Deficits and Inflation:
HANK meets FTPL*

George-Marios Angeletos† **Chen Lian**‡ **Christian K. Wolf**§

Conference draft – preliminary and incomplete

July 8, 2024

Abstract

In HANK models, fiscal deficits drive aggregate demand and thus inflation because households are non-Ricardian; in the Fiscal Theory of the Price Level (FTPL), they do so via assumptions about equilibrium selection. *Because* of this difference, the mapping from deficits to inflation in HANK is robust to many of the controversies surrounding the FTPL. *Despite* this difference, a benchmark HANK model predicts as much inflation as the FTPL when fiscal adjustment is sufficiently slow. This is true even in the most extreme FTPL scenario, in which inflationary pressure is so large that it fully finances the deficit. In practice, however, unfunded fiscal deficits are likely to trigger a persistent boom in real economic activity and thus the tax base, substituting for debt erosion from inflation. In empirically disciplined HANK models, this arrests the inflationary effects of unfunded deficits by about one half relative to the textbook FTPL arithmetic.

*For helpful comments and suggestions we thank Manuel Amador, Marco Bassetto, Larry Christiano, Jordí Gali, Joao Guerreiro, Joel Flynn, Mikhail Golosov, Greg Kaplan, Emi Nakamura, Jón Steinsson, Mike Woodford, and seminar participants at numerous institutions. Chen Lian thanks the Alfred P. Sloan Foundation for financial support, and Christian Wolf acknowledges that this material is based upon work supported by the NSF under Grant #2314736.

†Northwestern University and NBER; angeletos@northwestern.edu.

‡UC Berkeley and NBER; chen_lian@berkeley.edu.

§MIT and NBER; ckwolf@mit.edu.

1 Introduction

Do fiscal deficits drive inflation? And if so, how, and how much? One answer is provided by the Fiscal Theory of the Price Level (FTPL): fiscal deficits not backed by commensurate future surpluses *must* be accompanied by an increase in nominal prices and a corresponding drop in the real value of the outstanding nominal debt—or else the government's intertemporal budget constraint cannot hold.[1] This theory has received much attention following the recent inflationary episode (Bianchi et al., 2023; Anderson and Leeper, 2023; Barro and Bianchi, 2023), yet remains controversial, because of its reliance on inherently untestable assumptions regarding equilibrium selection (e.g., Kocherlakota and Phelan, 1999; Buiter, 2002; Niepelt, 2004; Atkeson et al., 2010).[2] Another answer is provided by a much more mainstream, Keynesian, logic: if Ricardian Equivalence fails as a result of liquidity constraints, finite lives, or imperfect foresight, then fiscal deficits will naturally stimulate aggregate demand, driving up real output and thus inflation.[3] This mechanism is absent in the textbook New Keynesian model, because households there are Ricardian, but lies at the heart of both the old IS-LM framework and the modern Heterogeneous Agent New Keynesian (HANK) literature.

The core contribution of this paper is to build a bridge between the FTPL and Keynesian stories of how deficits drive inflation. Our main theoretical result is that, *despite* the different mechanisms at work, a benchmark HANK model with sufficiently delayed fiscal adjustment can feature just as much inflation as what is predicted by the FTPL. At the same time, *because* of the difference in mechanism, HANK naturally sidesteps the controversies surrounding the FTPL. We complement these lessons with another, more practical, takeaway. The simplest textbook FTPL arithmetic stipulates that unfunded deficits induce an exactly offsetting increase in nominal prices. Although the theoretical result stated above holds even in this extreme scenario, in practice unfunded deficits are likely to at least in part finance themselves by triggering a boom in real economic activity and thus the tax base (Angeletos et al., 2024). In a variety of empirically disciplined quantitative exercises, this channel cuts down deficit-driven inflation by around one half relative to the simple FTPL arithmetic.

Environment. For our baseline analytical results we consider the textbook New Keynesian model (Woodford, 2003; Galí, 2008), augmented with household mortality risk (as in Blanchard, 1985) as a convenient proxy for occasionally-binding borrowing constraints (e.g., as used in Farhi and Werning, 2019; Angeletos et al., 2024). When mortality risk is zero, our model nests the standard Representative

[1]This basic prediction holds in both the flexible-price version of the FTPL (Leeper, 1991; Sims, 1994; Woodford, 1995; Bassetto, 2002; Cochrane, 2005; Kaplan et al., 2023) as well as its sticky-price incarnations (Bianchi and Ilut, 2017; Bianchi et al., 2023; Cochrane, 2017, 2018, 2023). In this paper we are concerned exclusively with the latter.

[2]Bassetto (2002) and Cochrane (2005, 2011) have argued against this perspective. One of our main insights will be that, under alternative and natural assumptions on consumer behavior, this debate can actually be sidestepped entirely.

[3]We use the terms "Ricardian" and "non-Ricardian" in their classical sense: households are Ricardian if and only if they satisfy the same conditions as those in Barro (1974)—infinite horizons, no borrowing constraints, and full rationality.

Agent New Keynesian (RANK) model; when it is positive, the model emulates HANK. The economic environment reduces to three equations: the first describes consumption-savings optimality and thus aggregate demand; the second summarizes aggregate supply, i.e., the familiar New Keynesian Phillips Curve (NKPC); and the third is the government's intertemporal budget constraint.

We close the model with a pair of feedback rules for how the monetary and fiscal authorities adjust their policy instruments—i.e., the nominal rate of interest and taxes, respectively—as a function of the state of the economy (similarly to Leeper, 1991). Fiscal policy is parameterized by two key coefficients, $\tau_d \in [0,1)$ and $\tau_y \in [0,1)$. The former measures how much, and how fast, future taxes adjust to current deficits, conditional on real output; the latter parameterizes the automatic feedback from economic activity to the tax base, as studied in Angeletos et al. (2024). The FTPL conventionally assumes the absence of fiscal adjustment, i.e., $\tau_d = 0$ ("active" fiscal policy); fiscal policy is instead said to be "passive" if fiscal adjustment is always sufficient to stabilize government debt. The monetary authority sets nominal and thus (expected) real interest rates; its policy rule is "active" if real interest rates increase following a boom in real activity and inflation, and "passive" otherwise.

RANK-FTPL, and its controversies. The modern, sticky-price version of the FTPL (Bianchi and Ilut, 2017; Cochrane, 2017, 2018, 2023) is nested in our environment by assuming away mortality risk and by imposing an "active-fiscal, passive-monetary" policy mix. The first assumption yields RANK, while the second induces a unique equilibrium, henceforth referred to as the FTPL equilibrium, in which fiscal deficits trigger a joint boom in real economic activity and nominal prices. The sharpest version of this equilibrium obtains when real rates are kept fixed—similarly to Barro and Bianchi (2023)—and $\tau_d = \tau_y = 0$. In this special case, which we refer to as the simple FTPL arithmetic, the entirety of any fiscal innovation is absorbed by an exactly offsetting movement in nominal prices and, thereby, in the real value of the outstanding public debt—e.g., a one-period deficit shock of one per cent of steady-state output induces a price jump equal to the inverse of the debt-to-GDP ratio. More generally, the FTPL equilibrium allows the fiscal deficit to be financed not only by debt erosion, but also by monetary accommodation as well as a boom in the tax base. Either way, the active fiscal policy induces an equilibrium in which, even though households are Ricardian in the classical textbook sense of Barro (1974), deficits end up driving economic activity and thus inflation, and they do so by just enough to ensure that the government's intertemporal budget constraint holds.

The FTPL's account of how—and how much—fiscal deficits drive inflation has been subject to controversies. Some work has argued that an active fiscal policy amounts to an untestable, off-equilibrium, threat to "blow up the government budget" or to induce non-existence of a continuation equilibrium (Kocherlakota and Phelan, 1999; Buiter, 2002; Atkeson et al., 2010), with rebuttals in Bassetto (2002) and Cochrane (2005). The fragility of the FTPL equilibrium has also been highlighted: its prediction that fiscal deficits drive inflation is not robust to either (i) fiscal adjustment in the far-ahead future or (ii) appropriate noise in the spirit of the global games literature (Angeletos and Lian, 2023). An integral part of our contribution will be to show that HANK—where fiscal deficits drive inflation simply because households are non-Ricardian, and not because of equilibrium selection in the presence of Ricardian households—completely sidesteps these debates.

HANK meets FTPL, without the controversies. We next turn to the main part of our analysis: how, and by how much, deficits drive inflation in HANK. Since households now are non-Ricardian, fiscal deficits influence aggregate demand and thereby real output and inflation even if monetary policy is "active"—leaning against the boom through real rate hikes—and fiscal policy is "passive"—stabilizing public debt through eventual tax hikes. Our first main result is that, *despite* this difference in mechanism, HANK can actually replicate the FTPL's core prediction about how fiscal deficits induce debt erosion through inflation. We begin our analysis with the special case of short-term public debt and a monetary authority that stabilizes the (expected) real rate of interest. We then establish that, in this setting, a deficit shock induces a price jump—and thus debt erosion—that is increasing with the delay in fiscal adjustment (i.e., decreasing in τ_d). Furthermore, as fiscal adjustment is delayed further and further (in the sense of $\tau_d \to 0$), the initial price jump converges to its FTPL counterpart. This result holds independently of the strength of the tax base channel (τ_y); if this channel is absent ($\tau_y \to 0$) or if prices are very flexible ($\kappa \to \infty$), then the jump in prices *entirely* finances the fiscal deficit, just as in the simplest FTPL arithmetic mentioned earlier. Finally, the equivalence extends to richer monetary policy in the following sense: conditioning on the same path of real rates, the FTPL equilibrium and our HANK equilibrium again imply the exact same date-0 price jump.

We next show that, *because* of the difference in mechanism, HANK changes the deficit-inflation mapping in two important ways relative to the FTPL.

1. *Robustness.* HANK sidesteps the theoretical debates reviewed above. Formally, we show that the price response to a fiscal deficit shock in HANK is entirely unaffected by sufficiently far-ahead fiscal financing and an eventually active monetary reaction. It thus in particular follows that the deficit-inflation nexus in HANK is entirely robust to textbook "passive" fiscal policy—all that is required is that fiscal adjustment is sufficiently, though not perpetually, delayed. At the heart of this result are the short horizons of the non-Ricardian households in HANK, discounting any sufficiently far-ahead fiscal adjustment and monetary policy reaction.

2. *Front-loading.* Yet again because of household short horizons, the fiscally-induced demand boom—and so inflation—in HANK is necessarily more front-loaded and short-lived than in FTPL. We show that this front-loading is further reinforced by empirically relevant assumptions on the distribution

of government bond holdings. If bonds are predominantly held by households with low marginal propensities to consume (MPC), while fiscal transfers are received by high-MPC households, then the demand boom and thus inflation burst get front-loaded even more.

Our final theoretical contribution is to extend the "HANK-meets-FTPL" results to the empirically relevant case of long-term government debt. Deficit-relevant debt erosion is now given by the *cumulative* inflation response, discounted at the debt maturity rate $\delta \geq 0$. If real interest rates are fixed and the tax base channel is weak (either $\tau_y \to 0$ or $\kappa \to \infty$), then this cumulative inflation response is equal to the inverse of the debt-to-GDP ratio, in both FTPL and, with sufficiently delayed fiscal adjustment, HANK. If the tax base channel is instead operative, then the cumulative inflation response in HANK is below that in FTPL. The reason is front-loading: the inflation burst in HANK is more short-lived, and so, with long-term debt, there is less scope for debt erosion relative to tax base self-financing.

Quantification. The theory so far has revealed that the simple FTPL arithmetic—i.e., of prices jumping by exactly enough to finance an unfunded fiscal deficit—can also, under appropriate assumptions, emerge in HANK. Our final contribution is a quantitative analysis, aimed at ascertaining the empirical plausibility of this simple arithmetic. The main result is that inflationary pressures in practice are likely to be materially weaker, for two main reasons: first, the tax base "self-financing" channel studied in Angeletos et al. (2024), which endogenizes fiscal surpluses; and second, the moderating effects of inflation front-loading in the presence of long-term debt, as discussed above.

Our quantitative evaluation is based on a richer, and empirically disciplined, variant of the model in our theoretical analysis. It features: intertemporal marginal propensities to consume (iMPCs) consistent with empirical evidence (e.g., following Fagereng et al., 2021); heterogeneous fiscal transfer incidence and government debt holdings; an empirically estimated hybrid NKPC (following Barnichon and Mesters, 2020); government debt with a maturity profile close to that observed in the data; and a meaningful feedback from real activity to fiscal surpluses. Within this model, we consider the same policy experiment as in the theoretical part of our paper: a one-off fiscal deficit without (or with very delayed) future tax hikes, together with a monetary authority that fixes the expected real rate. Our headline finding is that, even for a relatively steep NKPC, debt erosion through inflation is cut down by around half relative to the textbook FTPL arithmetic. In addition, most of the inflation response is concentrated on the first few years after the initial stimulus, consistent with our front-loading results. These conclusions extend to a range of model variants, including: alternative assumptions on transfer receipt and government debt holdings; more aggressive monetary policy reactions and faster fiscal adjustment; alternative assumptions on the shape of the NKPC and the government debt maturity profile; cognitive discounting among consumers; and a full-blown quantitative heterogeneous-agent ("HANK") model, as in Auclert et al. (2023) or Kaplan et al. (2018).

Finally, we connect our results with the post-covid inflationary episode. To this end, we scale the size of the initial deficit shock to be in line with federal covid-related spending, and then fix the path of real interest rates as in the data. Our model predicts a meaningful—yet, relative to the textbook FTPL arithmetic, materially dampened—as well as front-loaded burst in inflation, consistent with actually observed inflation dynamics, as for example estimated in Barro and Bianchi (2023).

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

First, we connect with the literature on the FTPL, discussed throughout this introduction. As explained, the fundamental difference between HANK and the modern, sticky-price version of the FTPL is the origin of the failure of Ricardian Equivalence and of the resulting effects of deficits on inflation: conventional, empirically verifiable assumptions regarding consumption behavior in HANK; and assumptions regarding equilibrium selection in FTPL. Our key contribution is to study if and how this difference in the underlying theoretical mechanisms translates in different tangible predictions about the relation between deficits and inflation.

Second, our analysis is situated in a very large literature on the effects of fiscal policy when households are non-Ricardian and prices are sticky. Prior work has investigated the effects of fiscal stimulus in such environments—on both real activity as well as inflation—in analytical (Galí et al., 2007; Bilbiie, 2020; Angeletos et al., 2024) as well as in quantitative (Auclert et al., 2023; Kaplan et al., 2018; Hagedorn et al., 2019; Eusepi and Preston, 2018) environments alike. Relative to that literature, our main contribution is to shed new light on the differences and similarities to the FTPL, along the lines explained above. A secondary contribution is to provide further analytical characterizations of the deficit-inflation mapping. Our analytical expressions furthermore reveal the key determinants of the deficit-inflation mapping—determinants that we, for our quantitative analysis and in particular post-covid application, discipline by the best available empirical evidence.

Third, much recent work has tried to shed light on the role played by fiscal stimulus (and monetary accommodation) in the post-covid inflationary episode. While much interest has centered on the FTPL (e.g., Bianchi et al., 2023; Anderson and Leeper, 2023; Barro and Bianchi, 2023), some commentators have instead proposed Keynesian mechanisms along the lines suggested here (Blanchard, 2021; Summers, 2021; Blanchard and Bernanke, 2023). Our contribution here is to discipline the Keynesian mechanism with the best available evidence on its ingredients, derive quantitative implications, and explore connections with the simple FTPL arithmetic.

Outline. Section 2 presents the model framework. Section 3 reviews the predictions of the textbook FTPL and some of the controversies that surround it. Section 4 contains our main results on characterizing the deficits-inflation nexus in HANK, with Section 5 discussing some further model extensions. Quantitative explorations follow in Section 6 and we conclude in Section 7.

2 Environment

For our main theoretical analysis we consider a perpetual-youth, overlapping-generations (OLG) version of the New Keynesian model. Similarly to Del Negro et al. (2015), Farhi and Werning (2019), and Angeletos and Huo (2021), finite lives serve as a convenient proxy for liquidity frictions.

The model environment builds on Angeletos et al. (2024). Since microfoundations for the core model relations are already discussed there, we will here whenever possible work directly with log-linearized dynamics, presenting full relations in levels only when necessary.[4] Levels will throughout be indicated with uppercase variables, while lowercase variables denote log-deviations from the economy's deterministic steady state.[5] We log-linearize around a deterministic steady state in which inflation is zero, real allocations are given by their flexible-price counterparts, and the real government debt burden is constant at some level. Time is discrete, indexed by $t \in \{0, 1, ...\}$.

2.1 Aggregate demand

The economy is populated by a unit continuum of households. A household survives from one period to the next with probability $\omega \in (0, 1]$, so that $1 - \omega$ is the mortality rate. Whenever a household dies, it is replaced by a new household. Households have standard separable preferences over consumption and hours worked, receive labor and dividend income, pay taxes, and can save and borrow through an actuarially fair, risk-free nominal annuity. They furthermore, as in Angeletos et al. (2024), make contributions to a social fund whose proceeds are distributed to newborn households; this ensures that all cohorts enjoy the same wealth and hence consumption in steady state, and also guarantees a steady-state real interest rate of $R^{ss} = \beta^{-1}$, where $\beta \in (0, 1)$ is the household discount factor. Finally, to facilitate aggregation, we assume that all households receive the same dividend payments, pay the same taxes, and supply the same (union-intermediated) total amount of hours worked. A detailed discussion of the household decision problem and its solution is presented in Appendix A.

Deriving the (log-linearized) consumption function of each individual household, and then aggregating across households, we obtain the following aggregate consumption function:

$$c_t = (1 - \beta\omega)\left(\tilde{a}_t + \mathbb{E}_t\left[\sum_{k=0}^{\infty} (\beta\omega)^k (y_{t+k} - t_{t+k})\right]\right) - \beta\left(\sigma\omega - (1 - \beta\omega)\frac{D^{ss}}{Y^{ss}}\right)\mathbb{E}_t\left[\sum_{k=0}^{\infty} (\beta\omega)^k r_{t+k}\right], \tag{1}$$

where c_t is aggregate household consumption, $\tilde{a}_t$ is real household wealth (also, the real value of government debt, by asset market clearing), y_t is real household income (labor income plus dividends),

[4] To ensure that the analysis in the present paper is self-contained, we nevertheless go through the full detailed set-up again in Appendix A.

[5] To accommodate the case of zero debt (which is conceptually useful even if empirically irrelevant), all fiscal and household wealth variables will be measured in terms of absolute deviations from the steady state, scaled by steady-state output.

t_t is real tax payments, r_t is the expected real rate of interest, σ is the elasticity of intertemporal substitution, and D^{ss}/Y^{ss} is the steady state wealth-to-income ratio.

Equation (1) is an intuitive generalization of the usual permanent-income hypothesis (PIH) aggregate consumption function. First, the marginal propensity to consume (MPC) out of current wealth and current disposable income equals $1-\beta\omega$. Thus, whenever $\omega < 1$, this MPC is elevated relative to the PIH benchmark. Second, future income is not just discounted at the steady-state rate of interest, but additionally discounted at rate ω, reflecting finite household lives (or, less literally, the bite of liquidity constraints, as in Farhi and Werning, 2019). Third, the response to interest rate movements reflects both the usual intertemporal substitution motive and wealth effects, with smaller ω dampening intertemporal substitution and strengthening wealth effects.

Connection to HANK. The three properties discussed above—elevated MPCs out of current income and current liquid wealth, extra discounting of the future, as well as a smaller sensitivity to real rates—are shared by a large class of Heterogeneous Agent New Keynesian (HANK) models (e.g., Kaplan et al., 2018; Auclert et al., 2023). In particular, it is well-known that the *intertemporal* marginal propensities to consume (iMPCs) induced by an aggregate consumption function like (1) agree closely with those in quantitative HANK models (e.g., Wolf, 2021). Our OLG model should thus not be interpreted literally, but rather as a tractable representation of HANK-type models.

The main limitation of the demand block considered here is that it abstracts from realistic cross-sectional heterogeneity in wealth, marginal propensities to consume, and exposure to fiscal transfers. As we will see later, however, these abstractions are not driving our main lessons about HANK-FTPL equivalence and the additional robustness of HANK: in Section 5.2 we verify that our main theoretical results extend to a variant model that accommodates a stylized version of such cross-sectional heterogeneity; and in Section 6 we show that our quantitative findings are only marginally affected if we move to a fully-fledged HANK model.

2.2 Aggregate supply

The production side of the economy is the same as in the textbook New Keynesian model: there is a unit-mass continuum of monopolistically competitive retailers, who set prices subject to the standard Calvo friction, hire labor on a spot market, produce according to a technology that is linear in labor, and pay out all their profits as dividends back to the households. Given our assumptions on household preferences, union labor supply rationing, and income taxes, it follows that the supply block of our economy reduces exactly to the standard New Keynesian Phillips curve (NKPC):

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t[\pi_{t+1}], \tag{2}$$

where $\kappa = \frac{(1-\theta)(1-\beta\theta)\left(\frac{1}{\varphi}+\frac{1}{\sigma}\right)}{\theta} > 0$ and $1-\theta \in (0,1)$ is the Calvo reset probability. Equivalently,

$$\pi_t = \kappa \sum_{k=0}^{\infty} \beta^k \mathbb{E}_t\left[y_{t+k}\right]. \tag{3}$$

A detailed discussion of the production side and the derivation of (2) is provided in Appendix A. In Section 5.3 we show that the essence of our analysis remains the same if (2) is replaced with a more empirically relevant, hybrid NKPC.

2.3 Fiscal policy

The government issues non-contingent, short-term, nominal debt; the extension to long-term debt is provided in Section 5.1. We let B_t denote nominal public debt outstanding at the beginning of period t and $D_t \equiv B_t/P_t$ denote the *real* value of public debt (which, by market clearing, is also the private sector's real, net, financial wealth). In log-linearized terms, the government flow budget is

$$d_{t+1} = \underbrace{\frac{1}{\beta}(d_t - t_t) + \frac{D^{ss}}{Y^{ss}} r_t}_{\text{expected debt burden tomorrow}} - \underbrace{\frac{D^{ss}}{Y^{ss}}(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}])}_{\text{debt erosion due to inflation surprise}}, \tag{4}$$

where t_t denotes the taxes collected in period t (and here also the primarily surplus), r_t is the (expected) real rate between dates t and $t+1$, and $\frac{D^{ss}}{Y^{ss}}$ is the steady-state debt-to-GDP ratio. Finally, the initial condition for real government debt is given by

$$d_0 = -\frac{D^{ss}}{Y^{ss}} \pi_0. \tag{5}$$

The government must satisfy both the above flow constraint (at each t) as well as the familiar non-Ponzi condition ($\lim_{k\to\infty} \mathbb{E}_t\left[\beta^k d_{t+k}\right] = 0$). We can thus go back and forth between the infinite sequence of flow budget constraints and the corresponding intertemporal budget constraint:

$$d_t = \mathbb{E}_t\left[\sum_{k=0}^{\infty} \beta^k \left(t_{t+k} - \beta \frac{D^{ss}}{Y^{ss}} r_{t+k}\right)\right]. \tag{6}$$

Equation (6) is commonly referred to as the "intertemporal government budget constraint;" Cochrane (2005) instead interprets equation (6) as a "valuation equation" instead of a "constraint." No matter the interpretation, equations (4) and (6) always represent a set of joint constraints on fiscal policy and equilibrium outcomes.

Tax rule. It remains to specify the evolution of total tax revenue t_t. For our main analysis, we will restrict attention to the subclass of tax policy rules given as:

$$t_t = \underbrace{-\varepsilon_t}_{\text{deficit shock}} + \underbrace{\tau_d(d_t + \varepsilon_t)}_{\text{fiscal adjustment}} + \underbrace{\tau_y y_t}_{\text{tax base}}, \tag{7}$$

for some $\tau_d, \tau_y \in [0,1)$.

The tax rule (7) mirrors those commonly assumed in prior work. Its first component, ϵ_t, reflects exogenous innovations in taxes and hence in deficits. In particular, we interpret ε_t as an unexpected lump-sum transfer given to households in period t—the "deficit shock". The second component captures how much lump-sum taxes adjust in response to the outstanding debt burden. In particular, similarly to Leeper (1991), $\tau_d \in [0,1)$ is a scalar that parameterizes the speed and strength of fiscal adjustment following a deficit shock: tax financing of the initial deficit is instantaneous for $\tau_d \to 1^-$, and absent for $\tau_d \to 0^+$. An important policy question, and one central to the remainder of our analysis, is what values of τ_d guarantee that d_t converges back to 0 following any given deficit shock. Finally, the third term indicates how much tax revenue covaries with aggregate income for a given tax schedule. For concreteness, we here interpret $\tau_y \in [0,1)$ as a proportional tax on total household income, with this tax distortionary but the rate τ_y time-invariant. This automatic feedback term—which gives the tax base "self-financing" channel studied in Angeletos et al. (2024)—will in particular loom large in our quantitative analysis in Section 6.

2.4 Monetary policy

The monetary authority sets I_t, the nominal rate of interest between dates t and $t+1$, according to the following Taylor-type policy rule:

$$\frac{I_t}{\Pi^{ss}} = R^{ss}\mathbb{E}_t\left[\frac{\Pi_{t+1}}{\Pi^{ss}}\right]\left(\frac{Y_t}{Y^{ss}}\right)^{\phi}, \tag{8}$$

where Π_t is the rate of inflation between t and $t+1$, Π^{ss} and R^{ss} are the steady-state values of inflation and of the real interest rate, respectively, and $\phi \in \mathbb{R}$ is a policy coefficient. Log-linearizing, we can write the above as a rule for the (expected) real interest rate:[6]

$$r_t = \phi y_t, \tag{9}$$

Monetary policy is thus parameterized by whether monetary policy implements higher or lower real rates (respectively, whether $\phi > 0$ or $\phi < 0$) in response to any demand-driven boom in real economic activity and thereby in inflation.[7]

[6] An integral part of the upcoming analysis will be to study *how* and *why* equilibrium depends on policy. Because the economy admits multiple equilibria for some policies but not for others, it is well understood (e.g., see King, 2000; Atkeson et al., 2010; Cochrane, 2011) that feedback rules such as (7) and (9) can generally serve two conceptually distinct functions: the one is to induce a particular comovement between policy instruments and economic activity along a given equilibrium; the other is to select a particular equilibrium by committing to take certain actions *off* equilibrium. We still study the role of both functions, in both FTPL and in HANK.

[7] Note that here we depart very slightly from the common practice of parameterizing monetary policy by the response of the nominal rate to inflation. In this case, the "Taylor principle" (or "active" monetary policy) translates to $\phi > 0$, rather to the more familiar requirement of a response coefficient greater than one. We consider this alternative set-up for reasons of analytical convenience; our quantitative analysis in Section 6 will also consider alternative monetary rules.

We will pay particular, though not exclusive, attention to the special case of $\phi = 0$—i.e., fixed real rates in response to fiscal shocks, as, e.g., in Woodford (2011), Auclert et al. (2023) and Angeletos et al. (2024). This special case can be interpreted as a "neutral" monetary policy that neither leans against a fiscally-led boom nor accommodates deficits by letting real rates (and so the government's cost of borrowing) fall. We do so because setting $\phi = 0$ has three distinct advantages for the purposes of our analysis. First, it lets monetary policy be "passive", thus making space for the FTPL. Second, it makes sure that deficits do not indirectly feed into real rates, which in turn could influence aggregate demand via the intertemporal substitution margin. Together, these two properties help sharply isolate both the FTPL mechanism when $\omega = 1$ and the HANK alternative when $\omega < 1$. Finally, fixed real rates facilitates the comparison between our results and those of Barro and Bianchi (2023), who also abstract from any effects of deficits on real interest rates. Notwithstanding these points, we will of course also discuss what happens when $\phi < 0$—a textbook "passive" monetary rule, as in much prior FTPL work.

2.5 Equilibrium definition and roadmap

A standard equilibrium definition combines (i) optimality for consumers; (ii) optimality for firms, (iii) market-clearing, and (iv) budget balance for the government, together with the no-Ponzi condition. With the policy rules we have assumed here, and with the conventional requirement that $\{c_t, y_t, \pi_t\}_{t=0}^{\infty}$ is bounded, this boils down to the following.[8]

Definition 1. An equilibrium is a stochastic path $\{c_t, y_t, \pi_t, \tilde{a}_t, d_t, t_t, r_t\}_{t=0}^{\infty}$ for consumption, output, inflation, the real values of household wealth and public debt, total tax revenue, and real interest rates that satisfies all of the following: aggregate demand (1), aggregate supply (2), market clearing $y_t = c_t$ and $\tilde{a}_t = d_t$, the boundedness of $\{c_t, y_t, \pi_t\}_{t=0}^{\infty}$, the law of motion for public debt (4), along with the initial condition (5) and the no-Ponzi condition $\mathbb{E}_t\left[\lim_{k\to\infty} \beta^k d_{t+k}\right] = 0$, and the fiscal and monetary policy rules (7) and (8).

The remainder of the paper studies the deficit-inflation mapping in two scenarios: RANK-FTPL, nested with $\omega = 1$ and $\tau_d = 0$, and HANK, corresponding to $\omega < 1$ and arbitrary $\tau_d \in [0, 1]$.

3 A review of RANK-FTPL

In this section, we restrict $\omega = 1$, which gives the textbook, representative agent, New Keynesian model (RANK). We review the indeterminacy problem of this model, characterize its conventional and FTPL

[8]Unlike Leeper (1991) and much of the subsequent applied literature on monetary-fiscal interactions, we do not a priori require d_t to be bounded, in order to avoid a confusion: the FTPL does *not* require boundedness of public debt, but it maintains the no-Ponzi condition, or equivalently equation (6). That said, as explained in Footnote (10) below, nothing of essence changes if we a priori require boundedness of d_t.

solutions, and highlight how the latter breaks Ricardian Equivalence despite the assumption of an infinitely lived, fully liquid, and fully rational household, as in Barro (1974).

3.1 Equilibrium characterization

We begin our analysis of equilibria in RANK by noting that, in any equilibrium and regardless of policy, the path of inflation is uniquely pinned down by the path of real output, via the NKPC (2).[9] Next recall that, when $\omega = 1$, the aggregate consumption function (1) boils down to the familiar, representative-agent, Euler equation:

$$c_t = -\sigma r_t + \mathbb{E}_t[c_{t+1}]$$

Combining with market clearing (i.e., that $y_t = c_t$) and the assumed rule for monetary policy (8), we arrive at the following equilibrium restriction on the path of output:

$$y_t = \frac{1}{1+\sigma\phi}\mathbb{E}_t[y_{t+1}]. \tag{10}$$

Studying (10), we see that two cases emerge, depending on the value of ϕ. When $\phi > 0$ ("active monetary policy"), then the unique bounded solution to equation (10) is $y_t = 0$, i.e., the economy rests in steady state. When instead $\phi \leq 0$ ("passive monetary policy"), then this remains a solution, but it is not the only one: the set of bounded solutions is now given by

$$y_t = \left(1+\sigma\phi\right) y_{t-1} + \eta_t, \tag{11}$$

for any bounded random variable η_t such that $\mathbb{E}_t\left[\eta_{t+1}\right] = 0$, and with the convention $y_{-1} = 0$.

Either way, given a solution to (10), characterizing an equilibrium is straightforward. First, π_t is pinned down by the NKPC (2). Consumer optimality, firm optimality, and market clearing are thus all satisfied. Second, to complete an equilibrium, the only thing remaining is to check that the government's intertemporal budget constraint (6) is also satisfied. When $\tau_d > 0$, then this is trivially true for both active as well as passive monetary policy. When instead $\tau_d = 0$, as in the FTPL, then this is true only under a tight set of conditions: monetary policy has to be passive, y_t has to follow (11), and η_t has to equal a specific multiple of ε_t, the concurrent fiscal innovation. The remainder of this section studies this logic in greater detail, first in an instructive special case, and then in general terms.

The simple FTPL arithmetic. We begin with the instructive special case of $\phi = 0$ and $\tau_y = \tau_d = 0$—i.e., no adjustment in either real interest rates or future tax revenue; we will refer to this case as the

[9] While this observation is trivial, it is worth emphasizing, because it helps clarify how the equilibrium indeterminacy found in the New Keynesian model is distinct from that found Sargent and Wallace (1975) and flexible-price versions of the FTPL (Leeper, 1991; Sims, 1994; Cochrane, 2005); also see Hagedorn (2024). There, a nominal interest rate peg—or more generally a passive monetary policy—leaves the nominal price level indeterminate *conditional* on the path of real output. Here, indeterminacy instead obtains in output itself, due to the Keynesian feedback between spending and income.

simple FTPL arithmetic, given its centrality in the FTPL literature. (11) then reduces to a random walk for output, which via the NKPC (2) also yields a random walk for inflation:

$$y_t = y_{t-1} + \eta_t \qquad \text{and} \qquad \pi_t = \frac{\kappa}{1-\beta} y_t = \pi_{t-1} + \frac{\kappa}{1-\beta}\eta_t.$$

At the same time, the government's intertemporal budget constraint reduces to

$$-\frac{D_{ss}}{Y_{ss}}\left(\pi_t - \mathbb{E}_{t-1}\left[\pi_t\right]\right) = \varepsilon_t,$$

which simply states that the inflation surprise and the resulting debt erosion must fully offset the concurrent deficit innovation. Combining the two relations above pins down η_t, the innovation in output in (11), as follows:

$$\frac{D_{ss}}{Y_{ss}}\frac{\kappa}{1-\beta}\eta_t = \varepsilon_t,$$

To summarize, the deficit shock ϵ_t results in an innovation η_t of exactly the right size to induce output and thus inflation surprises that finance the initial fiscal outlay.

The general case. We now show how our discussion of the equilibrium characterization—and in particular the FTPL logic—generalizes to the entire policy space under consideration.

Proposition 1. *Suppose that* $\omega = 1$.

1. *When* $\phi \leq 0$ *and* $\tau_d > 0$ *(i.e., both policies are passive), then there is a continuum of equilibria. In particular, an output path* $\{y_t\}_{t=0}^{\infty}$ *is part of an equilibrium if and only if it follows* (11)*, with* η_t *being an arbitrary innovation such that* $\mathbb{E}_t\left[\eta_{t+1}\right] = 0$*, possibly correlated with* ε_t.

2. *When* $\phi > 0$ *and* $\tau_d > 0$ *(i.e., active monetary policy and passive fiscal policy), then there exist a unique equilibrium, given by* (11) *with* η_t *identically zero. Equivalently,* $y_t = \pi_t = 0$ *for all t and all states of nature. That is, deficit shocks do not affect output and inflation.*

3. *When* $\phi \leq 0$ *and* $\tau_d = 0$ *(i.e., active fiscal policy and passive monetary policy), then there exists a unique equilibrium, given by* (11) *with* $\eta_t = \frac{1-\beta(1+\sigma\phi)}{\tau_y + (\kappa - \beta\phi)\frac{D^{ss}}{Y^{ss}}}\varepsilon_t$*. That is, deficit shocks trigger persistent fluctuations in output and inflation.*

4. *When* $\phi > 0$ *and* $\tau_d = 0$ *(i.e., both policies are active), then an equilibrium does not exist.*

Proposition 1 is essentially the sticky-price version of Leeper (1991).[10] Its first part shines light on the basic indeterminacy problem of the New Keynesian model: when both policy rules are passive,

[10] There are two inessential differences. First, we parameterized the monetary policy rule by the responsiveness of the real rate to output, instead of that of the nominal rate to inflation. This explains why our derivations are simpler than the textbook, New Keynesian, treatment of the Taylor principle: we have reduced the determinacy question from a two-dimensional problem in (y_t, π_t) space to a one-dimensional problem in y_t space, without affecting any of the economic essence. Second, in our equilibrium definition, we have not required boundedness of d_t; instead, we only impose boundedness for $\{y_t, \pi_t\}_{t=0}^{\infty}$ together with the weaker Ponzi condition $\mathbb{E}_t\left[\lim_{k\to\infty}\beta^k d_{t+k}\right] = 0$. This is why, in our discussion,

then the economy is ridden with multiple equilibria. In particular, there is an equilibrium in which Ricardian equivalence holds and the economy rests in steady state; however, there is also a continuum of other equilibria, in which output and inflation fluctuate with deficit shocks as well as pure sunspot shocks, and so Ricardian equivalence can fail. The second part is the "conventional" equilibrium, in which monetary policy selects the first equilibrium, preserving Ricardian equivalence. The third case is the FTPL, where active fiscal policy selects one particular equilibrium from the second class—one in which Ricardian Equivalence fails *exactly as much as necessary* for the government's intertemporal budget to be satisfied. In the "simple FTPL arithmetic" reviewed above (with $\phi = 0$ and $\tau_y = 0$), the entirety of the deficit shock is absorbed by an increase in the nominal price level and a reduction in the real value of the outstanding nominal debt. More generally, the FTPL selects an equilibrium in which the deficit is financed through (i) a reduction in the government cost of borrowing (when $\phi < 0$), (ii) a boom in income and thus tax revenue (when $\tau_y > 0$), and (iii) an inflation surprise.

We conclude our analysis with a useful corollary, characterizing the impulse response of inflation to a fiscal deficit shock in the FTPL equilibrium.

Corollary 1. *Suppose that $\omega = 1$, $\phi \leq 0$, and $\tau_d = 0$. The impulse response of inflation to a deficit shock is given by*

$$\pi^{FTPL}_{\varepsilon,k} \equiv \frac{d\mathbb{E}_t\left[\pi_{t+k}\right]}{d\epsilon_t} = \frac{\kappa\left(1+\sigma\phi\right)^k}{\tau_y + \left(\kappa - \beta\phi\right)\frac{D^{ss}}{Y^{ss}}} \qquad \forall\, t, k \geq 0. \tag{12}$$

The characterization in Corollary 1 will serve as a key reference point in our subsequent analysis of the deficit-inflation nexus in HANK ($\omega < 1$). In particular, we will pay special attention to $\pi^{FTPL}_{\varepsilon,0}$, which captures the initial price jump and thus, with one-period government debt, the resulting erosion in the real value of that debt.

3.2 The FTPL's controversies

The FTPL helps capture two universally accepted ideas: that fiscal deficits may contribute to inflationary pressures; and that the market value of outstanding public debt may naturally fluctuate with any gap between current deficits and expected future surpluses. What is controversial, however, is the FTPL's specific *formalization* of these two ideas.

Early criticisms of the FTPL emphasized that an active fiscal policy amounts to an untestable, off-equilibrium threat to "blow up the government's budget" or to induce non-existence of a continuation equilibrium (Kocherlakota and Phelan, 1999; Buiter, 2002; Atkeson et al., 2010)—a criticism

"active" fiscal policy is equated with $\tau_d = 0$, rather than $\tau_d \in [0, 1-\beta]$. Intuitively, $\tau_d > 1-\beta$ is necessary and sufficient for d_t not to explode for any path of y_t that solves (10), whereas the weaker restriction that $\tau_d > 0$ is necessary and sufficient for d_t to satisfy the government's non-Ponzi condition, or, equivalently, the consumer's transversality condition.

that proponents of the FTPL have refuted (Bassetto, 2002; Cochrane, 2005). It has also been argued that the FTPL's account of how deficits drive inflation is unduly fragile to changes in the economic environment and the assumed policy mix. First, as already discussed and further elaborated in Section 4.3, fiscal adjustment—even in the very far future—together with a monetary authority that leans against inflationary booms suffices to remove any effect of deficits on inflation. Second, borrowing insights from the global-games literature, Angeletos and Lian (2023) have argued that, with even a small friction to social memory, the unique remaining rational-expectations equilibrium in the RANK environment is the "conventional" solution of $y_t = 0$ (and so $\pi_t = 0$), regardless of whether monetary policy is active or passive. Both of these thought experiments highlight the sensitivity of the FTPL to subtle assumptions about beliefs and coordination in the very long run.

The merits of this ongoing debate are outside the scope of the present paper. Instead, our contribution will simply be to establish that the deficit-inflation mapping in HANK is robust to all of these controversies, for one straightforward reason: fiscal deficits now affect inflation because households are non-Ricardian, and not because the policy mix selects an equilibrium in which Ricardian equivalence fails. Appealingly, then, HANK will allow us to capture the same widely accepted ideas that are at the heart of FTPL, but now with a solid empirical foundation—i.e., high household MPCs—replacing untestable assumptions about equilibrium selection.

4 HANK meets FTPL

This section presents our main theoretical results on the deficit-inflation mapping in HANK ($\omega < 1$), and in particular its relation to the FTPL as reviewed in Section 3. We begin in Section 4.1 with the instructive special case of fixed real rates, establishing HANK-FTPL equivalence when fiscal adjustment is sufficiently delayed in the former. Section 4.2 extends that conclusion to more general monetary feedback rules. Finally, in Section 4.3, we show that the difference in underlying mechanism between HANK and FTPL means that, in HANK, the deficit-inflation mapping now becomes entirely robust to the controversies surrounding FTPL.

4.1 An equivalence result with fixed real rates

In this section we will restrict attention to the special case of fixed (expected) real rates, that is, we set $\phi = 0$ in (8). In addition to helping capture one of the "textbook" cases of the FTPL (e.g., see Barro and Bianchi, 2023), the assumption of fixed real rates is also pedagogically useful: it clearly separates the direct failure of Ricardian Equivalence—whether that due to equilibrium selection in FTPL or that due to short household horizons in HANK—from any indirect effects via the monetary reaction.

This section will establish that, despite the difference in mechanism, if fiscal adjustment in HANK is sufficiently delayed (i.e., $\tau_d \to 0$), then the equilibrium price jump in HANK is identical to the FTPL outcome reviewed above.[11]

Equilibrium characterization. We begin with a characterization of equilibrium outcomes. For now, the characterization is general, applying for all $\tau_d \in [0,1)$.

Proposition 2. *Suppose that $\omega < 1$, $\tau_y > 0$, and $\phi = 0$. There exists a unique bounded equilibrium in the HANK economy. In this equilibrium,*

$$y_t = \chi\left(d_t + \varepsilon_t\right), \qquad \mathbb{E}_t\left[d_{t+1}\right] = \rho_d\left(d_t + \varepsilon_t\right), \qquad \text{and} \qquad \pi_t = \frac{\kappa\chi}{1-\beta\rho_d}\left(d_t + \varepsilon_t\right) \tag{13}$$

for some scalars $\chi > 0$ and $\rho_d \in (0,1)$ that are continuous functions of $\beta, \omega, \tau_y, \tau_d$. The impulse response of inflation to a fiscal deficit shock is then given by

$$\pi_{\varepsilon,k} \equiv \frac{d\mathbb{E}_t\left[\pi_{t+k}\right]}{d\epsilon_t} = \frac{\frac{\kappa\chi}{1-\beta\rho_d}\rho_d^k}{1 + \frac{\kappa\chi}{1-\beta\rho_d}\frac{D^{ss}}{Y^{ss}}} \qquad \forall k \geq 0, \tag{14}$$

and the corresponding debt erosion due to the inflation surprise in (4) *is thus*

$$d_t - \mathbb{E}_{t-1}\left[d_t\right] = -\frac{\kappa\chi\frac{D^{ss}}{Y^{ss}}}{1-\beta\rho_d + \kappa\chi\frac{D^{ss}}{Y^{ss}}}\varepsilon_t. \tag{15}$$

The equilibrium characterization in Proposition 2 embeds a classical Keynesian two-way feedback between aggregate demand and fiscal conditions. Deficits increase demand and so output as well as inflation ($\chi > 0$), which in turn helps stabilize debt via higher tax revenue and debt erosion ($\rho_d < 1$). Equations (14) - (15) then give the full inflation impulse response and the debt erosion term—our main objects of interest—as continuous functions of χ, ρ_d, and other model primitives.

The deficit-inflation mapping in HANK summarized in Proposition 2 contrasts markedly with our discussion of the FTPL. For FTPL, recall from Section 3 that it requires an "active" fiscal policy (i.e., $\tau_d = 0$). By contrast, in HANK, the effects of fiscal deficits are continuous in $\tau_d \in [0,1)$. In this sense, and unlike what is predicted by FTPL, here there is no essential difference between a fiscal authority that "passively" commits to adjusting future surpluses, if slowly—i.e., $\tau_d > 0$ but small—and one that is "active" in the sense of committing to *never* adjust future taxes ($\tau_d = 0$). Similarly, as we will discuss later, the HANK equilibrium is continuous with respect to the monetary policy rule coefficient ϕ at both sides of $\phi = 0$, yet again contrasting with FTPL. These differences are a manifestation of the different underlying economic mechanisms at work.

[11] For most of the analysis in this paper, we associate *delayed* fiscal adjustment with smaller values of τ_d. An alternative notion of delays in fiscal adjustment is a fiscal rule that features no tax adjustment for a finite number of periods after the deficit shock, and then complete adjustment, perfectly stabilizing government debt. The equivalence between these two notions of delayed fiscal adjustment is established in Angeletos et al. (2024), so in this paper we will focus on the former case, which facilitates relating our results to prior work. We return to this point in Section 4.3.

HANK meets FTPL. Despite this difference in the underlying economic mechanism, we find that our HANK equilibrium can—if fiscal adjustment is sufficiently delayed—replicate the core predictions of the FTPL, of prices increasing enough to finance fiscal deficits.

Proposition 3. *Let $\omega < 1$, $\tau_y > 0$, and $\phi = 0$. Then the initial price jump in response to a fiscal deficit shock in the HANK economy decreases in the strength of fiscal adjustment $\tau_d \in [0,1)$. When $\tau_d = 0$, the initial price jump in the HANK economy is exactly same as the FTPL counterpart:*

$$\pi_{\varepsilon,0}\Big|_{\tau_d=0} = \frac{\kappa}{\tau_y + \frac{D^{ss}}{Y^{ss}}\kappa} = \pi_{\varepsilon,0}^{FTPL}. \tag{16}$$

The HANK-FTPL equivalence result in Proposition 3 holds independently of the strength of the tax base channel. If this channel is absent ($\tau_y \to 0$) or if prices are very flexible ($\kappa \to \infty$), then the jump in prices *entirely* finances the fiscal deficit, just as in the simplest textbook FTPL arithmetic.

Corollary 2. *Let $\omega < 1$, $\tau_d = 0$, and $\phi = 0$. If either $\kappa \to \infty$ or $\tau_y \to 0$, then the initial price jump in response to a fiscal deficit shock in the HANK economy converges to the textbook FTPL counterpart:*

$$\pi_{\varepsilon,0} \to \left(\frac{D^{ss}}{Y^{ss}}\right)^{-1} \tag{17}$$

Conversely, the more potent the tax base self-financing channel, the smaller the initial price jump; in particular, if prices are very rigid ($\kappa \to 0$), then the fiscal deficit entirely finances itself through an increase in output (Angeletos et al., 2024).

A visual representation of the equivalence results in Proposition 3 and Corollary 2 is provided in Figure 1.[12] We see that, the weaker fiscal adjustment (in the sense of smaller τ_d), the larger the impact inflation response to a fiscal deficit shock, converging to the FTPL limit as $\tau_d \to 0$. If $\tau_y \to 0$, then this limit is the textbook FTPL arithmetic of prices jumping by exactly enough to fully finance the deficit; if $\tau_y > 0$, then the price jump is strictly smaller. The intuition underlying these results is straightforward, and is rooted in the finite household horizons of HANK. If fiscal adjustment is sufficiently weak and delayed, then, since households are non-Ricardian, an initial fiscal deficit increases demand. If the tax base channel is absent ($\kappa \to \infty$ or $\tau_y \to 0$), then in general equilibrium prices must increase by just enough to arrest this increase in demand; the price increase that does so is $\left(\frac{D^{ss}}{Y^{ss}}\right)^{-1}$, redistributing back from households to the government. If instead the tax base channel is operative, then some flow of funds back towards the government instead occurs via the tax base increase, dampening the price jump required to achieve convergence in general equilibrium.

Our intuitive discussion in the preceding paragraph heavily leveraged the HANK model property of short household horizons. This model property will also lie at the heart of the robustness of the

[12] For this visual illustration, we set $\omega = 0.8$—a meaningful departure from Ricardian equivalence—and $\kappa = 0.1$—a steep NKPC. For our later quantitative analysis we will consider actually empirically disciplined variants of our model.

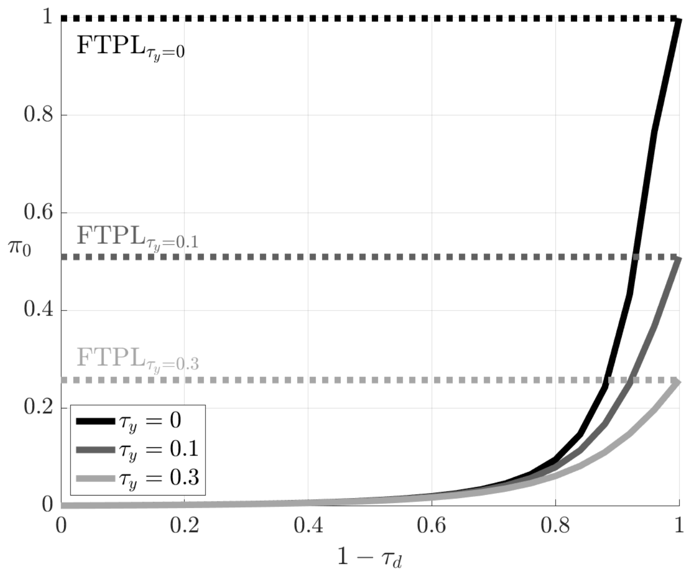

Figure 1: Date-0 inflation response to a fiscal deficit shock in HANK (solid), for different τ_d and τ_y. The dashed lines show the corresponding inflation response in the FTPL equilibrium. The size of the shock is normalized to give a date-0 FTPL inflation response of 1 per cent for $\tau_y = 0$.

HANK conclusions, distinguishing it from the fragility of the FTPL reviewed in Section 3.2. Before substantiating these claims in Section 4.3, we will first dig deeper into the dynamic time profile of the inflation response in HANK (in the remainder of this section) and then extend to our results to more general monetary policy reactions (in Section 4.2).

Front-loading. While identical in their predictions on the initial price jump and so the total amount of debt erosion, FTPL and HANK do differ in their predictions on the *timing* of the induced inflation response. To make this point precise, we consider the following measure of the front-loadedness of the inflation impulse response:

$$\pi^\dagger \equiv \frac{\pi_{\varepsilon,0}}{\sum_{k=0}^{\infty} \beta^k \pi_{\varepsilon,k}} \tag{18}$$

$\pi^\dagger$ is the initial impact relative to the cumulative inflation response. We then have the following result.

Proposition 4. *Let $\omega < 1$, $\tau_y > 0$, and $\phi = 0$. The inflation impulse response to a fiscal deficit shock in the HANK economy is more front-loaded (according to the measure (18)) the larger the departure from permanent-income behavior (i.e., the smaller ω). Furthermore, $\pi^\dagger$ in the HANK economy is bounded from below by its FTPL counterpart,*

$$\pi^\dagger > \pi^{\dagger,FTPL} = 1 - \beta, \tag{19}$$

with the distance between the two vanishing only when $\omega \to 1$.

The intuition underlying Proposition 4 is straightforward, and yet again rooted in short household horizons. Because of those short horizons, the demand boom—and thus the inflationary pressure that it causes—in HANK is necessarily short-lived. Formally, in the case of fixed real rates considered here, output and thus also inflation in RANK-FTPL follow a random walk, while in HANK the demand boom is transitory, i.e., $\rho_d \in (0,1)$.

4.2 Extension to interest rate feedback

We next extend our equivalence results to more general interest rate feedback rules.

Assumptions on policy. We now consider the more general monetary policy rule (9), restated here for convenience.

$$r_t = \phi y_t \tag{20}$$

Throughout this section, and consistent with the results in Section 4.1, we will restrict attention to the special case of no fiscal adjustment, i.e., $\tau_d = 0$. Given this fiscal feedback rule, we will characterize equilibria for $\phi \in [-1/\sigma, \bar{\phi})$, where the upper bound $\bar{\phi}$ is defined as $\bar{\phi} = \frac{\frac{(1-\beta\omega)(1-\omega)}{\omega}\tau_y}{\sigma(1-\beta)+\frac{(1-\beta\omega)(1-\omega)}{\omega}\beta\frac{D^{ss}}{Y^{ss}}}$. Beyond this upper bound, no bounded equilibrium in our HANK economy exists; at the lower bound $-1/\sigma$, the deficit-induced boom is so large that debt is stabilized immediately.

HANK meets FTPL. We characterize the equilibrium in HANK under the stated assumptions on policy, and then contrast outcomes with those of an analogous FTPL experiment. In defining that analogous experiment we need to deal with two challenges. First, for $\phi > 0$ (i.e., active monetary policy), bounded equilibria can exist in HANK (since $\bar{\phi} > 0$), but of course not in the FTPL. Second, because of the front-loading property discussed above, a common monetary feedback rule in HANK and FTPL would lead to differentially large interest rate changes and thus different amounts of monetary accommodation. A straightforward way of dealing with those two challenges and providing the natural "apples-to-apples" comparison is to first find the equilibrium in HANK, and then to study the unique FTPL equilibrium with the same path of real interest rates, and thus the same amount of monetary policy accommodation. Proposition 5 summarizes the results of this exercise.

Proposition 5. *Suppose that $\omega < 1$, $\tau_y > 0$, $\tau_d = 0$ and $\phi \in [-1/\sigma, \bar{\phi})$. There exists a unique bounded equilibrium in the HANK economy. Now consider the analogous FTPL economy with an exogenous path of (expected) real interest rates equal to its HANK counterpart. Then, the comparisons established in Propositions 3 and 4 continue to hold. That is, we still have*

$$\pi_{\varepsilon,0} = \pi_{\varepsilon,0}^{FTPL}.$$

as well as

$$\pi^{\dagger} > \pi^{\dagger,FTPL}.$$

Our two conclusions of Section 4.1—the HANK-FTPL equivalence result, and the front-loading—thus extend without any change. Intuitively, general equilibrium movements in real interest rates provide further redistribution between households and government, thus affecting the Keynesian demand loop in HANK; *conditioning* on the path of real interest rates, however, inflation and output must ultimately move by the same overall amounts in HANK and in FTPL, by exactly the same reason as discussed in the previous section.

4.3 The robustness of HANK

In this section we establish how the difference in underlying mechanism between HANK and FTPL robustifies the former against the controversies surrounding the later.

A change in policy. We begin our analysis with a change to the contemplated policy experiment. That alteration is designed to allow us to connect as transparently as possible to the core FTPL controversies reviewed in Section 3.2. Specifically, we assume that, at some future date H—possibly far ahead, but finite—the monetary authority will hike real rates to lean against inflation (turn "active"), and the fiscal authority will increase taxes to stabilize debt (turn "passive"). While this change to the policy experiment eliminates the FTPL equilibrium in RANK, it has essentially no effect on the equilibrium response of output and inflation to fiscal deficits in HANK.

Proposition 6. *Let H be finite but arbitrarily large future date. Suppose that the fiscal and monetary authorities to follow the rules (7) and (8) with $\phi = \tau_d = 0$ for $t < H$, but then switch to $r_t = \phi' y_t$ and $t_t = \tau'_d d_t$ at $t \geq H$, for arbitrary $\phi' > 0$ and $\tau'_d \in (0,1]$. Then:*

1. *In RANK ($\omega = 1$), there exists a unique equilibrium, and it has $y_t = \pi_t = 0$ for all t and all realizations of uncertainty.*

2. *In HANK ($\omega < 1$), there exist a unique equilibrium and it is such that, for any $T > 0$, as $H \to \infty$, $\{y_t, \pi_t\}_{t=0}^{T}$ converge to their counterparts in Propositions 2, for all realizations of uncertainty.*

Figure 2 plots date-0 inflation response to a date-0 fiscal deficit for different H and provides a visual illustration of Proposition 6. The left panel corresponds to the first part: for any finite H, in RANK, fiscal deficits have no effect on output and inflation. It is only with $H = \infty$—i.e., there is *never* any fiscal adjustment, and monetary policy throughout remains passive—that prices jump in response to the fiscal deficit. The right panel then corresponds to the second part: in HANK, as H is increased, fiscal deficits become more and more inflationary, eventually converging to the FTPL price jump.[13]

[13] The contrast between the two panels reflects our earlier discussion of continuity: in RANK, the effects of deficits on inflation are discontinuous at $\phi = 0$ and $\tau_d = 0$; in HANK, on the other hand, they vary continuously in ϕ and τ_d.

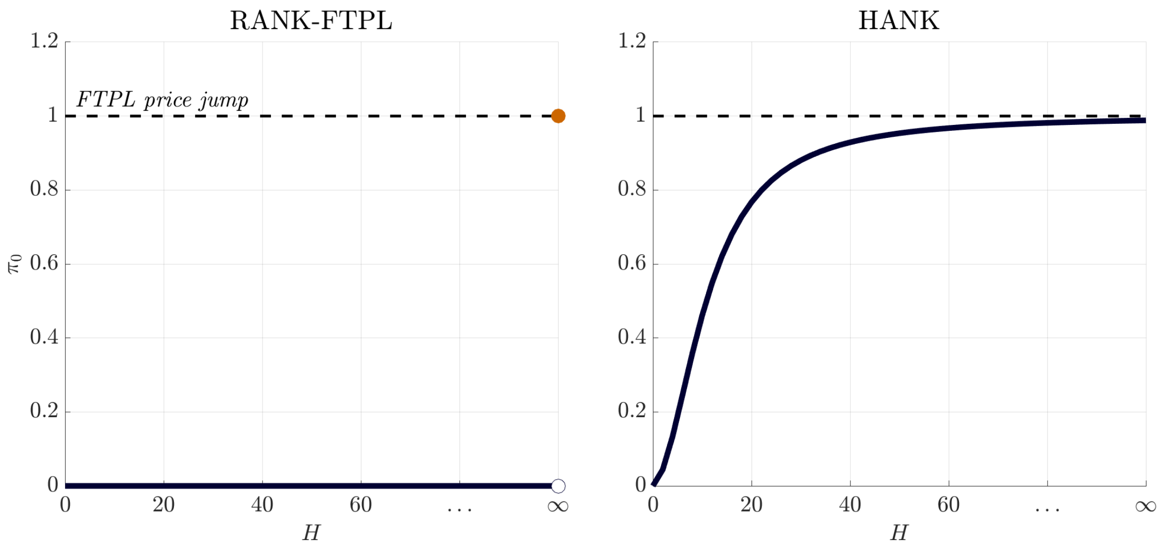

Figure 2: Date-0 inflation response to a date-0 fiscal deficit shock in RANK (left panel) and HANK (right panel) for different H. The size of the shock is normalized to give a date-0 FTPL price jump of 1 per cent.

At the heart of this result lie the finite horizons of HANK: because of those finite horizons, fiscal adjustment and a monetary reaction that occur with a significant delay are simply too late to arrest the inflationary effects of the initial fiscal deficit. In contrast, with the infinite horizons of consumers in RANK, such fiscal and monetary reactions are never too late.

Far-ahead beliefs. While Proposition 6 considered a change of far-ahead policy, it is important to note that the underlying idea is in fact more general. The key to the fragility of the FTPL equilibrium documented in part (i) of Proposition 6 is not the policy change per se, but rather the fact that this change in policy helps rationalize the belief that the economy will return to steady state at H. To see this, note that, if $y_t = 0$ at $t = H$, then iterating (10) backward delivers $y_t = 0$ for all $t \leq H$, regardless of the policies assumed before H. The same logic underlies the results of Angeletos and Lian (2023): by adding small but appropriate noise (as in the global games literature), they anchor far-ahead beliefs, returning the RANK model's "conventional" solution—here $y_t = 0$ for all t—as the unique equilibrium even when monetary policy is *always* passive.

Summary. Our analysis in this section has revealed that HANK can replicate the FTPL's core predictions about inflation and debt erosion, all while completely sidestepping its theoretical controversies. The reason is simple: the FTPL assumes Ricardian households, thus requiring delicate assumptions for fiscal deficits to nevertheless matter, and for Ricardian equivalence to fail. This fragility is absent in HANK simply because households there are non-Ricardian—given their short horizons, deficits already meaningfully drive inflation if fiscal adjustment is only delayed, rather than entirely absent.

In the remainder of the paper we make two additional points. First, in Section 5, we document how our conclusions extend to richer model environments. Second, in Section 6, we provide a quantitative analysis of the deficit-inflation mapping, in particular relating our results to the simple FTPL arithmetic that has loomed large so far.

5 Extensions

We now discuss how our results on the deficit-inflation mapping in FTPL and in HANK extend to three important model extensions: long-term government debt in Section 5.1; heterogeneous household bond holdings as well as transfer receipts in Section 5.2; and a hybrid NKPC in Section 5.3. All of these extensions will feature prominently in our quantitative analysis in Section 6.

5.1 Long-term government debt

We allow for government debt to be long-term. In keeping with much of the FTPL literature, we consider the analytically tractable case of a geometric maturity structure (e.g., Cochrane, 2001). For simplicity we furthermore restrict attention to the special case of fixed (expected) real rates, with the extension to alternative monetary feedback relegated to Appendix B.

Environment. Nominal public debt is long-term, with its maturity parameterized by $\delta \in [0,1]$; the baseline case of short-term debt is nested as $\delta = 0$. The government flow budget then becomes

$$d_{t+1} = \frac{1}{\beta}(d_t - t_t) + \frac{D^{ss}}{Y^{ss}} r_t - \underbrace{\frac{D^{ss}}{Y^{ss}}\left(\pi^{\delta}_{t+1} - \mathbb{E}_t\left[\pi^{\delta}_{t+1}\right]\right)}_{\text{debt erosion due to inflation surprise}} - \underbrace{\frac{D^{ss}}{Y^{ss}}\left(r^{\delta}_{t+1} - \mathbb{E}_{t-1}\left[r^{\delta}_{t+1}\right]\right)}_{\text{debt revaluation due to real rate surprises}} \tag{21}$$

where

$$\pi^{\delta}_t \equiv \mathbb{E}_t\left[\sum_{k=0}^{\infty}(\beta\delta)^k \pi_{t+k}\right] \quad \text{and} \quad r^{\delta}_t \equiv \mathbb{E}_t\left[\sum_{k=0}^{\infty}(\beta\delta)^{k+1} r_{t+k}\right]. \tag{22}$$

We refer the reader to Appendix A for a detailed derivation of (21). The remainder of the model is exactly as in Section 2.

HANK meets FTPL. The economy with long-term debt features a weaker form of equivalence between inflation outcomes in HANK and FTPL, as summarized in Proposition 7. The comparison now concerns the maturity-adjusted *cumulative* inflation response $\text{NPV}^{\delta}_{\pi} \equiv \frac{d\mathbb{E}_t[\pi^{\delta}_t]}{d\epsilon_t}$—i.e., the summary statistic of the impact of inflation surprise on the government budget in (21). This object has received much attention in the FTPL literature (e.g., see Barro and Bianchi, 2023), because it is the object for which the FTPL makes the starkest prediction. We also note that, for the empirically relevant case of δ close to 1, it is very similar to the cumulative inflation impulse response (discounted by β).

Proposition 7. *Let $\omega < 1$, $\tau_y > 0$, $\tau_d = 0$, and $\phi = 0$. There exists a unique bounded equilibrium in the HANK economy. The quantity NPV_π^δ, which measures the degree of debt erosion or, equivalently, the maturity-adjusted cumulative inflation response to a fiscal deficit, is strictly lower in the HANK economy than its FTPL counterpart:*

$$NPV_\pi^\delta < NPV_\pi^{\delta,FTPL}, \tag{23}$$

with the distance between the two vanishing when $\tau_y \to 0$ or $\kappa \to \infty$.

We see that, if the tax base channel is weak (because either $\tau_y \to 0$ or $\kappa \to \infty$), then HANK and the FTPL again make the exact same prediction—there is just enough inflation to finance the initial fiscal deficit, i.e., the textbook FTPL arithmetic. If instead the tax base channel is operative, then the inflation response in HANK is *strictly below* its FTPL counterpart. The intuition is simple, reflecting the interaction of long-term debt with the inflation front-loading implied by HANK.[14] Since inflation is at all dates proportional to the present discounted value of future output responses, making any given output boom more front-loaded (while holding its present value fixed) will leave the impact inflation unchanged, but lower the subsequent inflation responses. This explains why, with $\delta > 0$, the debt erosion effect—and thus the overall cumulative inflation response—is smaller.

While Proposition 7 focuses on cumulative inflation, its proof in Appendix D provides a full characterization of equilibrium transition paths. These characterizations turn out to be almost identical to those discussed earlier for short-term debt in Proposition 2; the only effect of moving to long-term debt is to *re-scale* those impulse responses. Intuitively, changing debt maturity affects the amount of redistribution from households to the government associated with a given path of inflation, but otherwise does not alter the intertemporal Keynesian cross underlying the equilibrium characterization.

5.2 Heterogeneous distributional incidence

While analytically convenient as a device to break permanent-income consumption-savings behavior, the OLG model of Section 2 is not quite rich enough to serve as a quantitatively relevant, HANK-like model of consumer demand, for at least two reasons. First, the implied time profile of intertemporal MPCs does not mimic that of either the data or of richer HANK models; specifically, iMPCs decline too gradually (e.g., see Wolf, 2021). Second, the surprise inflation burst associated with fiscal deficits may redistribute between households with heterogeneous (nominal) bond positions. If bond hold-

[14]In the interest of parsimony, Proposition 7 does not repeat our front-loading results for the long-term debt economy. The fact that this property extends, however, is immediate from Proposition 7: we have that $\pi^\dagger$ defined in (18) satisfies

$$\pi^\dagger > \pi^{\dagger,FTPL}.$$

ings correlate with MPCs, then such redistribution will matter for aggregate demand. A single-type OLG model abstracts from such considerations. To address these shortcomings, we here consider a richer—but still tractable—model of aggregate demand.

Environment. We will study a "hybrid" model that combines our baseline OLG block with a margin of hand-to-mouth spenders, with $\mu \in (0,1)$ denoting the share of spenders. By Auclert et al. (2023) and Wolf (2021), such models can fit empirically observed and HANK-implied patterns of iMPCs relatively well; and since spenders do not hold any bonds, such a model can also capture—if in a crude way—redistributional effects related to (nominal) bond devaluation.

It is straightforward to establish that, in this generalized model, the aggregate demand relation in (1) generalizes to

$$y_t = \left(1-\beta\omega\right) d_t + \left(\mu + (1-\mu)(1-\beta\omega)\right)\left(\left(y_t - t_t\right) + \frac{(1-\mu)(1-\beta\omega)}{\mu+(1-\mu)(1-\beta\omega)}\mathbb{E}_t\left[\sum_{k=1}^{\infty}\left(\beta\omega\right)^k\left(y_{t+k}-t_{t+k}\right)\right]\right), \tag{24}$$

where we have for simplicity already imposed the assumption of a neutral monetary policy, i.e. $\phi = 0$ in (9). The remainder of the model is as in Section 2, though we will here also allow for long-term debt.

HANK meets FTPL. Even in this generalized model variant we continue to obtain similar comparison results between HANK and FTPL; there is exact equivalence when $\delta = 0$, and takes the form of an upper bound when $\delta > 0$, mirroring Propositions 3 and 7.

Proposition 8. *Let $\omega < 1$, $\tau_y > 0$, $\tau_d = 0$, $\phi = 0$, and $\mu \in (0,1)$. There exists a unique bounded equilibrium, and it has the following properties.*

1. *If $\delta = 0$, then the initial price jump in response to a fiscal deficit shock is exactly the same as its counterparts in our baseline HANK economy and FTPL:*

$$\pi_{\varepsilon,0} = \pi_{\varepsilon,0}^{Baseline} = \pi_{\varepsilon,0}^{FTPL},$$

   where $\pi_{\varepsilon,0}^{Baseline}$ is the inflation response in the baseline HANK economy with $\mu = 0$.

2. *If $\delta > 0$, then the cumulative inflation response to deficit a fiscal deficit shock, NPV_π^δ, decreases with the share of spenders μ, and is bounded above by its analogue in our baseline economy and hence by FTPL:*

$$NPV_\pi^\delta < NPV_\pi^{\delta,Baseline} < NPV_\pi^{\delta,FTPL}, \tag{25}$$

   where $NPV_\pi^{\delta,Baseline}$ is the maturity adjusted-cumulative inflation response in our baseline economy, with the distance between the three vanishing when $\tau_y \to 0$ or $\kappa \to \infty$.

In the generalized economy with heterogenous distributional incidence, fiscal deficits—through the inflation that they induce—redistribute from low-MPC households with large (nominal) bond holdings to high-MPC, hand-to-mouth households. This additional impetus to demand then front-loads the deficit-induced demand boom even more. If $\delta = 0$, then this additional front-loading is immaterial for the deficit-induced price jump; if instead debt is long-term ($\delta > 0$), then the front-loading further dampens inflationary pressures, by exactly the same reasoning as that discussed in Section 5.1.[15] Our quantitative analysis in Section 6 will consider an aggregate demand block that captures the key intuitions of—but, in terms of model richness, goes even further beyond—the two-type spender-saver construction considered here.

5.3 Hybrid NKPC

Our analysis so far featured a textbook forward-looking NKPC. We here ask how our results change for a more empirically relevant, hybrid NKPC, where price-setting is less forward-looking, and thus more consistent with empirical evidence on the sluggishness of inflation.

Environment. We replace (2) with the empirically more relevant Hybrid NKPC,

$$\pi_t = \kappa y_t + \xi \beta \pi_{t-1} + (1-\xi)\,\beta \mathbb{E}_t\left[\pi_{t+1}\right], \tag{26}$$

where $\xi \in [0,1]$ parameterizes the degree of backward-lookingness in price-setting. The remainder of the model is exactly as in Section 2; in particular, we will restrict attention to the case of short-term government debt, for reasons that will become clear shortly.

HANK meets FTPL. With a hybrid NKPC, exact equivalence between HANK and FTPL as before continues to obtain in the absence of the tax base channel. If it is present, then the interaction of HANK front-loading with the backward-looking shape of the hybrid NKPC now means that short-run inflationary pressures in HANK are *larger* than in FTPL, exactly the opposite of the long-term government debt effect studied above. Proposition 9 summarizes our insights.

Proposition 9. *Let $\omega < 1$, $\tau_y > 0$, $\tau_d = 0$, $\delta = 0$, $\phi = 0$ and $\xi \in (0,1]$ in the hybrid NKPC (26). There exists a unique bounded equilibrium in the HANK economy. The initial price jump in response to a fiscal deficit shock is strictly higher than the FTPL counterpart with the same hybrid NKPC:*

$$\pi_{\varepsilon,0} > \pi_{\varepsilon,0}^{FTPL},$$

with the distance between the two vanishing when $\tau_y \to 0$ or $\kappa \to \infty$.

[15] Proposition 8 omits a formal front-loading statement for expositional fluidity, but it can indeed be shown that, as expected, we have

$$\pi^{\dagger} > \pi^{\dagger,\text{Baseline}} > \pi^{\dagger,FTPL},$$

where $\pi^{\dagger,\text{Baseline}}$ is the degree of front-loadedness of the inflation response in our baseline economy.

The intuition is as follows. Compared to the textbook NKPC, its hybrid generalization (26) is less forward-looking, so current inflation depends more heavily on output in the immediate future. Since the output boom in HANK is front-loaded, this means that the initial inflation increase is larger. Note that this result assumes short-term government debt; if instead $\delta > 0$, then there would be two offsetting effects, with unclear sign. Our quantitative analysis in Section 6 will combine all of the model ingredients considered here to shed light on the empirically relevant deficit-inflation mapping.

6 Quantitative analysis

We now complement our theoretical results with a quantitative analysis of the deficit-inflation nexus in HANK. Our results so far suggest that even the predictions of the textbook extreme version of the FTPL—in which current deficits are financed entirely through a commensurate jump in prices—can emerge in HANK economies. The main takeaway of this section, however, is that, in practice, deficits are likely to be much less inflationary than predicted by the simple FTPL arithmetic.

To this end, we study the mapping from deficits to inflation in a version of our HANK model that is disciplined through direct evidence on the key ingredients of our theory. Section 6.1 describes model and calibration, Sections 6.2 and 6.3 contain the main results, and Section 6.4 closes with an application to post-covid inflation dynamics.

6.1 HANK model and calibration

We consider a variant of the model environment in Section 2, with three additions, following our discussion in Section 5. First, government debt is now long-term. Second, we allow for moderate household heterogeneity, with three types of households i, indexed by heterogeneous survival probabilities ω_i. This extension will allow the model to be simultaneously consistent with empirical evidence on (i) intertemporal marginal propensities to consume (e.g., Auclert et al., 2023) as well as (ii) household wealth holdings and transfer receipts. Third, we consider a hybrid NKPC, yielding more realistic inflation dynamics. The remainder of this section presents calibration details for all model blocks, with a summary provided in Table 1.

Throughout this section, and as in Sections 2 - 5, the policy experiment that we will consider is a one-off, surprise fiscal deficit increase at date 0, equal to one per cent of steady-state GDP.

Consumers. We extend the consumer block of Section 2.1 to allow for three types of households i, with respective population shares χ_i. Households differ in their death probabilities ω_i—or, less literally, in their probably of being subject to a binding borrowing constraint—, steady-state wealth holdings D_i^{SS}, and exposure to fiscal deficit shocks (i.e., transfer receipts). We choose population shares

Parameter	Description	Value	Target
Demand Block			
χ_i	Population shares	$\{0.218, 0.629, 0.153\}$	Fagereng et al.
ω_i	Survival rates	$\{0.972, 0.833, 0\}$	Fagereng et al.
D_i^{SS}	Wealth shares	$\{0.6, 0.4, 0\} \times D^{SS}$	See text
ε_i	Transfer receipt	$\{0.122, 0.706, 0.172\} \times \varepsilon$	See text
σ	EIS	1	Standard
β	Discount factor	0.998	Annual real rate
Supply Block			
κ	Slope of Hybrid NKPC	$\{0.006, 0.019, 0.056\}$	Hazell et al.; Cerrato and Gitti
ξ	Backward-lookingness	0.288	Barnichon and Mesters
Policy			
τ_y	Tax rate	0.33	Average Labor Tax
D^{ss}/Y^{ss}	Gov't debt level	1.04	Liq. wealth holdings
δ	Gov't debt maturity	0.95	Av'g debt maturity
τ_d	Tax feedback	0	Anderson and Leeper
ϕ	Inflation feedback	0	See text

Table 1: Quantitative model, calibration.

and death probabilities to match empirical evidence on average intertemporal marginal propensities to consume, from Fagereng et al. (2021). Wealth shares are set to roughly replicate the skewness of the U.S. wealth distribution (e.g., see Kaplan et al., 2018), with the bottom 15 per cent holding no wealth, and the top quantile holding 60 per cent of all wealth. Finally, consistent with U.S. policy practice, transfer receipts are somewhat more concentrated at the bottom. We also set $\sigma = 1$ (giving log preferences), and back out β to hit a steady-state real rate of interest of one per cent (annual).

Our alternative model variants in Section 6.3 will consider several alternative assumptions on the departure from Ricardian equivalence (the ω's), wealth holdings, and transfer receipts.

Nominal rigidities. We assume a hybrid NKPC, as discussed in Section 5.3. For the slope κ we consider three headline values: the shallow slope estimated by Hazell et al. (2022); a three-times steepening of that NKPC, as estimated in Cerrato and Gitti (2022) for the post-covid inflationary period; and yet another three-times steepening, as a plausible (if ad-hoc) upper bound. For the backward-forward split (ξ vs. $1-\xi$), we take the headline point estimates of Barnichon and Mesters (2020).

In our main quantitative analysis we will furthermore report results for an entire (and wide) range of κ's. The alternative model variants studied in Section 6.2 will also feature alternative assumptions

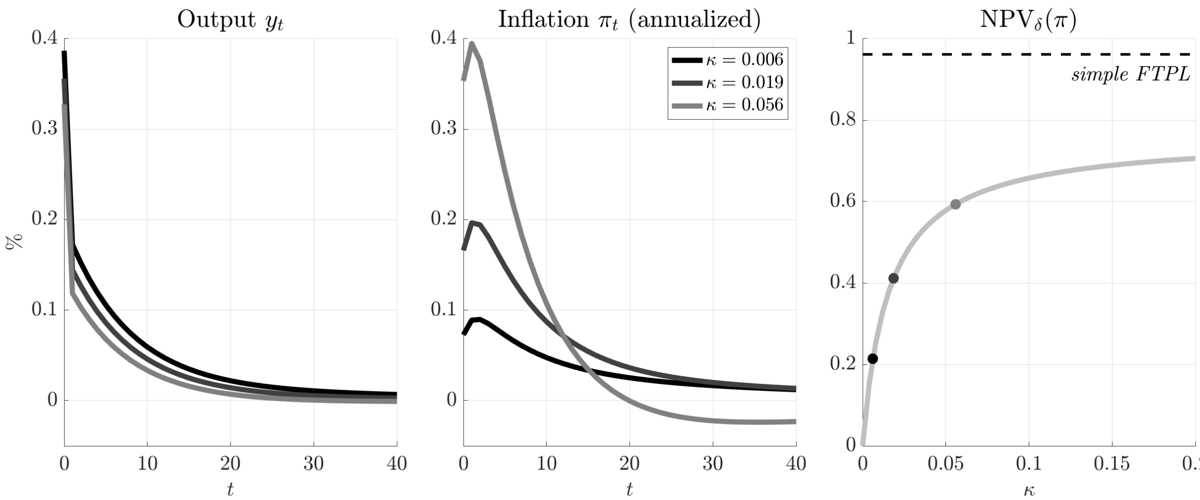

Figure 3: Output and inflation impulse responses to a date-0, 1% deficit shock for different values of κ (left and middle), and NPV^δ_π as a function of κ (right).

on the backward-forward split in the NKPC.

Policy. We set $\tau_y = 0.33$, implying meaningful—and empirically realistic—feedback from economic activity to primary surpluses. The total amount of government debt is set as in Kaplan et al. (2018), and we set δ to match an average debt maturity of five years. Consistent with legislative evidence on the post-covid fiscal stimulus (e.g., see the detailed discussion Anderson and Leeper, 2023), we consider an "unbacked" fiscal expansion, so $\tau_d = 0$. Finally, as in our main analysis, we set $\phi = 0$, corresponding again to fixed real-rate rule. We do so for two reasons. First, in that case, our simulations will be informative about the pure effect of the deficit, without any direct monetary offset or accommodation. Second, as discussed in Angeletos et al. (2024), this case is actually a reasonable approximation to many past fiscal stimulus episodes.

For our alternative model variants and the quantitative post-covid application we will pay particular attention to what happens under alternative assumptions on fiscal adjustment (τ_d) and on the monetary policy reaction (ϕ).

6.2 Benchmark model

We now study how, in our quantitative HANK environment, fiscal deficits transmit to inflation. Figure 3 shows impulse responses of aggregate output and inflation to the deficit shock (left and middle panel), for our three headline values of κ (shades of grey). The right panel then displays the cumulative (maturity-adjusted) inflation response NPV^δ_π as a function of κ, over a large range.

The main takeaway from the figure is that the inflationary pressures associated with the unfunded fiscal deficit shock are—while material—quite substantially weaker than predicted by the simplest textbook FTPL arithmetic, with the cumulative inflation response at least cut in half. According to the simple FTPL arithmetic, the one per cent deficit shock would translate to cumulative inflation of just below 1 (equal to the inverse of the debt-to-GDP ratio of 1.04). The right panel of Figure 3 reveals that, even for an NKPC three-times as steep as the pre-covid estimates of Hazell et al. (2022), the cumulative inflation response is actually less than half that. The left panel provides the (mechanical) answer: output booms, with a cumulative multiplier around 1.6, thus generating meaningful tax revenue through the feedback from economic activity to primary surpluses. Finally, the middle panel shows the time profile of the inflation response: consistent with our theoretical results, the inflation that does occur is front-loaded and relatively short-lived, with around a quarter of the entire inflation response already occurring over the first year. Given that government debt is long-term, this front-loading—which is further reinforced by the fiscal shock's distributional incidence—then further contributes to the dampening of the overall inflation response.

The remainder of this section extends our analysis in two ways. First, in Section 6.3, we go beyond the benchmark model parameterization and explore the effects of various possible model alterations. Second, in Section 6.4, we discuss implications of our results for the post-covid inflationary episode.

6.3 Alternative model variants

We now study the deficit-inflation mapping in several alternative variants of our quantitative model, allowing us to shed light both on the broader relevance of our conclusions as well as on the role played by the various model ingredients. Details for all variants are provided in Appendix C.

- *Consumers.* For a first set of experiments, we alter our empirically disciplined consumer block to feature no heterogeneity in bond holdings and dividend receipts ("iMPC"), heterogeneity only in bond holdings ("Het. B"), and heterogeneity only in transfer receipts ("Target"). Second, we consider what happens if households are instead behavioral, with a sticky information friction as in Auclert et al. (2020) ("Behavioral"). Finally, for a third set of experiments, we replace our consumer block by the one-type OLG structure of Section 2 ("OLG") and by a full-blown HANK structure ("HANK").

- *Nominal rigidities.* Our analysis in Section 6.2 already shed light on the role of NKPC κ. We here additionally consider what happens if our empirically disciplined hybrid NKPC is replaced by a simple textbook forward-looking one ("f-NKPC").

- *Policy.* To further illustrate our theoretical "robustness" discussion, we also investigate what happens with active monetary policy ("Active MP", $\phi = 0.5$) and with gradual (rather than ab-

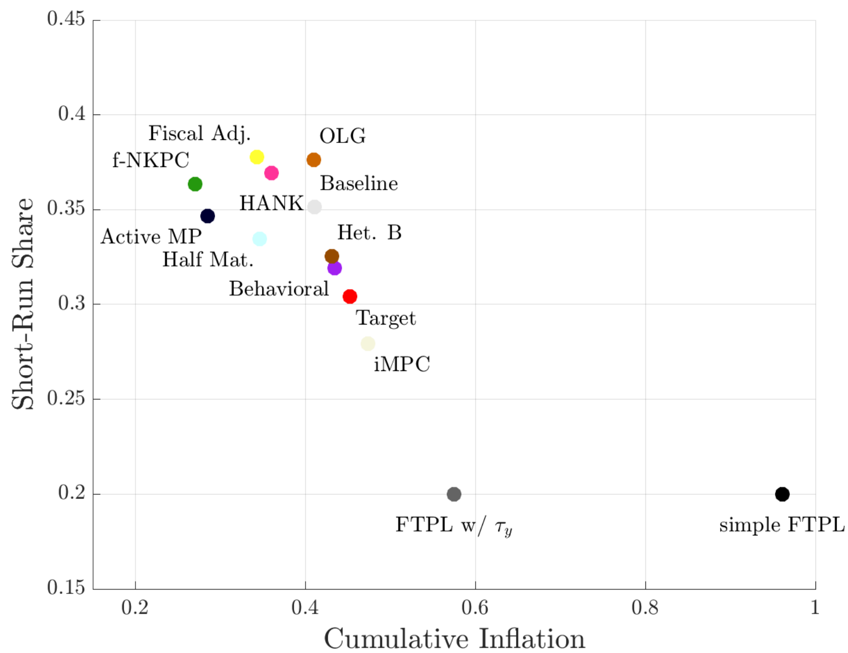

Figure 4: Cumulative inflation response and short-run response share to a date-0, 1% deficit shock for different model variants, indicated by dots.

sent) fiscal adjustment ("Fiscal Adjustment", $\tau_d = 0.02$). We also consider a model variant in which the average government debt maturity is halved ("Half Mat.", $\delta = 0.9$).

Our results are reported in Figures 4 and 5. Figure 4 shows the cumulative (maturity-adjusted) inflation response $\text{NPV}_{\pi}^{\delta}$ (x-axis) as well as the short-run inflation share (defined as the share of total inflation in the first year, relative to the first five years, y-axis) for the simple FTPL (black), a variant of the FTPL with tax base self-financing (dark grey), our baseline quantitative model (grey), as well as all other model variants (colored). The textbook simple FTPL arithmetic is in the bottom right of the figure: the cumulative inflation response is equal to the inverse of the debt-to-GDP ratio of 1.04, and inflation is highly persistent. Adding tax base self-financing does not affect the persistence of the inflation burst, but dampens its magnitude. Next, and consistent with the discussion in Section 6.2, our quantitative model is instead in the top left: cumulative inflation is now much smaller, and the short-run inflation share is larger. In particular, the further dampening of cumulative inflation relative to the FTPL with tax feedback reflects the interaction of long-term debt and inflation front-loading. Finally, we see that *all other* model dots are also in the top left: changing details of the model parameterization of course affects precise magnitudes, but does not change the core insight that inflation responses are smaller and more front-loaded than in the simplest FTPL benchmark.

Figure 5 shows full impulse responses for selected model variants, allowing us to dig deeper into

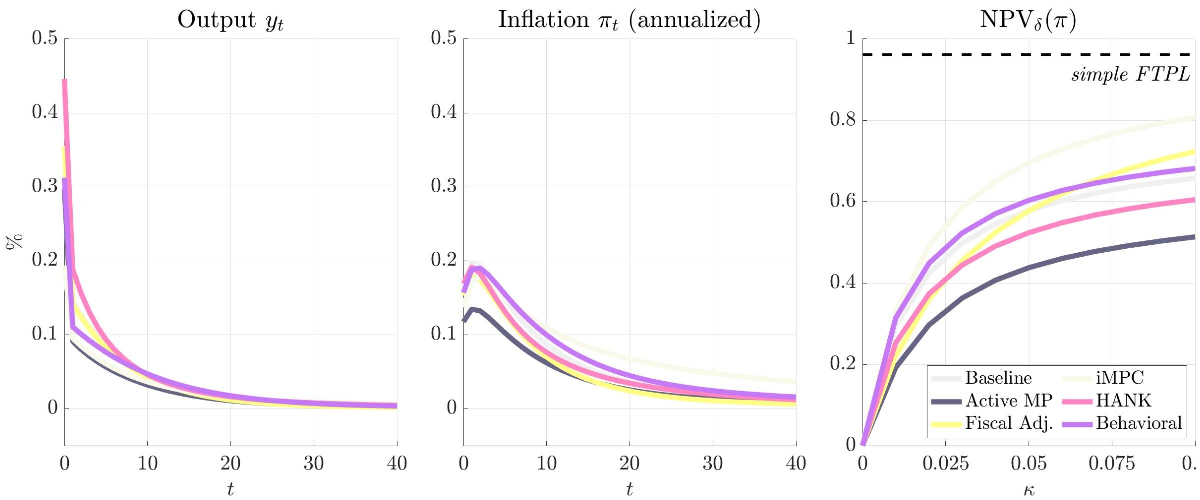

Figure 5: Output and inflation impulse responses to a date-0, 1% deficit shock (left and middle) and NPV_π^δ (right) for different model variants.

the role played by the various model alterations. First, with more aggressive monetary policy or faster fiscal adjustment, the inflation response is—as expected—dampened, but of course remains present, illustrating our theoretical results on the robustness of the deficits-inflation mapping in HANK-type models. Second, consistent with the discussion in Section 5.2, homogeneous wealth holdings and transfer incidence would lead to a larger inflation response. Third, in the less forward-looking behavioral model, the intertemporal Keynesian cross underlying the deficit-inflation mechanism in HANK plays out more slowly, and so the inflation spike is more delayed. And fourth, moving to a full-blown quantitative HANK model has very limited effect on our results, consistent with prior work establishing that analytical models of the sort provided here provide an excellent approximation to aggregate output and inflation dynamics in HANK.

Finally we also note that, while the results in Figure 4 assume a fixed real rate (except, of course, for the model variant with an active monetary rule), our results do not hinge on that assumption. Specifically, we in Appendix C repeat our analysis for a fiscal stimulus that is accompanied by monetary accommodation. In that case, the standard FTPL also features a front-loaded inflation response, since the real rate cut encourages households to front-load consumption. Crucially, however, in our model variants, and for the same real rate path, the inflation response is *even more* front-loaded, thus overall delivering the same picture as in Figure 4.

6.4 Application to post-covid inflation dynamics

We now use our quantitative model for an application to post-covid inflation dynamics. Results are reported in Figure 6, which shows output and inflation impulse responses as well as the discounted

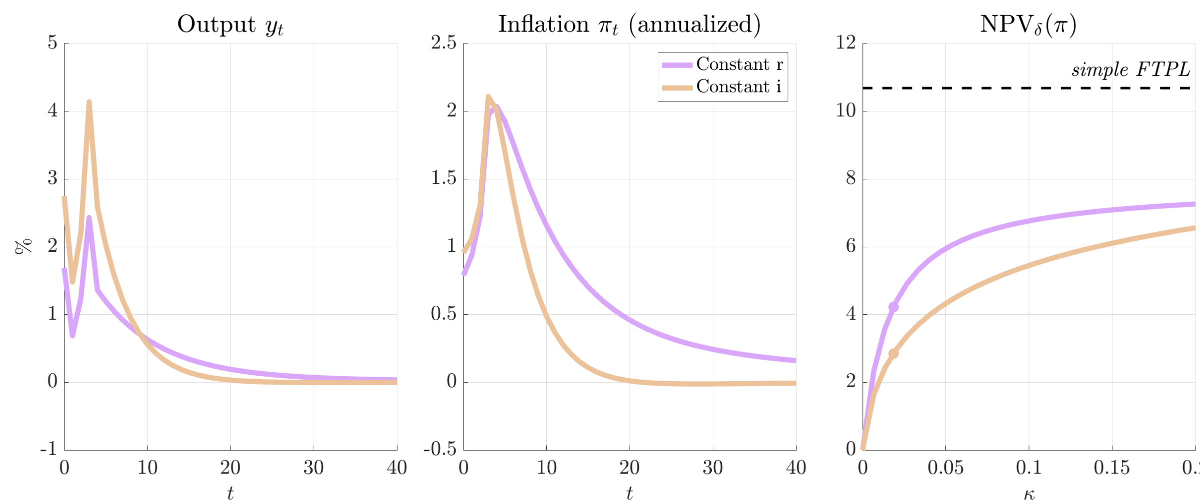

Figure 6: Output and inflation impulse responses to a date-0, \$2.35tr deficit shock (left and middle) and NPV$_\pi^\delta$ (right) in the baseline quantitative model, under two different assumptions on the monetary policy reaction: fixed real rates (purple) and fixed nominal rates (orange).

cumulative inflation response under different assumptions on policy.

Policy experiments. For all policy experiments, we consider a deficit shock of size \$2.35tr, corresponding to the share of total covid-related federal spending sent to households (e.g., see Table 2 in Anderson and Leeper, 2023). We furthermore assume that there is no fiscal adjustment ($\tau_d = 0$), consistent with actual legislation so far (e.g., see the review in Anderson and Leeper, 2023).

We then study impulse responses to the fiscal deficit shock under two different assumptions on the monetary policy reaction. First, we keep real rates fixed. The resulting impulse responses will identify the causal effect of the fiscal expansion *in isolation*; i.e., what is the incremental impetus to inflation, keeping the monetary policy stance—in terms of real rates—as observed in the data? Second, we keep nominal rates fixed. This counterfactual keeps the monetary stance in policy instrument space as in the data, and thus—since the fiscal deficit will be inflationary—embeds the effects of additional monetary accommodation.

Results. The results from our policy experiments are reported as the purple and orange lines in Figure 6. Consider first the overall magnitudes. With an initial debt-to-GDP ratio slightly above 1, and a fiscal stimulus of around 11% of GDP, the simple textbook FTPL accounting would predict a cumulative discounted inflation response just below 11%. We see that both policy experiments in our setting predict material dampening relative to that upper bound, consistent with our results in Sections 6.2 - 6.3. The burst in inflation is furthermore, in both cases, concentrated in the first couple of years after the fiscal deficit shock.

We next investigate further the role of the monetary policy response by contrasting the two sets of impulse responses. The counterfactual of nominal interest rates kept as in the data corresponds to additional monetary accommodation, and thus leads to a larger boom and, at least at the beginning, somewhat larger inflation. However, the implied real rate cut alleviates the budgetary shortfall, thus easing demand pressures in the medium run and thereby leading to *less* inflation. This finding is the flip-side of the classical "stepping-on-a-rake" effect of monetary contractions that are not accompanied by fiscal adjustment, as studied, for example, in Sims (2011).

7 Conclusion

Do deficits drive inflation? If so, by how much? We have addressed these questions in the New Keynesian framework, comparing and contrasting HANK and FTPL. *Despite* their very different underlying economic mechanisms, these two theories can agree—not just qualitatively, but quantitatively—on the inflationary effects of fiscal deficits. *Because* of the difference in mechanism, however, these common predictions are, in HANK, much more robust to extraneous and hard-to-test assumptions about far-ahead policy and beliefs. Our main quantitative contribution is then to work out just how inflationary fiscal deficits are actually likely to be in practice. We here find meaningful dampening relative to the textbook FTPL arithmetic, which stipulates that prices jump enough to fully finance any shock to the deficit. In our preferred quantitative environment, deficits are inflationary, but only by a factor of around half of that simple arithmetic.

Our analysis suggests at least two avenues for future research. First, our quantification of the empirically relevant deficit-inflation mapping is model-based, with empirical discipline provided indirectly in the form of evidence on individual model ingredients. It would be interesting to confront our predictions with more direct empirical evidence, e.g., as in Hazell and Hobler (2024). Second, the HANK perspective on how deficits drive inflation is likely to have very different normative implications for optimal fiscal debt management than the FTPL.

References

ANDERSON, J., AND E. M. LEEPER (2023): “A Fiscal Accounting of COVID Inflation,” Working Paper.

ANGELETOS, G.-M., AND Z. HUO (2021): “Myopia and Anchoring,” *American Economic Review*, 111, 1166–1200, 10.1257/aer.20191436.

ANGELETOS, G.-M., AND C. LIAN (2023): “Determinacy without the Taylor Principle,” *Journal of Political Economy*.

ANGELETOS, G.-M., C. LIAN, AND C. K. WOLF (2024): “Can Deficits Finance Themselves?”, Working Paper.

ATKESON, A., V. V. CHARI, AND P. J. KEHOE (2010): “Sophisticated monetary policies,” *Quarterly Journal of Economics*, 125, 47–89.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2020): “Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model.”

——— (2023): “The intertemporal keynesian cross.”

BARNICHON, R., AND G. MESTERS (2020): “Identifying modern macro equations with old shocks,” *The Quarterly Journal of Economics*, 135, 2255–2298.

BARRO, R. J. (1974): “Are government bonds net wealth?” *Journal of political economy*, 82, 1095–1117.

BARRO, R. J., AND F. BIANCHI (2023): “Fiscal Influences on Inflation in OECD Countries, 2020-2022,”Technical report, National Bureau of Economic Research.

BASSETTO, M. (2002): “A game–theoretic view of the fiscal theory of the price level,” *Econometrica*, 70, 2167–2195.

BIANCHI, F., R. FACCINI, AND L. MELOSI (2023): “A Fiscal Theory of Persistent Inflation,” *The Quarterly Journal of Economics*, qjad027.

BIANCHI, F., AND C. ILUT (2017): “Monetary/Fiscal Policy Mix and Agent’s Beliefs,” *Review of Economic Dynamics*, 26, 113–139.

BILBIIE, F. O. (2020): “The new keynesian cross,” *Journal of Monetary Economics*, 114, 90–108.

BLANCHARD, O. (2021): “In defense of concerns over the $1.9 trillion relief plan,” *Peterson Institute for International Economics, February*, 18, 2021.

BLANCHARD, O. J. (1985): “Debt, deficits, and finite horizons,” *Journal of political economy*, 93, 223–247.

BLANCHARD, O. J., AND B. S. BERNANKE (2023): “What caused the US pandemic-era inflation?”Technical report, National Bureau of Economic Research.

BUITER, W. H. (2002): “The Fiscal Theory of the Price Level: A Critique,” *The Economic Journal*, 112, 459–480.

CERRATO, A., AND G. GITTI (2022): “Inflation since covid: Demand or supply,” *Available at SSRN 4193594*.

COCHRANE, J. (2023): “The fiscal theory of the price level,” in *The Fiscal Theory of the Price Level*: Princeton University Press.

COCHRANE, J. H. (2001): “Long-term debt and optimal policy in the fiscal theory of the price level,” *Econometrica*, 69, 69–116.

——— (2005): “Money as stock,” *Journal of Monetary Economics*, 52, 501–528, https://doi.org/10.1016/j.jmoneco.2004.07.004.

——— (2011): “Determinacy and identification with Taylor rules,” *Journal of Political economy*, 119, 565–615.

——— (2017): “The new-Keynesian liquidity trap,” *Journal of Monetary Economics*, 92, 47–63.

——— (2018): “Michelson-Morley, Fisher, and Occam: The radical implications of stable quiet inflation at the zero bound,” *NBER Macroeconomics Annual*, 32, 113–226.

DEL NEGRO, M., M. GIANNONI, AND C. PATTERSON (2015): “The Forward Guidance Puzzle,” *FRB of New York mimeo*.

EUSEPI, S., AND B. PRESTON (2018): “Fiscal foundations of inflation: imperfect knowledge,” *American Economic Review*, 108, 2551–89.

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2021): “MPC heterogeneity and household balance sheets,” *American Economic Journal: Macroeconomics*, 13, 1–54.

FARHI, E., AND I. WERNING (2019): “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 109, 3887–3928.

GALÍ, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*: Princeton University Press.

GALÍ, J., J. D. LÓPEZ-SALIDO, AND J. VALLÉS (2007): “Understanding the effects of government spending on consumption,” *Journal of the european economic association*, 5, 227–270.

HAGEDORN, M. (2024): “The failed theory of the price level.”

HAGEDORN, M., I. MANOVSKII, AND K. MITMAN (2019): “The fiscal multiplier,” *NBER Working Paper*.

HAZELL, J., J. HERRENO, E. NAKAMURA, AND J. STEINSSON (2022): “The slope of the Phillips Curve: evidence from US states,” *The Quarterly Journal of Economics*, 137, 1299–1344.

HAZELL, J., AND S. HOBLER (2024): “Do Deficits Cause Inflation? A High Frequency Narrative Approach,”Technical report, Technical report.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary policy according to HANK,” *American Economic Review*, 108, 697–743.

KAPLAN, G., G. NIKOLAKOUDIS, AND G. L. VIOLANTE (2023): “Price Level and Inflation Dynamics in Heterogeneous Agent Economies,” *NBER Working Paper*.

KING, R. G. (2000): “The new IS-LM model: language, logic, and limits,” *Federal Reserve Bank of Richmond Economic Quarterly*, 86, 45–104.

KOCHERLAKOTA, N., AND C. PHELAN (1999): “Explaining the fiscal theory of the price level,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 23, 14–23.

LEEPER, E. M. (1991): “Equilibria under "active" and "passive" monetary and fiscal policies,” *Journal of monetary Economics*, 27, 129–147.

NIEPELT, D. (2004): “The Fiscal Myth of the Price Level,” *Quarterly Journal of Economics*, 119, 277–300, 10.1162/003355304772839597.

SARGENT, T. J., AND N. WALLACE (1975): “"Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule,” *Journal of Political Economy*, 83, 241–254, https://doi.org/10.1086/260321.

SIMS, C. A. (1994): “A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy,” *Economic theory*, 4, 381–399.

——— (2011): “Stepping on a rake: The role of fiscal policy in the inflation of the 1970s,” *European Economic Review*, 55, 48–56.

SUMMERS, L. (2021): “The Biden stimulus is admirably ambitious. But it brings some big risks, too,” *Washington Post*, 4, 2021.

WOLF, C. K. (2021): “Interest rate cuts vs. stimulus payments: An equivalence result,” *NBER Working Paper*.

WOODFORD, M. (1995): “Price-level determinacy without control of a monetary aggregate,” in *Carnegie-Rochester conference series on public policy* Volume 43, 1–46, Elsevier.

——— (2003): “Interest and prices.”

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

Online Appendix for
Deficits and Inflation: HANK meets FTPL

This online appendix contains supplemental material for the article "Deficits and Inflation: HANK meets FTPL".

A Supplementary details

[Under construction]

B Additional theoretical results

[Under construction]

C Additional quantitative results

[Under construction]

D Proofs

[Under construction]