auclert et al. (2021). Consumers and firms in the home economy H consume and invest using a final good bundle that consists of both domestic and foreign goods (indexed by F), while the government consumes only the domestic good. The domestic economy is small, so domestic policies do not affect the rest of the world.

The domestic consumption basket aggregates the home and foreign final good:

$$c_{it} = \left[\phi^{\frac{1}{\eta_1}}(c_{it}^H)^{\frac{\eta_1-1}{\eta_1}} + (1-\phi)^{\frac{1}{\eta_1}}(c_{it}^F)^{\frac{\eta_1-1}{\eta_1}}\right]^{\frac{\eta_1}{\eta_1-1}}$$

Here ϕ is the degree of home bias and η_1 is the elasticity of substitution between home and foreign goods. I let the price of the total consumption bundle be the numeraire, and denote the real relative prices of the domestic and foreign final good by x_t^H and x_t^F, respectively. Log-linearized real relative prices thus satisfy

$$\phi\widehat{\tilde{x}}_t^H + (1-\phi)\widehat{\tilde{x}}_t^F = 0$$

For simplicity I assume that the investment bundle purchased by intermediate goods producers also consists of the domestic and foreign final goods, with the same steady-state home share ϕ and elasticity of substitution η_1. The problem of domestic intermediate goods producers is then unchanged relative to the baseline economy. For retailers we now get a price-NKPC in inflation of the domestic good:

$$\widehat{\tilde{\pi}}_t^H = \kappa_p(\widehat{\tilde{p}}_t^I - \widehat{\tilde{x}}_t^H) + \frac{1}{1+\bar{r}_b}\widehat{\tilde{\pi}}_{t+1}^H$$

where inflation and real relative prices are linked as

$$\widehat{\tilde{\pi}}_t^H = (\widehat{\tilde{x}}_t^H - \widehat{\tilde{x}}_{t-1}^H) + \widehat{\tilde{\pi}}_t$$

Let e_t denote the nominal exchange rate. Since foreign prices are fixed we have that

$$\widehat{\tilde{e}}_t - \widehat{\tilde{e}}_{t-1} = (\widehat{\tilde{x}}_t^F - \widehat{\tilde{x}}_{t-1}^F) + \widehat{\tilde{\pi}}_t$$

With nominal interest rates on foreign bonds also fixed, arbitrage dictates that

$$\widehat{\tilde{i}}_t = \widehat{\tilde{e}}_{t+1} - \widehat{\tilde{e}}_t \tag{B.8}$$

Finally, foreign consumer and firm demand for the domestic final good satisfies

$$\begin{aligned}\widehat{\tilde{c}}_t^{H*} &= -\eta_2(\widehat{\tilde{x}}_t^H - \widehat{\tilde{x}}_t^F) \\ \widehat{\tilde{i}}_t^{H*} &= -\eta_2(\widehat{\tilde{x}}_t^H - \widehat{\tilde{x}}_t^F)\end{aligned}$$

where η_2 elasticity of substitution between home and foreign goods in the foreign bundle.

I consider the same monetary policy rule as before, so the central bank stabilizes inflation in the domestic consumption bundle. The fiscal authority consumes only the domestic good, so in the government budget constraint we have $p_t^g = x_t^H$. Finally, the domestic bond market-clearing condition is dropped for the arbitrage relation (B.8). The model is closed by requiring domestic output market-clearing, which dictates that[41]

$$c_t^H + c_t^{H*} + i_t^H + i_t^{H*} + g_t = y_t$$

PARAMETERIZATION. I set $\phi = 0.89$, matching the domestic consumption share of the U.S. economy. For the elasticities of substitution I set $\eta_1 = \eta_2 = 2$, in line with previous work. All other model parameters are kept exactly as in the baseline HANK model of Section 4.1.

B.3.3 Two-asset model

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

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

such that

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

and

$$b_{it}^h \geq \underline{b}, \quad a_{it}^h \geq 0$$

[41] Note that my estimated model sets $\kappa_b = 0$, and so I here drop the intermediation cost term from the market-clearing condition.

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

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

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

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

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

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

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

The rest of the economy is unchanged; in particular, firms still discount at $\frac{1+i_{t-1}^b}{1+\pi_t}$, which in the absence of aggregate risk is equivalent to discounting at $\frac{1+i_{t-1}^m}{1+\pi_t} = \frac{1+i_{t-1}^h}{1+\pi_t}$. The only change to Definition 2 is the new asset market-clearing condition:[42]

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

I solve the model following the computational strategy of Bayer & Luetticke (2020).

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

[42]The output market-clearing condition now additionally features the liquid asset financial intermediation cost, given as $\kappa_b \int_0^1 b_{it-1}^h di$.

Parameter	Description	Value	Target	Model	Data
Households					
η	Probability of Adjustment	0.15	A/Y	11.29	10.64
β	Discount Rate	0.98	B/Y	1.49	1.04
r^h	Return	0.015	Upper Bound		
κ_b	Liquid Wedge	0.0125	Upper Bound		
Firms					
δ	Depreciation	0.009	Firm Valuation		

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

displayed in Table B.1 for the benchmark one-asset model.

To provide a stringent test of the demand equivalence approximation, I set the wedge between returns on household deposits and government debt to be 1.25 per cent *per quarter*. Given this large difference, I then choose the adjustment probability η to ensure a reasonable fit to total liquid and illiquid wealth in the U.S. economy.

B.3.4 Weak wealth effects

Relative to the baseline model, the environment without unions but with weak wealth effects in labor supply differs in three respects. First, the economy is now populated by a *double* unit continuum of households—a unit continuum of *families* $f \in [0,1]$, and a unit continuum of households $i \in [0,1]$ for each f. Each family is a replica of the unit continuum of households in the benchmark model, but shock exposures may be heterogeneous across families. I will explain the purpose of this artificial construction momentarily. Second, there are no unions—each household decides on its own labor supply. Third, I change household preferences. Similar to Jaimovich & Rebelo (2009) and Galí et al. (2012), I assume that

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

where the preference shifter θ_{ift} satisfies

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

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

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

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

$$\chi \ell_{ift}^{\frac{1}{\varphi}} = w_i x_{ft}^{-\gamma} \tag{B.9}$$

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 2 will pick up direct earnings responses. In particular, let $\boldsymbol{\ell}^h = \boldsymbol{\ell}^h(\mathbf{w}, \mathbf{c})$ denote the mapping from wages and family consumption into family labor supply induced by (B.9). The micro regression estimand in (9) then satisfies

$$\widehat{\mathbf{c}}_{\tau}^{PE} = \left(I - \frac{\partial \mathbf{c}}{\partial \boldsymbol{\ell}} \times \frac{\partial \boldsymbol{\ell}^h}{\partial \mathbf{c}} \right)^{-1} \times \left(\frac{\partial \mathbf{c}}{\partial \boldsymbol{\tau}} \cdot d\boldsymbol{\tau} \right) \tag{B.10}$$

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

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

B.3.5 Productive government spending

In the model variant with productive government spending, the intermediate goods production function is generalized to take the form

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

where $k_t^g \equiv (1-\delta)k_{t-1}^g + g_t$.

Government purchases thus endogenously and gradually improve the productive capacity of the economy. Note that, for simplicity, the assumed depreciation rate of the stock of total "government capital" is identical to the rate of depreciation of private capital.

PARAMETERIZATION. I set $\alpha_g = 0.4$, giving a two-year cumulative government spending multiplier that is around 35 per cent larger than the unit multiplier in the baseline model. Such a two-year cumulative multiplier for government investment is roughly consistent with the empirical evidence reviewed in Leduc & Wilson (2013) and Gechert (2015).

C Further results on demand equivalence

This appendix collects several supplementary theoretical results. In Appendix C.1 I show that my arguments apply without change to perturbations around arbitrary transition paths. Appendices C.2 and C.3 emphasize the generality of consumption demand equivalence by considering a larger family of shocks and models, and in Appendix C.4 I illustrate the range of general equilibrium outcomes consistent with exact equivalence. Finally, in Appendix C.5, I show that many popular heterogeneous-firm models of investment are nested by the investment demand equivalence result.

C.1 General transition paths

All equivalence results in this paper are stated for transition paths starting at the deterministic steady state. However, it is immediate from the proof of Proposition 1 (and similarly that of Proposition A.1) that nothing 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. The equivalence results would then apply to deviations from the unshocked transition path of the economy back to steady state. These deviations need not agree with impulse responses at steady state, but they remain tied together across different kinds of demand shocks.

C.2 Generic consumption demand shifters

I argued in Section 2.3 that demand equivalence also extends to *generic* shifters of consumption demand. To establish this claim I augment the baseline model to feature fluctuations in household patience as a simple reduced-form stand-in for various more plausibly structural shocks to consumer spending (e.g. changes in borrowing constraints, redistribution, ...).

The discount factor of every household is now subject to an additional common shifter ζ_t, with $\boldsymbol{\zeta} = \boldsymbol{\zeta}(\boldsymbol{\varepsilon}_\nu)$, giving preferences as

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

Note that impatience shocks—shocks that just shift the intertemporal profile of private consumption spending—necessarily induce consumption response paths $\widehat{\boldsymbol{c}}_{\nu}^{PE}$ with zero net present value. As a result, a government spending shock satisfying requirement (i)—i.e., $\widehat{\boldsymbol{g}}_g = \widehat{\boldsymbol{c}}_{\nu}^{PE}$—also has zero net present value, and so need not be financed through any change in taxes or transfers. By the same argument as in the proof of Proposition 1 we can thus again have $\boldsymbol{\tau}_g^e = \boldsymbol{\tau}_{\nu}^e$, and so the rest of the proof applies without change to give

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

Importantly, the observed tax responses to the two shocks now reflect *only* general equilibrium feedback to taxes (e.g., through changes in inflation or labor tax revenue).

C.3 Exact equivalence beyond the baseline model

The proof of the consumption demand equivalence result applies to any model that satisfies the set of semi-structural exclusion restrictions in Section 2.3. This section briefly discusses some prominent examples of such models.

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

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

such that

$$c_{it} + d_{it}^h + b_{it}^h = (1-\tau_\ell) w_t e_{it} \ell_{it} + \left(\frac{1 + i_{t-1}^b}{1 + \pi_t} + \kappa_b \mathbf{1}_{b_{it-1}^h < 0} \right) b_{it-1}^h + (1-\delta) d_{it-1}^h + \tau_{it} + d_{it} + \phi_d(d_{it-1}^h, d_{it}^h)$$

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 (e.g., as in 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.[43]

Crucially, I assume that the common final good y_t can be costlessly turned into either the durable or the non-durable consumption good, as reflected in the absence of relative price terms in the household budget constraint. Additionally imposing Assumption 1, the aggregate resource constraint then becomes

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

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

Corollary C.1. *Extend the structural model of Section 2.1 to feature durable goods, as in Problem (C.3). Consider a stimulus check policy $\boldsymbol{\varepsilon}_\tau$, and suppose that Assumptions 1 to 3 hold. Then, for a fiscal spending policy $\boldsymbol{\varepsilon}_g$ such that (i) $\widehat{\mathbf{g}}_g = \widehat{\mathbf{e}}_\tau^{PE}$ (identical net excess demand) and (ii) $\widehat{\boldsymbol{\tau}}_g^e = \widehat{\boldsymbol{\tau}}_\tau^e$ (identical tax response), we have that, to first order,*

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

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

Preferences. My baseline structural model assumes time separability in household preferences. It is, however, immediate that general forms of time non-separability are similarly nested: as long as the consumption block of the model admits aggregation to some aggregate consumption function $\mathbf{c}(\bullet)$, the equivalence proof goes through unchanged. My approximate

[43] In particular, this implies that the model can in principle be consistent with empirical evidence suggesting MPCs close to zero for some households (non-adjusters) and in excess of one for others (those pushed towards durables adjustment).

equivalence results for the model of Justiniano et al. (2010)—with habit formation as a very simple form of non-separability—illustrate this claim.

Valued government spending. In my baseline model, households do not value government expenditure. However, it is immediate from the proof strategy for consumption demand equivalence that this assumption is stronger than necessary—the key restriction is that the aggregate consumption function $\mathbf{c}(\bullet)$ does not directly depend on government consumption. A possible sufficient condition is that government spending enters the per-period felicity function in an additively separable fashion,

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

This is for example the case under a CES preference specification

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

with $\rho = \gamma$.

Expectation formation. All of the models considered in this paper impose rational expectation formation for households and firms. An attractive alternative is the sticky information structure in Auclert et al. (2019). For a simple example, suppose that only the consumption-savings problem of households is subject to a sticky information friction, with a fraction $1-\theta$ of households updating their information at each point in time t. Then, for every input $\boldsymbol{p}$ to the consumption-savings problem, the sticky information consumption derivative map $\mathcal{C}_p \equiv \frac{\partial \boldsymbol{c}}{\partial \boldsymbol{p}}$ is related to the original rational expectations map $\mathcal{C}_p^*$ via

$$\mathcal{C}_p = \begin{pmatrix} \mathcal{C}_p^*(1,1) & (1-\theta)\mathcal{C}_p^*(1,2) & (1-\theta)\mathcal{C}_p^*(1,3) & \dots \\ \mathcal{C}_p^*(2,1) & (1-\theta)\mathcal{C}_p^*(2,2) + \theta\mathcal{C}_p^*(1,1) & (1-\theta)\mathcal{C}_p^*(2,3) + \theta(1-\theta)\mathcal{C}_p^*(1,2) & \dots \\ \mathcal{C}_p^*(3,1) & (1-\theta)\mathcal{C}_p^*(3,2) + \theta\mathcal{C}_p^*(2,1) & \vdots & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

Since the proof of Proposition 1 relies only on the existence of these linear maps (and not their shape), it follows immediately that all results extend without change to such behavioral model economies.

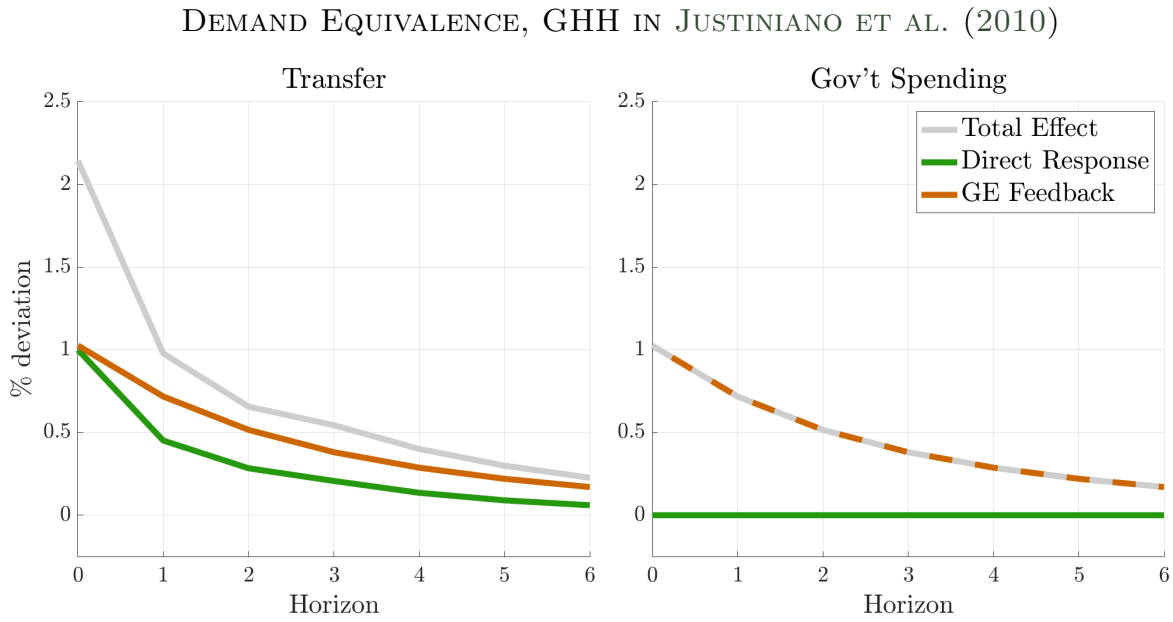

Figure C.1: Consumption impulse response decompositions after impatience and government spending shocks in the model of Justiniano et al. (2010) with GHH preferences, where the two shocks are selected to have identical effects on net excess demand, displayed as the green line in the left panel. The direct response and the indirect general equilibrium feedback are then computed following Definition 1.

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

Proposition 1 asserts that private and public spending shocks induce the same general equilibrium effects, but is silent on the strength of this common general equilibrium feedback. In this section I give two extreme examples, one with full general equilibrium crowding-out, and one with strong general equilibrium amplification.

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

The second example is quantitative. I consider the estimated New Keynesian business-

cycle model of Justiniano et al. (2010), but now assume that preferences are as in Greenwood et al. (1988). Results are reported in Figure C.1.

It is immediate that this model satisfies all assumptions in Proposition 1, and so exact demand equivalence holds. 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 et al., 2020, for an analytical characterization).

C.5 Nested models for investment demand equivalence

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

First, in all of those models, capital adjustment costs are internal to the firm, so Assumption A.1 holds. Second, each model is closed with a simple representative household with linear labor disutility, so Assumptions A.2 and A.3 hold. Finally, since none of these models feature nominal rigidities, Assumption A.4 is irrelevant.[44]

[44]Well-known heterogeneous-firm models with nominal rigidities include Ottonello & Winberry (2018) and Koby & Wolf (2020). In both cases Assumption A.4 is satisfied. Furthermore, and as discussed in Appendix A.2, this assumption anyway just affects the interpretation of the results: the demand equivalence approximation yields an investment demand shock counterfactual valid for the same interest rate response as that of the originally identified fiscal spending shock.

D Approximation accuracy

In this appendix I provide supplementary details to my assessment of the demand equivalence approximation in Section 4.

First, complementing the discussion in Section 4.1, Appendices D.1 and D.2 consider alternative parameterizations of the baseline HANK model as well as other canonical quantitative business-cycle models. Second, in Appendices D.3 to D.7, I present detailed results for the accuracy checks of Section 4.2. Finally, in Appendix D.8 I add one further experiment: I construct the proposed demand equivalence approximation under the assumption that the demand matching in (11) is not exact. In all experiments I report the population estimands of the demand equivalence methodology, effectively assuming that the econometrician has access to infinitely large cross-sectional and time series samples.

D.1 Random parameter draws

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

Parameter	**Description**	**Lower Bound**	**Upper Bound**
ϕ_p	Price Calvo Parameter	0.15	0.95
ζ	Capacity Utilization	0.5	10
κ	Investment Adjustment Cost	0.5	10
ρ_m	Taylor Rule Persistence	0.15	0.95
ϕ_π	Taylor Rule Inflation	1.1	2.5
ϕ_y	Taylor Rule Output	0	0.5
ϕ_{dy}	Taylor Rule Output Growth	0	0.5

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

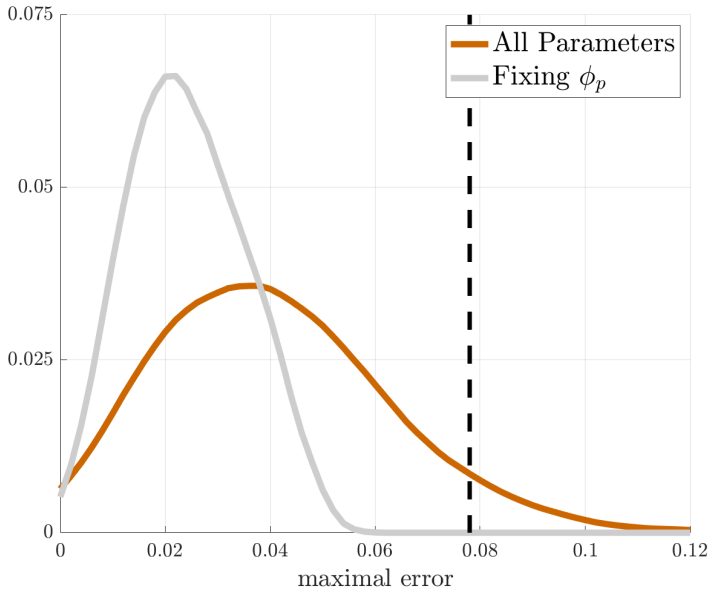

Figure D.1: Kernel estimate of maximal error distribution, with parameters drawn randomly according to Table D.1 (orange, 95th percentile black dashed). The grey lines show the same kernel density estimate when ϕ_p is fixed at its estimated posterior mode.

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

D.2 Other estimated business-cycle models

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

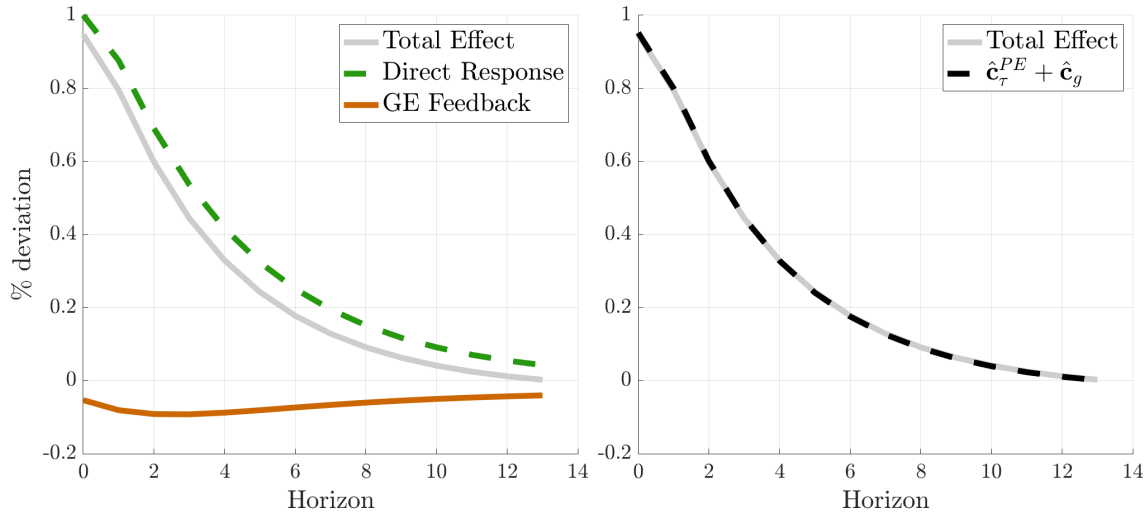

Figure D.2: Consumption impulse response decompositions and demand equivalence approximation in the model of Justiniano et al. (2010), solved at the posterior mode and for an impatience shock with persistence $\rho_b = 0.1$. The direct response and the indirect general equilibrium feedback are computed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

Justiniano et al. (2010). In the estimated model of Justiniano et al. (2010), consumption demand equivalence fails only because Assumption 3 is not satisfied: wealth effects in labor supply are not zero, and hours worked are not fully demand-determined. However, prices and wages are estimated to be very sticky, and so—consistent with Christiano (2011a)—hours worked are still largely demand-determined, at least in the short run. This discussion suggests that demand equivalence should hold nearly exactly. Figure D.2 shows that this is indeed the case: for a transitory consumption demand (impatience) shock, the error associated with the demand equivalence approximation is barely visible.

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

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

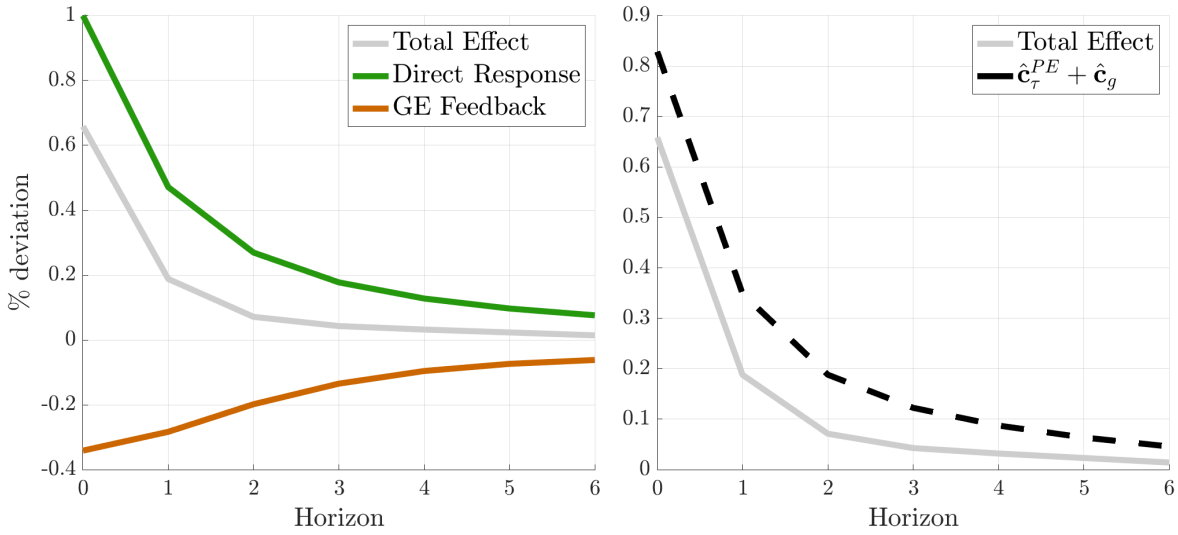

Figure D.3: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with (nearly) flexible prices & wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

where

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

and

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

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

D.3 Labor supply

The "flexible price" error line in Figure 3 corresponds to an economy in which prices and wages re-set every 1.25 quarters on average (i.e., I set $\phi_p = \phi_w = 0.2$). Figure D.3 shows the corresponding full decomposition of aggregate impulse responses: there is significant general equilibrium crowding out, and the simple demand equivalence approximation misses a large fraction of that crowding-out.

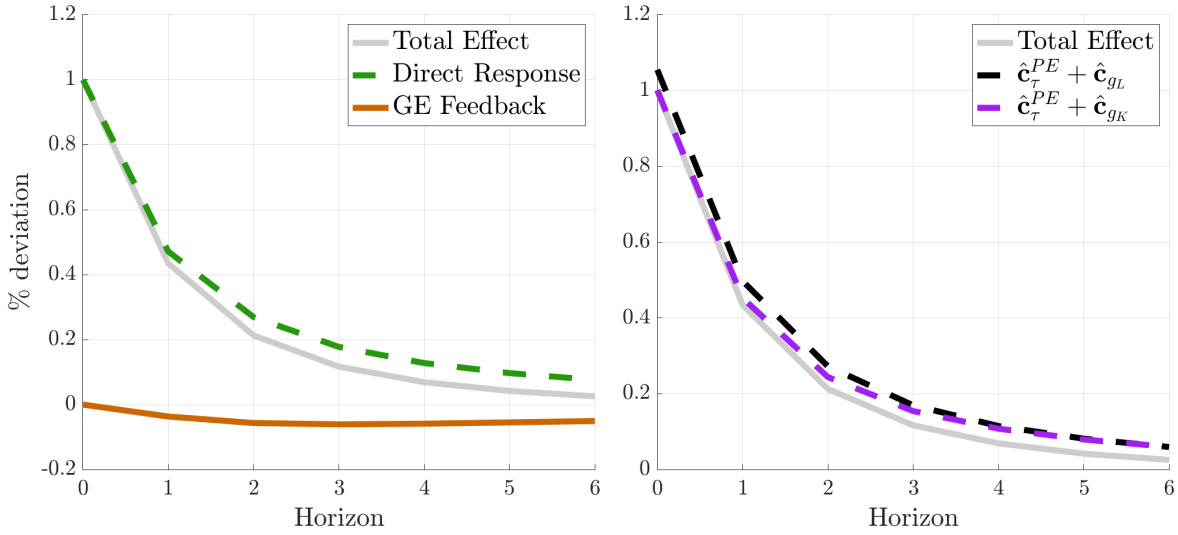

Figure D.4: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with multiple goods. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

D.4 Multi-sector economy

Figure D.4 shows impulse response decompositions and the demand equivalence approximation for government purchases of (i) the labor-intensive consumption good (black) and (ii) the capital-intensive consumption good (purple). In both, relative price effects increase the approximation error relative to the baseline economy. General equilibrium MPC multiplier effects, however, *increase* the error for purchases of the labor-intensive good, and *decrease* it for purchases of the capital-intensive good. Figure 3 considers government purchases of the labor-intensive good (which give a larger bias), and is thus conservative.

D.5 Productive government spending

Full results for the exercise with productive government spending are reported in Figure D.5. Since the productive benefits of government purchases in this model variant are large, the demand equivalence approximation is relatively poor, in particular at long horizons. This reflects the fact that those productive benefits accrue gradually and are long-lasting, persisting long beyond the initial demand stimulus itself.

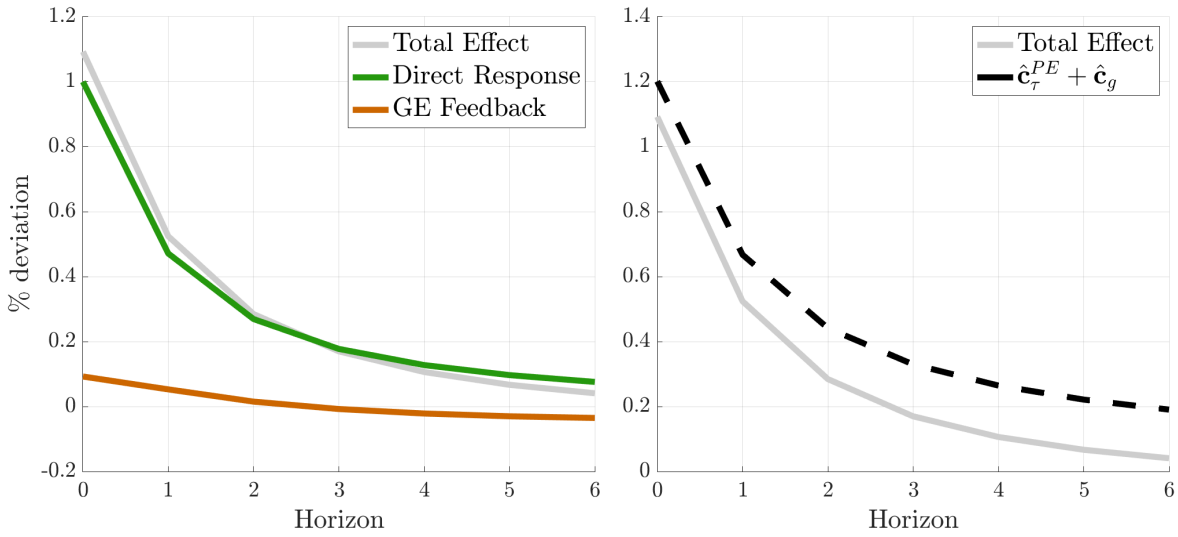

Figure D.5: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with productive government spending. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

D.6 Open Economy

Full results for the open-economy extension are reported in Figure D.6. I emphasize that the error displayed there is small only because my model economy is fairly closed, with $\phi = 0.89$. While fitting for the U.S., such a calibration is clearly not approrpriate for all economies; for example, setting $\phi = 0.5$ (which corresponds to typical calibrations for small, very open economies), the peak demand equivalence error is much larger, at around 15 per cent.

D.7 Interest rates

In Figure D.7 I consider a two-asset HANK model with a penalty on household liquid savings. That model is a particularly stringent test of the demand equivalence approximation: it pushes the approximation error upwards, reinforcing the labor supply bias and thus giving me the largest possible error.[45] Even in that extreme case, however, the demand equivalence approximation overall remains quite accurate.

[45]In previous drafts of this paper I have instead considered model variants with large borrowing wedges (e.g., reflecting credit card debt). With indebted households facing large effective rates of return, the interest rate channel in this case imparts a (small) negative bias, largely offsetting the small positive bias of the labor supply channel.

Approximate Demand Equivalence, Open Economy

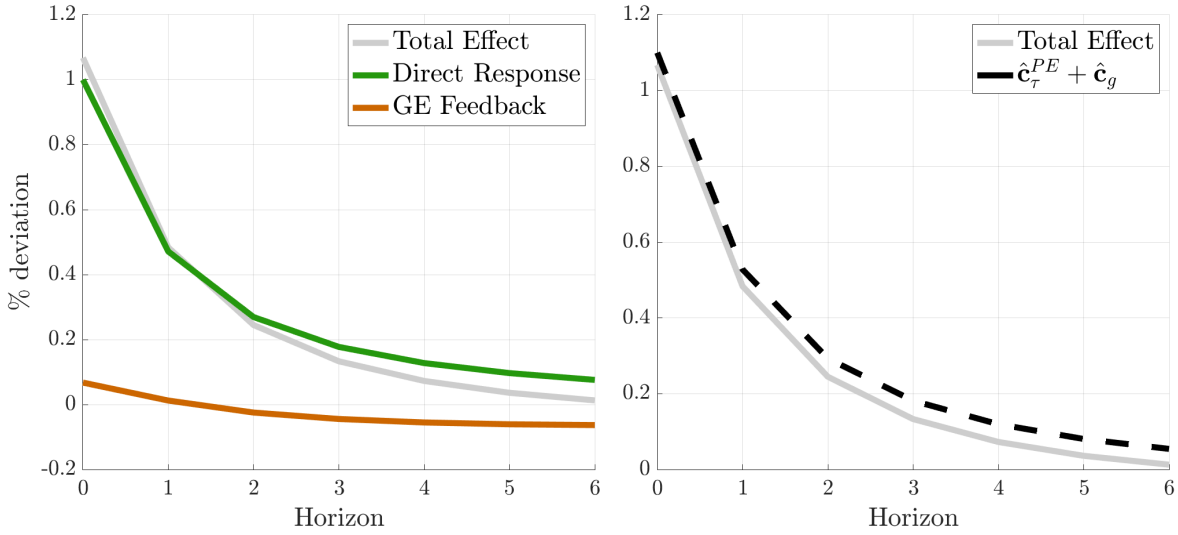

Figure D.6: Consumption impulse response decompositions and demand equivalence approximation for the open-economy HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Approximate Demand Equivalence, Two-Asset HANK Model

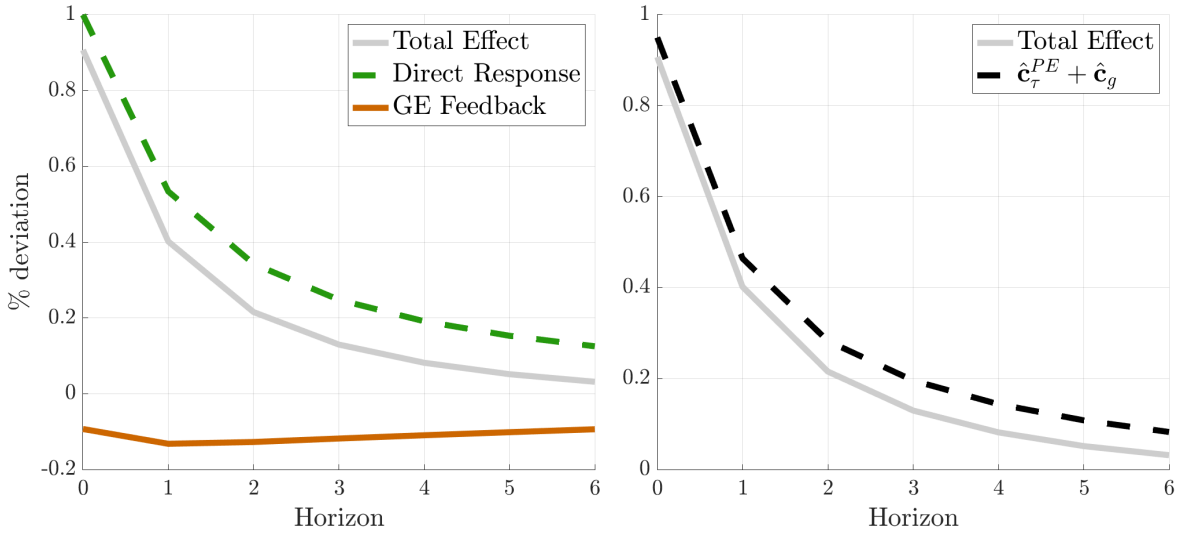

Figure D.7: Consumption impulse response decompositions and demand equivalence approximation for the two-asset HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

D.8 Imperfect demand matching

The net excess demand path in Figure F.1 is matched well, but of course not perfectly. With imperfect demand matching, my demand equivalence aggregation procedure will be correct up to the general equilibrium effects of a shock that induces an aggregate net excess demand path equal to the matching error path. To gauge the distortions associated with moderate mis-matching of the kind observed in my empirical applications, 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 $\widehat{\mathbf{g}}_g$ with

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

where $\nu_t \sim N(0, \sigma_\nu^2)$. I set σ_ν^2 to get average errors comparable in size to those displayed in Figure F.1; this gives $\sigma_\nu^2 = 0.123$.

I then construct the demand equivalence approximation for 1,000 draws of the error sequence $\boldsymbol{\nu}$ in (D.1), and for each compute the maximal prediction error relative to the peak true consumption response. I find that 95 per cent of all errors lie below 9.6 per cent, and so the approximation remains quite accurate.[46] The intuition is quite transparent: since the model only features relatively moderate general equilibrium amplification, prediction errors for consumption can only be large if the error in demand path matching itself is substantial. This error, however, is by construction small, and thus so are the overall approximation errors. To illustrate, Figure D.8 shows the quality of the demand equivalence approximation for one particular draw of the error sequence $\boldsymbol{\nu}$, with the implied government spending net excess demand path displayed in purple.

[46] Most of the large approximation errors come from draws in which the $\boldsymbol{\nu}$'s are so far from 0 that demand matching is clearly violated, so the results displayed here are actually a quite conservative upper bound on likely inaccuracies.

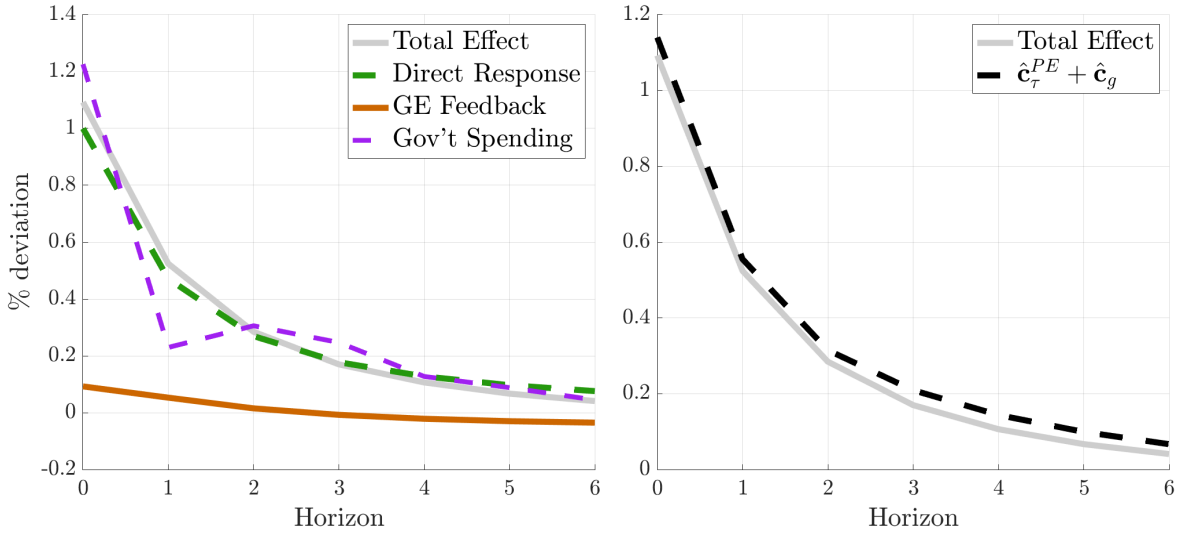

Figure D.8: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, with imperfect demand matching, following (D.1). The direct response and the indirect general equilibrium feedback are computed following Definition 1.

E Empirical appendix

This appendix presents the empirical results that I use as an input to my applications in Appendix F. Appendix E.1 discusses estimates of the direct (partial equilibrium) consumption response to stimulus checks, Appendix E.2 does the same for the investment response to tax credits, and finally Appendix E.3 presents the time series estimates of government spending shock transmission that I use to recover the "missing intercept."

E.1 Cross-sectional consumption elasticities

I review empirical evidence on direct consumption responses to stimulus check receipt. First, I begin by giving conditions under which the regressions of Parker et al. (2013) for the 2008 stimulus check experiment can indeed be interpreted as giving such direct responses. Second, I discuss various alternative estimates.

Baseline estimates. Proposition 2 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 heavily on Kaplan & Violante (2014), I here discuss how to interpret their regression estimands. Parker et al. estimate a differenced version of (9):

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

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

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

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. We can then

follow exactly the same steps as in the proof of Proposition 2 to show that, to first order,

$$\beta_0 = MPC_{0,0}, \quad \beta_1 = MPC_{1,0} - MPC_{0,0}$$

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 $t = 0$.

For my baseline analysis in Appendix F.1, I will indeed make the strong assumption that rebate receipt was a surprise for all households, or equivalently that anticipation effects are negligible. While strong, this assumption is at least broadly consistent with results reported in Broda & Parker (2014), Kueng (2018), Ganong & Noel (2019) and Baugh et al. (2021). Under this assumption I can interpret the estimates of Parker et al. (2013) and Broda & Parker (2014) as giving $MPC_{0,0} = 0.5$ and $MPC_{1,0} = 0.2$. I then extrapolate assuming a geometric rate of decay from $t = 1$ onwards, together with the requirement that

$$\sum_{t=0}^{\infty} \left(\frac{1}{1 + \bar{r}^b} \right)^{t-1} MPC_{t,0} = 1. \tag{E.3}$$

This gives the green line in Figure F.1, with $\bar{r}^b = 0.01$.[47]

Alternative estimates. I briefly discuss two alternative strategies to estimating the required direct consumer spending responses $MPC_{t,0}$. Both indicate somewhat lower MPCs; I construct stimulus check counterfactuals using those alternative estimates in Appendix F.2.

First, consumer spending responses to lump-sum income gains may be estimated through surprise lottery wins, as done in Fagereng et al. (2018) for Norwegian data. Such studies give the required $MPC_{t,0}$ without any further assumptions on expectation formation, but are of course less directly informative about my actual policy experiment of interest—the 2008 stimulus check experiment in the U.S. Most notably, household balance sheets in the U.S. in 2008 are likely to have been quite different from household balance sheets in Norway in normal times. Overall, once translated to quarterly frequency, the study of Fagereng et al. suggests MPCs around 0.2-0.3 (see Auclert et al., 2019, for a discussion of the interpolation).

Second, the original econometric specification in Parker et al. (2013) is subject to recent concerns about two-way fixed effects estimators in the presence of heterogeneous treatment effects (see Orchard et al., 2022, for a review of this literature), leading to a potential upward

[47] A constant rate of decay of intertemporal MPCs is (roughly) consistent with standard incomplete-market models (Auclert et al., 2018; Wolf, 2021) as well as other empirical evidence (e.g., Fagereng et al., 2018). The adding-up condition (E.3) on the other hand simply follows from household budget constraints.

bias in the MPC estimates. Correcting for this source of bias, Orchard et al. find an impact MPC estimate of around 0.3.

E.2 Cross-sectional investment elasticities

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

$$\widehat{i}_{jt+h} = \alpha_j + \delta_t + \beta_{qh} \times z_{n(j),t} + u_{jt} \tag{E.4}$$

where $z_{n(j),t}$ is the size of the bonus depreciation investment stimulus for industry $n(j)$ of firm j. They estimate this regression on a quarterly Compustat sample spanning the years 1993–2017; this sample period features the two bonus depreciation episodes of 2001-2004 and 2008-2010, exactly as in Zwick & Mahon (2017). They then give sufficient conditions under which the estimands $\{\beta_{qs}\}$ are interpretable as the direct partial equilibrium response of investment to a one-time bonus depreciation stimulus. Given their estimated partial equilibrium path $\{\widehat{i}^{PE}_{qt}\}_{t=0}^{3}$, I recover the full partial equilibrium investment response $\widehat{\boldsymbol{i}}^{PE}_{q}$ by simply fitting a single Gaussian basis function, following Barnichon & Matthes (2018).

E.3 Time series fiscal policy estimates

My time series analysis of aggregate government spending shock propagation closely follows Ramey (2011) and Blanchard & Perotti (2002).[48]

Ramey (2011). The first approach to the identification of aggregate government spending shocks relies on professional forecast errors. These forecast errors are treated as a valid IV for structural government spending shocks, and I study their propagation by ordering them first in a recursive VAR (Plagborg-Møller & Wolf, 2021). This approach to identification amounts to assuming a conventional timing restriction—i.e., that government spending does not directly, within the quarter, respond to any other macroeconomic shocks—but now defines innovations with respect to a larger information set—that of the forecasters, and not the reduced-form VAR itself. This promises to sidestep potential non-invertibility concerns.

My benchmark VAR consists of the log real per capita quantities of total government

[48] In a previous version of this paper I also leveraged the identification strategy of Caldara & Kamps (2017) which gives a government spending shock similarly persistent to that of Blanchard & Perotti. Results based on this shock series are available upon request.

spending, total output (GDP), total (non-durable, durable and services) consumption, private fixed investment, total hours worked, and a measure of the federal average marginal tax rate (Alexander & Seater, 2009).[49] All variables are defined and then measured as in Ramey (2011). As further robustness checks, I also consider alternative specifications with (i) Greenbook defense spending forecast errors in lieu of professional forecaster errors (taken from Drautzburg, 2020), (ii) a measure of log total government debt (taken from Ramey, 2016), (iii) the federal funds rate as a measure of the monetary policy stance, and (iv) a measure of the real relative price of the government consumption bundle.[50,51]

I estimate all VARs in levels, with a quadratic time trend and four lags. 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.[52]

Figure E.1 shows the impulse responses of government spending, output, consumption, investment, the marginal tax rate, total federal debt, the real relative price of the government bundle, and the federal funds rate. As in most existing structural VAR work, I construct 16th and 84th percentile confidence bands; the output and tax responses, however, remain significant at the more conventional 95 per cent level. In line with most of the previous literature I find a significant positive output response (corresponding to around a unit multiplier), and a largely flat reaction of consumption, with some delayed crowding-out. 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. Finally, I find that neither the relative price of the government bundle nor the nominal rate respond much.[53]

My central results—the 1-1 increase in output and the limited crowding-out of private expenditure—are robust to various changes in model specification. First, I have experimented

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

[50]I obtain the debt series from the tax shock replication data for Ramey (2016), deflating *pubdebt* by *pgdp*. For the real relative price series, I divide the implicit price deflator for federal government consumption expenditures and gross investment by the GDP deflator.

[51]For demand matching I need to re-scale public and private demand shocks to be in dollar terms. This can be done using information on the GDP shares of consumption, investment, and government expenditure.

[52]My use of a recursive VAR follows the finite-sample recommendations of Li et al. (2021).

[53]Interestingly, the absence of a nominal interest rate reaction suggests a violation of the Taylor principle. This could nevertheless be consistent with equilibrium determinacy if either (i) the Taylor principle is satisfied off-equilibrium (e.g., via a King-type rule, see Cochrane (2011)) or (ii) the Taylor principle is not necessary for determinacy, e.g. due to departures from full-information rational expectations (Angeletos & Lian, 2022).

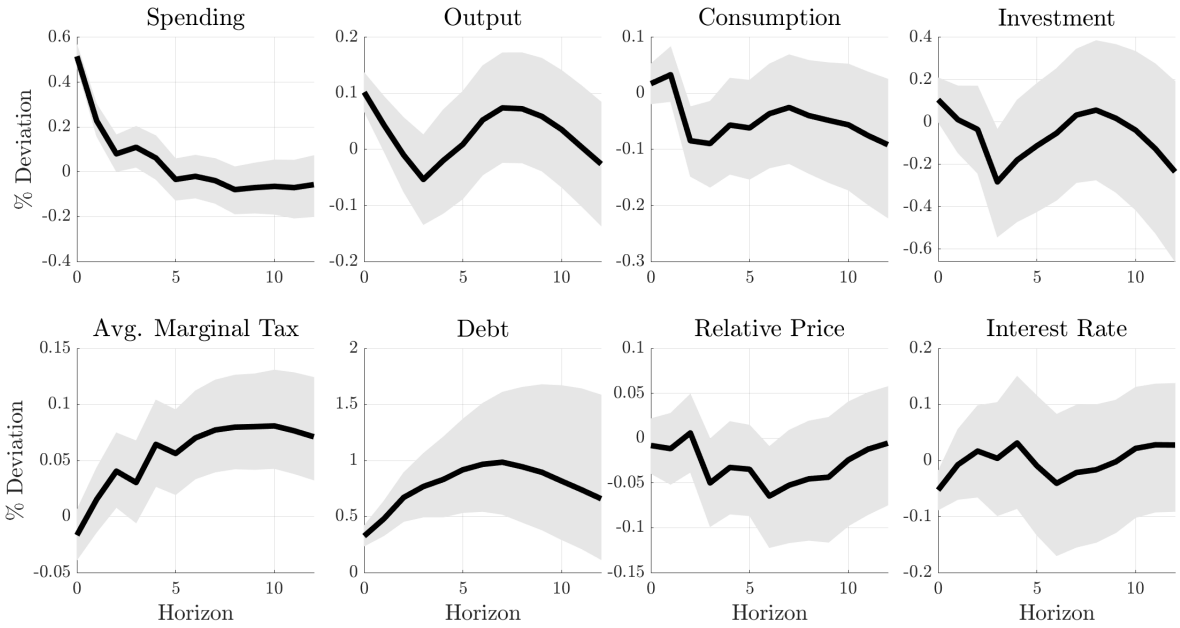

Figure E.1: Impulse responses after a one standard deviation innovation to the forecast error, quarterly frequency. The grey areas correspond to 16th and 84th percentile confidence bands, constructed using 10,000 draws from the posterior distribution of the reduced-form VAR parameters.

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. Depending on the set of included controls, the undershooting of consumption and investment is, in this earlier sample, usually more pronounced (similar to Ramey, 2011). However, the undershooting then goes hand-in-hand with a similar undershooting of spending itself, invalidating the required alignment of excess demand paths.[54] For earlier versions of this paper I also considered longer samples, going to 2015Q1. Results in this expanded sample suggested that impact consumption crowding-in is actually slightly stronger, broadly consistent with standard intuition on zero lower bound effects. These results are, however, not particularly robust, similar to Ramey & Zubairy (2018) and Debortoli et al. (2019). Second, replacing my benchmark measure of government spending forecast errors with Greenbook defense spending forecast errors gives similar output and consumption responses (see Figure F.3), with the

[54]Note, however, that this dynamic under-shooting of consumption and output is—depending on the set of controls, and on when exactly the sample starts—insignificant or only barely statistically significant at the 95 per cent level. Furthermore, again depending on the set of controls, I in those specifications also find delayed (and again largely insignificant) over-shooting of output.

main difference being that now there is some evidence of impact consumption crowding-in. Overall this similarity in the results suggests that either (i) the benchmark exercise itself is largely picking up the response to military spending forecast errors or (ii) multipliers are invariant to the spending type (similar to the conclusion in the meta study of Gechert (2015)). The high correlation between my baseline forecast error series and the defense forecast error series provides some important evidence in favor of the former interpretation. Third, dropping the quadratic time trend has some effects on far-out impulse responses, but not on the short-run responses that I emphasize.

Blanchard & Perotti (2002). While my main application in Appendix F.1 uses only the Ramey (2011) shock, the other applications rely on a second shock measure, constructed following Blanchard & Perotti (2002). I estimate the same VAR as above, but now additionally impose the assumption that the innovation to the equation for government spending g_t itself is *also* a structural fiscal shock. This identification scheme (which is identical to Blanchard & Perotti) assumes that the equation for g_t is in fact the correctly specified government spending rule; by the discussion in Section 3.2, the residual innovation to that rule is then likely to reflect a combination of contemporaneous and news policy shocks. Indeed, and consistent with prior work, I find that this alternative approach identifies a government spending innovation that induces a somewhat more persistent response of fiscal purchases than the professional forecast errors. Qualitatively, the impulse responses of other aggregates—in particular output, consumption and investment—look very similar to those for my benchmark Ramey identification. The most notable difference is that I now find moderate consumption and investment crowding-in, consistent with the findings of Caldara & Kamps (2017) for even more persistent innovations in government spending.[55]

Importantly, because both sets of impulse responses—for the Ramey (2011) shock and the Blanchard & Perotti (2002) shock—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 (11).

[55]One possible rationalization for these empirical findings is provided in Dupaigne & Fève (2016). More persistent government spending shocks lead to a more persistent boom and so in particular a more persistent expansion in aggregate employment, which in turn may prompt firms to increase their capital stock. We can thus see investment crowding-*in* and, with sufficiently high MPCs, even moderate consumption crowding-in. The overall result is a cumulative multiplier slightly above one, consistent with the results I document.

F Applications

This appendix presents several applications of the demand equivalence methodology. I begin in Appendices F.1 and F.2 with stimulus checks—the headline application of Section 3. I then study bonus depreciation in Appendix F.3 and income redistribution in Appendix F.4.

F.1 Stimulus checks

I implement my empirical method exactly as promised in Section 3.3: I first characterize the policy's direct effects on spending, then estimate a fiscal spending shock with the required properties, and finally construct the desired aggregate counterfactual.

Direct effect. As discussed in Section 3.3 the empirical estimates of Parker et al. (2013) suggest that the stimulus check policy of 2008 increased consumption demand by around 1.5 per cent of total personal consumption expenditure on impact, and then a further 0.6 per cent in the following quarter. In the left panel of Figure F.1, the two green squares show those two direct consumption responses $\widehat{c}^{PE}_{\tau,0}$ and $\widehat{c}^{PE}_{\tau,1}$; the solid green line then extrapolates those first two MPCs as discussed in Appendix E.1.

General equilibrium aggregation. Given the direct spending response $\widehat{\boldsymbol{c}}^{PE}_{\tau}$, the next step of my methodology requires the researcher to find a fiscal spending shock that satisfies requirements (i)–(iii): a similar time profile, deficit financing, and monetary accommodation.

One suitable candidate to satisfy these very particular and demanding requirements are fiscal shocks identified through professional forecast errors of government purchases (following Ramey, 2011), presented in Appendix E.3. As discussed there, identification through such professional forecast errors relies on the implicit assumption that fiscal purchases react to aggregate macroeconomic conditions only with a lag.[56] Appealingly, such errors embed the large information set of professional forecasters, thus alleviating concerns related to possible shock non-invertibility (Leeper et al., 2013). Finally, and most importantly for my purposes, this fiscal shock (i) leads to a transitory uptick in fiscal purchases, (ii) is deficit-financed, and (iii) is not followed by a meaningful monetary policy reaction. Detailed estimation results for this shock are reported in Appendix E.3.

[56] To the extent that those forecast errors are dominated by military spending (as suggested by the results in Figure F.3), it may alternatively be argued that the forecast errors plausibly measure structural government spending shocks because military spending is exogenous to wider macroeconomic conditions.

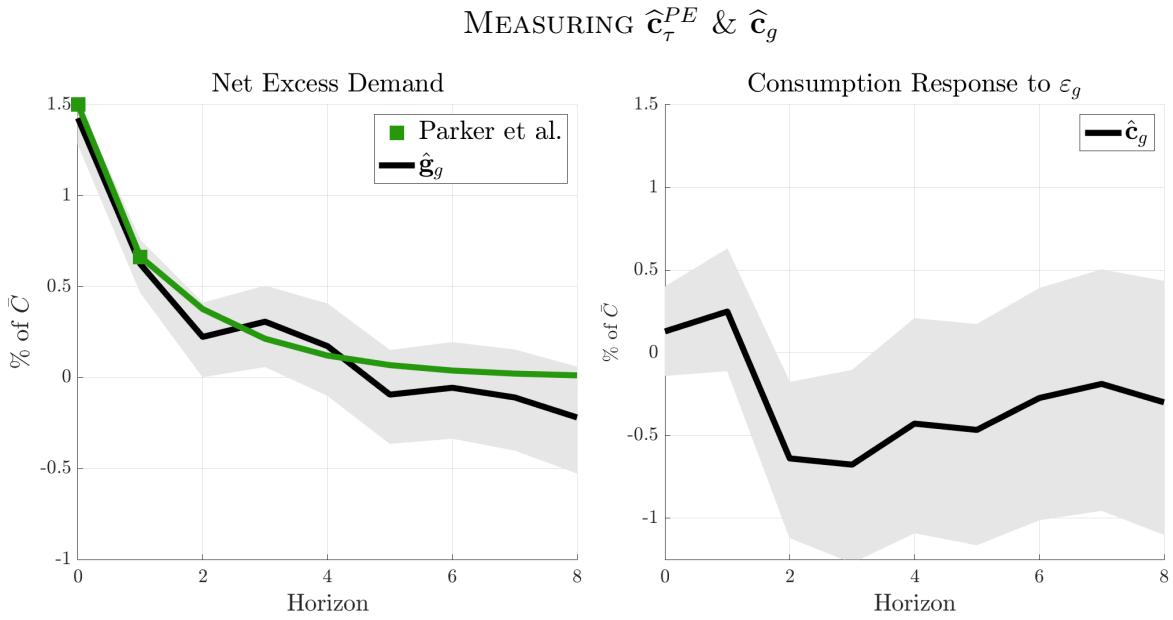

Figure F.1: The left panel shows direct consumption responses to the stimulus check (green) vs. direct government spending response to identified spending shock (black), with 16th and 84th percentile confidence bands (grey), quarterly frequency. Estimated consumption responses from Parker et al. (2013) and Broda & Parker (2014), extrapolated for horizons beyond $t = 1$. The right panel shows the response of consumption to the fiscal spending shock.

The left panel of Figure F.1 reveals that, as required, the estimated increase in fiscal purchases closely mirrors the spending expansion implied by the stimulus check policy, with the targeted $\widehat{\mathbf{c}}_\tau^{PE}$ always remaining within the confidence bands for the estimated $\widehat{\mathbf{g}}_g$.[57] Furthermore, the corresponding estimates for government debt and taxes reported in Appendix E.3 reveal the increase in fiscal purchases to be rather persistently deficit-financed. Finally, I there also show that the spending expansion was indeed largely accommodated by the monetary authority, with nominal interest rates responding very little. It follows that the consumption response to the fiscal experiment $\widehat{\mathbf{c}}_g$, displayed in the right panel of Figure F.1, promises to at the same time tell us about the missing general equilibrium effects of a deficit-financed, one-off stimulus check policy with little monetary offset—exactly the kind of counterfactual relevant for the 2008 stimulus check experiment.

AGGREGATE COUNTERFACTUAL. Figure F.2 puts all the pieces together to present full general equilibrium counterfactuals for stimulus checks. The left panel begins by implement-

[57] In Appendix D.8 I use a structural model to study the inaccuracy associated with demand matching errors of similar magnitude to those observed in Figure F.1.

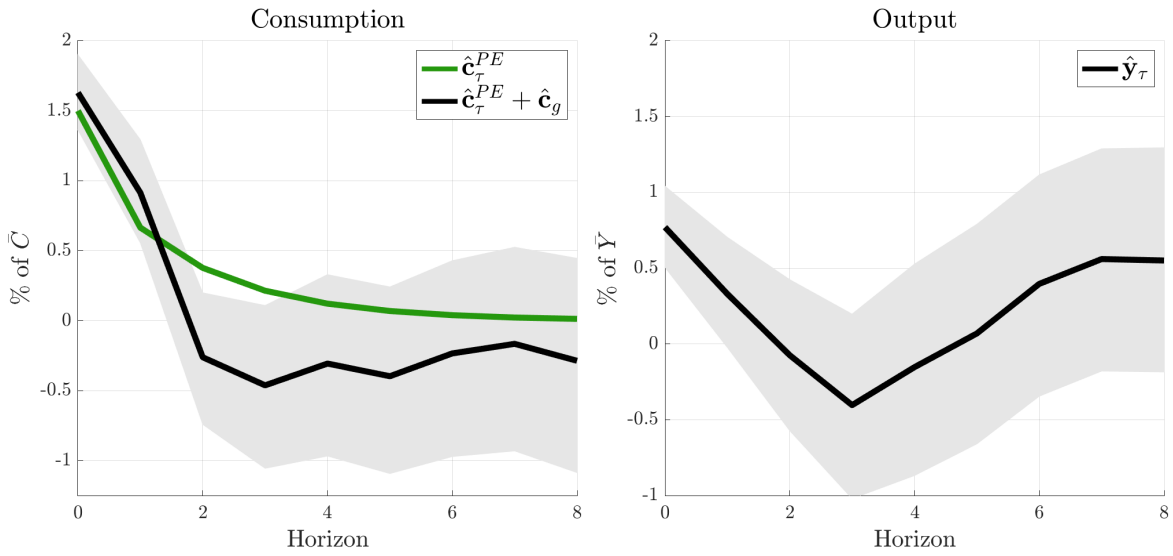

Figure F.2: Consumption and output responses to a stimulus check shock, quarterly frequency. The full consumption response is computed following the exact additive decomposition of Proposition 1, while the output response is simply equal to the response after a government spending shock. The grey areas again correspond to 16th and 84th percentile confidence bands.

ing the demand equivalence decomposition in (6), summing a) the micro-estimated direct spending response $\widehat{\mathbf{c}}_\tau^{PE}$ and b) the response of consumption to the fiscal shock $\widehat{\mathbf{c}}_g$. Since the direct spending effect is large, while the response of private consumption to the fiscal spending expansion is muted, the estimated *aggregate* effect of the policy turns out to be close to the micro-estimated direct effect—my headline takeaway in Section 3.3. The right panel shows the corresponding response of output, which by demand equivalence is the same for the fiscal spending expansion and the stimulus checks. Here I find a significant (if short-lived) total response, with output on impact rising by somewhat less than 1 per cent, and then returning to baseline. Finally I note that I find very similar results—though with slightly stronger initial consumption crowding-in—using instead the defense spending forecast error series from Drautzburg (2020). This series is arguably more plausibly exogenous than the baseline series, but also more obviously subject to concerns about differences in private and public consumption baskets. Results are reported in Figure F.3.

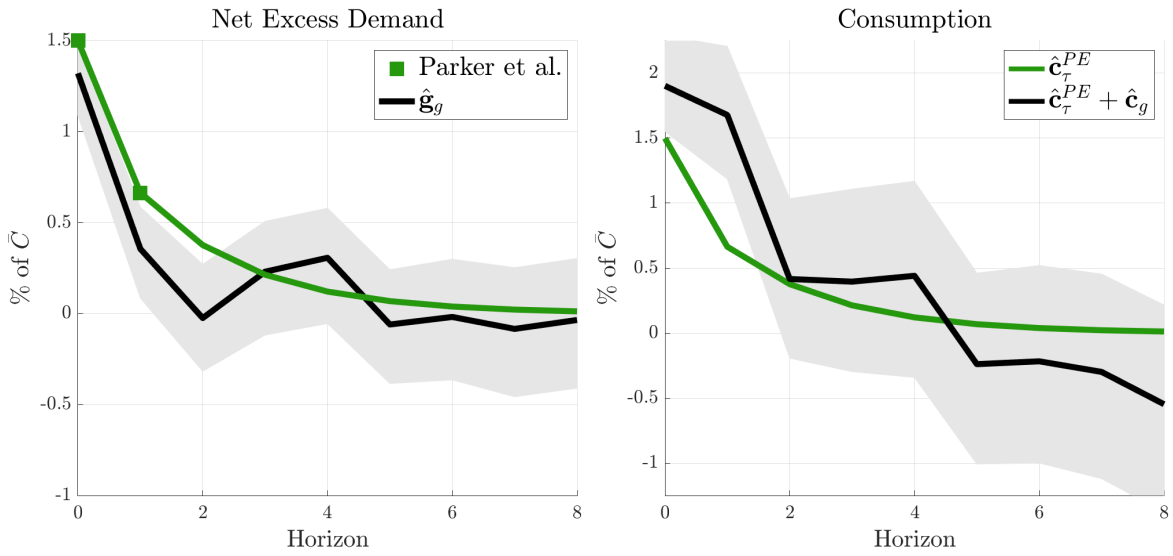

Figure F.3: See the captions for Figure F.1 and Figure F.2. I now use the defense forecast spending error series from Drautzburg (2020).

F.2 Alternative check estimates based on Orchard et al. (2022)

In an important recent contribution, Orchard et al. (2022) study the 2008 stimulus check policy, exactly as I do in my headline application in Section 3.3 and Appendix F.1. They argue for an impact general equilibrium multiplier of stimulus checks of around 0.2 (compared to around 0.5 in my analysis, see Figure F.2). The two main ingredients required to arrive at this conclusion are: (i) substantially lower MPC estimates than in the work of Parker et al. (2013); and (ii) a moderate amount of further general equilibrium dampening.

In this section I discuss the extent to which these results are consistent with my counterfactuals reported in Appendix F.1. I proceed in two steps. First, I apply my aggregation methodology not to the high MPC estimates of Parker et al. used in Appendix F.1, but instead to the lower MPC estimates of Orchard et al.. Second, I discuss the general equilibrium crowding-out effects present in the structural model of Orchard et al..

Direct response. Orchard et al.'s preferred estimate of the immediate spending response is 30 cents for every dollar of stimulus; that is, $MPC_{0,0} = 0.3$. Since they do not report any further dynamics, I restrict the ratio $MPC_{1,0}/MPC_{0,0}$ to be as in my baseline experiment based on Parker et al., and then as in Appendix E.1 impose a constant rate of decay in intertemporal MPCs from date $t = 1$ onwards, with the rate of decay chosen to ensure that

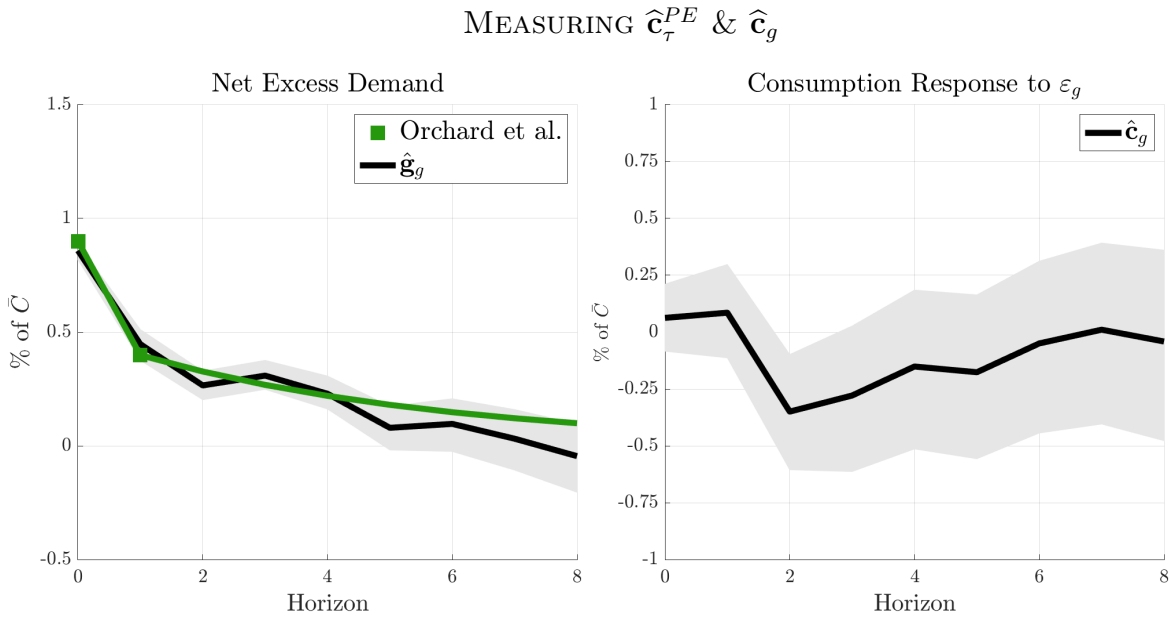

Figure F.4: The left panel shows direct consumption responses to the stimulus check (green) vs. direct government spending response to identified spending shock (black), with 16th and 84th percentile confidence bands (grey), quarterly frequency. Estimated consumption responses from Orchard et al. (2022), extrapolated for horizons beyond $t = 1$ (see text). The right panel shows the response of consumption to the matched linear combination of fiscal spending shocks.

discounted MPCs sum to 1.

The resulting direct spending response is displayed as the green line in the left panel of Figure F.4. Compared to Figure F.1 two main changes are evident. First, the *level* of the impact response is much lower—the lower MPC estimate of Orchard et al. relative to Parker et al.. Second, the dynamic *time profile* is different: since the initial MPC is smaller, the overall spending response is now necessarily more persistent.

Aggregation: the missing intercept. General equilibrium aggregation through demand equivalence now requires empirical evidence on an aggregate fiscal spending shock that: (i) induces a similar spending profile to the green line in Figure F.4; (ii) is persistently deficit-financed; and (iii) is accommodated by the monetary authority. Differently from Appendix F.1 and following my discussion in Section 3.2, I now construct this counterfactual through a *linear combination* of the forecast error and Blanchard & Perotti (2002) fiscal policy shocks discussed in Appendix E.3; intuitively, the former is a little bit too transitory to match a spending profile like that displayed in Figure F.4, while the latter is too persistent. A linear combination thus seems promising, and indeed the figure reveals that a well-chosen

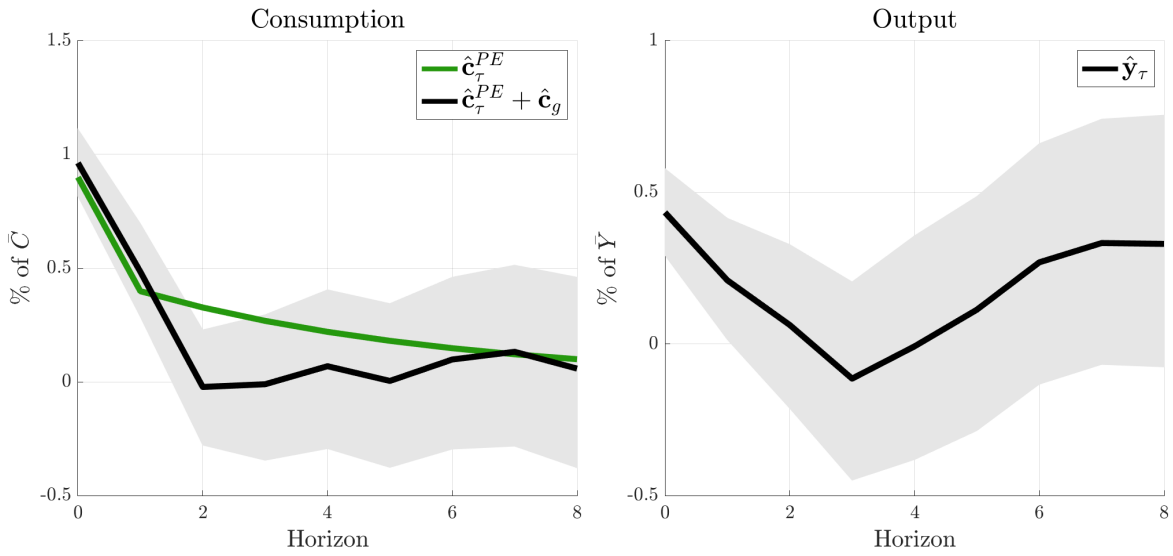

Figure F.5: Consumption and output responses to a stimulus check shock, quarterly frequency. The full consumption response is computed following the exact additive decomposition of Proposition 1, while the output response is simply equal to the response after the matched linear combination of government spending shock. The grey areas again correspond to 16th and 84th percentile confidence bands.

combination (that turns out to mostly load on the forecast error shock) can closely mimic the required spending path. Furthermore, and as in my baseline analysis, I find that this spending expansion is persistently deficit-financed and that there is little monetary offset.

The right panel of Figure F.4 shows the consumption response to this combination of fiscal spending shocks. We see that the time profile is still quite similar to Figure F.1, just scaled down in magnitude. Intuitively, both of the identified fiscal shocks used here to induce the required path of spending lead to only a moderate response of household consumption.

Macro counterfactuals. Aggregate counterfactuals are reported in Figure F.5. The conclusions are qualitatively identical to those of Figure F.2: the general equilibrium consumption response is close to the direct cross-sectional spending estimate on impact, with some slight additional crowding-out in the following quarters. Similarly output increases on impact before returning to trend. The important *quantitative* difference, however, is that all responses are scaled down, consistent with the smaller direct effect of Orchard et al..

Discussion. Orchard et al. argue, based on a careful narrative account of events in 2008, that large aggregate causal effects of the 2008 stimulus checks are inconsistent with aggregate time series evidence. They show that, in a particular structural model with non-Ricardian households, an impact MPC of 0.5—as estimated by Parker et al. and as used in my headline exercise in Appendix F.1—would lead to additional general equilibrium amplification of more than 50 per cent, and so to an aggregate consumption increase of almost 2.5 per cent (which they deem implausible). Their preferred aggregate effect, corresponding to an MPC of 0.3 and some further general equilibrium crowding out, instead amounts to an impact consumption increase of around 0.65 per cent.

In this paper I instead aggregate micro MPC estimates through time series evidence on fiscal multipliers. In this section I have combined the MPC estimates of Orchard et al. with a fiscal spending multiplier of 1, giving me a total consumption response around 1 per cent. This in turn corresponds to a general equilibrium transfer multiplier of around 0.3 (= micro MPC, vs. 0.2 in Orchard et al.) thus lying at the larger end of the range deemed "plausible" by those authors. The additional general equilibrium dampening in the model of Orchard et al. —which moves us from 0.3 to 0.2—mainly comes from the following two channels.

1. *Monetary policy response.* They consider a responsive monetary authority, with a coefficient on inflation in the policy rule of 1.5 (see their Table 1). As discussed previously, my estimates instead correspond to a fiscal expansion that is largely *accommodated* by the monetary authority, with little response of nominal interest rates. It is thus unsurprising that my estimates suggest less general equilibrium crowding-out. Furthermore, for 2008, a counterfactual with monetary accommodation is arguably the more relevant one.

2. *Relative durables price response.* Orchard et al. consider a two-good model with durables and non-durables. In their setting general equilibrium crowding-out is strong because: (i) a very large share of stimulus checks is spent on durables; (ii) relative durable good prices can move substantially since relative prices are flexible and supply is relatively inelastic; and (iii) durable goods demand is highly price-elastic. Importantly, if the government on the other hand purchases a good with less responsive prices, then my demand equivalence approach to aggregation will miss some general equilibrium crowding-out and thus give an upper bound of actual causal effects, as discussed in Section 4.

   How important is this relative price channel likely to have been for the 2008 rebate policy? Note that, because of features (i) - (iii), the structural model of Orchard et al. predicts quite large and persistent increases in the relative price of durable goods (see their Figure

B.3). These large and persistent predicted price responses seem somewhat hard to square with price data for 2008.[58]

To summarize, the most important and empirically relevant difference between my results in Section 3.3 and those of Orchard et al. are related to differences in estimated micro MPCs, not to any further second-round general equilibrium effects. The analysis in Figure F.5 has demonstrated that my methodology can be applied just as well to their preferred MPC estimates, resulting in general equilibrium counterfactuals that do not imply the pronounced "V-shapes" in consumer spending deemed implausible by Orchard et al.. Future work that further improves measurement of the direct micro MPCs would be very welcome.

F.3 Bonus depreciation

I use the demand equivalence approach to estimate counterfactuals for aggregate investment, output and consumption following an expansionary bonus depreciation stimulus policy.

Direct response. My estimates of the direct response of investment rely on Zwick & Mahon (2017) and Koby & Wolf (2020), who exploit cross-sectional firm-level heterogeneity in the exposure to bonus depreciation investment stimulus. Koby & Wolf estimate dynamic regressions akin to (9) and give sufficient conditions under which the regression estimands are identical to or at least informative about the desired direct investment spending responses $\widehat{\mathbf{i}}_q^{PE}$. The discussion is largely analogous to that in Proposition 2 (see Appendix E.2).

With the direct investment spending response $\widehat{\mathbf{i}}_q^{PE}$ thus measured, it remains to recover the corresponding output path $\widehat{\mathbf{y}}_q^{PE}$. In the absence of direct measurement of this path, I propose to construct it by imposing the same production function—a simple Cobb-Douglas production function in capital and labor, potentially with decreasing returns to scale—and the same competition structure—the separation of intermediate goods producers, retailers and aggregators—as in my baseline structural model. Under those assumptions we get

$$\widehat{\tilde{y}}_{qt}^{PE} = \frac{\alpha\nu}{1-(1-\alpha)\nu} \times \widehat{\tilde{k}}_{qt-1}^{PE} \tag{F.1}$$

[58]The BLS price series for new vehicles (series ID `CUSR0000SETA01`) shows a decline -0.91 per cent in the half-year between April and September 2008, while the CPI less food and energy (series ID `CUSR0000SA0L1E`) shows an increase of 1 per cent. For reference, in the prior half year, car prices declined by -0.65 per cent, while the CPI increased by 1.1 per cent, so if anything relative prices go in the wrong direction, even accounting for trends. Such an absence of relative price responses for transitory shocks is also consistent with my findings in Figure E.1.

I set $\alpha = 0.2$, $\nu = 1$ and $\delta = 0.016$, in line with standard modeling practice in general and my estimated HANK model in particular.[59]

I take the regression estimates of $\widehat{\mathbf{i}}_{qt}^{PE}$ for $t = 0, 1, 2, 3$ straight from Koby & Wolf (2020, Table 1). The green squares in the investment panel of Figure F.6 show the estimated path of direct investment spending responses to a one-quarter bonus depreciation shock worth around 8 cents, a shock similar in magnitude to the stimulus of 2008-2010, and applied to all investment. The solid green line extrapolates the empirical estimates to a full investment demand response path $\widehat{\mathbf{i}}_{q}^{PE}$, as discussed further in Appendix E.2.

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.

Aggregation: the missing intercept. Following (17), it remains to replicate the estimated net excess demand path through a suitable list of government spending shocks:

$$\widehat{\mathbf{i}}_{q}^{PE} - \widehat{\mathbf{y}}_{q}^{PE} \quad = \quad \sum_{k=1}^{n_k} \gamma_k \times \widehat{\mathbf{g}}_{g_k} \tag{F.2}$$

It is unlikely that any single estimated government spending shock can replicate the reversal documented in Figure F.6. Encouragingly, much previous work on fiscal multipliers actually estimates the effects of *delayed* increases in government spending (Blanchard & Perotti, 2002; Caldara & Kamps, 2017)—that is, government spending news shocks. In principle, combining these delayed spending responses with the immediate Ramey (2011) spending effect estimated in the first application in Appendix F.1 should allow me to replicate the net demand effects of the investment tax credit.

I operationalize this insight by jointly identifying the forecast error shock as well as the

[59]Note that (F.1) heavily leverages the fact that competition among intermediate goods producers is perfect, so nominal rigidities only matter in *general equilibrium*, via feedback through intermediate goods prices p_t^I. This assumption is popular in structural modeling (e.g. Ottonello & Winberry, 2018). A more robust approach, of course, would be to directly measure $\widehat{\mathbf{y}}_{q}^{PE}$. I leave such an extension to future work.

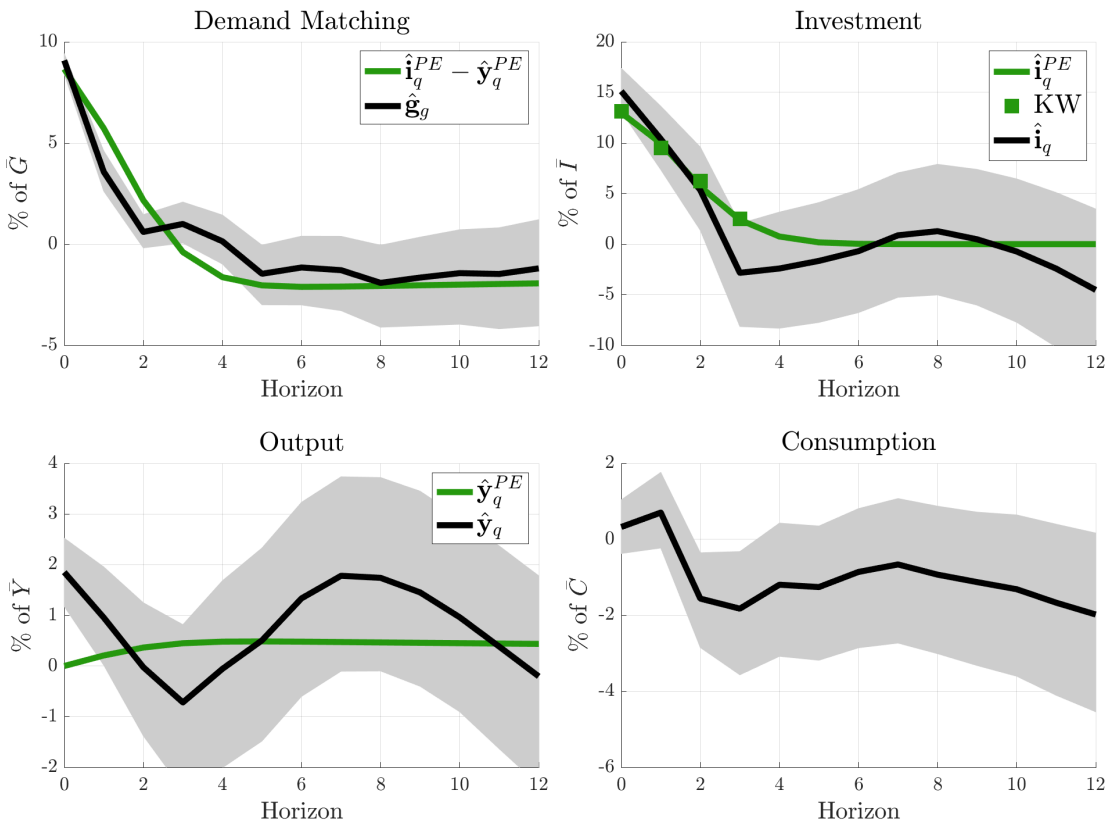

Figure F.6: Investment, output and consumption responses to an investment tax incentive shock, quarterly frequency, with the partial equilibrium net excess demand path matched to a linear combination of government spending shocks. "KW" refers to Koby & Wolf (2020); details are given in Appendix E.2. The investment and output responses are computed in line with (17) - (18), while the consumption response is just the response after the identified combination of government spending shocks. The grey areas again correspond to 16th and 84th percentile confidence bands.

Blanchard & Perotti shock in a single VAR, as discussed in Appendix E.3 and as also done in Appendix F.2. Since the effects of the Blanchard & Perotti shock are more delayed, a linear combination of the two shocks allows me to match the implied net excess demand path of the investment demand shock, as shown in the top left panel of Figure F.6. Note that further details on the implementation—in particular the construction of standard errors for general equilibrium feedback—are provided in Appendix E.3.

Macro counterfactuals. All results for general equilibrium counterfactuals are displayed in Figure F.6. With the requirement that $\widehat{\mathbf{g}}_g = \widehat{\mathbf{i}}_q^{PE} - \widehat{\mathbf{y}}_q^{PE}$ satisfied, the investment

and output panels implement the additive decompositions in (17) and (18), respectively. My main finding is that the substantial partial equilibrium investment demand responses estimated in Zwick & Mahon (2017) and Koby & Wolf (2020) also survive in general equilibrium. The increase in investment demand is accommodated through a sharp immediate increase in output as well as a smaller and somewhat delayed drop in consumption.[60] Taken together, the large direct investment spending responses estimated in micro data as well as prior evidence on the transmission of aggregate government spending shocks suggest that bonus depreciation investment incentives provide a sizable macroeconomic stimulus.

F.4 Income redistribution

As my final application I use the demand equivalence approach to estimate the response of aggregate consumption to a short-lived increase in (labor) income inequality. My analysis here builds on the important prior contribution of Auclert & Rognlie (2018). Those authors first use a partial equilibrium model of the consumption-savings decision to recover the direct effect of the shock, and then aggregate using a general equilibrium closure of that model. I follow the first step of their analysis, but then use *empirical* evidence on aggregate government spending shocks to provide the general equilibrium aggregation.[61]

Direct response. Similar to Auclert & Rognlie (2018), I recover the direct response of aggregate consumer net excess demand through a structural, partial equilibrium consumption-savings problem—the consumption-savings problem of my quantitative HANK model of Section 4.1. Specifically, I change the household budget constraint to

$$c_{it} + b^h_{it} = (1-\tau_\ell) w_t e_{it} \ell_{it} (1 + \varepsilon_{z,t} z_{it}) + \frac{1+i^b_{t-1}}{1+\pi_t} b^h_{it-1} + \tau_{it} + d_{it}, \quad b^h_{it} \geq \underline{b}$$

[60] Note that, with technology fixed and capital pre-determined, the impact increase in output must reflect an increased use of the other factor of production: labor. This is entirely consistent with my assumption of weak wealth effects in the first part of the paper: for example, with Assumption A.3 holding because of sticky wages, hours worked will increase in *general equilibrium* because the intermediate goods price has increased, pushing up firm labor demand at $t = 0$.

[61] As discussed in Section 3.1, the first step could also come from data—in this case cross-sectional evidence on the spending response to income redistribution. However, empirical evidence here is less clear than for stimulus checks, so I follow a structural approach instead. This has the added benefit of illustrating a second way of implementing my methodology: the direct response comes from a *partial equilibrium* model, while the difficult question of general equilibrium aggregation is fully addressed through empirical evidence.

Figure F.7: Consumption, debt and tax responses to a redistribution shock, with the partial equilibrium net excess demand path matched to a linear combination of government spending shocks. The consumption response is computed in line with Proposition 1. The plot also shows the required demand matching and financing alignment. The dashed lines again correspond to 16th and 84th percentile confidence bands.

where $\varepsilon_{z,t}$ is an aggregate inequality shock, $\int_0^1 e_{it} z_{it} di = 0$, and $z_{it} \propto e_{it}$. A one-off shock $\varepsilon_{z,0} > 0$ thus leads to a one-period increase in labor income inequality, with more productive households receiving a larger share of total aggregate income. Solving the partial-equilibrium consumption-savings problem for all households i given $\varepsilon_{z,0}$ and then aggregating, we recover the direct response path $\widehat{\mathbf{c}}_z^{PE}$ for the inequality shock $\boldsymbol{\varepsilon}_z$. I scale the shock to lead to a decline in consumption demand of 1 per cent on impact; the green lines in the top panels of Figure F.7 show the implied full consumption response path.

AGGREGATION: THE MISSING INTERCEPT. Following the discussion in Section 5.1 it remains to replicate the net excess demand path $\widehat{\mathbf{c}}_z^{PE}$ through a suitable list of government spending shocks:

$$\widehat{\mathbf{c}}_z^{PE} \quad = \quad \sum_{k=1}^{n_k} \gamma_k \quad \times \quad \widehat{\mathbf{g}}_{g_k} \tag{F.3}$$

Note that, since the increase in inequality is a non-policy shock, the path $\widehat{\mathbf{c}}_z^{PE}$ has zero net present value, so any list of government spending shocks satisfying (F.3) also necessarily has zero net present value. The spending change itself thus in principle can be—and for our purposes needs to be, by condition (ii)—purely deficit-financed. I proceed as in Appendices F.2 and F.3 and use a combination of transitory and persistent fiscal spending changes to align the excess demand paths (top left panel). The two bottom panels of Figure F.7 show that, as required, the fiscal expansion is deficit-financed, with little response of taxes.[62]

MACRO COUNTERFACTUALS. The top right panel of Figure F.7 shows the desired aggregate general equilibrium consumption counterfactual. The macro-equivalent fiscal contraction leaves consumption largely unchanged on impact, before then leading to an increase in spending. We thus conclude that the temporary increase in inequality leads to a significant contraction of consumption on impact (mirroring the direct spending effect), before then leading to a delayed boom.

[62]Of course taxes can and generally will respond to both shocks *in general equilibrium*. With condition (i) of aligned demand paths satisfied, we know that the path $\boldsymbol{g}_g$ requires no *direct* financing. Encouragingly, the bottom right panel of Figure F.7 gives no reason to believe that there was any kind of important direct tax financing.

G Further proofs and auxiliary lemmas

G.1 Proof of Lemma A.1

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

1. Given $(\mathbf{i}^b, \mathbf{y})$, the Taylor rule of the monetary authority allows us to back out the path of inflation $\boldsymbol{\pi}$. Thus all inputs to the firm problem are known,[63] so indeed $\mathbf{s}^f = \mathbf{s}^f(\mathbf{x})$. We thus obtain $\mathbf{y}$, $\mathbf{i}$ and $\boldsymbol{\ell}^f$. Setting $\boldsymbol{\ell} = \boldsymbol{\ell}^f$ and since $\boldsymbol{\tau}^e \in \mathbf{x}$, all inputs to the household problem are known, so indeed $\mathbf{s}^h = \mathbf{s}^h(\mathbf{x})$. We can thus also solve for the path of consumption, so that indeed $\mathbf{s}^u = \mathbf{s}^u(\mathbf{x}; \boldsymbol{\varepsilon})$, and we finally recover union labor supply.

2. Optimal household, firm and government behavior is assured by assumption. It thus remains to check that (i) all markets clear; (ii) the input path of output is consistent with firm production; and (iii) the lump-sum tax 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 endogenous quantities for the government budget constraint are $(\mathbf{i}^b, \boldsymbol{\pi}, \mathbf{w}, \boldsymbol{\ell})$—ensures that the government budget constraint holds period-by-period and that $\lim_{t\to\infty} \widehat{b}_t = 0$, by definition of $\boldsymbol{\tau}^e(\bullet)$.

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

G.2 Proof of Proposition 2

The proof proceeds in three steps. First, I show that *aggregate* impulse responses to the heterogeneous shocks $\{\varepsilon_{\tau i0}\}$ are identical to impulse responses to the common aggregate shock $\varepsilon_{\tau 0} \equiv \int_0^1 \varepsilon_{\tau i0}$. Second, I prove that $\widehat{\mathbf{c}}_{\tau i} - \widehat{\mathbf{c}}_\tau = (\xi_{\tau i0} - 1) \times \widehat{\mathbf{c}}_\tau^{PE} + \boldsymbol{\zeta}_i$, where $\int_0^1 (\xi_{\tau i0} - 1)\boldsymbol{\zeta}_i di = \mathbf{0}$. And third, I exploit standard properties of fixed-effects regressions to complete the argument. As in the proof of Proposition 1, 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.

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

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

$$\widehat{\mathbf{c}}_\tau^{PE} \equiv \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} \times \xi_{\tau i 0} \times \boldsymbol{\varepsilon}_\tau di$$

where $\mathbf{c}_i(\bullet)$ is the consumption function of individual i, defined analogously to the aggregate consumption function $\mathbf{c}(\bullet)$. Since $\int_0^1 \xi_{\tau i 0} di = 1$ and since $\xi_{\tau i 0}$ is assigned randomly across households (and so does not correlate with $\frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} \times \boldsymbol{\varepsilon}_\tau$ at any t), we have that

$$\widehat{\mathbf{c}}_\tau^{PE} = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_{\tau_0}} \times \boldsymbol{\varepsilon}_\tau di \times \left[1 + \int_0^1 (\xi_{\tau i 0} - 1) di\right] = \int_0^1 \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} \times \boldsymbol{\varepsilon}_\tau di$$

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

2. Consumption of household i along the transition path satisfies

$$\widehat{\mathbf{c}}_{i\tau} = \frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} \times \widehat{\mathbf{x}} + \frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} \times \xi_{\tau i 0} \times \boldsymbol{\varepsilon}_\tau$$

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

$$\widehat{\mathbf{c}}_{\tau i} - \widehat{\mathbf{c}}_\tau = (\xi_{\tau i 0} - 1) \times \underbrace{\frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau} \times \boldsymbol{\varepsilon}_\tau}_{\widehat{\mathbf{c}}_\tau^{PE}} + \underbrace{\left(\frac{\partial \mathbf{c}_i}{\partial \mathbf{x}} - \frac{\partial \mathbf{c}}{\partial \mathbf{x}}\right) \times \widehat{\mathbf{x}} + \xi_{\tau i 0} \left(\frac{\partial \mathbf{c}_i}{\partial \boldsymbol{\varepsilon}_\tau} - \frac{\partial \mathbf{c}}{\partial \boldsymbol{\varepsilon}_\tau}\right) \times \boldsymbol{\varepsilon}_\tau}_{\equiv \boldsymbol{\zeta}_i}$$

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

3. By the standard properties of fixed-effects regression, we can re-write regression (9) as

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

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

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

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

□

G.3 Proof of Proposition 3

Following the same steps as in the proof of Proposition 1, but without imposing Assumption 3, we get the two direct shock responses as

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

The general equilibrium response paths of consumption thus now satisfy

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

By properties (i) and (ii) of the fiscal spending shock, we can combine the two expressions above to get

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

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

as claimed. □

G.4 Auxiliary lemma for Proposition A.1

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

$$\mathbf{c} + \mathbf{i}(\mathbf{s}^f(\mathbf{c}); \boldsymbol{\varepsilon}) + \mathbf{g}(\boldsymbol{\varepsilon}) = \mathbf{y}(\mathbf{s}^f(\mathbf{c}); \boldsymbol{\varepsilon}) \tag{G.2}$$

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

To prove Lemma G.1 I as before proceed in two steps. First, I show that all relevant inputs to the firm problem can be obtained as functions only of $\mathbf{c}$ and $\boldsymbol{\varepsilon}$. Second, I show sufficiency of the aggregate market-clearing equation.

1. By Assumptions A.2 and A.3, the household block admits aggregation to a single representative household with period felicity function $u(c) - v(\ell)$. Given $\mathbf{c}$, the Euler equation of the representative household allows us to back out the path of real interest rates $\mathbf{r}$. Given $\mathbf{r}$, the Fisher equation and the Taylor rule of the monetary authority (by Assumption A.4) allow us to recover the paths of nominal interest rates $\boldsymbol{i}^b$ and aggregate inflation $\boldsymbol{\pi}$, and so by the NKPC of retailers we recover $\mathbf{p}^I$. Next, given Assumption A.3, the wage-NKPC allows us to recover the path of real wages $\mathbf{w}$. Together with $\boldsymbol{\varepsilon}$ we thus have all inputs to the firm problem, and in particular indeed $\mathbf{s}^f = \mathbf{s}^f(\mathbf{c})$, as claimed.

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

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

G.5 Proof of Proposition A.1

By Lemma G.1, a perfect foresight equilibrium is, to first order, a solution to

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

As before, we thus in general have

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

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

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

and

$$\widehat{\mathbf{i}}_g \quad = \quad \mathbf{0} \quad + \quad \frac{\partial \mathbf{i}}{\partial \mathbf{c}} \times \mathcal{H} \times \widehat{\mathbf{g}}_g$$

respectively. This establishes (17). The equations for output are exactly analogous. □

G.6 Proof of Corollary C.1

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

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

where $\mathbf{e}$ is now the aggregated optimal household expenditure function for durable and non-durable consumption. Applying the same steps as in the proof of Proposition 1 to this new system, the result follows. □