Tax on Inflation Policy at the Zero Lower Bound*

Damien Capelle and Yang Liu

January 7, 2026

Abstract

This paper assesses the effectiveness and robustness of a Tax on Inflation Policy (TIP) in increasing welfare at the Zero Lower Bound (ZLB). When the ZLB arises from a decline in the neutral rate of interest, a negative TIP can implement the constrained-efficient allocation by mitigating the deflationary spiral and reducing the output gap. In conventional calibrations, TIP performs equally well or better than forward guidance and government spending. When the ZLB arises from self-fulfilling expectations, TIP can rule out the deflationary equilibrium. In contrast to other policies at the ZLB, a sufficiently responsive inflation-targeting TIP is robust to the source of the liquidity trap. Simulations of a medium-scale model calibrated for Japan suggest significant stabilization gains from TIP.

JEL classification: E31, E52, E64

Keywords: Deflation, Zero Lower Bound, Liquidity Trap, Tax on Inflation, Tax-based Incomes Policies, Externality

*Capelle: International Monetary Fund (email: dcapelle[at]imf.org), Liu: Boston College (email: yang.liu.7[at]bc.edu). We thank Rui Mano and seminar participants at IMF for helpful comments. All errors are our own. The views expressed therein are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

Once a Japan-specific issue, liquidity traps have become a global challenge since the Global Financial Crisis. During these episodes, output remained persistently below potential and inflation fell short of its medium-run target. Although the post-pandemic surge in inflation briefly lifted interest rates above the zero lower bound (ZLB), a return to a low-interest-rate environment remains likely (IMF, 2023). In June 2025, for instance, the Swiss National Bank cut its policy rate back to zero to fight deflation. Similarly, inflation in China has remained exceptionally low and aggregate demand persistently weak since 2023. Designing effective stabilization policies at the ZLB therefore remains an urgent and globally relevant challenge.

At the ZLB, conventional monetary policy becomes constrained and largely ineffective, leaving other instruments such as forward guidance and fiscal interventions to play a central role. In this spirit, we evaluate the effectiveness of a *Tax on Inflation Policy* (TIP)—a fiscal instrument that shapes firms' incentives to adjust prices—as a tool to improve welfare in a liquidity trap. Prior work has shown that TIP effectively regulates inflation outside of the liquidity trap (Capelle and Liu, 2025). We hypothesize that, by encouraging firms to raise prices, a negative TIP can also mitigate the deflationary spirals characteristic of liquidity traps. To test this hypothesis, we embed a TIP in a standard New Keynesian model with a zero lower bound and multiple sources of liquidity traps. We derive the optimal TIP, analyze its effectiveness across different types of liquidity traps, and compare it with other policies proposed to lift the economy out of the ZLB.

We start our investigation with a small-scale NK model that incorporates TIP and analyze how TIP can be used to improve upon the *laissez-faire* when the economy is in a persistent ZLB. We consider three potential sources of the liquidity trap: a decline in the neutral rate of interest, pessimistic self-fulfilling expectations, or shifts in the Phillips curve. While shifts in the Phillips curve have not been considered in the literature as a potential driver of the recent liquidity trap episodes, they are a theoretical possibility and our quantitative exercise provides some empirical support for it. An appealing feature of the small-scale NK model is that we are able to provide closed-form expressions for the decentralized and socially optimal allocations in each of the three cases.

In a liquidity trap caused by a decline in the neutral rate of interest, a negative TIP effectively curbs the deflationary spiral and implement the constrained-efficient

allocation in which inflation is positive and output is below its full-employment level. In the *laissez-faire*, firms instead cut prices in a competition for limited demand, which generates deflation, raises the real interest rate and further depresses demand.[1] Firms don't internalize the pecuniary externality of their private decisions to decrease prices, which further raises the real interest rate above its neutral level. By giving incentives to firms to increase prices, a negative TIP aligns the private with the social benefit of price increases in a liquidity trap, which is to relax the zero lower bound constraint. While TIP can implement the constrained-efficient allocation, it cannot however restore the first best in which inflation is at target and output is at its potential.

We then compare the welfare gains implied by TIP with those of forward guidance and of government spending, two commonly proposed policies in this type of liquidity traps. Like TIP, these two other policies are second best. Forward guidance entails future costs, since it relies on overheating the economy after the liquidity trap has ended. Government spending is also costly because it crowds out private consumption and distorts the relative size of private and public consumptions. Numerical simulations of a calibrated small-scale model suggest that, in the empirically-relevant part of the parameter space, TIP performs as well as government spending and better than forward guidance. Importantly, these comparisons abstract from the fiscal costs of government spending and implied distortionary taxation as well as from from the commitment problems and other channels that limit the power of forward guidance (McKay et al., 2016).

When there are multiple equilibria and a liquidity trap arises because of a coordination failure and pessimistic self-fulfilling expectations, we show that TIP can restore the first-best allocation. It does so by ruling out the existence of the deflationary equilibrium, ensuring the coordination of private agents on the unique first best equilibrium with inflation at target and no output gap. Surprisingly, eliminating the liquidity trap equilibrium requires a positive and large enough TIP rate, *i.e.* a tax on price increases as opposed to a tax on price decreases. Compared to active fiscal policies, a commonly proposed policy to address this type of ZLB, we see two advantages in using TIP (Benhabib et al., 2001a). First, the effectiveness of active

[1] In practice, this mechanism may be compounded by an increase in real wages and debt deflation, due to sticky wages and financial frictions, which would amplify deflation. This paper abstracts from these additional effects, see Eggertsson and Krugman (2012) for an analysis incorporating debt deflation.

fiscal policies hinges on the commitment of governments to not fund debt with future taxes, and on the credibility of this commitment with the public. Second, if this credibility is weak and the policy doesn't create the expected inflationary boom, the accumulated government debt may trigger fiscal issues and a debt crisis.

We then consider a liquidity trap caused by a shift in the Phillips curve. We find that in this case, TIP can perfectly restore the first-best allocation with full employment and inflation at target. This result is consistent with Capelle and Liu (2025) which shows that TIP can perfectly stabilize macroeconomic fluctuations in the face of cost-push shocks. In this case, the fundamental source of inefficiency—the cost-push shock—is the wedge between the private and social valuations of price increases. TIP can perfectly address this inefficiency and re-align private and social benefits.

Given the uncertainty surrounding the underlying drivers of liquidity traps in practice, it is crucial to design policy rules that are robust to all potential sources. Like other policies analyzed in the literature and used in practice—forward guidance, government spending, unfunded fiscal expansion—, the simple constant Tax on Inflation Policy studied above isn't robust. To investigate the set of robust policies, we set up the problem of a welfare-maximizing planner who learns from observing the *laissez-faire* equilibrium but remains uncertain about the model's parameters and shocks. The planner maximizes over the set of inflation-targeting rules for TIP. Our investigation of robust policies follows three common approaches to optimal decisions under model uncertainty, ambiguity aversion and learning: (i) a planner with a prior on the model's parameters maximizes expected social utility, (ii) a planner with multiple priors maximizes social welfare in the worst possible case (Gilboa and Schmeidler, 1989) and (iii) a planner facing ambiguity and with a reference model in mind maximizes multiplier preferences (Hansen and Sargent, 2001).

We show that an inflation-targeting rule that is sufficiently responsive to deflation and with a negative intercept is a robust optimal policy. We derive this result both analytically under certain restricted informational structure and numerically for more general ones. In liquidity traps arising from a low neutral rate of interest, welfare increases monotonically with the strength of the response of TIP to deflation. TIP needs to react sufficiently strongly to deflation to eliminate the self-fulfilling trap and increase welfare in a Phillips-curve-induced trap.

These results connect to the "paradox of flexibility" according to which increasing

price stickiness at the ZLB can improve welfare (Eggertsson and Krugman, 2012). This is because an inflation-targeting TIP increases the degree of *aggregate* price stickiness. Although this paradox was derived in the context of a liquidity trap caused by a decline in the neutral rate of interest, our results suggest that if the increase in price stickiness is strong enough it improves welfare in all types of liquidity traps. However, small increases in price stickiness worsen the allocation in a self-fulfilling equilibrium. An important policy implication is that the strength of the response of TIP to inflation needs to be appropriately calibrated. Finally, it is worth stressing that TIP increases aggregate price stickiness without increasing price stickiness at the firm-level. This is because TIP is a linear tax which doesn't impede the adjustment of relative prices, contrary to price controls (Capelle and Liu, 2025).

We then evaluate the stabilization properties of TIP in a medium-scale DSGE model à la Smets and Wouters (2007) calibrated to Japan. Japan provides an ideal laboratory for this exercise, having experienced prolonged deflation and a policy rate constrained by the zero lower bound since the early 2000s. Drawing on Hirose (2020), the model includes a variety of important features and realistic frictions. We consider two different calibrations of the model, corresponding to the different explanations of the persistent liquidity trap, following the approach by Cuba-Borda and Singh (2024). We also introduce multiple temporary shocks driving macroeconomic fluctuations.

Using the calibrated model, we find that TIP could have been effective at lifting Japan out of its deflationary equilibrium during 1998–2012. In the self-fulfilling liquidity trap calibration, a sufficiently aggressive reaction of TIP to inflation—corresponding to 80 percentage points for each percentage point of additional deflation—would eliminate the self-fulfilling liquidity trap. In the r^* calibration of the model, deflation could have been substantially milder under TIP, with inflation positive in most periods and closer to the 1% target. On average, the rule raises inflation by about 1.5 percentage points, with negligible impact on output growth.

Literature. In a liquidity trap caused by a decline in the neutral rate of interest, we highlight a pecuniary externality in firms' pricing decisions at the zero lower bound. Blanchard and Kiyotaki (1987) first emphasized the aggregate demand externality, whereby agents fail to internalize how their individual consumption decisions affect overall demand and others' income. Later work focused on pecuniary externalities stemming from deleveraging or precautionary saving which depress asset prices and tighten borrowing constraints (Eggertsson and Krugman, 2012; Korinek and Simsek,

2016; Farhi and Werning, 2016). Our paper shows that an inefficiency arises even in a representative-agent economy without leverage or heterogeneity: when the ZLB binds, firms ignore that cutting their prices exacerbates aggregate deflation, raises real interest rates, and depresses demand. The resulting inefficiency can be mitigated by TIP, which realigns private and social incentives to adjust prices.

We contribute to the large literature on the fiscal-monetary interactions and fiscal policies at the ZLB. Papers have analyzed the role of tax increases aimed at restricting supply (Eggertsson and Woodford, 2006), tax cuts aimed by stimulating demand (Eggertsson, 2011), investment subsidies (Benigno and Fornaro, 2017), and temporary government spending (Woodford, 2011; Schmidt, 2013; Nakata, 2016) when the ZLB arises from a fundamental decline in the neutral rate of interest. When it arises from self-fulfilling expectations, other papers have emphasized the role of cuts in marginal labor tax rates that boost confidence (Mertens and Ravn, 2014), or non-ricardian fiscal expansion (Benhabib et al., 2001b, 2002). Finally, well-designed paths of consumption and labor taxes, import and export tariffs can also help by mimicking the same relative price of final goods over time as in the first-best allocation (Correia et al., 2013; Farhi et al., 2014).

Our contribution is to analyze the stabilization properties of TIP. We highlight how it can address the deflationary spiral at the ZLB and compare its effectiveness with other policies proposed at the ZLB, such as forward guidance and government spending. In contrast to these papers which focus on a specific driver, we analyze the effectiveness of TIP to all potential sources of the liquidity trap and highlight the conditions of its robustness.

We build on and contribute to the literature on robust policies in the context of model uncertainty. We build on seminal works on model uncertainty, ambiguity and robust policies by Gilboa and Schmeidler (1989), Hansen and Sargent (2001), Hansen and Sargent (2008) and Adam and Woodford (2021). In the context of liquidity traps, Tillmann (2021) and Özcan and Traficante (2024) analyze robust optimal policies when policymakers are uncertain about the location of the ELB. By contrast, we focus on uncertainty about the entire structure of the economy, i.e. the slopes and intercepts of the Euler equation and the Phillips curve. An important novelty is that we allow the policymakers to learn about the model's structure from the observed *laissez-faire* as in Sargent et al. (2006). Our findings reveal that a simple rule for TIP with a sufficiently aggressive reaction to inflation combined with a negative intercept

is a robust policy at the ZLB, in contrast with alternative policies used in practice. These results are robust to various approaches to robustness.

Our paper also contributes to the literature on TIP which originates in the 1960s and 1970s with Wallich and Weintraub (1971), Seidman (1978) and Seidman (1979). More recently, Capelle and Liu (2025) analyze the effectiveness of TIP outside of the liquidity trap, showing that TIP could effectively correct externalities in firms' pricing decisions, tackle excessive inflation and reduce output volatility, without exacerbating price distortions. The contribution of this paper is to focus on the role of TIP at the liquidity trap.

Our paper builds on recent studies quantifying the sources and implications of the zero lower bound in Japan (Hirose, 2020; Cuba-Borda and Singh, 2024). Our quantitative assessment draws on these approaches by studying the effectiveness of TIP in a medium-scale NK model. We find that TIP could have provided stimulus at the zero lower bound, highlighting that TIP is a robust policy.

The remainder of this paper is organized as follows. Section 2 introduces a conventional sticky price model augmented with a TIP and a ZLB. Section 3 analyzes the optimal design of TIP when the ZLB arises from a fundamental decline in the neutral rate of interest and Section 4 compares it to forward guidance and government spending. Section 5 analyzes the effectiveness of TIP when the ZLB arises from self-fulfilling expectations or a shift to the Phillips curve and shows that a simple inflation-targeting rule is a robust policy. Section 6 provides a quantitative analysis of TIP at the zero lower bound in Japan. Section 7 concludes.

2 A New-Keynesian Model with TIP and a ZLB

We start with embedding TIP into an otherwise conventional New Keynesian model with a zero lower bound on the nominal interest rate. We briefly sketch the model, focusing on the implication of TIP. We then derive a three-equation log-linear representation of the macroeconomic dynamics which we use to analyze the positive and normative properties of TIP in the three types of liquidity traps in the following sections.

2.1 Household.

There is a representative infinitely-lived household. They seek to maximize their discounted sum of CRRA utility over consumption and labor:

$$\mathcal{U}(B_{t-1}) = \max_{C_t, N_t, B_t} \left\{ \frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1+\psi}}{1+\psi} + \beta_t E_t \mathcal{U}(B_t) \right\} \tag{1}$$

where C_t, N_t and B_t denote consumption, labor supply and nominal wealth in period t, $\beta_t \in (0,1)$ the time discount factor, σ and ψ the inverse elasticity of intertemporal substitution of consumption and labor respectively.

Importantly, the time discount factor, β_t, is allowed to vary stochastically over time. In equilibrium, it will be tightly related to the neutral rate of interest. An increase in the discount factor, corresponding to a decline in the neutral rate of interest, will be one of the three potential causes of the liquidity trap (Eggertsson and Woodford, 2006; Bilbiie, 2019).

The problem of the household is to maximize their lifetime utility subject to the sequence of flow budget constraints $P_t C_t + Q_t B_t = B_{t-1} + W_t N_t + T_t$ and a no-ponzi condition $\lim_{T \to +\infty} \Pi_{j=0}^{T} Q_j B_T \geq 0$. For all $t = 0, 1...$, the optimality conditions are given by

$$\frac{W_t}{P_t} = C_t^{\sigma} N_t^{\psi} \tag{2}$$

$$Q_t = E_t \left[\beta_t \left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{P_t}{P_{t+1}} \right]. \tag{3}$$

The first equation determines the optimal labor supply N_t given consumption C_t and the real wage W_t. The second equation is the traditional Euler equation which determines the optimal path of consumption given the nominal returns on bonds $1/Q_t$, the path of inflation, P_{t+1}/P_t, and the discount factor, β_t.

2.2 Final Good Competitive Firms

The final good is produced competitively by a continuum of firms. The production technology combines a continuum of varieties of intermediate goods supplied by

monopolistic firms with a CES technology:

$$Y_t = \left(\int_0^1 Y_{ti}^{1-1/\epsilon_t} di \right)^{\frac{\epsilon_t}{\epsilon_t - 1}} \tag{4}$$

where ϵ_t is the elasticity of substitution across varieties. Taking the price of the final good P_t and the prices of inputs $\{P_{ti}\}_i$ as given, final good firms maximize profits $\max_{Y_{ti}} P_t Y_t - \int_0^1 P_{ti} Y_{ti} di$ subject to the technology (4).

2.3 Intermediate Good Monopolistic Firms.

There is a continuum of mass one of firms specializing in the production of a single variety of good which they sell to the final good firms. The technology to produce goods displays decreasing marginal returns

$$Y_{ti} = A_t N_{ti}^{1-\alpha} \tag{5}$$

where A_t denotes total factor productivity, which is common across all firms, and $1-\alpha$ is the elasticity of output to labor. Firms are in monopolistic competition and choose their price, P_{ti}, taking into account the final good firms' demand schedule. Following Rotemberg (1982), they also face quadratic adjustment costs to price changes $\mathcal{C}_t\left(P_{t-1i}, P_{ti}\right) = \frac{\theta}{2}\left(\frac{P_{ti}}{P_{t-1i}} - 1\right)^2 P_t Y_t$.

The key novelty is that firms pay a tax proportional to the increase in their price, $\tau_t(P_{ti} - P_{t-1i})$, where $\tau_t \in \mathbb{R}$ is the tax rate which is allowed to vary over time. In addition, the total tax payment scales with the firm's output, Y_{ti}. Hence profits net of taxes are given by

$$\Pi(P_{t-1i}, P_{ti}) = P_{ti} Y_{ti} - W_t N_{ti} - \tau_t(P_{ti} - P_{t-1i}) Y_{ti} - \mathcal{C}_t\left(P_{t-1i}, P_{ti}\right). \tag{6}$$

where labor demand N_{it} is given by the technology (5) and output Y_{ti} by the demand schedule from final goods' firms. In recursive form, the problem of the firm is given by

$$V(P_{t-1i}) = \max_{P_{ti}} \Pi(P_{ti}, P_{t-1i}) + E_t\left[Q_t V(P_{ti})\right] \tag{7}$$

After substituting for labor using the production function, and for output using the

demand schedule, one can take the first order condition for the optimal price. After imposing that in equilibrium all firms are symmetric, $P_{ti} = P_t$, and denoting the rate of inflation $\pi_t = P_t/P_{t-1} - 1$, the optimality condition is given by

$$(\epsilon_t - 1)\left(\mathcal{M}_t MC_t - 1\right) + E_t\left[Q_t \frac{Y_{t+1}}{Y_t}\left(\tau_{t+1} + \theta(\pi_{t+1}+1)\pi_{t+1}\right)\right] = \tau_t\left(1 - \epsilon_t \frac{\pi_t}{1+\pi_t}\right) + \theta\pi_t(\pi_t + 1) \tag{8}$$

where the real marginal cost is given by $MC_t = \frac{W_t}{P_t(1-\alpha)} \frac{Y_t^{\frac{\alpha}{1-\alpha}}}{A_t^{\frac{1}{1-\alpha}}}$ and the ideal markup in the flexible price equilibrium is given by $\mathcal{M}_t = \frac{\epsilon_t}{\epsilon_t - 1}$.

A negative cost-push shock will be one of the three causes of a liquidity trap we analyze this paper. Although it has not been considered in the literature, we show that it is a theoretical possibility and the quantitative exercise provides some empirical support for it. The cost-push is a reduced-form way to model conflicting aspirations of workers and firms over relative wages and prices, as in Werning and Lorenzoni (2023). It can also be interpreted as a shock to inflation expectations due, for instance, to better anchoring of these expectations.

2.4 Central bank: inflation targeting and a ZLB

To close the model, we specify the behavior of the central bank which sets the nominal interest rate. Outside of the zero lower bound the central bank follows an inflation-targeting interest rate rule. However, it is constrained by a zero lower bound on interest rate. Its policy rate is thus given by

$$i_t = -\log(Q_t) = \max(0, \phi_\pi \pi_t - \log \beta_t). \tag{9}$$

where $-\log \beta_t$ is the equilibrium neutral rate of interest. In addition, monetary policy is active, in the sense that the policy rate reacts more than one for one to changes in the rate of inflation:

$$\phi_\pi > 1. \tag{10}$$

Benhabib et al. (2001a) show that the nonlinearity introduced by the zero lower bound on nominal interest rates, when combined with a steep Phillips curve, can give rise to multiple rational expectations equilibria, including a deflationary trap alongside the intended target equilibrium. This self-fulfilling deflationary equilibrium is the third and final potential cause of the liquidity trap.

We close the model with a balanced government budget that ensures that the net transfers to firms due to TIP and the wage bill subsidy are funded by lump-sum taxes on households. In addition, the markets for each intermediate goods, the final good, labor and bonds should clear. We give a full definition of the general equilibrium in Appendix A.1.

2.5 Three-equation Representation

Following the literature and to obtain analytical results, we focus on the linearized version of the model around the efficient equilibrium. We denote $\hat{y}_t = \log(y_t/y_t^e)$ the log-deviation of output from its efficient value. In the absence of TFP shocks, the steady-state equilibrium is also the flexible price equilibrium. We assume that a production subsidy corrects for the steady-state markup distortion. Capelle and Liu (2025) show that a steady-state TIP can also restore efficiency. The model boils down to three equations: an Euler equation, a Phillips curve and a Taylor rule. The following lemma gives their expression.

Lemma 1. *The linearized Euler equation, the Phillips curve and the Taylor rule are given by:*

$$\hat{y}_t = E_t\hat{y}_{t+1} - \frac{1}{\sigma}\left(i_t - E_t\pi_{t+1} + \log\beta_t\right) \tag{11}$$

$$\pi_t = \beta E_t\pi_{t+1} + \kappa\hat{y}_t + \frac{1}{\theta}\left[\beta E_t\tau_{t+1} - \tau_t + u_t\right] \tag{12}$$

$$i_t = \max(0, \phi_\pi\pi_t - \log\beta_t). \tag{13}$$

where $\kappa = \frac{\epsilon-1}{\theta}\left(\sigma + \frac{\psi+\alpha}{1-\alpha}\right)$ *and* u_t *is the cost-push shock.*

Capelle and Liu (2025) show that Calvo-type frictions would lead to the same first-order representation.

3 Optimal TIP in a r^*-Driven Liquidity Trap

We start our investigation with looking at a ZLB arising from a decline in the neutral rate of interest (r^*). While interest rate policy is adequate to address shifts to the neutral rate of interest outside of the liquidity trap, it becomes impotent once the policy rate reaches its ZLB. A natural question is whether TIP can improve welfare in a liquidity trap. This driver of the ZLB has received significant attention (Eggertsson and Woodford, 2006; Eggertsson, 2011; Bilbiie, 2019). This interpretation is consistent with a number of facts, such as the observed secular decrease in the real interest rate documented across many rich countries (Rachel and Summers, 2019).

3.1 Liquidity Trap: the *Laissez-Faire*

We model the decline in the neutral rate of interest as a discount factor shock.[2] The discount factor can take two values $\beta_t \in \{\beta, \bar{\beta}\}$, with $\beta < 1 < \bar{\beta}$ and a ZLB may arise when the discount factor is high $\beta_t = \bar{\beta}$. To make the liquidity trap persistent, we assume that every period there is a constant probability, p, that the discount factor remains high in the following period and that the economy remains in the liquidity trap. With complementary probability $(1-p)$, the economy exits the ZLB. Once the economy leaves the liquidity trap, it stays there forever, *i.e.* β is an absorbing state.

An appealing feature of this stochastic structure is that we can obtain closed-form expressions for the equilibrium in the log linearized model. The equilibrium outside of the ZLB is characterized by $\pi = \hat{y} = 0$ and $i = -\log \beta$. In the ZLB equilibrium, by definition the nominal interest rate is at its lower bound $i = 0$ and the inflation rate π_t^L and the output gap $\hat{y}_t^L$ are given by the Euler equation (11) and the Phillips curve (12). Consistent with the Markov structure of the underlying discount factor shock, we assume that the equilibrium is also Markov and thus omit the time subscripts. The equilibrium is then solution to the following system of two equations in two unknowns, π^L and $\hat{y}^L$:

$$\hat{y}^L = p\hat{y}^L + \frac{1}{\sigma}\left(p\pi^L - \log\bar{\beta}\right) \tag{14}$$

$$\pi^L = p\beta\pi^L + \kappa\hat{y}^L + \frac{1}{\theta}\left[\beta p\tau - \tau\right]. \tag{15}$$

[2]The discount factor shock is also a reduced-way of modeling a deleveraging shock as in Eggertsson and Krugman (2012).

We restrict our attention to the range of parameters that are consistent with a ZLB equilibrium with negative output gap and deflation in the absence of TIP, $\tau = 0$.[3]

Lemma 2 (Decentralized Equilibrium - Decline in Neutral Rate of Interest)**.** *There exists a unique equilibrium with deflation $\pi^L < 0$ and negative output gap $\hat{y}^L < 0$ when $\beta_t = \bar{\beta}$ if and only if*

1. *The discount factor is high enough,* $\log \bar{\beta} > 0$,

2. *The Phillips curve is less steep than the Euler equation:*

$$\frac{1 - p\beta}{\kappa} > \frac{p}{\sigma(1-p)}. \tag{16}$$

The decentralized equilibrium is given by

$$\pi^L = -\frac{1}{(1-p\beta)\sigma(1-p) - \kappa p}\left[\kappa \log \bar{\beta} + \frac{(1-p\beta)\sigma(1-p)}{\theta}\tau\right] \tag{17}$$

$$\hat{y}^L = \frac{1}{\sigma(1-p)}\left(p\pi^L - \log \bar{\beta}\right) \tag{18}$$

Figure 1 plots the Phillips curve and the Euler equation in and outside the liquidity trap and shows how a liquidity trap can arise when the discount factor is high $\beta_t = \bar{\beta}$.

3.2 Constrained-efficient Allocation and Pecuniary Externality

To benchmark the decentralized allocations implemented by TIP, we first derive the constrained-efficient allocation. The social planner looks for an allocation, *i.e.* a path of output, labor, bonds, prices and bond prices $\{C_t, N_t, B_t, Y_t, N_t, P_t, Q_t\}_{t=0,1,2..}$ that maximizes the household utility (1), subject to the technological constraints (4), (5) and the resource constraints (45), (46) as well as the implementation constraint

$$E_t\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\sigma}\frac{P_t}{P_{t+1}}\right] \leq 1. \tag{19}$$

[3]The persistence of the liquidity trap, parametrized by p, gives rise to an upward sloping Euler equation in the liquidity trap of the economy. This is an important property to obtain a persistent liquidity trap equilibrium. Other microfoundations such as wealth in the utility (Michaillat and Saez, 2021; Cuba-Borda and Singh, 2024) or cognitive discounting (Gabaix, 2020) lead to the same reduced-form representation. Our results are robust to the specific microfoundation.

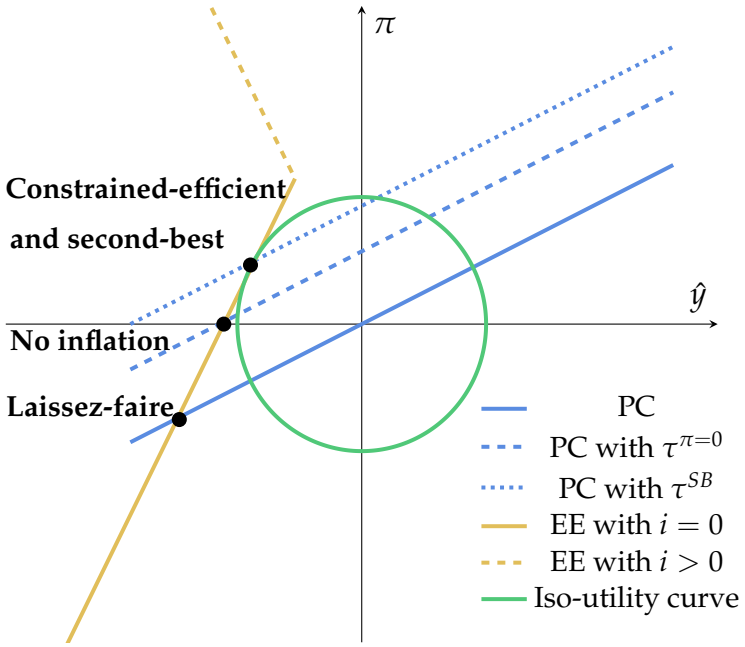


Figure 1: TIP in an r^*-driven liquidity trap

Note: The blue lines represent Phillips curves und er different policies. The orange line represents the Euler equation. The iso-utility curve illustrates preferences over inflation-output pairs.

Instead of solving for the full non-linear solution of this problem, we reformulate it as a linear-quadratic (LQ) optimal policy problem, following the literature (Rotemberg and Woodford, 1999; Eggertsson and Woodford, 2003). The social welfare objective consists in minimizing the discounted weighted sum of welfare losses, given by squared inflation and output gaps:

$$\mathcal{L} = \sum_{t=0}^{\infty} E_0 \beta^t \left[\pi_t^2 + \eta_y \left(\hat{y}_t \right)^2 \right] \tag{20}$$

for some $\eta_y \geq 0$. A second-order approximation of the household welfare loss around the efficient steady-state would deliver this objective function with $\eta_y = \frac{1}{\theta}\left(\sigma + \frac{\alpha+\psi}{2(1-\alpha)}\right)$. The problem of the social planner is to minimize this objective subject to the linearized version of the inequality ZLB constraint (19). By definition, this constraint binds only in the liquidity trap equilibrium. In the LQ version of the problem, the constraint is given by equation (18).

We make two restrictions to the set of allocations. First we assume that as soon

as the discount factor reverts back to its lower value, β, and the economy exits the liquidity trap, the social planner implements the first-best allocation, with $\pi = \hat{y} = 0$. The main purpose of this assumption is to exclude allocations of the forward guidance type, in which the social planner manufactures a boom after the economy exits the liquidity trap. Second, consistent with our focus on Markov equilibria in the decentralized economy, we look for Markov equilibria in the constrained-efficient economy. This allows us to remove the time subscript on π^L and $\hat{y}^L$.

Our analysis of the constrained-efficient allocation in the liquidity trap complements prior work by Eggertsson and Woodford (2003). In this paper, they provide a detailed analysis of the time-varying constrained-efficient allocation and show how forward guidance can implement it. Our approach differs in two important ways. First, we don't assume that the social planner is constrained by the optimal price-setting behavior of private agents as they do. As a result, our constrained-efficient allocation is a significantly more demanding benchmark to attain than theirs and it allows us to completely relax the constraint on inflation dynamics, consistent with our focus on the deflationary spiral. Second, as explained above, we prevent the social planner from generating a boom after exiting the ZLB. This assumption is made for a couple of reasons. These allocations may prove difficult to implement in practice as they rely on the ability to credibly commit to overheat the economy after exiting the ZLB and on implausibly strong responses of private agents to future inflation and output growth (McKay et al., 2016; Gabaix, 2020). In addition, these intertemporal mechanisms have been extensively studied in prior work and we focus instead on the deflationary spiral and the pecuniary externality.[4]

The following proposition characterizes the constrained-efficient allocation. We denote with an "LF" superscript, the variables in the *laissez-faire*. Proofs can be found in Appendix A.2.

Proposition 1 (Constrained-efficient Allocation)**.** *In the constrained efficient (CS) allocation, the optimal output gap and inflation rate are given by*

$$\hat{y}^{LF} < \hat{y}^{CS} = -\frac{1}{\sigma^2(1-p)^2 + p^2\eta_y} \log \bar{\beta} < 0 \tag{21}$$

$$\pi^{LF} < 0 < \pi^{CS} = -\frac{\eta_y p}{\sigma(1-p)} \frac{1}{\sigma^2(1-p)^2 + p^2\eta_y} \log \bar{\beta}. \tag{22}$$

[4]It is worth noting that TIP could also be used to manufacture booms after exiting the ZLB.

As illustrated in Figure 1, the social planner chooses a positive inflation rate and a negative output gap. By boosting inflation, the social planner relaxes the constraint on the zero lower bound, reduces the real interest rate and thus increases output. Around the optimal point where inflation is positive, the social planner faces a trade-off between limiting the cost of inflation and further reducing the output gap. At the optimum, the marginal cost of increasing inflation should be equal to the marginal gains of reducing the output gap.

An important implication of this lemma is that there is scope for improving upon the *laissez-faire*, due the presence of a pecuniary externality at the zero lower bound. Although the ZLB is the *causa sine qua non* of the inefficiency in equilibrium, it is compounded by a deflationary spiral stemming from the behavior of firms competing to attract scarce demand when output is below potential. Firms fail to take into account the fact that by decreasing their individual price, they contribute to aggregate deflation which in turn increases the real rate of interest and depresses output, which aggravates the imbalance between aggregate demand and output potential.

3.3 Ramsey Allocation: Welfare-maximizing TIP

To correct this pecuniary externality, giving rise to excessive deflation, policymakers can use TIP to directly affect the price-setting behavior of firms. We now investigate the optimal setting of TIP by solving for the Ramsey allocation. The Ramsey optimal allocation problem consists in choosing a TIP rate τ^L and an allocation to maximize the household utility (1), subject to the constraint that the allocation is a decentralized equilibrium of the economy.[5] In particular, it is subject to the household' optimal intertemporal allocation (3) and the firms' optimal pricing decision (8). The fact that the social planner's problem is subject to the firms' optimal pricing decision is the key difference with the constrained-efficient allocation.

By giving firms incentives to increase their prices, we find that TIP can implement the constrained-efficient allocation. To see this, consider the graphic representation of the social planner's problem. The circle in Figure 1 represents the indifference locus such that $\pi^2 + \eta_y \hat{y}^2 = \bar{U}$, for a given level of household utility $\bar{U}$. By changing TIP, the social planner is able to shift the Phillips curve. The point at which the Phillips

[5]Consistent with our approach in the decentralized and constrained-efficient economy, we also restrict attention to Markov policy rules and on the LQ Ramsey optimal problem.

curve is tangent to the indifference locus corresponds the second-best allocation. At the optimum, the social planner chooses a level of TIP that implements an optimal balance between closing the output gap and the costs of higher inflation. Proposition 2 formalizes the result that the second-best allocation implemented by TIP coincides with the constrained-efficient allocation.

Proposition 2 (Implementation of Constrained-efficient Allocation with TIP)**.** *In the Ramsey optimal second best allocation (SB), the output gap and inflation rate coincide with their constrained-efficient levels*

$$(\pi^L)^{SB} = (\pi^L)^{CS} \quad \textit{and} \quad (\hat{y}^L)^{SB} = (\pi^L)^{CS} \tag{23}$$

and the optimal TIP is given by

$$\tau^L = -\frac{\log \bar{\beta}}{(1-p\beta)\sigma(1-p)}\left[\theta\kappa - \theta\eta_y p\frac{(1-p\beta)\sigma(1-p)-\kappa p}{(1-p)^2\beta^2+\eta_y p^2}\right]. \tag{24}$$

By signaling the social cost of price changes, TIP makes firms internalize the pecuniary externality of their decentralized behavior and can thus break the deflationary spirals. By addressing the deflationary spiral, TIP can then raise output by lowering the real rate of interest. Although TIP corrects the amplification mechanism and addresses the perverse deflationary incentives that firms face when setting their price, it cannot restore the first-best allocation.

3.4 Remarks

First best allocation. The first-best allocation at the ZLB is the same as outside the ZLB: inflation is at zero $\pi^L = 0$ and output is at potential $\hat{y}^L = 0$. TIP isn't able to achieve this first best allocation when the liquidity trap stems from a decline in the neutral rate of interest because it cannot address the root cause of the problem, i.e. the impossibility to bring the nominal interest rate to the negative value of the natural rate of interest.

A necessary condition for any policy to implement the first best is that it changes the relative price of consumption in the present and in the future, *i.e.* it must shift the Euler equation. An appropriate path of increasing consumption taxes, in combination with labor subsidies, has been shown to do exactly that (Correia et al., 2013). By

contrast, TIP doesn't affect the relative price of consumption across periods. Instead it affects the marginal cost of changing prices, and it shows up in the Phillips curve.

Deflation-minimizing TIP, $\eta_y = 0$. One concern with the second best allocation, which requires to sustain inflation above target, is that it wouldn't be politically feasible given the salience of inflation. A more realistic approach would be to assume that policy-makers prioritize price stability over output stabilization. To see how this affects the optimal allocation, we now consider the extreme case in which they have lexicographic preferences, $\eta_y = 0$.

In this case, the optimal TIP achieves zero losses by re-inflating the economy, $\pi^L = 0$. Proposition 3 solves the TIP that can implement such an allocation. Incidentally this level of TIP reduces the negative output gap, although less so than in the second best. We denote $x^{\pi=0}$, the value of variable x in the equilibrium with TIP such that there is no deflation.

Proposition 3. *The TIP that sets inflation to 0, $\pi^L = 0$, at the ZLB is given by*

$$\tau^{SB} < \tau^{\pi=0} = -\frac{\kappa\theta}{1-p\beta}\frac{\log\bar{\beta}}{\sigma(1-p)} < 0. \tag{25}$$

This level of TIP implies a reduction in the output gap:

$$\hat{y}^{LF} < \hat{y}^{\pi=0} = -\frac{\log\bar{\beta}}{\sigma(1-p)} < \hat{y}^{SB} < 0. \tag{26}$$

The optimal level of TIP in this case is less negative than the second best level of TIP, because inflation needs not be as high as in the second best. The cost of a zero inflation is however that the output gap is larger than in the second best. This highlights again the key trade-off policy-makers face when using TIP in a liquidity trap driven by a decline in the natural rate of interest. Figure 1 represents the optimal point when preferences are lexicographic.

3.5 Numerical Illustration

To illustrate this finding and get a sense of the magnitudes of the optimal TIP given by equation (24) and (25), we now calibrate the small-scale New Keynesian model.

Baseline calibration. A period is a quarter. The list of parameters is in Table 1 and consists in the elasticity of output to labor $(1-\alpha)$, the discount factor β, the elasticity of intertemporal substitution σ, the Frisch elasticity ψ and the elasticity of substitution ϵ. In our baseline calibration, all parameters are within the range of standard values in the literature.

Parameters	Description	Value
α	One minus the elasticity of output to labor	0.25
β	Time discount factor	0.99
σ	Elasticity of intertemporal substitution	1
ψ	Inverse Frish elasticity of labor	5
ϵ	Elasticity of substitution across varieties	6
θ	Adjustment cost	372.8
$\bar{\beta}$	Discount factor shock	1.001
p	Persistence of liquidity trap	.7

Table 1: Model parameters

We calibrate the Rotemberg parameter, θ, to match the slope of the Phillips curve in Galí (2015), $\kappa = 0.17$. We choose the discount factor $\bar{\beta} = 1.001$ which amounts to a neutral rate of interest of -.4% per annum. We calibrate the persistence probability to $p = .7$.[6] The weight on output in the welfare loss function is given by its household welfare-consistent expression: $\eta_y = \frac{1}{\theta}\left(\sigma + \frac{\alpha+\psi}{2(1-\alpha)}\right)$.

Results and robustness. In the *laissez-faire*, this parametrization generates an output gap in the liquidity trap of -1.8% and an inflation gap of -.6%. In the second-best allocation, the output gap equal to -.3% and inflation is above target by .01%. The level of TIP that implements this allocation is equal to -44%. This subsidy implies a small fiscal cost for the government of .004% of firms' sales.

We then look at how the optimal TIP given by the formula (24) and (25) vary with four important parameters: the persistence of the liquidity trap p, the price adjustment cost θ, the intertemporal elasticity of substitution σ and the discount factor shock $\bar{\beta}$. Figure 2 shows the results. The dashed black vertical line corresponds to the baseline value of the parameter.

Consistent with equations (24) and (25), the optimal level of TIP needs to be

[6] This value is slightly below that used in Bilbiie (2019). Higher values of p would violate condition (32).

more negative as the discount factor shock increases, *i.e* as the neutral rate of interest declines. This is because a decline in neutral rate of interest worsens the output and the inflation gaps. The degree of price stickiness has a negligible impact on the optimal of TIP, which is the result of two offsetting effects. On the one hand, stickier prices attenuate the negative effects of the decline in the neutral rate of interest, but on the other it makes a given decrease in TIP less effective at boosting inflation. Finally, a higher persistence parameter and a smaller the IES worsen the recession and the deflation, and they thus require a more aggressive decrease in TIP.

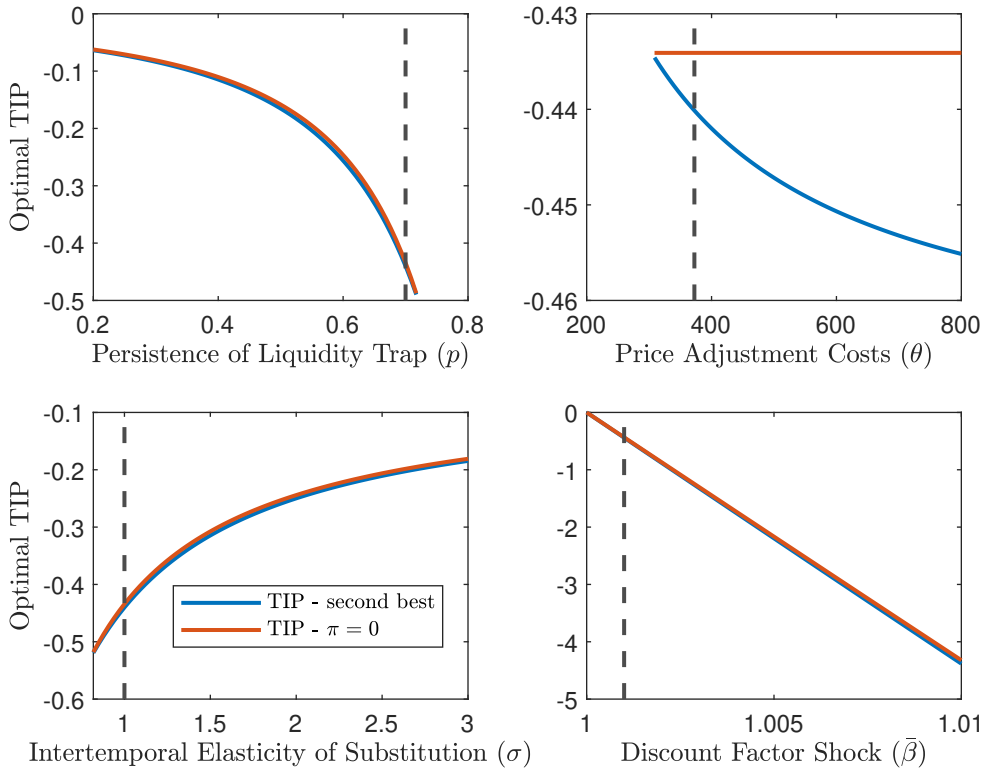


Figure 2: Optimal TIP in a r^*-driven Liquidity Trap

4 TIP and Other Policies in a r^*-Liquidity Trap

Several instruments have been proposed to improve welfare at the zero lower bound. Among them, forward guidance and government spending have received particular attention. In this section, we first compare the welfare gains from forward guidance

and TIP, and then turn to a comparison with government spending. Our aim is not to offer a quantitative ranking—given the simplicity of our small-scale New Keynesian model—but rather to provide a proof of concept that TIP can perform at least as well as traditional policy tools in a standard calibrated framework.

4.1 TIP and Forward Guidance

Forward guidance has been extensively studied in previous literature, including in Eggertsson and Woodford (2003) and Bilbiie (2019). We closely follow the approach proposed in the latter paper in our modeling of forward guidance and assume that when the economy exits the liquidity trap, the central bank commits to keep its interest rate at 0 for some time. As argued in this paper, this second best policy is very well approximated by a time-invariant policy consisting in keeping the interest rate at 0 with some probability q every period.

There are thus three states: a liquidity trap state "L" (where $\beta_t = \bar{\beta}$ and $i_t = 0$), a forward guidance state "FG" (where $\beta_t = \beta$ and $i_t = 0$) and a normal state, "N" (where $\beta_t = \beta$ and $i_t = -\log\beta > 0$). The economy remains in the "L" state with probability p every period. Once it is in the "FG" state, it remains there with probability q every period. Once it exits this second state, it stays in the "N" forever, *i.e.* the normal state is absorbing.

In Appendix B we derive the equilibrium conditions in the liquidity trap and in the forward guidance state. We also derive the expression for the policymakers' objective in these two states. The key trade-off faced by policymakers choosing the persistence of forward guidance q is between boosting the economy in the liquidity trap and overheating the economy after exiting it. The optimal choice of q balances these two forces.

Both TIP and forward guidance are second-best policies. Whether TIP offers better stabilization gains than forward guidance is a complex matter that depends on many parameters. For this reason, we now compare numerically the welfare losses implied by the second best forward guidance and the second best TIP. The calibration is the same as in the previous section.

Results. To compare systematically the welfare losses of forward guidance to those implied by TIP, we vary four important parameters, one at a time: the persistence of the liquidity trap p, the slope of the Phillips curve κ, the intertemporal elasticity

of substitution σ and the discount factor shock $\bar{\beta}$. Figure 3 shows the results. The dashed black vertical line corresponds to the baseline value of the parameters.

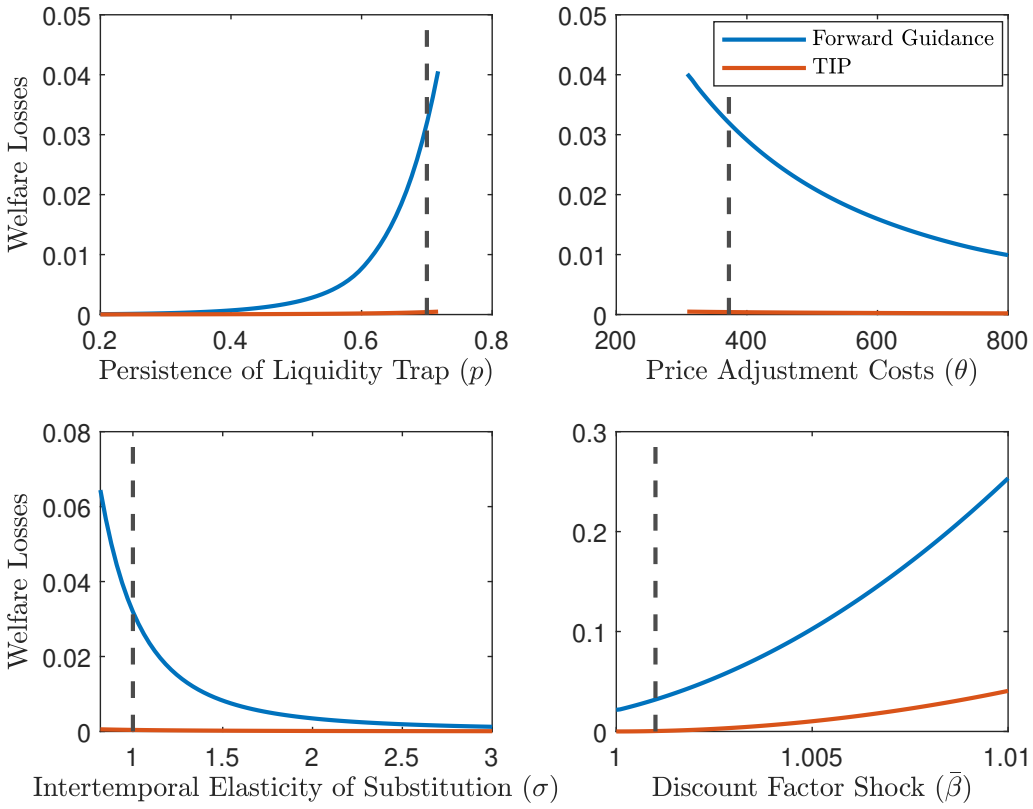


Figure 3: Welfare Losses under Optimal Forward Guidance and TIP

It is interesting to see that welfare losses are smaller under TIP than under forward guidance in most of the parameter space. This is because forward guidance needs to generate a boom in the future. This boom, by distorting the future allocation, entails welfare costs. TIP on the contrary doesn't work through an intertemporal trade-off, but through an intra-temporal trade-off: it boosts inflation today (above its first-best level) to decrease the real rate and increase output. It is worth nothing that, although we abstract from this possibility by construction, TIP could also be used to manufacture a boom after the economy exits the liquidity trap.

As Eggertsson et al. (2019) argue, forward guidance becomes less and less effective when the persistence of the liquidity trap increases. This is true as well in our setting as can be seen in the top left panel of Figure 3. It turns out that TIP remains equally effective even as the persistence of the liquidity trap increases.

Importantly, these comparisons abstract from additional costs and challenges associated with forward guidance. First, its effectiveness relies on the central bank's ability to credibly commit to keeping interest rates low and overheating the economy after exiting the ZLB. Such a policy is inherently time-inconsistent and may prove difficult to implement in practice. Second, the standard model implies implausibly strong responses of private agents to future interest rates, a feature known as the forward guidance puzzle. This arises because the model omits mechanisms that dampen these effects in reality, such as liquidity constraints and bounded rationality (McKay et al., 2016; Gabaix, 2020). Accounting for these limitations would further increase the relative attractiveness of TIP.

4.2 TIP and Government Spending

Government spending is another policy tool that has attracted a lot of attention (Bilbiie et al., 2019; Woodford, 2011; Schmidt, 2013; Nakata, 2016). In this section, we compare the welfare gains associated with government spending and TIP.

We follow the approach proposed in these papers and assume that utility is separable in government spending with intertemporal elasticity of substitution σ_G:

$$\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1+\psi}}{1+\psi} + \frac{G_t^{1-\sigma_G}}{1-\sigma_G} \tag{27}$$

Outside of the liquidity trap, the level of government spending is chosen to equalize the marginal utility of private consumption and government consumption. In the liquidity trap, the optimal choice of government spending balances the gains from a smaller output gap, since government spending boosts output and employment, with its costs. The costs stem from the distortion of the intratemporal allocation across sectors, since government spending increases output but crowds out private consumption. This distortion shows up as a wedge in the household optimal intratemporal condition between consumption and employment which shifts the Phillips curve.[7]

Like for the comparison with forward guidance, whether TIP offers higher or lower welfare gains than government spending depends on many parameters. We

[7] In Appendix C we derive the equilibrium conditions in the liquidity trap. We also derive the expression for the second-order approximation of the welfare loss function, which we take as the policymakers' objective.

thus turn to numerical simulations of the welfare losses implied by both policies for different parameters. We do this comparison along the same four important parameters. Consistent with the literature and the size of government consumption in advanced economies, we set $\sigma_G = 1.5$ and $w_c = .2$.

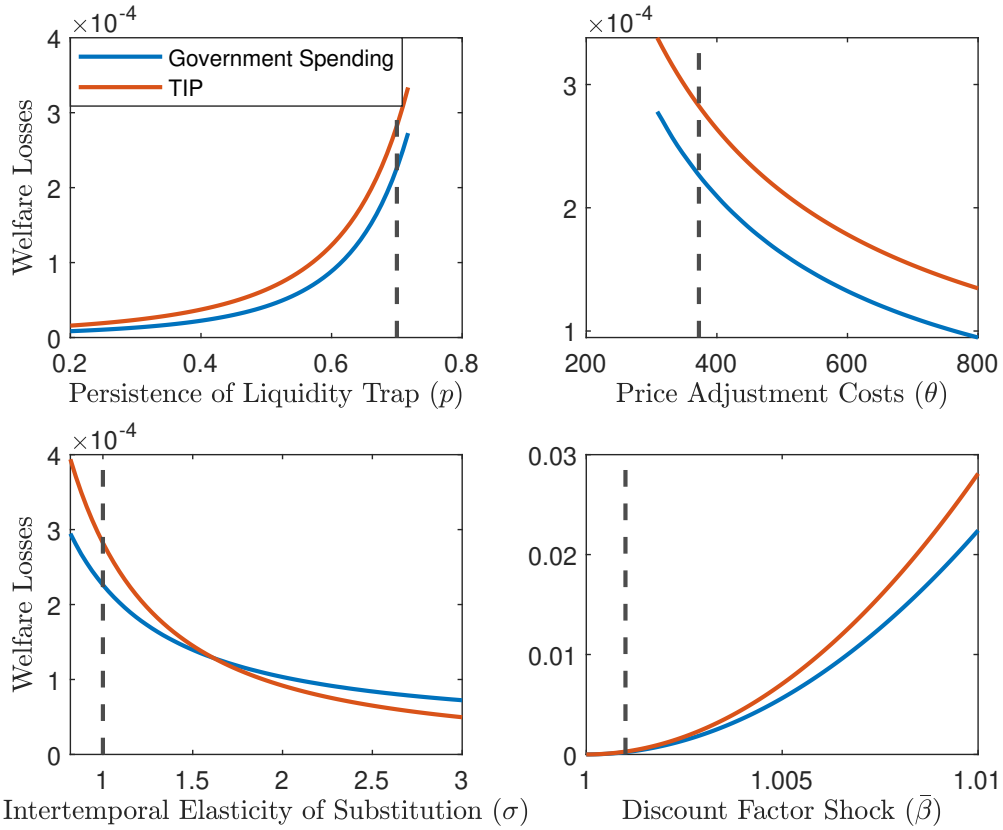


Figure 4: Welfare Losses under Optimal Government Spending and TIP

As illustrated in Figure 4, welfare losses are very similar under TIP and under government spending in most of the parameter space. These comparisons abstract from additional costs and challenges associated with government spending. In particular, they assume spending is financed through costless lump-sum taxes, omitting potential distortions associated with taxation and the fiscal burden resulting from higher government debt. In practice, distortionary taxation and concerns about fiscal sustainability would further limit the attractiveness of government spending.

5 A Robust Rule at the Zero Lower Bound

The previous section focused on the role of fundamental discount factor shocks as drivers of the liquidity trap. We now look at two other potential sources of a liquidity trap: self-fulfilling expectations and shifts in the Phillips curve. A key finding of this section is that a simple inflation-targeting TIP rule is a robust policy to increase welfare in a liquidity trap arising from all three sources, provided it responds sufficiently strongly to inflation.

5.1 A Self-fulfilling Liquidity Trap

An important stream of the literature has shown that liquidity traps could be one of multiple steady-state equilibria when monetary policy follows an active policy and the Phillips curve is sufficiently steep (Benhabib et al., 2001b, 2002; Mertens and Ravn, 2014). In this situation, self-fulfilling pessimistic expectations can lead an economy to the zero lower bound. Cuba-Borda and Singh (2024) finds support for the self-fulfilling hypothesis in Japan.

Model and *laissez-faire*. In contrast to Section 3, the discount factor is fixed at its value below one, $\beta < 1$, and the liquidity trap is one of two possible equilibria. Whenever there exist two equilibria, we assume that a sunspot selects the equilibrium. The sunspot can thus take two values $x_t \in \{L, FB\}$ where "L" stands for liquidity trap, while "FB" stands for first best. In the spirit of the previous section, we assume that the economy is initially in the liquidity trap equilibrium, *i.e.* $x_t = L$, the sunspot stays at its low value with constant probability p every period. With complementary probability, $(1-p)$, it is equal to $x_t = FB$. Like before, the "FB" state is absorbing.

In absence of TIP, the equilibrium outside of the ZLB is characterized by $\pi = \hat{y} = 0$ and $i = -\log\beta$. By contrast, in the liquidity trap equilibrium, the nominal interest rate is at 0, $i = 0$ and the inflation rate and the output gap are solutions to the system of equations given by the Phillips curve (15) and the Euler equation with a constant discount factor below one, $\beta < 1$:

$$\hat{y}^L = p\hat{y}^L + \frac{1}{\sigma}\left(p\pi^L - \log\beta\right). \tag{28}$$

We restrict our attention to the range of parameters that are consistent with the

existence of multiple steady-state equilibria with one displaying negative output gap and deflation in the absence of TIP, $\tau = 0$.

Lemma 3 (Decentralized allocation - Self-fulfilling Expectations)**.** *There exists a ZLB equilibrium with deflation and negative output gap if*

1. *The Taylor coefficient is large enough,* $\phi_\pi > \frac{\kappa p - (1-p\beta)\sigma(1-p)}{\kappa}$,
2. *The Phillips curve is steeper than the Euler equation*

$$\frac{1-p\beta}{\kappa} < \frac{p}{\sigma(1-p)}. \tag{29}$$

The expression for the equilibrium inflation rate and output gap are given by (17) *and* (18)*.*

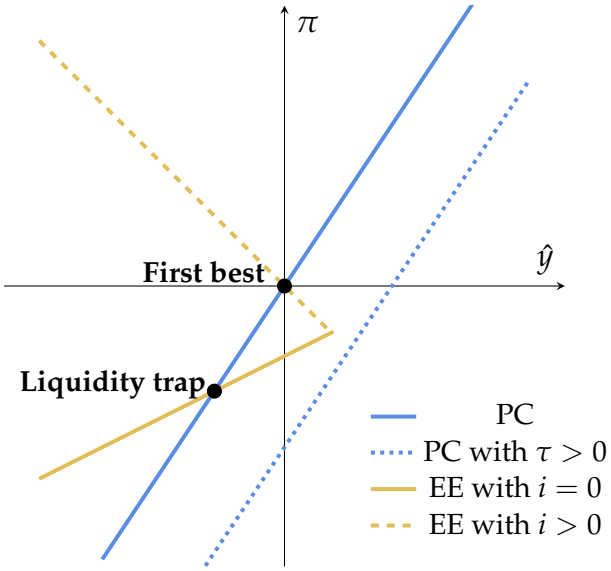


Figure 5: TIP in a self-fulfilling liquidity trap

Note: The blue lines represent Phillips curves under different policies. The orange line corresponds to the Euler equation at and outside the ZLB.

Figure 5 illustrates the multiplicity of equilibria under this parametrization.

Constrained-efficient and first-best allocations. In this type of liquidity traps, both the constrained-efficient and first-best allocations feature zero inflation and no output gap, $\pi^L = \hat{y}^L = 0$. The social planner can easily achieve this allocation

by selecting the good equilibrium. This is implementable even in the constrained-efficient problem because the key source of inefficiency in the decentralized economy is a coordination failure among private agents.

Optimal TIP. In an expectation-driven liquidity trap, can TIP help coordinate agents on the first-best equilibrium? Figure 5 offers an answer: by increasing TIP whenever agents coordinate on the liquidity trap equilibrium, policymakers can eliminate it. Due to the relative slopes of the Euler equation and the Phillips curve, an increase in TIP leads to increases in output and in inflation. At some point, the level of inflation is high enough that the central bank is able to set a positive interest rate, which allows the economy to exit the liquidity trap.

By implementing a sufficiently large increase in TIP should agents coordinate on the deflationary equilibrium, a state-contingent policy ensures that this equilibrium never occurs on the equilibrium path. Denoting τ^L the level of TIP in the liquidity trap and τ^{FB} its level in the first-best equilibrium, the following proposition characterizes the optimal Markov policy. In particular, it gives an expression for the minimum level of TIP needed in the liquidity trap equilibrium to rule out its existence.

Proposition 4 (Optimal TIP in a Self-Fulfilling Liquidity Trap)**.** *There exists a minimum level of TIP τ^{SF} such that if $\tau^L > \tau^{SF}$ and $\tau^{FB} = 0$, the first best equilibrium is unique. In addition,*

$$\tau^{SF} = \frac{\theta}{(1-p\beta)(1-p)\sigma}\left[\frac{(\kappa p - (1-p\beta)\sigma(1-p))\log\beta}{\phi_\pi} - \kappa\right]\log\beta \tag{30}$$

Discussion. It is somewhat surprising that to eliminate the deflationary equilibrium one needs a policy that worsens deflationary pressures. This is another instance of the inversion of intuitive results in a self-fulfilling liquidity trap. For example, it has been shown that an increase in government spending worsens the allocation in that case or that increases in the nominal interest rate could allow the economy to escape the liquidity trap through neo-Fisherian effects (Schmitt-Grohé and Uribe, 2014).

We will show later in this section that a more intuitive inflation-targeting rule for TIP can also eliminate the liquidity trap equilibrium, but through a different mechanism. The appeal of this rule is that it also improves welfare in all types of liquidity traps.

5.2 Shift in the Phillips Curve

We now consider a third possible explanation for an economy to fall into a liquidity trap: a persistence negative cost-push shock that shifts the Phillips curve. Although such an explanation for the ZLB hasn't received much attention in the literature, it is a theoretical possibility and Section 6 provides some empirical support for it. A negative cost-push shock can be interpreted as a decline in the bargaining power of workers or in the firms markup or a shock to inflation expectations due, for instance, to better anchoring of these expectations. In this section, we show that in this type of liquidity trap, TIP can implement the first-best allocation with no inflation and no output gap.

Model and *laissez-faire*. To model a shift in the Phillips curve, we augment the New Keynesian model used previously to allow for a cost-push shock, u_t. Consistent with our approach in the previous sections, we assume that this shock is persistent and can take two values $u_t \in \{\underline{u}, 0\}$ with $\underline{u} < 0$. Every period there is a constant probability p that the cost-push shock remains negative, $u = \underline{u}$, in the following period. Once the shock reverts back to zero, it stays there forever.

Given this stochastic structure, the equilibrium outside of the ZLB is characterized by $\pi = \hat{y} = \tau = u = 0$ and $i = -\log \beta$. By definition of the ZLB equilibrium $i = 0$ and the Phillips curve is given by

$$\pi^L = p\beta\pi^L + \kappa\hat{y}^L + \frac{1}{\theta}\left[\beta p\tau^L - \tau^L + \underline{u}\right]. \tag{31}$$

We restrict our attention to the range of parameters that are consistent with the notion that a negative cost-push shock can lead to a ZLB equilibrium with negative output gap and deflation in the absence of TIP $\tau = 0$.

Lemma 4 (Decentralized Equilibrium - Shifts in the Phillips Curve)**.** *There exists a unique ZLB equilibrium with deflation and negative output gap with $\underline{u} < 0$ if and only if*

1. *The decrease in u is large enough, $\underline{u} < \frac{\theta\kappa}{\sigma(1-p)} \ln \beta$,*
2. *The Phillips curve is less steep than the Euler equation*

$$\frac{1 - p\beta}{\kappa} > \frac{p}{\sigma(1-p)}. \tag{32}$$

In the decentralized equilibrium, output is given by Equation (18) *and inflation by*

$$\pi^L = -\frac{1}{(1-p\beta)\sigma(1-p)-\kappa p}\left[\kappa \log\beta + \frac{\sigma(1-p)}{\theta}\left((1-p\beta)\tau^L - \underline{u}\right)\right] \tag{33}$$

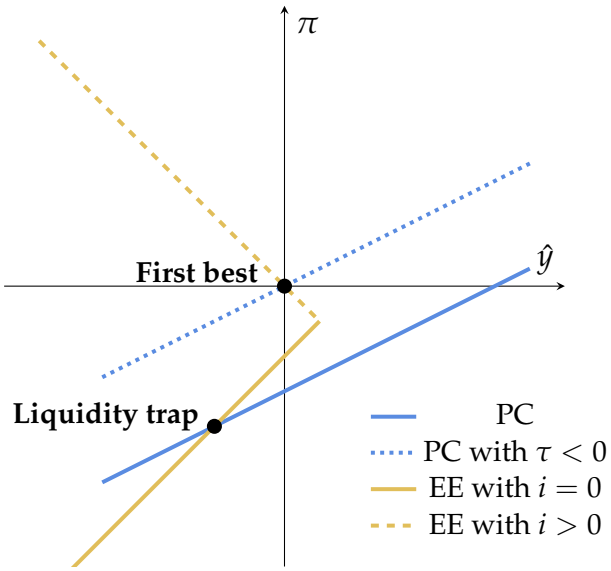


Figure 6: TIP in a Phillips-curve driven liquidity trap

Note: The blue lines represent Phillips curves under different policies. The orange line represents the Euler equation at and outside the ZLB.

Figure 6 plots the Phillips curve and the Euler equation when these conditions are met and shows how a liquidity trap can arise due to a downward shift in the Phillips curve.

Constrained-efficient and first-best allocations. Both the constrained-efficient and the first-best allocations feature zero inflation and no output gap, $\pi^L = \hat{y}^L = 0$. The key source of inefficiency in the decentralized economy is the cost-push shock. This shock opens a wedge between the private and the social returns to decreasing prices: firms' private valuations of price decreases exceed their social counterparts.

Importantly, this externality in the pricing decisions of firms should be distinguished from the pecuniary externality analyzed in the case of a r^*-driven liquidity trap in Section 3. In this earlier case, the fundamental source of inefficiency lied in the combination of zero lower bound constraint and a decline in the neutral rate of interest and the pecuniary externality acted as an amplifying force. The

constrained-efficient social planner could address this amplifying mechanism, but the fundamental source of inefficiencies remained. In the case of a shift in the Phillips curve, the source of inefficiencies is the wedge introduced by the cost-push shock itself. While the same amplification forces are operating in general equilibrium, the important difference is that the social planner is not subject to the cost-push shock. As a result, the constrained-efficient social planner can implement the first best allocation.

Optimal TIP. It turns out that a very simple negative TIP can restore the first best allocation at the liquidity trap. The following proposition gives the expression of the optimal level of TIP.

Proposition 5 (Optimal TIP in a Phillips-curve induced Liquidity Trap)**.** *There exists a Markov policy for TIP such that the allocation is always first-best,* $\pi^L = \hat{y}^L = 0$*. It is given by* $\tau^{FB=0} = 0$ *and*

$$\tau^L = \frac{\underline{u}}{1 - p\beta} < 0. \tag{34}$$

Given that the cost-push shock is negative, $\underline{u}$, the optimal level of TIP should also be negative. A negative TIP closes the wedge between the private and the social returns to decreasing prices and therefore correct excessive aggregate deflation. Compared to the r^*-driven liquidity trap in Section 3 where TIP could only address the pecuniary externality and the amplifying deflationary forces, in the present case, TIP is able to address the root cause of the problem: the shift in the Phillips curve.

Figure 6 shows how TIP can shift the Phillips curve upward, which increases inflation and output. Mathematically, it is also easy to see how TIP works: the shift in the Phillips curve $\underline{u}$ can be exactly undone with a shift in TIP since they both enter linearly in the Phillips curve.

5.3 A Robust Inflation-Targeting Rule

The previous analysis has shown that in a regime in which policymaker sets the level of TIP without targeting specific macroeconomic variables, the optimal policy depends on the sources of the liquidity traps. This lack of robustness of the optimal

TIP is problematic when the planner is uncertain about the true state of the economy.[8]

An inflation-targeting rule for TIP. Given the uncertainty surrounding the source of the liquidity trap in practice and the lack of robustness of a simple TIP, we consider a generalized specification for the design of TIP: an inflation-targeting rule. This rule has an intercept τ_0, as in the previous section, and a feedback term with coefficient φ_π :

$$\tau = \tau_0 + \varphi_\pi \pi. \tag{35}$$

Substituting this rule for TIP gives the following Phillips curve:

$$\pi^L = \frac{\kappa \hat{y}^L + (\underline{u} - (1 - p\beta)\tau_0)/\theta}{(1 - p\beta)\left(1 + \frac{\varphi_\pi}{\theta}\right)}. \tag{36}$$

It is clear from this equation that increasing the response of TIP to inflation φ_π decreases the sensitivity of inflation to the output gap and to the cost-push shock. In other words a stronger response of TIP to inflation flattens the Phillips curve.

To build intuition for the results in this section, we first show how the equilibrium allocation varies with the response of TIP to inflation φ_π in each source of the liquidity trap. Let's define three sets of parameters, $\Omega_s = (\sigma^s, \theta^s, \underline{u}^s, \epsilon_\beta^s)$ for each type of liquidity traps $s = r^*, SF, PC$. In the r^*-driven trap, $\epsilon_\beta > 0$, $\underline{u} = 0$ and the relative slopes are given in Lemma (2); in the SF-trap the slopes are given in Lemma (3) and $\epsilon_\beta = \underline{u} = 0$; and in the PC-trap, $\underline{u} < 0$, $\epsilon_\beta = 0$ and the parameters are such that the assumptions in Lemma (4) are true. All three sets of parameters are consistent with the same *laissez-faire* allocation. The remaining parameters are the same across types of traps. Proofs are given in Appendix F.

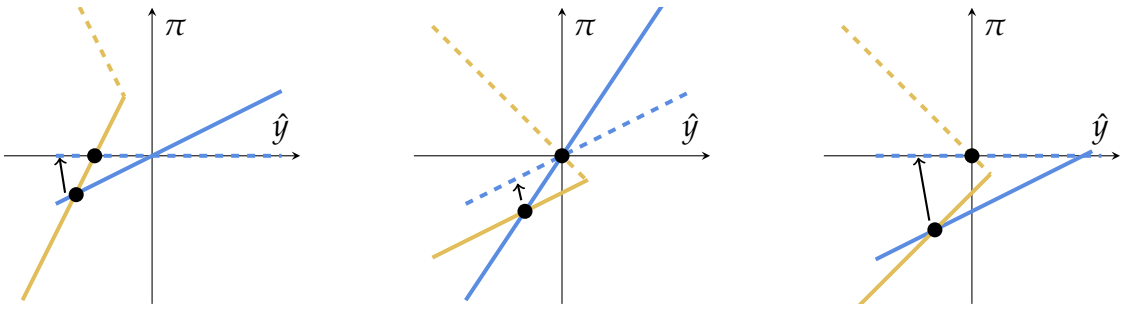
Lemma 5 (Comparative statics - φ_π)**.**

1. *In an r^* liquidity trap, the output gap, inflation and welfare increase monotonically with φ_π. In the limit $\varphi_\pi \to +\infty$, inflation converges to 0 and the output gap to $\hat{y}^L = -\frac{1}{\sigma(1-p)} \log(\beta + \epsilon_\beta)$.*

[8] Other remedies to the ZLB face similar problems: forward guidance and government spending can increase welfare in a r^*-liquidity trap, but they worsen the macroeconomic allocation when the liquidity trap comes from self-fulfilling expectations. Conversely, classical remedies to the expectation-driven liquidity trap such as active fiscal policies have detrimental effects in a r^*-liquidity trap.

2. *In an SF liquidity trap, the first-best equilibrium $\pi = \hat{y} = 0$ is unique if and only if $\varphi_\pi \geq \varphi_\pi^{SF}$ with $\varphi_\pi^{SF} = \theta\left(\frac{\kappa p}{(1-p\beta)\sigma(1-p)} - 1\right) > 0$.*

3. *In a PC liquidity trap, inflation increases monotonically to 0 with φ_π. The output gap increases with φ_π for $\varphi_\pi < \varphi_\pi^{PC}$ (defined in Appendix) then decreases for $\varphi_\pi > \varphi_\pi^{PC}$ and converges monotonically to 0 as $\varphi_\pi \to +\infty$.*



(a) Decline in r^* (b) Self-fulfilling (c) Shift in Phillips curve

Figure 7: Effects of Increasing the Response of TIP to Inflation, φ_π.

The solid blue line represents the Phillips curve in the Laissez-Faire equilibrium, $\varphi_\pi = 0$. The dashed blue line represents the Phillips curve for a stronger response of TIP to inflation, $\varphi_\pi > 0$. In Panel (a) and (c), $\varphi_\pi \to +\infty$. In panel (b), φ_π is such that the Phillips curve and the Euler equation are parallel and the ZLB equilibrium doesn't exist.

Graphically, in a r^*-driven liquidity trap, a clockwise rotation of the Phillips curve around the point where it crosses the x-axis in Panel (a) of Figure 7 implies a lower inflation and a lower output gap. In the expectation-driven liquidity trap, a large enough rotation of the Phillips curve in Panel (b) of Figure 7 eliminates the intersection of the Euler equation and the Phillips curve at the liquidity trap point, leaving only the first-best equilibrium.

Finally, a rotation of the Phillips curve in Panel (c) in Figure 7, for the cost-push shock explanation of the ZLB, leads to a reduction in deflation. The effect on the output gap is however non-monotonic: at first the output gap increases as the real rate decreases with the rise of inflation, but then once the economy exits the zero lower bound, further increases in inflation leads to an increase in the real rate as the central bank reacts strongly to inflation.

Lemma 5 formalizes how welfare change as the response of TIP increases starting for $\varphi_\pi > 0$. When $\varphi_\pi < 0$, further decreases in φ_π makes the Phillips curve more vertical. In the *SF*-trap, it leads to a decline in the output gap, and in deflation, which

is welfare-enhancing. However for a strong enough response, no equilibrium exists in the r^* and PC cases. This is thus not a policy a welfare-maximizing planner would contemplate. For these reasons, we focus on positive response to inflation in the rest of the section, $\varphi_\pi > 0$.

A corollary of Lemma 5 is that when $\varphi_\pi > 0$, it is always weakly welfare maximizing to set the highest possible φ_π. In addition, in the r^*- and $PC-$traps, a negative constant $\tau_0 < 0$ is also optimal.

Corollary 1. *Assume $\varphi_\pi \in [0, \overline{\varphi_\pi})$ with $\overline{\varphi_\pi} > \max(\varphi_\pi^{SF}, \varphi_\pi^{PC})$.*

- *For all $s \in \{r^*, SF, PC\}$, setting $\varphi_\pi = \overline{\varphi_\pi}$ is (weakly) welfare-maximizing.*
- *For $s = r^*, PC$, there exists $c^S < 0$ such that welfare increases as τ_0 decreases for $\tau_0 \in (c^s, 0]$.*

Corollary 1 characterizes the optimal response φ_π if there are three distinct possible states, Ω_s for $s = r^*, SF, PC$. Of course, in practice, policymakers face a higher degree of uncertainty and in the rest of the section, we explore the implication of this uncertainty for the optimal setting of (τ_0, φ_π).

Model uncertainty. To model the planner's (P) uncertainty about the source of the liquidity trap, we assume that she is uncertain about the elasticity of intertemporal substitution σ, the degree of price stickiness θ, the cost-push shocks $\underline{u}$, and the discount factor shocks $\bar{\beta}$. This is a very general approach since these four parameters control the slopes and the intercepts of both curves. The remaining parameters are known with certainty. A vector of these four parameters, $\Omega_s = (\sigma^s, \theta^s, \underline{u}^s, \epsilon_\beta^s)$, thus describes the state s of this economy.

To investigate the set of robust policies, we consider three complementary approaches: (i) P has a prior distribution over the states of the economy and maximizes expected social welfare using these priors, (ii) P has non-unique priors, is averse to ambiguity, and maximizes social welfare in the worst possible case as in Gilboa and Schmeidler (1989); (iii) P also faces ambiguity, fears that the data are generated by an unknown perturbation of his model, and maximizes multiplier preference utility as in Hansen and Sargent (2008) and Woodford (2010).

In all three cases, P also understands that the final allocation $(\pi^L, \hat{y}^L)$ is a decen-

tralized equilibrium of the economy:

$$\log(\beta + \epsilon_\beta) = p\pi^L - \hat{y}^L\sigma(1-p) \tag{37}$$

$$\underline{u} = \theta\left[\pi^L(1-\beta p) - \kappa\hat{y}^L\right] + (1-\beta p)(\tau_0 + \varphi_\pi\pi). \tag{38}$$

Importantly, P initially observes the *laissez-faire* (LF) allocation, $(\pi^{LF}, \hat{y}^{LF})$. This pair of signals provides information to P and imposes restrictions on the set of possible structural parameters and shocks:

$$\log(\beta + \epsilon_\beta) - p\pi^{LF} + \hat{y}^{LF}\sigma(1-p) = 0 \tag{39}$$

$$\underline{u} - \theta\left[\pi^{LF}(1-\beta p) - \kappa\hat{y}^{LF}\right] = 0. \tag{40}$$

By combining equations (37) and (39) on the one hand and (38) and (40) on the other, we can eliminate $\log(\beta + \epsilon_\beta)$ and $\underline{u}$ and reduce the number of constraints to two.

Finally, allocations should be consistent with a liquidity trap, so we exclude those with high inflation, $\pi^L < \bar{\pi}$ for some $\bar{\pi} > 0$ or large positive output gaps $\hat{y} < \bar{\hat{y}}$ for some $\bar{\hat{y}} > 0$.[9]

Maximizing expected social welfare with a prior. In the first approach, P chooses the intercept, $\tau_0 \in [\underline{\tau_0}, \overline{\tau_0}]$ and the slope of the rule $\varphi_\pi \in [0, \overline{\varphi_\pi}]$ to minimize expected losses (20) under her prior distribution subject to the constraints (39), (40), (37), and (38). Given that we can eliminate $\log(\beta + \epsilon_\beta)$ and $\underline{u}$, the expectation is taken over the joint prior distribution of θ and σ, denoted g:

$$W(\tau_0^*, \varphi_\pi^*) = \max_{\tau_0, \varphi_\pi} \int \mathcal{L}(\theta, \sigma, \tau_0, \varphi_\pi) g(d\theta, d\sigma) \tag{41}$$

We rely on numerical simulations to find a solution to this problem. We assume that θ and σ are log-normally distributed with mean (standard deviation) 372 (80) and 2 (1), respectively. For parameters that are common knowledge and for the *laissez-faire* allocation — $\pi^{LF} = -.06$, $\hat{y}^{LF} = -1.6$ — we follow the calibration in Section 3. The set of possible policy parameters encompasses the range of reasonable values, $[\underline{\tau_0}, \overline{\tau_0}] = [-1, 1]$, and $\overline{\varphi_\pi} = 1500$.

[9] It is worth noting that our simulation results are not sensitive to this last pair of constraints.

A max-min planner. Our second approach introduces ambiguity and ambiguity aversion following the seminal work by Gilboa and Schmeidler (1989): we define a robust policy as the outcome of a maxmin problem, or the outcome of a two-stage game between the Planner (P) and a malign Agent (A). In the first stage, P chooses the intercept, $\tau_0 \in [\underline{\tau_0}, \overline{\tau_0}]$ and the slope of the rule $\varphi_\pi \in [0, \overline{\varphi_\pi}]$ to maximize the household welfare (20) anticipating the response of A.

In the second stage of the game, A chooses four parameters to minimize the household welfare (20), given the instruments (τ_0, φ_π) set by P in the previous stage: the underlying structural drivers of the ZLB—the cost-push shock, $\underline{u} \in \mathbb{R}$, and the neutral rate of interest, $\bar{\beta} \in \mathbb{R}$, as well as the parameters of the economy $\sigma \in [\underline{\sigma}, \bar{\sigma}], \theta \in [\underline{\theta}, \bar{\theta}]$. Given the nature of the problem, we need to restrict the sets of possible θ, σ to closed sets $[\underline{\sigma}, \bar{\sigma}]$ and $[\underline{\theta}, \bar{\theta}]$, otherwise A would choose infinite values. A makes these choices subject to the constraints (39), (40) (37), and (38). The fact that P observes the *laissez-faire* allocation $(\pi^{LF}, \hat{y}^{LF})$ in absence of TIP $(\tau_0, \varphi_\pi) = (0, 0)$ adds some two linear restrictions to the set of possible parameters, as explained above.

Denoting G the set of probability distributions consistent with these restrictions, P's objective is given by

$$W(\tau_0^*, \varphi_\pi^*) = \max_{\tau_0, \varphi_\pi} \min_{g \in G} \int \mathcal{L}(\theta, \sigma, \tau_0, \varphi_\pi) g(d\theta, d\sigma) \tag{42}$$

Given the non-convexity of the problem, we rely on numerical simulations to compute the model's solution. We calibrate these sets to encompass the range of reasonable values $[\underline{\sigma}, \bar{\sigma}] = [.5, 10]$ and $[\underline{\theta}, \bar{\theta}] = [50, 700]$.

A planner with multiplier preferences. One concern with the max-min framework is that it delivers an overly cautious objective. Our third approach follows instead Woodford (2010) and Hansen and Sargent (2008). Like in the max-min situation, P faces ambiguity. But in contrast to it, P has a reference prior in mind, g^*, and fears that the data are generated by an unknown perturbation of her approximating model. In this context, P has multiplier preferences and seek to maximize the following objective

$$W(\tau_0^*, \varphi_\pi^*) = \max_{\tau_0, \varphi_\pi} \min_{g \in G} \int \mathcal{L}(\theta, \sigma, \tau_0, \varphi_\pi) g(d\theta, d\sigma) + \nu \int \ln\left(\frac{dg}{dg^*}\right) dg \tag{43}$$

subject to the constraints (39), (40) (37), and (38), but with unbounded support for θ and σ: $(\theta, \sigma) \in \mathbb{R}^+ \times \mathbb{R}^+$. The second terms in the P's objective denotes the relative entropy of the prior g with respect to a reference prior g^*. The calibration is the same as for the case of a planner with a prior distribution.

Results. In Appendix F we plot welfare, inflation and output as a function of (τ_0, φ_π) in all three cases. In the case of the maxmin, we also report the optimal policy rules of θ and σ by A. All three approaches lead to the same conclusions: a strong positive reaction of TIP to inflation, $\varphi_\pi = \bar{\varphi}_\pi$, combined with a negative intercept $\tau_0 = \underline{\tau}_0$ are optimal. The intuition for these results is given by Lemma 5 and Corollary 1.

The paradox of flexibility revisited. The finding that an inflation-targeting TIP can reduce deflation and the output gap echoes the paradox of flexibility according to which reducing price flexibility at the ZLB can improve welfare (Eggertsson and Krugman, 2012). An inflation-targeting TIP indeed increases the degree of aggregate price stickiness.

Our results provide important refinements of this paradox. We show that this paradox, which was derived in the context of a liquidity trap caused by a decline in the neutral rate of interest, holds across several types of liquidity traps. However, it holds in the *SF*- and *PC*-traps only for strong enough increases in aggregate price stickiness. By contrast, small increases in price stickiness unambiguously worsen the allocation in a self-fulfilling equilibrium. An important policy implication is that the strength of the response of TIP to inflation needs to be appropriately calibrated.

Relative price distortions. It is important to stress, however, that an inflation-targeting TIP increases the aggregate stickiness without increasing price stickiness at the firm-level. This is because TIP is a linear tax, which doesn't impede the adjustment of relative prices (Capelle and Liu, 2025). This is in sharp contrast with price controls, which act as implicit convex taxes.

6 A Quantitative Assessment for Japan

We previously showed qualitatively that TIP can enhance welfare at the ZLB. This section uses a medium-scale DSGE model to quantify these stabilization gains in the context of Japan's persistent liquidity trap. Japan provides an ideal laboratory for this exercise, having experienced prolonged deflation and a policy rate constrained by the zero lower bound since the early 2000s.

6.1 A Quantitative Model

Model structure. We extend the small-scale New Keynesian model considered in the previous sections to include a variety of important features and frictions. More specifically, we allow for wealth in the utility, habits in consumption, investment in physical capital subject to adjustment costs, sticky nominal wage and indexation of prices and wages, variable capital utilization and fixed costs in production, following closely Hirose (2020) and Cuba-Borda and Singh (2024). We also allow for multiple temporary shocks driving fluctuations of the economy around its steady state: a price markup, wage markup, risk premium, investment-specific technology, and monetary policy shocks.

We consider two versions of the economy corresponding to the r^* and the expectation explanation of the ZLB. These two economies differ in the values of the utility of wealth and the price stickiness parameters, as in Cuba-Borda and Singh (2024). We describe the calibration strategy in detail below. To address the local indeterminacy around the steady state of the expectation-driven ZLB economy, we introduce a sunspot variable, which indexes inflation expectation errors following closely Cuba-Borda and Singh (2024).[10]

In both versions of the economy, the probability of remaining in the liquidity trap is assumed to be one, $p = 1$. While this parametric assumption would make the Euler equation horizontal in the model of Section 2, the presence of wealth in the utility ensures that it remains upward sloping. The slope of the Euler equation is shaped by the wealth in the utility parameter. Finally, the zero lower bound is assumed to bind in both steady states and at all times, $i_t = 0$ for all t.

[10]There are two types of indeterminacy in the expectation version of the model. First there are multiple steady-state equilibria, a feature we have analyzed in Section 5. Second, there is local indeterminacy around the liquidity trap steady state, which we address by introducing a sunspot shock on inflation expectation errors.

Finally, we introduce the following inflation targeting rule for TIP

$$\tau_t = \rho\tau_{t-1} + \tau_0 + \varphi_\pi(\mathbb{E}_{t-1}\pi_t - \pi^*). \tag{44}$$

This rule is a more realistic version of the simple rule used in Section 2 for three reasons. First policy-makers prefer smoother paths of policies to more volatile ones, as evident from the smooth paths of central bank's interest rates. This motivates the introduction of the backward-looking term, $\rho\tau_{t-1}$. Given that inflation fluctuates significantly at the quarterly level (the frequency of the our analysis), a rule without this backward-looking term would imply significant volatility of the TIP rate. Second we allow the target of inflation, π^*, to be non-zero, but equal to a strictly positive number. We calibrate it to be consistent with stated objectives of the Bank of Japan. Third policymakers set the TIP based on their expectations of inflation during the quarter, $\mathbb{E}_{t-1}\pi_t$, not based on its realized value.

Calibration. Our calibration strategy closely follows Cuba-Borda and Singh (2024) and parameter values are given in Table 2. As in CBS, we calibrate the marginal utility of bonds, δ, to be 0.0405 in the expectation-driven scenario and 0.0415 in the r^*-driven scenario, consistent with a steady state r^* of 0% in the former case and -1.1% in the latter. Our main deviation from CBS is that we assume a much lower degree of price stickiness in the expectation-driven scenario. Specifically, we set the price adjustment cost, ψ_p, to 300 instead of 800 as in CBS, yielding a Phillips curve slope of 0.02. This value is closer to empirical estimates in the literature. For instance, Hirose (2020) estimates the NKPC slope to be around 0.02, and Iiboshi et al. (2022) find it to be around 0.05. As in CBS, ψ_w is determined endogenously to ensure that steady-state inflation equals -1.06%. Given $\psi_p = 300$, wage stickiness is 1783.5—again consistent with Hirose (2020), who estimates that wages are approximately seven times stickier than prices. In the r^*-decline scenario, we set ψ_p at 1600, slightly below the 1800 assumed by CBS. Because this scenario requires a high degree of price stickiness, we cannot bring the NKPC slope closer to empirical estimates. The corresponding ψ_p is 1143.4. Finally, we estimate both scenarios and report estimated parameters that govern the shock processes in Tables 3 and 4.

Note that, because wage stickiness is determined endogenously, the joint degree of price and wage stickiness remains roughly unchanged conditional on the same steady-state inflation. As a result, choosing higher or lower price stickiness has no

Table 2: Calibrated Parameters

Parameters	Description	Value (self-fulfilling / r^* decline)
β	Time discount factor	0.942
δ	Marginal utility of bonds	0.0398 / 0.0752
σ	Elasticity of intertemporal substitution	1
ψ	Inverse Frish elasticity of labor	2.27
ϵ_p	Elasticity of substitution across varieties	6
ϵ_w	Elasticity of substitution across workers	6
α	Capital share	0.37
δ_k	Capital depreciation rate	0.015
h	Consumption habit	0.358
ω	Labor disutility	0.588
ω	Capital utilization elasticity	2.246
ψ_I	Investment adjustment cost	5.16
ψ_p	Price adjustment cost	373.0 / 370.2
ψ_w	Wage adjustment cost	628.6 / 640.1
ι_p	Price indexation	0.225
ι_w	Wage indexation	0.295
g	Autonomous spending	1.333
G_z	TFP trend growth	0.26
G_z	Behavioral discounting	0.998 / 0.945

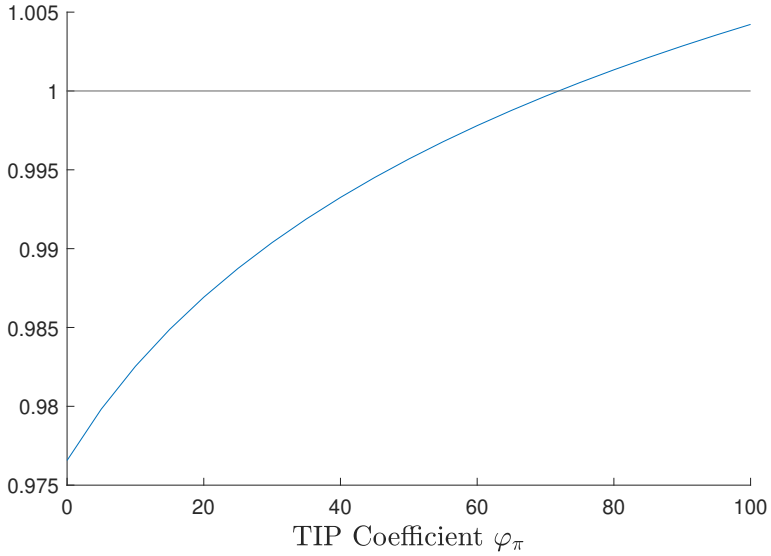
quantitatively important effect on the value of φ_π in (44) that is required to eliminate the self-fulfilling equilibrium.

6.2 Counterfactuals

We now turn to counterfactual analyses quantifying the effectiveness of TIP in lifting Japan out of its deflationary equilibrium during 1998–2012. We consider two scenarios: in the self-fulfilling liquidity trap, we ask what value of φ_π is sufficient to rule out expectation-driven stagnation. In the r^*-driven liquidity trap*, we evaluate how effective TIP is in mitigating deflation.

Self-fulfilling liquidity trap. We quantify the strength of the TIP response to inflation, φ_π, required to eliminate the self-fulfilling trap by computing the set of eigenvalues for alternative values of φ_π. The trap disappears once one of the eigenvalues below 1 crosses above 1. Figure 8 plots this eigenvalue as a function of φ_π.

Figure 8: Eigenvalue as a function of φ_π



Quantitatively, we find that the eigenvalue flips at values of φ_π just below 80. Hence, any φ_π above this threshold is sufficient to eliminate the self-fulfilling liquidity trap. To give a sense of magnitude, when $\varphi_\pi = 80$, a one-percentage-point increase in inflation raises the TIP rate by 80 percentage points.

Liquidity trap caused by a decline of r^* We now turn to the second question: how effective is TIP at reducing deflation in an r^*-driven liquidity trap? As shown in Figure 10, during the period from 1998 to 2013 in Japan, inflation was consistently negative and the policy rate was at the zero lower bound. Our main exercise solves for the policy rule for TIP that brings average inflation back to the 1% target. For this exercise, we start from the calibrated parameters associated with the "decline in r^*" case. We use the model to estimate the underlying structural shocks. We then assume that TIP is used and follows the inflation-targeting rule given by equation (44). We then specify the inertia parameter, $\rho = .9$ and the strength of the reaction to inflation $\varphi_\pi = 100$, which is strictly higher but close to the value that ensures uniqueness in the self-fulfilling calibration. We then solve for the value of the intercept τ_0 that brings inflation back to target.

We find that the value of the intercept that achieves this objective is $\tau_0 = -2.73\%$. The fact that it is negative is consistent with the analysis above that an upward shift of the Phillips curve can boost inflation and output at the ZLB. Figure 9 shows that deflation would have been eliminated under TIP, with inflation positive and close to

the 1% target on average. The rule raises inflation by about 2.5 percentage points, with little variation across years. Despite its strong effect on inflation, however, TIP has only a negligible impact on output growth. This reflects the low sensitivity of output to the real interest rate in our calibration, and thus to inflation.

The TIP rate that implements such a reduction in deflation is negative in most periods, fluctuating between –70% and 30%, as shown in Figure 10. Although these figures may appear to be high, they translate into low effective fiscal cost (or low fiscal revenues when the TIP rate is positive): the fiscal cost fluctuates between -0.15% of GDP and .35% of GDP. The effective transfers are small because the tax base—the change in prices—is not only small to start with but also shrinks as the tax rate rises.

Figure 9: Inflation and Output Growth

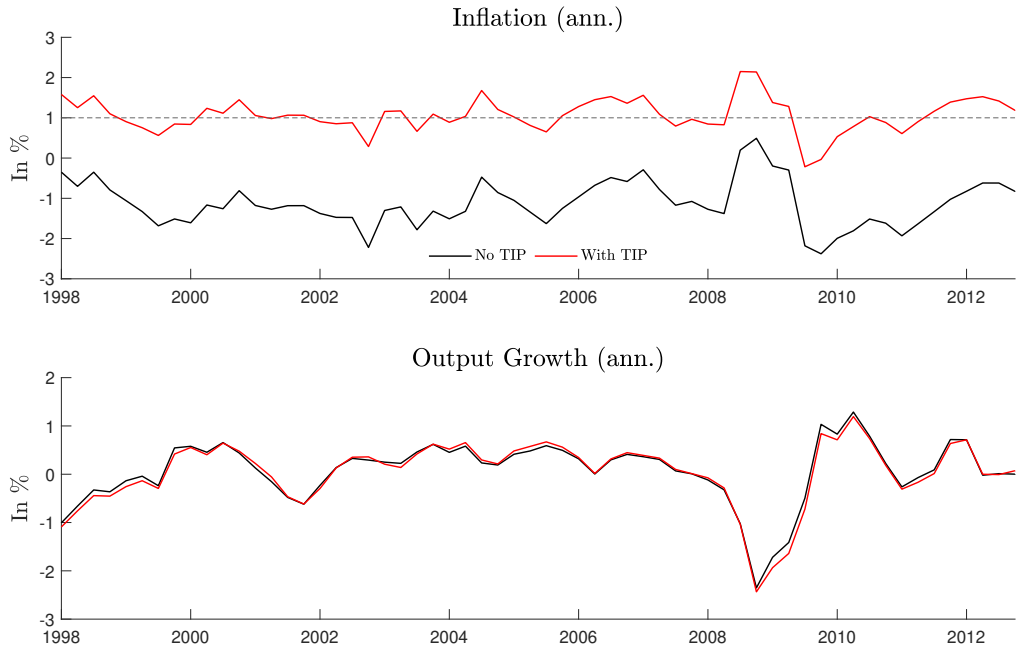
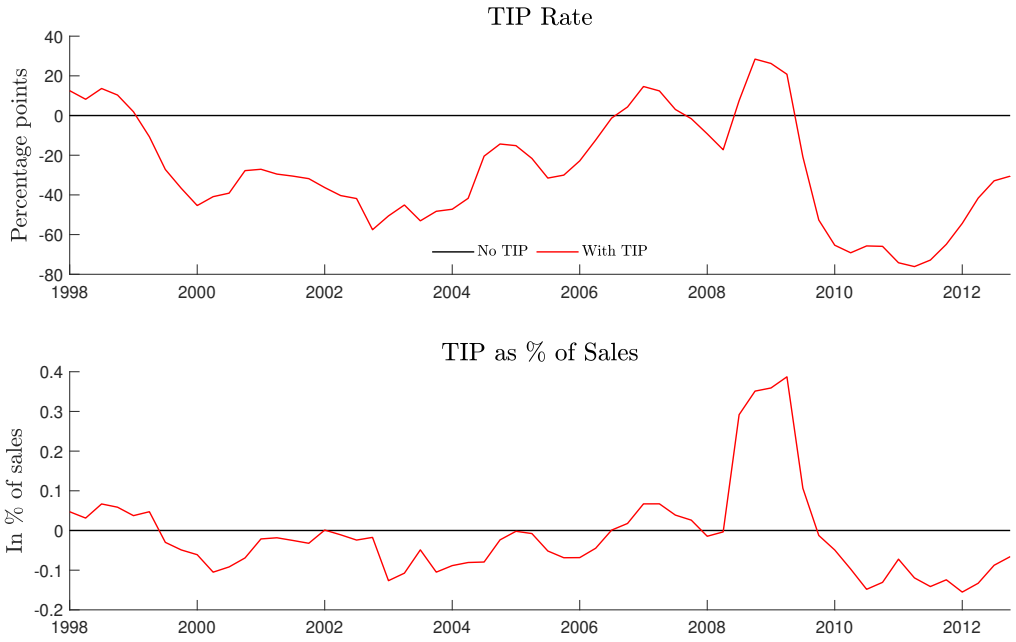


Figure 10: TIP Rate and Implied Tax Burden/Subsidy



7 Conclusion

This paper has assessed the effectiveness of a Tax on Inflation Policy (TIP)—a fiscal instrument that alters firms' incentives to adjust prices—as a tool to lifting the economy out of the ZLB and improve welfare. In a liquidity trap caused by a decline in r^*, by encouraging firms to raise prices, a negative TIP can implement the constrained-efficient allocation by mitigating the deflationary spiral. It can thus increase inflation and boost output. In conventional calibrations, TIP performs equally well or better than forward guidance and government spending. In a liquidity trap caused by coordination failures and self-fulfilling expectations, a large enough TIP can eliminate the ZLB equilibrium. In a ZLB caused by a shift in the Phillips curve, a negative ZLB can restore the first best. Simulations of a medium-scale model calibrated for Japan suggest significant stabilization gains from TIP. We show that an inflation-targeting rule that is sufficiently responsive to deflation and with a negative intercept is a robust optimal policy.

Our paper opens avenues for future research. First, a few implementation issues deserve a more in-depth quantitative inquiry. Tax avoidance, for example, warrants more attention. The main risk is that firms relabel old products as seemingly new ones, or that they shrink their quality. Quantifying the effects of alternative designs

of TIP in a framework with endogenous product creation, information asymmetries about product quality and costly monitoring by the tax administration is an important next step. Second, delving into the political economy of TIP is another important avenue for future research. What are the risks that TIP be used for objectives other than macroeconomic stabilization? Could it lead to less independent monetary policy in countries with weak institutional frameworks?

References

Adam, K. and Woodford, M. (2021). Robustly optimal monetary policy in a new keynesian model with housing. *Journal of Economic Theory*, 198:105352.

Benhabib, J., Schmitt-Grohé, S., and Uribe, M. (2001a). Monetary policy and multiple equilibria. *American Economic Review*, 91(1):167–186.

Benhabib, J., Schmitt-Grohé, S., and Uribe, M. (2001b). Monetary policy and multiple equilibria. *American Economic Review*, 91(1):167–186.

Benhabib, J., Schmitt-Grohé, S., and Uribe, M. (2002). Avoiding liquidity traps. *Journal of Political Economy*, 110(3):535–563.

Benigno, G. and Fornaro, L. (2017). Stagnation traps. *The Review of Economic Studies*, 85(3):1425–1470.

Bilbiie, F. O. (2019). Optimal forward guidance. *American Economic Journal: Macroeconomics*, 11(4):310–45.

Bilbiie, F. O., Monacelli, T., and Perotti, R. (2019). Is government spending at the zero lower bound desirable? *American Economic Journal: Macroeconomics*, 11(3):147–73.

Blanchard, O. J. and Kiyotaki, N. (1987). Monopolistic competition and the effects of aggregate demand. *American Economic Review*, 77(4):647–666.

Capelle, D. and Liu, Y. (2025). Optimal Taxation of Inflation. IMF Working Papers 2023/254, International Monetary Fund.

Correia, I., Farhi, E., Nicolini, J. P., and Teles, P. (2013). Unconventional Fiscal Policy at the Zero Bound. *American Economic Review*, 103(4):1172–1211.

Cuba-Borda, P. and Singh, S. R. (2024). Understanding persistent zlb: Theory and assessment. *American Economic Journal: Macroeconomics*, 16(3):389–416.

Eggertsson, G. B. (2011). What Fiscal Policy Is Effective at Zero Interest Rates? In *NBER Macroeconomics Annual 2010, volume 25*, NBER Chapters, pages 59–112. National Bureau of Economic Research, Inc.

Eggertsson, G. B. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach*. *The Quarterly Journal of Economics*, 127(3):1469–1513.

Eggertsson, G. B., Mehrotra, N. R., and Robbins, J. A. (2019). A model of secular stagnation: Theory and quantitative evaluation. *American Economic Journal: Macroeconomics*, 11(1):1–48.

Eggertsson, G. B. and Woodford, M. (2003). The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity*, 34(1):139–235.

Eggertsson, G. B. and Woodford, M. (2006). Optimal Monetary and Fiscal Policy in a Liquidity Trap. In *NBER International Seminar on Macroeconomics 2004*, NBER Chapters, pages 75–144. National Bureau of Economic Research, Inc.

Farhi, E., Gopinath, G., and Itskhoki, O. (2014). Fiscal Devaluations. *Review of Economic Studies*, 81(2):725–760.

Farhi, E. and Werning, I. (2016). A theory of macroprudential policies in the presence of nominal rigidities. *Quarterly Journal of Economics*, 131(1):223–261.

Gabaix, X. (2020). A Behavioral New Keynesian Model. *American Economic Review*, 110(8):2271–2327.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition*. Number 10495 in Economics Books. Princeton University Press.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Hansen, L. and Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2):60–66.

Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press, stu - student edition edition.

Hirose, Y. (2020). An Estimated Dsge Model With A Deflation Steady State. *Macroeconomic Dynamics*, 24(5):1151–1185.

Iiboshi, H., Shintani, M., and Ueda, K. (2022). Estimating a nonlinear new keynesian model with the zero lower bound for japan. *Journal of Money, Credit and Banking*, 54(6).

IMF (2023). *Chapter 2 The Natural Rate of Interest: Drivers and Implications for Policy*, page CH002. International Monetary Fund, USA.

Korinek, A. and Simsek, A. (2016). Liquidity trap and excessive leverage. *American Economic Review*, 106(3):699–738.

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

Mertens, K. R. S. M. and Ravn, M. O. (2014). Fiscal Policy in an Expectations-Driven Liquidity Trap. *The Review of Economic Studies*, 81(4 (289)):1637–1667.

Michaillat, P. and Saez, E. (2021). Resolving new keynesian anomalies with wealth in the utility function. *The Review of Economics and Statistics*, 103(2):197–215.

Nakata, T. (2016). Optimal fiscal and monetary policy with occasionally binding zero bound constraints. *Journal of Economic Dynamics and Control*, 73:220–240.

Rachel, L. and Summers, L. H. (2019). On Secular Stagnation in the Industrialized World. *Brookings Papers on Economic Activity*, 50(1 (Spring):1–76.

Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.

Rotemberg, J. J. and Woodford, M. (1999). Interest Rate Rules in an Estimated Sticky Price Model. In *Monetary Policy Rules*, NBER Chapters, pages 57–126. National Bureau of Economic Research, Inc.

Sargent, T., Williams, N., and Zha, T. (2006). Shocks and government beliefs: The rise and fall of american inflation. *American Economic Review*, 96(4):1193–1224.

Schmidt, S. (2013). Optimal monetary and fiscal policy with a zero bound on nominal interest rates. *Journal of Money, Credit and Banking*, 45(7):1335–1350.

Schmitt-Grohé, S. and Uribe, M. (2014). Liquidity Traps: an Interest-rate-based Exit Strategy. *Manchester School*, 82(S1):1–14.

Seidman, L. S. (1978). Tax-Based Incomes Policies. *Brookings Papers on Economic Activity*, 9(2):301–361.

Seidman, L. S. (1979). The Role of a Tax-Based Incomes Policy. *American Economic Review*, 69(2):202–206.

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

Tillmann, P. (2021). Robust monetary policy under uncertainty about the lower bound. *The B.E. Journal of Macroeconomics*, 21(1):309–321.

Wallich, H. C. and Weintraub, S. (1971). A Tax-Based Incomes Policy. *Journal of Economic Issues*, 5(2):1–19.

Werning, I. and Lorenzoni, G. (2023). Inflation is conflict. Technical report.

Woodford, M. (2010). Robustly optimal monetary policy with near-rational expectations. *American Economic Review*, 100(1):274–303.

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

Özcan, G. and Traficante, G. (2024). Optimal robust monetary and fiscal policy under uncertainty on the lower bound. *Journal of Macroeconomics*, 81:103605.

Appendix

A Model - details

A.1 Equilibrium

Market Clearing. In equilibrium, the markets for each intermediate good Y_{ti} and for the final good should clear

$$Y_t = C_t + \int_0^1 \frac{\theta}{2}\left(\frac{P_{ti}}{P_{t-1i}} - 1\right)^2 Y_t di \tag{45}$$

The sum of labor hired in all firms should be equal to the supply of labor by the household:

$$N_t = \int_0^1 N_{ti} di \tag{46}$$

With no government debt, holdings of bonds by the household are zero: $B_t = 0$. Finally, transfers received by the household include profits and tax receipts:

$$T_t = \int_0^1 \Pi_{it} + \tau_t(P_{ti} - P_{t-1i})Y_{ti} di. \tag{47}$$

Equilibrium. An equilibrium is a path of output, labor, bonds, wages, price level, bond prices and TIP $\{C_t, N_t, B_t, \{Y_{ti}, N_{ti}, P_{ti}\}_i, W_t, P_t, Q_t, \tau_t\}_{t=0,1,2..}$, such that

- Taking TIP as given, intermediate firms maximize their discounted sum of profits (7) subject to the definition of profits (6), the demand schedule from final goods firms, and the technology (5).
- Taking prices as given, final good firms maximize their profits subject to the technology (4).
- Taking prices and transfers (47) as given, the household maximizes their utility (1) subject to their budget constraint, no-ponzi condition and demand for labor.
- The markets for final good (45), intermediate goods, labor (46) and bonds clear.

A.2 Proof of proposition 1 and 2

Constrained-efficient allocation. The Lagrangian is given by

$$\mathcal{L} = \frac{(\pi^L)^2 + \eta_y(\hat{y}^L)^2}{1-\beta p} + \mu\left[\hat{y}^L - \frac{1}{\sigma(1-p)}\left(p\pi^L - \log\bar{\beta}\right)\right]$$

The F.O.C. with respect to $\pi, \hat{y}$ are given by

$$\frac{2\pi}{1-\beta p} - \mu\frac{p}{\sigma(1-p)} = 0$$
$$\frac{\eta_y 2\hat{y}}{1-\beta p} + \mu = 0$$

which gives

$$\pi = -\eta_y\hat{y}\frac{p}{\sigma(1-p)}.$$

Replacing this into the constraint gives

$$\hat{y}^{SB} = -\frac{\sigma(1-p)}{\sigma^2(1-p)^2 + p^2\eta_y}\log\bar{\beta}.$$

Since this is negative, it implies, together with the optimality condition of the policy-makers, that inflation should be positive $\pi = -\eta_y\hat{y}\frac{p}{\sigma(1-p)} > 0$.

Ramsey allocation. The Lagrangian for the Ramsey optimal allocation problem is given by

$$\mathcal{L} = \frac{\pi^2 + \eta_y\hat{y}^2}{1-\beta p} + \mu\left[\hat{y} - \frac{1}{\sigma(1-p)}\left(p\pi - \log\bar{\beta}\right)\right]$$
$$+ \lambda\left[\pi + \frac{1}{(1-p\beta)\sigma(1-p) - \kappa p}\left[\kappa\log\bar{\beta} + \frac{(1-p\beta)\sigma(1-p)}{\theta}\tau\right]\right]$$

The F.O.C. with respect to $\pi, \hat{y}, \tau$ are given by

$$\begin{aligned} 2\pi + \lambda - \mu \frac{p}{\sigma(1-p)} &= 0 \\ \eta_y 2\hat{y} + \mu &= 0 \\ \lambda &= 0 \end{aligned}$$

The solution to this problem is thus exactly the same as the constrained-efficient allocation. We can then solve for the level of TIP that implements this allocation:

$$\eta_y \frac{\sigma(1-p)}{\sigma^2(1-p)^2 + p^2\eta_y} \log\bar{\beta} \frac{p}{\sigma(1-p)} = -\frac{1}{(1-p\beta)\sigma(1-p) - \kappa p}\left[\kappa \log\bar{\beta} + \frac{(1-p\beta)\sigma(1-p)}{\theta}\tau\right]$$

$$-\frac{(1-p\beta)\sigma(1-p)}{\theta\left((1-p\beta)\sigma(1-p) - \kappa p\right)}\tau = \log\bar{\beta}\left[\eta_y \frac{\sigma(1-p)}{\sigma^2(1-p)^2 + p^2\eta_y}\frac{p}{\sigma(1-p)} + \frac{\kappa}{(1-p\beta)\sigma(1-p) - \kappa p}\right]$$

$$\tau = -\log\left(\bar{\beta}\right)\theta \frac{\left[\eta_y \frac{p((1-p\beta)\sigma(1-p) - \kappa p)}{\sigma^2(1-p)^2 + p^2\eta_y} + \kappa\right]}{(1-p\beta)\sigma(1-p)}$$

A.3 Proof of proposition 4

Proof. To rule out the liquidity trap equilibrium, TIP has to be such that inflation in that equilibrium is above the level below which the nominal interest rate is zero. This implies the following condition

$$\begin{aligned} \pi^L = \frac{1}{\kappa p - (1-p\beta)\sigma(1-p)}\left[\kappa\log\beta + \frac{(1-p\beta)\sigma(1-p)}{\theta}\tau\right] &> \frac{\log\beta}{\phi_\pi} \\ \iff \kappa\log\beta + \frac{(1-p\beta)\sigma(1-p)}{\theta}\tau &> \frac{(\kappa p - (1-p\beta)\sigma(1-p))\log\beta}{\phi_\pi} \end{aligned}$$

Rearranging gives directly the results. □

B Model with Forward Guidance

We follow Bilbiie (2019) in that we assume that when the economy exits the liquidity trap, the interest rate is maintained at 0 with probability q every period. With this,

we can define the inflation rate and output gap in the forward guidance state "F".

$$\hat{y}^F = \frac{1}{\sigma(1-q)}\left(q\pi^F - \log\beta\right)$$
$$\pi^F = \frac{-\log\beta}{(1-q\beta)\sigma(1-q)-\kappa q}$$

We can then solve for the inflation rate and output gap in the liquidity trap state "L" by backward induction.

$$\pi^L = \frac{\beta(1-p)q\pi^F + \kappa\hat{y}^L}{1-\beta p}$$
$$\hat{y}^L = \frac{(1-p)q\hat{y}^F + \frac{1}{\sigma}\left(p\pi^L + (1-p)q\pi^F - \log\bar{\beta}\right)}{1-p}$$

Solving for π^L gives

$$\pi^L = \frac{\beta(1-p)q\pi^F + \kappa\frac{(1-p)q\hat{y}^F + \frac{1}{\sigma}\left(p\pi^L + (1-p)q\pi^F - \log\bar{\beta}\right)}{1-p}}{1-\beta p}$$
$$\Rightarrow \pi^L = \frac{\left(\beta(1-p)^2 q + \kappa\frac{(1-p)q}{\sigma}\right)\pi^F + \kappa\left[(1-p)q\hat{y}^F - \frac{1}{\sigma}\log\bar{\beta}\right]}{(1-p)(1-\beta p) - \kappa\frac{p}{\sigma}}$$

We can also compute welfare losses conditional on starting in the liquidity trap

$$X^F = (\pi^F)^2 + \eta_y(\hat{y}^F)^2$$
$$X^L = (\pi^L)^2 + \eta_y(\hat{y}^L)^2$$
$$W^F = E_F\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2 + \eta_y\left(\hat{y}_t^e\right)^2\right] = X^F + \beta q Y^F = \frac{X^F}{1-\beta q}$$
$$W^L = E_L\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2 + \eta_y\left(\hat{y}_t^e\right)^2\right] = X^L + \beta\left(pW^L + (1-p)qW^F\right) = \frac{X^L + \beta(1-p)qW^F}{1-\beta p}$$

C Model with Government Spending

The model is given by the following system of linearized equations:

$$\hat{c}_t = E_t\hat{c}_{t+1} - \frac{1}{\sigma}\left(i_t - E_t\pi_{t+1} + \log(\beta)\right) \tag{48}$$

$$\hat{w}_t - \hat{p}_t = \sigma\hat{c}_t + \psi\hat{n}_t \tag{49}$$

$$\hat{y}_t = (1-\alpha)\hat{n}_t \tag{50}$$

$$\hat{mc}_t = \hat{w}_t - \hat{p}_t + \frac{\alpha}{1-\alpha}\hat{y}_t - \log(1-\alpha) \tag{51}$$

$$\hat{y}_t = w_c\hat{c}_t + (1-w_c)\hat{g}_t \tag{52}$$

$$\pi_t = \beta E_t\pi_{t+1} + \frac{\epsilon-1}{\theta}\hat{mc}_t + \frac{1}{\theta}\left(\beta E_t\tau_{t+1} - \tau_t\right) + u_t \tag{53}$$

Combining the definition of marginal cost, the goods market clearing condition and the intratemporal condition of the household gives

$$\hat{mc}_t = \hat{y}_t\left(\frac{\sigma}{w_c} + \frac{\psi}{1-\alpha} + \frac{\alpha}{1-\alpha}\right) - \frac{\sigma(1-w_c)}{w_c}\hat{g}_t$$

Using this expression to substitute for the marginal cost in the Phillips curve and using the market clearing condition to substitute for consumption in the Euler equation gives

$$\hat{y}_t - (1-w_c)\hat{g}_t = E_t\left(\hat{y}_{t+1} - (1-w_c)\hat{g}_{t+1}\right) - \frac{w_c}{\sigma}\left(i_t - E_t\pi_{t+1} + \log(\beta)\right) \tag{54}$$

$$\pi_t = \beta E_t\pi_{t+1} + \tilde{\kappa}\hat{y}_t - \frac{\epsilon-1}{\theta}\frac{\sigma(1-w_c)}{w_c}\hat{g}_t + \frac{1}{\theta}\left(\beta E_t\tau_{t+1} - \tau_t\right) + u_t \tag{55}$$

with $\tilde{\kappa} = \frac{\epsilon-1}{\theta}\left(\frac{\sigma}{w_c} + \frac{\psi}{1-\alpha} + \frac{\alpha}{1-\alpha}\right)$.

In the liquidity trap without TIP, the Euler equation and the Phillips curve become

$$\hat{y}^L = (1-w_c)\hat{g}^L + \frac{w_c}{(1-p)\sigma}\left(p\pi^L - \log(\bar{\beta})\right) \tag{56}$$

$$\pi^L = \frac{1}{1-\beta p}\left[\tilde{\kappa}\hat{y}^L - \frac{\epsilon-1}{\theta}\frac{\sigma(1-w_c)}{w_c}\hat{g}^L + u_t\right] \tag{57}$$

Abstracting from cost-push shocks, and substituting y^L into the second equation, we can solve for π^L :

$$\pi^L\left(1-\beta p-\tilde{\kappa}\frac{w_c p}{(1-p)\sigma}\right)=\left[\tilde{\kappa}\left((1-w_c)\hat{g}^L-\frac{w_c}{(1-p)\sigma}\log(\bar{\beta})\right)-\frac{\epsilon-1}{\theta}\frac{\sigma(1-w_c)}{w_c}\hat{g}^L\right] \tag{58}$$

We thus obtain:

$$\pi^L=-\frac{w_c\tilde{\kappa}}{(1-p)\sigma(1-\beta p)-\tilde{\kappa}pw_c}\log\bar{\beta}+\Gamma_g\hat{g}_t \tag{59}$$

$$\Gamma_g=\frac{\tilde{\kappa}(1-w_c)-\frac{\epsilon-1}{\theta}\frac{\sigma(1-w_c)}{w_c}}{\left(1-\beta p-\tilde{\kappa}\frac{w_c p}{(1-p)\sigma}\right)} \tag{60}$$

$$\hat{y}^L=(1-w_c)\hat{g}^L+\frac{w_c}{(1-p)\sigma}\left(p\pi^L-\log(\bar{\beta})\right). \tag{61}$$

D Second-order approximation of welfare loss

We next show that the welfare loss function of the household depends on the inflation rate, the output gap and the government spending gap.

$$\mathcal{L}=\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2-\frac{1-\sigma}{\theta w_c}\left(\hat{y}_t-(1-w_c)\hat{g}_t\right)^2+\frac{1+\psi}{\theta(1-\alpha)}\left(\hat{y}_t\right)^2-\frac{(1-w_c)}{\theta}(1-\sigma_G)\hat{g}_t^2-\frac{\Phi}{\theta}\hat{y}_t\right]$$

Without government spending and without steady-state distortion this simplifies to the usual welfare loss function

$$\mathcal{L}=\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2+\frac{1}{\theta}\left(\sigma+\frac{\alpha+\psi}{2(1-\alpha)}\right)\left(\hat{y}_t\right)^2\right]$$

The log deviation of the representative household is given by

$$\frac{U_t-U}{U_cC}=\left(\hat{c}_t+\frac{1-\sigma}{2}\hat{c}_t^2\right)+\frac{U_NN}{U_cC}\left(\hat{n}_t+\frac{1+\psi}{2}\hat{n}_t^2\right)+\frac{U_gG}{U_cC}\left(\hat{g}_t+\frac{1-\sigma_G}{2}\hat{g}_t^2\right)$$

.

We next use a second order approximation of the goods market clearing condition

to get an approximation for the first bracket

$$w_c\hat{c}_t + (1-w_c)\hat{g}_t = \hat{y}_t - \frac{1}{2}\theta\pi_t^2.$$

with $w_c = \frac{C^{SS}}{Y^{SS}}$.

We next use the labor market clearing condition to get an approximation for the second bracket:

$$\hat{n}_t = \frac{1}{1-\alpha}\left(\hat{y}_t - \hat{a}_t\right)$$

We next use the optimality of spending on government and private consumption in steady-state $U_c = U_g$.

Combining these three conditions and keeping only the terms of first and second order, we obtain

$$\begin{aligned}\frac{U_t - U}{U_c C} = &\left(\frac{\hat{y}_t - (1-w_c)\hat{g}_t - \frac{1}{2}\theta\pi_t^2}{w_c} + \frac{1-\sigma}{2}\left(\frac{\hat{y}_t - (1-w_c)\hat{g}_t}{w_c}\right)^2\right)\\ &+ \frac{U_N N}{U_c C}\left(\frac{1}{1-\alpha}\left(\hat{y}_t - \hat{a}_t\right) + \frac{1+\psi}{2}\left(\frac{1}{1-\alpha}\right)^2\left(\hat{y}_t - \hat{a}_t\right)^2\right)\\ &+ \frac{G}{C}\left(\hat{g}_t + \frac{1-\sigma_G}{2}\hat{g}_t^2\right)\end{aligned}$$

Denote Φ the steady-state distortion implicitly defined by $-\frac{U_N N}{U_C(C+G)(1-\alpha)} = 1-\Phi$. We obtain:

$$\begin{aligned}\frac{U_t - U}{U_c C} = &\left(\frac{\hat{y}_t - (1-w_c)\hat{g}_t - \frac{1}{2}\theta\pi_t^2}{w_c} + \frac{1-\sigma}{2}\left(\frac{\hat{y}_t - (1-w_c)\hat{g}_t}{w_c}\right)^2\right)\\ &- \frac{1-\Phi}{w_c}\left(\left(\hat{y}_t - \hat{a}_t\right) + \hat{\delta}_t + \frac{1+\psi}{2}\left(\frac{1}{1-\alpha}\right)\left(\hat{y}_t - \hat{a}_t\right)^2\right)\\ &+ \frac{1-w_c}{w_c}\left(\hat{g}_t + \frac{1-\sigma_G}{2}\hat{g}_t^2\right)\end{aligned}$$

We can simplify the terms in $\hat{g}_t$.

$$\frac{U_t - U}{U_c C} = \left(\frac{\hat{y}_t - \frac{1}{2}\theta\pi_t^2}{w_c} + \frac{1-\sigma}{2}\left(\frac{\hat{y}_t - (1-w_c)\hat{g}_t}{w_c}\right)^2 \right) - \frac{1-\Phi}{w_c}\left((\hat{y}_t - \hat{a}_t) + \frac{1+\psi}{2}\left(\frac{1}{1-\alpha}\right)(\hat{y}_t - \hat{a}_t)^2 \right) + \frac{1-w_c}{w_c}\frac{1-\sigma_G}{2}\hat{g}_t^2$$

Abstracting from TFP shocks and under the "small distortion" assumption, we can ignore the product of Φ with a second-order term

$$w_c \frac{U_t - U}{U_c C} = \Phi\hat{y}_t - \frac{1}{2}\theta\pi_t^2 + \frac{1-\sigma}{2w_c}(\hat{y}_t - (1-w_c)\hat{g}_t)^2 - \frac{1+\psi}{2(1-\alpha)}(\hat{y}_t)^2 + (1-w_c)\frac{1-\sigma_G}{2}\hat{g}_t^2 + t.i.p$$

E Additional Proofs

E.1 Self-fulfilling liquidity trap

Proof. It has to be the case that inflation, given by $-\frac{1}{(1-p\beta)\sigma(1-p)-\kappa p}[\kappa \log\beta]$, is negative without TIP, which requires $\frac{1-p\beta}{\kappa} < \frac{p}{\sigma(1-p)}$ given that $\log\beta < 0$. It also has to the the case that the nominal interest rate that would be implemented following the rule is indeed negative,

$$\begin{aligned} \phi_\pi \pi^L - \log\beta &< 0 \\ \pi^L &< \frac{\log\beta}{\phi_\pi} \\ -\frac{1}{(1-p\beta)\sigma(1-p)-\kappa p}[\kappa\log\beta] &< \frac{\log\beta}{\phi_\pi} \\ \phi_\pi &> \frac{\kappa p - (1-p\beta)\sigma(1-p)}{\kappa} \end{aligned}$$

where we used the fact that $\log\beta < 0$. □

E.2 Shift in the Phillips curve

Proof. For inflation to be negative, it has to be that

$$\begin{aligned}
&\pi^L = -\frac{1}{(1-p\beta)\sigma(1-p)-\kappa p}\left[\kappa\log\beta - \frac{\sigma(1-p)}{\theta}\underline{u}\right] < 0 \\
\iff &\kappa\log\beta - \frac{\sigma(1-p)}{\theta}\underline{u} > 0 \quad \textit{and} \quad \frac{1-p\beta}{\kappa} > \frac{p}{\sigma(1-p)} \\
\iff &\underline{u} < \frac{\kappa\theta}{\sigma(1-p)}\log\beta \quad \textit{and} \quad \frac{1-p\beta}{\kappa} > \frac{p}{\sigma(1-p)}
\end{aligned}$$

□

E.3 Robust rule

Proof. 1. r^*-driven liquidity trap. The proof is direct.

2. Self-fulfilling liquidity trap. Recall that for a liquidity trap equilibrium to exist, we need $-\frac{1}{(1-p\beta)\sigma(1-p)-\kappa p}\left[\kappa\log\beta\right] < 0$. It thus has to be for a liquidity trap to not exist with an inflation-targeting TIP that $-\frac{1}{(1-p\beta)\sigma(1-p)\left(1+\frac{\varphi_\pi}{\theta}\right)-\kappa p}\left[\kappa\log\beta\right] > 0$. This immediately implies $\varphi_\pi > \theta\left(\frac{\kappa p}{(1-p\beta)\sigma(1-p)} - 1\right)$.

3. Phillips curve.

   Let's first consider the case where $\pi < \frac{\log\beta}{\phi_\pi}$ which corresponds to the case where the economy is at the ZLB with $i = 0$. In that case, inflation and output are given by

$$\pi^L = -\frac{1}{(1-p\beta)\sigma(1-p)\left(1+\frac{\varphi_\pi}{\theta}\right)-\kappa p}\left[\kappa\log(\beta) - \frac{\sigma(1-p)}{\theta}\underline{u}\right] \tag{62}$$

$$\hat{y}^L = \frac{1}{\sigma(1-p)}\left(p\pi^L - \log(\beta)\right) \tag{63}$$

   It is clear that π is monotonically increasing (decreasing in absolute terms) in φ_π and that $\hat{y}$ is monotonically increasing in φ_π.

   Let's then consider the case where $\pi > \frac{\log\beta}{\phi_\pi}$. In that case inflation and output

are given by

$$\pi^L = \frac{1}{(1-p\beta)\sigma(1-p)\left(1+\frac{\varphi_\pi}{\theta}\right)+\kappa(\phi_\pi-p)}\frac{\underline{u}\sigma(1-p)}{\theta} \tag{64}$$

$$\hat{y}^L = \frac{1}{\sigma(1-p)}\left(p-\phi_\pi\right)\pi^L \tag{65}$$

The threshold φ_π^{PC} is solution to

$$\pi = \frac{\log\beta}{\phi_\pi} = -\frac{1}{(1-p\beta)\sigma(1-p)\left(1+\frac{\varphi_\pi^{PC}}{\theta}\right)-\kappa p}\left[\kappa\log(\beta)-\frac{\sigma(1-p)}{\theta}\underline{u}\right]$$

$$\iff \varphi_\pi^{PC} = \theta\left[\frac{\frac{\phi_\pi}{\log\beta}\frac{\sigma(1-p)}{\theta}\underline{u}-\kappa(\phi_\pi-p)}{(1-p\beta)(1-p)\sigma}-1\right]$$

□

F Figures - Robust Rules

F.1 Prior distribution

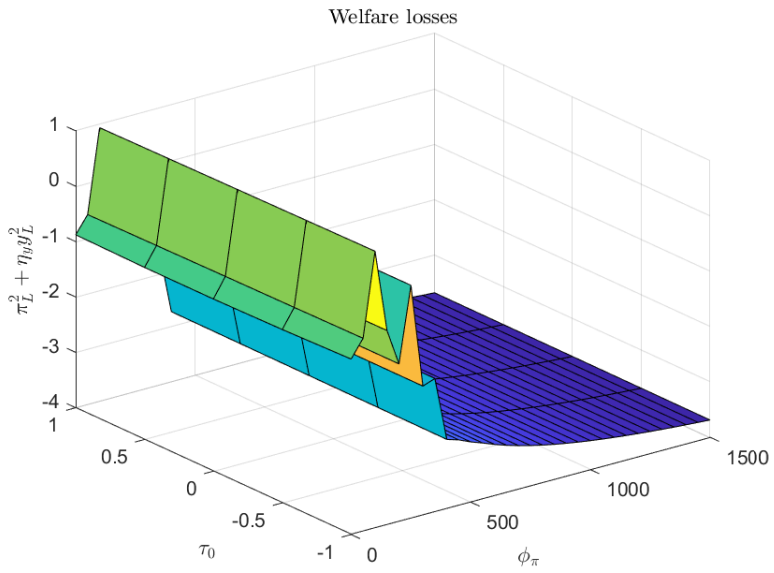


Figure 11: Welfare Losses - Prior Distribution

F.2 Maxmin

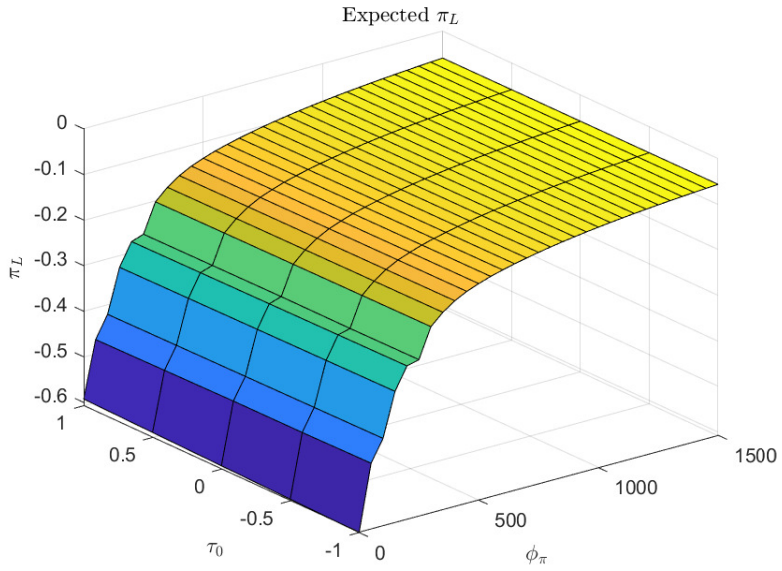


Figure 12: Inflation - Prior Distribution

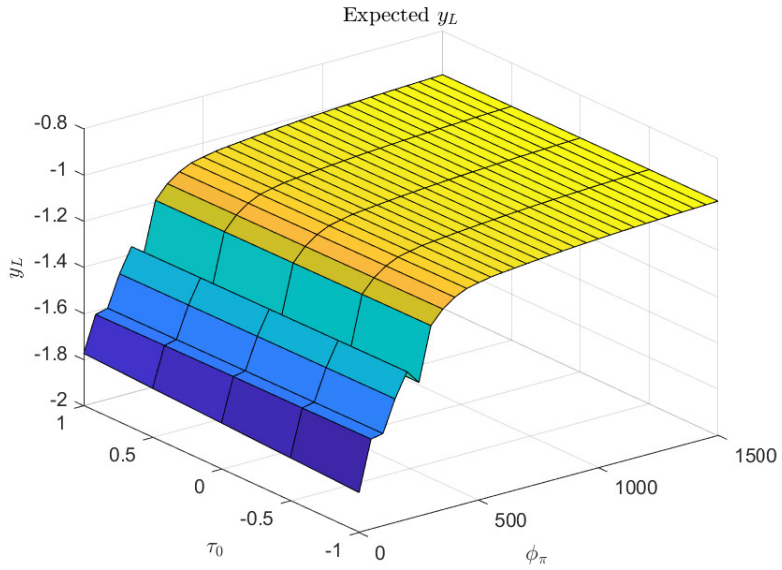


Figure 13: Output Gap - Prior Distribution

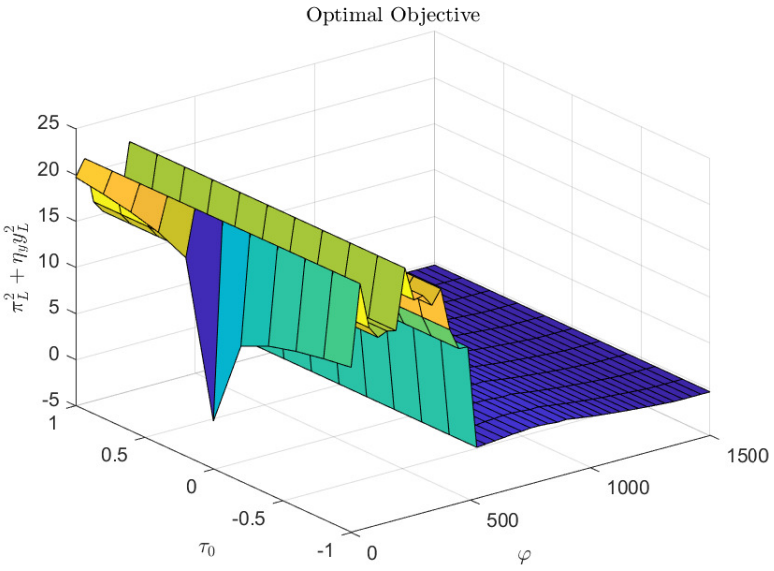


Figure 14: Welfare losses - Maxmin

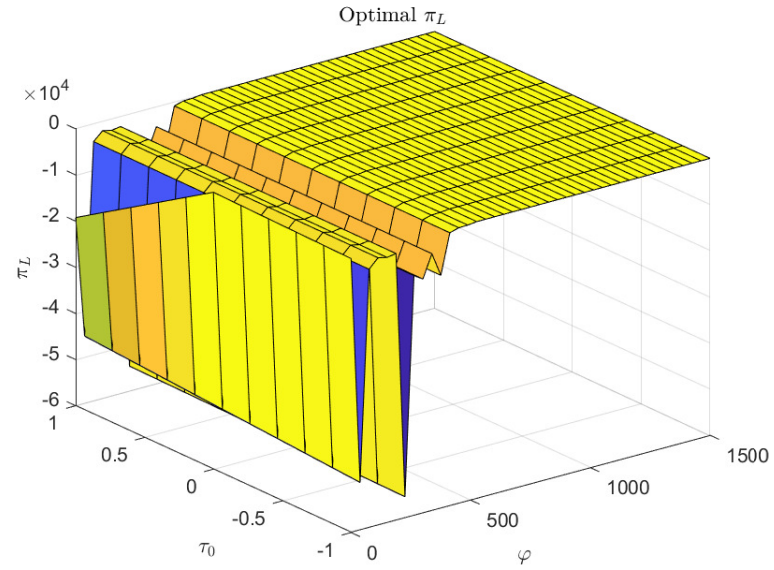


Figure 15: Inflation - Maxmin

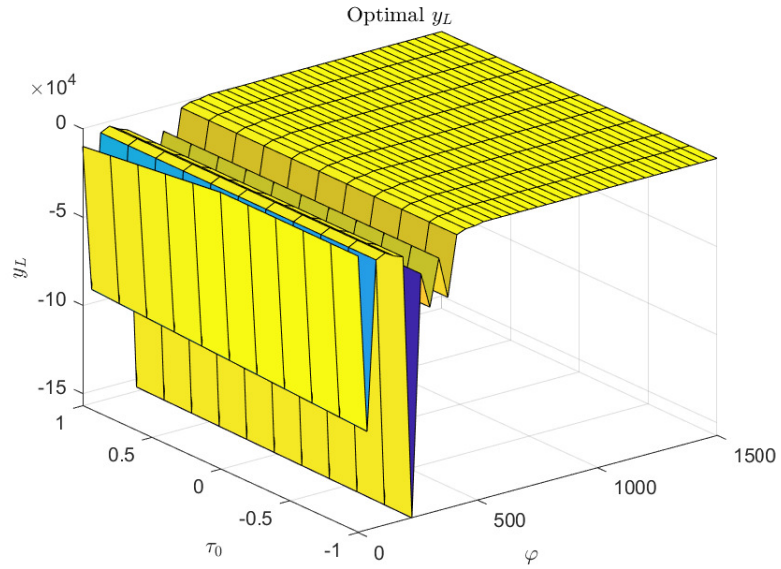


Figure 16: Output gap - Maxmin

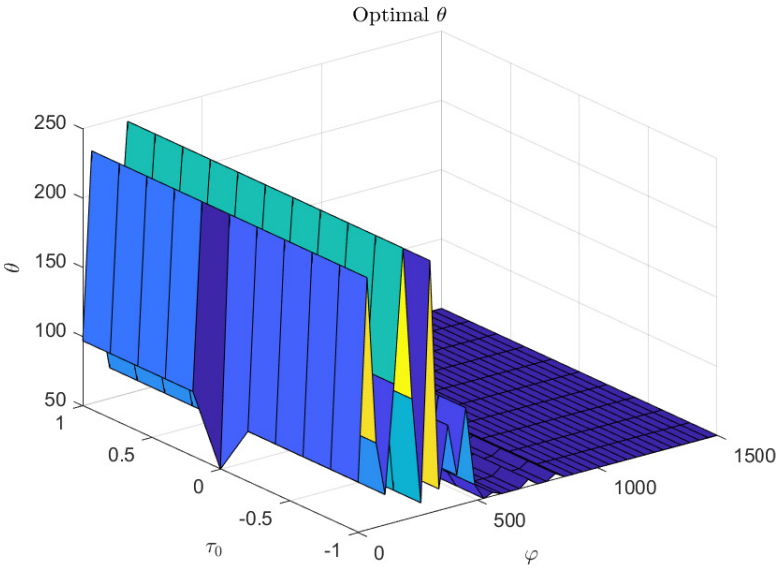


Figure 17: Optimal θ - Maxmin

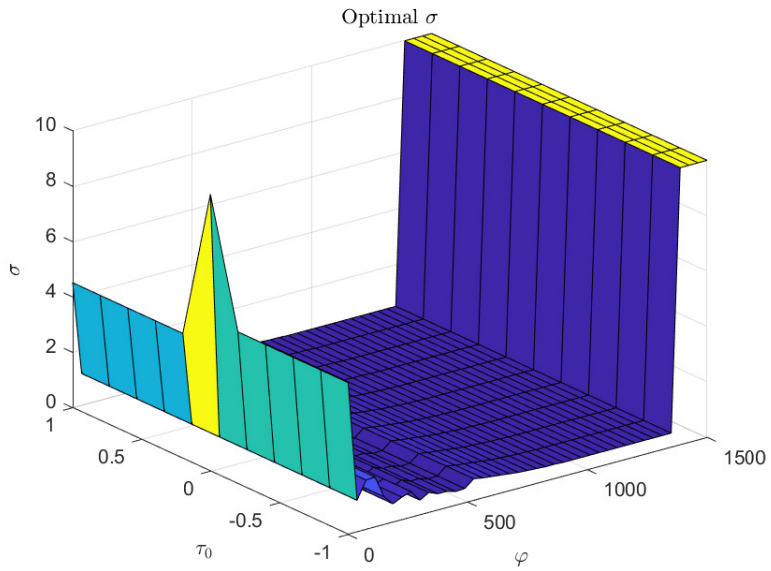


Figure 18: Optimal σ - Maxmin

F.3 Multiplier preference

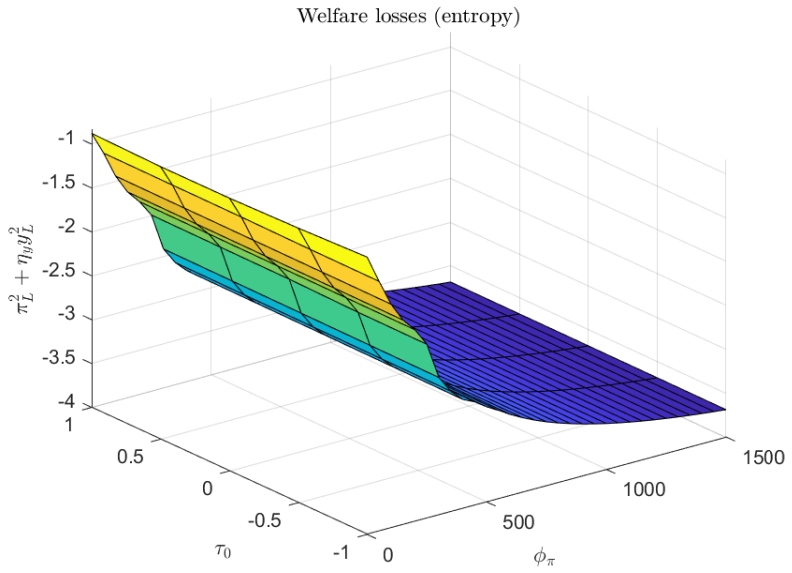


Figure 19: Welfare losses - Multiplier Preference

G Tables - Quantitative Model

Table 3: r^* decline

	Prior			Posterior			
	Dist.	Mean	Stdev.	Mean	Stdev.	HPD inf	HPD sup
e_w	invg	0.005	Inf	0.005	0.0004	0.0040	0.0055
e_p	invg	0.005	Inf	0.005	0.0004	0.0048	0.0062
e_g	invg	0.005	Inf	0.007	0.0006	0.0065	0.0084
e_μ	invg	0.005	Inf	0.138	0.0144	0.1143	0.1622
e_z	invg	0.005	Inf	0.014	0.0011	0.0124	0.0161
e_b	invg	0.005	Inf	0.053	0.0084	0.0372	0.0652
ρ_w	beta	0.500	0.1500	0.167	0.0577	0.0694	0.2557
ρ_p	beta	0.500	0.1500	0.129	0.0359	0.0667	0.1869
ρ_g	beta	0.500	0.1500	0.908	0.0222	0.8734	0.9452
ρ_μ	beta	0.500	0.1500	0.260	0.1150	0.0912	0.4514
ρ_b	beta	0.500	0.1500	0.770	0.0301	0.7267	0.8287
ρ_z	beta	0.500	0.1500	0.153	0.0458	0.0796	0.2316

Table 4: Sell-fulfilling

	Prior			Posterior			
	Dist.	Mean	Stdev.	Mean	Stdev.	HPD inf	HPD sup
e_w	invg	0.005	Inf	0.005	0.0024	0.0023	0.0086
e_p	invg	0.005	Inf	0.008	0.0013	0.0058	0.0096
e_g	invg	0.005	Inf	0.007	0.0005	0.0059	0.0075
e_μ	invg	0.005	Inf	0.160	0.0229	0.1234	0.1954
e_z	invg	0.005	Inf	0.015	0.0010	0.0128	0.0162
e_b	invg	0.005	Inf	0.072	0.0618	0.0011	0.1479
e_{sunspot}	invg	0.005	Inf	0.007	0.0029	0.0038	0.0110
ρ_w	beta	0.500	0.1500	0.499	0.2172	0.1734	0.7601
ρ_p	beta	0.500	0.1500	0.084	0.0444	0.0182	0.1520
ρ_g	beta	0.500	0.1500	0.940	0.0235	0.9022	0.9772
ρ_μ	beta	0.500	0.1500	0.237	0.0886	0.0955	0.3762
ρ_b	beta	0.500	0.1500	0.790	0.1488	0.5707	0.9748
ρ_Z	beta	0.500	0.1500	0.230	0.0562	0.1400	0.3196