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:

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

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,

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

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:

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

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

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

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,*

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

the bond market clears,

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

and the labor market clears,

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

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

B.2 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:

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

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

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

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

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

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,*

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

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.

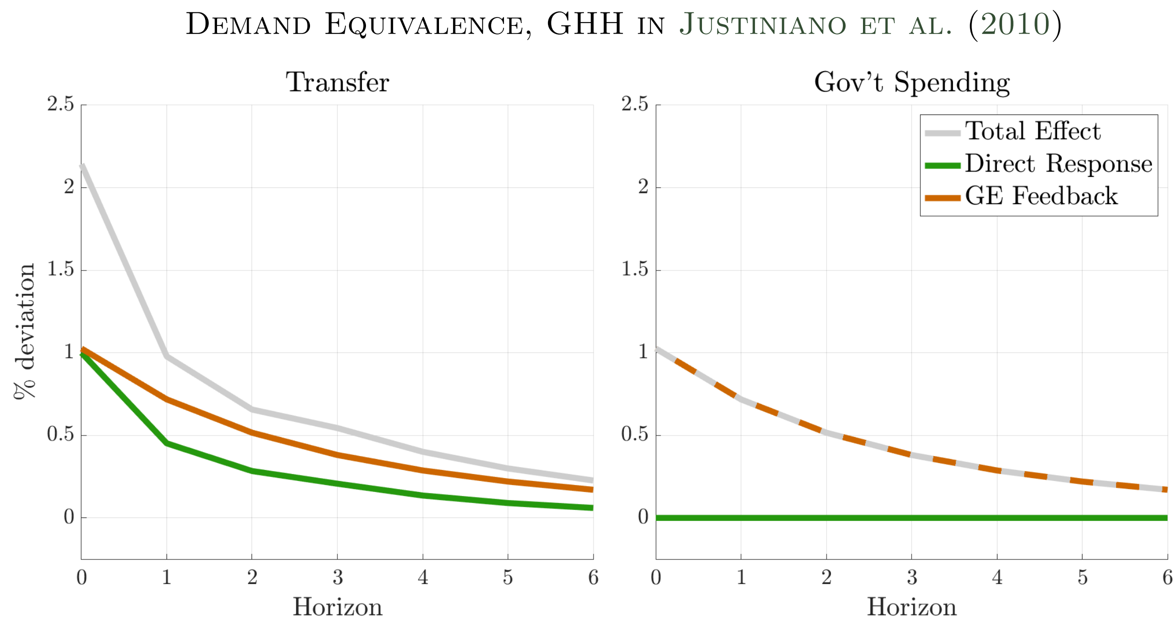

Figure C.1: Consumption impulse response decompositions after impatience and government spending shocks in the model of Justiniano et al. (2010) with GHH preferences, where the two shocks are selected to have identical effects on net excess demand, displayed as the green line in the left panel. The direct 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 $(\widehat{\boldsymbol{\tau}} = \widehat{\boldsymbol{\tau}}^x + \widehat{\boldsymbol{\tau}}^e(\widehat{\boldsymbol{\tau}}^x, \widehat{\mathbf{r}}), \widehat{\mathbf{r}})$ 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.

Demand Equivalence Error Distribution, Random Draws

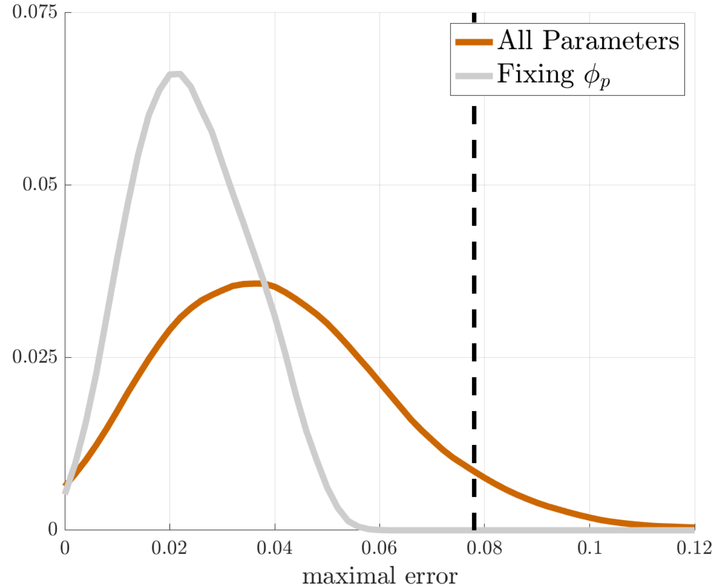

Figure D.1: Kernel estimate of maximal error distribution, with parameters drawn randomly according to Table D.1 (orange, 95th percentile black dashed). The grey lines show the same kernel density estimate when ϕ_p is fixed at its estimated posterior mode.

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.

Approximate Demand Equivalence, Justiniano et al. (2010)

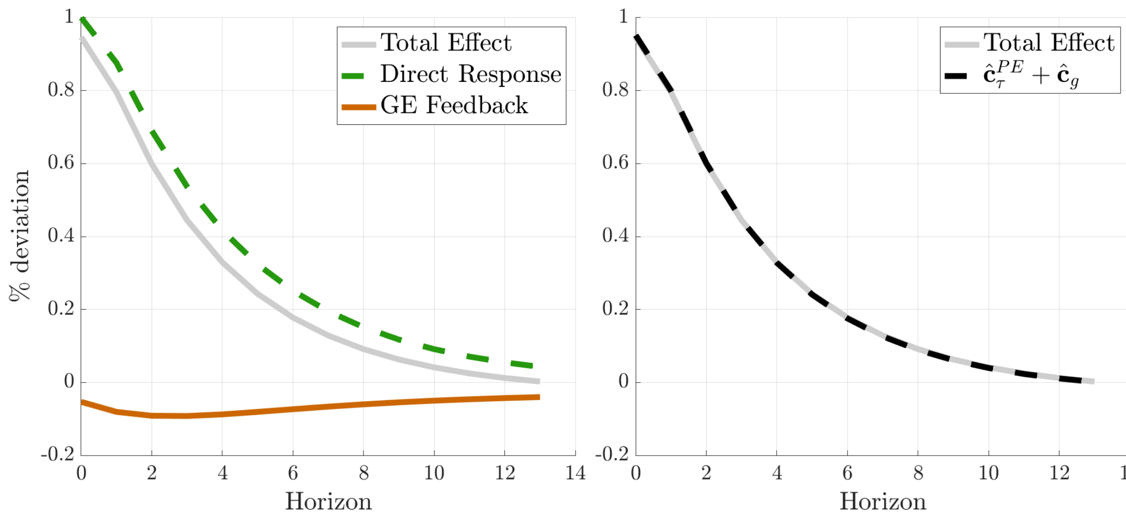

Figure D.2: Consumption impulse response decompositions and demand equivalence approximation in the model of Justiniano et al. (2010), solved at the posterior mode and for an impatience shock with persistence $\rho_b = 0.1$. The direct response and the indirect general equilibrium feedback are computed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

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

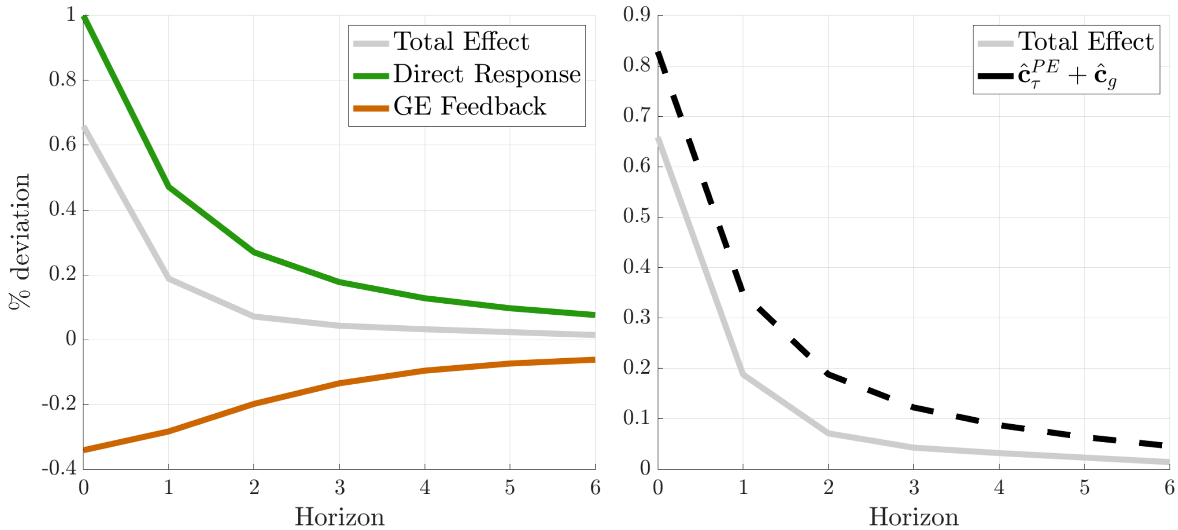

Figure D.3: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with (nearly) flexible prices & wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

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.

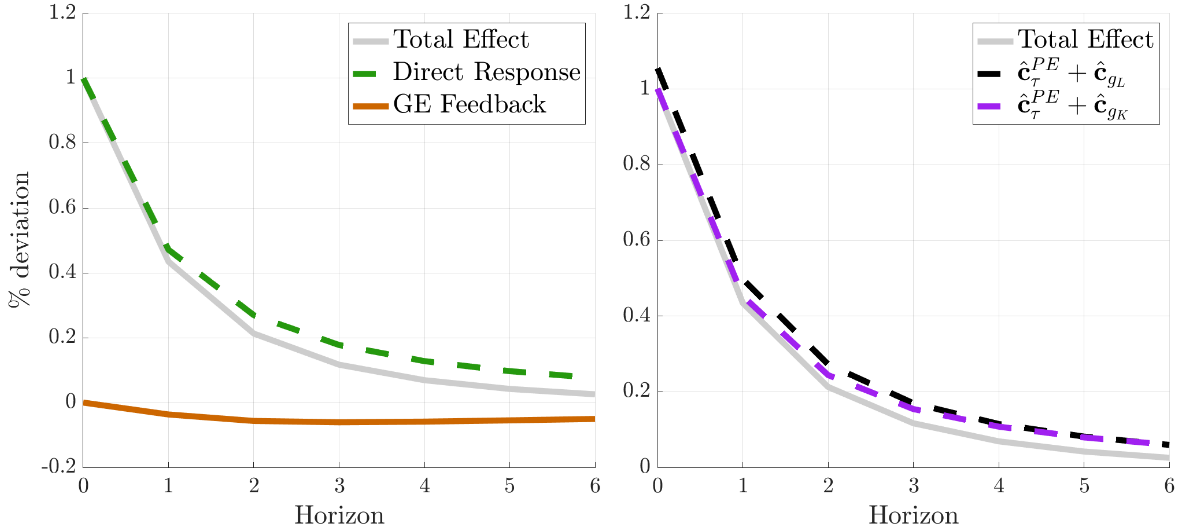

Figure D.4: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with multiple goods. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

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.

Approximate Demand Equivalence, Productive Government Spending

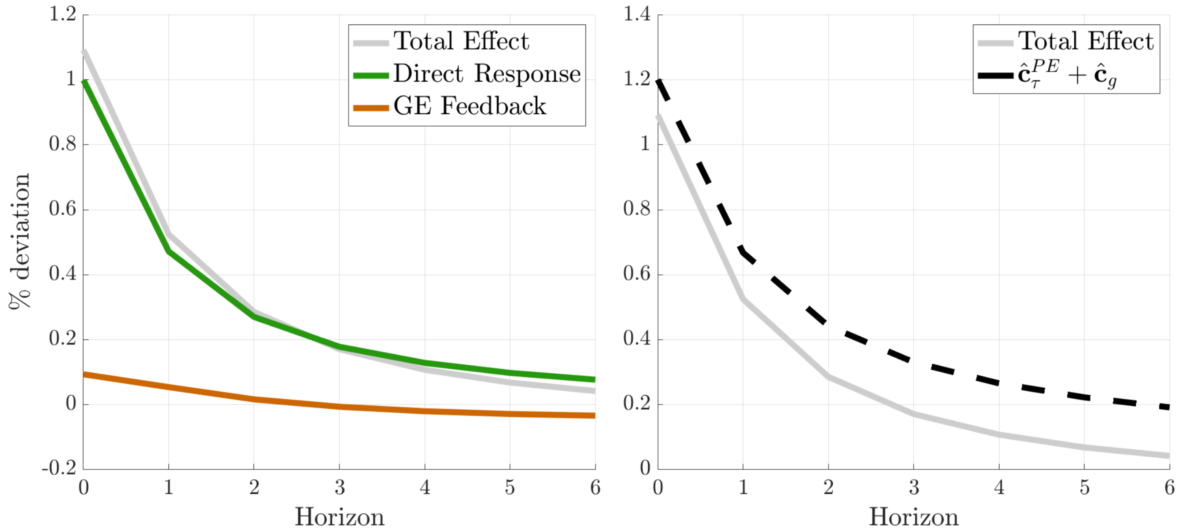

Figure D.5: Consumption impulse response decompositions and demand equivalence approximation for the HANK model with productive government spending. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

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

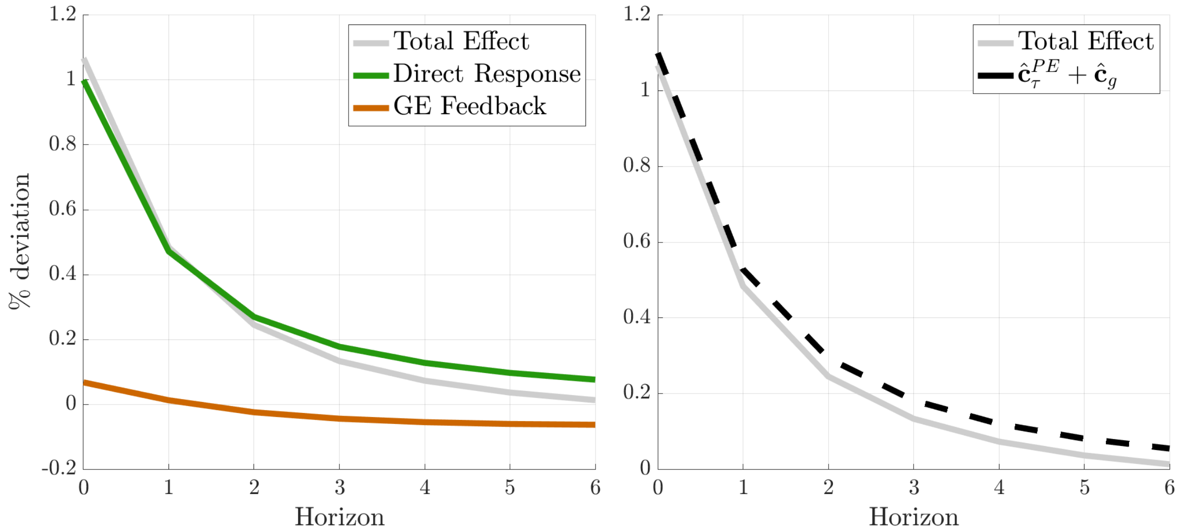

Figure D.6: Consumption impulse response decompositions and demand equivalence approximation for the open-economy HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Approximate Demand Equivalence, Two-Asset HANK Model

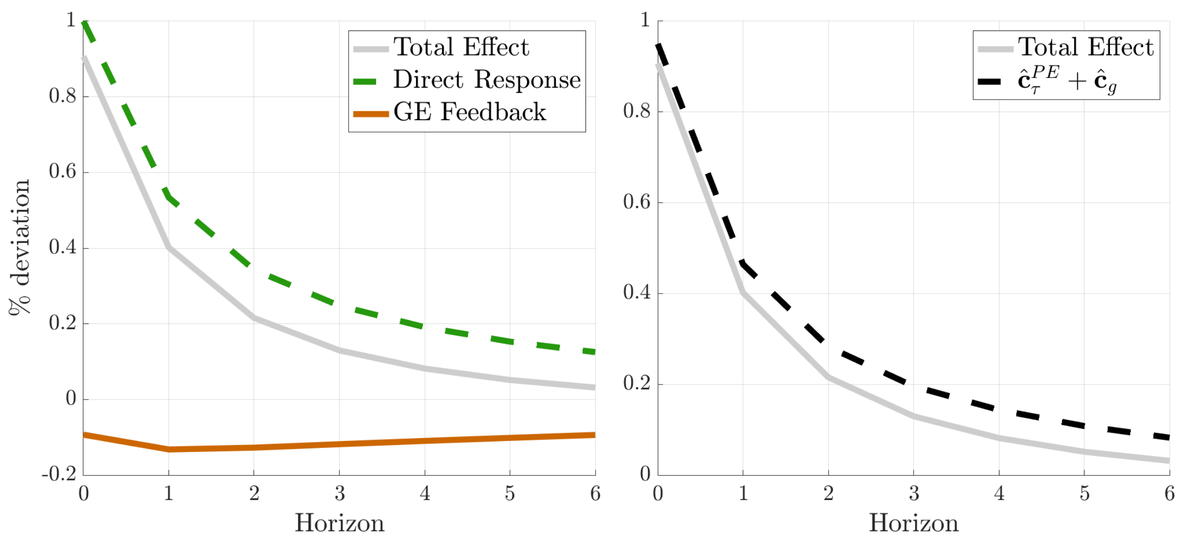

Figure D.7: Consumption impulse response decompositions and demand equivalence approximation for the two-asset HANK model. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

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.

Approximate Demand Equivalence, Imperfect Matching

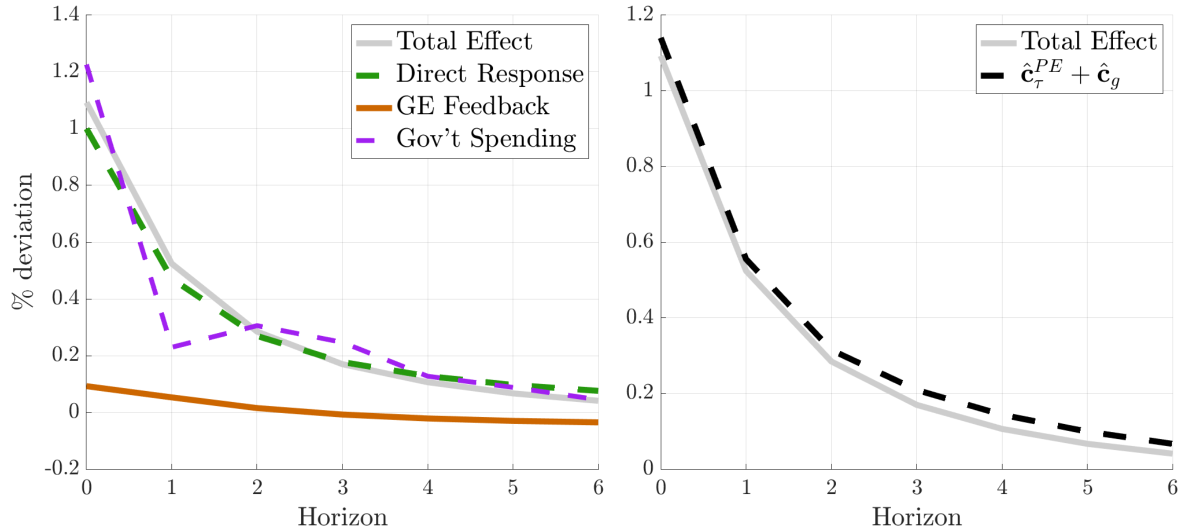

Figure D.8: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, with imperfect demand matching, following (D.1). The direct response and the indirect general equilibrium feedback are computed following Definition 1.

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 $\{\widehat{i}_{qt}^{PE}\}_{t=0}^{3}$, I recover the full partial equilibrium investment response $\widehat{\boldsymbol{i}}_q^{PE}$ 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).

Ramey (2011) Government Spending Shock, VAR IRFs

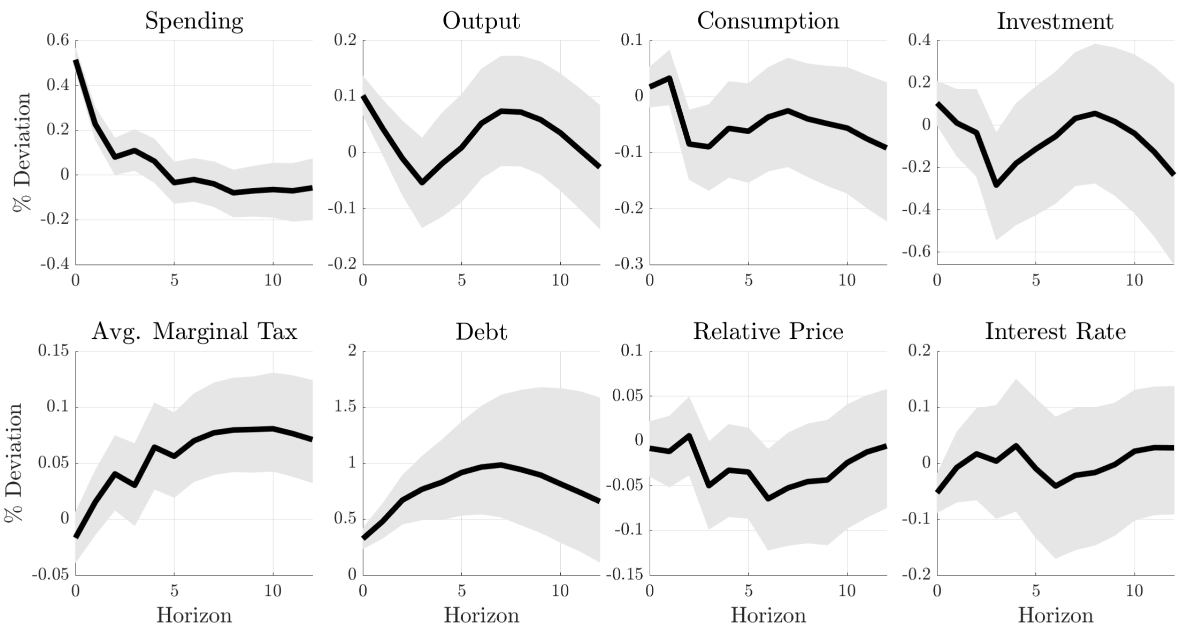

Figure E.1: Impulse responses after a one standard deviation innovation to the forecast error, quarterly frequency. The grey areas correspond to 16th and 84th percentile confidence bands, constructed using 10,000 draws from the posterior distribution of the reduced-form VAR parameters.

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.

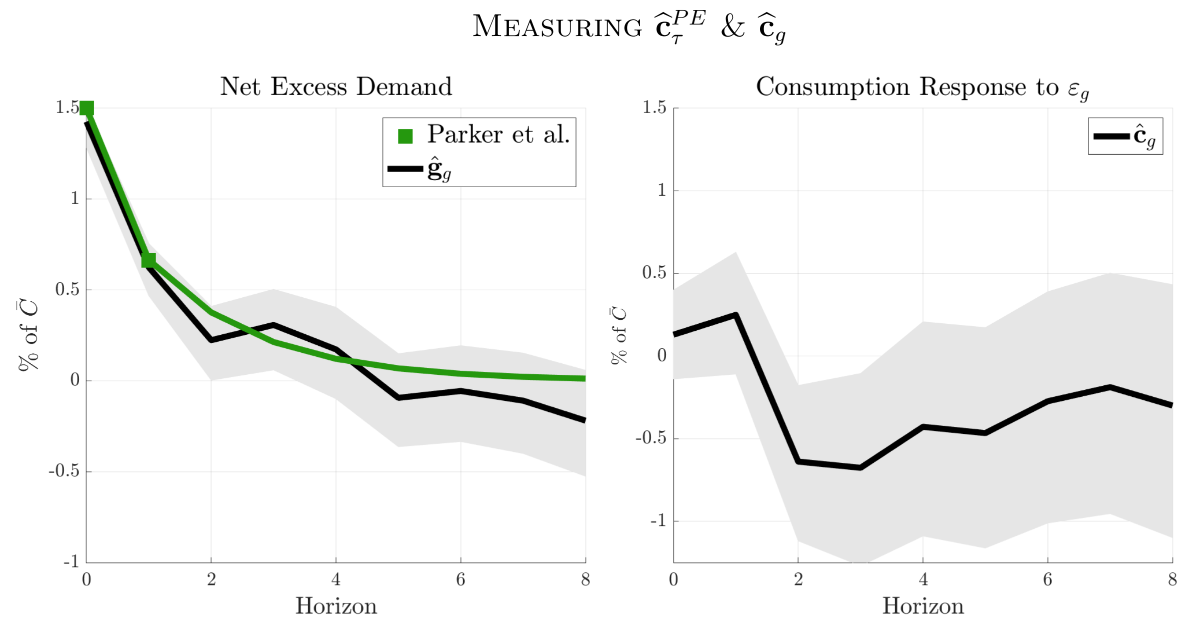

Figure F.1: The left panel shows direct consumption responses to the stimulus check (green) vs. direct government spending response to identified spending shock (black), with 16th and 84th percentile confidence bands (grey), quarterly frequency. Estimated consumption responses from Parker et al. (2013) 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.

Stimulus Checks, Aggregate Impulse Responses

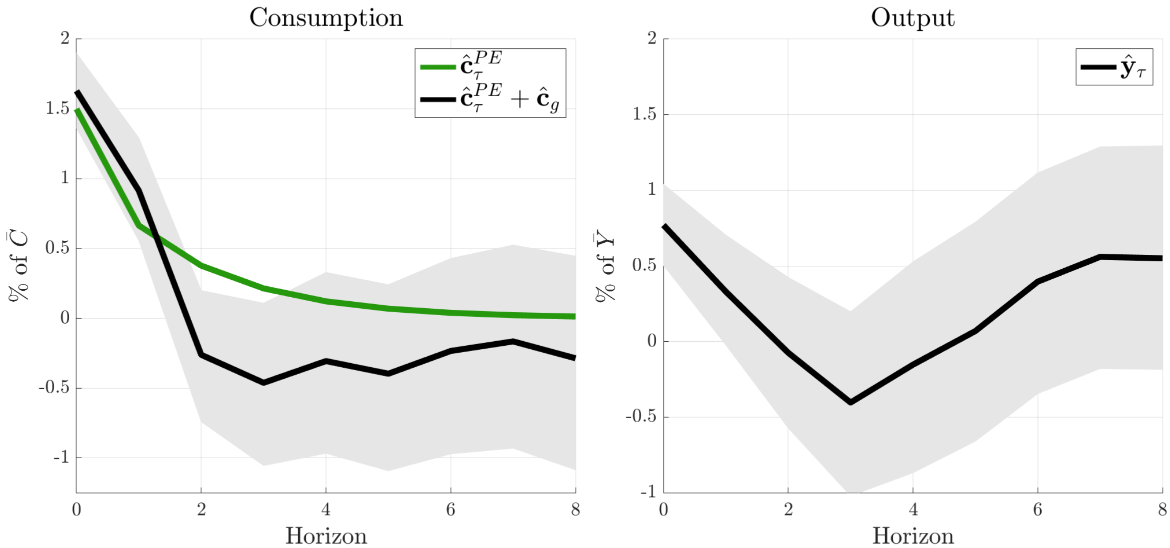

Figure F.2: Consumption and output responses to a stimulus check shock, quarterly frequency. The full consumption response is computed following the exact additive decomposition of Proposition 1, while the output response is simply equal to the response after a government spending shock. The grey areas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

Stimulus Checks Counterfactuals, Defense Forecast Errors

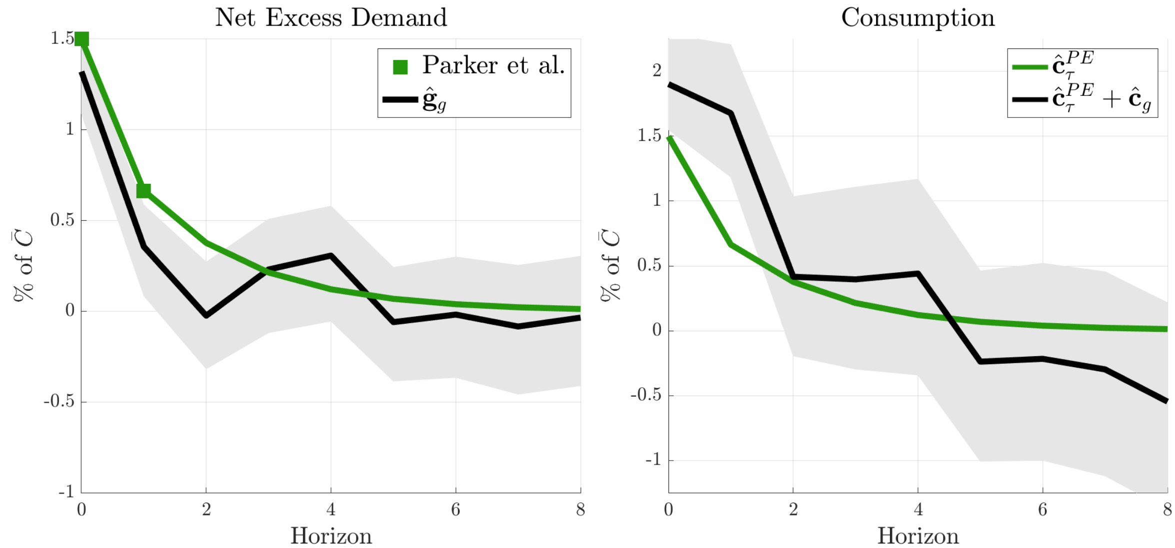

Figure F.3: See the captions for Figure F.1 and Figure F.2. I now use the defense forecast spending error series from Drautzburg (2020).

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

Measuring $\widehat{\mathbf{c}}_{\tau}^{PE}$ & $\widehat{\mathbf{c}}_{g}$

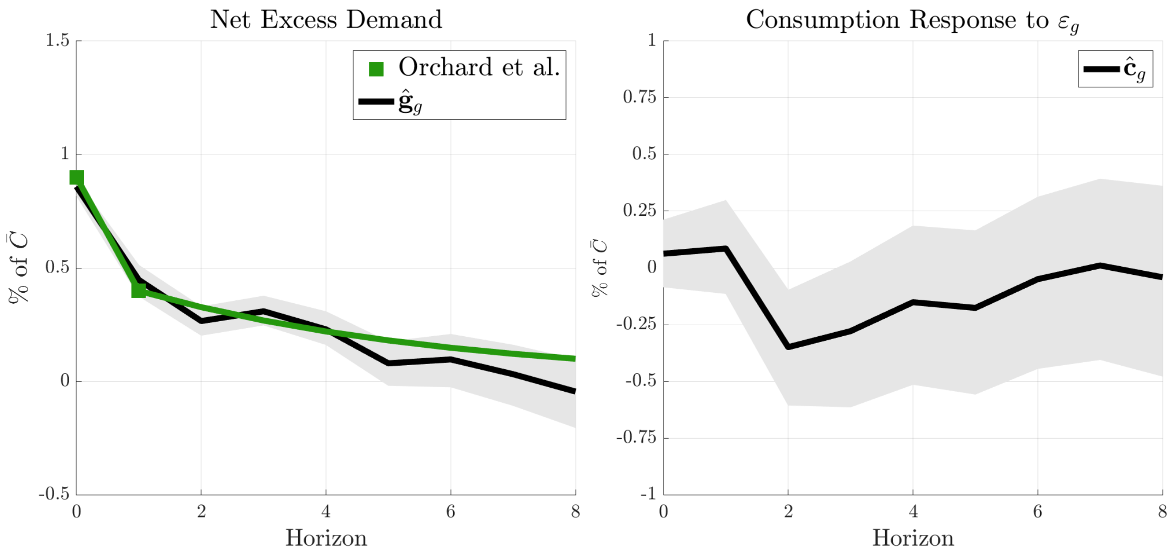

Figure F.4: The left panel shows direct consumption responses to the stimulus check (green) vs. direct government spending response to identified spending shock (black), with 16th and 84th percentile confidence bands (grey), quarterly frequency. Estimated consumption responses from Orchard et al. (2022), 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

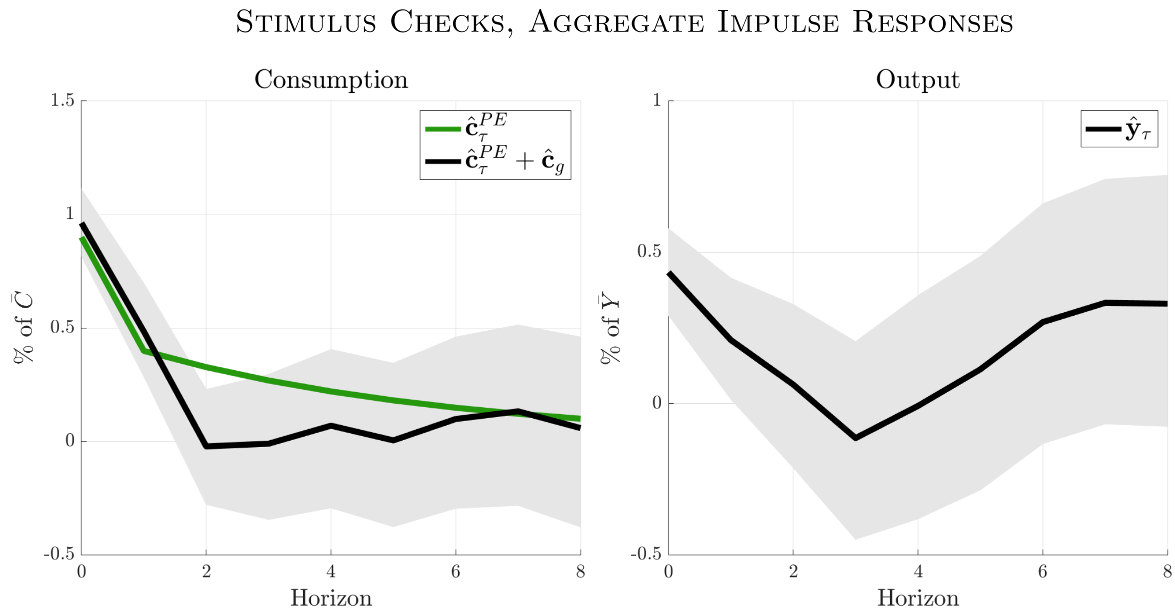

Figure F.5: Consumption and output responses to a stimulus check shock, quarterly frequency. The full consumption response is computed following the exact additive decomposition of Proposition 1, while the output response is simply equal to the response after the matched linear combination of government 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

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

[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:

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

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.

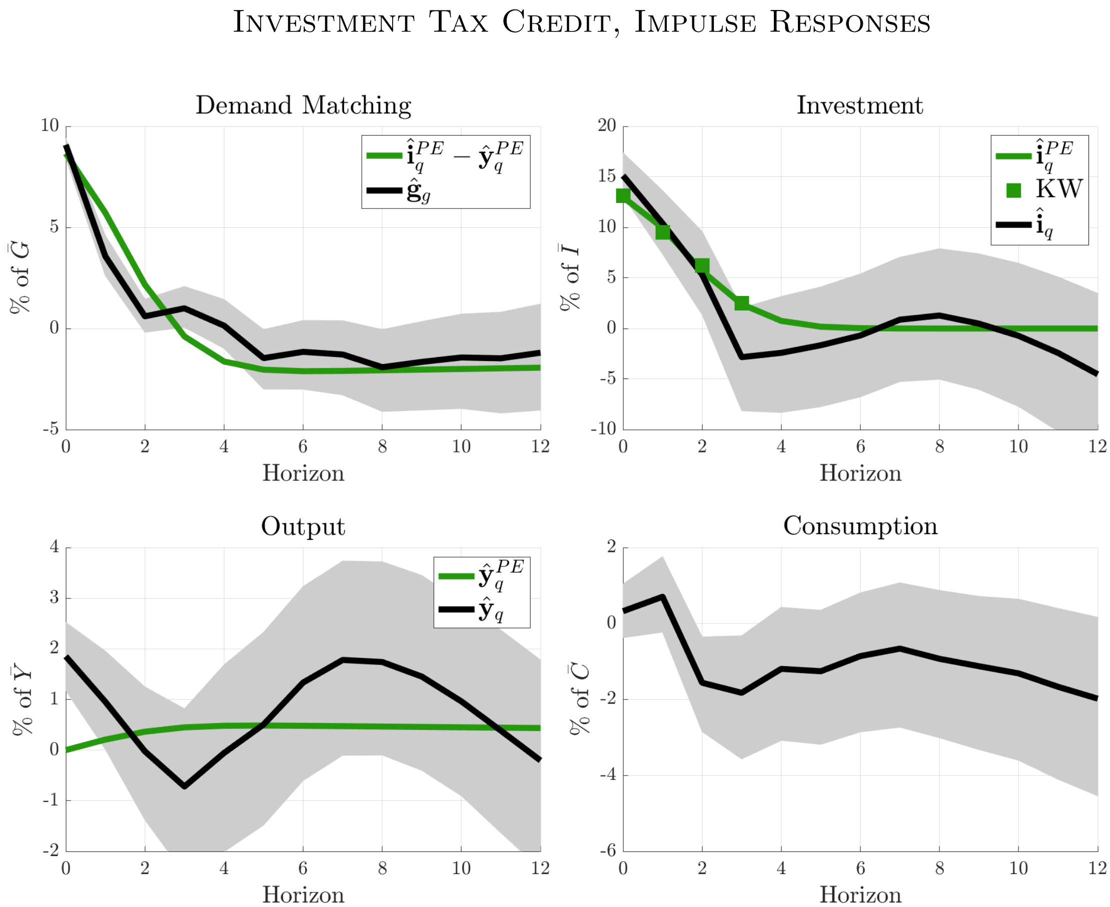

Figure F.6: Investment, output and consumption responses to an investment tax incentive shock, quarterly frequency, with the partial equilibrium net excess demand path matched to a linear combination of government spending shocks. "KW" refers to Koby & Wolf (2020); details are given in Appendix E.2. 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

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

[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.

Redistribution Shock, Impulse Responses

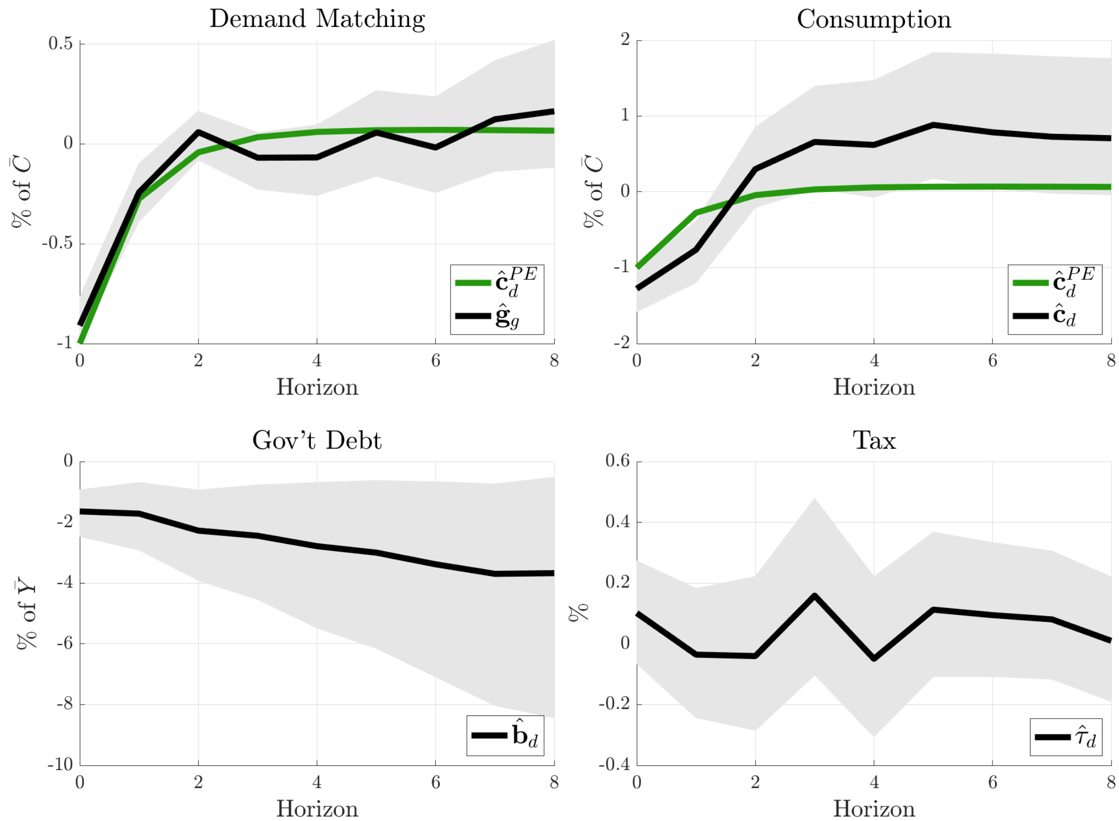

Figure F.7: Consumption, debt and tax responses to a redistribution shock, with the partial equilibrium net excess demand path matched to a linear combination of government spending shocks. The consumption response is computed in line with Proposition 1. The plot also shows the required demand matching 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. $\square$

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

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

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

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

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

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

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

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

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

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

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

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

$\square$

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. □