NBER WORKING PAPER SERIES

THE MISSING INTERCEPT: A DEMAND EQUIVALENCE APPROACH

Christian K. Wolf

Working Paper 29558
http://www.nber.org/papers/w29558

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2021

This paper is a subtantially revised version of Chapter 1 of my PhD dissertation. I am very grateful to my advisors Ben Moll, Mikkel Plagborg-Møller, Gianluca Violante and Mark Watson for their continued guidance and support. I have particularly benefited from detailed comments by Mark Aguiar, Adrien Auclert, Markus Brunnermeier, Dave Donaldson, Thorsten Drautzburg, Peter Ganong, Erik Hurst, Oleg Itskhoki, Greg Kaplan, Alisdair McKay, Emi Nakamura, Jonathan Payne, Monika Piazzesi, Richard Rogerson, Matt Rognlie, Martin Schneider, Jón Steinsson, Ludwig Straub, Chris Tonetti, Iván Werning, Tom Winberry, as well as four anonymous referees. I also thank seminar participants at several venues for their insights, and acknowledge generous support from the Alfred P. Sloan Foundation and the Macro Financial Modeling Project, as well as the Washington Center for Equitable Growth. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The Missing Intercept: A Demand Equivalence Approach
Christian K. Wolf
NBER Working Paper No. 29558
December 2021
JEL No. E20,E32,E62

ABSTRACT

I give conditions under which changes in private spending are accommodated in general equilibrium exactly like changes in aggregate fiscal expenditure. Under such demand equivalence, researchers can use time series evidence on fiscal multipliers to recover the general equilibrium "missing intercept" of shocks to private spending identified in the cross section. I apply this method to deficit-financed stimulus checks, and find a) a large direct consumption spending response, and b) a fiscal multiplier of one and so a missing intercept close to zero. I also discuss the robustness of this aggregation approach to empirically plausible violations of demand equivalence.

Christian K. Wolf
MIT Department of Economics
The Morris and Sophie Chang Building
50 Memorial Drive
Cambridge, MA 02139
and NBER
ckwolf@mit.edu

An online appendix is available at http://www.nber.org/data-appendix/w29558

1 Introduction

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

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

This paper offers a hybrid perspective: I use structural models to gauge the informativeness of fiscal spending shocks for the general equilibrium effects of private demand shifters, clarify the conditions under which those fiscal shocks solve the missing intercept problem, and finally assess the plausibility of these conditions. The analysis proceeds in three steps. First, in the context of a relatively general structural macro model, I give a set of restrictions on economic primitives ensuring "demand equivalence"—that is, identical changes in private and public demand eliciting identical general equilibrium feedback. Second, I leverage this equivalence result to propose a measurement strategy. Under demand equivalence, time series estimates of the aggregate effects of a given change in fiscal expenditure for free allow us to aggregate *particular* shifters of private spending: those that (i) induce the same time path of aggregate net excess demand; (ii) are financed just like the fiscal spending change; and (iii) occur in the same macroeconomic environment. To illustrate my measurement strategy, I combine a) cross-sectional evidence on household consumption following stimulus check receipt and b) time series evidence on fiscal spending shocks to estimate the *aggregate* effects

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

of stimulus check policies. My estimates of a fiscal multiplier of one suggest a "missing intercept" close to zero, and so a macro counterfactual close to cross-sectional spending responses. Third, I discuss the robustness of this approach to violations of demand equivalence. I find that most plausible violations lead to an upward bias: aggregation via fiscal multipliers understates the degree of general equilibrium crowding-out, and so overstates the macro causal effect. Using structural model simulations, I present general conditions under which the bias is likely to be small, and find them to be satisfied in my stimulus check application.

In the first part of the paper I identify conditions under which shocks to private spending—e.g., stimulus checks or credit tightenings—propagate in general equilibrium just like changes in public expenditure. The key building block result is that, in standard business-cycle models, linearized impulse responses to aggregate shocks can be characterized implicitly as solutions to a linear infinite-horizon system of market-clearing conditions. Private and public spending shocks thus induce identical general equilibrium effects as long as they perturb the same market-clearing conditions by the same amount. In a rich heterogeneous-agent model, this abstract exclusion restriction maps into three substantive economic assumptions. First, households and government consume the same final good. If so, identical changes in private or public spending lead to identical excess demand for that common good. Second, households and government borrow and lend at the same interest rate. The identical expansions in private and public demand then induce the same fiscal deficit (in net present value terms), and so can in principle be financed using identical paths of future taxes. In particular, if the net excess demand path has zero net present value (as is the case for non-policy shifters like credit tightenings), then the identical change in public spending can be *purely* deficit-financed, with no direct tax response. Third, household labor supply does not respond differentially to the two shocks; sufficient conditions for this are either the absence of wealth effects in labor supply or fully demand-determined employment. Under these three restrictions, for any given shock to private spending, I prove that *there exists* a public spending shock that solves the missing intercept problem—that is, the response of aggregate consumption to the private demand shifter is equal to the sum of a) the shifter's *direct* effect on consumer spending and b) the *total* response of consumption to the public spending shock. The constructive proof reveals the properties of this public spending shock: it must (i) induce the same path of excess demand and (ii) be associated with the same path of future taxes as the private demand shifter. My focus on linearized equilibria furthermore automatically imposes a third condition (iii): both shocks occur in the same macro environment, including in particular the same response of the monetary authority.

Next, I show how to combine this theoretical demand equivalence insight with empirical time series and cross-sectional evidence on public and private spending shock transmission. Cross-sectional regressions of private spending on a demand shifter (e.g., check receipt) recover that shifter's *direct* effect on private demand. Semi-structural time series methods on the other hand have been applied widely to estimate the aggregate effects of fiscal purchases. Demand equivalence tells us how to put the two together: time series fiscal multiplier estimates for free give us the general equilibrium effects of *any* shifter of private demand that satisfies my conditions (i) - (iii)—same net excess demand, same financing, and same macro environment. Can we in practice find cross-sectionally and time series identified shocks that align in this very particular way? I first of all present two strategies to ensure (i), the alignment of net excess demand paths.

1. Given any cross-sectionally identified demand path, researchers can search the set of fiscal spending experiments studied in previous work to find the one—or a linear combination of several—that induces as similar an aggregate excess demand path as possible.

2. Given any public spending path estimate from time series experiments, researchers can try to figure out what kind of shock to private demand would induce that same spending path. I show that one can do so by either running a large number of of cross-sectional regressions or (more plausibly) by specifying a *partial equilibrium* model of private spending behavior.

With demand paths matched using either of these strategies, the next step is to compute the sum of the a) cross-sectional and b) time series estimated consumption impulse responses. Conditions (ii) and (iii) then tell us how to interpret this sum: under demand equivalence, the sum is a valid general equilibrium counterfactual for a private demand shifter that (ii) was financed just like the identified fiscal shock and (iii) occurred in the same macro environment.

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

I first showcase the method's applicability through a study of "stimulus check" policies. While increasingly popular as a policy stimulus tool, relatively little is known about their aggregate effects, simply because of a lack of plausibly exogenous time series variation. Cross-sectional micro evidence suggests that one-off checks lead to a large but short-lived expansion in private spending (Parker et al., 2013). I connect this finding to the time series evidence on a similarly short-lived expansion in public spending, following Ramey (2011). The expansion is deficit-financed, largely accommodated by the monetary authority, and approximately moves output one-for-one, with only small effects on private spending. By demand equivalence, I put these two pieces of evidence together and conclude that one-off, deficit-financed checks briefly but significantly stimulate aggregate consumption, with the overall response close to the direct effect estimated using micro data alone. To illustrate my second strategy for the alignment of net excess demand paths, I then consider the fiscal proxy-SVAR of Caldara & Kamps (2017). Their methodology identifies a persistent, deficit-financed increase in fiscal purchases, again largely accommodated by the monetary authority. I combine cross-sectional estimates of household spending behavior with some basic consumer theory (following Wolf, 2021) to conclude that the identified spending path can be equivalently induced via a string of stimulus checks sent to households. I then aggregate as before, and again find a missing intercept path close to zero throughout.

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

I conclude with a brief discussion of the scope of the demand equivalence approach. First, while my main application is to uniform stimulus checks, the method applies without change to many other shifters of consumer demand, including targeted transfers (Andreolli & Surico, 2021), household deleveraging (Guerrieri & Lorenzoni, 2017) or increases in inequality (Auclert & Rognlie, 2018). To illustrate, I show that a fiscal contraction today offset by an expansion in the future identifies the missing intercept of a temporary increase in earnings inequality. Second, I extend the equivalence theory to shifters of investment demand and provide an application to bonus depreciation stimulus.

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

First, the demand equivalence approach to the missing intercept problem connects two empirical literatures. A fast-growing line of work uses variation at the individual or regional level to estimate spending responses to policy changes and other shocks (e.g. Mian & Sufi, 2009; Parker et al., 2013; Zwick & Mahon, 2017). As all of these studies control for macro fluctuations through time fixed effects, they are silent on any possible general equilibrium feedback. I give formal conditions under which a second literature—that on the aggregate effects of changes in government spending—can be informative about this "missing intercept." Comprehensive literature summaries are Hall (2009) and Ramey (2018).

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

Third, the proposed methodology naturally complements existing strategies for the estimation of (policy) counterfactuals in macroeconomics. In its reliance on general exclusion restrictions rather than parametric structural models, it is semi-structural in the same way as conventional Structural Vector Autoregressive (SVAR) analysis (Sims, 1980). The general idea also connects with the "sufficient statistics" approach that is common in public finance (Chetty, 2009) and increasingly widespread in macroeconomics (Auclert et al., 2018; Nakamura & Steinsson, 2018). I show that, across a particular family of structural models, certain moments—the cross-sectional and time series estimates required by my methodology—are fully informative about the desired macroeconomic counterfactual, thus obviating the need for model solution (Andrews et al., 2018).

OUTLINE. Section 2 establishes the consumption demand equivalence result. In Section 3, I show how to connect this theoretical result with cross-sectional and time series estimates of dynamic causal effects, and present an application to stimulus checks. Section 4 critically assesses the output of the proposed procedure. Applications to other shifters of consumer demand as well as the extension to investment are discussed in Section 5. Section 6 concludes, and supplementary details, proofs and further applications are all relegated to the appendix.

2 Consumption demand equivalence

This section builds on the familiar Keynesian cross logic to develop an equivalence result for the general equilibrium propagation of shocks to private consumption and to public spending. In Sections 2.1 and 2.2, I discuss the restrictions on economic primitives needed for demand equivalence in a standard structural business-cycle model. Section 2.3 complements this particular model-based analysis with a general abstract formulation of shock equivalence as a set of exclusion restrictions on (linearized) aggregate equilibrium conditions.

2.1 Model

I begin by presenting a particular model environment, general enough to nest many seminal contributions to quantitative business-cycle analysis.[2] The purpose of the model is twofold: First, it allows me to present economically interpretable sufficient conditions for demand equivalence in a familiar, canonical environment. Second, the model will form the backbone of my critical assessment of the demand equivalence methodology in Section 4.

Time is discrete and runs forever, $t = 0, 1, \ldots$. The economy is populated by households, firms, and a government. There is no aggregate uncertainty, but households and firms are allowed to face idiosyncratic risk. I study perfect foresight transition paths back to steady state after one-time unexpected aggregate innovations at time 0; for vanishingly small innovations, these transition paths are equivalent to standard impulse response functions

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

computed from the first-order perturbation solution of an otherwise identical model with aggregate risk.[3] Anticipating my main application, I will focus on two such innovations: first, stimulus checks sent to households, and second, a transitory expansion in government spending. Section 2.3 shows how the equivalence result extends to generic policy and non-policy shifters of consumption demand (e.g., changes in borrowing constraints).

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

Households. A unit continuum of households $i \in [0,1]$ has preferences over consumption c_{it} and labor ℓ_{it}. The real relative price of the consumption bundle in terms of the economy's numeraire is denoted p_t^c. Households are subject to idiosyncratic productivity risk e_{it}, and can self-insure by investing in liquid nominal bonds b_{it}^h, with nominal returns i_t^b and subject to a borrowing constraint $\underline{b}$. Borrowing incurs an additional penalty $\kappa_b \geq 0$. Income consists of labor earnings as well as (potentially type-specific) lump-sum transfers τ_{it} and dividend income d_{it}. Total hours worked ℓ_{it} are determined by demands of a unit continuum $k \in [0,1]$ of price-setting labor unions, as in Erceg et al. (2000); the problem of labor unions will be considered later. Given a path of prices, transfers, dividends, hours worked and inflation (π_t), the consumption-savings problem of household i is thus

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

such that

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

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

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

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

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

sold at the aggregate wage index w_t, and where ϵ_w denotes the elasticity of substitution between different types of labor. Union k chooses its wage rate w_{kt} subject to wage-setting adjustment costs, and satisfies the corresponding demand for its labor services. I assume that it does so by demanding a common amount of hours worked from its members.[4] Since the wage-setting problem is standard, I relegate details to Appendix B.1. For the purposes of the analysis here, it suffices to note that union behavior can be summarized through a wage New Keynesian Phillips curve—effectively, an aggregate labor supply relation.

Fiscal policy. The fiscal authority consumes a bundle with real relative price p_t^g. Fiscal consumption g_t and total lump-sum transfers $\tau_t \equiv \int_0^1 \tau_{it} di$ are financed through debt issuance and taxes on labor income. The government flow budget constraint is

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

I assume that total government spending $\mathbf{g} = \mathbf{g}(\boldsymbol{\varepsilon}_g)$ and the discretionary part of stimulus checks $\boldsymbol{\tau}^x = \boldsymbol{\tau}^x(\boldsymbol{\varepsilon}_\tau)$ follow exogenous processes. A government debt financing rule is mapping from spending targets $(\boldsymbol{g}, \boldsymbol{\tau}^x)$, initial nominal debt b_{-1} and a path of prices and quantities $(\mathbf{w}, \boldsymbol{\ell}, \mathbf{i}^b, \boldsymbol{\pi}, \boldsymbol{p}^g)$ into the endogenous part of transfers $\boldsymbol{\tau}^e$ such that $\boldsymbol{\tau} = \boldsymbol{\tau}^e + \boldsymbol{\tau}^x$, the flow government budget constraint holds at all periods t, and $\lim_{t\to\infty} \widehat{b}_t = 0$. That is, lump-sum taxes adjust in response to fiscal outlays—both outright expenditure and stimulus checks—to ultimately return government debt to its steady-state level. I emphasize that all results below extend without change to the alternative assumption of outlays financed with time-varying

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

distortionary taxes τ_ℓ; the key restriction for demand equivalence will be that stimulus checks and spending increases are financed using *identical* paths of taxes, distortionary or not.

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

The last remaining entity in the model is the monetary authority. This monetary authority sets nominal rates on liquid bonds i_t^b following some (Taylor-type) rule.

Equilibrium. I assume that there exists a unique deterministic steady state.[5] To allow interpretation of perfect foresight transition paths as conventional first-order perturbation solutions, I impose that the economy is indeed initially in steady state, and then study perfect foresight transition equilibria back to the initial deterministic steady state. The definition of equilibrium perfect foresight transition paths is then standard (see Appendix B.1); I discuss an extension to transition paths with other starting points in Appendix C.1.

2.2 The equivalence result

I now formalize the Keynesian cross intuition of a commonality in general equilibrium propagation. A precise statement of such equivalence first of all requires a definition of direct (or partial equilibrium) responses and indirect (or general equilibrium) effects.

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

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

total impulse response of consumption to the shock path $\boldsymbol{\varepsilon}$ is simply

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

I decompose this aggregate impulse response into two parts: a direct "partial equilibrium" impulse and an indirect "general equilibrium" feedback part.[6]

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

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

Similarly, let the indirect (general equilibrium) feedback be

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

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

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

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

The remainder of this section establishes properties of the decomposition (5)—the desired formalization of the simple Keynesian cross intuition. Section 3 will then connect theory and measurement, linking the components of (5) to measurable objects and so in particular to the "missing intercept" (or aggregation) problem of cross-sectional causal effect estimates.

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

The first assumption restricts goods bundles in the economy.

[6] My definition of the partial equilibrium consumption response abstracts from endogenous partial equilibrium adjustments in earnings. I do so for three reasons. First, many empirical estimates of household spending responses to sudden income changes are actually interpretable as such netted spending elasticities (e.g. see Auclert, 2019, footnote 34). Second, in models with union-intermediated labor supply, replicating cross-sectional regressions differences out labor responses (see Proposition 2). Third, microeconomic evidence suggests that short-run wealth effects are weak (Cesarini et al., 2017). Nevertheless, in Appendix E.3, I repeat my analysis in an alternative model without unions, but with a non-standard preference parameterization allowing for weak short-run wealth effects (as in Jaimovich & Rebelo, 2009; Galí et al., 2012).

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

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

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

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

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

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

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

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

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

Demand Equivalence Illustration, Sticky-Wage HANK Model

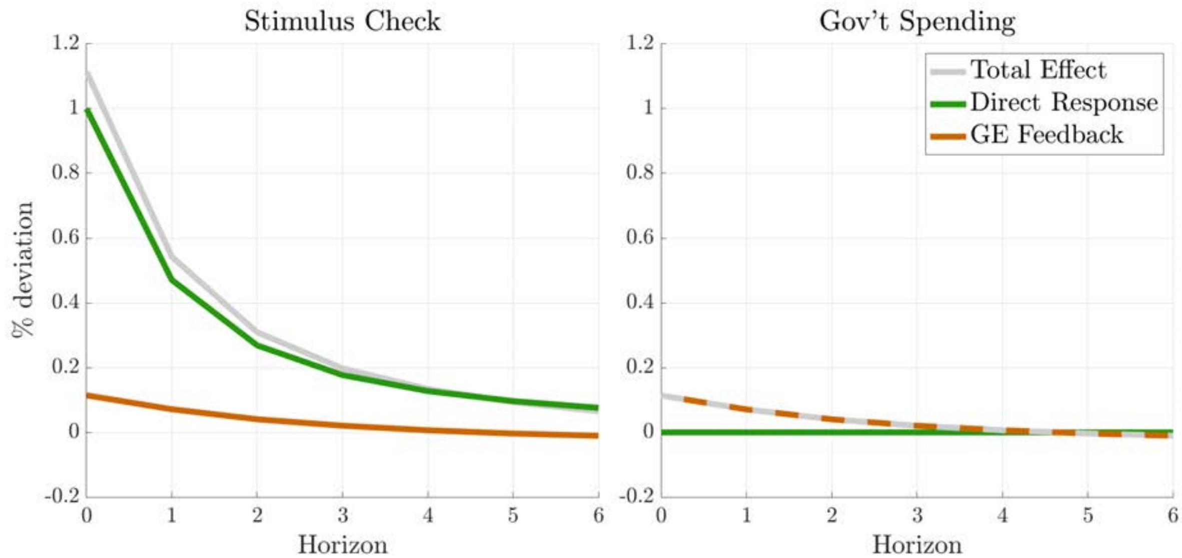

Figure 1: Consumption impulse response decompositions after stimulus check and government spending shocks in the estimated HANK model of Section 4.1, but with fully rigid wages. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Finally, I emphasize that my focus on linearized equilibria also implicitly imposes a further condition (iii) on the two shocks: they occur in the same macroeconomic environment, including for example the same monetary policy response. This assumption will be important in communicating the results of my measurement strategy in Section 3.

Proof sketch. My proof of the demand equivalence decomposition in (6) leverages the "sequence-space" approach to equilibrium characterization developed in Auclert & Rognlie (2018) and Auclert et al. (2018).

The basic idea of the argument is the following. Equilibria even in the rich model of Section 2.1 can be characterized as a system of several aggregate prices and quantities adjusting to clear several markets. Assumptions 1 to 3 simply turn out to be sufficient to ensure that, for any given private spending shock, a public expenditure shock with properties (i) and (ii) will perturb the same market-clearing conditions by the same amount, thus eliciting the same general equilibrium adjustment and so implying (6). First, Assumption 1—in conjunction with requirement (i) on the fiscal shock, $\widehat{\mathbf{g}}_g = \widehat{\mathbf{c}}_\tau^{PE}$—ensures that the private and public demand shocks lead to the same excess demand pressure for the common final good. Second, since by Assumption 2 households and governments borrow and lend at identical rates, this common excess demand path can *in principle* be financed using identical paths of future taxes. Property (ii) of the fiscal spending shock ensures that this is indeed the case. Third, Assumption 3—a restriction on household behavior—ensures that the consumption increase induced by stimulus checks does not lead to any *direct* adjustment in hours worked.[7]

Overall, my statement of demand equivalence in Proposition 1 offers two key insights relative to previous work. First, it explicitly characterizes the properties of the fiscal shock required for the decomposition in (6) to hold. These properties will take center stage in the connection of theory to measurement in Section 3. Second, it reveals that not all assumptions necessary to arrive at a Keynesian cross-type equilibrium characterization (as e.g. in Auclert et al., 2018) are also necessary for a demand equivalence result. Notably, neither the absence of investment nor the assumption of a fixed real rate of interest are required.

2.3 Extension to general exclusion restrictions

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

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

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

it follows immediately that, to first order,

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

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

exactly as in Proposition 1.[8] Condition (7) is a general exclusion restriction on the equilibrium system: both shocks must enter all equilibrium equations symmetrically. The proof of Proposition 1 works because, under my imposed restrictions, the equilibrium can be cast in a form consistent with (7).[9] In Appendices C.2 and C.3 I give examples of other shocks and models that can be written in this form. First, I extend the model of Section 2.1 to allow for time preference shocks as a generic non-policy consumption shifter. Equivalence obtains under the same restrictions as those discussed in Section 2.2; in particular, since non-policy private demand shifters $\boldsymbol{\varepsilon}_d$ necessarily induce spending paths $\widehat{\boldsymbol{c}}_d^{PE}$ with zero net present value, requirement (ii) in the equivalence proof now dictates that the equivalent public spending shock $\boldsymbol{\varepsilon}_g$ is *purely deficit-financed*, with $\boldsymbol{\tau}^e$ moving only because of general equilibrium effects, and not directly because of the spending shock $\boldsymbol{\varepsilon}_g$. Second, I consider several examples of popular models beyond the familiar New Keynesian tradition—including for example models with non-rational expectation formation of firms and households—, and show that they still all fit into the general semi-structure of (7).

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

The appeal of the demand equivalence result is instead that it tells us how to *measure* those common general equilibrium effects, however strong (or weak) they may turn out to be. I now turn to this empirical measurement strategy.

3 Solving the missing intercept problem

This section shows how the demand equivalence result together with time series evidence on the dynamic causal effects of fiscal purchases can be used to solve the "missing intercept" aggregation problem for cross-sectionally identified consumption demand shifters. Section 3.1

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

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

begins by tying the theoretical decomposition in (8) to empirically measurable objects. Section 3.2 uses this mapping between theory and empirics to propose a measurement strategy, and Section 3.3 applies the method to estimate the aggregate effects of stimulus checks.

3.1 From theory to measurement

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

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

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

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

where α_i and δ_t are individual and time fixed effects. It is straightforward to show that, under my assumptions, regressions such as (9) estimate average household-level causal effects that are interpretable as direct partial equilibrium shock responses, as claimed.[10,11]

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

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

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

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

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

Fiscal multipliers. The dynamic causal effects of changes in fiscal purchases on macroeconomic outcomes can be estimated using the conventional semi-structural macroeconometric toolkit (see Ramey, 2016, for a review). Under demand equivalence, such time series analysis promises to identify part b): the "missing intercept" of general equilibrium effects.

Most previous work has relied on one of three possible sets of identifying assumptions. First, researchers may have access to direct, "narrative" measures of aggregate fiscal policy shocks (Ramey, 2011). Second, outside information—either in the form of direct zero or sign restrictions (Blanchard & Perotti, 2002; Mountford & Uhlig, 2009) or via measures of other shocks (Caldara & Kamps, 2017)—can help researchers pin down government spending rules and so identify spending shocks. Third, professional forecast errors of government spending may be interpreted as orthogonal to any (known) rules-based spending components, and so again serve as an instrumental variable for fiscal shocks (Ramey, 2011; Drautzburg, 2020). In all of these cases, the desired counterfactuals can then be estimated using Vector Autoregressions (VARs) or Local Projections (LPs). Importantly, across this list of identifying assumptions, estimated fiscal multipliers tend to lie at around 1 (Ramey, 2018).

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

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

To make this potentially powerful insight operational, however, we must confront an important challenge: the demand equivalence result only allows us to combine *particular* cross-sectional and time series identified shocks—those that satisfy the conditions (i) - (iii). This is challenging; for example, nothing guarantees that any given time series and cross-sectional causal effect estimates give demand paths that are aligned as required by (i). The next subsection presents two solutions to this problem of aligning the shocks.

3.2 Matching time series & cross-sectional evidence

I begin with the first requirement: how can we find net excess demand paths $\widehat{\mathbf{g}}_g$ and $\widehat{\mathbf{c}}_d^{PE}$ that are at least approximately aligned? Linearity gives one degree of freedom, so we can slightly relax the matching requirement to the proportionality condition $\widehat{\mathbf{g}}_g \propto \widehat{\mathbf{c}}_d^{PE}$. I consider two possible approaches to ensuring this alignment of spending paths.

1. *From cross-section to time series.* The first approach begins with some given cross-sectional spending estimate $\widehat{\mathbf{c}}_d^{PE}$, and then tries to find the fiscal spending experiment that solves the "missing intercept" problem for that particular shock. To do so, suppose that a researcher has used one or several of the semi-structural time series identification approaches discussed in Section 3.1 to estimate the aggregate effects of a menu of n_k different kinds of government spending shocks $\{\varepsilon_{g_k}\}_{k=1}^{n_k}$ with implied spending paths $\widehat{\mathbf{g}}_{g_k}$. A linear projection of $\widehat{\mathbf{c}}_d^{PE}$ on the space spanned by those shock paths gives

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

   If the error is sufficiently small, we may consider the weighted average

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

   as a promising candidate to learn about the "missing intercept" $\widehat{\mathbf{c}}_d^{GE}$. Of course, by construction, using this weighted average of fiscal shocks will only solve the missing intercept problem up to the aggregate effects of a shock that equals the error term.[12] Section 3.3 will provide a concrete example of a case in which this matching error is indeed very small, implying that the proposed approach to aggregation is promising.

[12]See Appendix E.8 for a further discussion of this point. I there also use a structural model to study the inaccuracy associated with demand matching errors of similar magnitude to those that I observe in my empirical stimulus check application in Section 3.3.

2. *From time series to cross-section.* The second approach proceeds the other way around, beginning with a path $\widehat{\mathbf{g}}_g$ identified in the time series, and then asking what kind of private demand shock—e.g., what sequence of stimulus checks—would induce the same spending response. Heuristically, given a private demand shock ε_d and the consumption function $\mathbf{c}(\bullet)$ defined in Section 2.2, we can in principle recover the required shock path as

$$\boldsymbol{\varepsilon}_d \quad \equiv \quad \mathcal{C}_d^{-1} \quad \times \quad \widehat{\mathbf{g}}_g \tag{13}$$

where $\mathcal{C}_d \equiv \frac{\partial \mathbf{c}(\bar{\mathbf{s}}^h;\boldsymbol{\varepsilon})}{\partial \boldsymbol{\varepsilon}_d}$, assumed to be invertible. It follows from Proposition 2 that cross-sectional analysis is informative about (weighted averages of) columns of $\mathcal{C}_d$. Running a large number of cross-sectional regressions for different paths of private demand shocks $\boldsymbol{\varepsilon}_d$ would thus in principle allow researchers to construct this inverse mapping. More realistically, researchers may combine empirical regression estimates with theoretical restrictions on the shape of $\mathcal{C}_d$. Importantly, such restrictions would only require the researcher to take a stance on a *partial equilibrium* model of household consumption decisions—all general equilibrium effects would still be captured semi-structurally by the time series estimates of $\widehat{\mathbf{c}}_g$. My second application in Section 3.3 will provide an example of this approach.

Having addressed the demand path matching problem through either of these two strategies, I propose to proceed and construct the demand equivalence approximation as

$$\widehat{\mathbf{c}}_d \quad = \quad \underbrace{\widehat{\mathbf{c}}_d^{PE}}_{\text{PE response}} \quad + \quad \underbrace{\sum_{k=1}^{n_k} \gamma_k \quad \times \quad \widehat{\mathbf{c}}_{g_k}}_{\text{GE feedback}} \tag{14}$$

with the weights γ_k as in (11) under the first approach, or we set $\widehat{\mathbf{c}}_d = \widehat{\mathbf{c}}_d^{PE} + \widehat{\mathbf{c}}_g$ with $\widehat{\mathbf{c}}_d^{PE} = \widehat{\mathbf{g}}_g$ for the second approach, giving a counterfactual for the shock $\boldsymbol{\varepsilon}_d$ derived in (13).

Conditions (ii) and (iii) then affect the *interpretation* of this construction. Under demand equivalence, my sum is interpretable as a valid general equilibrium counterfactual for a shock $\boldsymbol{\varepsilon}_d$ that is associated with the same aggregate movements in taxes as the identified public demand shock(s) *and* occurred in the same macroeconomic environment. For example, if the private demand shifter is stimulus checks, and the matched time series $\boldsymbol{\varepsilon}_g$ is persistently deficit-financed and largely accommodated by the monetary authority, then the estimated counterfactual will apply to similarly deficit-financed, accommodated stimulus checks.[13]

[13]Note that, for a non-policy demand shifter, $\widehat{\mathbf{c}}_d^{PE}$ necessarily has zero net present value, so the relevant comparison is a change in fiscal purchases financed by a reversal tomorrow, with taxes changing *only* because of general equilibrium effects. I consider an example of such a shock in Appendix F.2.

3.3 Two applications to stimulus checks

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

In the remainder of this section I construct macro counterfactuals for two kinds of stimulus check experiments: one-off checks and a more persistent sequence of checks. The micro and macro parts of my analysis leverage canonical contributions to the relevant cross-sectional and time series literatures: Parker et al. (2013) and Broda & Parker (2014) for the response of consumption to lump-sum income receipts, and Ramey (2011) and Caldara & Kamps (2017) for the aggregate general equilibrium effects of fiscal purchases.

One-off stimulus checks. For a one-off, one-quarter stimulus check sent to everyone, the policy's direct effect on partial equilibrium net excess spending is given as

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

where

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

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

Several recent studies have used rich household spending data to estimate objects that are either exactly or approximately interpretable as the desired average MPC (e.g. Johnson et al., 2006; Parker et al., 2013; Jappelli & Pistaferri, 2014; Fagereng et al., 2018). A common finding in this literature is that households spend most of a (small) one-time income receipt on impact, and that the spending response then decays back to zero quickly. In particular, the point estimates of Parker et al. (2013) and Broda & Parker (2014) suggest that, following a one-off stimulus check, total consumption expenditures increase by around 50 cents on the dollar on impact and 20 cents in the subsequent quarter. Translated to the size of the 2008 stimulus check policy, this corresponds to around 1.5 per cent of total personal consumption expenditure on impact, and 0.6 per cent in the following quarter. In the left panel of Figure 2,

MEASURING $\widehat{\mathbf{c}}_{\tau}^{PE}$ & $\widehat{\mathbf{c}}_g$

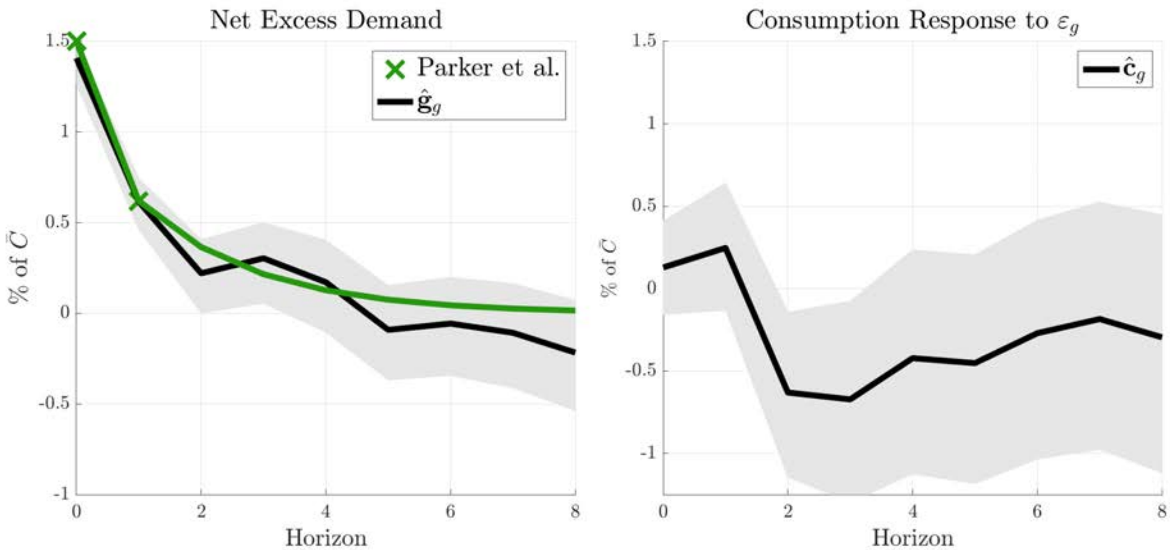

Figure 2: The left panel shows direct consumption responses to the stimulus check (green) vs. direct government spending response to identified spending shock (black), with 16th and 84th percentile confidence bands (grey), quarterly frequency. Estimated consumption responses from Parker et al. (2013) and Broda & Parker (2014), extrapolated for horizons beyond $t = 1$. The right panel shows the response of consumption to the fiscal spending shock.

the two green x's show those two direct consumption responses $\widehat{c}_{\tau,0}^{PE}$ and $\widehat{c}_{\tau,1}^{PE}$; the solid green line extrapolates those MPCs simply by imposing a constant geometric rate of decay for MPCs from $t = 1$ onwards, together with the restriction that (discounted) MPCs sum to 1, consistent with the lifetime household budget constraint.[14]

Taking as given that net excess demand path, we now need to find a similarly transitory fiscal spending expansion. Most time series experiments feature rather persistent increases in fiscal purchases, with the notable exception of the professional forecast errors studied in Ramey (2011). Those forecast errors are appealing as a measure of fiscal expenditure shocks because (1) professional forecasts are implicitly controlling for the rules-based component of fiscal expenditure and (2) the large information set of forecasters limits concerns of shock non-invertibility (Leeper et al., 2013). Furthermore, and most importantly for the purposes of the analysis here, they are known to induce a very transitory uptick in total fiscal purchases.

[14] I discuss the mapping between empirical regressions of transfer stimulus receipt and my notion of a direct spending response in Appendix D.1. That appendix also elaborates on my extrapolation. As I emphasize there, a constant rate of decay of 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).

Stimulus Checks, Aggregate Impulse Responses

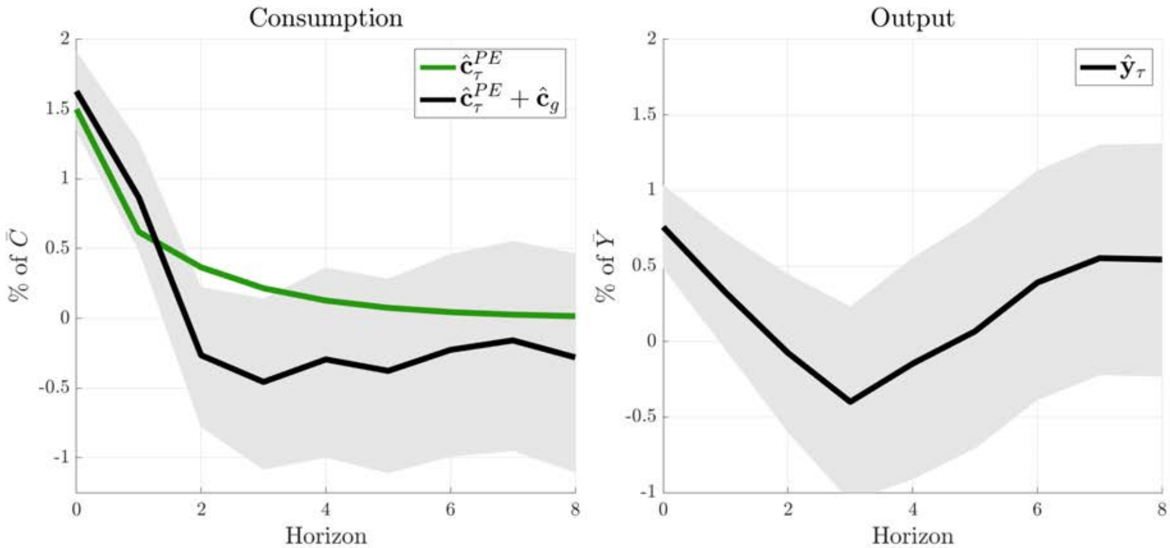

Figure 3: Consumption and output responses to a stimulus check shock, quarterly frequency. The full consumption response is computed following the exact additive decomposition of Proposition 1, while the output response is simply equal to the response after a government spending shock. The grey areas again correspond to 16th and 84th percentile confidence bands.

I study their propagation to the rest of the economy using a specification very similar to that in Ramey (2011), on a sample from 1981:Q3 (when forecast errors are first available) to 2008:Q4 (to ensure a stable macroeconomic regime). I then embed the shock in a recursive VAR, in line with the population results of Plagborg-Møller & Wolf (2021) and the finite-sample recommendations of Li et al. (2021). Since my choice of further macro controls, variable definitions, data construction and estimation procedure are all standard and in line with previous work, I relegate further details to Appendix D.3.

The left panel of Figure 2 reveals that, as required by the demand matching condition, 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 $\widehat{\mathbf{g}}_g$. Furthermore, the corresponding estimates for government debt and taxes reported in Appendix D.3 reveal the increase in fiscal purchases to be persistently deficit-financed. Finally, I there also show that the spending expansion was 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 2, 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.

Figure 3 puts all the pieces together to present full general equilibrium counterfactuals for stimulus checks. The left panel begins by implementing the demand equivalence decomposition in (6), simply 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 (with some crowding-out over time), the estimated *aggregate* effect of the policy turns out to be close to the micro-estimated direct effect. Thus, perhaps surprisingly, the various price and multiplier effects cited in previous empirical and theoretical work seem to roughly cancel. The right panel shows the corresponding 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) response, with output on impact rising by somewhat less than 1 per cent, and then returning to baseline. Overall, my estimates suggest that deficit-financed stimulus checks provide meaningful stimulus to aggregate consumption and output.

A STRING OF STIMULUS CHECKS. For my second illustration I proceed the other way around, beginning with the time-series fiscal multiplier estimates of Caldara & Kamps (2017). Those authors estimate a conventional fiscal policy VAR and show that conclusions on fiscal multipliers hinge critically on the systematic feedback from economic activity to fiscal purchases. The key contribution is to estimate this feedback not through outright restrictions on the fiscal rule (as in Blanchard & Perotti, 2002), but through proxy variables for *other* aggregate shocks as valid instrumental variables. I replicate their baseline specification, and show the response of fiscal purchases to the identified fiscal spending shock $\widehat{\mathbf{g}}_g$ in the left panel of Figure 4.[15] Fiscal purchases increase persistently for several quarters, and then gradually return to baseline. I normalize the impact spending response to equal 1 per cent of steady-state (quarterly) consumption, or around \$150 per household.

It remains to do the inversion step: what sequence of stimulus check payments to households would induce a net excess spending path $\widehat{\mathbf{c}}_\tau^{PE}$ akin to that in the left panel of Figure 4? By (13), we need to learn about $\mathcal{C}_\tau^{-1}$: the *inverse* of the mapping from sequences of transfers into sequences of consumer spending. I construct an estimate of this mapping by combining the empirical evidence in Parker et al. (2013) and Broda & Parker (2014) on spending responses to checks with consumer theory reviewed in Auclert et al. (2018) and Wolf (2021).

Standard models of the household consumption-savings decision with binding liquidity constraints imply that lump-sum stimulus checks today are spent gradually. I show in Wolf

[15] Details of the replicated VAR specification are provided in Appendix D.3.

Stimulus Check News Shock

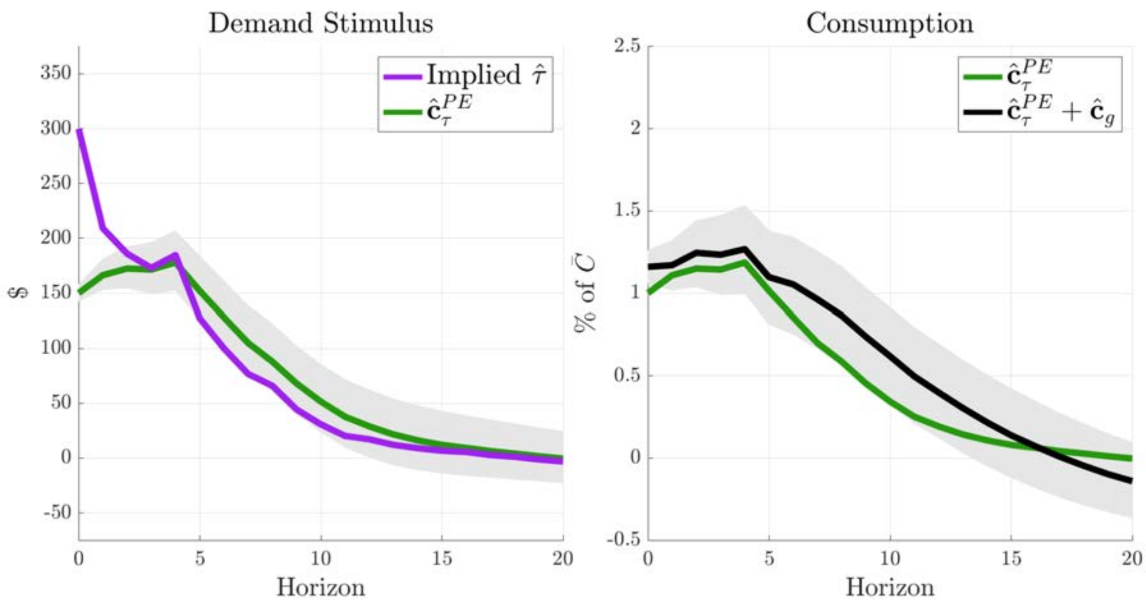

Figure 4: Direct spending response and general equilibrium consumption counterfactuals for a sequence of stimulus checks. The transfer shock path (purple line) is recovered as $\mathcal{C}_\tau^{-1} \times \widehat{\mathbf{g}}_g$, with $\mathcal{C}_\tau$ as in (15), while the full consumption response is computed following the exact additive decomposition of Proposition 1. The grey areas again correspond to 16th and 84th percentile confidence bands.

(2021) that this gradual spending response is very well-described by a simple three-parameter profile: $(\omega, \omega\xi\theta, \omega\xi\theta^2, \omega\xi\theta^3, \dots)$. Matching this spending profile to the evidence reported by Broda & Parker gives $\{\omega = 0.5, \xi = 0.7, \theta = 0.59\}$—the first column of $\mathcal{C}_\tau$.[16] Furthermore, since Broda & Parker fail to find significant pre-treatment effects, and also consistent with results in other settings (e.g., Kueng, 2018; Ganong & Noel, 2019; Baugh et al., 2021), I for my baseline analysis assume the absence of any anticipation effects related to *future* transfer receipt. The direct household spending response map $\mathcal{C}_\tau$ is then simply given as

$$\mathcal{C}_\tau \quad = \quad \omega \quad \times \quad \begin{pmatrix} 1 & 0 & 0 & \dots \\ \xi\theta & 1 & 0 & \dots \\ \xi\theta^2 & \xi\theta & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{15}$$

[16] As discussed above, Broda & Parker estimate the first two entries of the first column to be around 0.5 and 0.2. I then discipline the rate of decay by requiring a lifetime MPC of 1 (i.e., the column sums to 1, assuming zero rate of interest), exactly as I did for my MPC extrapolation in the first application.

While my discussion in the remainder of this section will rely on (15), I in Appendix D.1 present results for different degrees of household foresight, following Gabaix (2020), Auclert et al. (2019) and Wolf (2021).[17]

The purple line in Figure 4 shows the implied equivalent transfer path. Inverting $\mathcal{C}_\tau$, we find it as

$$\widehat{\tau}_t \;=\; \frac{1}{\omega} \times \left[\widehat{c}^{PE}_{\tau,t} \;-\; \xi \sum_{\ell=1}^{t} \theta^\ell (1-\xi)^{\ell-1} \widehat{c}^{PE}_{\tau,t-\ell}\right]$$

The displayed path has the expected shape: it scales in $1/\text{MPC} = 1/\omega$, features a significant news component, but is less persistent than the spending path $\widehat{\mathbf{c}}^{PE}_\tau$ that it engineers. In particular, given the assumed absence of anticipatory spending effects, the time-0 check is equal to $1/\text{MPC} \times \$150 = \$300$, i.e., exactly twice as large as the desired impact spending response. For stronger anticipation effects the impact stimulus check can be somewhat smaller, as I show in the supplementary discussion of Appendix D.1.

Finally, it remains to construct the full general equilibrium counterfactual for aggregate consumption, reported in the right panel of Figure 4. Caldara & Kamps (2017) estimate fiscal multipliers slightly in excess of one; as I show in Appendix D.3, their estimates correspond to persistently deficit-financed spending expansions with limited monetary offset. Summing a) the direct spending impact and b) the response of consumption to the fiscal spending shock, we see that the transfer shock depicted in the left panel induces a persistent increase in household consumption, with the full response slightly larger than the direct effect.

Summary. This section has illustrated the practical feasibility of the demand equivalence approach with applications to an increasingly popular policy tool: stimulus checks. By putting together standard pieces of cross-sectional and time series evidence, I conclude that such policies significantly stimulate aggregate consumption, with the estimated causal effect quite close to the direct micro estimates—a "missing intercept" close to zero.

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

[17] An alternative interpretation of (15)—consistent with forward-looking households—would be the following. On the one hand, unconstrained households behave in line with the permanent income hypothesis, so their consumption responds by very little. On the other hand, constrained households are in fact *fully* constrained, so they cannot respond to stimulus check *news*, and spend realized income gradually.

4 How plausible is demand equivalence?

The methodology and results presented in Section 3 are *exactly* valid only under the strong conditions required for demand equivalence. These conditions, however, are unlikely to hold in practice. I thus in this section apply the proposed methodology to models violating demand equivalence, and ask whether the estimated counterfactuals at least *approximately* equal the true model-implied causal effects.[18]

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

4.1 Estimated business-cycle models

My first model laboratory is an estimated HANK model, rich enough to feature many of the frictions popular in the quantitative business-cycle literature (e.g. Smets & Wouters, 2007; Justiniano et al., 2010). Models of this sort are routinely used for quantitative policy evaluation, and so in particular are a natural candidate for a fully structural solution to the aggregation (missing intercept) problem.

I build on the general framework of Section 2.1, and continue to impose Assumptions 1 and 2, but now relax Assumption 3. Demand equivalence in this generalized setting thus fails *only* because of the labor supply channel.[19]

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

$$\widehat{\mathbf{c}}_\tau \quad = \quad \widehat{\mathbf{c}}_\tau^{PE} \quad + \quad \widehat{\mathbf{c}}_g \quad + \quad error\left(\widehat{\boldsymbol{\ell}}_\tau^{PE}\right) \tag{16}$$

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

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

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

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

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

The household block is as described in Section 2.1, while the rest of the economy is designed to be as close as possible to the canonical model of Justiniano et al. (2010). First, I allow for investment adjustment costs, variable capacity utilization, and a rich monetary policy rule. Second, I extend the economy to be subject to a standard menu of aggregate shocks: to total factor productivity and the marginal efficiency of investment, to household patience, to wage mark-ups, to government spending, and to monetary policy. The only purpose of these additional shocks is to allow the model to fit aggregate U.S. business-cycle dynamics reasonably well, opening the door for a conventional likelihood-based estimation approach (An & Schorfheide, 2007). I calibrate the model's steady state using targets familiar from the HANK literature (e.g. Kaplan et al., 2018). Importantly, because household self-insurance is severely limited, the average MPC is high, at around 30% quarterly out of a lump-sum 500$ income gain. Model parameters governing dynamics are then estimated using likelihood methods on a standard set of macroeconomic aggregates.[20] The key exception is the degree of wage stickiness which—in light of its centrality to my results—is directly calibrated to be consistent with recent micro evidence (Grigsby et al., 2019; Beraja et al., 2019), with wage re-sets every 2.5 quarters on average. Most of the results in the remainder of this section refer to the estimated model's posterior mode.

Results. I subject the economy to a one-off stimulus check policy, and consider a researcher that uses the methodology of Section 3 to estimate that policy's aggregate effects.[21]

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

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

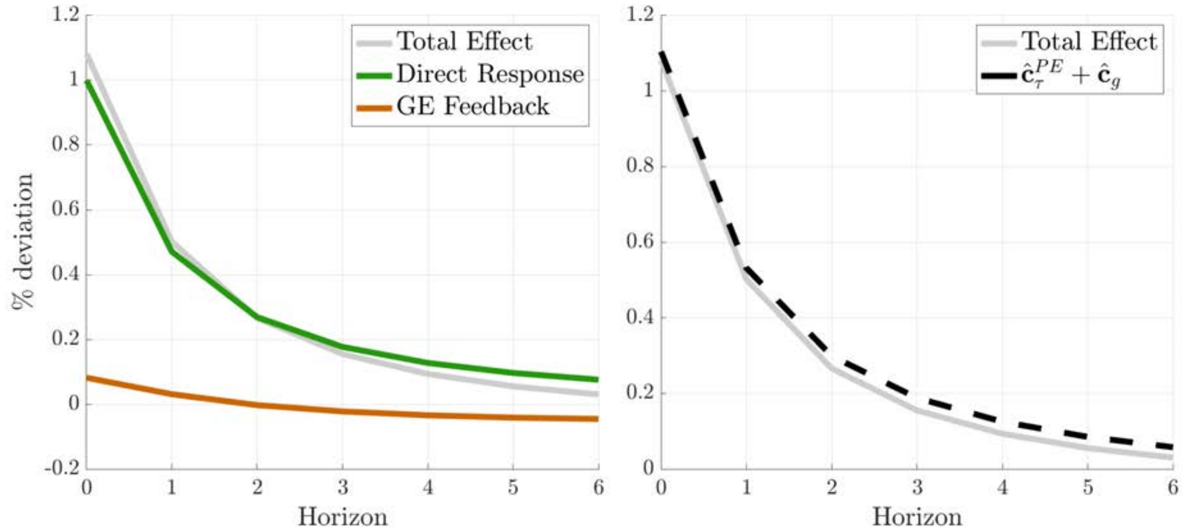

Figure 5: Consumption impulse response decompositions and demand equivalence approximation in the estimated HANK model, with details on the parameterization in Appendix B.2. The direct response and the indirect general equilibrium feedback are computed following Definition 1.

Results are displayed in Figure 5: the left panel decomposes the response of aggregate consumption to the stimulus check into direct partial equilibrium (green) and indirect general equilibrium (orange) effects, while the right panel compares the actual model-implied causal effect (grey) with the output of the my procedure (dashed black),

The main take-away from Figure 5 is that the demand equivalence approximation remains excellent, with the grey and black dashed lines in the right panel close to each other throughout.[22] The left panel first of all reveals that general equilibrium effects in the estimated model are small throughout, reflecting largely offsetting interest rate, tax financing, and Keynesian amplification effects. Following a similarly short-lived and deficit-financed fiscal spending expansion, the same forces imply that aggregate consumption barely moves, giving the small approximation error displayed in the right panel. The intuition for the sign and magnitude of that approximation error is simple. Following receipt of the stimulus check, households consume more. Given their lower marginal utility of consumption, they would optimally like to work less, thus in general equilibrium depressing aggregate output and consumption. This labor supply channel is absent after an increase in (unvalued) fiscal purchases, so the demand equivalence approximation *overstates* the response of consumption

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

to the transfer stimulus. However, even with strong wealth effects, this channel is largely irrelevant quantitatively: as long as wages are at least moderately sticky, labor is mostly demand-determined in the short run, so transitory shifts in labor supply do not matter much. This finding is consistent with conventional wisdom in the business-cycle literature (e.g. Christiano, 2011a,b): at least for relatively transitory fluctuations, hours worked in conventional (New Keynesian) business-cycle models are largely demand-determined.

Extensions & other models. The results in Figure 5 are neither special to the posterior mode of my estimated model, nor to the particular HANK setting considered here.

First, in Appendix E.1, I *randomly* draw model parameters from large supports, solve the implied model, and compute the approximation accuracy. The analysis reveals that, of all estimated parameters, only the degrees of price and wage rigidity have a material impact on the accuracy of the approximation, as expected. Second, in Appendix E.2, I use the demand equivalence approximation to construct counterfactuals for private demand shocks in the popular model of Justiniano et al. (2010), solved at the posterior mode. Since wages are even stickier there, the approximation is in fact better than in my estimated HANK model, with the approximation error now barely visible.

Discussion. The analysis in this section has demonstrated that demand equivalence is, at least approximately, a feature of standard quantitative models of business-cycle fluctuations. Intuitively, the features added to such models to ensure agreement with time series aggregates—for example investment adjustment costs, variable capacity utilization or price- and wage-indexing—are entirely consistent with demand equivalence. In particular, the common goods and financing assumptions are imposed regularly, and the labor supply channel is generally found to be quantitatively insignificant. We can thus strengthen the conclusions of Section 3: conventional business-cycle models, when estimated to be consistent with a) cross-sectional evidence on consumer spending behavior, b) time-series evidence on aggregate fiscal multipliers, and c) the behavior of standard U.S. time series aggregates, are likely to be in close agreement with my semi-structural stimulus check causal effect estimates presented in the two applications in Section 3.3.

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

4.2 Breaking equivalence

In this section I consider a large number of model extensions, each designed to challenge the quality of the demand equivalence approximation by breaking one or several of Assumptions 1 to 3. For each model variant, I begin with the baseline estimated HANK model of Section 4.1, and then add another friction to further break demand equivalence. My discussion here will be largely focussed on the sign and size of the bias induced by each of those model extensions, with model and calibration details relegated to Appendix B. Results are reported in Figure 6, which plots the demand equivalence error for each of my experiments, defined as

$$\text{error} \quad = \quad \frac{(\widehat{\mathbf{c}}_\tau^{PE} + \widehat{\mathbf{c}}_g) - \widehat{\mathbf{c}}_\tau}{\widehat{c}_{\tau 0}} \tag{17}$$

Note that (17) does *not* normalize the approximation errors to be positive. The fact that the errors displayed in Figure 6 all turn out to be positive is thus not an artifact of normalization, but in fact a key result.

Labor supply & wealth effects. Four experiments—the baseline HANK model, fixed wages, a model with flexible prices and wages, and a model with household preferences that imply weak wealth effects—illustrate the role of Assumption 3 on labor supply in breaking demand equivalence. As shown previously in Figures 1 and 5, in the estimated HANK model, the demand equivalence approximation is very accurate even with moderately sticky wages, and exact in the case of fully rigid wages. The purple line in Figure 6 shows that, with (nearly) flexible prices and wages, the quality of the approximation deteriorates sharply: because of quite strong wealth effects in labor supply, households cut hours worked following transfer receipt, and so the demand equivalence approximation—which misses these wealth effects—significantly overstates the response of aggregate consumption.

How material is this particular threat to the demand equivalence approach? I have already emphasized that, for relatively transitory shocks (such as one-off stimulus checks), even moderately sticky wages are enough to mute the labor supply channel. Other pieces of macro and micro evidence suggest the same conclusion. First, on the macro side, standard time series estimation exercises usually call for near-zero wealth effects in labor supply (Schmitt-Grohé & Uribe, 2012; Born & Pfeifer, 2014). Second, on the micro side, quasi-experimental evidence at the household level suggests that, at least in response to moderately sized lump-sum transfer receipts, hours worked and earnings drop by an order of magnitude less than

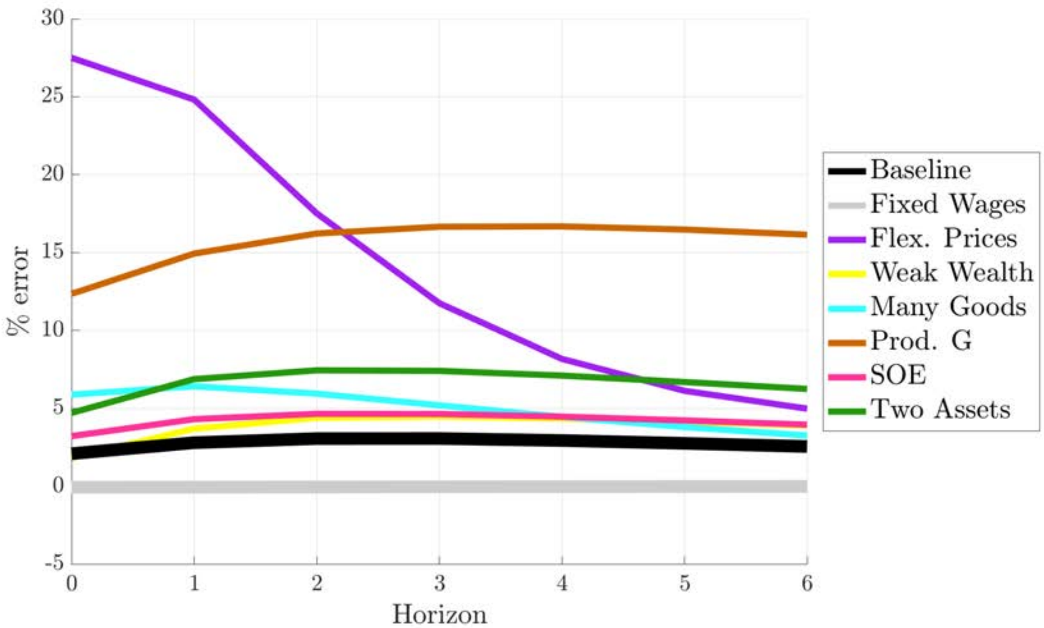

Figure 6: Errors (relative to the true impact consumption response) of the demand equivalence approximation in several model extensions. Details for all extensions and their parameterizations are relegated to Appendix B.

spending increases (e.g. Cesarini et al., 2017; Fagereng et al., 2018).[23] The yellow line in Figure 6 shows that, with fully flexible wages but household preferences adjusted to feature such data-consistent weak wealth effects, the approximation error is again small.

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

[23] Coibion et al. (2020) document similarly small earnings responses after stimulus payments in the COVID-19 recession. Mogstad et al. (2021) in contrast estimate significantly larger marginal propensities to earn, but their estimates come from much larger lottery wins (>$30,000). Consistent with this finding, I in Section 4.3 recommend that researchers should only apply the demand equivalence approximation to moderately sized spending shocks.

necessarily be shock-specific (Alonso, 2017; Baqaee, 2015).

To gauge the importance of these channels, I construct the demand equivalence approximation in a multi-good model in which: goods differ in their labor intensity; real relative prices fluctuate in response to (sectoral) shocks; and government expenditure is concentrated on the relatively more labor-intensive good. The cyan line in Figure 6 shows that these model extensions further reinforce the (still-present) wealth effect baseline error, with the positive bias now even more pronounced. The logic is as follows. First, the real relative price of the private consumption bundle naturally increases by more after stimulus checks than after an increase in fiscal purchases. Demand equivalence thus misses one channel of general equilibrium crowding-out. Second, since in my model MPCs out of labor income exceed those out of capital income, fiscal purchases have larger general equilibrium multipliers.

While positive throughout, the error is again moderate, at around double of the baseline model. First, even with prices adjusting every 2.5 quarters on average, transitory spending shocks induce only small relative price fluctuations, so the price channel is almost completely irrelevant. Second, and consistent with both Alonso (2017) and Baqaee (2015), I find that plausible differences in MPCs and factor incidence are not enough to yield sizable differences in multipliers. In the data, the average consumption good has a labor share of around of 0.4, while the network-adjusted labor share of government consumption is around 0.65. Even assuming an average quarterly MPC out of labor income of around 0.4, and an MPC out of any residual income of 0.1, the resulting second-round demand difference from spending on the two goods would only be around 7.5 cents for every dollar of spending.[24]

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

The orange line in Figure 6 reveals that productive benefits of government purchases can quite materially undermine the quality of the demand equivalence approximation. The approximation error is positive throughout, reflecting the fact that government purchases increase the economy's productive capacity and so crowd-in consumption—an amplification

[24] Arguably, this is an upper bound for the likely size of the effect, since heterogeneity in MPCs by skill implies the opposite conclusion: Government expenditure is concentrated on relatively high-skilled labor (Baqaee, 2015); if MPCs out of skilled labor are smaller, then the gap displayed in Figure 6 shrinks.

channel missing for transfer stimulus. In fact, since the productive benefits are long-lived, the error remains quite persistently elevated even for transitory shocks.

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

The pink line in Figure 6 shows the approximation error for an open economy with a home bias of 0.89 (matching the U.S.). As expected, openness increases the approximation error relative to the baseline economy: because of the lack of demand leakage, government spending is cheaper than the equivalent stimulus check, so taxes rise by less, leading to less general equilibrium crowding-out. However, given the substantial degree of home bias, it is not surprising that the error remains close to the baseline model throughout.

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

To get a sense of the likely magnitude of the implied approximation error, the green line in Figure 6 shows results for a two-asset model in which households pay an intermediation fee on liquid deposits, giving a positive bias and thus reinforcing the always-present labor supply error. I find that, even for an (implausible) *quarterly* return gap of 1.25 per cent, the error remains moderate, peaking at around 7 per cent of the true impact consumption response. To see why, suppose that, in response to a transfer stimulus, direct (partial equilibrium) household spending increases by $1 for one year. My approximation compares the aggregate effects of this shock to those of an identical expansion in public spending. Even if annual household and government discount rates differ by $4 \times 1.25 = 5$ per cent, the difference in present discounted values of the two spending expansions is just five cents—small compared to the initial size of the stimulus. The implied difference in tax financing is thus also small, and so the approximation remains quite accurate. Thus, even for large return gaps (in either direction), the bias is comparable in magnitude to the (small) labor supply-related bias.

4.3 Summary and recommendations for practice

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

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

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

5 Extensions

My theoretical and empirical analysis so far has been largely restricted to the stimulus check application. This section discusses the wider scope of the demand equivalence approach. First, in Section 5.1, I give examples of other shocks and policies covered by the consumption equivalence result. Second, in Section 5.2, I discuss the conditions required for an analogous *investment* demand equivalence result, and apply it to study another well-known fiscal stimulus policy: investment bonus depreciation.

5.1 Other consumer spending shocks

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

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

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

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

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

The direct response is thus estimable using information on household MPCs in the targeted sub-group, and so general equilibrium counterfactuals can be estimated as in Section 3.3. The demand equivalence approach can thus be used to gauge the effects of stimulus check policies even if the desired targeting has no historical precedent—we only need the corresponding net excess demand path to have been studied before in a time series experiment.[25]

[25] For the inverse step from fiscal multiplier estimates to equivalent targeted stimulus, the researcher would need to use the response matrix $\mathcal{C}_\tau^{\mathcal{T}}$ for the targeted sub-population.

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

I apply the demand equivalence approach to temporary increase in (labor) income inequality in Appendix F.2. The analysis proceeds in two steps, leveraging my second approach to demand matching discussed in Section 3.2. First, I construct a linear combination of fiscal spending experiments that has zero net present value and so can be entirely deficit-financed; that is, I consider a contraction of government spending today offset through an equivalent increase in the future. Second, I use a standard partial equilibrium consumption-savings problem to show that a short-lived increase in labor income inequality would engineer a very similar time path of aggregate net excess demand. My analysis thus naturally complements that in Auclert & Rognlie (2018): I use a partial equilibrium model to predict the direct effect on demand (like they do), but then I leverage demand equivalence together with time series fiscal policy shock evidence to solve the missing intercept problem, rather than relying on a full structural model. Overall my findings echo those of the stimulus check application: general equilibrium effects are relatively weak, so the total response of consumption is close to the shock's direct effect.

5.2 Investment

The demand equivalence logic can also be leveraged to estimate general equilibrium counterfactuals for shifters of *investment* demand. In this section I first sketch the conditions required for investment demand equivalence and then discuss an application to bonus depreciation stimulus. Details for theory and application are provided in Appendices A.2 and F.1.

Theory: investment demand equivalence. I again use the model of Section 2.1. Anticipating the empirical application, I augment the model to feature investment tax credit shocks $\boldsymbol{\varepsilon}_q$—shocks that reduce the cost of capital purchases by intermediate goods producers at time t by an amount $\tau_{qt} = \tau_{qt}(\boldsymbol{\varepsilon}_q)$.[26] I define direct (partial equilibrium) responses and indirect (general equilibrium) feedback for firm investment exactly analogously to Definition 1, using the implied aggregate investment function $\mathbf{i}(\bullet)$. As before, the question is: under what restrictions on primitives does the response of investment to a suitably chosen fiscal spending experiment give the "missing intercept" $\widehat{\mathbf{i}}_q^{GE}$?

The proof strategy is identical to that in Section 2: I characterize equilibrium response paths as a system of market-clearing conditions, and then impose enough restrictions on this system to ensure that the investment tax credit as well as a suitable fiscal experiment perturb the same equations by the same amount. In my model, the investment tax credit policy has three direct effects. First, investment responds; since investment invariably boosts the future productive capacity of the economy, production also increases, so the induced partial equilibrium net excess demand path for the final output and investment good is $\widehat{\mathbf{i}}_q^{PE} - \widehat{\mathbf{y}}_q^{PE}$. Second, the policy may be redistributive: the cost of financing is borne by taxed households, but the benefits accrue to households receiving dividend payments. The two groups need not be the same. Third, more investment and so more capital will increase the marginal product of labor, pushing up firm labor demand.[27]

Matching the first effect is straightforward: we simply need to consider a fiscal spending expansion with

$$\widehat{\mathbf{g}}_g = \widehat{\mathbf{i}}_q^{PE} - \widehat{\mathbf{y}}_q^{PE} \tag{18}$$

For the other two effects I require additional exclusion restrictions. To rule out heterogeneous distributional implications of tax financing and dividend payments, I assume that household income risk is perfectly insurable, thus effectively imposing a representative-household structure. This restriction also implies that Ricardian equivalence holds, so the precise timing of the policy financing is irrelevant. Next, to ignore the labor demand response, I assume that labor supply is perfectly flexible, either because of a large Frisch elasticity of labor supply, or again because labor is fully demand-determined. Under those two additional restrictions on primitives, a fiscal experiment satisfying (18) indeed gives the "missing intercept" of the investment response, in the sense that

$$\widehat{\mathbf{i}}_q \quad = \quad \widehat{\mathbf{i}}_q^{PE} \quad + \quad \widehat{\mathbf{i}}_g \tag{19}$$

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

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

Analogously we can also recover the output response as

$$\widehat{\mathbf{y}}_q \quad = \quad \widehat{\mathbf{y}}_q^{PE} \quad + \quad \widehat{\mathbf{y}}_g \tag{20}$$

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

Application: bonus depreciation. I apply the investment demand equivalence result to estimate general equilibrium counterfactuals for bonus depreciation stimulus—that is, the ability to tax-deduct investment expenditure at a faster rate.[28] I focus on bonus depreciation for three reasons. First, it is popular: with conventional monetary policy constrained by an effective lower bound, bonus depreciation has arguably become the go-to tool for investment stimulus. Second, previous empirical work has leveraged heterogeneity in firm exposure to the stimulus to estimate its direct effect on investment, $\widehat{\mathbf{i}}_q^{PE}$ (Zwick & Mahon, 2017; Koby & Wolf, 2020)—the key empirical input needed for my approach. Third, given the endogeneity of bonus depreciation policies to wider macroeconomic conditions, I am not aware of any studies that credibly identify the aggregate causal effects of such policies.

I only provide a brief summary of the results here, with details provided in Appendix F.1. Overall, my findings closely echo those of Section 3.3. Micro data indicate a large response of investment: following a one-quarter bonus depreciation shock worth around 8 cents (a shock similar in magnitude to the stimulus of 2008-2010), investment demand increases by around 16 per cent on impact, and then remains elevated. Combining the investment demand equivalence result in (19) with evidence on fiscal multipliers, I conclude that the increase in investment demand is accommodated through a sharp immediate increase in output as well as a smaller and somewhat delayed drop in consumption.

6 Conclusion

How can researchers learn about the "missing intercept" of cross-sectionally identified shifters of private spending? In this paper I ask whether—consistent with a simple Keynesian cross

[28]In the absence of firm-level financial frictions, such accelerated bonus depreciation schedules are isomorphic to the investment tax credits covered by the investment equivalence result (see Winberry, 2018).

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

I leave several avenues for future research. First, to be widely applicable, the demand equivalence approach requires time series estimates for a wide *menu* of fiscal spending experiments with different persistence and financing. More research on fiscal multipliers is thus welcome: it promises to not only tell us narrowly about those fiscal multipliers, but also about the propagation of many other shocks and policies. By the same token, running more cross-sectional regressions will help researchers construct the inverse mapping from demand path to primitive private spending shock. Second, other interesting macro shocks covered by the demand equivalence approach include firm uncertainty (Bloom, 2009; Bloom et al., 2018), shocks to firm credit conditions (Khan & Thomas, 2013) and household debt relief (Auclert et al., 2019). The methodology developed here could be applied to all of them. Third, I emphasize that the general conceptual idea of this paper—to leverage equivalence in the general equilibrium propagation of different shocks and policies—is not necessarily limited to fiscal spending and demand amplification, and so may help to solve the missing intercept problem in other contexts as well.

A Appendix

A.1 Proof of consumption demand equivalence

I begin by writing the equilibrium of the full baseline model as a dynamic system of market-clearing conditions.

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

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

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

Proof. See Appendix G of the Online Appendix. □

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

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

where J_i denotes the infinite-dimensional generalization of the selection matrix selecting the ith entry of a vector x_t. Assuming equilibrium existence and uniqueness,[29] there exists a unique linear map $\mathcal{H}$ such that

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

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

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

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

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

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

The impulse response paths of consumption thus satisfy

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

respectively. By Assumption 2 we know that, if $\widehat{\mathbf{c}}_\tau^{PE} = \widehat{\mathbf{g}}_g$, then setting $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$ is consistent with $\lim_{t\to\infty} \widehat{b}_t = 0$, since $\widehat{\boldsymbol{\tau}}_\tau^x$ and $\widehat{\mathbf{g}}_g$ by construction have the same net present value.[30] This establishes that, if $\widehat{\mathbf{c}}_\tau^{PE} = \widehat{\mathbf{g}}_g$ and $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$, then $\widehat{\mathbf{c}}_\tau^{GE} = \widehat{\mathbf{c}}_g$. (6) then follows. Finally note that, given the assumed fiscal financing rule $\boldsymbol{\tau}^e(\bullet)$, $\widehat{\boldsymbol{\tau}}_g^{e,PE} = \widehat{\boldsymbol{\tau}}_\tau^{e,PE}$ also implies $\widehat{\boldsymbol{\tau}}_g^e = \widehat{\boldsymbol{\tau}}_\tau^e$—the stated property (ii) of the fiscal spending shock. □

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

A.2 Details on investment demand equivalence

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

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

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

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

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

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

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

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

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

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

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

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

Proof. See Appendix G of the Online Appendix. □

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

References

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

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

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

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

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

Andreolli, M. & Surico, P. (2021). Less is more: Consumer spending and the size of economic stimulus payments.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, *113*(1), 1–45.

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

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

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

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

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

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

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

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review*, *110*(8), 2271–2327.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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