The Intertemporal Keynesian Cross

Adrien Auclert[*] Matthew Rognlie[†] Ludwig Straub[‡]

August 2018

Abstract

We demonstrate the importance of intertemporal marginal propensities to consume (iMPCs) in disciplining general equilibrium models with heterogeneous agents and nominal rigidities. In a benchmark case, the dynamic response of output to a change in the path of government spending or taxes is given by an equation involving iMPCs, which we call the intertemporal Keynesian cross. Fiscal multipliers depend only on the interaction between iMPCs and public deficits. We provide empirical estimates of iMPCs and argue that they are inconsistent with representative-agent, two-agent and one-asset heterogeneous-agent models, but can be matched by models with two assets. Quantitatively, models that match empirical iMPCs predict deficit-financed fiscal multipliers that are larger than one, even if monetary policy is active, taxation is distortionary, and investment is crowded out. These models also imply larger amplification of shocks that involve private borrowing, as we illustrate in an application to deleveraging.

[*]Stanford University and NBER. Email: aauclert@stanford.edu.

[†]Northwestern University and NBER. Email: matthew.rognlie@northwestern.edu.

[‡]Harvard University. Email: ludwigstraub@fas.harvard.edu.

We thank Sushant Acharya, Luigi Bocola, Chris Carroll, John Cochrane, Martin Eichenbaum, Mark Gertler, Dan Greenwald, Marcus Hagedorn, Greg Kaplan, Pete Klenow, Gisle Natvik, Jorge Miranda-Pinto, Kurt Mitman, Ben Moll, Martin Schneider, Morten Ravn, Gianluca Violante, Iván Werning, and Christian Wolf for helpful comments and suggestions, as well as Andreas Fagereng, Martin Holm and Gisle Natvik for generously providing us with empirical estimates of iMPCs. Remaining errors are our own.

1 Introduction

Within the last decade, quantitative macroeconomics has made significant advances in modeling household behavior. Modelers can now contemplate a wealth of alternatives to the traditional representative-agent framework. Popular options include simplified models with two agents, and models with a full distribution of agents holding either one or several assets. This abundance of options raises a number of key questions. Which model features are the most important in determining the economy's response to shocks and policies? What empirical evidence can we use to discipline these features? In turn, how does this evidence inform our choice of models?

A partial answer is provided by a recent literature which has argued that marginal propensities to consume (MPCs) are important moments for *partial equilibrium* effects. This was shown, for instance, by Kaplan and Violante (2014) in the context of fiscal policy, by Auclert (2017) for monetary policy, and by Berger et al. (2018) for house price changes.

In this paper, we propose a new set of moments—intertemporal MPCs, or iMPCs—and argue that they are essential for *general equilibrium* effects. We provide estimates of iMPCs in the data and find that, among typical modeling choices, only heterogeneous-agent models with multiple assets can match our estimates. We demonstrate that a model that matches iMPCs has distinct predictions for deficit-financed fiscal policy (our main application) and household deleveraging shocks.

We begin by setting up a benchmark framework in section 2 which nests a variety of common models from the literature. In line with our focus on household behavior, we keep the supply side simple at first and assume no capital, sticky wages, and a constant-real-rate monetary policy rule. We study fiscal policy, which sets paths for government spending G_t and aggregate tax revenue T_t, raised according to a progressive tax schedule.

In this framework, aggregate household behavior is entirely captured by an aggregate consumption function $\mathcal{C}_t\left(\{Y_s - T_s\}\right)$. $\mathcal{C}_t$ depends only on the path $\{Y_s - T_s\}$ of after-tax income in every time period s. Goods market clearing at each date then implies a fixed point equation in the path for output, $Y_t = \mathcal{C}_t\left(\{Y_s - T_s\}\right) + G_t$. Building on this fixed point, we show, to first order, that the impulse response of output $d\mathbf{Y} = (dY_t)$ to a change in fiscal policy $d\mathbf{G} = (dG_t)$, $d\mathbf{T} = (dT_t)$ solves a Keynesian-cross-like equation

$$d\mathbf{Y} = d\mathbf{G} - \mathbf{M}d\mathbf{T} + \mathbf{M}d\mathbf{Y} \tag{1}$$

Since this equation characterizes the entire dynamic path of output and stems from a microfounded model, we refer to it as the *intertemporal Keynesian cross*. The central object in the intertemporal Keynesian cross is the matrix $\mathbf{M} = (M_{t,s})$ of partial derivatives $M_{t,s} \equiv \partial\mathcal{C}_t/\partial Y_s$. For given dates t and s, $M_{t,s}$ captures the response of consumption at date t to an aggregate income shock at date s. Since the $M_{t,s}$ capture spending patterns over time, we refer to them as *intertemporal MPCs*.

Given their central role for fiscal policy, it is important to know what iMPCs are in the data and in our models. For measurement, in section 3, we focus on $M_{t,0}$—the dynamic response to an unanticipated income shock—which is where we have the best data. Our evidence on $M_{t,0}$ comes from two independent sources: the dynamic response to lottery earnings from Norwegian administrative data, as reported by Fagereng, Holm and Natvik (2018), and the distribution of self-reported marginal propensities to consume from the 2016 Italian Survey of Household Income and Wealth (Jappelli and Pistaferri 2018). Both sources confirm the common finding in the literature that the average impact MPC $M_{0,0}$ is high—above 0.4 at an annual level. The key new fact we uncover is that iMPCs in subsequent years are sizable as well, with $M_{1,0}$ lying above 0.14 according to both sources.

What models can match these patterns? Representative-agent (RA) and standard heterogeneous-agent models (HA-std) fail immediately on the grounds that they cannot match the high impact MPC. This is not an issue for two-agent (TA) models, which are sufficiently flexible to allow for arbitrary impact MPCs, by changing the share of hand-to-mouth agents. Yet, TA models predict that $M_{1,0}$ is around 0.02, almost an order of magnitude below our estimate.[1] The model with the best fit turns out to be a suitably calibrated heterogeneous-agent model with illiquid assets (HA-illiq). It is able to match both the high impact MPC as well as sizable subsequent iMPCs.[2]

iMPCs, therefore, are a useful device for distinguishing models. And from (1), we know that they characterize the general equilibrium response to a fiscal shock. But do the distinct patterns of iMPCs across models translate into equally distinct predictions for the impact of fiscal policy? As we explain in section 4, the answer turns out to depend crucially on the degree of deficit financing. When fiscal policy runs a balanced budget, iMPCs are in fact irrelevant: we derive a fiscal multiplier of exactly one in our benchmark framework, irrespective of iMPCs. Our result generalizes Gelting (1941) and Haavelmo (1945), who derived a balanced-budget multiplier of one in a static IS-LM model. It also generalizes the unit multiplier obtained in Woodford (2011) for the RA model, and thus provides a case where household heterogeneity is irrelevant—a fiscal policy analogue to Werning (2015)'s landmark result on monetary policy.

In contrast, iMPCs play a pivotal role for deficit-financed fiscal shocks. We show that for any fiscal policy that involves deficits, the fiscal multiplier is determined entirely by the interaction between iMPCs and the path of primary deficits. When iMPCs are (approximately) flat (RA, HA-std) the fiscal multiplier is (approximately) equal to 1. When the impact MPC is matched, but subsequent iMPCs are too low (TA), the impact fiscal multiplier dY_0/dG_0 can now lie significantly above 1; cumulative multipliers, however, are still equal to 1, pointing to a short-lived output response. Only when both impact and subsequent iMPCs are matched, as in the illiquid-asset model (HA-illiq), can impact and cumulative multipliers significantly exceed 1.

These findings suggest that matching iMPCs is important quantitatively. Our benchmark framework, however, is restrictive in several dimensions. To explore the role of iMPCs more gen-

[1] We also consider models with bonds in the utility (BU) which, conditional on matching $M_{0,0}$, overpredict $M_{1,0}$.

[2] Another model we find that fits the data is a two-agent version of a model with bonds in the utility ("TABU").

erally, we relax the benchmark's main limitations in section 5 and introduce capital, sticky prices, and active monetary policy. Since real rates now react to fiscal policy, important dampening forces appear, such as the crowding out of investment and the disincentive effects of distortionary taxation. We simulate the model for various degrees of deficit financing and confirm the role of the interaction between iMPCs and deficit financing. While all our models predict similar dynamics in the case of balanced-budget policies, the TA and HA-illiq models predict sizable impact multipliers under deficit financing. As in our benchmark framework, only the HA-illiq model predicts sizable cumulative multipliers that can lie above 1 for deficit-financed spending—now despite the addition of dampening feedback from interest rates.

To demonstrate the generality of our methodology, we extend our analysis to other shocks in section 6. We show that (1) also characterizes the transmission of these shocks. In that sense, iMPCs continue to be central, and it is important for models to match them. As before, this is particularly relevant when shocks involve deficit-financed spending—but now, we show that *private* deficits matter in addition to public deficits. We apply this general principle to two illustrative cases: deleveraging shocks, and fiscal shocks with lump-sum rather than progressive taxation. For the former, we find that the HA-illiq model predicts a \$3 drop in output on impact for every \$1 of deleveraging, as well as a negative cumulative output response. In contrast, the HA-std model barely amplifies the deleveraging shock, and the TA model features zero cumulative drop in output. For the latter, we show that the adverse redistributive effects of lump-sum taxation tend to reduce the multiplier, and that this can be understood as the result of smaller private deficits incurred by more heavily constrained taxpayers.

There is a large literature studying fiscal multipliers (see Hall 2009, Ramey 2011, and Ramey 2018 for surveys). Early theoretical analyses used the framework of the IS–LM model (Haavelmo 1945, Blinder and Solow 1973). The development of macroeconomic models with microfoundations enabled a quantification of mechanisms in the context of representative-agent models, from the role of the neoclassical wealth effect on labor supply (Aiyagari, Christiano and Eichenbaum 1992, Baxter and King 1993) to the role of monetary policy (Christiano, Eichenbaum and Rebelo 2011, Woodford 2011). Our benchmark partials out this role of monetary policy to focus on other factors likely to affect multipliers. Building on the Campbell and Mankiw (1989) "saver-spender" metaphor, Galí, López-Salido and Vallés (2007) introduced two-agent models to explain positive consumption multipliers in the data. As Coenen et al. (2012) documents, this class of models is the dominant paradigm for the study of fiscal policy in central banks today.

A recent literature has analyzed fiscal policy with heterogeneous agents. Oh and Reis (2012) was an early paper studying the effect of fiscal transfers. McKay and Reis (2016) focus instead on the role of automatic stabilizers. Ferriere and Navarro (2017) stress heterogeneous labor supply responses to changes in taxes in a model with flexible prices. Closest to our work is Hagedorn, Manovskii and Mitman (2017), who also study the effect of government spending in a model with nominal rigidities similar to ours. Their analysis is based on a different equilibrium selection criterion that relies on a long-run nominal debt anchor, following Hagedorn (2016). In contrast to our paper, which studies policy at the margin around the steady state, they focus on nonlinearities and the state dependence of multipliers. Our work is different in that we show the importance of iMPCs, provide analytical results in a benchmark case, and elicit why our model differs from RA and TA models. Both our studies conclude that heterogeneous-agent models differ from two-agent models, that deficit-financed fiscal multipliers can be significantly larger than one, and that balanced budget fiscal multipliers tend to be less than one, especially when taxes are raised lump sum.

There is also a vast empirical literature on fiscal multipliers based on aggregate macroeconomic evidence. As surveyed by Ramey (2018), this literature points to output multipliers in the range of 0.6–0.8, though the data does not reject multipliers as high as 1.5 (Ramey 2011, ben Zeev and Pappa 2017). The literature testing state dependence has mostly focused on the prediction from the representative-agent literature that multipliers differ depending on the extent of the monetary policy response (Auerbach and Gorodnichenko 2012, Ramey and Zubairy 2018). A robust prediction of our heterogeneous-agent model is that they also depend on the extent to which spending is deficit-financed. While the empirical literature acknowledges the potential importance of deficits, this prediction has not been subject to extensive testing.

Finally, our paper builds upon several lines of research that seek to discipline macroeconomic models with heterogeneity. A rapidly emerging literature identifies sufficient statistics for partial equilibrium effects (see, for instance, Kaplan and Violante 2014, Auclert 2017 and Berger et al. 2018). Auclert and Rognlie (2018) note that these partial equilibrium sufficient statistics can be converted into general equilibrium effects using numerical multipliers, while Farhi and Werning (2017) and Kaplan, Moll and Violante (2018) decompose aggregate consumption outcomes between the underlying inputs to the consumption function. Our paper combines these insights to show that the structure of general equilibrium itself can be reduced to a limited set of moments, intertemporal MPCs. To the best of our knowledge, these constitute the first set of sufficient statistics informing the general equilibrium propagation of shocks and policies.[3]

2 The intertemporal Keynesian Cross

In this section, we introduce our benchmark framework for the study of fiscal policy. The framework nests most of the common New Keynesian models in use in the literature, including those with heterogeneous agents. For this section, we make three simplifying assumptions that allow us to derive analytical results showing the central role of intertemporal marginal propensities to consume. As is standard in the New Keynesian literature, we abstract away from capital. We deviate a little from convention by assuming sticky wages, but flexible prices. Our main simplifying assumption is a constant-real-rate rule for monetary policy. This assumption allows us to partial out the effects of monetary policy so that we can focus on the potential effects of heterogeneity

[3] In recent contributions, Koby and Wolf (2018) use this methodology to study aggregate investment in models with firm heterogeneity, while Guren, McKay, Nakamura and Steinsson (2018) apply it in reverse, to convert general equilibrium estimates to partial equilibrium effects.

itself. We discuss the consequences of relaxing these assumptions at the end of this section, and we show that our main conclusions are robust to introducing capital, sticky prices and alternative monetary policy rules in section 5.

2.1 General framework

Time is discrete and runs from $t = 0$ to ∞. The economy is populated by a unit mass of agents, or households, who face no aggregate uncertainty, but may face idiosyncratic uncertainty. Agents vary in their idiosyncratic ability state e, which follows a Markov process with fixed transition matrix Π. We assume that the mass of agents in idiosyncratic state e is always equal to $\pi(e)$, the probability of e in the stationary distribution of Π. The average ability level is normalized to be one, so that $\sum_e \pi(e)\, e = 1$. If agents are permanently different, Π is the identity matrix and π the initial distribution over e.

Agents. In period t, agent i enjoys the consumption of a generic consumption good c_{it} and gets disutility from working n_{it} hours, leading to a time-0 utility of

$$\mathbb{E}\left[\sum_{t\geq 0} \beta^t \left\{u\left(c_{it}\right) - v\left(n_{it}\right)\right\}\right] \tag{2}$$

Pretax labor income is subject to a log-linear retention function as in Bénabou (2000) and Heathcote, Storesletten and Violante (2017).[4] This retention function is indexed to real wages, so that if P_t is the nominal price of consumption goods, W_t is the nominal wage per unit of ability, and e_{it} is the agent's current ability, real after-tax income is

$$z_{it} \equiv \tau_t \left(\frac{W_t}{P_t} e_{it} n_{it}\right)^{1-\lambda} \tag{3}$$

We nest standard models by allowing for various market structures. Agent i may trade in multiple assets a_i^j and face state- and asset-specific portfolio restrictions, so that the following constraints apply each period:

$$c_{it} + \sum_j a_{it}^j = z_{it} + (1 + r_{t-1}) \sum_j a_{it-1}^j \tag{4}$$

$$a_{it}^j \in \mathscr{A}_{e_{it}}^j \tag{5}$$

Agent i maximizes (2) by choice of c_{it} and a_{it}^j, subject to (4) and (5). By contrast, due to frictions in the labor market, all agents take their hours worked n_{it}—and therefore total after-tax income z_{it}—as given.[5] Hours are determined in general equilibrium.

[4] Heathcote et al. (2017) show that this provides a good approximation to the income tax schedule in practice.

[5] One advantage of this formulation is that it is consistent with weak wealth effects on labor supply in the short run, in line with empirical evidence on marginal propensities to earn. See Auclert and Rognlie (2017).

Labor market. Following standard practice in the New Keynesian sticky-wage literature, labor hours n_{it} are determined by union labor demand (Erceg, Henderson and Levin 2000, Schmitt-Grohé and Uribe 2005). Specifically, we assume that every worker i provides n_{ikt} hours of work to each of a continuum of unions indexed by $k \in [0,1]$. Total labor effort for person i is therefore

$$n_{it} \equiv \int_k n_{ikt} dk$$

Each union k aggregates efficient units of work into a union-specific task $N_{kt} = \int e_{it} n_{ikt} di$. A competitive labor packer then packages these tasks into aggregate employment services using the constant-elasticity-of-substitution technology

$$N_t = \left(\int_k N_{kt}^{\frac{\epsilon-1}{\epsilon}} dk \right)^{\frac{\epsilon}{\epsilon-1}}$$

and sells these services to final goods firms at price W_t.

We assume that there are quadratic utility costs of adjusting the nominal wage W_{kt} set by union k, by allowing for an extra additive disutility term $\frac{\psi}{2} \int_k \left(\frac{W_{kt}}{W_{kt-1}} - 1 \right)^2 dk$ in household utility (2). In every period t, union k sets a common wage W_{kt} per efficient unit for each of its members, and calls upon its members to supply hours according to a uniform rule, so that $n_{ikt} = N_{kt}$. The union sets W_{kt} to maximize the average utility of its members given this allocation rule.

In this setup, all unions choose to set the same wage $W_{kt} = W_t$ at time t and all households work the same number of hours, equal to

$$n_{it} = N_t \tag{6}$$

so efficiency-weighted hours worked $\int e_{it} n_{it} di$ are also equal to aggregate labor demand N_t. Observe that the combination of (6) with the retention function (3) implies that changes in N_t affect households' after-tax incomes z_{it} in proportion.

Appendix C.1 shows that the dynamics of aggregate nominal wage inflation $1 + \pi_t^w \equiv \frac{W_t}{W_{t-1}}$ are described by the following nonlinear[6] New Keynesian Phillips Curve:

$$\pi_t^w \left(1 + \pi_t^w\right) = \frac{\epsilon}{\psi} \int N_t \left(v'\left(n_{it}\right) - \frac{\epsilon - 1}{\epsilon} \frac{\partial z_{it}}{\partial n_{it}} u'\left(c_{it}\right) \right) di + \beta \pi_{t+1}^w \left(1 + \pi_{t+1}^w\right) \tag{7}$$

According to (7), conditional on future wage inflation, unions set higher nominal wages when an average of marginal rates of substitution between hours and consumption for households $v'\left(n_{it}\right) / u'\left(c_{it}\right)$ exceeds a marked-down average of marginal after-tax income from extra hours $\frac{\partial z_{it}}{\partial n_{it}}$.[7]

[6] As we show in appendix C.1, a linearized version of (7) takes a standard form $\pi_t^w = \kappa^w \left(\frac{1}{\phi} \frac{dN_t}{N} + \frac{1}{\nu} \frac{dC_t^*}{C^*} - \left(\frac{dZ_t}{Z_t} - \frac{dN_t}{N_t} \right) \right) + \beta \pi_{t+1}^w$ involving only aggregate hours N_t, after-tax income Z_t, and a "virtual consumption aggregate" C_t^* that captures all the effects of distributional changes on wage inflation.

[7] This term includes the distortions from labor income taxes, which are important for fiscal multipliers (Uhlig 2010).

Final goods producers. We assume a simple linear aggregate production technology

$$Y_t = N_t \tag{8}$$

Due to perfect competition and flexible prices, the final goods price is given by

$$P_t = W_t \tag{9}$$

and profits are zero, justifying why dividends do not enter households' budget constraints in (4). The real wage per efficient hour is constant and equal to $\frac{W_t}{P_t} = 1$. Thus, goods price inflation and wage inflation are equal at all times, $\pi_t \equiv \frac{P_t}{P_{t-1}} - 1 = \pi_t^w$.

Government. The government sets an exogenous plan for spending G_t and tax revenue T_t. Assuming initial government debt B_{-1}, the sequences $\{G_t, T_t\}$ must satisfy the intertemporal budget constraint[8]

$$(1 + r_{-1})B_{-1} + \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t-1} \frac{1}{1 + r_s} \right) G_t = \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t-1} \frac{1}{1 + r_s} \right) T_t \tag{10}$$

In each period $t \geq 0$, the government implements this plan by issuing or retiring debt as needed. Its outstanding debt at the end of period t is

$$B_t = (1 + r_{t-1})\, B_{t-1} + G_t - T_t \tag{11}$$

To raise tax revenue T_t, the government adjusts the coefficient τ_t on the labor income retention function according to

$$\int \left\{ \frac{W_t}{P_t} e_{it} n_{it} - \tau_t \left(\frac{W_t}{P_t} e_{it} n_{it} \right)^{1-\lambda} \right\} di = T_t \tag{12}$$

Given the path for goods prices P_t and a rule for the nominal interest rate i_t, the real interest rate on assets at t (the price of date-$t+1$ goods in units of date-t goods) is equal to

$$1 + r_t \equiv \frac{1 + i_t}{1 + \pi_{t+1}} \tag{13}$$

In this section, monetary policy sets the nominal interest rate i_t by following a simple rule, according to which the real interest rate is a positive constant

$$r_t = r > 0 \tag{14}$$

equal to the flexible-price steady state interest rate r. This is a special case of a Taylor rule, with a coefficient of 1 on expected inflation. Such a rule delivers a "neutral" monetary policy response to

[8]Observe that government debt B_t is specified in real terms and any plan must respect the intertemporal budget constraint, which rules out both the fiscal theory of the price level and equilibrium adjustment based on nominal bonds as in Hagedorn (2016).

fiscal shocks, in the sense that nominal interest rates rise exactly enough to offset the expected inflation these shocks create.[9] This allows us to focus our analysis on forces orthogonal to monetary policy before we consider more general Taylor rules in section 5.

Definition 1. Given an initial nominal wage W_{-1}, initial government debt B_{-1}, an initial distribution $\Psi_{-1}\left(\left\{a^j\right\}, e\right)$ over assets a^j, idiosyncratic states e, and exogenous sequences for fiscal policy $\{G_t, T_t\}$ that satisfy the intertemporal budget constraint (10), a general equilibrium is a path for prices $\{P_t, W_t, \pi_t, \pi_t^w, r_t, i_t\}$, aggregates $\{Y_t, N_t, C_t, B_t, G_t, T_t\}$, individual allocation rules $\left\{c_t\left(\left\{a^j\right\}, e\right), a_t^j\left(\left\{a^j\right\}, e\right)\right\}$, and joint distributions over assets and productivity levels $\left\{\Psi_t\left(\left\{a^j\right\}, e\right)\right\}$, such that households optimize, unions optimize, firms optimize, monetary and fiscal policy follow their rules, and the goods and bond markets clear

$$G_t + \int c_t\left(\left\{a^j\right\}, e\right) d\Psi_t\left(\left\{a^j\right\}, e\right) = Y_t \tag{15}$$

$$\sum_j \int a^j d\Psi_t\left(\left\{a^j\right\}, e\right) = B_t \tag{16}$$

Note that all assets pay the same equilibrium rate of return and there exists a unique market clearing condition for assets.

2.2 Nested models

Our formulation nests four major classes of models used to study fiscal policy.

Most of the models considered to date feature only one asset ($J = 1$). The standard *representative-agent model* (RA) is a model with a single productivity state $e = 1$ and without any portfolio constraints. The *two-agent model* (TA) features two permanent productivity states $\{e_1, e_2\}$ with equal productivity, $e_1 = e_2 = 1$, but different portfolio constraints: a mass $\pi(e_1) = \mu$ of fully constrained agents, $\mathscr{A}_{e_1} = \{0\}$, and a mass $\pi(e_2) = 1 - \mu$ of unconstrained agents. The *standard heterogeneous-agent model* (HA-std) works with many idiosyncratic states e, a unique stationary distribution π and a borrowing constraint, $\mathscr{A} = [\underline{a}, \infty)$.

A recent literature has studied two-asset models ($J = 2$). In this paper, we consider a simplified version of these models, which we call the *illiquid-asset heterogeneous-agent model* (HA-illiq). Agents face idiosyncratic risk, trade in a liquid asset on which they face a borrowing constraint, $\mathscr{A} = [\underline{a}, \infty)$, and also all hold an entirely illiquid asset $\mathscr{A}^{illiq} = \left\{a^{illiq}\right\}$, whose returns accrue to their liquid account.[10] This formulation allows our model to simultaneously match high average MPCs and a high level of aggregate wealth while retaining the tractability of a one-asset model. While

[9] Woodford (2011) uses a similar rule in a representative-agent model, as do McKay, Nakamura and Steinsson (2016) in a heterogeneous-agent model. One advantage of this rule is that the Phillips curve (7) only affects nominal quantities. A drawback in representative-agent models is that it can lead to indeterminacy. It turns out that our heterogeneous-agent model is locally determinate despite this rule (see Auclert, Rognlie and Straub 2018 for a determinacy result, which was also included in earlier drafts of this paper).

[10] Following (16), the aggregate combined holdings of liquid and illiquid assets in each period equal government debt B_t. In section 5, when we introduce monopolistically competitive firms and capital, the illiquid household asset will also include equity.

Kaplan et al. (2018) and Lütticke (2018) have shown that the possibility of trade between liquid and illiquid assets can matter for monetary policy, we believe that keeping illiquid holdings fixed represents a useful approximation for the study of fiscal policy.[11]

2.3 The aggregate consumption function

We now show that each of these models admits a simple representation of aggregate household behavior in general equilibrium. Starting from (3) and the fact that, in equilibrium, (6) and (9) hold, we can write aggregate after-tax income as

$$Z_t \equiv \int z_{it} di = \tau_t N_t^{1-\lambda} \int e_{it}^{1-\lambda} di \tag{17}$$

Combining (12) and $\int_i e_{it} n_{it} di = N_t = Y_t$, we can also write $Z_t = Y_t - T_t$. From (17), we see that individual after-tax income z_{it} is just a fraction of the aggregate

$$z_{it} = \frac{e_{it}^{1-\lambda}}{\int e_{it}^{1-\lambda} d\iota} Z_t \tag{18}$$

Substituting (18) into the household budget constraint (4), we see that the path of optimal policy rules $\{c_t(\{a^j\}, e), a_t^j(\{a^j\}, e)\}$ is entirely determined by the sequence of aggregate after-tax incomes $\{Z_t\}$. Thus, given the initial distribution $\Psi_{-1}(\{a^j\}, e)$, which we assume to be the ergodic steady-state distribution, aggregate consumption can be written entirely as a function of $\{Z_t\}$, that is,

$$\int_i c_{it} di = \mathcal{C}_t(\{Z_s\}) = \mathcal{C}_t(\{Y_s - T_s\})$$

We call $\mathcal{C}_t$ the *aggregate consumption function*.[12] Its existence relies only on the facts that in general equilibrium, household incomes are determined by the paths of macroeconomic aggregates through their effects on individual incomes, and that real interest rates are held constant by monetary policy. $\mathcal{C}_t$ encapsulates the potentially complex interactions between heterogeneity, macroeconomic aggregates, and the wealth distribution featured in our framework. Specifically, from the point of view of aggregate equilibrium behavior, the *entire* difference between the four models (RA, TA, HA-std and HA-illiq) is captured by differences in their aggregate consumption function. We now build on this observation to derive a simple representation of equilibrium.

2.4 The intertemporal Keynesian cross

A key condition in definition 1 is goods market clearing. Using Walras' law, it is simple to show that, given any path $\{G_t, T_t\}$ satisfying the government's intertemporal budget constraint (10), a

[11] See Fagereng et al. (2018) for evidence that households leave their illiquid asset positions almost entirely unchanged in response to income shocks.

[12] Similar aggregate consumption functions have been derived in Kaplan et al. (2018) and Farhi and Werning (2017), among others.

path of output $\{Y_t\}$ is part of an equilibrium if, and only if, it satisfies the equation

$$Y_t = G_t + \mathcal{C}_t\left(\{Y_s - T_s\}\right) \tag{19}$$

at all time periods t (see appendix A.1 for a proof). This fixed point equation contains all the complexity of general equilibrium.

Totally differentiating (19), we find that the first-order response of output $\{dY_t\}$ to a change in fiscal policy $\{dG_t, dT_t\}$ must satisfy

$$dY_t = dG_t + \sum_{s=0}^{\infty} M_{t,s}\left(dY_s - dT_s\right) \tag{20}$$

where we have defined the *intertemporal marginal propensities to consume*, or *iMPCs* for short, as

$$M_{t,s} \equiv \frac{\partial \mathcal{C}_t}{\partial Z_s} \tag{21}$$

We can collect the iMPCs in a matrix $\mathbf{M} \equiv (M_{t,s})$ whose s-th column $M_{\cdot,s}$ captures the dynamic response of aggregate consumption to an additional unit of after-tax income Z_s at date s. Since budget constraints must hold, each such additional unit of income is eventually spent. In other words, the present value of $M_{\cdot,s}$ is always equal to one, $\sum_{t=0}^{\infty} \frac{M_{t,s}}{(1+r)^{t-s}} = 1$.

Equation (20) is readily expressed in vector form. Defining $d\mathbf{Y} \equiv (dY_0, dY_1, \ldots)'$, and similarly $d\mathbf{G} \equiv (dG_0, dG_1, \ldots)'$ and $d\mathbf{T} \equiv (dT_0, dT_1, \ldots)'$, we obtain the following proposition.

Proposition 1 (The intertemporal Keynesian cross)**.** *If the first-order response of output $d\mathbf{Y}$ to a fiscal policy shock $\{d\mathbf{G}, d\mathbf{T}\}$ exists, it solves the* intertemporal Keynesian cross

$$d\mathbf{Y} = d\mathbf{G} - \mathbf{M}d\mathbf{T} + \mathbf{M}d\mathbf{Y} \tag{22}$$

If $\mathcal{M}$ is a linear map (defined on the space of summable sequences) with $(\mathbf{I} - \mathbf{M})\,\mathcal{M} = \mathbf{I}$ and $d\mathbf{G}$, $d\mathbf{T}$ are summable, then a solution to (22) is $d\mathbf{Y} = \mathcal{M}\left(d\mathbf{G} - \mathbf{M}d\mathbf{T}\right)$.

This proposition shows that our model gives rise to a Keynesian-cross-like relationship between output and government spending: $d\mathbf{Y}$ is given by the sum of government spending $d\mathbf{G}$ and the implied consumption response to the (endogenous) change in after-tax income $d\mathbf{Y} - d\mathbf{T}$. Unlike the traditional Keynesian cross, however, (22) is derived from a microfounded model, and crucially is a vector-valued equation, which captures intertemporal spending responses by agents through optimal borrowing and savings decisions.

Intertemporal MPCs as sufficient statistics. Note that the $\mathbf{M}$ matrix encapsulates the entire heterogeneity and micro structure of any model that matches the framework of section 2. Through its place in (22), $\mathbf{M}$ governs the effects of fiscal policy on output. Up to multiplicity in $\mathcal{M}$, knowledge of the iMPCs is therefore sufficient to compute $d\mathbf{Y}$ for any possible path $\{d\mathbf{G}, d\mathbf{T}\}$.

Determinacy. Our model may admit multiple equilibria. This is due to the presence of nominal rigidities, which are well known to lead to indeterminacy. The nature of indeterminacy is that there might be several linear maps $\mathcal{M}$ satisfying $(\mathbf{I}-\mathbf{M})\,\mathcal{M}=\mathbf{I}$. Below, we confine our attention to temporary and summable policies, implying that $\lim_{t\to\infty} dG_t = \lim_{t\to\infty} dT_t = 0$, and to the unique map $\mathcal{M}$ ensuring that $\lim_{t\to\infty} dY_t = 0$ for such policies. In fact, for the models with heterogeneous agents, HA-std and HA-illiq, this map gives the unique bounded solution $d\mathbf{Y}$ to equation (22), corresponding to the locally determinate equilibrium.

2.5 Extensions

We now briefly discuss how extensions of the environment alter the intertemporal Keynesian cross (22). Our approach turns out to be quite general. Across all of the following extensions, we obtain a generalized version of (22),

$$d\mathbf{Y} = d\mathbf{G} - \mathbf{M}^T d\mathbf{T} + \mathbf{M}^Y d\mathbf{Y} \tag{23}$$

where $\sum_{t=0}^{\infty} \frac{M_{t,s}^T}{(1+r)^{t-s}} = \sum_{t=0}^{\infty} \frac{M_{t,s}^Y}{(1+r)^{t-s}} = 1$ for all s. The sufficient statistics for the output response to fiscal policy are now the two matrices $\mathbf{M}^T$ and $\mathbf{M}^Y$ that reflect the response of aggregate demand to changes in taxes and income, respectively.

Alternative tax incidence. If the government finances its marginal expenses $d\mathbf{G}$ using alternative tax instruments that are not captured by (12), this requires a more general aggregate consumption function, $\mathcal{C}_t\left(\{Y_s; T_s\}\right)$, which separately depends on income and taxes. Thus, we obtain equation (23) by defining $M_{t,s}^T \equiv -\frac{\partial \mathcal{C}_t}{\partial T_s}$ and $M_{t,s}^Y \equiv \frac{\partial \mathcal{C}_t}{\partial Y_s} = M_{t,s}$ (see appendix B.1).

Durable goods. Suppose households also purchase durable goods, produced by a similar linear technology. In that case, the intertemporal Keynesian cross (23) holds with both $\mathbf{M}^T$ and $\mathbf{M}^Y$ now corresponding to intertemporal marginal propensities to spend, rather than to consume. We formally develop a simple model with durables along these lines in appendix E.

Investment. One can include investment by modifying the production technology to include both capital and labor. Maintaining the monetary policy rule (14), there now also exists an equilibrium *investment function* $\mathcal{I}_t\left(\{Y_s\}\right)$ that depends solely on the path for output. Intuitively, this path affects employment and therefore the prospective path for its marginal product of capital, determining investment decisions. The goods market clearing equation (19) is replaced by $Y_t = G_t + \mathcal{C}_t\left(\{Y_s; T_s\}\right) + \mathcal{I}_t\left(\{Y_s\}\right)$, where income and taxes no longer enter symmetrically into the consumption function due to revaluation effects. We obtain (23) with $\mathbf{M}^T = \mathbf{M}$ and $M_{t,s}^Y \equiv \frac{\partial \mathcal{C}_t}{\partial Y_s} + \frac{\partial I_t}{\partial Y_s}$, where the latter now contains the impulse responses of both consumption and investment to a unit increase in output at date s. Details can be found in appendix B.2.

Sticky prices. It is well known that sticky prices lead to countercyclical redistribution from wages to profits. This is especially important in heterogeneous-agent models since, depending on the distribution rule for profits, wage-earners and profit-earners do not necessarily coincide (see e.g. Werning 2015, Broer, Hansen, Krusell and Öberg 2016, Debortoli and Galí 2017).[13] However, in the natural case where agents earn profits in proportion to their current productivity e, these redistributive effects are neutralized and we obtain our benchmark equation (1) with $\mathbf{M}^T = \mathbf{M}^Y = \mathbf{M}$. See appendix B.3 for details.

Limitations of our approach. The commonality behind these extensions is that they can be reduced to a fixed point equation in the path for output delivered by the goods market clearing condition. We now briefly discuss when this approach fails to apply.

The main limitation of our approach is that it cannot easily handle the case of other monetary policy rules, or sticky prices with a distribution rule different from above. In the former case, the real interest rate responds to inflation; in the latter case, real earnings respond to inflation. Both these outcomes affect consumption, but wage inflation (7) is itself affected by consumption, leading to a fixed point problem that makes it more difficult to solve for $\mathbf{M}^T$ and $\mathbf{M}^Y$. In light of this limitation, the approach we follow in this paper is first to study the constant-real-rate, sticky-wage benchmark as a way to identify the relevant micro moments for fiscal policy, then to compare those moments to the data, and finally to demonstrate that the same moments are still relevant in a full-fledged quantitative model with sticky prices and alternative monetary policy rules.

3 Intertemporal MPCs in the models and the data

The intertemporal Keynesian cross in the previous section highlighted the crucial importance of iMPCs in determining the effects of fiscal policy. This raises an obvious question: how can we measure iMPCs in the data, and which models can match the evidence?

To answer this question, we proceed in three steps. We first collect the best available evidence on $\mathbf{M}$. Due to data limitations, this is unfortunately restricted to the first column of $\mathbf{M}$: the dynamic response to an unanticipated increase in income. This then allows us to distinguish between models to find those that are most consistent with the evidence. Finally, since the intertemporal Keynesian cross requires a complete matrix $\mathbf{M}$, we use the models most consistent with the evidence to fill in the later columns of $\mathbf{M}$.

3.1 Evidence on the response to unexpected income shocks

To estimate the first column of $\mathbf{M}$, we observe that it can be expressed as an average of individual responses to an unexpected income shock, $\partial c_{it}/\partial z_{i0}$, weighted by pretax income in the year of the

[13] In fact, this is one reason why we prefer to work with sticky wages in our benchmark model, since the interaction of these distributional effects with the countercyclicality of profits in the sticky-price New Keynesian model can have erratic consequences.

income shock,[14]

$$M_{t,0} = \int \frac{z_{i0}}{\int z_{i0} di} \cdot \frac{\partial c_{it}}{\partial z_{i0}} di \tag{24}$$

We propose two sources of evidence for the path of individual responses $\partial c_{it}/\partial z_{i0}$.[15]

Norwegian lottery evidence. Our first source of evidence comes from Norwegian administrative data, as analyzed in Fagereng et al. (2018). The data includes comprehensive information on consumption and uses the random winnings of lotteries to identify the dynamic consumption responses to income shocks. The authors' main estimating equation is

$$c_{it} = \alpha_i + \delta_t + \sum_{t=0}^{5} \gamma_k \text{lottery}_{i,t-k} + \theta X_{it} + \epsilon_{it} \tag{25}$$

where c_{it} is consumption of individual i in year t, α_i an individual fixed effect, δ_t a time effect, X_{it} are household characteristics, and $\text{lottery}_{i,t}$ is the amount household i wins in year t. The authors provided us with regression results weighted by after-tax incomes at the time of the lottery win.[16] Since lottery wins are not forecastable and are disbursed at the time they are announced, the estimated $\widehat{\gamma_k}$ precisely correspond to the weighted average in (24) and thus the first column of the iMPC matrix $\mathbf{M}$.

The black dots in figure 1 represent the point estimates for $\widehat{\gamma_0}$ through $\widehat{\gamma_5}$, together with 99% confidence intervals. Consistent with a large empirical literature, the annual MPC out of a one-time transfer is large, at about 0.55. What the literature has not stressed as much, but clearly appears in the Norwegian data, is that the iMPC in the year following the transfer is also fairly large, at around 0.16. This data point will turn out to be crucial to discriminate between models. After this point, the iMPCs slowly decay and become statistically insignificant around year 4.

A lower bound from Italian survey evidence. Our second source of evidence is a lower bound estimate for $\partial c_{it}/\partial z_{i0}$ constructed from survey data on MPCs. We implement it using the latest version of the Italian Survey of Household Income and Wealth (SHIW 2016), which asks survey respondents to report their annual contemporaneous MPC, $\partial c_{i0}/\partial z_{i0}$.

We obtain a point estimate for $M_{0,0}$ by weighting MPCs by income. To inform the later elements $M_{t,0}$ for $t > 0$, we propose the following idea based on the assumption that the distribution of MPCs remains the same over time.[17] Consider $M_{1,0}$. How small could this year-1 iMPC possibly

[14] For a proof, see appendix A.1. This approach also allows to measure $\mathbf{M}^Y$ and $\mathbf{M}^T$ separately by choosing weighting functions in line with the incidence of aggregate income and taxes.

[15] The existing literature mostly focuses on estimating contemporaneous marginal propensities to consume (which is helpful to inform $M_{0,0}$ in our notation), e.g. Shapiro and Slemrod (2003), Johnson, Parker and Souleles (2006), Blundell, Pistaferri and Preston (2008), Jappelli and Pistaferri (2014), and Fuster, Kaplan and Zafar (2018).

[16] Our reference estimates are their weighted full sample estimates, including responses to all sizes of lottery winnings up to \$150,000. An alternative would have been to restrict the sample to only small winnings. However, MPC estimates in this sample tend to be even larger than full sample estimates, do not sum to one over time, and are inherently imprecisely estimated due to the large noise-signal ratio.

[17] In appendix D.2 we validate this assumption by comparing the 2010 and 2016 distributions of MPCs.

Figure 1: iMPCs in the Norwegian and Italian data.

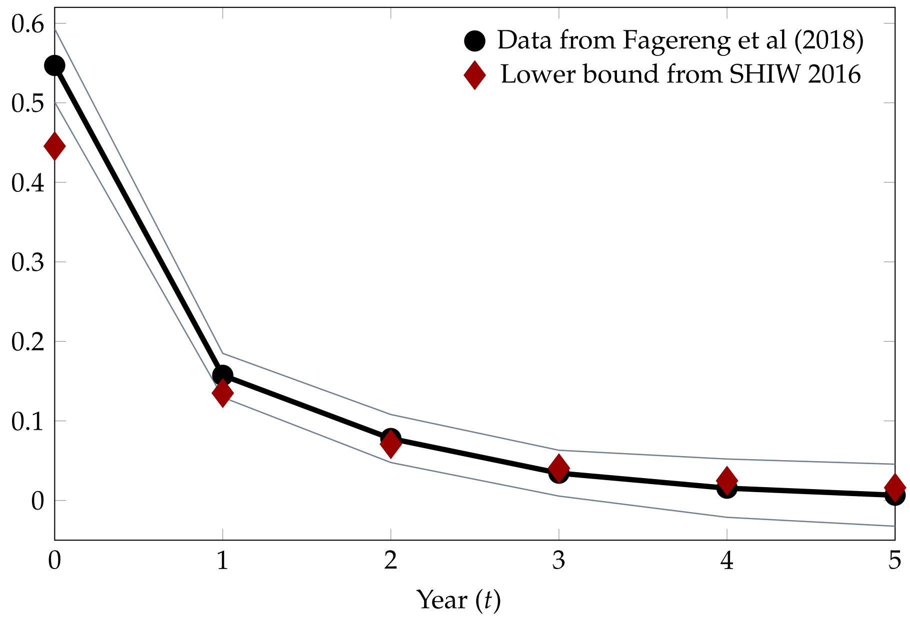

be for a given distribution of MPCs? It is smallest precisely when those households that save the most in year 0 are also the ones who save the most in year 1. In other words, a weighted average of $(1 - MPC_i) \cdot MPC_i$ delivers a lower bound on the true value of $M_{1,0}$. We extend this approach in appendix D.2 to all future iMPCs $M_{t,0}$ for $t > 0$.

The red diamonds in figure 1 show the lower bound estimates. The results are remarkably consistent with those obtained from the Norwegian administrative data. While the weighted contemporaneous MPC is slightly lower, at 0.42, the subsequent lower bound estimates are closely aligned with those obtained from the Norwegian data. The year-1 lower bound, in particular, is equal to 0.14 and thus close to the Norwegian estimate of 0.16. Recall that this point is a weighted average of $(1 - MPC_i) \cdot MPC_i$, so it is entirely accounted for by individuals in the sample that report *intermediate* MPCs, not too close to 0 or 1. Applying this logic in reverse suggests that matching our iMPC estimates will require models that generate an entire distribution of MPCs, including an important role for intermediately-constrained agents. This is what we confirm next.

3.2 Model discrimination

To assess the ability of the models described in section 2 to match the evidence reported in figure 1, it is necessary to calibrate them. Our calibration procedure follows literature standards and maintains maximal comparability across models and across sections of this paper.

Table 1 summarizes parameter estimates across models. In all models we consider, we assume that the economy is initially at a steady state. Households have constant CES utility over consumption $u(c) = \frac{c^{1-\nu^{-1}}}{1-\nu^{-1}}$ with an EIS of $\nu = \frac{1}{2}$, and a power disutility from labor $v(n) =$

Table 1: Calibrating the benchmark models.

Parameters	Description	Values					
		HA-illiq	RA	TA	HA-std	BU	TABU
ν	Elasticity of intertemporal substitution	0.5			(same across all models)		
ϕ	Frisch elasticity of labor supply	1			(same across all models)		
r	Real interest rate	5%			(same across all models)		
λ	Retention function curvature	0.181			(same across all models)		
G/Y	Government spending to GDP	0.2			(same across all models)		
A/Z	Wealth to after-tax income ratio	8.2			(same across all models)		
β	Discount factor	0.80	0.95	0.95	0.92	0.90	0.90
B/Z	Liquid assets to after-tax income	0.26	8.2	8.2	8.2	8.2	8.2
$\underline{a}$	Borrowing constraint	0			0		
μ	Share of hand-to-mouth households			52%			36%

$\gamma n^{1+\phi^{-1}}/(1+\phi^{-1})$ with Frisch elasticity $\phi = 1$. We set the curvature parameter of the retention function to $\lambda = 0.181$ as in Heathcote et al. (2017), assume that government spending is $\frac{G}{Y} = 20\%$ of output, and set γ to normalize steady-state output. We assume that steady-state inflation is $\pi = 0$ and that the steady state real interest rate is $r = 5\%$. We set β to match a wealth to after-tax (labor) income ratio of $\frac{A}{Z} = 8.2$ at that interest rate. While this number is larger than typically assumed for models without capital, it more accurately reflects the amount of effective liquidity in quantitatively realistic models with capital, and allows us to continue using the same household calibration once we introduce capital in section 5.

For the two models with idiosyncratic income risk and borrowing constraints, HA-std and HA-illiq, we follow standard practice in the literature and assume that gross income follows an AR(1) process. We use Floden and Lindé (2001)'s estimates of the persistence of the US wage process, equal to 0.91 yearly, set the variance of innovations to match the standard deviation of log gross earnings in the US of 0.92 as in Auclert and Rognlie (2018),[18] and discretize this process as an 11-point Markov chain. Following McKay et al. (2016), we also assume that households cannot borrow, $\underline{a} = 0$.

There has been recent interest in the ability of various tractable models to mimic properties of heterogeneous-agent models (see e.g. Debortoli and Galí 2017). While the TA model is the poster child for this approach, a recent promising alternative proposed by Kaplan and Violante (2018), Michaillat and Saez (2018) and Hagedorn (2018) is to introduce bonds in the utility function of a representative agent. To explore the consequences of such a model for iMPCs, we add this BU model to the set of models we consider.[19]

Our calibration targets fully specify both the RA and the HA-std model, while leaving one degree of freedom in the TA, the HA-illiq and the BU model. We use this extra degree of freedom to

[18] This is the same value as in Kaplan et al. (2018), and somewhat higher than the value of 0.70 implied by Floden and Lindé (2001), in order to capture the "new view" of household idiosyncratic risk.

[19] See appendix A.5 for details on the BU model.

Figure 2: iMPCs in the Norwegian data and several models.

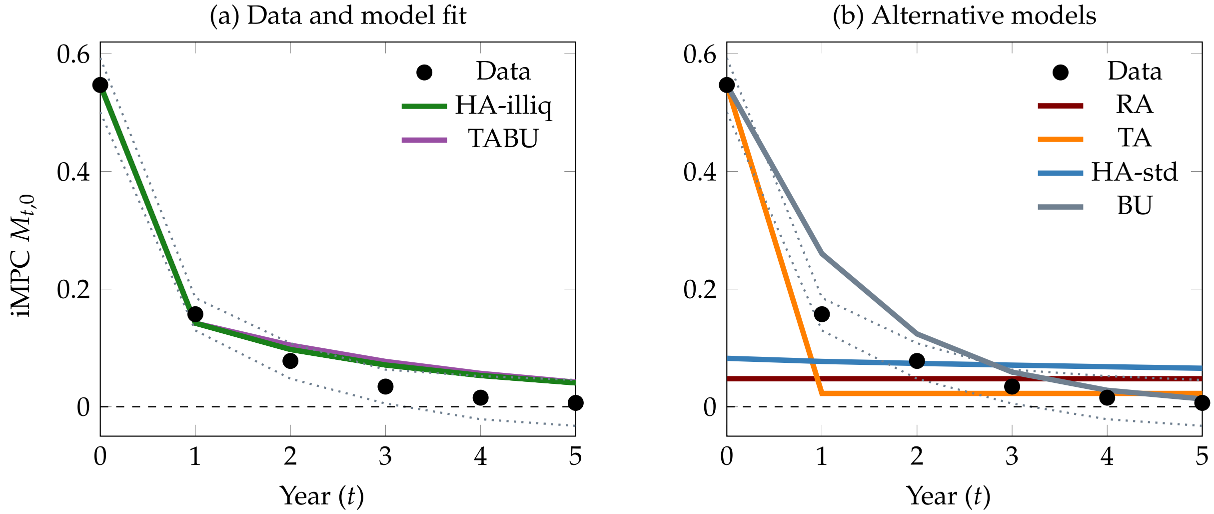

target the contemporaneous MPC $M_{0,0}$ of the Norwegian evidence in figure 1. The extra parameter of the TA model is the share of hand-to-mouth households μ. The extra parameter for the BU model is the curvature of the utility function over bonds. The extra parameter for the HA-illiq model is the amount of liquid bonds B. We find liquid bonds to be a fraction of $\frac{B}{Z} = 27\%$ of steady state after-tax income. This is somewhat lower than in Kaplan et al. (2018), who calibrate liquid assets to 26% relative to GDP, mostly because our calibration goal is to match the contemporaneous MPC estimate.

Figure 2 compares the model iMPCs to their counterparts in the Norwegian data. Panel (a) shows that despite our single degree of freedom B, the HA-illiq calibration matches the entire shape of estimated iMPCs relatively well. In particular, it is able to correctly reproduce the relatively high values of $M_{1,0}$ and $M_{2,0}$ suggested by both sources of evidence.

Panel (b) in figure 1 shows that the iMPCs implied by our alternative models all fail to match at least one important dimension of the estimated iMPCs. The iMPCs of the RA model are flat at a low level close to the real rate r, reflecting the permanent-income behavior of agents and entirely inconsistent with the data. The iMPCs of the TA model are also flat, except for the high impact MPC that the model is calibrated to match. Due to the absence of intermediately constrained agents, the TA model cannot generate elevated iMPCs in year one and later, which are a key characteristic of the data. The iMPCs of the standard heterogeneous-agent model HA-std are much closer to those of the RA model than to those of our HA-illiq model, echoing the approximate aggregation result of Krusell and Smith (1998). Finally, the BU model tends to deliver iMPCs for year 1 and 2 that are too large relative to the data, since its iMPCs decay exponentially over time.

Squinting at the iMPCs for the TA and the BU models suggests an alternative model that combines an agent with bonds in the utility with a fraction of hand-to-mouth agents. Such a *TABU* model has two degrees of freedom that can be calibrated to match the contemporaneous MPC as

Figure 3: Columns of the iMPC matrix in the HA-illiq model: $M_{\cdot,s}$ for $s = 0, 5, 10, 15, 20$.

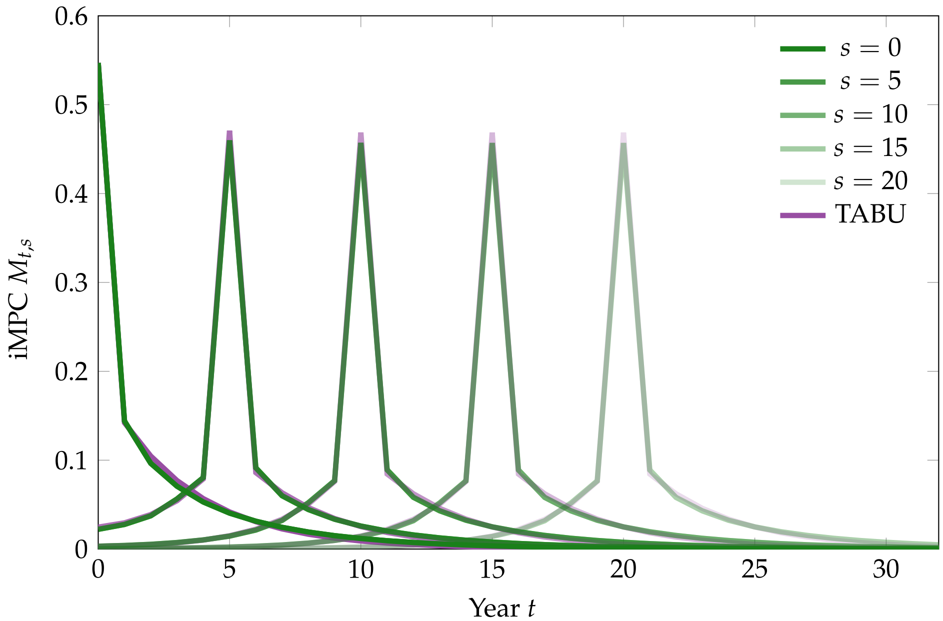

Note: The green lines show intertemporal MPCs implied by a two-asset heterogeneous-agent model (HA-illiq) that was calibrated to match empirical estimates of the first column of intertemporal MPCs. The purple lines show intertemporal MPCs of a two-agent bonds-in-the-utility model (TABU).

well as the subsequent iMPC $M_{1,0}$ of the HA-illiq model. The purple line in panel (a) shows the outcome of this procedure: the overall iMPC patterns are extremely similar.

3.3 The response to expected income shocks

With unlimited data, we would also estimate other columns of the iMPC matrix **M** directly. Unfortunately, there currently exists very limited information on consumption responses to anticipated changes in income one year out or later. Thus, we have to content ourselves with matching the first column and relying on models for extrapolation to other columns.

The green lines in figure 3 display the implications of our main model, HA-illiq, for the entire iMPC matrix. Each tent-shaped graph in the figure represents a column of the iMPC matrix—the response of aggregate consumption to an increase in aggregate after-tax income at some future date. The tent shape is a common feature of heterogeneous-agent models. The peaks of the tents decline for further-out income shocks, because income is spent partly in anticipation of its receipt. However, the declines of the tents to the right of their peak mirror the empirically-confirmed decline in first column iMPCs, which seems reasonable. Thus, our model is able to match the first column of the iMPC matrix directly and has intuitively reasonable implications for responses to future anticipated income shocks.

Consistency with existing empirical evidence. The limited evidence on consumption responses to anticipated income shocks generally confirms the pattern predicted by our model and visible in figure 3. For example, in their survey, Fuster et al. (2018) find that a few households would cut spending immediately in response to the news of a $500 loss one quarter ahead, indicative of some anticipation effects. Agarwal and Qian (2014) find evidence of a spending response between the announcement of a cash payout in Singapore and its disbursement two months afterward, and Di Maggio et al. (2017) find some evidence of one-quarter-ahead new car spending in expectation of a predictable reduction in mortgage payments.[20] However, this evidence is typically quarterly, not yearly as required by our model, and is too imprecise for us to confidently use as a model input.

Alternative models: TABU and durable goods. In the absence of direct empirical evidence, one way to confirm the predictions of our model for the later columns of $\mathbf{M}$ is to consider what alternative models with the same predictions for the first column would predict. We first consider the TABU model, which matches well the first column of the HA-illiq $\mathbf{M}$ matrix. Figure 3 shows that this model has almost identical predictions for later columns. This result makes us confident that the information contained in the impulse response to unexpected income shocks is informative about the impulse response to expected income shocks, since two very different models, once calibrated to match the former, also agree on the latter.

In appendix E, we also consider the predictions from a model with frictionless durables. This is an important exercise since the spending response that we target in the data includes durable spending, and agents have more scope for intertemporal substitution in durable spending. When calibrating the model to match the response of durables in the Norwegian evidence, we find very similar patterns for future iMPCs, except that spending is a little less elevated in the year after the income receipt as households decumulate some of their durables.

4 Fiscal policy in the benchmark model

We now solve the intertemporal Keynesian Cross to elicit the relationship between iMPCs and the impulse response to fiscal policy in our benchmark model. This relationship depends crucially on the financing of fiscal policy. We first consider the case of balanced budget policy and then study the general case.

It is standard in the literature to summarize the effects of fiscal policy on output using "multiplier" statistics. We follow the convention of reporting both the *impact multiplier* dY_0/dG_0 and the *cumulative multiplier* $\sum_{t=0}^{\infty}(1+r)^{-t}dY_t/\sum_{t=0}^{\infty}(1+r)^{-t}dG_t$ (see Mountford and Uhlig 2009 and Ramey 2018). The latter is sometimes considered a more useful measure of the overall impact of policy, capturing propagation as well as amplification of fiscal shocks.[21]

[20] By contrast, Kueng (2018) finds limited evidence of anticipation effects from the Alaska Permanent Fund news.

[21] The literature also sometimes refers to intermediate objects such as $\sum_{t=0}^{T}(1+r)^{-t}dY_t/\sum_{t=0}^{T}(1+r)^{-t}dG_t$ for some $T>0$. This number typically lies somewhere between our impact and cumulative multipliers.

4.1 Balanced-budget fiscal policy

Our first result is a sharp characterization of the effects of balanced-budget fiscal policy.

Proposition 2 (Balanced-budget policies, Haavelmo 1945)**.** *Assume the fiscal policy $\{d\mathbf{G}, d\mathbf{T}\}$ has a balanced budget, that is, $d\mathbf{G} = d\mathbf{T}$. Then, the fiscal multiplier is 1 at every date, $d\mathbf{Y} = d\mathbf{G}$.*

This result can easily be verified using (22) and our determinacy assumption. It is nevertheless very striking. Economically, it reflects the cancellation of two forces. Holding pretax incomes fixed, an increase in spending financed by a contemporaneous increase in taxes has an effect on output, $(I - \mathbf{M})\, d\mathbf{G}$, that is less than one-for-one. This is because consumption falls in response to the additional taxes. In equilibrium, however, pretax incomes rise, pushing up consumption. Our assumptions imply that the increase in pretax income exactly offsets the increase in taxes for every household at every date and state.[22] Households' consumption decisions are therefore unchanged, implying an output multiplier of exactly one at every date, irrespective of the timing of spending.

Proposition 2 provides a heterogeneous-agent counterpart to Woodford (2011)'s seminal representative-agent result under constant real interest rates. We view this result as the fiscal policy equivalent of Werning (2015)'s powerful irrelevance result for monetary policy. As such, it can serve as a useful benchmark as the literature on fiscal multipliers in heterogeneous-agent models develops (see, e.g., Ferriere and Navarro 2017 and Hagedorn et al. 2017).

4.2 Deficit-financed fiscal policy

While iMPCs are irrelevant for balanced-budget fiscal policies, they are central with deficit financing, as the following proposition emphasizes.

Proposition 3 (Deficit-financed policies)**.** *The output response to a fiscal policy shock $\{d\mathbf{G}, d\mathbf{T}\}$ is the sum of the government spending policy $d\mathbf{G}$ and the effect on consumption $d\mathbf{C}$,*

$$d\mathbf{Y} = d\mathbf{G} + \underbrace{\mathcal{M} \cdot \mathbf{M} \cdot (d\mathbf{G} - d\mathbf{T})}_{d\mathbf{C}}. \tag{26}$$

The consumption response $d\mathbf{C}$ only depends on the path of primary deficits $d\mathbf{G} - d\mathbf{T}$. In particular, holding the deficit fixed, government spending has a greater effect on output than transfers.

Proposition 3 highlights that for non-balanced-budget policies, the consumption response is entirely driven by the *interaction* between iMPCs—which determine $\mathcal{M}\mathbf{M}$—and primary deficits $d\mathbf{G} - d\mathbf{T}$. One implication of this is a clear relationship between government *spending* and *transfer* multipliers (see e.g. Giambattista and Pennings 2014 and Mehrotra 2014).

Consider a government spending plan that leads to a given path of primary deficits (or, equivalently, government debt). Under a constant-r rule, this has the same effect on aggregate con-

[22] Gelting (1941) and Haavelmo (1945) were the first to spell out this logic in the context of a static IS-LM model.

sumption as a transfer program that generates the same trajectory for government debt. Therefore, the government spending multiplier is equal to the sum of the transfer multiplier and the direct effect of government spending. Ramey (2018) argues that empirically, transfer multipliers as in Romer and Romer (2010) tend to be larger, in absolute value, than the output multiplier minus one. Our observation implies that this difference should be traceable to different monetary responses and/or different tax rules.

Multipliers in the RA and TA models. In two special cases, we can explicitly characterize the consumption and output responses. The first special case is the RA model.

Proposition 4 (Fiscal policy in the benchmark RA model)**.** *In the benchmark RA model, $d\mathbf{Y} = d\mathbf{G}$ irrespective of $d\mathbf{T}$. In particular, impact and cumulative multipliers are equal to 1.*

The reason for this stark result is that Ricardian equivalence holds in the RA model, so any policy is equivalent to a balanced-budget policy (proposition 2) and thus carries a unit multiplier. This result was first noted by Woodford (2011).

The second special case for which the solution of (26) is tractable is the TA model. This class of models has been very influential for the study of fiscal policy (see Galí et al. 2007, Bilbiie and Straub 2004, Coenen et al. 2012 and Bilbiie, Monacelli and Perotti 2013), yet to the best of our knowledge the following simple result has not been noted before.

Proposition 5 (Fiscal policy in the benchmark TA model)**.** *In the benchmark TA model, $d\mathbf{Y} = d\mathbf{G} + \frac{\mu}{1-\mu}(d\mathbf{G} - d\mathbf{T})$. The impact multiplier is equal to $\frac{1}{1-\mu} - \frac{\mu}{1-\mu}\frac{dT_0}{dG_0}$, but the cumulative multiplier is 1.*

This model is no longer Ricardian, and therefore generally produces non-unitary multipliers when $d\mathbf{G} \neq d\mathbf{T}$. In particular, as proposition 5 illustrates, output each period is determined by a static traditional Keynesian cross where μ, the share of constrained agents, plays the role of the standard MPC. The outcome is an impact multiplier of $1/(1-\mu)$ for a spending policy that is entirely deficit-financed, $dT_0 = 0$. Interestingly, however, the model *still* generates unitary cumulative multipliers, since consumption declines as soon as deficits are turned into surpluses. In this sense, the iMPCs of the TA model have implications similar to those of the RA model.

What happens in models whose iMPCs are further away from the RA model, as suggested by the micro evidence of section 3? This is what we study next.

Comparison across all models. To compare the output responses to fiscal policy outside of our two special cases, we consider a specific fiscal policy shock. We assume that government spending declines exponentially at rate ρ_G, $dG_t = \rho_G^t$. Taxes are chosen such that the path of public debt is given by $dB_t = \rho_B(dB_{t-1} + dG_t)$. In this formulation, ρ_B is the degree of deficit financing: if $\rho_B = 0$, the policy keeps a balanced budget, while for greater ρ_B, the policy leads to a greater deficit. We simulate the responses to this shock for various degrees of deficit financing and for the main models considered in figure 2, and compute the corresponding impact and cumulative multipliers.

Figure 4: Multipliers across the benchmark models.

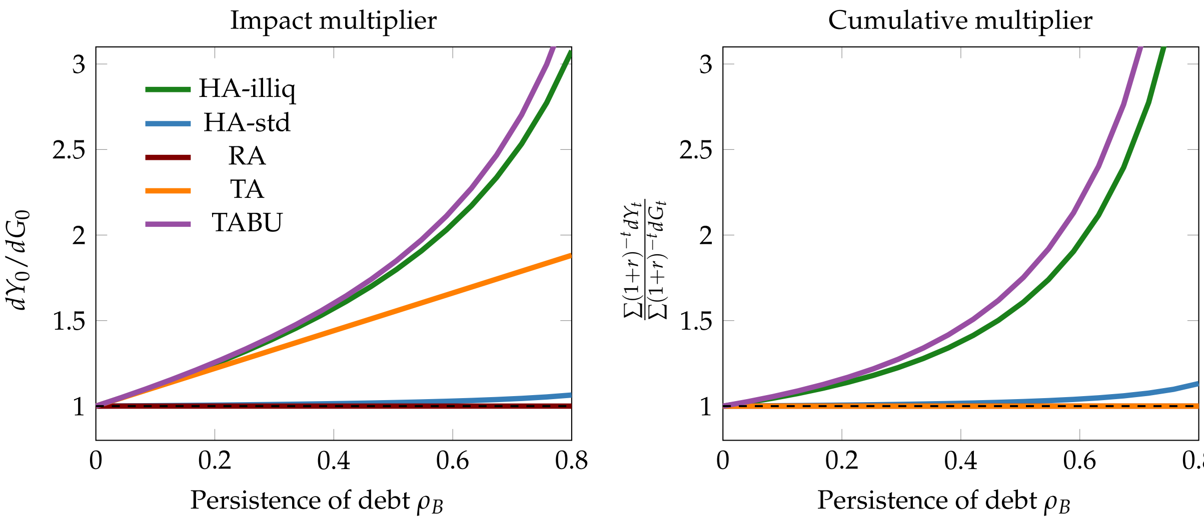

Figure 4 displays these multipliers. As per proposition 2, both impact (left panel) and cumulative multipliers (right panel) are exactly equal to 1 when fiscal policy balances the budget, irrespective of iMPCs. As the degree of deficit financing ρ_B rises, however, the models separate. The two models that provide a good match to the iMPCs in the data—the two-asset heterogeneous-agent model (HA-illiq) and the two-agent bonds-in-the-utility model (TABU)—lead to sizable impact *and* cumulative multipliers.[23] The models that imply approximately flat iMPCs—the RA model and the standard HA model—predict multipliers close to or exactly equal to one, in line with proposition 4. The TA model, as highlighted in proposition 5, lies somewhere in between: while it generates large impact multipliers under deficit financing, it always predicts a unitary cumulative multiplier.

4.3 Discussion

Role of distribution of tax policy. In section 2.5, we described a straightforward extension of the basic intertemporal Keynesian cross (22), which allows for fiscal policy with arbitrary tax incidence. What role does tax incidence play for the effects of fiscal policy on output? In section 6, we consider a case where the tax burden to finance a government spending shock is raised entirely with lump-sum taxes, as in Hagedorn et al. (2017). We show that this results in lower impact and cumulative multipliers, and we trace this outcome to the tighter constraints faced by the average taxpayer, which prevent much borrowing to smooth the cost of taxation. These smaller private deficits reduce the multiplier, mirroring the close connection we have found between public deficits and the multiplier.

[23] Note that these two models generate very similar multipliers for any ρ_B, since they have very similar iMPCs.

Are these multipliers similar to regional multipliers? The advent of well-identified cross-sectional evidence on the regional effects of government spending has raised the question of the relationship between regional "transfer" multipliers and closed-economy multipliers (see Nakamura and Steinsson 2014, Farhi and Werning 2016, and Chodorow-Reich 2017 for a recent survey). More specifically, could cross-sectional evidence distinguish between the vastly different models and closed-economy multipliers displayed in figure 4?

In a currency-union version of the models in this section, assuming that relative prices between regions are perfectly rigid (a natural benchmark), it is possible to obtain a clear dichotomy: cumulative regional multipliers are always equal to $1/(1-\text{home bias})$, independent of the model's iMPCs and the degree of deficit financing. Cumulative closed-economy multipliers, however, are independent of the home bias parameter and are entirely a function of the interaction of iMPCs and the degree of deficit financing. This dichotomy underscores the general difficulty in using the evidence on regional multipliers to inform the debate surrounding closed-economy multipliers.

5 Fiscal policy in the quantitative model

Figure 4 makes clear that models matching iMPCs imply large deficit-financed government spending multipliers. One reason is that we have thus far focused on a set of benchmark economies. We kept the "supply side" of these economies intentionally simple to focus on the role of intertemporal MPCs in determining the effects of fiscal policies. We now enrich the supply side by adding capital, sticky prices, and a Taylor rule for monetary policy. As expected, these modifications bring down multipliers across the board. Nevertheless, we show that intertemporal MPCs remain crucial in determining the overall effect of fiscal policy.

5.1 Extended model

Capital and sticky prices. To accommodate sticky prices, we now assume a standard two-tier production structure. Intermediate goods are produced by a mass one of identical monopolistically competitive firms, whose shares are traded and owned by households. All firms have the same production technology, now assumed to be Cobb-Douglas in labor and capital, $y_t = F(k_{t-1}, n_t) = k_{t-1}^{\alpha} n_t^{1-\alpha}$. Final goods firms aggregate intermediate goods with a constant elasticity of substitution $\mu/(\mu-1) > 1$. Capital is subject to quadratic capital adjustment costs, so that the costs arising from choosing capital stocks k_{t-1} and k_t in any period t are given by $\zeta\left(\frac{k_t}{k_{t-1}}\right) k_{t-1}$, $\zeta(x) \equiv x - (1-\delta) + \frac{1}{2\delta\epsilon_I}(x-1)^2$, where $\delta > 0$ denotes depreciation and $\epsilon_I > 0$ is the sensitivity of net investment to Tobin's Q. Finally, any firm chooses a price p_t in period t subject to Rotemberg (1982) adjustment costs $\xi(p_t, p_{t-1}) \equiv \frac{1}{2\kappa^p(\mu-1)}\left(\frac{p_t - p_{t-1}}{p_{t-1}}\right)^2$ where $\kappa^p > 0$.[24] An intermediate goods

[24] As in section 2.4, the impulse responses we compute are linearized to first order in aggregates. Hence, Rotemberg adjustment costs are equivalent to price setting à la Calvo (1983).

firm maximizes its value

$$J_t(k_{t-1}) = \max_{p_t,k_t,n_t} \left\{ \frac{p_t}{P_t} F(k_{t-1}, n_t) - \frac{W_t}{P_t} n_t - \zeta \left(\frac{k_t}{k_{t-1}} \right) k_{t-1} - \xi(p_t, p_{t-1}) Y_t + \frac{1}{1+r_t} J_{t+1}(k_t) \right\} \quad (27)$$

subject to the requirement that it satisfies final goods firms' demand in each period at its chosen price,

$$F(k_{t-1}, n_t) = Y_t \left(\frac{p_t}{P_t} \right)^{-\mu/(\mu-1)} \quad (28)$$

Since all these firms are identical, $k_t = K_t$, $n_t = N_t$, and $p_t = P_t$ in equilibrium. As we show in appendix C.2, this setup generates a nonlinear Phillips curve for price inflation

$$\pi_t (1 + \pi_t) = \kappa^p mc_t + \frac{1}{1+r_t} \frac{Y_{t+1}}{Y_t} \pi_{t+1} (1 + \pi_t) \quad (29)$$

where $mc_t \equiv \mu \frac{W_t/P_t}{F_{n,t}} - 1$ is the deviation between real marginal costs and their steady state value, as well as a set of standard Q theory equations for capital demand and the dynamics of investment.

Agents. Agents trade firms' stocks in addition to bonds. We denote by x_{it} agent i's date t total stock market position and allow for portfolio constraints $\mathcal{X}_e$. Defining dividends by $d_t \equiv J_t(K_{t-1}) - J_{t+1}(K_t)/(1+r_t)$, households face the budget constraint

$$c_{it} + \sum_j a^j_{it} + x_{it} = z_{it} + (1 + r_{t-1}) \sum_j a^j_{it-1} + x_{it-1} d_t$$

as well as portfolio constraints $a^j_{it} \in \mathscr{A}^j_{e_{it}}$ and $x_{it} \in \mathcal{X}_{e_{it}}$. The RA, TA, HA-std, and HA-illiq models are then extended as follows. The agent in the RA model is as always unconstrained. In the TA model, we assume the hand-to-mouth agent is also prevented from owning stocks, i.e. $\mathcal{X}_{e_1} = \{0\}$. In the HA-std model, we assume stocks cannot be used to borrow, $\mathcal{X} = [0, \infty)$.[25] And in the HA-illiq model, we assume stocks are held in illiquid accounts, $\mathcal{X} = \{1\}$.

Monetary and fiscal policy. The monetary authority now follows a Taylor rule,

$$i_t = r + \phi \pi_t \quad (30)$$

where the coefficient on inflation ϕ ensures determinacy[26] and r is the steady state interest rate. As before, the government follows an exponentially decaying spending policy, $dG_t = \rho^t_G$, with $\rho_G \in (0, 1)$. Taxes are chosen such that the path of public debt is given by $dB_t = \rho_B(dB_{t-1} + dG_t)$. As our baseline, we choose $\rho_G = 0.7$, which is in the range of usual estimates of the persistence of

[25] In the HA-std model, stocks and bonds are perfect substitutes and therefore the composition of agents' portfolio is indeterminate in steady state. Since it may nevertheless matter in response to shocks, we make the standard assumption that stocks are held in proportion to an agent's total asset position.

[26] As we show in Auclert et al. (2018) and in a previous version of this paper, the determinacy threshold may be above or below one. In our calibrated HA-illiq model, it is strictly below 1.

Table 2: Calibration of the quantitative supply side.

Parameter	Parameter	Value	Source
α	Capital share	0.33	Match $K/Y = 2.5$
B/Y	Debt-to-GDP	0.7	NIPA
K/Y	Capital-to-GDP	2.5	NIPA
μ	steady state markup	1.015	Match total wealth $= 3.5 \times GDP$
δ	Depreciation rate	0.08	NIPA
ϵ_I	Investment elasticity to Q	4	Literature on investment
κ^p	Price flexibility	0.1	Standard value
κ^w	Wage flexibility	0.1	Standard value
ϕ	Taylor rule coefficient	1.5	Standard value

US government spending (see, e.g., Davig and Leeper 2011 and Nakamura and Steinsson 2014), and $\rho_B = 0.7$, which corresponds to approximately 2 years of deficits following a fiscal shock, in line with typical responses of deficits in VAR evidence (see, e.g., Galí et al. 2007).

Equilibrium. Equilibrium is defined similarly to section 2.

Definition 2. Given initial values for the nominal wage W_{-1}, price level P_{-1}, government debt B_{-1}, capital K_{-1}, an initial distribution $\Psi_{-1}\left(\left\{a^j\right\}, e\right)$ over assets a^j and idiosyncratic states e, as well as exogenous sequences for fiscal policy $\{G_t, T_t\}$ that satisfy the intertemporal budget constraint (10), a general equilibrium is a path for prices $\{P_t, W_t, \pi_t, \pi_t^w, r_t, i_t\}$, aggregates $\{Y_t, K_t, N_t, C_t, B_t, G_t, T_t\}$, individual allocation rules $\left\{c_t\left(\left\{a^j\right\}, e\right), a_t^j\left(\left\{a^j\right\}, e\right)\right\}$, and joint distributions over assets and productivity levels $\left\{\Psi_t\left(\left\{a^j\right\}, e\right)\right\}$, such that households optimize, unions optimize, firms optimize, monetary and fiscal policy follow their rules, and the goods market clears

$$G_t + \int c_t\left(\left\{a^j\right\}, e\right) d\Psi_t\left(\left\{a^j\right\}, e\right) + I_t + \zeta\left(\frac{K_t}{K_{t-1}}\right) K_{t-1} + \xi(P_t, P_{t-1}) Y_t \quad = \quad Y_t$$

as do the asset market (16) and the stock market $\int x_{it} di = 1$.

Calibration. Our parameters and calibration targets are shown in Table 2. For each of our models, we maintain the same calibration of the steady-state household problem as the one we chose in section 3 by matching the same wealth to after-tax income ratio of 8.2 at the same interest rate $r = 0.05$. We stick to standard targets from the literature for the price flexibility (e.g. Schorfheide 2008, Kaplan et al. 2018), wage flexibility (Altig, Christiano, Eichenbaum and Lindé 2011), and investment elasticity parameters (Gilchrist and Himmelberg 1995), and show that our results are robust to the calibration in these parameters in appendix F. The model is simulated numerically using methods discussed in appendix G.

Figure 5: Response to a government spending shock in the quantitative HA-illiq model

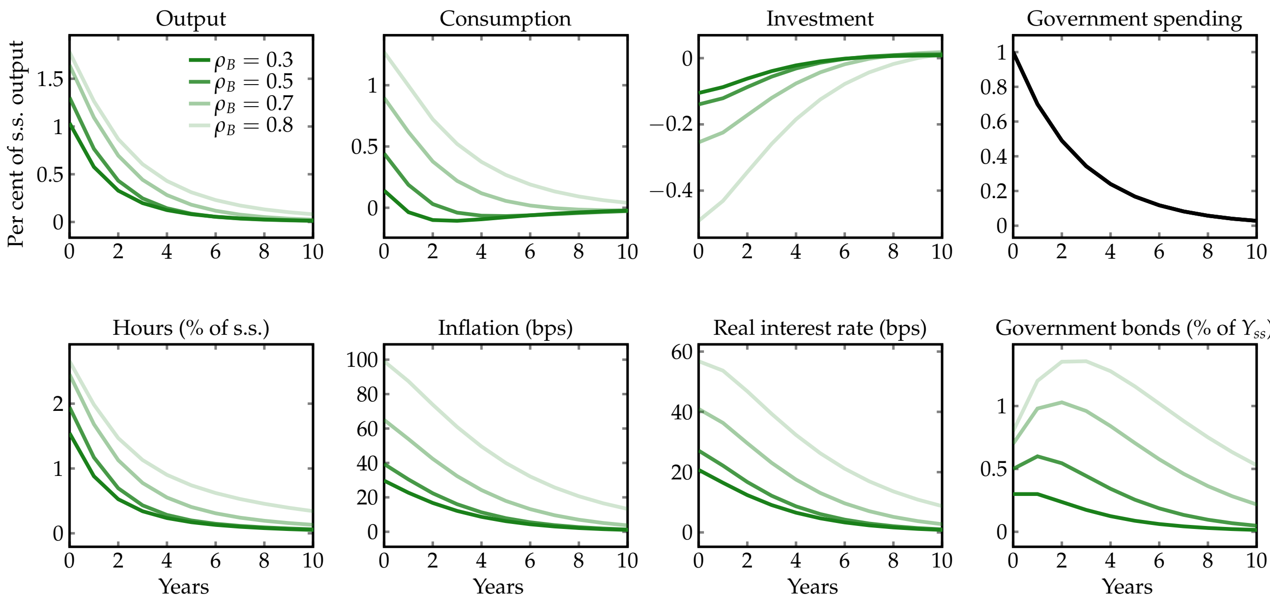

Note: The HA-illiq model is a two-asset heterogeneous-agent model which is calibrated to match evidence on intertemporal MPCs. The government spending shock declines exponentially at rate $\rho_G = 0.7$ and public debt follows $dB_t = \rho_B (dB_{t-1} + dG_t)$ for various values of ρ_B.

5.2 The quantitative effects of deficit-financed fiscal policy

What are the effects of fiscal policy with capital, a realistic monetary policy rule, and realistic iMPCs? Figure 5 shows the response of the HA-illiq economy to the government spending shock, for various degrees of deficit financing ρ_B.

The path of government spending is the same across all lines, but those with greater ρ_B lead to more elevated levels of public debt. Greater ρ_B is more stimulative and therefore leads to greater inflation. Since our Taylor rule coefficient is larger than one, real rates now rise in response. This crowds out investment, yet output still increases more than one for one vs. the spending shock, even for smaller values of ρ_B, as the consumption response is significantly positive. Thus, the quantitative version of the HA-illiq model still predicts significant output and consumption multipliers from deficit-financed government spending policies, despite rising real rates and the crowding out of investment.

One may wonder whether this conclusion holds up with somewhat larger Taylor rule coefficients ϕ, with greater investment sensitivities ϵ_I, and with greater price and wage flexibility parameters κ^p and κ^w. As we show in appendices 12–15, even as these parameters are modified, in most cases the quantitative version of the HA-illiq model predicts a greater than one-for-one output response and a positive consumption response.

Relative to the benchmark models in section 2, in this quantitative model consumption is no

Figure 6: Decomposing the consumption and investment responses

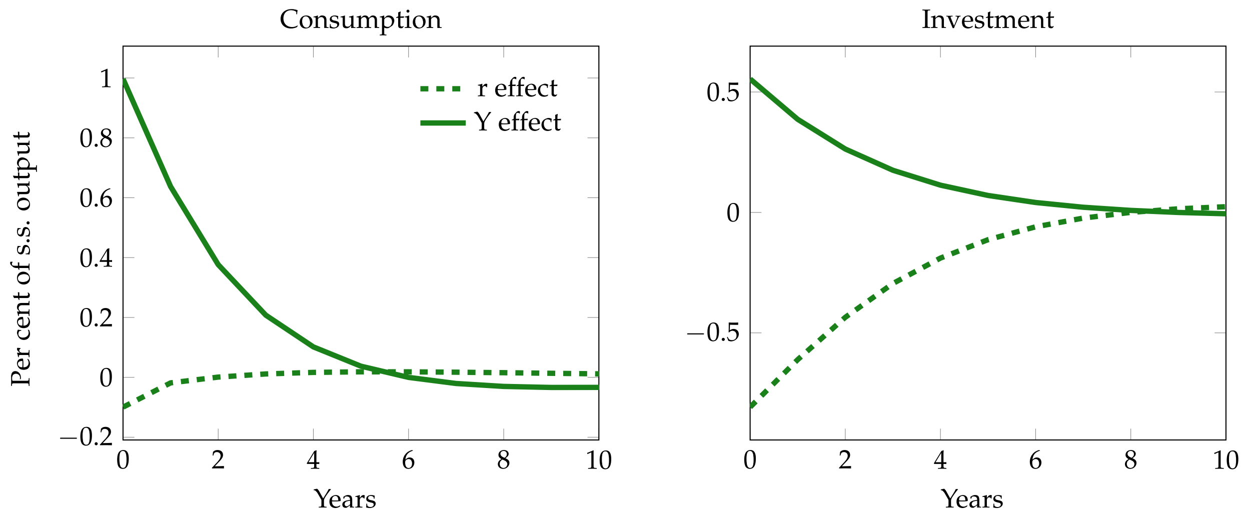

longer just a function of net incomes, but also depends on interest rates and dividends. Similarly, investment depends on aggregate incomes and interest rates. To isolate the contribution of iMPCs, figure 6 decomposes the responses of consumption and investment for $\rho_B = 0.7$ into two pieces: a Y *channel*, capturing the response to additional income and dividends; and an r *channel*, capturing the response to interest rates.[27] Both channels sum to the total effect in figure 5. For investment (right panel in figure 6), the two channels move in opposite directions. While the r channel leads to the familiar crowding out of investment, the Y channel turns out to crowd investment *in,* as greater demand for goods raises the marginal product of capital. This weakens the crowding-out effect significantly, even if it still dominates on net. For consumption (left panel), the Y channel clearly dominates the total response, hinting at the continued importance of iMPCs in the quantitative model. We explore the role of iMPCs more systematically next.

5.3 The role of intertemporal MPCs

A more direct way of inspecting the role of iMPCs is to compare the effects of fiscal policy predicted by our battery of models. Figure 7 repeats the exercise of figure 4 in this richer model. Observe first that all numbers are lower than their counterpart in the benchmark models. This is expected, given the additional dampening forces generated by an active monetary policy response and the attendant crowding out of investment.

Aside from this level shift, however, the results are closely in line with those of Section 3. The RA and standard HA models still predict multipliers close to independent of deficit financing ρ_B—in line with their almost flat iMPCs. The TA model still predicts impact multipliers that increase significantly in ρ_B, possibly above 1, yet its predicted effects are relatively short-lived,

[27] For similar decompositions, see Kaplan et al. (2018) and Kaplan and Violante (2018).

Figure 7: Multipliers in the quantitative models.

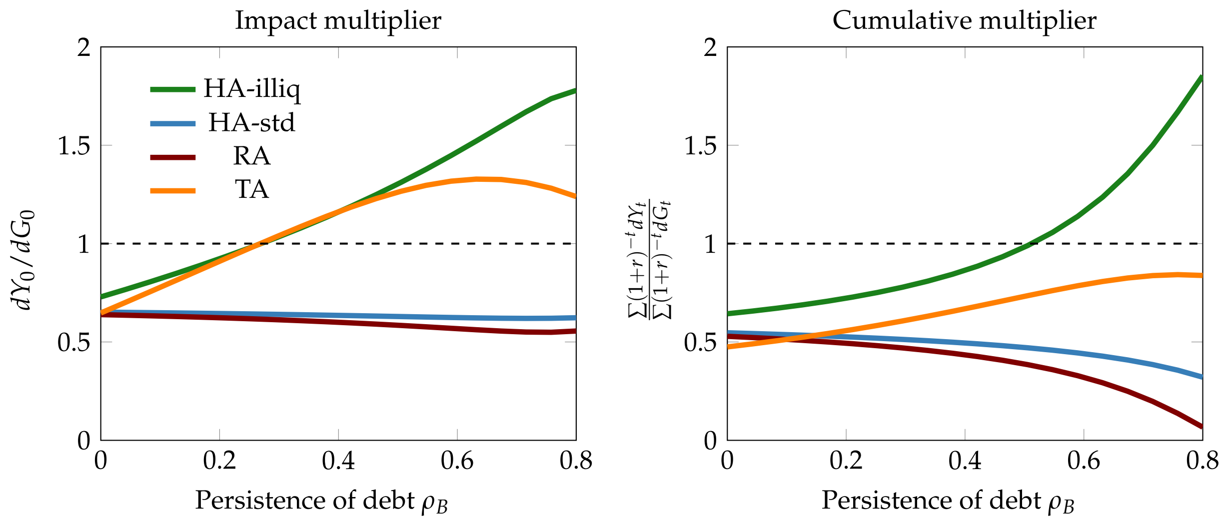

leading cumulative multipliers to be independent of ρ_B. This is in line with the TA model matching the high impact MPC but not any of subsequent intertemporal MPCs. Finally, the (two-asset) HA-illiq model predicts both sizable impact and cumulative multipliers with deficit financing—in line with it matching the high impact MPC *and* the high subsequent iMPCs.

5.4 Takeaway: Sizable multipliers with realistic iMPCs and deficit financing

Table 3 summarizes our main results for impact and cumulative multipliers, when computed using $\rho_B = 0.7$. The tables emphasize the *complementarity* between iMPCs and deficit financing: the combination of realistic iMPCs and deficit-financed fiscal policy predicts sizable multipliers above 1, both on impact and cumulatively.

The quantitative HA-illiq model's multipliers presented in Table 3 lie between 1.4 and 1.6 depending on the horizon, and are therefore toward the high end of the range typically estimated in aggregate data. The survey by Ramey (2018) concludes that the multiplier for temporary deficit-financed spending is "probably between 0.8 and 1.5", although reasonable people could argue that the data do not reject 0.5 or 2. There are two caveats, however, which complicate the comparison of our model-based conclusions with the data. First, in line with the theoretical literature, our economy was assumed to be entirely closed; we believe that openness would reduce multipliers somewhat. Second, the empirical literature typically characterizes a single type of "multiplier"; according to our model, however, the degree of deficit financing matters greatly for multipliers. This large dependence of cumulative multipliers on deficit financing constitutes a new testable prediction for the empirical literature, which to date has not been uncovered by either the RA or the TA literature. Third, we have confined our attention thus far to fiscal policies which adjust income taxes to raise tax revenues, without altering tax progressivity. However, the precise tax

Table 3: Complementarity between iMPCs and deficit financing: Multipliers across models.

Impact multipliers $\frac{dY_0}{dG_0}$

Fiscal rule	Model	RA	HA-std	TA	HA-illiq
bal. budget	benchmark	1.0	1.0	1.0	1.0
	quantitative	0.6	0.7	0.6	0.7
deficit-financed	benchmark	1.0	1.0	**1.8**	**2.5**
	quantitative	0.6	0.6	**1.3**	**1.6**

Cumulative multipliers $\frac{\sum_t (1+r)^{-t} dY_t}{\sum_t (1+r)^{-t} dG_t}$

Fiscal rule	Model	RA	HA-std	TA	HA-illiq
bal. budget	benchmark	1.0	1.0	1.0	1.0
	quantitative	0.5	0.5	0.5	0.6
deficit-financed	benchmark	1.0	1.1	1.0	**2.6**
	quantitative	0.2	0.4	0.8	**1.4**

Note. Simulated multipliers across four models: RA is a representative-agent model, HA-std is a heterogeneous-agent model with a standard calibration, TA is a two-agent model of whom one is hand-to-mouth, and HA-illiq is a two-asset heterogeneous-agent model calibrated to match evidence on intertemporal MPCs. "Benchmark" refers to a stylized model with linear technology and monetary policy implementing a constant real rate. "Quantitative" refers to a realistic model with capital and a Taylor rule. The government spending shock is an AR(1) with (annual) persistence $\rho_G = 0.7$ and an assumed persistence parameter of public debt of $\rho_B = 0.7$, corresponding to approximately 2 years of deficits.

instruments used can be crucial to multipliers. We highlight this idea as an example in the next section.

6 Other shocks

Up to now, our focus has remained firmly on the analysis of fiscal policy, where intertemporal MPCs and the intertemporal Keynesian cross can be understood intuitively as generalizing conventional MPCs and the traditional Keynesian cross. As we demonstrate in this section, however, both intertemporal MPCs and the intertemporal Keynesian cross also shape general equilibrium responses for a large class of other shocks. To make this point most tractably, this section is based on the benchmark model of section 2.

To include other kinds of shocks, allow the aggregate consumption function that we defined in section 2.3 to explicitly depend on an additional variable θ,

$$C_t = \mathcal{C}_t(\{Z_s\}, \theta) \tag{31}$$

Changes in θ can correspond to shifts in the interest rate (i.e. a monetary policy shock), in preferences, in the borrowing constraint (i.e. a deleveraging shock), or in the distribution of income and wealth across households.[28] Linearizing the goods market clearing condition (19) with (31) makes it clear how changes in θ can affect output:

$$dY_t = dG_t - \sum_{s=0}^{\infty} M_{t,s} dT_s + \mathcal{C}_{t,\theta} d\theta + \sum_{s=0}^{\infty} M_{t,s} dY_s \tag{32}$$

Relative to (20), this adds an additional term $\mathcal{C}_{t,\theta} d\theta$. This is the first-order effect of the shock $d\theta$ on consumption at date t, assuming no changes in the path $\{Z_s\}$ of aggregate after-tax income. In other words, it is the *direct* consumption effect of a shock to households, prior to any general equilibrium feedbacks.

Stacking $\partial\mathbf{C} \equiv (\mathcal{C}_{t,\theta} d\theta)$ and moving to vector notation, this allows us to generalize the intertemporal Keynesian cross with a simple extension of proposition 1.

Proposition 6 (Generalized intertemporal Keynesian cross)**.** *The first-order response of output $d\mathbf{Y}$ to a shock to $\{d\mathbf{G}, d\mathbf{T}\}$ and θ solves the* (generalized) intertemporal Keynesian cross

$$d\mathbf{Y} = d\mathbf{G} - \mathbf{M} d\mathbf{T} + \partial\mathbf{C} + \mathbf{M} d\mathbf{Y} \tag{33}$$

If $\mathcal{M}$ is a linear map (defined on the space of summable sequences) with $(I - \mathbf{M})\,\mathcal{M} = I$, and $d\mathbf{G}, d\mathbf{T}, \partial\mathbf{C}$ are summable, then a solution to (33) *is $d\mathbf{Y} = \mathcal{M}(d\mathbf{G} - \mathbf{M} d\mathbf{T} + \partial\mathbf{C})$.*

The generalized intertemporal Keynesian cross (33) is exactly the same as our earlier version (22), but with an added $\partial\mathbf{C}$ term on the right. Similarly, the solution is the same as in proposition 1, but with $\mathcal{M}$ multiplying a new $\partial\mathbf{C}$ term.

An important lesson of proposition 6 is that both fiscal shocks and this much broader family of shocks—to interest rates, preferences, borrowing constraints, and distribution—work through the same general equilibrium mechanisms. Indeed, this allows us to make predictions about transmission from partial to general equilibrium. For instance, if a deleveraging shock has a direct consumption effect of $\partial\mathbf{C}$, then a fiscal shock that perturbs the path of government spending by $d\mathbf{G} = \partial\mathbf{C}$ while leaving taxes unaffected will have exactly the same general equilibrium output effect, because $d\mathbf{G}$ and $\partial\mathbf{C}$ enter interchangeably in (33).[29]

We can gain additional insight by rewriting (6) as

$$d\mathbf{Y} = \underbrace{d\mathbf{G} - d\mathbf{T}}_{\text{public deficits}} + \underbrace{(I - \mathbf{M}) d\mathbf{T} + \partial\mathbf{C}}_{\text{PE private deficits}} + \mathbf{M} d\mathbf{Y} \tag{34}$$

The novel feature here is $(I - \mathbf{M}) d\mathbf{T} + \partial\mathbf{C}$, which we call the *partial equilibrium path of private deficits*. This combines net household spending $(I - \mathbf{M}) d\mathbf{T}$ from the change in taxes—the taxes themselves

[28] Since an interest rate shock also perturbs the government budget constraint, it comes together with some perturbation $d\mathbf{T}$ to the path of tax revenue, reflecting the government's plan to finance the marginal interest cost.

[29] Note that the government spending shock $d\mathbf{G} = \partial\mathbf{C}$ obeys the government budget constraint because the aggregated household budget constraint requires that $\partial\mathbf{C}$ have net present value zero. The equivalence between $d\mathbf{G}$ and $\partial\mathbf{C}$ is one reason we use fiscal shocks as our entry point to the intertemporal Keynesian cross: shocks to the path of government spending are a natural reference case for many other shocks.

$d\mathbf{T}$ minus the implied consumption decline $\mathbf{M}d\mathbf{T}$—with the direct effect $\partial\mathbf{C}$ of the shock on household consumption. It is "partial equilibrium" because it excludes general equilibrium adjustments in output. (34) tells us that the output effect of a shock is determined solely by the sum of public deficits and partial equilibrium private deficits.

This is a useful complement to proposition 3, where we derived a special representation (26) for fiscal shocks, relating consumption to public deficits. In contrast, in the more general representation (34), we relate output to *combined* public and partial equilibrium private deficits.

It is important to note that since we have a closed economy model, in general equilibrium the combined public and private deficit must be zero in every period. This is a consequence of goods market clearing. The sum in (34) is not zero because it captures only partial equilibrium private deficits, prior to any market-clearing general equilibrium adjustments in output. Indeed, the insight of (34) is that partial equilibrium deficits are what ultimately determine these general equilibrium outcomes.

Interaction with iMPCs: both public and private deficits. In section 4.2, we saw that the output effect of a fiscal shock was determined by the interaction between iMPCs and the path of public deficits $d\mathbf{G} - d\mathbf{T}$. Since public and partial equilibrium private deficits enter interchangeably in (34), this logic should apply equally to private deficits. Thus, any policy or shock that generates greater partial equilibrium private deficits is more stimulative in exactly the same way that greater public deficits are more stimulative.

We showcase this powerful logic in the following two examples.

Example 1: Deleveraging shock. One of the most natural shocks that directly shapes the path of private deficits is a deleveraging shock, in which a sudden tightening of borrowing constraints forces households near the debt limit to delever—in other words, to run negative deficits.

Up until now, we have worked with a constant borrowing constraint $\underline{a} = 0$ for the two HA models (HA-illiq and HA-std) as well as for the hand-to-mouth household in the TA and TABU models. We now model a deleveraging shock (as in e.g. Guerrieri and Lorenzoni 2017) such that the borrowing constraint $\underline{a}_t$ between periods t and $t+1$ tightens[30] for some time at rate ρ but eventually mean-reverts at rate $\rho_{\underline{a}} \in [0, 1)$:

$$\underline{a}_{t+1} = \rho_{\underline{a}} \left(\underline{a}_t + \rho^t \epsilon\right)$$

Figure 8 shows the direct consumption effect $\partial\mathbf{C}$ and the general equilibrium response of output $d\mathbf{Y}$ across models, for $\rho = 0.7$ and $\rho_{\underline{a}} = 0.7$. To focus on general equilibrium propagation, we ensure that the direct effects are similar across models by choosing ϵ to set $\partial C_0 = 1$ in each model.

[30] In order to use the same calibration as the rest of the paper, we retain $\underline{a} = 0$ as our steady state borrowing limit, and therefore our "deleveraging" shock forces households to hold a strictly positive asset position. Results are qualitatively and quantitatively similar if instead we start from a calibration with $\underline{a} < 0$.

Figure 8: The effects of deleveraging shocks.

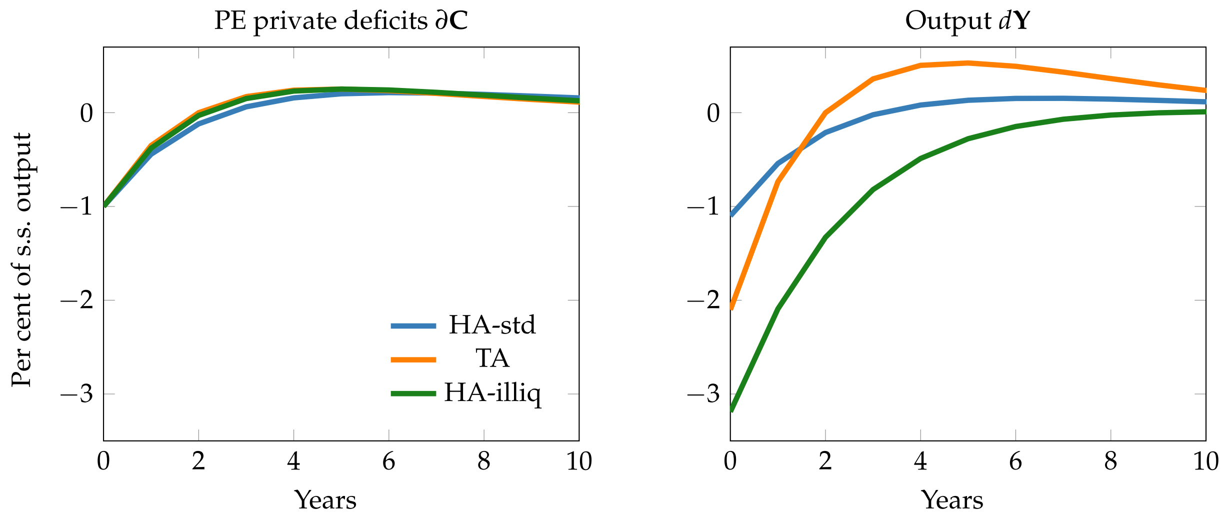

Note. To make responses comparable, ϵ is chosen to equalize the initial direct effect ∂C_0 across models. The persistence parameters are $\rho = 0.7$ and $\rho_{\underline{a}} = 0.7$.

This does not imply, however, that $d\mathbf{Y}$ is also similar across models—according to proposition 6, it is the interaction of $\partial\mathbf{C}$ and iMPCs that determines general equilibrium output.

Indeed, general equilibrium outcomes in figure 8 are quite different across models with different iMPCs. With the standard heterogeneous-agent model (HA-std), the output response essentially equals the direct effect, with no further amplification or persistence. The TA model does imply some amplification, but this amplification applies equally to the earlier negative and later positive terms of $\partial\mathbf{C}$—so that the net present value of $d\mathbf{Y}$ is zero, just as with the consumption effect of fiscal shocks. The HA-illiq model, uniquely, predicts an amplified *and* persistent response.

Figure 8 thus highlights that the same interaction between iMPCs and deficits discussed in section 4.2 also applies to entirely partial equilibrium *private* deficits—which, in this case, are just $\partial\mathbf{C}$. Shocks that reduce private deficits have large and persistent negative effects on output if our model matches the shape of iMPCs in the data, but not necessarily otherwise.

Example 2: Lump-sum financed government spending. Our second example is a government spending shock similar to those studied in section 4, with the exception that it is financed entirely using lump-sum taxes.[31] This is equivalent to our section 4 government spending shock financed by progressive taxes, if it is combined with an additional redistribution shock from low-productivity to high-productivity households in the periods of taxation. The effects of this redistribution shock are captured by our $\partial\mathbf{C}$ term in (33).

[31]This change in financing is only at the margin. To make sure that steady states are comparable, we retain our benchmark progressive fiscal rule otherwise. We could equivalently study this shock using the modified IKC from appendix B.1.

Figure 9: Comparing two ways to finance government spending: progressive vs. lump-sum taxation.

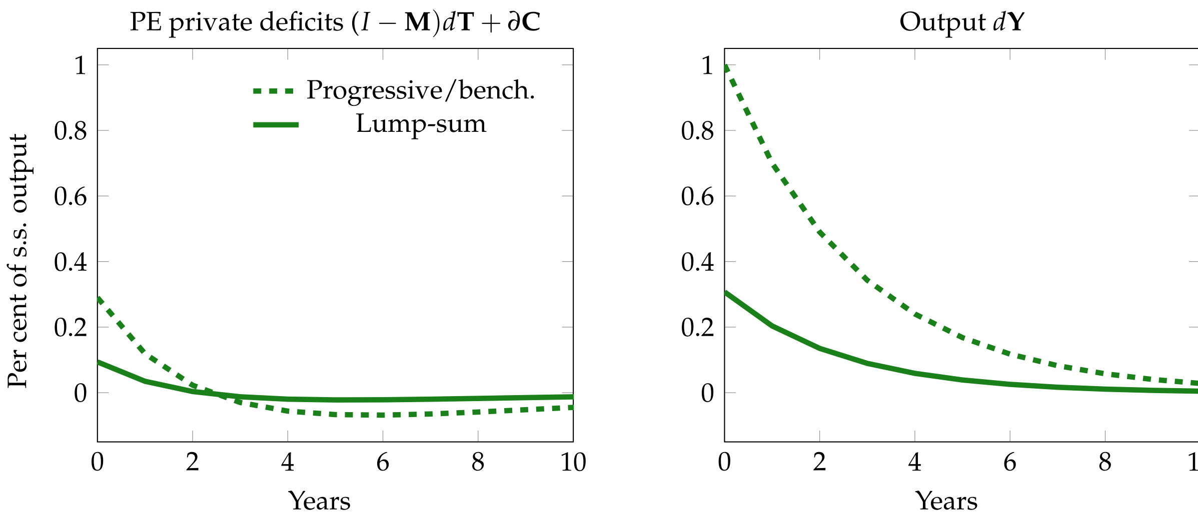

Figure 9 plots the partial equilibrium private deficit response $(I - \mathbf{M})d\mathbf{T} + \partial\mathbf{C}$ and the total output response $d\mathbf{Y}$ to a balanced-budget spending shock ($d\mathbf{G} = d\mathbf{T}$), for both types of taxation. The difference between the two paths for private deficits comes entirely from the redistribution term $\partial\mathbf{C}$. There are lower partial equilibrium private deficits on impact under lump-sum taxation, because this taxation targets many constrained households who have little ability to smooth the consumption effects of the tax by borrowing. By contrast, under our benchmark progressive tax, taxpayers are richer and relatively less constrained, and in response to taxes they partly offset the lack of public deficits by running substantial private deficits.

As figure 9 illustrates, these different implications for private deficits translate directly into different general equilibrium outcomes: as we move from large partial equilibrium private deficits under our benchmark taxation to small ones under lump-sum taxation, the impact multiplier falls from our benchmark of 1 to roughly 0.3. This decline is equal to the general equilibrium effect of the implied redistribution shock from poor to rich, for which $\partial\mathbf{C}$ is negative on impact—translating to an (amplified) negative response in general equilibrium, just as in our deleveraging example.

Overall, these results suggest an addendum to our earlier finding regarding the importance of deficits for fiscal shocks: it is not just public borrowing, but also private borrowing in response to taxes, that matters for general equilibrium amplification. Incidence of taxes is important—with the lump-sum case here being the most extreme example—but the role of distribution can be captured entirely by the induced effect on private deficits.

7 Conclusion

In this paper, we introduce a new set of moments, intertemporal MPCs. We argue that they are central to the general equilibrium transmission of shocks in models with heterogeneous agents and nominal rigidities, operating through a simple fixed-point equation, the intertemporal Keynesian cross. We provide estimates of intertemporal MPCs in the data and find that, within a set of commonly used models, only heterogeneous-agent models with multiple assets can match our estimates. Our key application is fiscal policy, where we generalize Haavelmo (1945)'s result of a unit multiplier for balanced-budget policies and argue that for deficit-financed policies, the shape of intertemporal MPCs is crucial. In particular, we find that, with a reasonable degree of deficit financing, a model matching intertemporal MPCs predicts impact and cumulative multipliers that lie above one, despite active monetary policy, distortionary taxation, and investment crowd-out.

Our paper provides a new approach to studying models with heterogeneity. Moving beyond the literature on sufficient statistics in partial equilibrium, we reduce the complexity of *general* equilibrium to a matrix of sufficient statistics, intertemporal MPCs, that can be disciplined empirically. This approach might be fruitfully extended to many other areas in macroeconomics, since the key insight—that agents interact in general equilibrium through a limited set of aggregates—applies to a wide variety of models.

References

Agarwal, Sumit and Wenlan Qian, "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, December 2014, *104* (12), 4205–4230.

Aiyagari, S.Rao, Lawrence J. Christiano, and Martin Eichenbaum, "The Output, Employment, and Interest Rate Effects of Government Consumption," *Journal of Monetary Economics*, October 1992, *30* (1), 73–86.

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Lindé, "Firm-Specific Capital, Nominal Rigidities and the Business Cycle," *Review of Economic Dynamics*, April 2011, *14* (2), 225–247.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *Manuscript*, May 2017.

_ **and Matthew Rognlie**, "A Note on Multipliers in NK Models with GHH Preferences," *Manuscript*, August 2017.

_ **and _** , "Inequality and Aggregate Demand," Working Paper 24280, National Bureau of Economic Research February 2018.

_ **, _ , and Ludwig Straub**, "A Note on Determinacy with Incomplete Markets," *Manuscript*, 2018.

Auerbach, Alan J and Yuriy Gorodnichenko, "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, May 2012, *4* (2), 1–27.

Baxter, Marianne and Robert G. King, "Fiscal Policy in General Equilibrium," *American Economic Review*, June 1993, *83* (3), 315–334.

ben Zeev, Nadav and Evi Pappa, "Chronicle of a War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks," *The Economic Journal*, August 2017, *127* (603), 1568–1597.

Bénabou, Roland, "Unequal Societies: Income Distribution and the Social Contract," *The American Economic Review*, 2000, *90* (1), 96–129.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, "House Prices and Consumer Spending," *The Review of Economic Studies*, July 2018, *85* (3), 1502–1542.

Bilbiie, Florin O. and Roland Straub, "Fiscal Policy, Business Cycles and Labor-Market Fluctuations," 2004.

_, Tommaso Monacelli, and Roberto Perotti, "Public Debt and Redistribution with Borrowing Constraints," *Economic Journal*, 2013, *123* (566), F64–F98.

Blinder, Alan S. and Robert M. Solow, "Does Fiscal Policy Matter?," *Journal of Public Economics*, January 1973, *2* (4), 319–337.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

Broer, Tobias, Niels-Jakob H Hansen, Per Krusell, and Erik Öberg, "The New Keynesian Transmission Mechanism: A Heterogenous-Agent Perspective," Working Paper 22418, National Bureau of Economic Research, July 2016.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, September 1983, *12* (3), 383–398.

Campbell, John Y. and N. Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," *NBER Macroeconomics Annual*, January 1989, *4*, 185–216.

Carroll, Christopher D., "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 2006, *91* (3), 312–320.

Chodorow-Reich, Gabriel, "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?," Working Paper 23577, National Bureau of Economic Research, July 2017.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo, "When Is the Government Spending Multiplier Large?," *Journal of Political Economy*, February 2011, *119* (1), 78–121.

Coenen, Günter, Christopher J Erceg, Charles Freedman, Davide Furceri, Michael Kumhof, René Lalonde, Douglas Laxton, Jesper Lindé, Annabelle Mourougane, Dirk Muir, Susanna Mursula, Carlos de Resende, John Roberts, Werner Roeger, Stephen Snudden, Mathias Trabandt, and Jan in't Veld, "Effects of Fiscal Stimulus in Structural Models," *American Economic Journal: Macroeconomics*, January 2012, *3* (4), 22–68.

Davig, Troy and Eric M. Leeper, "Monetary–Fiscal Policy Interactions and Fiscal Stimulus," *European Economic Review*, February 2011, *55* (2), 211–227.

Debortoli, Davide and Jordi Galí, "Monetary Policy with Heterogeneous Agents: Insights from TANK models," *Manuscript*, September 2017.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, October 2000, *46* (2), 281–313.

Fagereng, Andreas, Martib B. Holm, and Gisle J. Natvik, "MPC Heterogeneity and Household Balance Sheets," *Manuscript*, May 2018.

Farhi, Emmanuel and Iván Werning, "Chapter 31 - Fiscal Multipliers: Liquidity Traps and Currency Unions," in Harald Uhlig and John B. Taylor, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, 2016, pp. 2417–2492.

_ and _ , "Monetary Policy, Bounded Rationality, and Incomplete Markets," Working Paper 23281, National Bureau of Economic Research, March 2017.

Ferriere, Axelle and Gaston Navarro, "The Heterogeneous Effects of Government Spending: It's All About Taxes," *Manuscript*, November 2017.

Floden, Martin and Jesper Lindé, "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, April 2001, *4* (2), 406–437.

Fraumeni, Barbara, "The Measurement of Depreciation in the US National Income and Product Accounts," *Survey of Current Business*, 1997, *77*, 7–23.

Fuster, Andreas, Greg Kaplan, and Basit Zafar, "What Would You Do With $500? Spending Responses to Gains, Losses, News and Loans," Working Paper 24386, National Bureau of Economic Research March 2018.

Galí, Jordi, J. David López-Salido, and Javier Vallés, "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, 2007, *5* (1), 227–270.

Gelting, Jørgen, "Nogle Bemærkninger Om Finansieringen Af Offentlig Virksomhed," *Nationaløkonomisk Tidsskrift*, January 1941, *3.*

Giambattista, Eric and Steven Pennings, "When Is the Government Transfer Multiplier Large?," Technical Report, Manuscript, New York University April 2014.

Gilchrist, Simon and Charles P. Himmelberg, "Evidence on the Role of Cash Flow for Investment," *Journal of Monetary Economics*, December 1995, *36* (3), 541–572.

Guerrieri, Veronica and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *Quarterly Journal of Economics*, August 2017, *132* (3), 1427–1467.

Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson, "Housing Wealth Effects: The Long View," Working Paper 24729, National Bureau of Economic Research June 2018.

Haavelmo, Trygve, "Multiplier Effects of a Balanced Budget," *Econometrica*, 1945, *13* (4), 311–318.

Hagedorn, Marcus, "A Demand Theory of the Price Level," SSRN Scholarly Paper ID 2804542, Social Science Research Network, Rochester, NY June 2016.

_ , "Prices and Inflation when Government Bonds are Net Wealth," July 2018.

_ , Iourii Manovskii, and Kurt Mitman, "The Fiscal Multiplier," *Manuscript*, 2017.

Hall, Robert E., "By How Much Does GDP Rise If the Government Buys More Output?," *Brookings Papers on Economic Activity*, October 2009, *2009*, 183–231.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 2017, *132* (4), 1693–1754.

Jappelli, Tullio and Luigi Pistaferri, "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, October 2014, *6* (4), 107–136.

_ and _ , "Reported MPC and Unobserved Heterogeneity," Technical Report, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy 2018.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, December 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, July 2014, *82* (4), 1199–1239.

_ and _ , "Microeconomic Heterogeneity and Macroeconomic Shocks," *Journal of Economic Perspectives*, August 2018, *32* (3), 167–194.

_ , Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, March 2018, *108* (3), 697–743.

Koby, Yann and Christian Wolf, "Understanding Investment Stimulus: From Micro Data to Macro Outcomes," *Manuscript*, August 2018.

Krusell, Per and Anthony A. Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, *106* (5), 867–896.

Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers," *The Quarterly Journal of Economics*, 2018, *forthcoming.*

Lütticke, Ralph, "Transmission of Monetary Policy and Heterogeneity in Household Portfolios," *Manuscript*, June 2018.

Maggio, Marco Di, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging," *American Economic Review*, November 2017, *107* (11), 3550–3588.

McKay, Alisdair and Ricardo Reis, "The Role of Automatic Stabilizers in the U.S. Business Cycle," *Econometrica*, January 2016, *84* (1), 141–194.

_ , Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," *American Economic Review*, October 2016, *106* (10), 3133–3158.

Mehrotra, Neil, "Fiscal Policy Stabilization: Purchases or Transfers?," Technical Report, Manuscript, Brown University June 2014.

Michaillat, Pascal and Emmanuel Saez, "A New Keynesian Model with Wealth in the Utility Function," *Manuscript*, February 2018, p. 34.

Mountford, Andrew and Harald Uhlig, "What Are the Effects of Fiscal Policy Shocks?," *Journal of Applied Econometrics*, 2009, *24* (6), 960–992.

Nakamura, Emi and Jón Steinsson, "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, March 2014, *104* (3), 753–792.

Oh, Hyunseung and Ricardo Reis, "Targeted Transfers and the Fiscal Response to the Great Recession," *Journal of Monetary Economics*, December 2012, *59, Supplement*, S50–S64.

Ramey, Valerie A., "Can Government Purchases Stimulate the Economy?," *Journal of Economic Literature*, 2011, *49* (3), 673–685.

_ , "Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?," *Manuscript prepared for the NBER Global Financial Crisis @10*, July 2018.

_ **and Sarah Zubairy**, "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, March 2018, *126* (2), 850–901.

Romer, Christina D. and David H. Romer, "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, June 2010, *100* (3), 763–801.

Rotemberg, Julio J., "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies*, October 1982, *49* (4), 517–531.

Schmitt-Grohé, Stephanie and Martín Uribe, "Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model," *NBER Macroeconomics Annual*, January 2005, *20*, 383–425.

Schorfheide, Frank, "DSGE Model-Based Estimation of the New Keynesian Phillips Curve," *Federal Reserve Bank of Richmond Economic Quarterly*, 2008, *94* (4), 397–433.

Shapiro, Matthew D. and Joel Slemrod, "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys," in "Tax Policy and the Economy, Volume 17," MIT Press, 2003, pp. 83–110.

Uhlig, Harald, "Some Fiscal Calculus," *American Economic Review Papers and Proceedings*, May 2010, *100* (2), 30–34.

Werning, Iván, "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research, August 2015.

Woodford, Michael, "Simple Analytics of the Government Expenditure Multiplier," *American Economic Journal: Macroeconomics*, 2011, *3* (1), 1–35.

Appendix

A Proofs of propositions and additional results

A.1 General properties of M matrices and proof of proposition 1

We prove the following three results in this section.

Lemma 1. *A path for output $\{Y_t\}$ is part of an equilibrium if and only if*

$$Y_t = G_t + \mathcal{C}_t\left(\{Y_s - T_s\}\right) \tag{35}$$

for all periods t.

Proof. Suppose $\{Y_t\}$ is part of an equilibrium. Clearly, it satisfies the goods market clearing condition. Conversely, suppose that we are given a path for $\{Y_t\}$ that satisfies (35). Y_t immediately pins down Z_t and N_t. Given Z_t, we can find the individually optimal policy rules $\left\{c_t\left(\{a^j\},e\right), a_t^j\left(\{a^j\},e\right)\right\}$ and joint distributions over assets and productivity levels $\left\{\Psi_t\left(\{a^j\},e\right)\right\}$. Since $C_t = \int c_{it}di = \mathcal{C}_t\left(\{Y_s - T_s\}\right)$, goods market clearing holds and by Walras' law so does asset market clearing. The paths for nominal prices $\{P_t, W_t\}$ are determined by initial prices P_{-1}, W_{-1} and the Phillips curve (7). $\square$

Lemma 2. *The columns of* **M** *sum to 1 in present values,*

$$\sum_{t=0}^{\infty} \frac{M_{t,s}}{(1+r)^{t-s}} = 1 \tag{36}$$

Proof. Aggregating individuals' budget constraints (4) across agents i and over time t, we arrive at

$$\sum_{t=0}^{\infty} \frac{1}{(1+r)^t}\mathcal{C}_t\left(\{Z_s\}\right) = \sum_{t=0}^{\infty} \frac{1}{(1+r)^t} Z_t + (1+r)\sum_j a_{i,-1}^j$$

Taking derivatives with respect to Z_s, we arrive at

$$\sum_{t=0}^{\infty} \frac{1}{(1+r)^t} M_{t,s} = \frac{1}{(1+r)^s}$$

which is equivalent to (36). $\square$

Next, we prove proposition 1:

Proof of proposition 1. (22) follows directly from stacking (20). Note that **M** is non-negative and bounded, thus the matrix product with a summable vector such as $d\mathbf{T}$ is well-defined and itself summable. Existence of $d\mathbf{Y}$ implies that $\mathbf{M}d\mathbf{Y}$ must be well-defined, too. Let $\mathcal{M}$ be a linear map satisfying $\mathcal{M} = I + \mathbf{M}\mathcal{M}$. One can easily verify that $d\mathbf{G} - \mathbf{M}d\mathbf{T}$ is summable and $\mathcal{M}\left(d\mathbf{G} - \mathbf{M}d\mathbf{T}\right)$ solves (22). $\square$

Finally, we characterize the first column of $\mathbf{M}$ as a weighted average of individual iMPCs.

Lemma 3. *The first column of* $\mathbf{M}$ *can be written as*

$$M_{t,0} = \int \frac{z_{i0}}{\int z_{\iota 0} d\iota} \cdot \frac{\partial c_{it}}{\partial z_{i0}} di \tag{37}$$

Proof. To prove this result, we note that the consumption function $c_t(\{a^j\}, e_{it})$ is also the solution to the following sequential problem

$$\max \mathbb{E} \sum_{t\geq 0} \beta^t u(c_{it})$$

$$\begin{aligned} c_{it} + \sum_j a^j_{it} &= z_{it} + \Delta \cdot 1_{\{t=0\}} + (1+r) \sum_j a^j_{it-1} \\ a^j_{it} &\in \mathscr{A}^j_{e_{it}} \end{aligned}$$

where c_{it} is measurable with respect to time t information on the idiosyncratic process for ability, with $\Delta = 0$. More generally, the solution can be denoted by $c_t(\{a^j\}, e_{it}, \Delta)$ and we define *individual iMPCs* as

$$\frac{\partial c_{it}}{\partial z_{i0}} \equiv \frac{\partial c_t}{\partial \Delta}(\{a^j\}, e_{it}, \Delta)|_{\Delta=0}$$

Observe that due to (18), a change in Z_0 translates into changes in z_{i0} according to

$$z_{i0} = \frac{e_{i0}^{1-\lambda}}{\int e_{\iota 0}^{1-\lambda} d\iota} Z_0$$

and therefore

$$\frac{\partial c_{it}}{\partial Z_0} = \frac{e_{i0}^{1-\lambda}}{\int e_{\iota 0}^{1-\lambda} d\iota} \frac{\partial c_{it}}{\partial z_{i0}} = \frac{z_{i0}}{\int z_{\iota 0} d\iota} \frac{\partial c_{it}}{\partial z_{i0}}$$

which immediately implies (37) since $M_{t,0} = \frac{\partial}{\partial Z_0} \int c_{it} di$. □

A.2 Proofs of propositions 2 and 3

Proof of proposition 2. If $d\mathbf{G} = d\mathbf{T}$, then $d\mathbf{Y} = d\mathbf{G}$ solves (22) since with that guess

$$d\mathbf{G} - \mathbf{M} d\mathbf{T} + \mathbf{M} d\mathbf{Y} = d\mathbf{G} = d\mathbf{Y}$$

□

Proof of proposition 3. Rewriting (22) as

$$d\mathbf{Y} - d\mathbf{G} = \mathbf{M}(d\mathbf{G} - d\mathbf{T}) + \mathbf{M}(d\mathbf{Y} - d\mathbf{G})$$

and applying proposition 1 implies that the solution for $d\mathbf{Y} - d\mathbf{G}$ is given by

$$d\mathbf{Y} - d\mathbf{G} = \mathcal{M}\mathbf{M}\left(d\mathbf{G} - d\mathbf{T}\right)$$

which is equivalent to (26). $\square$

A.3 Results for the RA model and proof of proposition 4

The representative-agent (RA) model has a particularly simple consumption function,

$$\mathcal{C}_t(\{Z_s\}) = \frac{r}{1+r}\sum_{t=0}^{\infty}\frac{1}{(1+r)^t}Z_t + ra_{-1}$$

Linearizing around a steady state with $\beta = 1/(1+r)$, this leads to the following matrix of iMPCs

$$\mathbf{M} = \begin{pmatrix} 1-\beta & (1-\beta)\beta & (1-\beta)\beta^2 & \cdots \\ 1-\beta & (1-\beta)\beta & (1-\beta)\beta^2 & \cdots \\ 1-\beta & (1-\beta)\beta & (1-\beta)\beta^2 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{38}$$

Each column is constant in the RA model since the agent is a permanent-income consumer and thus consumes the constant annuity value of any increase in after-tax income.

Now consider a fiscal policy experiment $(d\mathbf{G}, d\mathbf{T})$. In line with proposition 1, the effect on output is described by the intertemporal Keynesian cross,

$$d\mathbf{Y} = d\mathbf{G} - \mathbf{M}d\mathbf{T} + \mathbf{M}d\mathbf{Y} \tag{39}$$

In this special case, the space of all solutions $d\mathbf{Y}$ that satisfy this equation is simply given by

$$d\mathbf{Y} = d\mathbf{G} + \eta d\mathbf{V} \tag{40}$$

for any real number η. Here, $d\mathbf{V} = (1, 1, \ldots)'$ is a constant vector. This is easily verified since, first, $d\mathbf{Y} = d\mathbf{G}$ is a particular solution to (39) (using (38) and the fact that $\mathbf{M}d\mathbf{T} = \mathbf{M}d\mathbf{G}$ due to the government's intertemporal budget constraint). Second, $d\mathbf{V} = \mathbf{M}d\mathbf{V}$, and indeed $d\mathbf{V}$ is the unique unit eigenvector of $\mathbf{M}$ by the Perron-Frobenius theorem.

Observe that our RA model is equivalent to a textbook New Keynesian model with a government sector and constant-r monetary policy. Thus, the degree of indeterminacy captured by the constant vector $d\mathbf{V}$ is the *exact* same as the one that is present in the Euler equation, according to which constant shifts in consumption (or the output gap) leave the equation satisfied in all periods.

Among all solutions satisfying (40), there is a unique one that ensures that $\lim_{t\to\infty} dY_t = 0$, namely $\eta = 0$ given our assumption that $\lim_{t\to\infty} dG_t = 0$. This proves proposition 4.

We can also express the solution using the linear map $\mathcal{M}$. Defining the map

$$\mathcal{M}d\mathbf{X} \equiv d\mathbf{X} - \lim_{t\to\infty} dX_t$$

it is easy to see that the unique solution to (39) with $\lim_{t\to\infty} dY_t = 0$ is given by

$$d\mathbf{Y} = \mathcal{M}\left(d\mathbf{G} - \mathbf{M}d\mathbf{T}\right) = d\mathbf{G}$$

A.4 Results for the TA model and proof of proposition 5

The TA model combines a mass $1-\mu$ of permanent-income agents with a mass μ of hand-to-mouth agents. The consumption function is thus a convex combination of both agents' consumption functions

$$\mathcal{C}_t(\{Z_s\}) = (1-\mu)\left\{\frac{r}{1+r}\sum_{t=0}^{\infty}\frac{1}{(1+r)^t}Z_t + ra_{-1}\right\} + \mu Z_t$$

The iMPC matrix $\mathbf{M}$ similarly is a convex combination of $\mathbf{M}^{RA}$, which from now on denotes the permanent-income iMPC matrix in (38), and the identity matrix, which captures the consumption response of the hand-to-mouth agents. Thus

$$\mathbf{M} = (1-\mu)\mathbf{M}^{RA} + \mu\mathbf{I} \tag{41}$$

Consider again a fiscal policy experiment $(d\mathbf{G}, d\mathbf{T})$. With the iMPC matrix given in (41), the intertemporal Keynesian cross (39) can be rewritten as

$$d\mathbf{Y} = \frac{1}{1-\mu}d\mathbf{G} - \mathbf{M}^{RA}d\mathbf{T} - \frac{\mu}{1-\mu}d\mathbf{T} + \mathbf{M}^{RA}d\mathbf{Y}$$

This is identical to equation (39) in the previous section, except that $d\mathbf{G}$ is now replaced by $\frac{1}{1-\mu}d\mathbf{G} - \frac{\mu}{1-\mu}d\mathbf{T} = d\mathbf{G} + \frac{\mu}{1-\mu}\left(d\mathbf{G} - d\mathbf{T}\right)$. Hence, the unique solution satisfying $\lim_{t\to\infty} dY_t = 0$ is

$$d\mathbf{Y} = d\mathbf{G} + \frac{\mu}{1-\mu}\left(d\mathbf{G} - d\mathbf{T}\right)$$

This proves proposition 5.

A.5 Definition and calibration of the BU model

The BU model is a version of the RA model in which the representative agent has utility over bonds. The agent's maximization problem is given by

$$\begin{aligned} \max \quad & \sum \beta^t \left\{u\left(c_t\right) - v(n_t) + \nu\left(a_t\right)\right\} \\ \text{s.t.} \quad & c_t + a_t = Z_t + \left(1 + r_{t-1}\right)a_{t-1} \end{aligned} \tag{42}$$

We parametrize the utility over bonds as $\nu(a) = \frac{\varphi}{1-\gamma} a^{1-\gamma}$. We calibrate φ to match the wealth to after-tax income ratio. We calibrate γ to match the impact MPC. Conditional on these targets, the choice of β turns out not to be very important. We choose it to lie between the parameters for the HA-std and the HA-illiq models. This yields: $\beta = 0.90$, $\varphi = 0.22^{-\gamma}$, and $\gamma = 220$. (Observe that this γ is rather extreme, suggesting that the BU model struggles to match our impact MPC.)

A.6 The generalized IKC and the proof of proposition 6

The derivation of the generalized IKC and the proof of proposition 6 follow in exactly the same way as the derivation and proof in appendix A.1, except that we carry the additional term $\partial \mathbf{C}$.

B Generalizations of the intertemporal Keynesian cross

B.1 General tax incidence

To allow for fiscal policy experiments which are financed using taxes with a different incidence, we introduce a new tax $\mathcal{T}_t$, of which a fraction $\mathcal{T}_t(e)$ is paid for by agents with ability e in period t, $\mathcal{T}_t = \sum_e \pi(e)\mathcal{T}_t(e)$. This changes the agents' budget constraints (4) to

$$c_{it} + \sum_j a_{it}^j \quad = \quad z_{it} - \mathcal{T}_t(e_{it}) + (1 + r_{t-1}) \sum_j a_{it-1}^j \tag{43}$$

and the government's intertemporal budget constraint (10) to

$$(1 + r_{-1})B_{-1} + \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t-1} \frac{1}{1 + r_s} \right) G_t = \sum_{t=0}^{\infty} \left(\prod_{s=0}^{t-1} \frac{1}{1 + r_s} \right) (T_t + \mathcal{T}_t)$$

Given (43), the consumption function now depends explicitly on the path of $\{\mathcal{T}_s\}$

$$C_t = \mathcal{C}_t \left(\{Z_s; \mathcal{T}_s\} \right)$$

When a given fiscal policy is financed using the arbitrary new tax instrument $\mathcal{T}_s$, output is determined by

$$dY_t = dG_t + \sum_{s=0}^{\infty} M_{t,s} dY_s - \sum_{s=0}^{\infty} M_{t,s}^T d\mathcal{T}_s$$

where $M_{t,s}^T \equiv \partial \mathcal{C}_t / \partial \mathcal{T}_s$. This equation is precisely of the form in (23).[32]

[32]Various other fiscal rules, not used in this paper, endogenize tax revenue in each period as a function of the path of output. To capture these rules in our framework, we must compose the matrix $\mathbf{M}^T$ with the matrix that (locally) maps the path of output to the path of endogenous tax revenue.

B.2 Investment

We first derive the existence of an investment function $\mathcal{I}_t(\{Y_s\})$ and then discuss the modified household consumption function $\mathcal{C}_t(\{Y_s; T_s\})$.

Investment function. We introduce a standard supply side with investment in appendix C.2 below. Here, we focus on a special case in which prices are flexible, $\kappa^p = \infty$ and there are no markups $\mu = 1$. In that case, the economy's capital stock is determined as solution to the following fixed point. Given a path for real wages $\{w_t\}$, the economy's capital stock K_t and labor supply N_t solve

$$J_t(K_{t-1}) = \max_{K_t, N_t} \left\{ F(K_{t-1}, N_t) - w_t N_t - \zeta\left(\frac{K_t}{K_{t-1}}\right) K_{t-1} + \frac{1}{1+r} J_{t+1}(K_t) \right\} \tag{44}$$

Equilibrium real wages are then pinned down to insure that

$$F(K_{t-1}, N_t) = Y_t \tag{45}$$

where $\{Y_t\}$ is a given path of output. Given K_{t-1}, denote by $\mathcal{N}(K_{t-1}, Y_t) = Y_t^{1/(1-\alpha)} K_{t-1}^{-\alpha/(1-\alpha)}$ the solution to (45). We assume that there exists a unique equilibrium path of capital stocks $\{K_{t-1}\}$.

We characterize this fixed point as follows.

Lemma 4. *$\{K_t\}$ is the equilibrium path of capital if and only if $\{K_t\}$ solves the following problem*

$$\mathcal{J}_t(K_{t-1}) = \max_{K_t, N_t} \left\{ \mathcal{F}(K_{t-1}, Y_t) - \zeta\left(\frac{K_t}{K_{t-1}}\right) K_{t-1} + \frac{1}{1+r} \mathcal{J}_{t+1}(K_t) \right\} \tag{46}$$

where we defined

$$\mathcal{F}(K, Y) \equiv \alpha Y \log K$$

Proof. The only difference between (44) and (46) are the terms $F(K_{t-1}, N_t) - w_t N_t$ and $\mathcal{F}(K_{t-1}, Y_t)$. The equilibrium marginal derivative of the former is given by

$$F_K(K_{t-1}, \mathcal{N}(K_{t-1}, Y_t)) = \alpha Y_t / K_{t-1}$$

which is precisely equal to $\mathcal{F}_K(K_{t-1}, Y_t)$. Thus, any equilibrium path satisfies the first order conditions of (46), and any solution to (46) satisfies the first order conditions of (44). Since there is a unique equilibrium path, this also implies that there can only be a single solution to (46). □

Lemma 4 is helpful since it immediately implies that there is a direct mapping from $\{Y_s\}$ to the equilibrium path of capital, and therefore also an investment function $\mathcal{I}_t(\{Y_s\})$.

Modified consumption function. Now, in response to a shock, there is a revaluation of capital at date 0. Household asset positions coming into period 0 will change depending on each household's holdings of capital, which are indeterminate since there is no aggregate uncertainty and capital and bonds have the same returns in steady state. We will generally resolve the indeterminacy by assuming that capital holdings are proportional to each household's overall assets (or the household's illiquid assets, in the HA-illiq model, where capital is held only in the illiquid account).

When the capital is revalued, it is worth $J_0(K_{-1})$. From (44) and the fact that both K_t and N_t can be written as functions of $\{Y_s\}$, we can rewrite this (with some abuse of notation) as a function of $\{Y_s\}$, and then the aggregate consumption function as

$$C_t = \mathcal{C}_t\left(\{Z_s\}, J(\{Y_s\})\right).$$

Thus, the consumption function needs to be slightly modified to a more general function of $\{Y_s\}$ and $\{T_s\}$, which we abbreviate as $\mathcal{C}_t\left(\{Y_s; T_s\}\right)$.

The intertemporal Keynesian cross. Having derived the investment function as well as the modified consumption function, the goods market clearing condition now reads

$$Y_t = G_t + \mathcal{C}_t\left(\{Y_s; T_s\}\right) + \mathcal{I}_t\left(\{Y_s\}\right)$$

based on which the derivation of (23) with $M^T_{t,s} = \frac{\partial \mathcal{C}_t}{\partial T_s}$ and $M^Y_{t,s} = \frac{\partial \mathcal{C}_t}{\partial Y_s} + \frac{\partial \mathcal{I}_t}{\partial Y_s}$ is straightforward.

B.3 Sticky prices

To explore the role of sticky prices, suppose the union sets wages perfectly flexibly, that is $\kappa^w = \infty$ in (7), and therefore

$$\int N_t \left\{ v'\left(n_{it}\right) - \frac{\epsilon - 1}{\epsilon} \frac{\partial z_{it}}{\partial n_{it}} u'\left(c_{it}\right) \right\} di = 0$$

but prices follow a standard Phillips curve such as

$$\pi_t = \kappa^p w_t + \frac{1}{1+r}\pi_{t+1}$$

Faced with demand Y_t, firms hire labor $N_t = Y_t$ and earn profits

$$\Pi_t = Y_t - w_t N_t$$

Suppose that profits are distributed according to a rule: an agent with idiosyncratic ability e receives share $\chi(e)$ of profits, so that agent i's date-t pretax income is now given by[33]

$$y_{it} = w_t N_t e_{it} + \chi(e_{it})\Pi_t$$

[33]The other common way to attribute profits to households is by allowing households to trade firms' shares. This is our assumption in the quantitative model of section 5.

Rewriting this, we see that

$$y_{it} = Y_t \left(w_t e_{it} + (1 - w_t) \chi(e_{it}) \right)$$

In the sticky-wage case, the real wage is always equal to 1 and aggregate income Y_t is entirely split according to ability e_{it}, i.e. $y_{it} = Y_t e_{it}$. In this model, the real wage fluctuates, affecting the way in which income is distributed. In booms, when the real wage is large, income is split more according to ability e_{it}, while in busts the distribution of profits $\chi(e_{it})$ matters more.

A natural benchmark case is the one where profits are distributed according to ability

$$\chi(e) = e \tag{47}$$

In that case, pretax income is exactly the same as it was in the sticky wage model

$$y_{it} = Y_t e_{it}$$

Thus, for the constant-r case, one may reinterpret our sticky wage model as a model with sticky prices together with the distribution rule for profits in (47).

C Model derivations and computation

C.1 Wage Phillips curve

In this section we derive the nonlinear wage Phillips curve. At any time t, union k sets its wage W_{kt} to maximize, on behalf of all the workers it employs,

$$\sum_{\tau \geq 0} \beta^{t+\tau} \left(\int \left\{ u\left(c_{it+\tau}\right) - v\left(n_{it+\tau}\right) \right\} d\Psi_{it+\tau} - \frac{\psi}{2} \left(\frac{W_{k,t+\tau}}{W_{k,t+\tau-1}} - 1 \right)^2 \right)$$

taking as given the initial distribution of households over idiosyncratic states Ψ_{it} as well as the demand curve for tasks emanating from the labor packers, which is

$$N_{kt} = \left(\frac{W_{kt}}{W_t} \right)^{-\epsilon} N_t$$

where $W_t = \left(\int W_{kt}^{1-\epsilon} dk \right)^{\frac{1}{1-\epsilon}}$ is the price index for aggregate employment services.

Each union is infinitesimal and therefore only takes into account its marginal effect on every household's consumption and labor supply. By (3), household total real earnings are

$$\begin{aligned} z_{it} &= \tau_t \left(\frac{W_t}{P_t} e_{it} n_{it} \right)^{1-\lambda} \\ &= \tau_t \left(\frac{1}{P_t} \int_0^1 W_{kt} e_{it} n_{ikt} dk \right)^{1-\lambda} \\ &= \tau_t \left(\frac{1}{P_t} \int_0^1 W_{kt} e_{it} \left(\frac{W_{kt}}{W_t} \right)^{-\epsilon} N_t dk \right)^{1-\lambda} \end{aligned}$$

The envelope theorem implies that we can evaluate indirect utility as if all income from the union wage change is consumed. In that case $\frac{\partial c_{it}}{\partial W_{kt}} = \frac{\partial z_{it}}{\partial W_{kt}}$, where

$$\begin{aligned} \frac{\partial z_{it}}{\partial W_{kt}} &= (1-\lambda)\, \tau_t \left(\frac{W_t}{P_t} e_{it} n_{it} \right)^{-\lambda} \frac{e_{it}}{P_t} \left\{ N_{kt} - W_{kt} \epsilon \left(\frac{1}{W_t} \right)^{-\epsilon} N_t W_{kt}^{-\epsilon-1} \right\} \\ &= (1 - MTR_{it}) \frac{e_{it}}{P_t} N_{kt} (1-\epsilon) \end{aligned}$$

where $MTR_{it} \equiv 1 - (1-\lambda)\, \tau_t \left(\frac{W_t}{P_t} e_{it} n_{it} \right)^{-\lambda}$ is household i's marginal tax rate at time t. On the other hand, household i's total hours worked are

$$n_{it} \equiv \int_0^1 \left(\frac{W_{kt}}{W_t} \right)^{-\epsilon} N_t dk$$

which falls when W_{kt} increases according to

$$\frac{\partial n_{it}}{\partial W_{kt}} = -\epsilon \frac{N_{kt}}{W_{kt}}$$

The first-order condition of the union with respect to W_{kt} is therefore

$$\begin{aligned} &\int N_{kt} \left\{ (1-\epsilon) \frac{e_{it}}{P_t} u'(c_{it}) (1 - MTR_{it}) + \frac{\epsilon}{W_{kt}} v'(n_{it}) \right\} d\Psi_{it} \\ &-\psi \left(\frac{W_{k,t}}{W_{k,t-1}} - 1 \right) \frac{1}{W_{k,t-1}} + \beta\psi \left(\frac{W_{k,t+1}}{W_{k,t}} - 1 \right) \left(\frac{W_{k,t+1}}{W_{k,t}} \right) \frac{1}{W_{k,t}} = 0 \end{aligned} \quad (48)$$

In equilibrium all unions set the same wage, so $W_{kt} = W_t$ and $N_{kt} = N_t$. Define wage inflation $\pi^w \equiv \frac{W_t}{W_{t-1}} - 1$. After multiplying (48) by W_t, and noting that

$$\frac{\partial z_{it}}{\partial n_{it}} = (1 - MTR_{it})\, e_{it} \frac{W_t}{P_t}$$

we obtain the aggregate wage Phillips curve

$$\pi_t^w(1+\pi_t^w) = \frac{\epsilon}{\psi}\int N_t\left\{v'\left(n_{it}\right) - \frac{\epsilon-1}{\epsilon}\frac{\partial z_{it}}{\partial n_{it}}u'\left(c_{it}\right)\right\}d\Psi_{it} + \beta\pi_{t+1}^w(1+\pi_{t+1}^w)$$

which is the formulation in (7).

Note that in equilibrium, enforcing $n_{it} = N_{kt} = N_t$, we simply have

$$\frac{\partial z_{it}}{\partial n_{it}} = (1-\lambda)\,\tau_t e_{it}^{1-\lambda}\left(\frac{W_t}{P_t}\right)^{1-\lambda}N_t^{-\lambda} = (1-\lambda)\,\frac{e_{it}^{1-\lambda}}{\int e_{it}^{1-\lambda}di}\frac{Z_t}{N_t}$$

where Z_t is average aftertax income. Hence equation (7) can be written in terms of aggregates π_t^w, Z_t, N_t together with a 'virtual aggregate consumption' term C_t^*

$$\pi_t^w(1+\pi_t^w) = \frac{\epsilon}{\psi}\left\{N_t v'\left(N_t\right) - \frac{\epsilon-1}{\epsilon}\left(1-\lambda\right)Z_t u'\left(C_t^*\right)\right\} + \beta\pi_{t+1}^w(1+\pi_{t+1}^w)$$

where we have defined C_t^* such

$$u'\left(C_t^*\right) = \int_i \frac{e_{it}^{1-\lambda}u'\left(c_{it}\right)}{\int e_{it}^{1-\lambda}di}d\Psi_{it}$$

Distribution matters for inflation dynamics only through its effects on the dynamics of C_t^*. Linearizing this expression around the zero inflation steady state yields a standard wage Phillips Curve

$$\pi_t^w = \kappa^w\left\{\frac{1}{\phi}\frac{dN_t}{N} + \frac{1}{\nu}\frac{dC_t^*}{C^*} - \left(\frac{dZ_t}{Z^*} - \frac{dN_t}{N}\right)\right\} + \beta\pi_{t+1}^w$$

where $\kappa^w = \frac{\epsilon}{\psi N v'(N)}$, ϕ is the Frisch elasticity of labor supply, and ν the elasticity of intertemporal substitution in consumption. The term $\frac{dZ_t}{Z^*} - \frac{dN_t}{N}$ captures the distortionary effects of taxation.

C.2 Firm problem and all FOCs

We rewrite the firm problem (27) as

$$J_t(k_{t-1}) = \max_{p_t,k_t,n_t}\left\{\frac{p_t}{P_t}F(k_{t-1},n_t) - \frac{W_t}{P_t}n_t - \zeta\left(1-\delta+\frac{\iota_t}{k_{t-1}}\right)k_{t-1} - \xi(p_t,p_{t-1})Y_t + \frac{1}{1+r_t}J_{t+1}(k_t)\right\}$$

subject to

$$\left(\frac{F(K_{t-1},L_t)}{Y_t}\right)^{\frac{1}{\mu}-1}Y_t = \frac{p_t}{P_t}Y_t$$

$$k_t = (1-\delta)k_{t-1} + \iota_t$$

Let η_t be the multiplier on the first constraint and Q_t be the multiplier on the second. Define

$$mc_t \equiv 1 + \eta_t \left(\frac{1}{\mu} - 1 \right) \tag{49}$$

The FOC for price p_t is then given by

$$(1 - \eta_t) \frac{1}{P_t} Y_t = \frac{1}{\kappa^p (\mu - 1)} Y_t \left(\frac{p_t - p_{t-1}}{p_{t-1}} \right) \frac{1}{p_{t-1}} + \frac{1}{1 + r_t} \frac{1}{\kappa^p (\mu - 1)} \left(\frac{p_{t+1} - p_t}{p_t} \right) \left(-\frac{p_{t+1}}{p_t^2} \right) Y_{t+1}$$

Noting that all firms are symmetric and thus $p_t = P_t$, defining price inflation as $\pi_t^p \equiv (P_t - P_{t-1})/P_{t-1}$ and using (49), this becomes

$$\pi_t^p (1 + \pi_t^p) = \kappa^p (\mu \cdot mc_t - 1) + \frac{1}{1 + r_t} \frac{Y_{t+1}}{Y_t} \pi_{t+1}^p (1 + \pi_{t+1}^p)$$

proving (29).

Optimality of investment requires in equilibrium

$$\frac{Q_{t+1}}{1 + r_t} = \zeta' \left(\frac{K_t}{K_{t-1}} \right)$$

where the law of motion of Q_t is given by

$$Q_t = mc_t F_{K,t} - \zeta \left(\frac{K_t}{K_{t-1}} \right) + \frac{K_t}{K_{t-1}} \frac{Q_{t+1}}{1 + r_t}$$

Optimality of labor requires

$$mc_t F_{n,t} = \frac{W_t}{P_t}$$

so in other words, mc_t is the date-t real marginal cost.

D Background on empirical evidence

D.1 Evidence from Norwegian administrative data

Our estimates on Norwegian iMPCs were generously provided to us by Andreas Fagereng, Martin Holm and Gisle Natvik. In Fagereng et al. (2018), they combine individual-level administrative income data and household-level wealth data from the Norwegian population, and residually impute a household-level consumption measure using a budget constraint approach. This is therefore a comprehensive measure of household expenditure, including durable and housing expenditures. However, the authors drop from their sample all households who record a housing market transaction, so that their iMPC estimates can be interpreted as including consumption and non-housing durable expenditures only.

The paper provides convincing evidence that the sample of gamblers is not selected: 70 percent of the population gambles, the population of winners is not significantly different from the rest of the population on observable characteristics including their consumption-income covariance over time, and gambling prizes are not predictable by prior household characteristics (Tables 1 and 2). To further limit the concern that iMPC estimates reflect the behavior of serial gamblers, the sample is limited to households who win only once.

The authors provided us with income-weighted estimates of regression (25). The regression includes all lottery wins below $150,000, and most prizes are below $20,000. As we discuss in footnote 16, their MPC estimates for a sample restricted to small gains are much larger than the full sample estimates, imprecisely estimated, and do not sum to one, so we prefer to use these full-sample estimates.

D.2 Evidence from the Italian Survey of Household Income and Wealth

The Italian Survey of Household Income and Wealth (SHIW) is a biannual survey, publicly available on the Bank of Italy website. In 2016, survey respondents were asked:

> "Imagine you unexpectedly receive a refund equal to the household's monthly income. How much of the sum would you save and how much would you spend? Indicate the percentage saved and the percentage spent."

In 2010, the same question was asked, except that the survey mentioned a "reimbursement" rather than a "refund". Given that answers are similar to those in the the 2012 survey which specified a timeframe "over the next 12 months" but had a slightly different wording, these answers are typically interpreted as annual MPCs. (The 2014 survey instead included a retrospective question about spending of the 2014 "Renzi bonus".)

We drop observations with zero or negative income, and are left with 7936 observations in the 2010 survey and 7367 observations in the 2016 survey.

Distributions of MPCs. Figure 10 displays the cumulative density functions of the distribution of MPCs in the 2010 and 2016 SHIW. As is apparent, these distributions are extremely similar. The largest distance between the two CDFs is 0.045, even though these distributions are measured six years apart. This justifies our assumption below that the distribution is not changing from one year to the next.

Construction of a lower bound. From the 2016 survey, we have a distribution of self-reported MPCs MPC_i as well as income net of taxes z_i. We can therefore construct $M_{0,0} = \mathbb{E}_I \left[\frac{z_i}{\mathbb{E}_I[z_i]} MPC_i\right]$ directly from the data.

Next, consider aggregate cumulative savings after T periods, equal to

$$A_T = (1+r)^T \, \mathbb{E}_I \left[\frac{z_i}{\mathbb{E}_I\left[z_i\right]} MPS_{0i} \cdot MPS_{1i} \cdots MPS_{Ti}\right] \tag{50}$$

Figure 10: MPC distribution in the 2010 and 2016 SHIW.

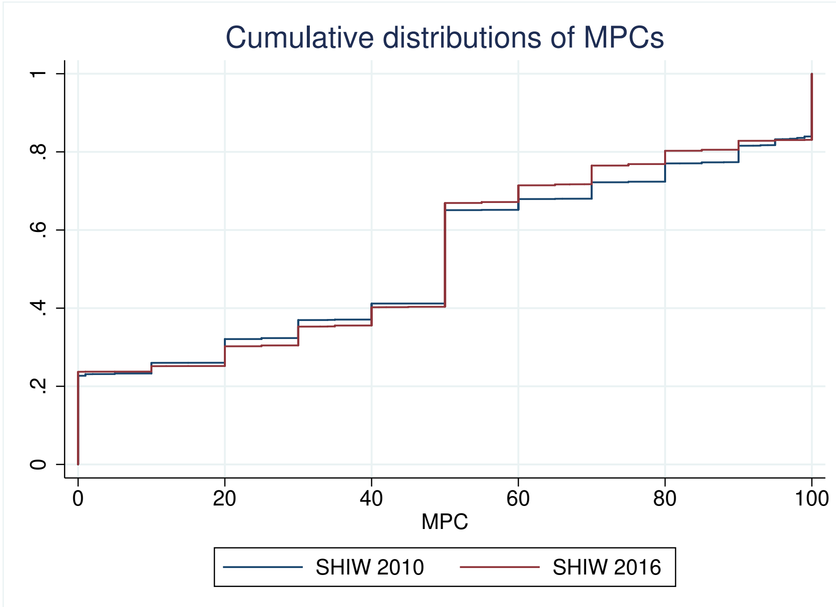

where MPS_{ti} is the marginal propensity to save of individual i at time t. Given our stationarity assumption, the distributions MPS_t are the same as the distribution MPS_0. The rearrangement inequality ensures that (50) is highest when the distributions have perfect correlation, so that individuals maintain the same MPS from year to year. This gives us an upper bound for cumulative saving,

$$\overline{A_T} = (1+r)^T \mathbb{E}_I \left[\frac{z_i}{\mathbb{E}_I [z_i]} (MPS_i)^T \right]$$

and therefore a lower bound for cumulative spending $\underline{C_T} = (1+r)^T - \overline{A_T}$.

Figure 1 reports the differences $\underline{C_T} - \underline{C_{T-1}}$, computed under our benchmark calibration for $r = 5\%$. For date $T = 1$ this is an exact lower bound for spending since $\underline{C_0} = M_{0,0}$. Intuitively, the worst case scenario for spending is the situation in which all individuals who saved in period 0 and therefore have the most remaining to spend still save in period 1, and this can simply be computed as $(1+r)\,\mathbb{E}\left[\frac{z_i}{\mathbb{E}_I[z_i]} (1 - MPC_i)\, MPC_i\right]$. For $T > 1$, this difference in lower bounds for savings is also a lower bound for spending in that period, unless a previous lower bound for cumulative spending is exceeded.

Note that, using the panel component of the SHIW, in principle we can refine this lower bound by seeing the extent to which individual MPCs change from year to year. Given that our lower bound is sufficient to reject most standard models already, we do not pursue this here.

E Durable goods

In this section, we amend our benchmark framework to include durable spending. We then show that the model generates an intertemporal Keynesian cross provided the consumption function now includes both nondurable and durable expenditure, as claimed in section 2.5. We explain how we match the $M_{t,0}$ columns to the Norwegian data, and discuss how other elements of the $\mathbf{M}$ matrix implied by this new calibration differ from that of our model without durables.

Model with durable goods. We introduce durables in the simplest possible way, by assuming homothetic durable demand and perfect collateralizability. Assuming that households can only trade in one other asset ($J = 1$) for ease of notation, and anticipating a constant-r monetary policy rule (14), the household problem is now

$$\begin{aligned}
\max \quad & \mathbb{E}\left[\sum_{t\geq 0} \beta^t \left\{u\left(c_{it}\right) + \kappa u\left(d_{it}\right)\right\}\right] \\
& c_{it} + d_{it} - \left(1-\delta_D\right) d_{it-1} + a_{it} = z_{it} + \left(1+r\right) a_{it-1} \\
& a_{it} + \frac{1-\delta_D}{1+r_t} d_{it} \in \mathcal{A}_{e_{it}}
\end{aligned}$$

where z_{it} is still taken as given and determined by labor demand in general equilibrium.

Observe that households can borrow against the undepreciated component of the next period durable stock. In particular, in this interpretation the TANK model is one in which constrained agents (for which $\mathcal{A}_{e_c} = \{0\}$) are perpetually up against a their collateral constraint. Redefining the overall asset position as

$$w_{it} \equiv a_{it} + \frac{1-\delta_D}{1+r} d_{it}$$

the problem rewrites as

$$\begin{aligned}
\max \quad & \mathbb{E}\left[\sum_{t\geq 0} \beta^t \left\{u\left(c_{it}\right) + \kappa u\left(d_{it}\right)\right\}\right] \\
& c_{it} + \frac{r+\delta_D}{1+r} d_{it} + w_{it} = z_{it} + \left(1+r\right) w_{it-1} \\
& w_{it} \in \mathcal{A}_{e_{it}}
\end{aligned}$$

where the user cost of durables $\frac{r+\delta_D}{1+r}$ appears. In this formulation, no matter whether the constraint on w_{it} is binding or not, there is a unique first order condition for the stock of durables d_{it} relative to consumption c_{it} that applies to every consumer, namely

$$\kappa u'\left(d_{it}\right) = u'\left(c_{it}\right)\left(\frac{r+\delta_D}{1+r}\right) \tag{51}$$

Equation (51) implies that the durable stock is a constant fraction of nondurable consumption at all times and for every consumer: $d_{it} = \upsilon c_{it}$ where $\upsilon = \left(u'\right)^{-1}\left(\frac{r+\delta_D}{1+r}\frac{1}{\kappa}\right)$. Further, given an initial level of wealth w_{-1} and a stochastic process for z_{it}, if we let c_{it}^{ND} be the path for nondurable consumption generated by our main model without durables, then the path for nondurable consumption in the model with durables is given, in every state and date, by $c_{it} = \frac{c_{it}^{ND}}{1+\frac{r+\delta}{1+r}\upsilon}$.Total expenditures

Figure 11: Columns of the iMPC matrix in the HA-illiq model with and without durables.

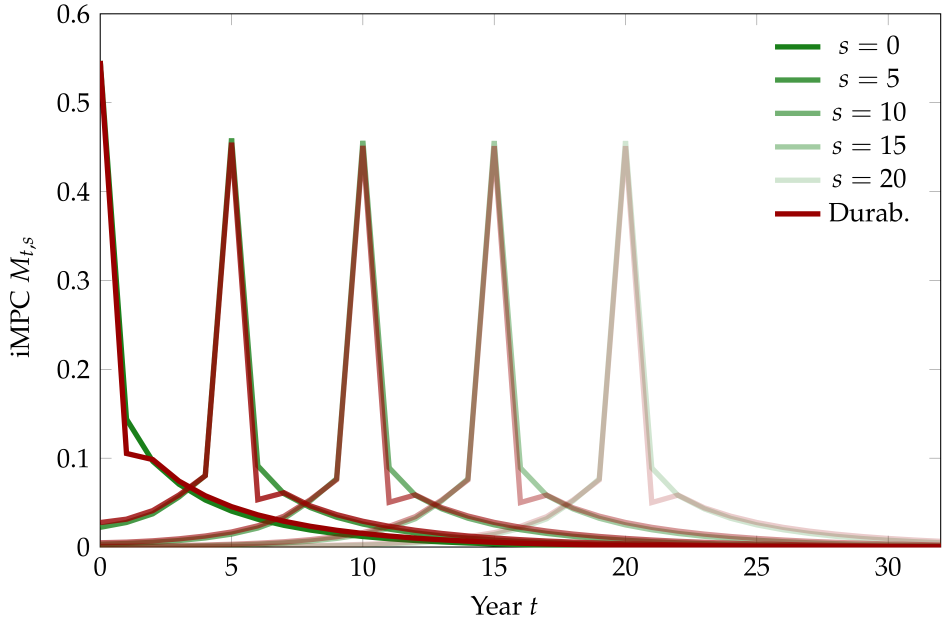

Note: This plot shows intertemporal MPCs implied by our benchmark two-asset heterogeneous-agent, and those implied by an enlarged two-asset model with durable goods.

$x_{it} \equiv c_{it} + d_{it} - (1-\delta_D)\, d_{it-1}$ in the enlarged model are therefore a simple lagged transformation of nondurable expenditures in the baseline model:

$$x_{it} = \frac{1+v}{1+\frac{r+\delta_D}{1+r}v} c_{it}^{ND} - \frac{(1-\delta_D)\, v}{1+\frac{r+\delta_D}{1+r}v} c_{it-1}^{ND}$$

Following the argument in section 2.3, in the aggregate this behavior defines an expenditure function $\mathcal{X}_t(\{Z_s\})$.

Intertemporal Keynesian Cross with durables. On the production side, we maintain our assumption that firms produce a unique good out of labor. The resource constraint for the economy is now

$$G_t + \mathcal{X}_t(\{Y_s - T_s\}) = Y_t \tag{52}$$

Totally differentiating (52), we obtain an intertemporal Keynesian Cross as in (33) where $M_{t,s}^Y = M_{t,s}^T \equiv \frac{\partial \mathcal{X}_t}{\partial Z_s}$, as claimed in section 2.5.

iMPC matching and extrapolation to other columns of M. We now interpret the Norwegian data as coming from a model with durables. We calibrate the durables depreciation rate to $\delta_D = 20\%$ (an average of durable depreciation rates from Fraumeni 1997). In Fagereng et al. (2018), the ratio of the marginal propensity to spend on cars and boats to the overall marginal propensity to spend is only 6%. We therefore conservatively set $v = 10\%$. We then recalibrate our HA-illiq model by changing the level of liquid assets so as to match the marginal propensities to spend on both durables and nondurables. This procedure yields $\frac{B}{Z} = 36\%$ of steady state after-tax income.

Figure 11 repeats figure 3, but now compares our benchmark model to this enlarged model model with durables. The extrapolation to later columns of the **M** matrix implied by the model with durables is still remarkably close to our benchmark. The tents have a similar peak. The main difference is that spending is not as elevated in the year immediately after the income receipt, as households decumulate some of their durables. We conjecture that the iMPCs of our durables model would be even closer to those of our main model if we assume some frictions to selling durables.

F Additional model simulations

Figures 12–15 present model comparative statics with respect to ϕ, κ^p, κ^w and ϵ_I.

Figure 12: Varying the Taylor rule coefficient ϕ

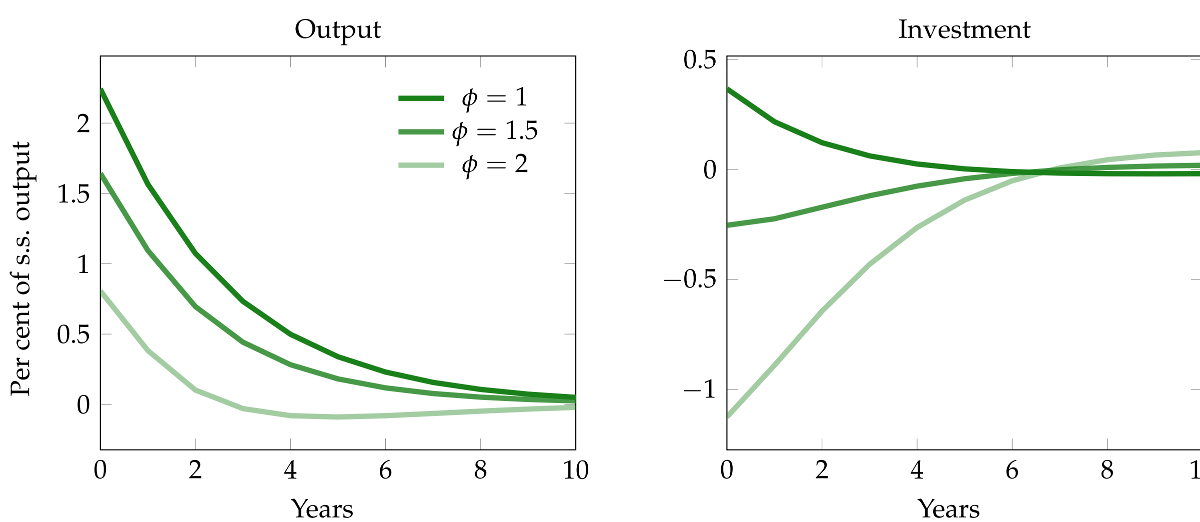

Note. This figure shows the consumption and investment responses of the quantitative illiquid-asset heterogeneous-agent model as the Taylor rule coefficient ϕ is varied. The calibration of the model can be found in section 5. Observe that for ϕ close to 1, where the real rate is approximately constant, investment is not crowded out, but rather crowded in, despite deficit financing. This is because investment responds positively to greater demand.

Figure 13: Varying the degree of price stickiness κ^p

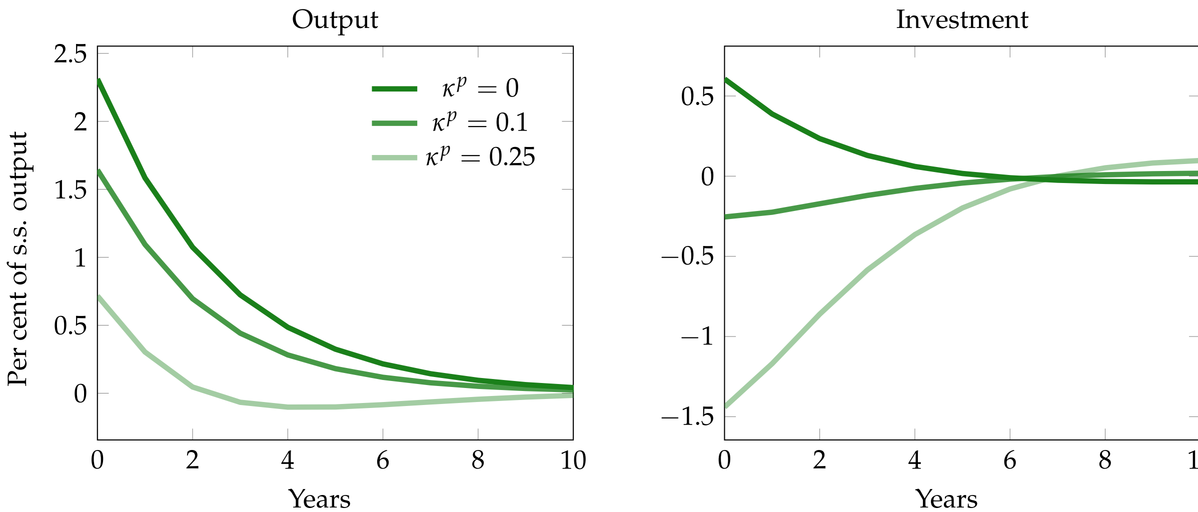

Note. This figure shows the consumption and investment responses of the quantitative illiquid-asset heterogeneous-agent model as the degree of price stickiness κ^p is varied. The calibration of the model can be found in section 5.

Figure 14: Varying the degree of wage stickiness κ^w

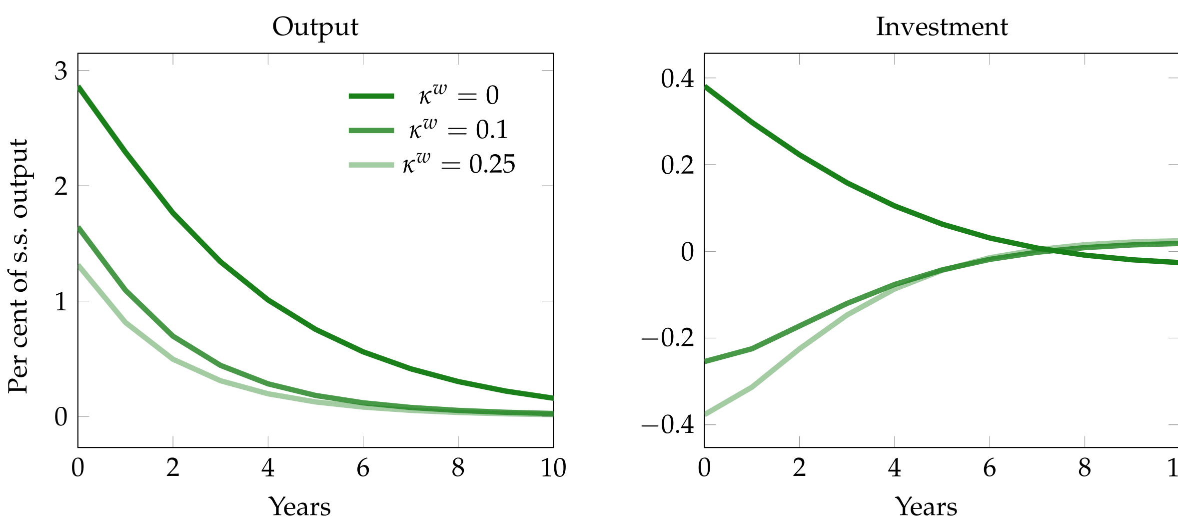

Note. This figure shows the consumption and investment responses of the quantitative illiquid-asset heterogeneous-agent model as the degree of wage stickiness κ^w is varied. The calibration of the model can be found in section 5.

Figure 15: Varying the investment-Q sensitivity ϵ_I

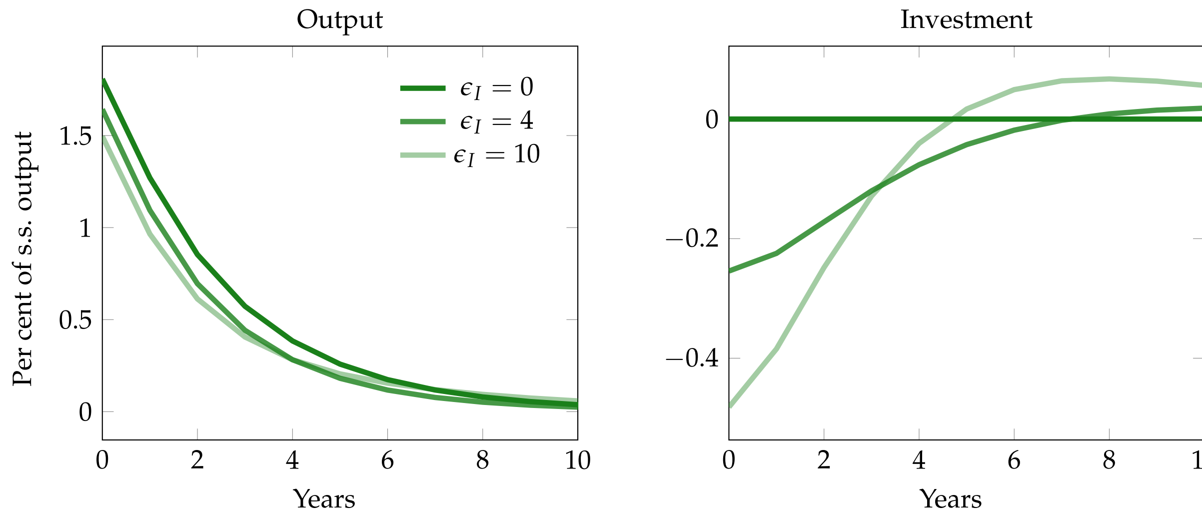

Note. This figure shows the consumption and investment responses of the quantitative illiquid-asset heterogeneous-agent model as the investment-Q sensitivity ϵ_I is varied. The calibration of the model can be found in section 5.

G Computational method

Our model is solved using tools we are currently developing to easily solve dynamic general equilibrium in models that include heterogeneous households or other advanced features. The basic idea is to represent equilibrium as a system of nonlinear equations in aggregate variables $\mathbf{X}_t$, which in general terms can be written as

$$\mathcal{H}_t \left(\{ \mathbf{X}_s \} \right) = 0 \tag{53}$$

In the code, we write these equations as a sequence of functions, which when composed map a set of unknown "inputs" to a set of "targets". To solve for equilibrium, our engine solves for the inputs such that the targets equal zero.

We use this approach to solve both for steady states and for the dynamic response to shocks. For the steady state, computation is relatively straightforward: the implied system of aggregate equations is small, and after the user provides an initial guess we apply standard nonlinear solvers. One nice feature of the design is that there is no built-in distinction between parameters and equilibrium values, nor is there any distinction between equilibrium conditions and calibration targets: they are all simply part of the system of steady-state nonlinear equations that we must solve.

For the computation of dynamics, we linearize and truncate (53) at some far-off horizon T. We then solve the equations using a form of Newton's method, but using interpolation and extrapolation to construct an approximate Jacobian for the heterogeneous-agent household side, since an exact Jacobian is prohibitively costly.

To compute the household side of the model, which enters into the aggregate equations in (53), we use the method of endogenous gridpoints of Carroll (2006), combined with customized code for rapid interpolation and exact forward propagation of the distribution. Each of our heterogeneous household specifications can be solved as a variant of the canonical one-asset consumption-savings problem: since they have the same return except for the date-0 shock, bonds and equity can be combined into a single asset in the HA-std model (which is revalued at date 0 as equity prices change in response to the shock), and the HA-illiq model can be implemented by altering households' exogenous income process.