Optimal Monetary Policy under Loose Commitment and the Zero Lower Bound

Nicolas Franz-Pattillo [*†‡]

November 12, 2026

Abstract

This paper studies optimal monetary policy when the central bank can renege on its plan with a fixed, exogenous probability (loose commitment), in a fully non-linear New Keynesian model with monopolistic competition, costly price adjustment, and an occasionally binding zero lower bound (ZLB). Solving the model with a global method, I trace welfare, inflation, and the incidence of the ZLB along the whole commitment line, from discretion to full commitment, holding the calibration fixed. Welfare improves monotonically with commitment, but the probability and duration of ZLB episodes do not: the probability of hitting the ZLB rises from 10% under discretion to 54% at intermediate commitment and falls to 39% near full commitment, as a falling inflation bias and increasingly credible forward guidance offset each other. As a consequence, higher average inflation does not always make the ZLB less likely, and the frequency or length of ZLB episodes cannot identify the central bank's level of commitment. When monopolistic competition creates an inflation bias, the ZLB lowers the welfare loss by restraining inflation; this gain vanishes under full commitment.

JEL Classification: E52, E58, E31, C63

Keywords: Optimal monetary policy; Zero lower bound; Loose commitment; Time inconsistency; Global solution methods

[*]nfranz@bcentral.cl

[†]The views expressed are those of the author and do not necessarily represent the views of the Central Bank of Chile or its board members.

[‡]Las opiniones expresadas son las del autor y no representan necesariamente las opiniones del Banco Central de Chile o de sus miembros del Consejo.

Resumen

Este trabajo estudia la política monetaria óptima cuando el banco central puede abandonar su plan con una probabilidad fija y exógena (compromiso imperfecto o *loose commitment*), en un modelo neokeynesiano completamente no lineal con competencia monopolística, costos de ajuste de precios y un límite inferior de cero para la tasa de interés nominal (ZLB) que se activa ocasionalmente. Resolviendo el modelo con un método global, describo el bienestar, la inflación y la incidencia del ZLB a lo largo de todo el rango de compromiso, desde la discreción hasta el compromiso total, manteniendo fija la calibración. El bienestar mejora de forma monótona con el compromiso, pero la probabilidad y la duración de los episodios de ZLB no lo hacen: la probabilidad de alcanzar el ZLB sube de 10% bajo discreción a 54% con un compromiso intermedio y cae a 39% cerca del compromiso total, porque la caída del sesgo inflacionario y la mayor credibilidad de las señales sobre la política futura (*forward guidance*) se compensan entre sí. En consecuencia, una mayor inflación promedio no siempre hace menos probable el ZLB, y la frecuencia o duración de los episodios de ZLB no permite identificar el grado de compromiso del banco central. Cuando la competencia monopolística genera un sesgo inflacionario, el ZLB reduce la pérdida de bienestar al contener la inflación; esta ganancia desaparece bajo compromiso total.

Clasificación JEL: E52, E58, E31, C63

Palabras clave: Política monetaria óptima; Límite inferior de cero; Compromiso imperfecto; Inconsistencia temporal; Métodos de solución global

1 Introduction

Optimal monetary policy has remained a central issue in macroeconomic theory for decades, and the debate acquired new urgency after the 2008 financial crisis pushed nominal interest rates in many advanced economies to zero. A near-zero policy rate limits a central bank's ability to stimulate activity through conventional means, and has renewed the question of whether higher average inflation — and therefore higher average nominal rates — would make future encounters with the zero lower bound (ZLB) less likely and less costly.

This paper asks a related but different question: rather than how ZLB risk changes with average inflation, I ask how it changes with the central bank's *credibility*. I develop and solve a fully non-linear small-scale New Keynesian model with monopolistic competition, costly price adjustment, and an occasionally binding ZLB, in which the monetary authority commits to a plan but is exogenously and randomly allowed to reset it — the "loose commitment" device of Schaumburg and Tambalotti (2007) and Debortoli and Nunes (2010), in the form used by Bodenstein et al. (2012). Because the model is solved with a global method rather than linearized, I can trace welfare, inflation, and the incidence of the ZLB continuously along the entire commitment line, from pure discretion to full commitment, holding every other parameter fixed — a comparison not documented thus far. Two results stand out. First, welfare losses fall monotonically as commitment increases, but the probability and duration of ZLB episodes do not: both are hump-shaped in the degree of commitment, peaking at intermediate credibility rather than at either extreme. Second, and for the same reason, higher average inflation does not uniformly make the ZLB less likely. Solving the model without linearizing is what allows these results to be seen at all: inflation's mean and volatility shift along the commitment line in ways a linear-quadratic approximation around a single steady state cannot capture, and it is precisely the shifting balance between an inflationary and a deflationary bias that generates the non-monotonicities above. Eggertsson and Singh (2019) find the error from log-linearizing at the ZLB to be modest, even in experiments calibrated to the Great Depression, but they approximate around a single zero-inflation steady state, with a two-state shock and policy fixed as a strict inflation target. In the comparisons below the point of approximation itself moves — average inflation and the distance to the bound both change with the level of commitment — so the approximation error need not be constant across the economies being ranked.

The starting point for this question is the time-inconsistency problem identified by Kydland and Prescott (1977): a policy that is optimal when announced need not remain optimal once private expectations have formed around it. Calvo (1978) established the same inconsistency for optimal policy in a monetary economy, Barro and Gordon (1983) formalized the inflation bias it produces — once the public understands the policymaker's objectives, surprise inflation cannot occur systematically, so discretion delivers excessive inflation with no systematic gain in activity — and Taylor (1993) argued that simple, transparent instrument rules provide a useful practical

benchmark for policy, which is why they — rather than the optimal-policy problem solved here — are the workhorse of applied monetary policy analysis. Blinder (2000) argued that credibility itself, built from central bank independence and, above all, a record of promises kept, is what makes commitment valuable in the first place; the continuous measure of commitment used in this paper is designed to speak directly to that idea.

A separate strand of the literature solves directly for optimal policy under a specific commitment assumption rather than a rule. Siu (2004) solves the Ramsey problem for monetary and fiscal policy in a sticky-price economy with distortionary taxation and finds that, for post-war spending processes, even small price rigidities make it optimal to use little state-contingent inflation as a fiscal shock absorber, and Yun (2005) characterizes optimal monetary policy under staggered price setting without relying on linearly approximated equilibrium conditions, finding complete stabilization of the price level optimal in the absence of initial price dispersion. Both study policy under commitment, although in Yun's economy, where a tax scheme removes the monopolistic distortion, the optimal policy is time-consistent and discretion delivers the same allocation; the opposite pole of time-consistent discretion is studied by Adam and Billi (2007), who — as Adam and Billi (2006) does for commitment — removes the inflationary bias created by monopolistic competition by assuming an output subsidy that makes the steady state efficient. Debortoli and Nunes (2010) instead let the policymaker renege on its plan with a fixed exogenous probability, the loose-commitment device adopted here. Since neither pole is realistic and imperfect credibility is hard to measure directly, Bodenstein et al. (2012) characterize optimal policy at the ZLB under imperfect credibility in a linearized model, and infer the credibility of the Federal Reserve and the Riksbank from the gap between their announced interest-rate paths and market expectations — cautioning that an announced period of low rates need not signal high credibility. My results sharpen that caution from the other side: in a fully non-linear setting the probability and duration of ZLB episodes are non-monotonic in commitment, so neither statistic identifies the commitment level on its own. Nunes et al. (2021) study welfare gains from credibility in a similarly linearized environment without a ZLB and find them to be approximately linear in the initial credibility level, and find the output–inflation stabilization trade-off to be non-monotonic in credibility; I find that non-linearity — and the ZLB in particular — is exactly what breaks that linearity once the lower bound can bind.

The ZLB literature itself has largely relied on linear-quadratic approximations around a zero-inflation steady state, with the bound itself handled globally. Eggertsson and Woodford (2003) show that a credible commitment to history-dependent policy — implemented as a time-varying price-level target — largely mitigates the distortions created by the bound, and Adam and Billi (2006) finds that the same promise of higher future inflation shortens and softens ZLB episodes under commitment, while Adam and Billi (2007) shows how much deeper the output losses and deflation become when the same economy is run under discretion. Nakov (2008), who solves the linearized model globally with an explicit occasionally binding non-negativity constraint rather than under

perfect foresight, finds that commitment reduces unconditional welfare losses to about a tenth of their discretionary level, and that constant price-level targeting recovers much — though not all — of that gain. More recently, Bonciani and Oh (2025) study the optimal mix of forward guidance and quantitative easing at the ZLB in a log-linearized, perfect-foresight model, again comparing only discretion and full commitment, and find the welfare losses under discretion to be substantially larger. Here, by contrast, the policy rate is the only instrument, and what varies is how credible the promise of future accommodation is. These results are informative about the mechanism but, by construction, cannot speak to how it changes away from the linearization point, which is exactly what happens as commitment and its associated inflation bias vary. Nakata (2016) solves a fully non-linear model of optimal fiscal and monetary policy at the ZLB, in which uncertainty creates a time-consistency problem precisely when the steady state is inefficient, and compares discretion and commitment directly, and shows that when monopolistic competition creates an inflation bias the ZLB raises welfare under discretion by restricting the options of the policymaker; it is the closest antecedent to this paper, and I extend the comparison to the continuum in between, holding the calibration fixed across commitment levels. On how costly inflation actually is away from zero, Blanco (2021) finds an optimal inflation target of 3.5% in a quantitative menu-cost model, because firm-level shocks make the effect of inflation on resource misallocation small, and Adam and Weber (2019) show that with systematic firm-level productivity trends the efficient rate of inflation generically differs from zero and responds to productivity disturbances — the result that motivates the free-of-cost rate of inflation π_f used in the calibration below. Gust et al. (2017) estimate a non-linear model with an occasionally binding bound and quantify how much of the fall in output it accounts for, underscoring that the welfare consequences of the ZLB are as much an empirical as a theoretical question. Taken together, this literature suggests both the inflationary and the deflationary bias studied here are quantitatively meaningful, but that their relative size along the commitment line cannot be settled without solving the model non-linearly.

I follow Rotemberg (1982a) in modeling price rigidity as a cost of adjusting prices, and much of this literature in using a small-scale New Keynesian model rather than a medium-scale one — the goal is to isolate how commitment interacts with the ZLB and monopolistic competition, not to fit the data; see Rupert and Šustek (2019) on the mechanics of the transmission mechanism in models of this class, and Auclert et al. (2024) on the first-order equivalence between menu-cost and Calvo pricing once the model is linearized. Solving the model without linearizing is, however, non-trivial: the loose-commitment equilibrium has no non-stochastic steady state, since the endogenous promise-state resets to zero every time the plan is reneged, so the relevant ergodic distribution is unknown a priori and depends on the commitment level itself. I solve this with time iteration (Coleman, 1991) on a grid whose nodes are placed according to an estimate of that ergodic distribution, refined iteratively as in Maliar and Maliar (2015); details are in the Appendix.

The main results follow from two competing forces along the commitment line. Higher commitment lowers the inflationary bias created by monopolistic competition, which, for a given real rate,

pushes nominal rates down and makes the ZLB more likely; at the same time, higher commitment makes forward guidance more credible, which makes the ZLB less likely and shortens the episodes that do occur. Close to discretion the first force dominates; close to full commitment the second does. The result is a ZLB probability, ZLB duration, and inflation volatility that all peak at intermediate commitment rather than moving monotonically between the discretion and commitment benchmarks. Average inflation, by contrast, falls monotonically from discretion to commitment, and welfare improves monotonically throughout, since the ZLB's own contribution to the welfare loss is second-order relative to the inflation bias it helps to offset under discretion. Close to discretion, forward guidance is not credible, so a more loosely committed economy instead relies on a higher average level of inflation to keep the ZLB rare and short-lived; under commitment, the central bank can instead keep normal-times inflation low and raise it only — but credibly — during ZLB episodes, which is what makes inflation volatility, and not just its level, vary non-monotonically with commitment.

The rest of the paper is organized as follows: Section 2 details the model, Section 3 the solution method, Section 4 the calibration, Section 5 the results, and Section 6 concludes.

2 The Model Economy

In this paper, I will adhere to the existing literature, employing a small-scale New Keynesian economy framework characterized by price rigidity, as established in Rotemberg (1982a), and closed to the rest of the world. The equilibrium of the private sector will yield an Euler equation, an intra-period labor consumption decision, a non-linear Phillips curve, and a resource constraint. The Central Bank will operate within this equilibrium, determining monetary policy by setting the nominal interest rate, which must remain above zero at all times. Given the time inconsistency of the optimal plan, I will adopt a loose commitment technology, following Debortoli and Nunes (2010) and Bodenstein et al. (2012).

The economy is populated by a representative household, final-good firms, a continuum of intermediate-good firms owned by the household, and a central bank. The only asset is a one-period nominal risk-free bond; with a representative household it is in zero net supply, so whether asset markets are complete plays no role. Relative to the standard small-scale model with an occasionally binding ZLB solved non-linearly (e.g., Nakata, 2016), the central bank operates under loose commitment instead of full commitment or discretion, the monopolistic-competition distortion is not offset by a subsidy in the benchmark, and the costless rate of price change π_f is a parameter, which some experiments use to shift average inflation.

2.1 Private Sector

In this section, I explain how private agents behave in the economy. I assume the private sector knows the central bank's behaviour and therefore sets its rational expectation taking the monetary

authority plan as given.

The private sector comprises a representative household, intermediate input firms, and final goods firms.

The uncertainty of the model comes from two shocks which I will analyze separately[1]: a total-factor-productivity shock, which will affect homogeneously the final goods firms, and a preferences shock, which changes the effective household discount factor. As we will see, these two shocks have different implications regarding the zero lower bound.

Households

The representative household is infinitely lived and chooses in each period how much to consume, work and save for the next period to maximize its expected discounted utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \left[\Pi_{k=0}^{t-1} \delta_k\right] \beta^t \left[\frac{c_t^{1-\sigma} - 1}{1-\sigma} - \chi \frac{n_t^{1+\eta}}{1+\eta}\right] \tag{1}$$

The variables c_t and n_t are consumption and total hours worked in period t, respectively. The parameter β represents the constant part of the discount factor, which will be effectively $\delta_t \beta$: the shock δ_t is known in period t and discounts utility between t and $t+1$ (the product is equal to one for $t = 0$). δ_t is one of the sources of uncertainty of the model and represents a shock that makes the household more (less) patient. I will assume this shock to be log-autoregressive, i.e. $\delta_t = \delta_{t-1}^{\rho_\delta} e^{\sigma_\delta \epsilon_t}$, with the innovation ϵ_t distributed as standard normal and $|\rho_\delta| < 1$.

In the expected utility above, σ corresponds to the inverse of elasticity of inter-temporal substitution, $1/\eta$ is the Frisch elasticity of labor supply, and χ scales the dis-utility from work. The maximization of (1) is constrained by the household budget given by (2)

$$p_t c_t + b_t = p_t w_t (1+\tau) n_t + R_{t-1} b_{t-1} + p_t d_t + T_t \tag{2}$$

Here p_t is the price of one unit of consumption. The gross inflation at time t is defined as $\pi_t = p_t / p_{t-1}$. Saving and borrowing are made through the nominal instrument b_t, which is risk-free and has a gross return R_t. The household does not face constraints in b_t. The aggregate labor supply of the household is distributed on a continuous set of intermediate firms indexed by j, each offering real wages $w(j)$. The household maximizes its real labor income $w_t n_t = \int_{j=0}^{1} w(j) n(j) \mathrm{d}j$ by choosing how much time to allocate in each firm. The household is the owner of the intermediate good firms which operate in a monopolistic competitive environment. The real profits generated are transferred each period via dividends $d_t = \int_{j=0}^{1} d(j) \mathrm{d}j$. The parameter τ is a labor input subsidy from the government with a balanced budget. The government revenues are collected through lump sum taxes T_t.

[1]Solving the loose-commitment problem jointly with both shocks active would add a further exogenous state dimension on top of the two endogenous "promise" states already required (Section 3), compounding the curse of dimensionality that motivates this paper's grid-construction method; see the Conclusion for further discussion.

The optimal labor allocation of the household is to work for the firm that offers the highest wage and would be indifferent between firms with equal salaries. In this case, it will be assumed symmetry across firms.

Upon integration, the first-order conditions are summarized by the following system.

$$
\begin{aligned}
R_t^{-1} &= \delta_t \beta \mathbb{E}t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{1}{\pi_{t+1}} \right] && (3) \\
\chi n_t^{\eta} &= c_t^{-\sigma} w_t (1+\tau) && (4) \\
c_t &= w_t n_t + \frac{R_{t-1} b_{t-1} - b_t}{p_t} + d_t && (5)
\end{aligned}
$$

Equation (3) is the standard Euler equation, (4) accounts for the intra-temporal consumption/leisure choice, and (5) is the budget constraint expressed in real terms. Here w_t is the highest wage in the market, and n_t is the integration of the hours worked for the firms offering it.

I define the nominal stochastic discount factor of the household as

$$
Q_{t+1|t} = \delta_t \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{1}{\pi_{t+1}} \tag{6}
$$

This discount factor will be relevant for the intermediate firms since the household owns them.

Final good firms

One firm that aggregates the imperfectly substitutable intermediate goods produces the consumption good. There is free entry to this market which implies perfect competition.

With these assumptions, the problem of this firm is:

$$
\max_{\{y_t(j)\}_{j=0}^{1}} p_t y_t - \int_0^1 y_t(j) p_t(j) \mathrm{d}j \tag{7}
$$

Where p_t and y_t are the price and quantity of the final good, and $p_t(j)$ and $y_t(j)$ are the price and quantity of the input j. I assume the inputs to be imperfect substitutes and use the technology proposed by Dixit and Stiglitz (1977) with constant elasticity of substitution ν^{-1}

$$
y_t = \left(\int_0^1 y_t(j)^{1-\nu} \mathrm{d}j \right)^{\frac{1}{1-\nu}} \tag{8}
$$

The solution of the problem in (7) generates the input demand (9), and the zero-profit condition of a competitive market yields the aggregate price in (10)

$$y_t(j) = \left[\frac{p_t(j)}{p_t}\right]^{-1/\nu} y_t \tag{9}$$

$$p_t = \left(\int_0^1 p_t(j)^{\frac{\nu-1}{\nu}} \mathrm{d}j\right)^{\frac{\nu}{\nu-1}} \tag{10}$$

Intermediate good firms

As intermediate goods are assumed to be imperfect substitutes for each other, firms producing them will have some degree of market power. In this monopolistic competition environment, firm j chooses the price and quantity of its unique good, subject to the available technology, demand, and costs.

Since the household owns the firm, its objective is to maximize the discounted flow of profits taking the discount factor in (6) as given.

The production technology of each intermediate good is assumed to follow $y_t(j) = z_t n_t(j)$, where $n_t(j)$ is the labor input provided by the household to firm j, and z_t is a technology shock that affects all the firms equally. This shock is assumed to be autoregressive following $z_t = z_{t-1}^{\rho_z} e^{\sigma_z \varepsilon_t}$, with the innovation ε_t distributed as standard normal and $|\rho_z| < 1$.

The labor market is assumed to be competitive, so wages are taken as given and $w(j) = w(i) = w_t$ for firms $i \neq j$.

Following Rotemberg (1982b), I assume the intermediate firms face a price adjustment cost relative to their scale. Ireland (2007) presents the most popular version of this adjustment cost, where the free-of-cost change in prices is given by $\pi_{t-1}^{\mu} \pi_f^{(1-\mu)}$, with $0 \leq \mu \leq 1$. To reduce the dimensionality of the firm's problem, I assume $\mu = 0$, which implies a purely forward-looking pricing function[2]. In particular, the adjustment costs are given by (11), where π_f is the "costless" price change. In the literature, this parameter is often set to one, implying that zero inflation is efficient (i.e., no loss of output). This might not be the case when richer exit and entrance firm dynamics are imposed[3]. As in Blanco (2021), π_f ensures that the "optimal" level of inflation is substantially above zero[4].

$$AC_t(j) = \frac{\phi}{2}\left[\frac{p_t(j)}{p_{t-1}(j)} - \pi_f\right]^2 y_t \tag{11}$$

The problem of firm j is

[2]Note that this decision makes inflation less sluggish. The benefit is to have a closed-form solution for the problem.
[3]See Adam and Weber (2019)
[4]Effectively, I will use π_f in experiments to exogenously control how often an economy falls into the ZLB.

$$\max_{\{p_{t+s}(j)\}_{s=0}^{\infty}} \mathbb{E}_t \sum_{s=0}^{\infty} Q_{t+s|t} \left[p_{t+s}(j) y_{t+s}(j) \right.$$
$$\left. - p_{t+s} w_{t+s} n_{t+s}(j) - p_{t+s} AC_{t+s}(j) \right] \tag{12}$$

Where $Q_{t+1|t}$ is given by (6), $y_t(j)$ by (9) and $n_t(j) = y_t(j)/z_t$ for all firm j and time t.

As usual, I focus on the symmetric equilibrium where all firms choose the same price. Therefore $p_t(j) = p_t$ and $y_t(j) = y_t$. In combination with the first-order conditions of (12), I get the non-linear Phillips curve:

$$1 - \nu - \frac{w_t}{z_t} + \phi\nu\pi_t \left(\pi_t - \pi_f\right) = \phi\nu\delta_t\beta\mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{y_{t+1}}{y_t} (\pi_{t+1} - \pi_f)\pi_{t+1} \right] \tag{13}$$

Note that ϕ controls the price rigidity in this economy. When this parameter is equal to zero, each firm sets the price of its good equal to the nominal marginal cost times a constant markup $(1-\nu)^{-1}$.

Market clearing, private sector equilibrium and efficient allocations

Given any monetary policy, the first-order conditions and the market clearing conditions imply that:

$$c_t^{-\sigma} R_t^{-1} = \delta_t \beta \mathbb{E}_t \left[\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right] \tag{14}$$

$$c_t^{-\sigma} y_t \left[1 - \nu - \frac{w_t}{z_t} + \phi\nu\pi_t \left(\pi_t - \pi_f\right) \right] = \phi\nu\delta_t\beta\mathbb{E}_t \left[c_{t+1}^{-\sigma} y_{t+1} (\pi_{t+1} - \pi_f)\pi_{t+1} \right] \tag{15}$$

$$\chi n_t^{\eta} = c_t^{-\sigma} w_t (1+\tau) \tag{16}$$

$$c_t = y_t \left[1 - \frac{\phi}{2} \left(\pi_t - \pi_f\right)^2 \right] \tag{17}$$

$$y_t = z_t n_t \tag{18}$$

$$z_t = z_{t-1}^{\rho_z} e^{\sigma_z \varepsilon_t} \tag{19}$$

$$\delta_t = \delta_{t-1}^{\rho_\delta} e^{\sigma_\delta \epsilon_t} \tag{20}$$

The first and second equations are the Euler and non-linear Phillips curves. These are the only

ones depending on private expectations.

The third equation is the labor-leisure decision of the household. The fourth is the aggregate resource constraint, obtained with the household budget constraint, the zero net supply of bonds, the government budget $\tau n_t = T_t$, and the integration of intermediate firms' dividends. The last three equations are the production technology, total factor productivity, and preferences processes.

It is important to emphasize that when there's no cost of adjusting prices ($\phi = 0$), the monopolistic competition is erased ($\nu = 0$), and there is no labor subsidy ($\tau = 0$), the allocation of labor and consumption is efficient and the monetary policy is irrelevant. Under these circumstances, labor and consumption allocations are given by:

$$n_{t,\text{flex, comp}} = \left[\frac{z_t^{1-\sigma}}{\chi}\right]^{1/(\eta+\sigma)} \tag{21}$$

$$c_{t,\text{flex, comp}} = \left[\frac{z_t^{1+\eta}}{\chi}\right]^{1/(\eta+\sigma)} \tag{22}$$

By adding back the monopolistic competition (i.e., flexible price equilibrium), the wedge generated by the lack of competition can be off-set by a specific labor subsidy level ($\tau = \frac{\nu}{1-\nu}$). The consumption-leisure allocation is no longer efficient without this peculiar subsidy — and subsequent lump sum transfer. The measure of the inefficiency is captured by the term $(1+\tau)(1-\nu)$:

$$n_{t,\text{flex}} = \left[\frac{(1-\nu)(1+\tau)z_t^{1-\sigma}}{\chi}\right]^{1/(\eta+\sigma)} \tag{23}$$

$$c_{t,\text{flex}} = \left[\frac{(1-\nu)(1+\tau)z_t^{1+\eta}}{\chi}\right]^{1/(\eta+\sigma)} \tag{24}$$

The optimal tax τ^* will always try to eliminate the distortion induced by ν, and the labor-consumption decision of the household will equalize marginal benefits and costs of working.

Finally, in the case of perfect competition ($\lim \nu \to 0$) but adjustment costs ($\phi > 0$), there is no distortion in the leisure-consumption relationship. Still, there's a level effect on output every time that π_t departs from π_f:

$$c_t^{-\sigma} R_t^{-1} = \beta \delta_t \mathbb{E}_t \left[\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right] \tag{25}$$

$$c_t = \left\{ \frac{z_t^{1+\eta}}{\chi} \left[1 - \frac{\phi}{2} \left(\pi_t - \pi_f \right)^2 \right]^{\eta} \right\}^{1/(\eta+\sigma)} \tag{26}$$

$$n_t = \left[\frac{c_t^{-\sigma} z_t}{\chi} \right]^{\frac{1}{\eta}} \tag{27}$$

Without a lower bound in the nominal interest rate, it's evident that the flexible price allocation can be attained by targeting $\pi_t = \pi_f$. In that case, the Euler equation defines the nominal interest rate the monetary authority should target —given private sector expectations. On the other hand, when the interest rate is occasionally constrained, the flexible price competitive equilibrium allocation is no longer available under rational expectations. To see this, note that $\pi_t = \pi_f$ is no longer supported for all states. I will refer broadly to the monetary policy in the next section.

2.2 Monetary Policy

I assume the monetary policy instrument is the nominal interest rate, which is constrained from below by the Zero Lower Bound (ZLB).

$$R_t \geq 1 \qquad \forall\ t \tag{28}$$

The problem of the central bank is to maximize the household welfare constrained by the set of attainable equilibria, which consists of a policy $\{R_t\}_{t=0}^{\infty}$, allocations $\{c_t, n_t, y_t, y_t(i), n_t(i)\}_{t=0}^{\infty}$, and prices $\{w_t, p_t, p_t(i), w_t(i)\}_{t=0}^{\infty}$ — where $i \in [0,1]$ — such that, for stochastic processes ε_t and ϵ_t, and initial level of prices, Total Factor Productivity, and preferences shock the system of equations (14) to (20) and (28) are satisfied.

This is a time-inconsistent problem since the monetary authority deals with three wedges that cannot be closed simultaneously all the time. The promise to close any of them is not credible since the policy's attention will drift to the other two once one wedge disappears. As a consequence, the original promise will change over time.

The first wedge is the consequence of price adjustment cost, the second due to monopolistic competition, and the third is induced by the Zero Lower Bound.

To understand how monopolistic competition and costly price adjustment interact, assume the Zero Lower Bound does not constrain the monetary policy. This implies that the Euler equation is not binding, i.e., the target interest rate can be achieved given any expectations. Monopolistic competition generates a loss since the allocation of consumption is lower than in perfect competition. Now assume the Central Bank promises to minimize the loss given by the price adjustment cost targeting next period inflation equal to the free-of-cost $\pi_{t+1} = \pi^f$. If the private sector believes

this promise and sets expectations accordingly, the Central Bank is incentivized to renege to offset the monopolistic loss. This time inconsistency generates, in equilibrium, higher inflation than π_f and, therefore, a positive inflation bias.

The zero lower bound wedges, on the other hand, generate a deflation bias. When the economy hits the bound, the Central Bank would like to promise higher future inflation to get out faster. Nonetheless, when the time comes, the incentives change to keep inflation at its free-of-cost level.

A "first best" allocation could be achieved if the monetary authority had the technology to set a once-and-for-all plan. This plan would internalize the shift in the authority's preferences for the wedges, and it would fix them at an initial time. If the technology does not exist, the literature typically analyses the opposite case: no promise is binding in the future, and the central bank is discretionary every period. The fact that we observe successful inflation targets in the world might be a good indicator of something in the middle. In the following subsections, I pay attention to all these alternatives.

For convenience, throughout the rest of the paper, I will call the set of endogenous variables in the model as $x_t = \{\pi_t, c_t, n_t, R_t, w_t, y_t\}$.

Discretion

As mentioned before, under a discretionary policy, the promises made at time t are not binding at time $t + j$ with $j > 0$. The Central Bank will re-calculate its problem each period with no implication to future states or expectations[5]. Therefore, the relevant part of the objective function will contain present terms only.

With this in mind, the problem of the Central Bank is to maximize the current household utility subject to the attainable equilibrium defined above (equations 14 to 20 and 28).

$$\max_{x_t} u(x_t) = \frac{c_t^{1-\sigma} - 1}{1-\sigma} - \chi \frac{n_t^{1+\eta}}{1+\eta} \tag{29}$$

As is commonly known in the literature[6], the equilibrium to the discretionary problem is not unique. Therefore I concentrate on a particular equilibrium concept: Markov-Perfect with the zero lower bound occasionally binding.[7]

Commitment

In the case of commitment, the central bank chooses a plan at the beginning of times for all possible TFP and preferences shock histories.[8] The plan maximizes (1), subject to the attainable

[5]Note the only state variables in the model with discretion are the exogenous total factor productivity and preferences processes.

[6]See Chari and Kehoe (1990)

[7]As pointed by Armenter (2008), an economy as the one I describe contains at least two Markov equilibria. One where the low inflation expectations create a trap that perpetuates the ZLB, and another where the ZLB binds occasionally.

[8]This means $\{x_t(z^t)\}_{t=0}^{\infty}$, where $z^t = \{z_0, \cdots, z_t\}$.

equilibrium (equations 14 to 20 and 28).

The Lagrangian associated with this problem can be written as:

$$\mathcal{L}(x_t, \lambda_t) : \mathbb{E}_0 \sum_{t=0}^{\infty} \left[\Pi_{k=0}^{t-1} \delta_k\right] \beta^t \Bigg\{ u(x_t) + \lambda_t^{\mathrm{E}} \left[c_t^{-\sigma} R_t^{-1} - \delta_t \beta \mathbb{E}_t \left(\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right) \right]$$

$$+ \lambda_t^{\mathrm{P}} \left[c_t^{-\sigma} y_t \left(1 - \nu - \frac{w_t}{z_t} + \phi \nu \pi_t \left(\pi_t - \pi_f \right) \right) \right. \tag{30}$$

$$\left. - \phi \nu \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} y_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1} \right] \right] + \lambda_t^{\mathrm{R}} \left[R_t - 1 \right] \Bigg\}$$

Where I omit the terms and Lagrange multipliers associated with equations holding with equality without loss of generality. λ_t^{E}, λ_t^{P}, and λ_t^{R} are the Lagrange multiplier associated with the Euler equation, Phillips curve, and Zero Lower Bound, respectively.

Marcet and Marimon (2019) shows that as long as (30) is well defined and has a non-empty solution set, it can be converted into a saddle-point functional equation given by

$$W(z, \delta, \mu) = \min_{\{\lambda\}} \max_{\{x\}} f(x, z, \lambda, \mu) + \delta \beta \mathbb{E} W(z', \delta', \mu') \tag{31}$$

where $\mu = \left\{\mu^E, \mu^P\right\}$, $\lambda = \left\{\lambda^E, \lambda^P\right\}$, $\mu' = \lambda$ with initial value of zero, and

$$\begin{aligned} f(x, z, \lambda, \mu) = {} & u(x) + \lambda^E \left(\frac{c^{-\sigma}}{R} \right) - \mu^E \left(\frac{c^{-\sigma}}{\pi} \right) \\ & + \lambda^P \left[c^{-\sigma} y \left(1 - \nu - \frac{w}{z} + \phi \nu \pi \left(\pi - \pi_f \right) \right) \right] \\ & - \mu^P \left[\phi \nu c^{-\sigma} y (\pi - \pi_f) \pi \right] \end{aligned} \tag{32}$$

The inherited multipliers enter with a negative sign because they carry the forward-looking terms of last period's constraints: since period-t utility is discounted by $\delta_{t-1}\beta$ relative to period $t-1$, the term $-\lambda_{t-1}\delta_{t-1}\beta B_t$ of the Lagrangian becomes $-\mu_t B_t$ in period t.

This is relevant because it points out how this problem can be solved — e.g., value function or time iteration.

To get a sense of how this saddle point functional equation works, notice that the promises and concessions of time t, reflected in equilibrium by the Lagrange multipliers of the Phillips and Euler equations, will become the state in time $t+1$, enforcing its fulfillment. The optimal plan tries to minimize such promises given the maximization of the objective function and its continuation value.

A Ramsey equilibrium is a time-invariant set of functions $W(z, \delta, \mu)$, $c^\star(z, \delta, \mu)$,$n^\star(z, \delta, \mu)$,

$y^\star(z,\delta,\mu)$, $\pi^\star(z,\delta,\mu)$, $w^\star(z,\delta,\mu)$, $R^\star(z,\delta,\mu)$, solving the problem (31) constrained by equations (14)-(20) and (28). Marcet and Marimon (2019) show that this is true as long as the solution satisfies their intertemporal consistency condition.[9] Nonetheless, as they point out in the working-paper version, this condition "only needs to be checked *ex-post*, not imposed *ex-ante*, which makes a difference computationally" (Marcet and Marimon, 2011, p. 21).

Loose Commitment

The imperfect or loose commitment is set up as in Debortoli and Nunes (2010) and Bodenstein et al. (2012). I assume there is an exogenous probability $1-\xi$ that the central bank reneges its promises and re-optimizes to generate a new contingent plan. Therefore, at time zero, the planner's problem is to set a contingent plan for a future history of commitments.

Define the random variable s_t which can take two values $s_t \in \{\mathrm{c},\mathrm{r}\}$, where "c" and "r" stand for *commit* and *renege* respectively. Assume s_t to be independent and identically distributed, with probability to draw *commit* $\text{prob}(\mathrm{c}) = \xi$, and the chances to draw *renege* being $\text{prob}(\mathrm{r}) = 1-\xi$.

Under the previous assumptions, the probability of having a history of commitment[10] at time t is ξ^t, and reneging at time t conditional on a history of commitment has probability $\xi^{t-1}(1-\xi)$.

Characterize x_t^{c} as the set of endogenous variables under the committed plan, and x_t^{r} its reneged version. Moreover, let V_t^{c} and V_t^{r} to denote the value of committing or reneging in period t, respectively.

Conditional on $s_0 = \mathrm{c}$ and given the expected value of reneging in the next period V_{t+1}^{r}, we can write the problem of the central bank as:

$$V_0^{\mathrm{c}} = \max_{\{x_t^{\mathrm{c}}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \left[\Pi_{k=0}^{t-1}\delta_k\right](\beta\xi)^t \left[u\left(x_t^c\right) + \delta_t\beta(1-\xi)\mathbb{E}_t V_{t+1}^{\mathrm{r}}\right] \tag{33}$$

With the maximization constrained by the attainable equilibrium characterized by equations 14 to 20 and 28.

The Central Bank's problem can be written as a saddle-point functional equation using the same rationale as in the previous section.

$$\begin{aligned} V^{\mathrm{c}}(z_t,\delta_t,\mu_t) = \min_{\{\lambda_t\}}\max_{\{x_t^c\}} \quad & g(x_t^{\mathrm{c}},z_t,\delta_t,\lambda_t,\mu_t) \\ & + \delta_t\beta\mathbb{E}_t\left[\xi V^{\mathrm{c}}(z_{t+1},\delta_{t+1},\mu_{t+1}) + (1-\xi)V^{\mathrm{r}}(z_{t+1},\delta_{t+1})\right] \end{aligned} \tag{34}$$

This problem is subject to the set of equations (14)-(20) and (28), $\mu = \left\{\mu^E,\mu^P\right\}$, $\lambda = \left\{\lambda^E,\lambda^P\right\}$,

[9] See the Corollary to Theorem 2 in Marcet and Marimon (2019), which also shows that a solution satisfying the condition always exists when the saddle-point functional equation has a solution.

[10] A history of commitment is defined as $s_i = \mathrm{c},\ \forall\ 0 \le i < t$

$\mu_{t+1} = \lambda_t$ with $\mu_0 = 0$, and the function $g(\cdot)$ defined as

$$\begin{aligned} & g(x_t^{\mathrm{c}}, z_t, \delta_t, \lambda_t, \mu_t) \equiv u\left(x_t^c\right) \\ & + \lambda_t^{\mathrm{E}}\left[\mathrm{EuR}\left(x_t^c\right) - (1-\xi)\delta_t\beta\mathbb{E}_t\mathrm{EuL}\left(x_{t+1}^{\mathrm{r}}\right)\right] \\ & + \lambda_t^{\mathrm{P}}\left[\mathrm{PhL}\left(x_t^{\mathrm{c}}\right) - (1-\xi)\delta_t\beta\mathbb{E}_t\mathrm{PhR}\left(x_{t+1}^{\mathrm{r}}\right)\right] \\ & - \mathbb{I}_\xi \times \left[\mu_t^{\mathrm{E}}\mathrm{EuL}\left(x_t^{\mathrm{c}}\right) + \mu_t^{\mathrm{P}}\mathrm{PhR}\left(x_t^{\mathrm{c}}\right)\right] \end{aligned} \tag{35}$$

Where $\mathbb{I}_\xi = 0$ if $\xi = 0$ and $\mathbb{I}_\xi = 1$ otherwise, and the functions inside the brackets are the current and forward-looking terms of the Euler equation and the Phillips curve:

$$\mathrm{EuR(x)} = c^{-\sigma}R^{-1} \tag{36}$$

$$\mathrm{EuL(x)} = c^{-\sigma}\pi^{-1} \tag{37}$$

$$\mathrm{PhL(x)} = c^{-\sigma}y\left[1 - \nu - \frac{w}{z} + \phi\nu\pi\left(\pi - \pi_f\right)\right] \tag{38}$$

$$\mathrm{PhR(x)} = \phi\nu c^{-\sigma}y(\pi - \pi_f)\pi \tag{39}$$

The terms multiplied by $(1-\xi)$ are the part of next period's expectations that falls on the reneged branch. The part that falls on the committed branch is enforced next period through the inherited multipliers $\mu_{t+1} = \lambda_t$, which enter g with a negative sign as in (32).

The imperfect commitment equilibrium consists in the set of time-invariant functions $V^C(z, \delta, \mu)$, $V^R(z, \delta)$, $x^{\mathrm{c}\star}(z, \delta, \mu)$, $x^{r\star}(z, \delta)$ that satisfies:

1. Given $\{x^{\mathrm{r}}\}_t$ and $V^{\mathrm{r}}(z_t, \delta_t)$, then $\{x_t^c\}$ solves the problem in (34)
2. $V^{\mathrm{r}}(z_t, \delta_t) = V^{\mathrm{c}}(z_t, \delta_t, \mathbf{0})$
3. Denote the optimal plan under commitment by $x^{\mathrm{c}\star}(z_t, \delta_t, \mu)$, then $x^{\mathrm{r}\star}(z_t, \delta_t) = x^{\mathrm{c}\star}(z_t, \delta_t, \mathbf{0})$

Every time the central bank reneges in this equilibrium, the promises inherited from the previous plan are erased: as conditions 2 and 3 above state, the inherited multipliers μ are reset to zero, and the central bank solves the same problem as a planner at time zero. The current multipliers λ are chosen optimally, as in any other period, and become next period's promises if the new plan survives.

One of the consequences of the statement above is that this equilibrium does not have a non-stochastic steady state. The endogenous state will constantly jump to zero every time reneging is allowed. Note that when the probability of committing approaches one ($\xi \to 1$), the imperfect commitment equilibrium converges to the one in the commitment section. Similarly, when the probability of committing approaches zero, the loose commitment equilibrium converges to the discretionary equilibrium in the discretionary world.

2.3 Welfare loss measure

I will use the commonly used "units of consumption" measure to compare welfare. i.e., how much consumption is required to compensate the household for not having the flexible price equilibrium (in terms of current consumption).

Formally, denote x_{flex} as the variable x in the flexible price equilibrium without monopolistic competition.

Denote x_{zlb} as the variable x in the model where the ZLB occasionally binds. Then the welfare loss as a percentage of current consumption (Λ) is computed as

$$u\left(c_{\text{zlb}}(1+\Lambda), n_{\text{zlb}}\right) = u\left(c_{\text{flex}}, n_{\text{flex}}\right) \tag{40}$$

which, in our utility function specification yields

$$\Lambda = \left[\left(\frac{c_{\text{flex}}}{c_{\text{zlb}}}\right)^{1-\sigma} - \chi\frac{1-\sigma}{c_{\text{zlb}}^{1-\sigma}}\left(\frac{n_{\text{flex}}^{1+\eta} - n_{\text{zlb}}^{1+\eta}}{1+\eta}\right)\right]^{1/(1-\sigma)} - 1 \tag{41}$$

Λ is a period-by-period measure: it is computed in every state and compares current utility with that of the efficient allocation. The figures and tables in subsection 5.2 report statistics of Λ across the simulated periods (its mean and standard deviation, or its mean conditional on the ZLB binding or not). The mean of Λ is in general different from the compensation that equates expected utilities, and it measures the average period loss in the stationary distribution, not the lifetime welfare of a household starting from a given state.

3 Solution Methods

The solution method used in this paper consists of a global method that discretizes the state and finds the policy functions using time iteration as in Coleman (1991). For the discretionary monetary policy, the state is exogenous and normally distributed; therefore, finding a discretization of the state is trivial (see Miranda and Fackler (2004), for example). Since the kink in the policies is located at the ZLB, a bisection algorithm is included to find this point. This ensures having a denser grid in the areas with more probability and curvature simultaneously. A detailed discussion can be found in the Appendix C.

The main difficulty for the imperfect commitment case is finding the endogenous state's ergodic distribution, considering the "curse of dimensionality". The number of grid points in each endogenous state must be low for the algorithm to be fast enough. This sacrifices accuracy. To minimize this trade-off, I implement an algorithm as in Maliar and Maliar (2015), which iterates over the state before finding the final solution. This algorithm ensures a higher density of non-symmetric nodes where the economy is more likely to be found. A more detailed discussion can be found in the Appendix C.

The resulting solutions are accurate. Residuals of the equilibrium and optimality conditions, evaluated away from the solution grid at simulated states, are small and of similar size at every level of commitment, and at least two orders of magnitude below the welfare losses reported below. Appendix C documents the accuracy tests in detail.

4 Calibration

In this section, I discuss the parameter space in which the model will be solved. Most of the benchmark figures are taken from the literature, which focuses on the United States market. I will follow this benchmark and explain what happens in a broader space.

Table 1 shows all the model parameters and describes the benchmark calibration. The first block of parameters is related to the household. β measures the household's patience and is set to 0.995 to account for a quarterly frequency with a steady-state real interest rate of around 2%. The inverse of the elasticity of inter-temporal substitution, σ, will be set to 1, as is usual in the literature. Some examples with a higher number will be explored.[11] The Frisch elasticity of labor supply is calibrated as in Nakata (2016) and therefore η equal to 0.5. Finally, the weight of labor in the household utility, χ, is set to one.

The second block corresponds to firm-related parameters. The free-of-cost level of inflation is set to one, as in most of the literature. Some experiments in the Appendix explore higher values for this parameter (e.g., 2% targets): since π_f shifts average inflation up at every level of commitment, raising it should push the average nominal rate further from zero and lower the probability of hitting the ZLB, though this is a mechanical implication of the calibration rather than a result solved for across the full commitment line. ν and ϕ are set to match an inflation bias of 2% in the steady state of the model without commitment.

The shocks of the model are calibrated to generate a probability of hitting the zero lower bound of 10% in the economy without commitment.

5 Results and Discussion

In this section, I discuss the results of solving the model using the previous parameter space. I divide this section in two since it is relevant to understand how the economy reacts in the discretionary equilibrium to assess the welfare gains when the Central Bank has higher degrees of commitment. In particular, I will show that when the monetary policy is discretionary, higher persistence in the TFP shocks will lead to an economy less prone to hit the zero lower bound. The opposite happens when the stochastic driver is the preferences shock: higher persistence will lead to a higher probability of hitting the ZLB. These results yield a different dispersion and levels of welfare loss across the exogenous state. The TFP shock implies a boom in activity and downward pressure

[11] See Appendix A.

Table 1: Benchmark Calibration

Parameter	Description
Household	
$\beta = 0.995$	Intertemporal discount factor.
$\sigma = 1.000$	Inverse of elasticity of intertemporal substitution.
$\eta = 0.500$	Inverse of Frisch elasticity of labor supply.
$\chi = 1.000$	Labor dis-utility weight.
Firms	
$\nu = 0.087$	Inputs' inverse elasticity of substitution.
$\pi_f = 1.000$	Free-of-cost inflation rate.
$\phi = 1{,}000$	Adjustment cost weight.
TFP Stochastic Process	
$\rho_z = 0.750$	Persistence factor
$\sigma_z = 0.004$	Variance of innovation.
Discount Factor Stochastic Process	
$\rho_\delta = 0.350$	Persistence factor
$\sigma_\delta = 0.007$	Variance of innovation.

in prices. This deflationary process pulls the nominal interest down, making the economy fall into the ZLB. Since this is counterfactual (ZLBs are associated with recessions), I will focus on the loose commitment results with the preferences shock. Concerning this, I will show that: (i) when the economy displays inflation bias due to monopolistic competition, the presence of the ZLB improves welfare. Nonetheless, when converging to commitment, whether the ZLB occasionally binds or not is irrelevant to the average welfare loss (the gap between the models with and without the ZLB closes by the time the Central Bank is fully committed). (ii) The gains in welfare are diminishing from discretion to commitment. (iii) The probability or duration of the ZLB cannot identify by itself the level of commitment. (iv) The loosely committed economy has lower levels but higher inflation volatility than discretion. Nonetheless, the volatility does not increase homogeneously with commitment: highly committed economies reduce the volatility by using more "bullets" during the ZLB, which becomes less likely than loosely committed ones.

5.1 Discretion

Figure 1 shows the policy function for inflation when the stochastic driver is the total factor productivity shock. For the same ergodic distribution, I show what happens when the persistence of the shock changes and when the zero lower bound restriction does not exist. Since monopolistic competition generates an inefficient loss that can be erased with a combination of labor subsidy

Figure 1: Optimal inflation under discretion and TFP shock

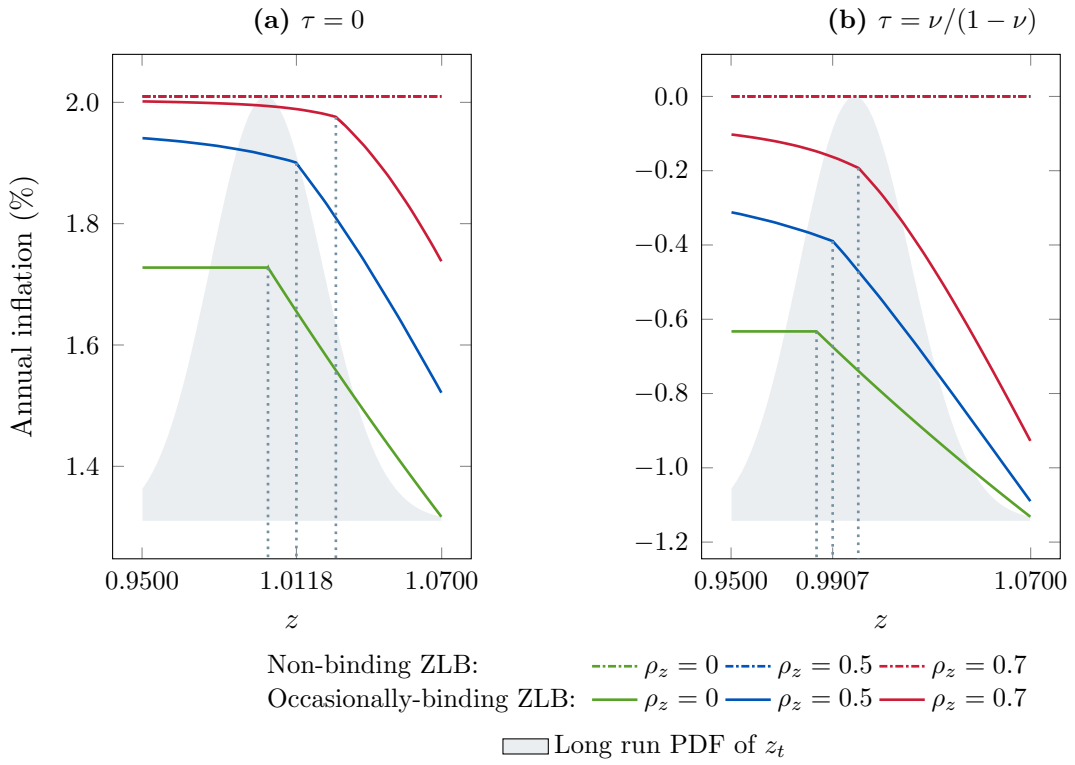

Note: The panels display the optimal discretionary inflation rates for different TFP shock persistences (ρ_z). The left panel presents results with monopolistic competition inefficiency, while the right panel shows the results with efficient labor subsidy. The dotted vertical lines indicate where the zero lower bound begins to bind for each ρ_z. The dot-dashed lines represent the inflation policies for the model without the ZLB, which overlap in both panels. The shaded area represents the ergodic probability density function of z for the selected abscissa. Appendix A contains the complete graph and other policy functions.

and lump sum tax, I show both cases[12]. Panel (1a) shows these results for the inefficient economy, while panel (1b) does it for the efficient one. Five conclusions can be drawn from Figure 1. First, the presence of monopolistic competition yields an average inflation bias. This is well known in the literature: once low-level inflationary expectations are set (e.g., π_f), the central bank has the incentives to increase inflation above that level to deal with the losses generated by the inefficiency. This level is bounded from above by the convex adjustment costs. Second, in both panels, inflation does not react to TFP shocks when the lower bound in nominal interest rates does not bind. An expansionary productivity shock increases the production of the consumption good and brings deflationary pressures. Since deflation harms the economy through the adjustment cost, it is in the interest of the Central Bank to counteract it. This would boost the economy even further since it would require a lower nominal interest rate and higher current consumption, thus increasing the incentives to offset the deflationary pressure. This is counterbalanced by the effects of this policy on future inflation and consumption and the cost of deviating from π_f. In the case of the panel (1b), the equilibrium is the zero inflation (since $\pi_f = 0$) at any state since it is the one that yields the efficient allocation and is time consistent. For panel (1a), however, the sign of this relationship is not general: it holds at the benchmark calibration because the inter-temporal elasticity of substitution is set to $\sigma = 1$. Repeating the same exercise with $\sigma = 2$ and everything else unchanged (Figure 12 and Figure 15 in the Appendix) flips the slope of the inflation policy function to positive. The discretionary welfare ranking is unaffected, however: the ZLB still lowers the welfare loss whenever the monopolistic-competition inefficiency is present, and still raises it in the efficient benchmark, exactly as under $\sigma = 1$; only the direction in which inflation itself responds to a TFP shock away from the ZLB is σ-dependent. Third, the zero lower bound brings a deflation bias in both economies. This negative bias in inflation partly offsets the positive one generated by discretion and monopolistic competition. Since the option to take advantage of the boom by lowering the nominal interest rate is not present everywhere, the expectations and the current level of inflation are lower. Fourth, whether the economy is efficient or not, more persistent shocks make the economy less prone to reach the zero lower bound. This is a consequence of three main drivers. First, more persistent shocks have less effect on the real interest rate since they introduce fewer alterations in the balance between consumption and savings. Second, to get the same ergodic distribution, the variance of the innovation in the *AR(1)* process has to be higher for lower persistence. This means that the probability of reaching an arbitrary area on the right side of the ergodic distribution, given that the economy is on the left side, is more than proportionately higher for the less persistent economy. Third, as I mentioned before, the fall into the ZLB is driven by a boom that cannot be reinforced by a loosening of the nominal interest rate; this generates lower inflation (current and expected), strengthening the fall into the ZLB. Finally, an *i.i.d.* shock implies a constant deflation bias when the economy is not in the ZLB. Similarly, when the persistence is higher than zero, the closer to the left tail of the ergodic distribution of the TFP shock, the smaller the deflationary pressure.

[12]I make this comparison since efficient taxation is ubiquitous in the literature where the model is linearized.

To understand this, note that inflation drops once the economy falls into the ZLB. Since current inflation depends positively on future inflation, the chances of falling into the ZLB are accounted for everywhere in the ergodic distribution of z. Thus, when the shock is *i.i.d.*, the probability of falling into the ZLB is independent of the current state and expected inflation is the same to the left of the ZLB threshold. This leads to a constant deflationary gap for the economy without ZLB. With a persistent shock, the gap is not constant because the expectations of falling into the ZLB now depend on the state. The closer to the ZLB, the more likely it is to happen, lower the expected inflation and, therefore, lower current inflation. This leads to an increasing gap from left to right.

Figure 2: Optimal inflation under discretion and preferences shock

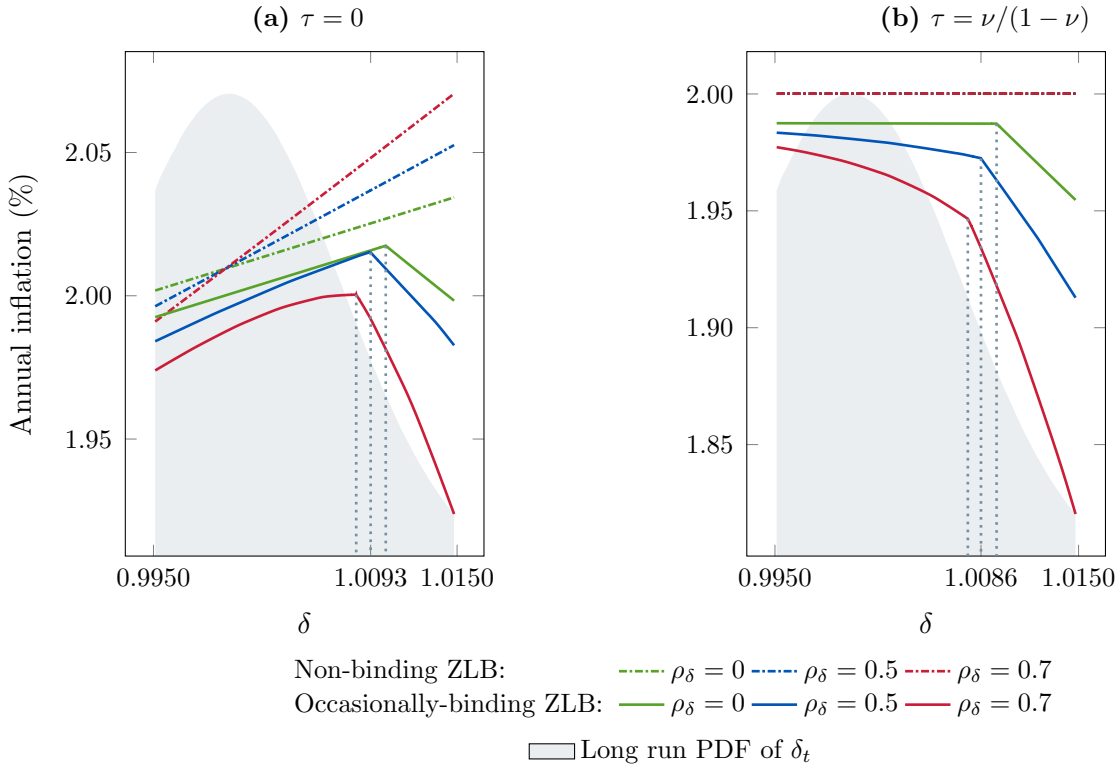

Note: The panels display the optimal discretionary inflation rates for different demand shock persistences (ρ_δ). The left panel presents results with monopolistic competition inefficiency, while the right panel shows the results with efficient labor subsidy and $\pi_f = 2\%$. The dotted vertical lines indicate where the zero lower bound begins to bind for each ρ_δ. The dot-dashed lines represent the inflation policies for the model without the ZLB, which overlap in panel (b). The shaded area represents the ergodic probability density function of δ for the selected abscissa. Appendix A contains the complete graph and other policy functions.

Figure 2 shows a similar exercise to the previous one but for the driver of uncertainty being the discount factor shock. This time I plot the efficient economy in panel (2a) with a level of *free-of-cost* inflation of 2% to avoid the ZLB being an absorbent state. In the TFP shock, the boom generated a price fall, decreasing the nominal interest rate. On the other hand, the increase in patience lowers the real interest rate, pushing its nominal counterpart down. When the zero lower bound

does not bind, the efficient economy – dot-dashed line in panel (2b)– has a constant inflation equal to the one *free-of-cost*. Similarly to the TFP shock, the central bank has no incentives to change its mind after the shocks are materialized and pursuing $\pi = \pi_f$ is optimal. Nonetheless, when the ZLB binds, or the inefficiency of the monopolistic competition is in place, optimal inflation reacts to preference shocks, and the equilibrium displays an average inflation (or deflation) bias. As before, the monetary policy explains the latter by the time inconsistency. The positive slope of (2a) dot-dashed lines' is the reaction that balances the reaction of current consumption and labor, and the cost of inflation given that average inflation is high. The shock increase implies a lower consumption and more hours worked, which should drive prices down given a fixed nominal interest rate. Nonetheless, part of the shock can be reverted with a monetary policy relaxation. This loosening increases inflation since it has to be aggressive enough to affect the decisions of the household, while also reducing the proportion of output available for consumption.[13] The difference of the economies with and without the ZLB in Figure 2 are explained similarly as in figure 1: since the loosening of the monetary policy is constrained in part of the distribution, the effects are transmitted through expectations to the rest of the policy function. Another notable feature is that inflation is more volatile in economies with more persistent shocks. For the economies without ZLB, average inflation is the same regardless of the persistence of the shock. Nonetheless, when the ZLB occasionally binds, the more persistent shocks imply lower average inflation. This is due to the deflation bias of the zero lower bound, and the more persistent the shock, the more likely the economy is to fall into it. An efficient economy has a higher probability of hitting the zero lower bound. This can be seen by comparing where the ZLB starts binding in the shock's PDF. All three vertical dotted lines in panel (2b) are closer to the centre than the comparable ones in panel (2a). The explanation for this is that the deflation bias (measured as the gap between inflation with and without the ZLB) is higher in the efficient economy, which is explained by the slope of the inflation policy function without the ZLB.

A result of the points above is that the relation between the length of the ZLB episodes and the unconditional probability of happening is not trivial. For the preference shocks, the more persistent shock is associated with more frequent ZLB episodes. Thus each episode is expected to last longer. For the TFP shock, however, this is not true. The more persistent shocks drive a lower probability of hitting the ZLB. Therefore, it needs to be clarified whether its duration will increase. Figure 3 shows this for a broader and more populated space of shock persistence. In both panels, the lines indicate the probability of hitting the zero lower bound (right axis), while the areas reflect the frequency of the durations of each episode. For example, panel (3b) shows that the likelihood of finding episodes longer than three periods for persistence lower than 0.2 ($\rho_\delta < 0.2$) is close to zero. Nonetheless, the more persistent the shock, the more the ZLB episodes last – e.g., when $\rho_\delta = 0.85$, the likelihood of finding ZLB episodes longer than six quarters is approximately 7%. This is not the case for the TFP shock, depicted in panel (3a): the more persistent shocks yield

[13] See A for a more formal explanation.

Figure 3: Probability and duration of the zero lower bound under discretion

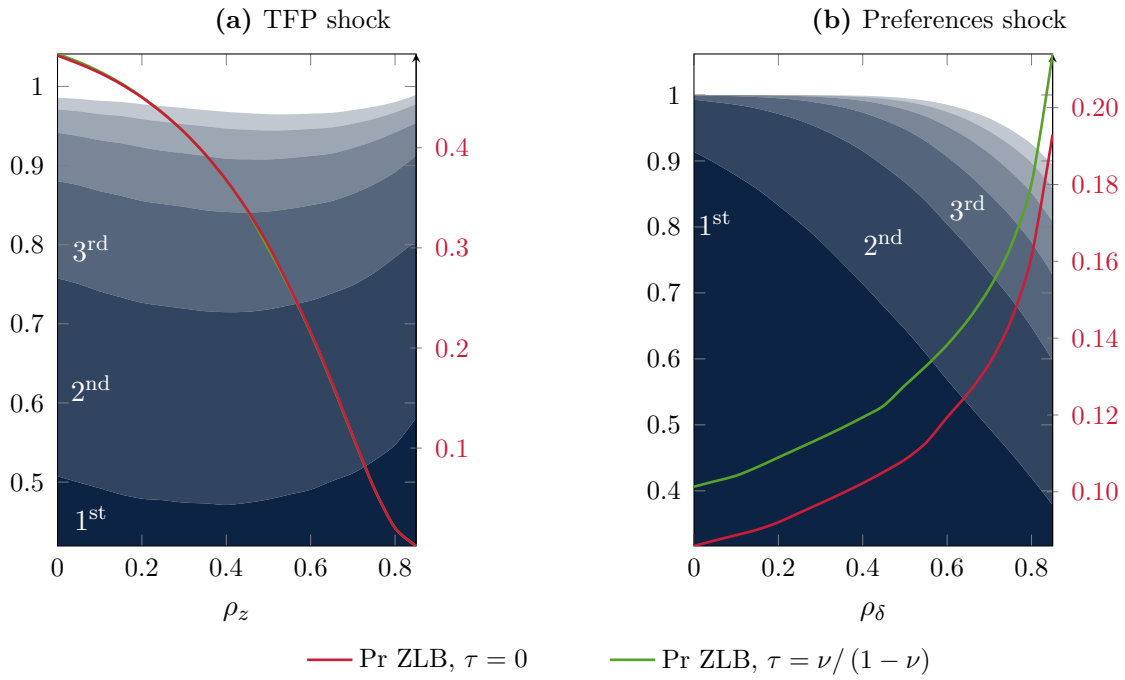

Note: The panels show how ZLB probability (right axis, lines) and duration (left axis, shaded areas) vary with shock persistence. Shaded areas indicate likelihood of ZLB episodes ending after 1st to 6th spell. Panels 3a and 3b depict TFP and preferences shocks, respectively. Both panels feature probability of ZLB with and without monopolistic inefficiency but same average inflation. These two lines overlap in panel 3a.

less likely ZLB episodes but similar durations. Figure 3 also shows the probability of hitting the ZLB for an economy that deals efficiently with the monopolistic competition but with free-of-cost inflation levels comparable with the inflation bias[14]. In the case of the TFP shock, the probabilities are identical, meaning that whether the positive inflation comes from inflation bias or efficient mechanism will not influence the chances of hitting the ZLB[15]. For the preferences shock, on the other hand, the deflation bias is higher in the efficient economy, leading to a homogeneous increase in the chances of hitting the ZLB. Again, this phenomenon can be tracked down to the difference in the slope of the inflation policy function for the efficient vs the inefficient economy. Although the reaction of inflation in the previous figures is optimal, its consequence in welfare loss is not trivial. The branch of the literature that linearizes the model taking the non-inflationary as the first-order approximation point, has to assume that the economy has a way to enforce the exact taxing scheme that erases the monopolistic competition. In this case, the ZLB always reduces the economy's welfare by generating a deflation bias. In any other case, an occasionally binding ZLB will be welcomed if the deflation bias is small enough to counteract the inflation bias caused by monopolistic competition without overpowering it. Figures 4 and 5 show the welfare losses of the model under discretion compared to the equilibrium with efficient allocations described in the model section. The dashed-dotted lines show the welfare loss for the economy without the ZLB, while the solid lines are equivalent for the economies with an occasionally binding one. The colors differentiate the results for different levels of persistence of the shocks. Figure 4 shows the welfare loss when the driver of uncertainty is the total factor productivity shock. Panel (4b) shows that when monopolistic competition inefficiency is taken care of, the ZLB reduces welfare everywhere. This is more pronounced in the economy with higher deflation bias and higher chances of hitting the ZLB ($\rho_z = 0$). For the economy with the monopolistic competition inefficiency (panel 4a), the welfare loss without the ZLB is constant as the reaction of inflation. Including the ZLB (i.e., comparing the solid and dot-dashed lines) yields an average more minor welfare loss but not everywhere. Falling too deep into the ZLB is associated with increased welfare loss. This is more pronounced in the *i.i.d.* shock (green line) at the right tale of the TFP shock distribution. Despite inflation getting closer to the free-of-cost one, it is not enough to overcome the amount of consumption not gained due to the inability to loosen the monetary policy. The Phillips curve stays put since inflation expectations do not change with the state. When the shock is persistent, inflation is higher to the left tail of the distribution. This implies that ceteris paribus, the Phillips curve shifts to lower inflation when the economy starts getting close to the ZLB and falls faster to the distribution's right tail. This generates a local minimum welfare loss.

[14]The figures for the duration are almost identical, therefore I omit the comparison in this dimension

[15]Note that this result depends on the slope of the policy function, which in turn depends on the elasticity of inter-temporal substitution.

Figure 4: Loss function under discretion and TFP shock

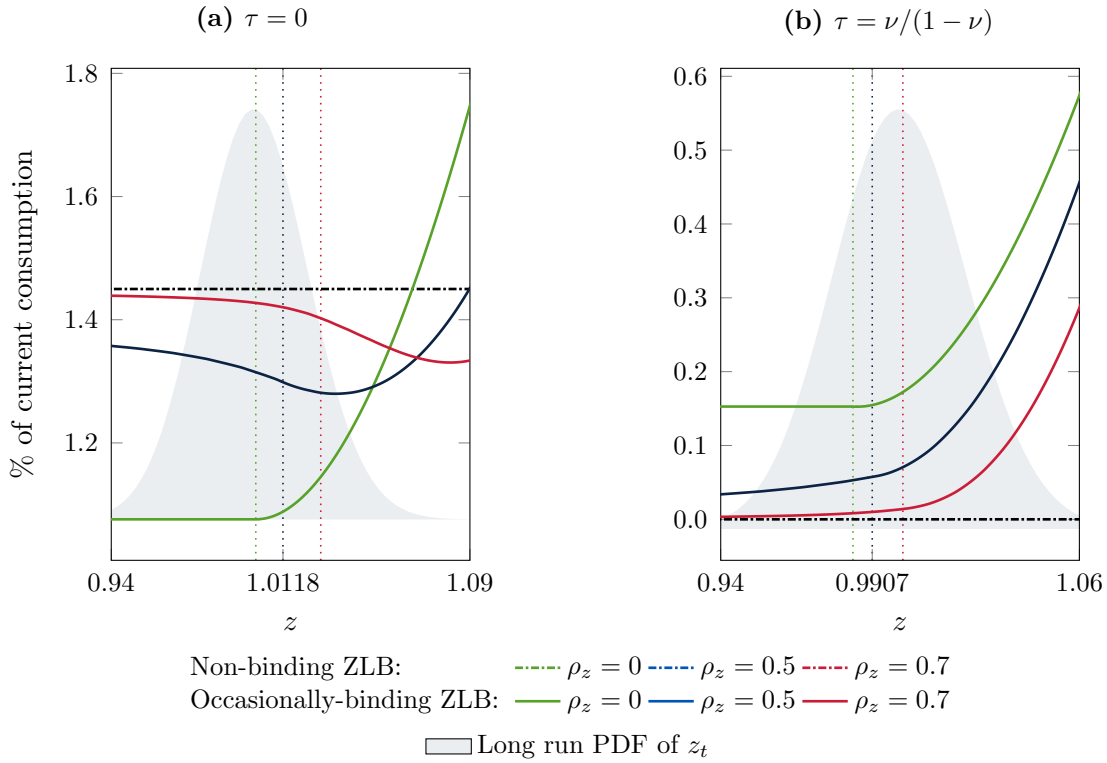

Note: The panels show the discretionary welfare loss for selected levels of persistance of the TFP shock ρ_z. The left panel shows the results with the monopolistic competition inefficiency, while the right one does it with factor inputs as in perfect competition and $\pi_f = 0\%$. The dotted vertical lines shows where the zero lower bound starts binding for each ρ_z. The dot-dashed lines show the inflation policy for the model without the ZLB, and the shaded area is the ergodic PDF of δ for the selected abscissa. The complete graph and other policy functions are available in the appendix.

The analysis for the demand shock is similar. Figure 5 shows the welfare loss for the economy where the discount factor shock is the driver of uncertainty. Again, when the effects of the monopolistic competition are taken care of, and the zero lower bound does not bind (panel 5b, dot-dashed line), the monetary policy is time consistent, and the efficient allocations can be reached, generating a zero welfare loss. When the zero lower bound is added back, the larger the deflation bias is (red line), the more significant the welfare loss. Note that here, since the free-of-cost inflation is set to a higher level, and therefore the ZLB starts binding further to the right —compared to the TFP shock— the effects of the ZLB away from it are minuscule. When the economy cannot solve the problem of monopolistic competition exogenously (panel 5a), the losses when the zero lower bound does not bind reflect the inflation bias from the optimal policy. Adding the ZLB will lower the average welfare loss since it introduces the inflation bias that brings the average inflation closer to its free-of-cost figures. Again, this does not happen everywhere, although the crossing between lines and dot-dashed lines happens in a non-significant area of the ergodic distribution of the shock.

Figure 5: Loss function under discretion, Preferences shock

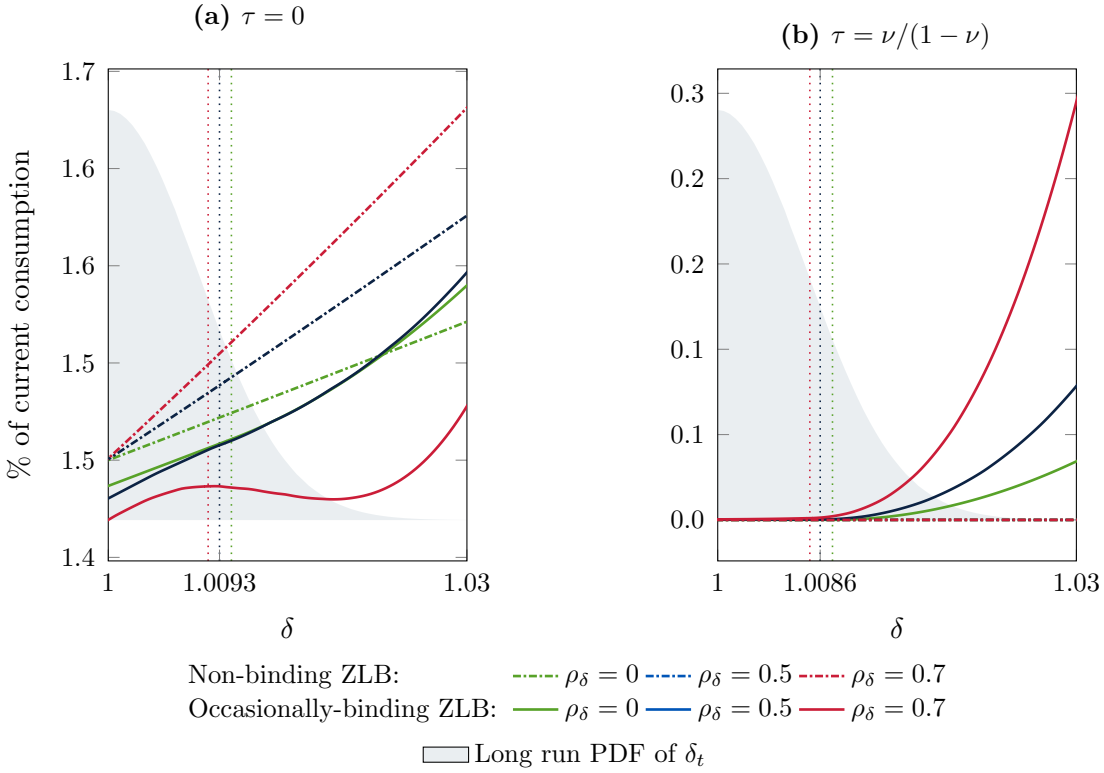

Note: The panels show the discretionary welfare loss for selected levels of persistance of the preference shock ρ_δ. The left panel shows the results with the monopolistic competition inefficiency, while the right one does it with factor inputs as in perfect competition but $\pi_f = 2\%$. The dotted vertical lines shows where the zero lower bound starts binding for each ρ_δ. The dot-dashed lines show the inflation policy for the model without the ZLB, and the shaded area is the ergodic PDF of δ for the selected abscissa. The complete graph and other policy functions are available in the appendix.

These results are consistent with higher monopolistic power or elasticity of inter-temporal substitution. Although inflation in the presence of monopolistic competition gets a positive slope, the analysis of welfare, average and dispersion of inflation is similar to the one described here.[16]

5.2 Imperfect commitment

Knowing the interactions of the ZLB and the monopolistic competition in the discretionary equilibrium helps to analyze what happens in the equilibrium where the Central Bank is allowed to reset its plan with a fixed probability. From the figures above and the selected calibration, the role of the ZLB is small compared to that of monopolistic competition. Nonetheless, when the latter exists, the ZLB appears as a mechanism to lower inflation expectations, reducing the welfare loss induced by the inflation bias. This happens along the commitment line. Figure 6 shows the average welfare loss (the mean of Λ across the simulated periods) with the parameters calibrated in

[16] See Appendix A

Figure 6: Unconditional welfare loss

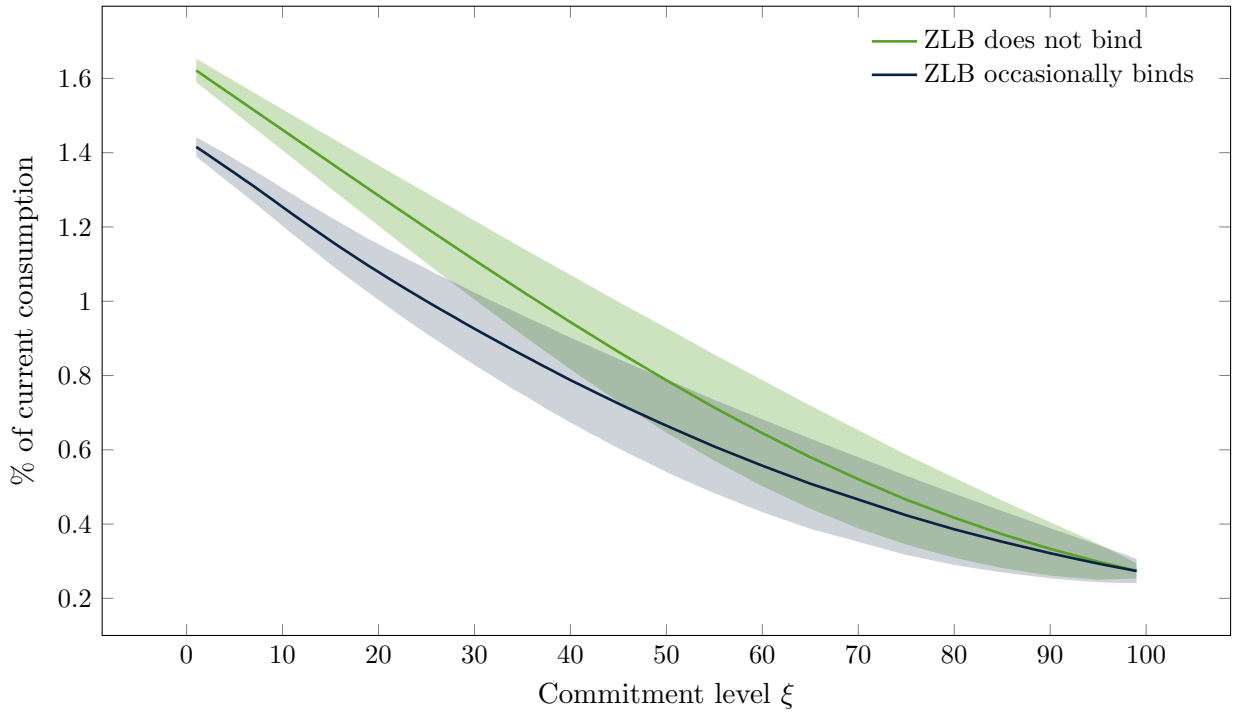

Note: The figure shows the average welfare loss in terms of current consumption for every level of commitment. The green line is the welfare loss in the model when the ZLB never binds, and the blue line shows the same when the ZLB occasionally binds. The shaded areas show the mean plus and minus one standard deviation of the simulated loss.

Table 1. The green line corresponds to the economy where the ZLB does not bind, while the blue represents the welfare loss of the occasionally binding constraint. The areas show the mean plus and minus one standard deviation of the simulated loss. In the abscissa is the level of commitment ξ. Since the discretionary equilibrium has an inflation level higher than its efficient counterpart, the ZLB generates a lower welfare loss. This is carried out along the commitment line until the 100% commitment is reached. At this point, the bulk of loss is caused by monopolistic competition, and the ZLB plays a minor role. This is because all the instances where the Central Bank can improve welfare by taking advantage of private expectations are taken care of. Since most of the welfare loss induced by the ZLB comes through this mechanism, its consequences are canceled. Another feature of Figure 6 is the negative slope of the loss without the ZLB appears to be constant closer to discretion, while this is not the case for the economy with occasionally binding ZLB: the latter reduces losses faster. The explanation is composed of the reduction inflation bias, which makes both slopes negative, and increasing numbers of ZLB events. Higher probability of hitting the ZLB makes the welfare loss smaller by reducing the inflation bias faster. This only happens in economies more committed but close to discretion. Closer to commitment, the likelihood of hitting the ZLB diminishes since the efficacy of forward guidance increases, making welfare gains to slow down. Figure 7 shows the probability of hitting the ZLB and its duration. Starting from discretion

Figure 7: Probability and Duration of the ZLB

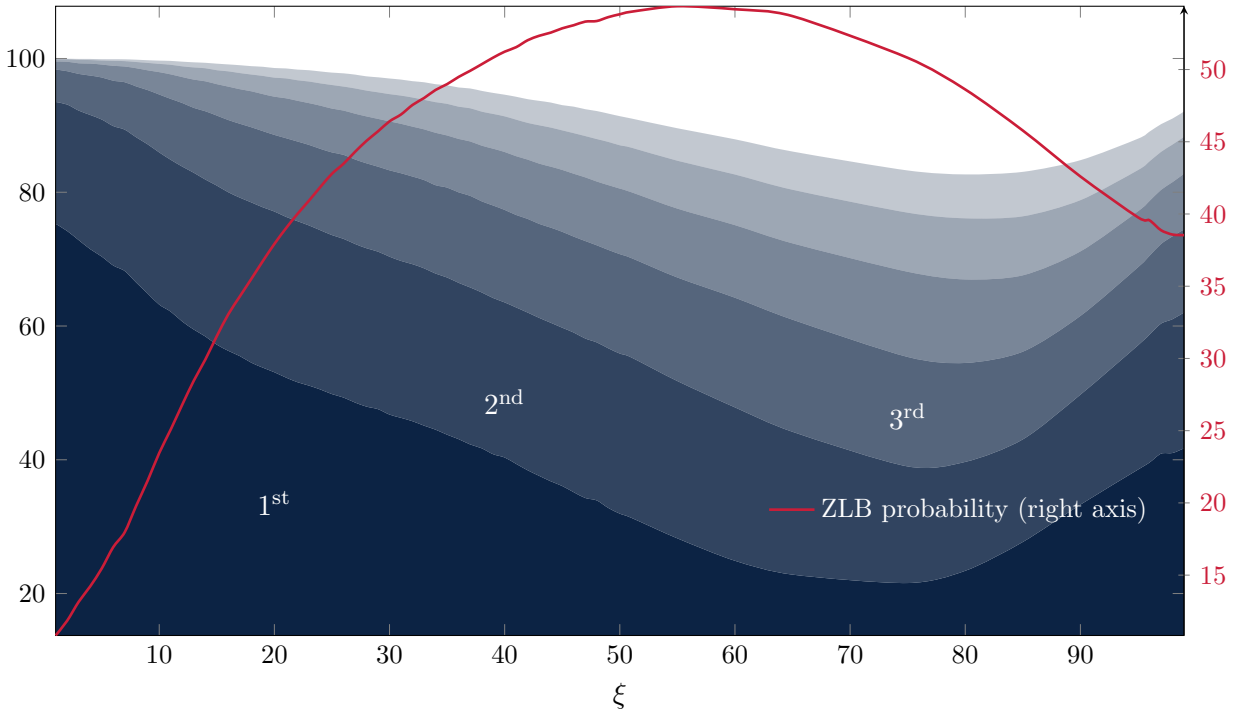

Note: The figure shows how ZLB probability and duration vary with commitment level. Shaded areas (left axis) indicate likelihood of ZLB episodes ending after 1st to 6th spell. The line (right axis) indicates the unconditional probability of reching the ZLB.

towards commitment, the initial chance of hitting the ZLB is 10% and reaches a maximum of 54% when the promises are kept 55% of the time; it stays above 50% for commitment levels roughly between 40% and 75%. After that, the probability of hitting the ZLB falls to 39% for an almost fully committed economy ($\xi = 99\%$). This implies that the likelihood of hitting the ZLB is not a good indicator to identify the monetary authority's commitment level. To understand why the ZLB likelihood across the commitment line is not homogeneous note that, on one hand, higher levels of commitment imply lower inflation bias. This makes nominal interest rates lower conditional on the same real interest rate. Thus, the probability of hitting the ZLB should increase with higher levels of commitment. On the other hand, a more committed plan can minimize the chances of the ZLB by increasing inflation expectations and delivering it when the time comes (forward guidance). According to this mechanism, the ZLB should be less likely. What Figure 7 depicts, is the interplay of both mechanisms. The duration of the ZLB events is not uniquely pinned down by the level of commitment either. Figure 7 shows that the discretionary ZLB episodes are short, with spells lasting less than four periods almost all the time. The duration of the spells starts rising with the increase of the probability of reaching the ZLB but reaches a maximum around the 70–75% commitment level. The chances of experiencing short spells of ZLB increase again towards total commitment. Nonetheless, at this point, it is possible to find long spells.

To analyze the characteristics of the results under loose commitment, Table 2 shows the results for selected levels of commitment and separates them between all the simulated samples (*all*), the observations in the ZLB (*zlb*) and the ones out of it (*non zlb*). It displays the average and standard deviation of the annualized inflation (π), the percentage distance of consumption (c) and labor (n, equal to output in this economy) to their efficient counterparts, and the welfare loss (*loss*, the period loss Λ defined in subsection 2.3) as displayed in Figure 6. The losses in and out of the ZLB are similar when the monetary policy is more discretionary. Nonetheless, when the economy gains levels of commitment, most of the loss is generated in periods where the ZLB does not bind [17]. Moreover, greater levels of volatility in the welfare loss are obtained out of the ZLB for lower levels of commitment. This reinforces the idea that the ZLB works as a mechanism that lowers the inflation generated by the inefficiencies of the monopolistic competition, therefore reducing its losses. This mechanism dies out with higher levels of commitment. In Table 2, for example, the monetary authority committed 90% of the time generates an overall loss of 0.32% distributed as 0.29% in ZLB periods and 0.35% in periods out of it. The volatility of the loss is in the same order of magnitude. For an economy with a Central Bank committed 50% of the time, on the other hand, the overall loss is 0.67% with an average of 0.63% in ZLB times and 0.71% in regular times. Moreover, the volatility of the loss out of the ZLB is one magnitude higher.

The intermediately committed economies fulfill lower welfare loss through lower but more volatile inflation. This is also true for highly committed economies, although the possibility of using its credibility (e.g., forward guidance) eliminates the necessity of adjusting inflation so often. Ideally, in regular times inflation is kept low, and during the ZLB, the Central Bank raises it to meet the ZLB constraint. The less committed economies do not need to do this since, in normal times, the expected inflation is already high, making the necessity of increasing it during the ZLB less active. This is observed in Table 2: overall inflation decreases as commitment rises. Moreover, volatility increases from $\xi = 10$ to $\xi = 75$, and falls from there to $\xi = 90$. Finally, the bulk of inflation volatility of $\xi = 90$ is generated in ZLB episodes, while the opposite is true for the rest.

Dynamics around the ZLB and reneging events

In this section, a simulated event study is provided to gain a more in-depth understanding of the model's implications[18]. The key situations generated by the model include the onset of the Zero Lower Bound (ZLB) and the resetting of the plan. A Zero Lower Bound (ZLB) event starts when the nominal interest rate changes from non-binding to binding in the current period ($t = 0$). This means that when $R = 1$ for several periods in a row, we consider it as one ZLB event. In Figure 8, we can see how a ZLB event affects inflation, the output gap, and welfare loss in economies with different levels of commitment. Economies with higher commitment tend to have more fluctuating inflation. As shown in Panel (a), after a ZLB event starts, inflation usually goes up and often stays

[17] See the complementary figure in the appendix

[18] A simulation of output and inflation is presented in Figure 23.

Table 2: Results for selected levels of commitment

		$\xi = 10$		$\xi = 30$		$\xi = 50$		$\xi = 65$		$\xi = 75$		$\xi = 90$	
		$\mathbb{E}$	σ	$\mathbb{E}$	σ	$\mathbb{E}$	σ	$\mathbb{E}$	σ	$\mathbb{E}$	σ	$\mathbb{E}$	σ
all	π	1.83	*0.077*	1.46	*0.153*	1.09	*0.236*	0.78	*0.294*	0.57	*0.315*	0.23	*0.273*
	c	6.08	*1.068*	6.00	*1.701*	5.94	*2.026*	5.91	*2.052*	5.89	*1.922*	5.88	*1.374*
	n	5.10	*1.152*	5.36	*1.826*	5.58	*2.159*	5.70	*2.170*	5.77	*2.020*	5.85	*1.431*
	loss	1.25	*0.051*	0.93	*0.098*	0.67	*0.125*	0.51	*0.121*	0.42	*0.107*	0.32	*0.068*
zlb	π	1.84	*0.017*	1.54	*0.037*	1.22	*0.091*	0.96	*0.145*	0.78	*0.192*	0.42	*0.281*
	c	6.11	*0.371*	5.38	*0.543*	4.97	*0.803*	4.76	*0.978*	4.69	*1.129*	5.05	*1.414*
	n	5.12	*0.393*	4.68	*0.576*	4.52	*0.862*	4.48	*1.048*	4.50	*1.202*	4.97	*1.478*
	loss	1.26	*0.014*	0.91	*0.020*	0.63	*0.028*	0.46	*0.037*	0.37	*0.047*	0.29	*0.067*
non zlb	π	1.82	*0.087*	1.40	*0.184*	0.93	*0.260*	0.57	*0.282*	0.35	*0.260*	0.10	*0.163*
	c	6.07	*1.203*	6.53	*2.127*	7.08	*2.395*	7.24	*2.165*	7.14	*1.776*	6.51	*0.944*
	n	5.10	*1.299*	5.95	*2.277*	6.81	*2.535*	7.12	*2.268*	7.08	*1.848*	6.50	*0.975*
	loss	1.25	*0.057*	0.94	*0.130*	0.71	*0.172*	0.57	*0.155*	0.48	*0.122*	0.35	*0.055*

Note: This table summarizes the loose commitment results for selected levels of commitment ξ (in percent). The columns $\mathbb{E}$ refers to the average of the simulation, while σ to the standard deviation. The variable π refers to the annualized inflation rate, c and n to the deviations in percentages from the frictionless steady state values of consumption and labor, respectively; loss refers to the welfare loss measure explained above. The sample *all* refers to all the observations in the simulation, while *zlb* and *non zlb* refers to the observations in the zero lower bound and out of it, respectively.

high for a while. This is common in plans with high commitment, which tend to keep the output gap closer to the efficient allocation and result in lower welfare losses after the event. The situation is different when there's less commitment.

Figure 8: ZLB simulated events

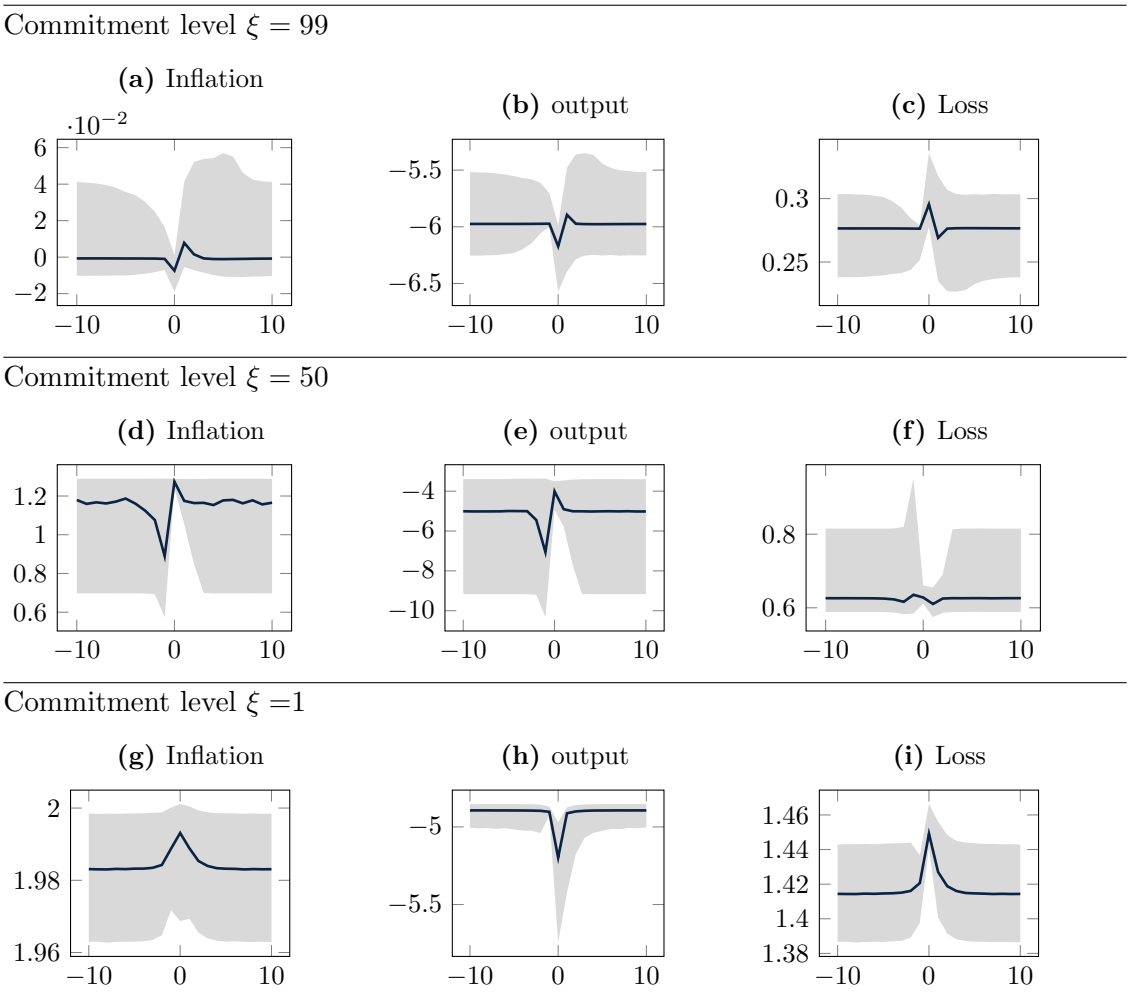

Note: The panels display inflation, output, and welfare losses around the ZLB (Zero Lower Bound) events. Panels (a), (b), and (c) show these variables for $\xi = 99$; (d), (e), and (f) for $\xi = 50$; and (g), (h), and (i) for $\xi = 1$. The solid lines represent the 50th percentile, while the shaded area covers the 10th to 90th percentile. Output is measured as the gap with respect to the efficient economy, while loss is the percentage amount of current consumption that leaves the household indifferent with respect to the efficient allocation.

In Panel (g), we see that at the start of ZLB episodes, inflation is usually below the average (2%). But when the episode begins, inflation goes up to meet the limit set by the nominal interest rate. This increase doesn't last, so it doesn't fully offset the crisis, leading to bigger losses as shown in Panel (i). Economies with some level of commitment show a mix of the patterns we see in fully committed and less committed economies. Now I pay attention to reneging events. Under commitment, there's no plan reset, while in discretion, plan resets are frequent. However, because the chance of a plan reset is independent, I simulate many scenarios starting from the reset state, and force the monetary policy to stick to the same plan for 10 periods. This rarely happens in less committed economies. For instance, when $\xi = 0.5$, the likelihood of it occurring

is less than 0.1%.[19] Still, this simulation helps highlight why economic factors fluctuate more in loosely committed settings. When a committed economy resets the plan, inflation quickly goes up, along with a significant increase in output that results in a gain in welfare. However, when a discretionary monetary authority has to stick to one plan, it tends to reduce inflation quickly and cause a recession that outweighs the benefits of lower inflation. In these rare cases, welfare losses skyrocket. These effects are shown in Figure 9. The situation for mid-level commitment schemes is mixed. A plan reset brings an initial welfare gain, but it doesn't last. Like in the discretionary case, sticking to the plan for too long worsens welfare rapidly.

Figure 9: Plan re-setting simulated events

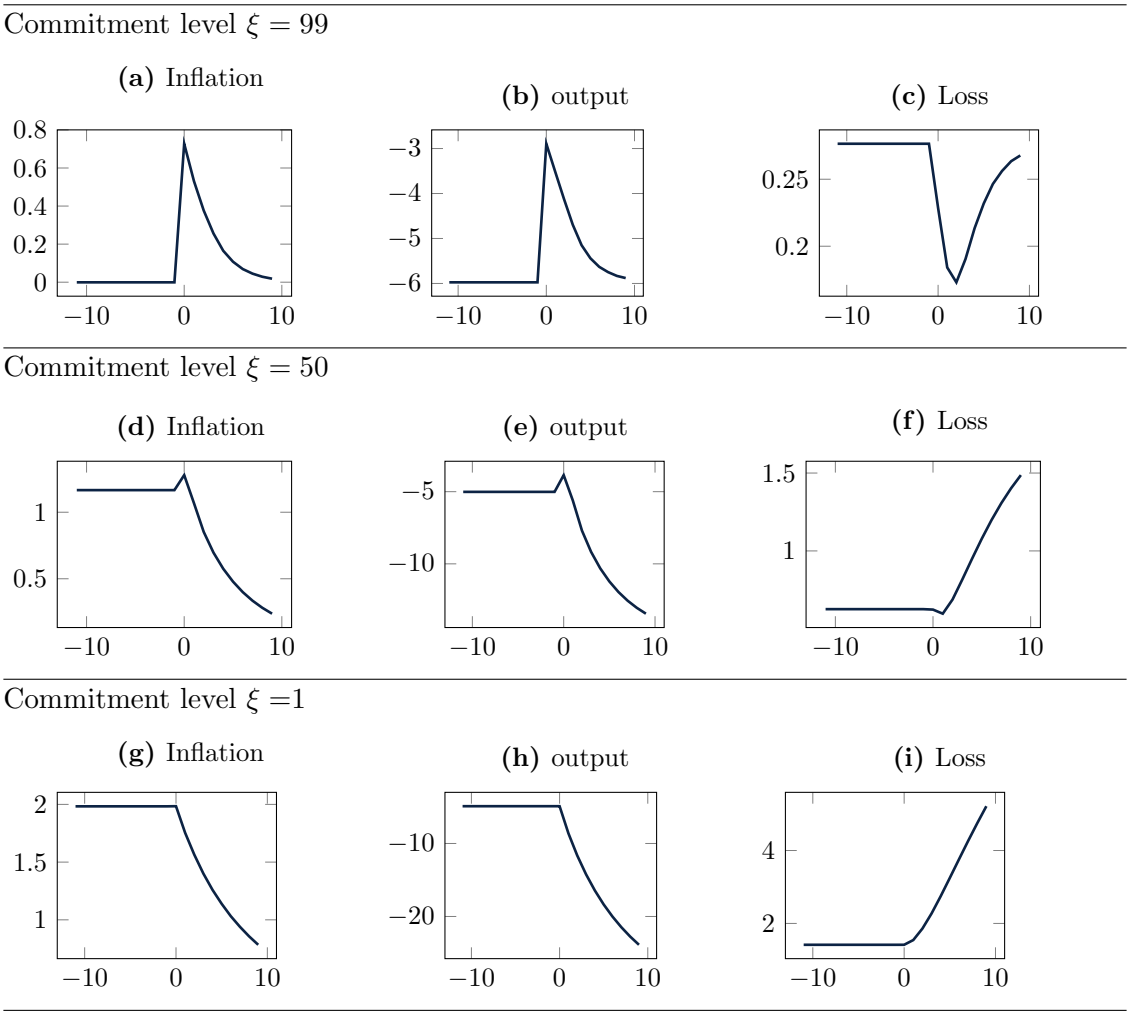

Note: The panels display inflation, output, and welfare losses around the plan resetting events. Panels (a), (b), and (c) show these variables for $\xi = 99$; (d), (e), and (f) for $\xi = 50$; and (g), (h), and (i) for $\xi = 1$. The solid lines represent the 50th percentile. Output is measured as the gap with respect to the efficient economy, while loss is the percentage amount of current consumption that leaves the household indifferent with respect to the efficient allocation. The events assume, for all levels of commitment, that the central bank cannot renege from time zero until the end of the event.

[19] See Figure 22 in Appendix A to get a sense of the chances of this happening for other commitment levels.

6 Conclusion

This paper solves for optimal monetary policy in a fully non-linear New Keynesian model with monopolistic competition, costly price adjustment, and an occasionally binding ZLB, when the central bank keeps its plan with probability ξ and reneges otherwise. Varying ξ from discretion to full commitment, with the rest of the calibration fixed, shows how two biases interact along the commitment line: the inflation bias created by monopolistic competition, which the central bank cannot commit to avoid, and the deflation bias created by the ZLB, which only a credible promise of higher future inflation can offset.

In the paper, I showed how both inefficiencies interact when the monetary policy is discretionary. The main conclusion is that when the effects of monopolistic competition are present, the deflationary pressure of the zero lower bound is welcomed in terms of welfare, as shown by Nakata (2016) for the polar cases. In this case, the ZLB becomes a mechanism that disciplines the Central Bank with a more downward inflation bias. How this happens on average and along the exogenous state depends on the calibration of the model, particularly how persistent the shocks are. Regardless, the welfare gap between the economy with and without the ZLB closes when the monetary policy is conducted with higher levels of commitment, closing entirely in the fully committed equilibrium (i.e., making the effects of the ZLB negligible). Along the commitment line, the chances of hitting the zero lower bound start increasing with higher levels of commitment. The main driver here is the inflation bias falling faster than the benefits of mid/low-level increases in credibility. This trend is not homogeneous since the probability starts falling when the peak is reached, in the calibration I choose, at around 55% of commitment. The likelihood of hitting the ZLB is roughly four times larger in commitment than discretion, and the peak is five times larger than discretion. The length of the ZLB spells also changes along the commitment space. In particular, the relatively low probability of hitting the ZLB closer to discretion is associated with shorter episodes (less than four periods). In contrast, more extended periods with a nominal interest rate equal to zero are expected with higher levels of commitment. Again, this trend is not homogeneous with shorter spells in the fully committed economy than the committed 95% of the time. As I mentioned above, inflation is higher in discretion but less volatile. The main reason is the lack of need to increase inflation against a negative preferences shock, given that there is a relatively sizeable inflationary bias. With the increase in commitment, inflation becomes more volatile since the ZLB are more frequent and need higher inflation levels than regular times. In low to mid levels of commitment, the bulk of the inflation volatility comes from periods where the ZLB does not bind. This trend is reverted once the economy gets closer to being fully committed. Here the Central Bank chooses low and less volatile inflation in regular times and higher and more volatile inflation when hit by the ZLB. The evidence from the full non-linear model suggests that it is difficult to identify the level of commitment by looking at the length and probability of the ZLB episodes.

Throughout the paper, the two exogenous shocks are analyzed one at a time rather than jointly.

This is a deliberate scope choice rather than an oversight: loose commitment already requires tracking two endogenous "promise" state variables in addition to whichever exogenous shock is active, and Section 3's solution method is itself built around estimating a non-regular grid to keep that three-dimensional problem tractable. Adding a second simultaneous exogenous shock would add a further dimension on top of an already curse-of-dimensionality-constrained state space, at a computational cost that I leave for future work.

References

Adam, K. and R. M. Billi (2006): "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit and Banking*, 38, 1877–1905. pages 4

——— (2007): "Discretionary monetary policy and the zero lower bound on nominal interest rates," *Journal of Monetary Economics*, 54, 728–752. pages 4

Adam, K. and H. Weber (2019): "Optimal Trend Inflation," *American Economic Review*, 109, 702–737. pages 5, 9

Amidror, I. (2002): "Scattered data interpolation methods for electronic imaging systems: a survey," *Journal of Electronic Imaging*, 11, 157–76. pages 67

Armenter, R. (2008): "A General Theory (and Some Evidence) of Expectation Traps in Monetary Policy," *Journal of Money, Credit and Banking*, 40, 867–895. pages 13

Auclert, A., R. Rigato, M. Rognlie, and L. Straub (2024): "New Pricing Models, Same Old Phillips Curves?" *The Quarterly Journal of Economics*, 139, 121–186. pages 5

Barro, R. J. and D. B. Gordon (1983): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91, 589–610. pages 3

Blanco, A. (2021): "Optimal Inflation Target in an Economy with Menu Costs and a Zero Lower Bound," *American Economic Journal: Macroeconomics*, 13, 108–141. pages 5, 9

Blinder, A. S. (2000): "Central-Bank Credibility: Why Do We Care? How Do We Build It?" *American Economic Review*, 90, 1421–1431. pages 4

Bodenstein, M., J. Hebden, and R. Nunes (2012): "Imperfect credibility and the zero lower bound," *Journal of Monetary Economics*, 59, 135–149. pages 3, 4, 6, 15

Bonciani, D. and J. Oh (2025): "Optimal monetary policy mix at the zero lower bound," *Journal of Economic Dynamics and Control*, 170, 105001. pages 5

Calvo, G. A. (1978): "On the Time Consistency of Optimal Policy in a Monetary Economy," *Econometrica*, 46, 1411–1428. pages 3

Chari, V. V. and P. J. Kehoe (1990): "Sustainable Plans," *Journal of Political Economy*, 98, 783–802. pages 13

Coleman, II, W. J. (1991): "Equilibrium in a Production Economy with an Income Tax," *Econometrica*, 59, 1091–1104. pages 5, 17, 66, 70

Debortoli, D. and R. Nunes (2010): "Fiscal policy under loose commitment," *Journal of Economic Theory*, 145, 1005–1032. pages 3, 4, 6, 15

Dixit, A. K. and J. E. Stiglitz (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297–308. pages 8

Eggertsson, G. B. and S. R. Singh (2019): "Log-linear Approximation versus an Exact Solution at the ZLB in the New Keynesian Model," *Journal of Economic Dynamics and Control*, 105, 21–43. pages 3

Eggertsson, G. B. and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34, 139–235. pages 4

Gust, C., E. Herbst, D. López-Salido, and M. E. Smith (2017): "The Empirical Implications of the Interest-Rate Lower Bound," *American Economic Review*, 107, 1971–2006. pages 5

Ireland, P. N. (2007): "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," *Journal of Money, Credit and Banking*, 39, 1851–1882. pages 9

Judd, K. L. (1992): "Projection methods for solving aggregate growth models," *Journal of Economic Theory*, 58, 410–452. pages 69

Kydland, F. E. and E. C. Prescott (1977): "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 473–491. pages 3

Maliar, L. and S. Maliar (2015): "Merging simulation and projection approaches to solve high-dimensional problems with an application to a New Keynesian model," *Quantitative Economics*, 6, 1–47. pages 5, 17, 66, 69

Marcet, A. and R. Marimon (2011): "Recursive Contracts," CEP Discussion Papers 1055, Centre for Economic Performance, LSE, revised March 2017. pages 15

——— (2019): "Recursive Contracts," *Econometrica*, 87, 1589–1631. pages 14, 15, 62

Miranda, M. J. and P. L. Fackler (2004): *Applied Computational Economics and Finance*, vol. 1, The MIT Press. pages 17

Nakata, T. (2016): "Optimal fiscal and monetary policy with occasionally binding zero bound constraints," *Journal of Economic Dynamics and Control*, 73, 220–240. pages 5, 6, 18, 34

Nakov, A. (2008): "Optimal and Simple Monetary Policy Rules with Zero Floor on the Nominal Interest Rate," *International Journal of Central Banking*, 4, 73–127. pages 4

Nunes, R., D. Park, and L. Rondina (2021): "Imperfect credibility, sticky wages, and welfare," *Journal of Macroeconomics*, 70. pages 4

Rotemberg, J. J. (1982a): "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies*, 49, 517–531. pages 5, 6

——— (1982b): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211. pages 9, 61, 63

Rupert, P. and R. Šustek (2019): "On the mechanics of New-Keynesian models," *Journal of Monetary Economics*, 102, 53–69. pages 5

Schaumburg, E. and A. Tambalotti (2007): "An investigation of the gains from commitment in monetary policy," *Journal of Monetary Economics*, 54, 302–324. pages 3

Siu, H. E. (2004): "Optimal fiscal and monetary policy with sticky prices," *Journal of Monetary Economics*, 51, 575–607. pages 4

Taylor, J. B. (1993): "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214. pages 3

Yun, T. (2005): "Optimal Monetary Policy with Relative Price Distortions," *American Economic Review*, 95, 89–109. pages 4

Appendix

A Additional Results

A.1 Discretion

In the main body of the paper, I show selected parts of the inflation and loss functions because they depict the most important facts of the model. For completeness, in this section, I show all the policy and loss functions.I divide the section in two. The first shows the results for the total factor productivity including some exercises with alternative parameters — less competitive markets and higher elasticity of intertemporal substitution. In the second, I show the complete results for the preferences shock.

Total Factor Productivity shock

Figure 10 shows the discretionary policy functions for different levels of persistence of z_t (color) for the benchmark calibration. The panels display the optimal discretionary inflation, interest, consumption, and labor for different TFP shock persistence (ρ_z). The left panels present results with monopolistic competition inefficiency, while the right panels show the results with efficient labor subsidy and $\pi_f = 0\%$. The dotted vertical lines indicate where the zero lower bound begins to bind for each ρ_z. The dot-dashed lines represent the inflation policies for the model without the ZLB, which overlaps in some panels. The shaded area represents the ergodic probability density function of z.

Figure 10: Policy functions under discretion and TFP shocks, benchmark calibration

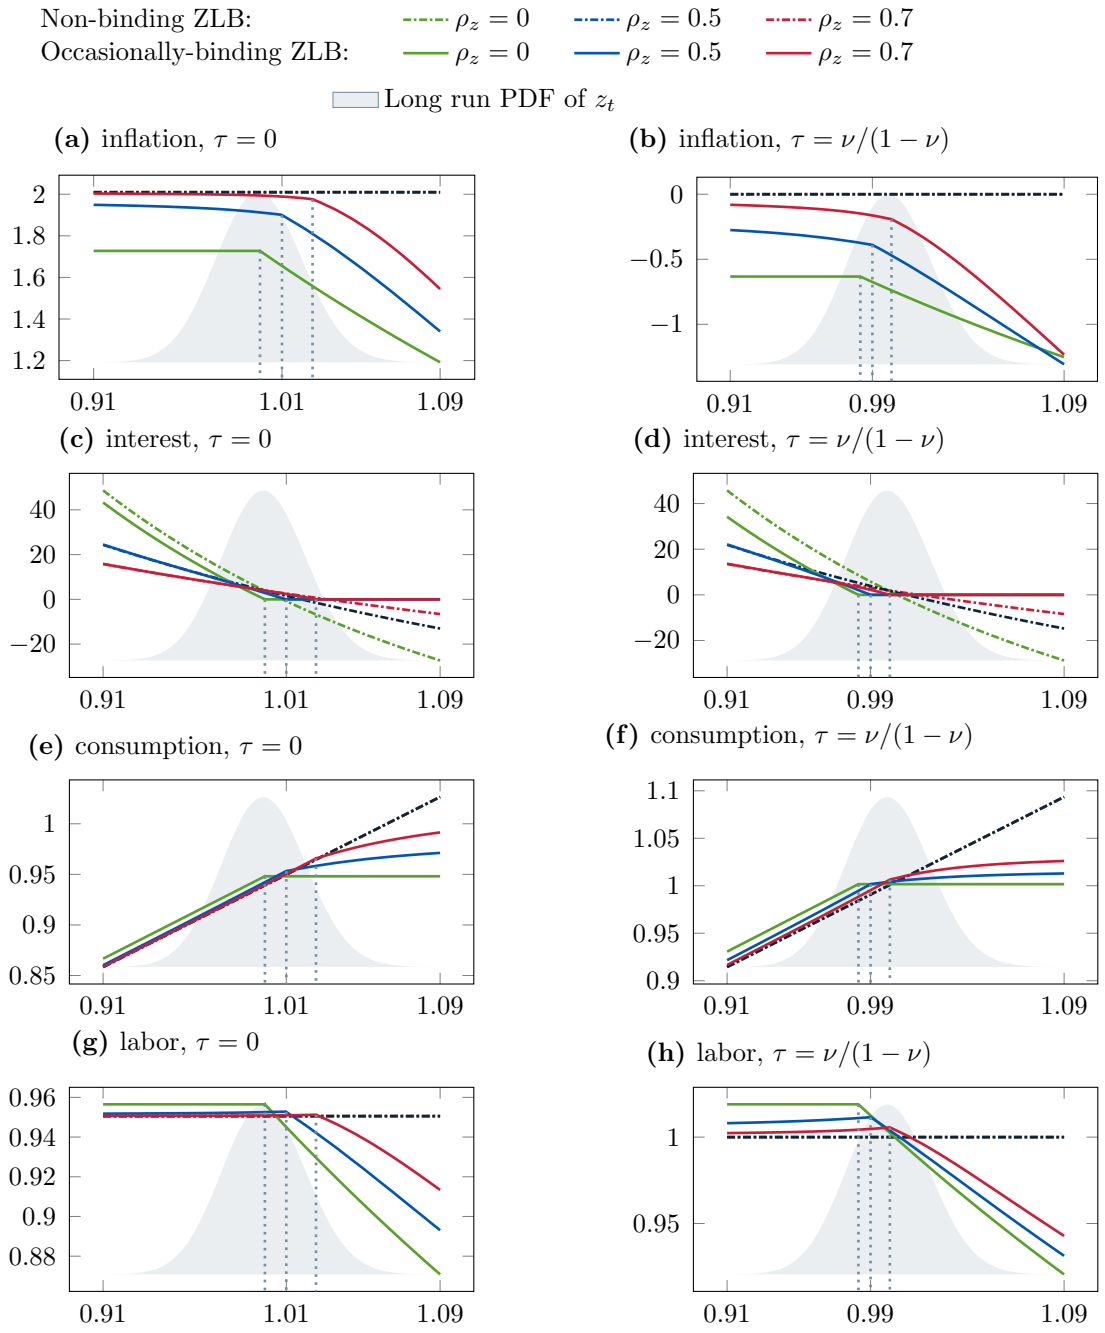

Figure 11 shows the discretionary policy functions for the same parameters space as in the benchmark calibration but with a less competitive market. In particular $\nu = 1/6$. The panels display the optimal discretionary inflation, interest, consumption, and labor for different TFP shock persistence (ρ_z). The left panels present results with monopolistic competition inefficiency, while the right panels show the results with efficient labor subsidy and $\pi_f = 0\%$. The dotted vertical lines indicate where the zero lower bound begins to bind for each ρ_z. The dot-dashed lines represent the inflation policies for the model without the ZLB, which overlaps in some panels. The shaded area represents the ergodic probability density function of z.

Figure 11: Policy functions under discretion and TFP shocks, $\nu = 0.167$

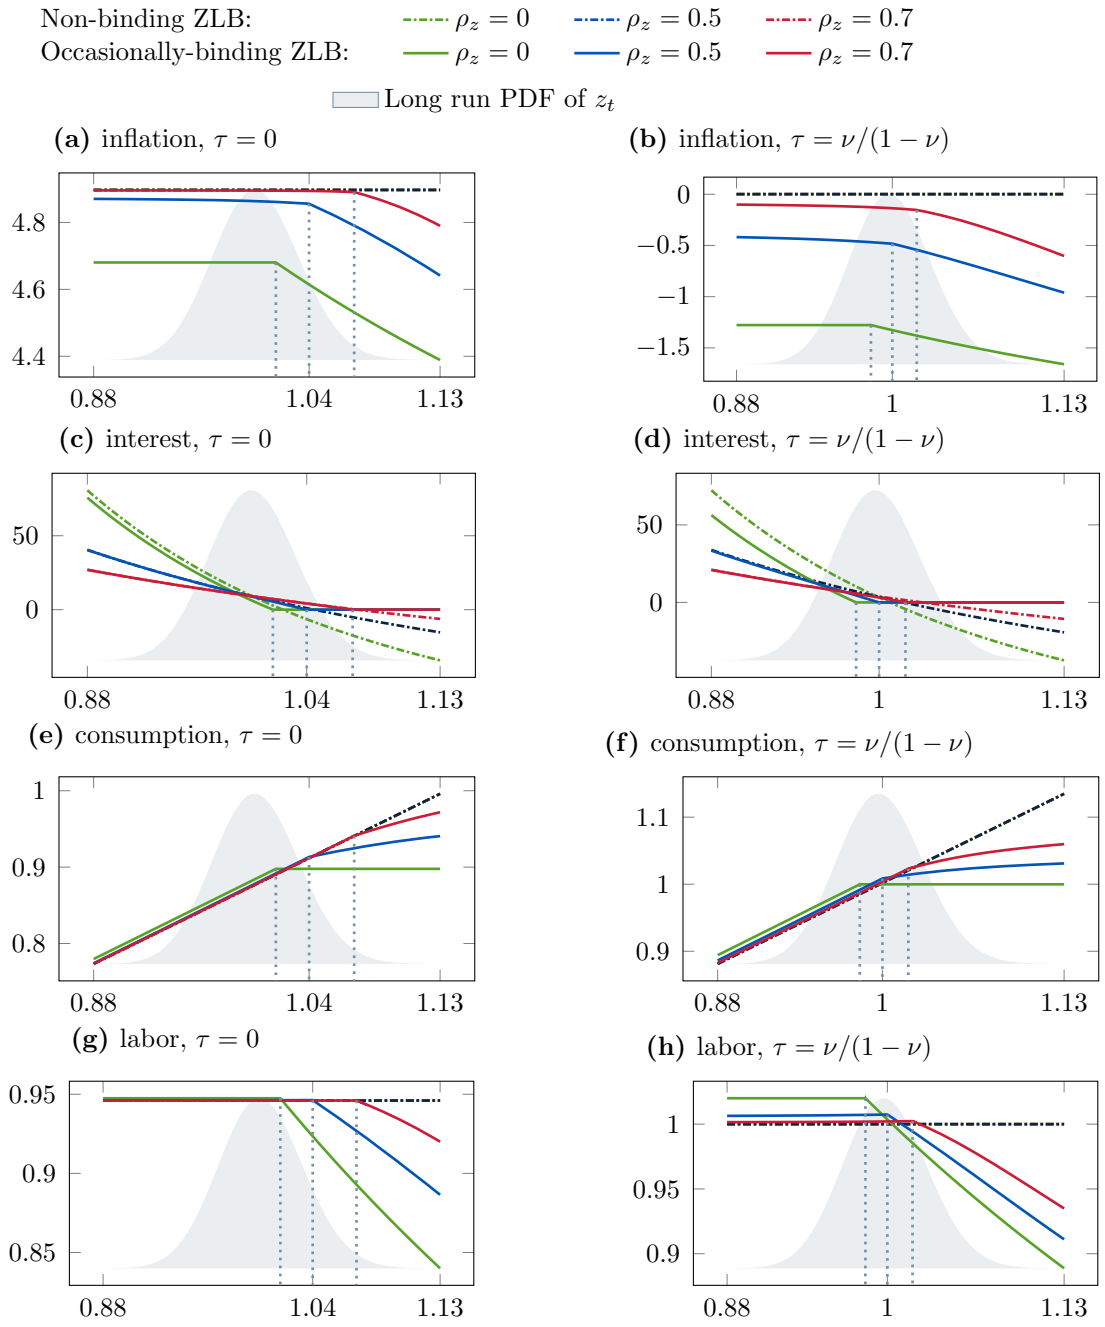

Figure 12 shows the discretionary policy functions for the same parameters space as in the benchmark calibration but with $\sigma = 2$. The panels display the optimal discretionary inflation, interest, consumption, and labor for different TFP shock persistence (ρ_z). The left panels present results with monopolistic competition inefficiency, while the right panels show the results with efficient labor subsidy and $\pi_f = 0\%$. The dotted vertical lines indicate where the zero lower bound begins to bind for each ρ_z. The dot-dashed lines represent the inflation policies for the model without the ZLB, which overlaps in some panels. The shaded area represents the ergodic probability density function of z.

Figure 12: Policy functions under discretion and TFP shocks, $\sigma = 2$

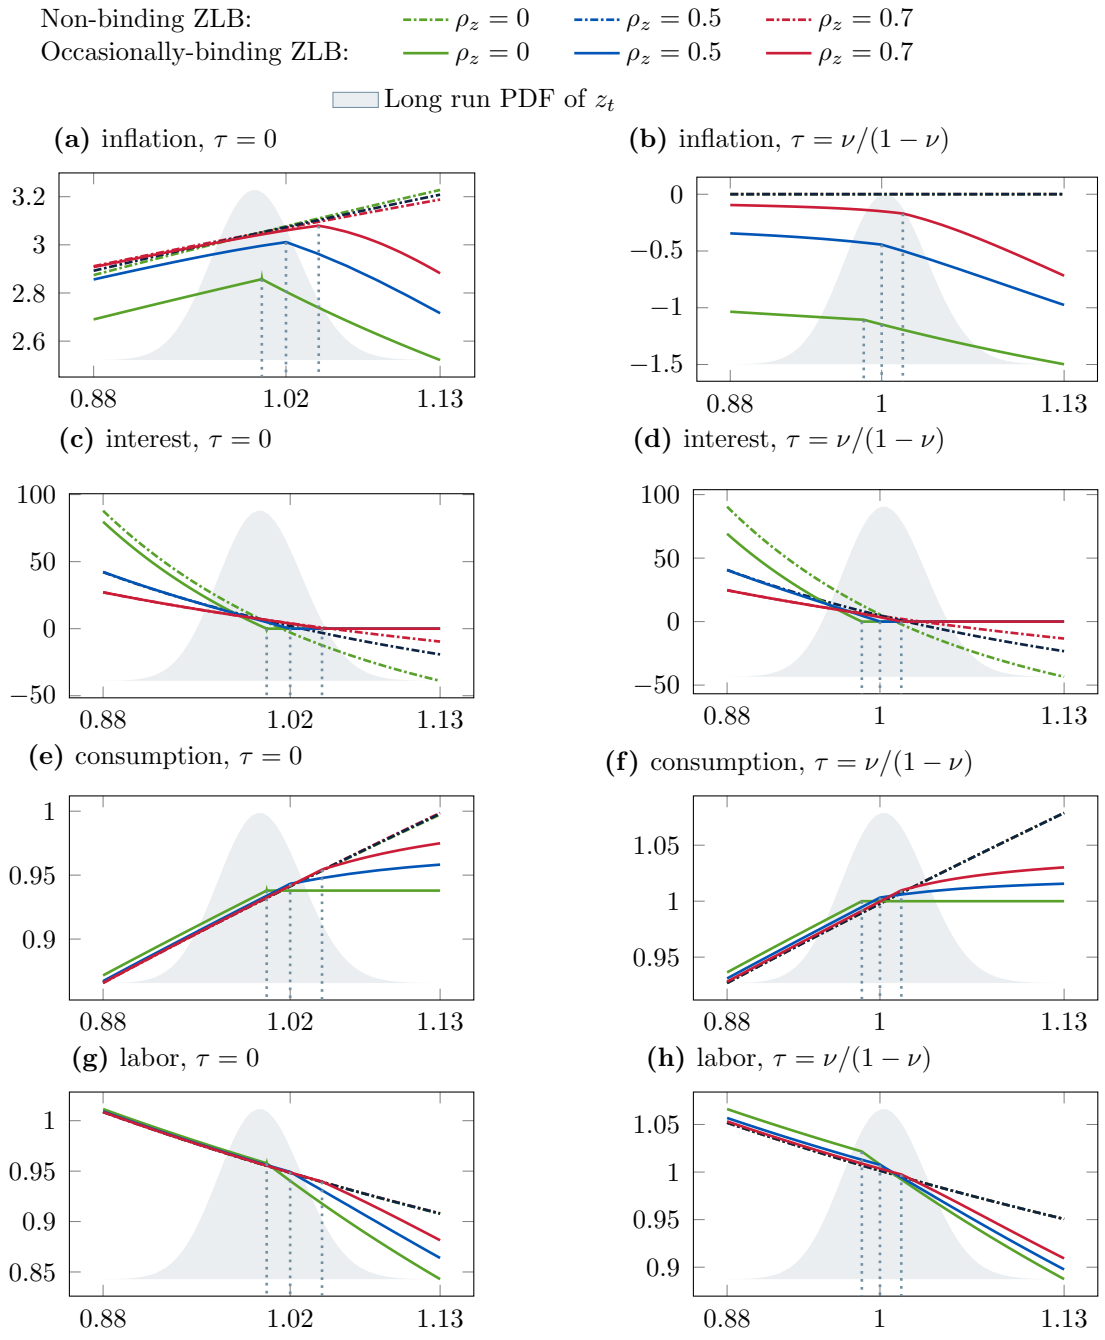

Figure 13 shows the loss function for the benchmark calibration and TFP shocks. The panels show the discretionary welfare loss for selected levels of persistence of the TFP shock ρ_z. The left panels show the results with the monopolistic competition inefficiency, while the right ones do it with factor inputs as in perfect competition. The dotted vertical lines show where the zero lower bound starts binding for each ρ_z. The dot-dashed lines show the inflation policy for the model without the ZLB, and the shaded area is the ergodic PDF of z.

Figure 13: Loss function under discretion and TFP shocks, benchmark calibration

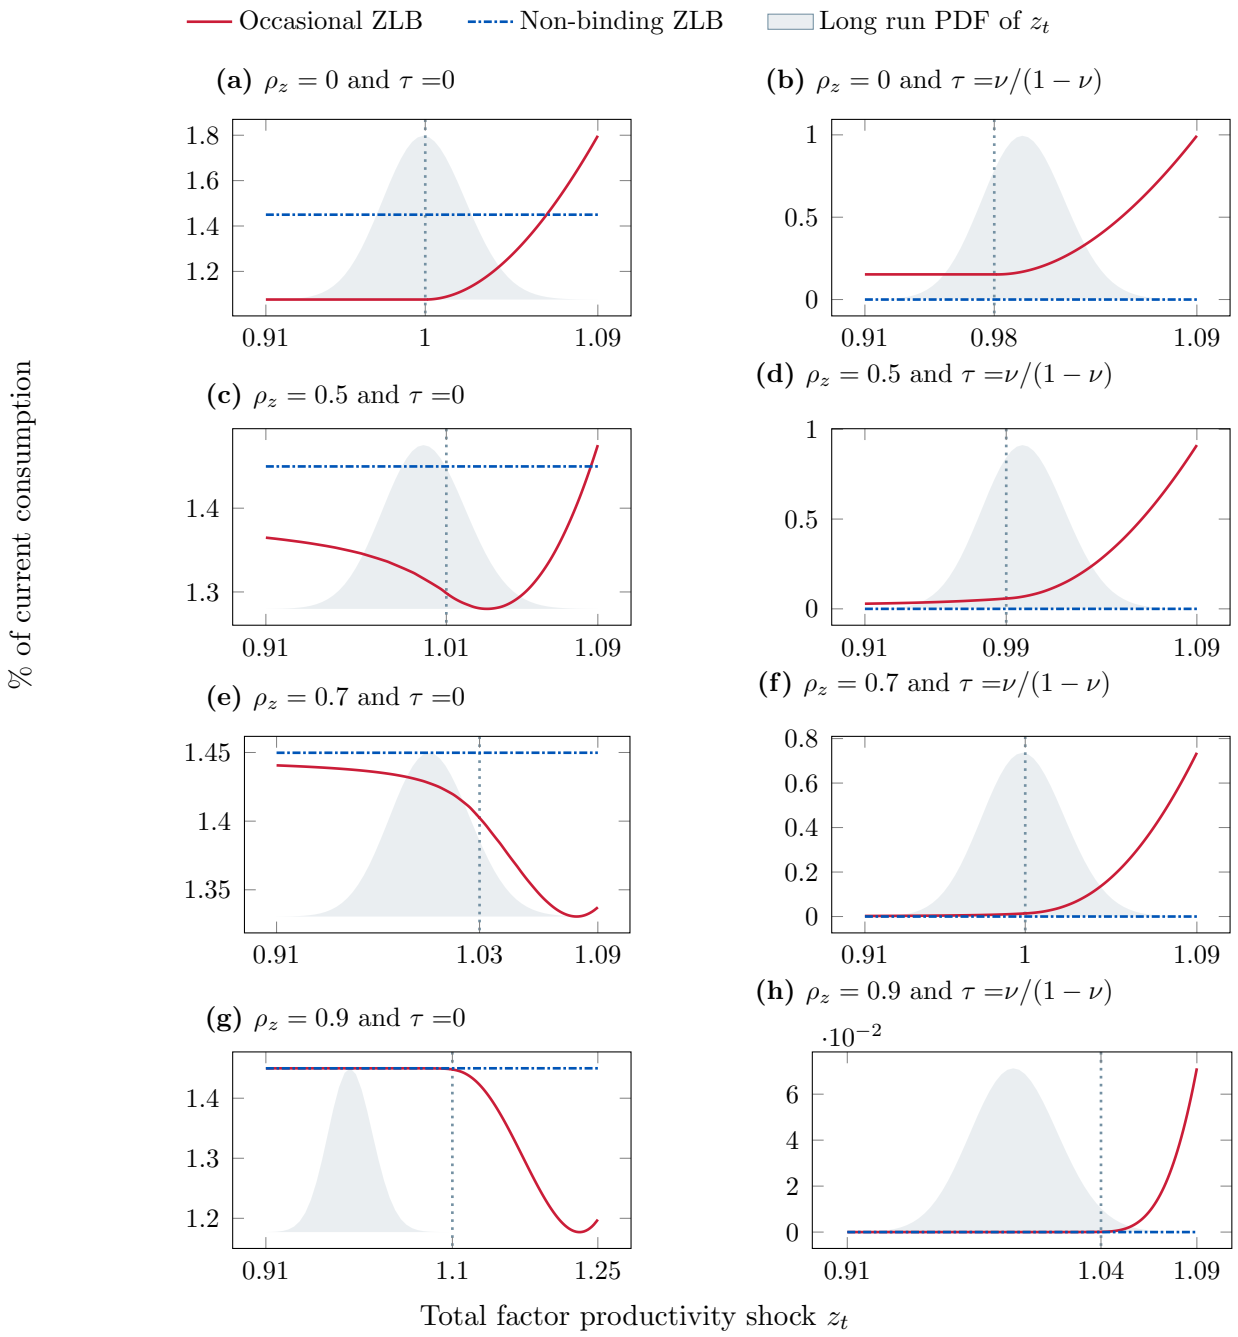

Figure 14 shows the loss function for a less competitive economy than the benchmark calibration. In particular, the elasticity of substitution among intermediate goods is equal to 6 (i.e. $\nu = 0.167$). The panels show the discretionary welfare loss for selected levels of persistence of the TFP shock ρ_z. The left panels show the results with the monopolistic competition inefficiency, while the right ones do it with factor inputs as in perfect competition. The dotted vertical lines show where the zero lower bound starts binding for each ρ_z. The dot-dashed lines show the inflation policy for the

model without the ZLB, and the shaded area is the ergodic PDF of z.

Figure 14: Loss function under discretion and TFP shocks, $\nu = 0.167$

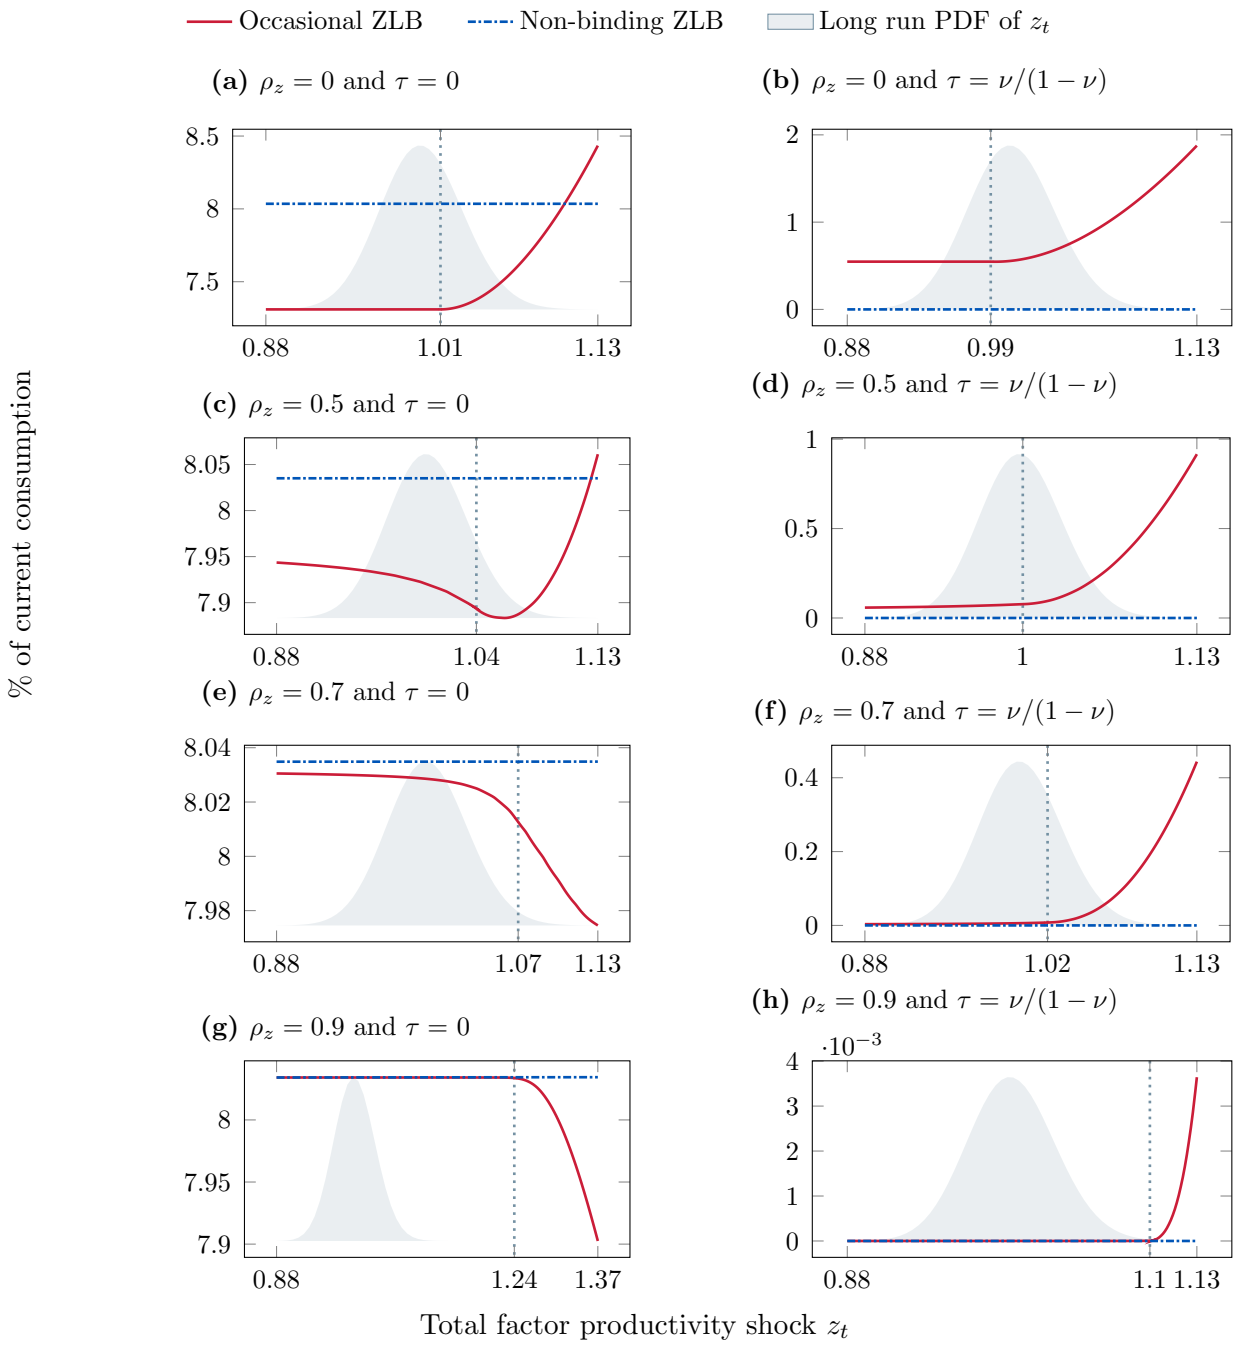

Figure 15 shows the loss function for $\sigma = 2$. The panels show the discretionary welfare loss for selected levels of persistence of the TFP shock ρ_z. The left panels show the results with the monopolistic competition inefficiency, while the right ones do it with factor inputs as in perfect competition. The dotted vertical lines show where the zero lower bound starts binding for each ρ_z. The dot-dashed lines show the inflation policy for the model without the ZLB, and the shaded area is the ergodic PDF of z.Comparing Figure 12 to the benchmark $\sigma = 1$ case in Section 5.1, the slope of the discretionary inflation policy function with respect to the TFP shock flips sign, confirming that this particular comparative static is not general. The welfare ranking in Figure 15, however, is unchanged from the benchmark: the ZLB still lowers the discretionary welfare loss when the monopolistic-competition inefficiency is present (left panels) and still raises it in the efficient benchmark (right panels). This robustness check is only available for the discretionary equilibrium; the loose-commitment results in Section 5.2 have not been re-solved at $\sigma = 2$.

Figure 15: Loss function under discretion and TFP shocks, $\sigma = 2$

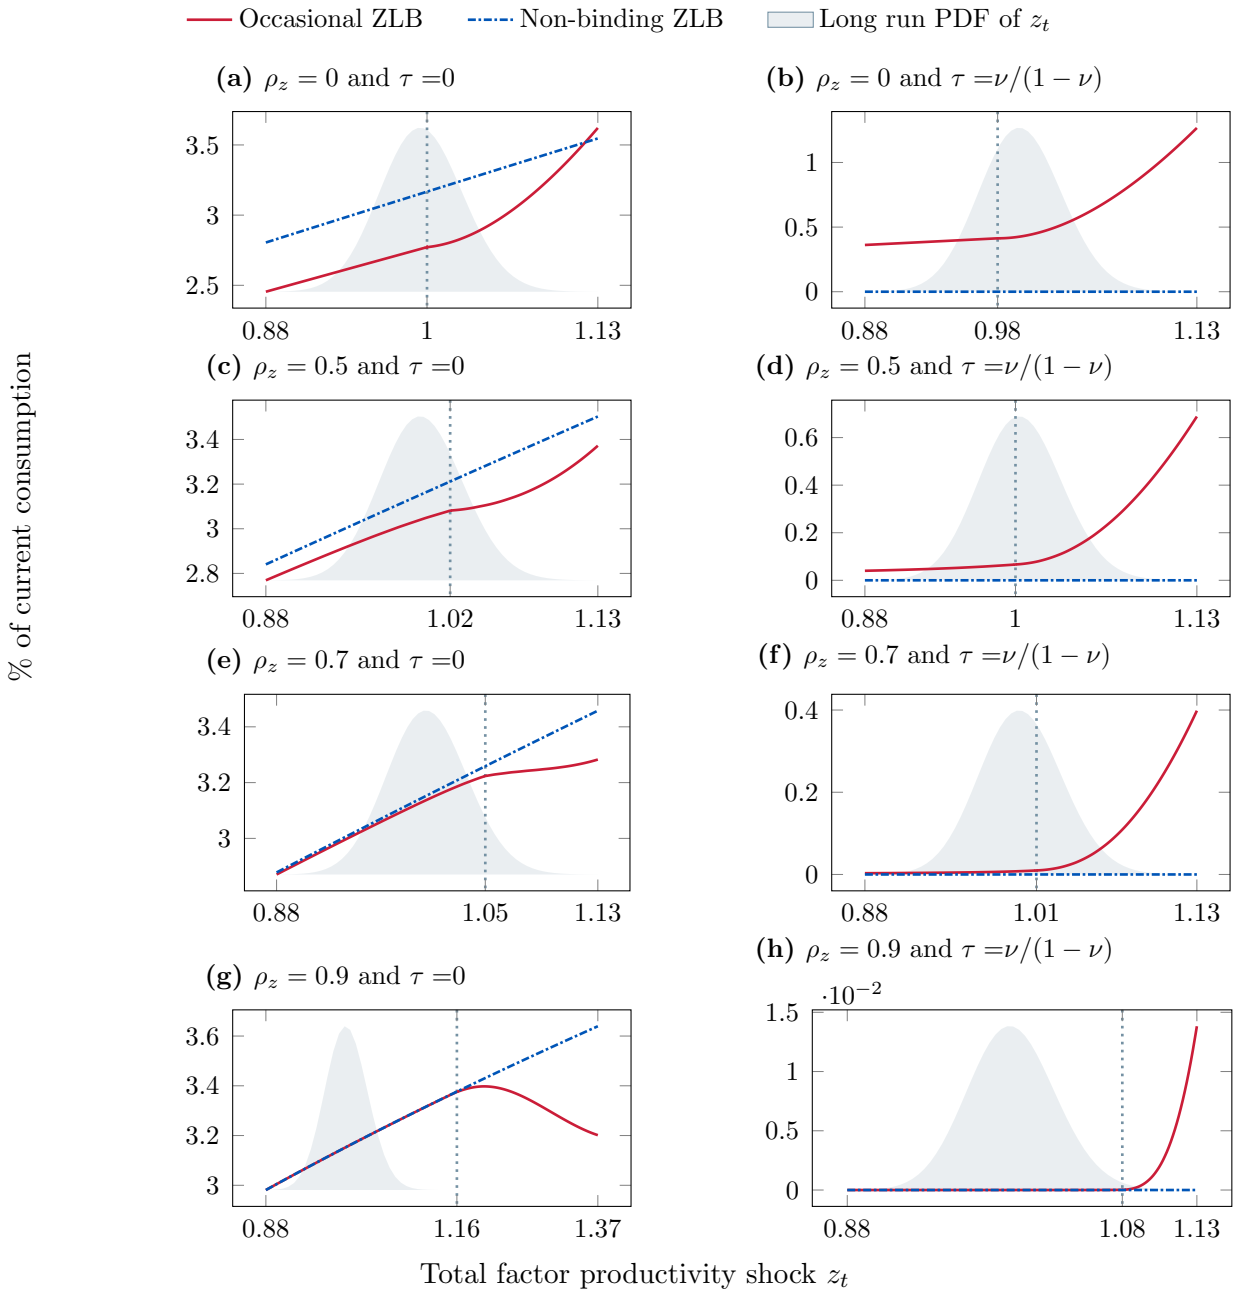

Preferences shock

In this section, I show the complete policy and loss functions when the driver of uncertainty is the discount factor shock, δ_t. Figure 16 shows policy functions for the benchmark calibration. The panels display the optimal discretionary inflation, interest, consumption, and labor for different demand shock persistence (ρ_δ). The left panels present results with monopolistic competition inefficiency, while the right panels show the results with efficient labor subsidy and $\pi_f = 2\%$. The dotted vertical lines indicate where the zero lower bound begins to bind for each ρ_δ. The dot-dashed

lines represent the inflation policies for the model without the ZLB, which overlaps in some panels. The shaded area represents the ergodic probability density function of δ.I choose a higher level of $\pi_f = 2\%$ to reduce the probability of hitting the ZLB. Otherwise, the ZLB becomes too dominant and the algorithm fails to find a solution.

Figure 16: Policy functions under discretion and preferences shocks, benchmark calibration

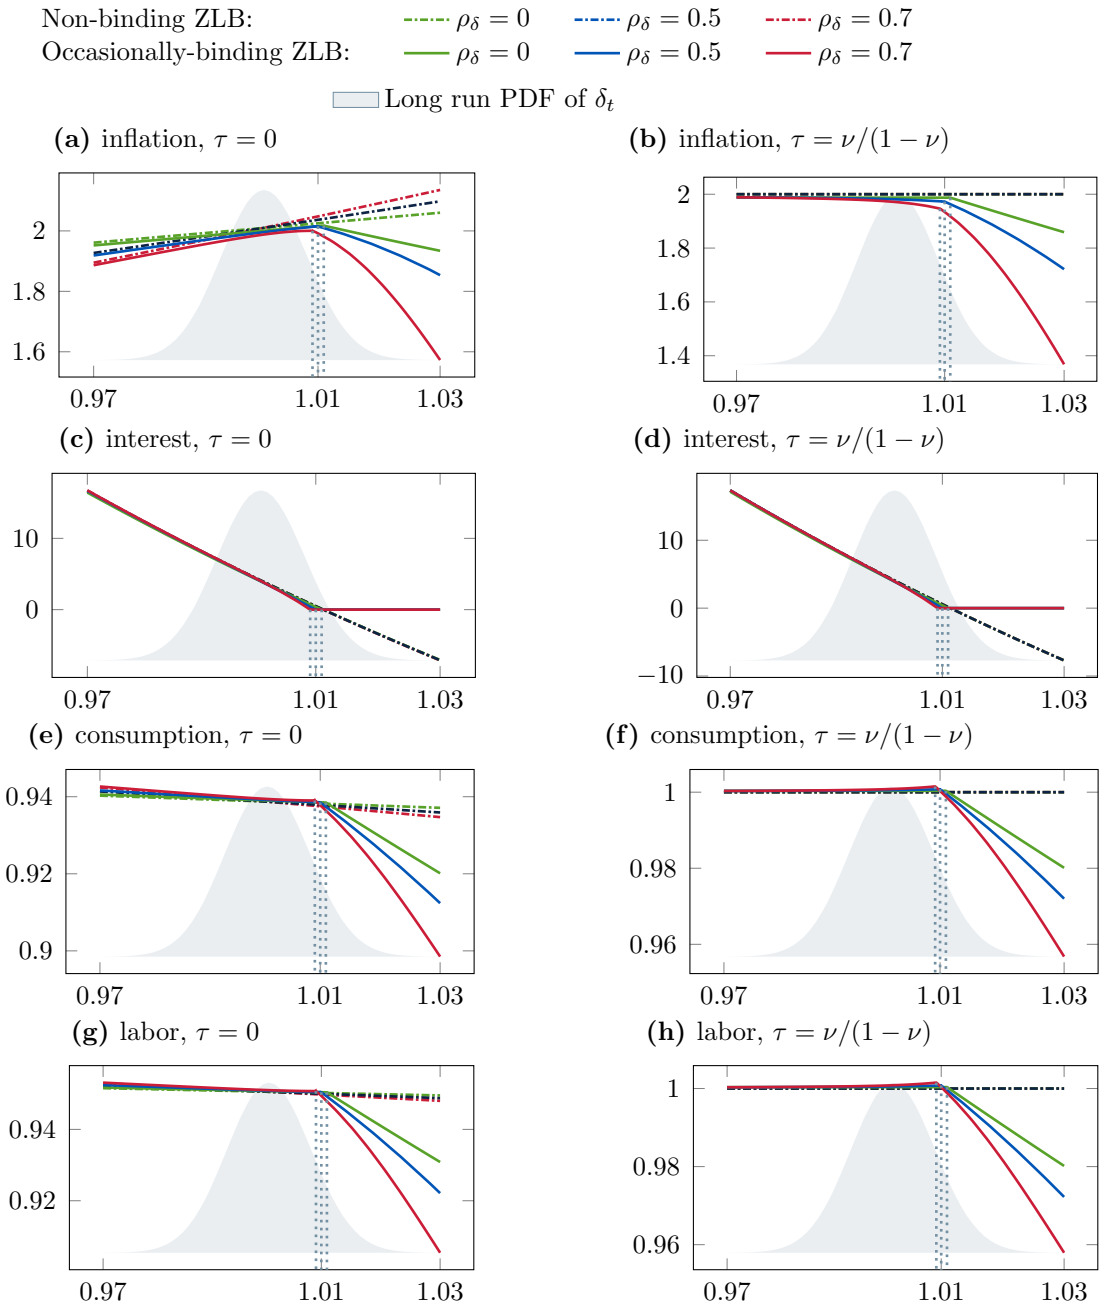

Figure 17 shows the loss function for the model with the benchmark calibration and demand shocks. The panels show the discretionary welfare loss for selected levels of persistence of the demand shock ρ_δ. The left panels show the results with the monopolistic competition inefficiency, while the right ones do it with factor inputs as in perfect competition. The dotted vertical lines show where the zero lower bound starts binding for each ρ_δ. The dot-dashed lines show the inflation policy for the model without the ZLB, and the shaded area is the ergodic PDF of δ.

Figure 17: Loss function under discretion and preferences shock, benchmark calibration

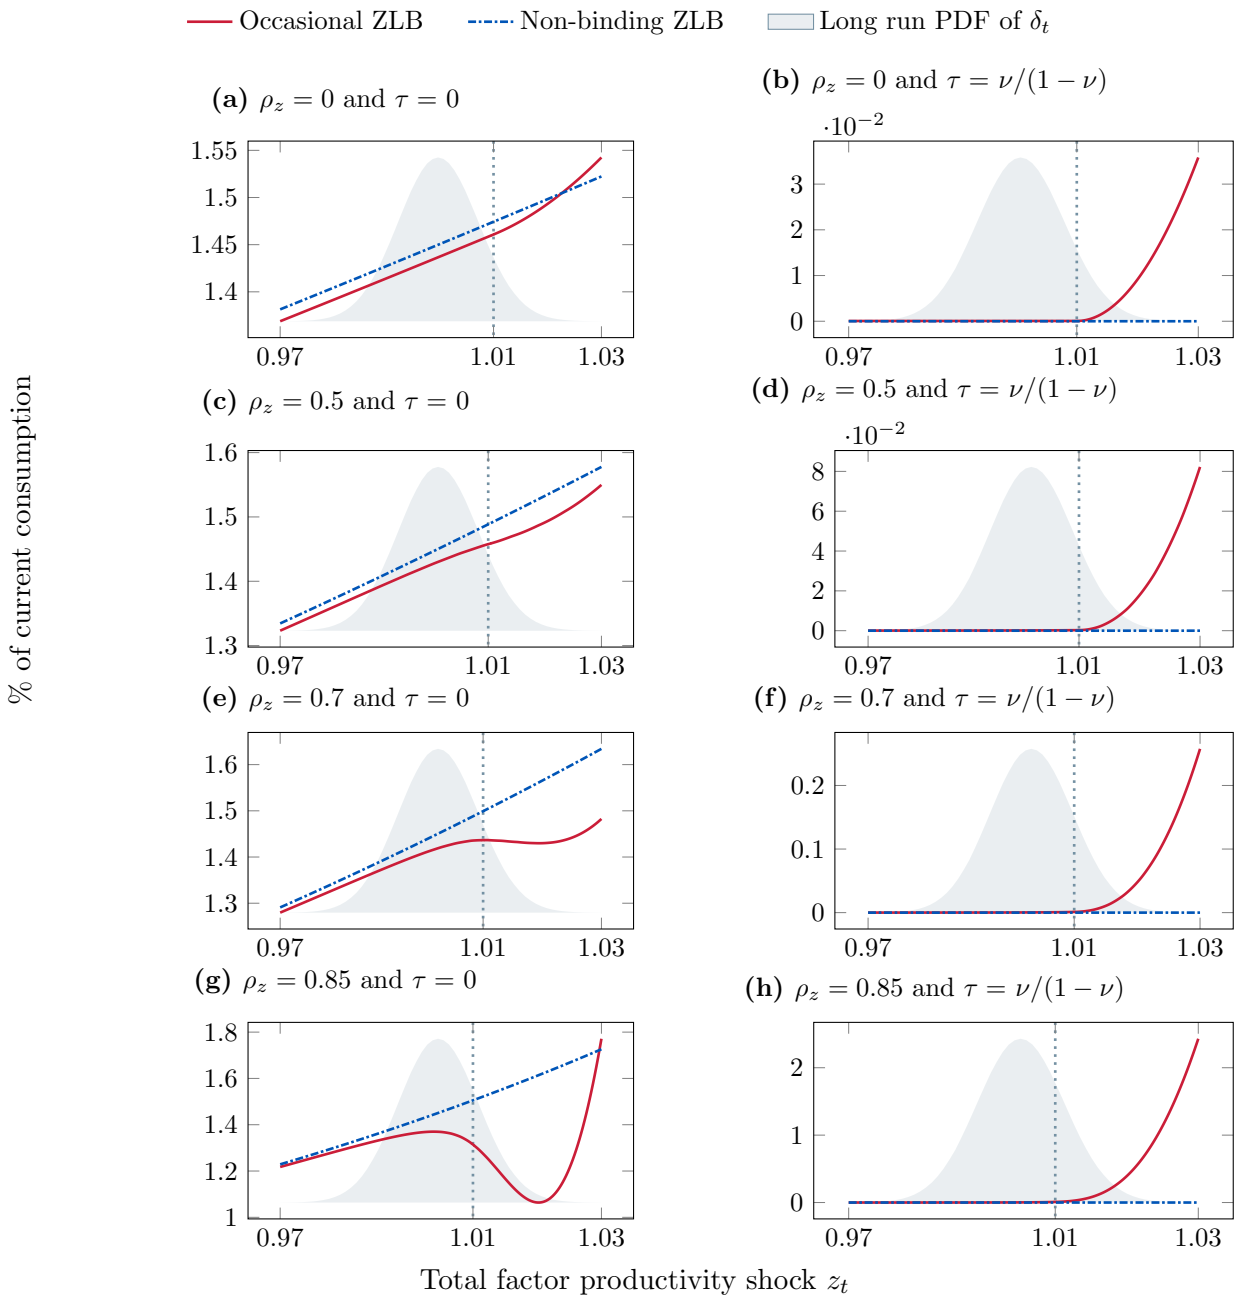

The slope of inflation respect to δ

In the paper, inflation increases with the preferences shock. This is due to the preferences of the household but also amplified by the inflation bias. To see why this happens, use the resource constraint in (43.3) and take total derivatives denoting $\pi_t - \pi_f = B_t$

$$\phi B_t n_t \frac{d\pi_t}{d\delta} = \frac{dc_t}{d\delta} - \frac{dn_t}{d\delta}\left(1 - \frac{\phi}{2}B_t\right)$$

Figure 18: Average Welfare loss in and out the ZLB

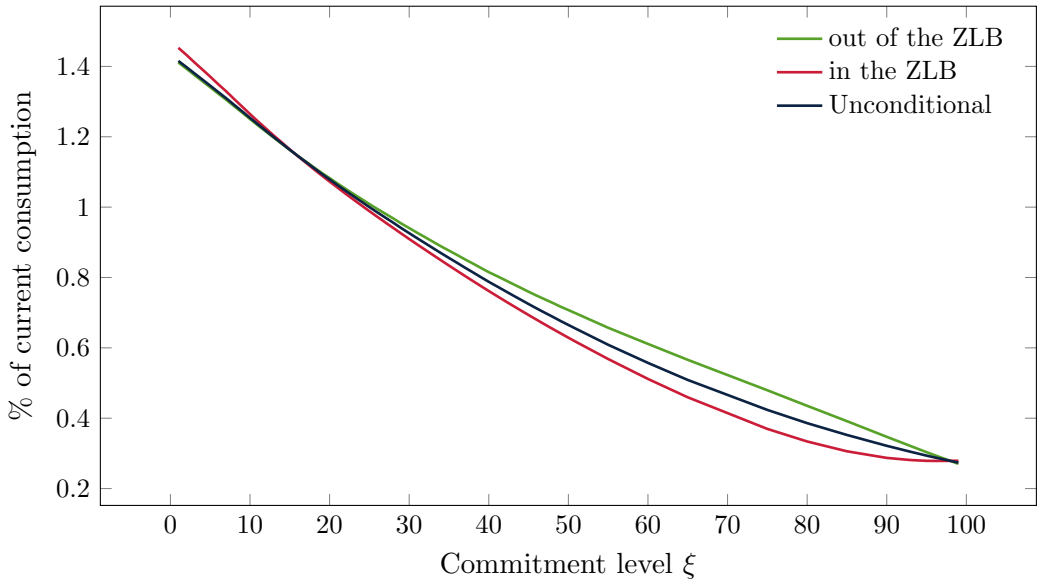

Note: the figure shows the average welfare loss in terms of current consumption for every level of commitment, conditional on the economy being in the ZLB or not.

Note that inflation bias implies $B_t > 0$ and that ϕ and n are positive. With this $\frac{d\pi_t}{d\delta} > 0$ requires that

$$\frac{dc_t}{d\delta} > \frac{dn_t}{d\delta}\left(1 - \frac{\phi}{2}B_t\right)$$

Which is more likely with increasing inflation bias.

A.2 Loose Commitment

Figure 18 shows that the economy suffers more in the ZLB than out of it only on lower levels of commitment.

Figure 19: Average Welfare loss in the first commitment spell

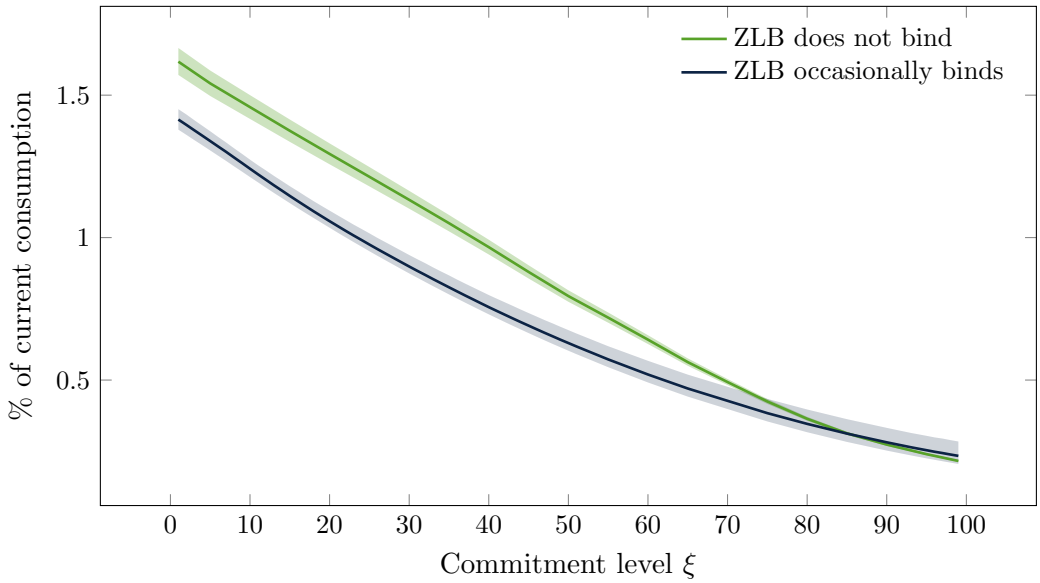

Note: the figure shows the average welfare loss in terms of current consumption for every level of commitment, conditional on the monetary policy being just reneged.

Figure 20: Average Welfare loss conditional on the commitment spell

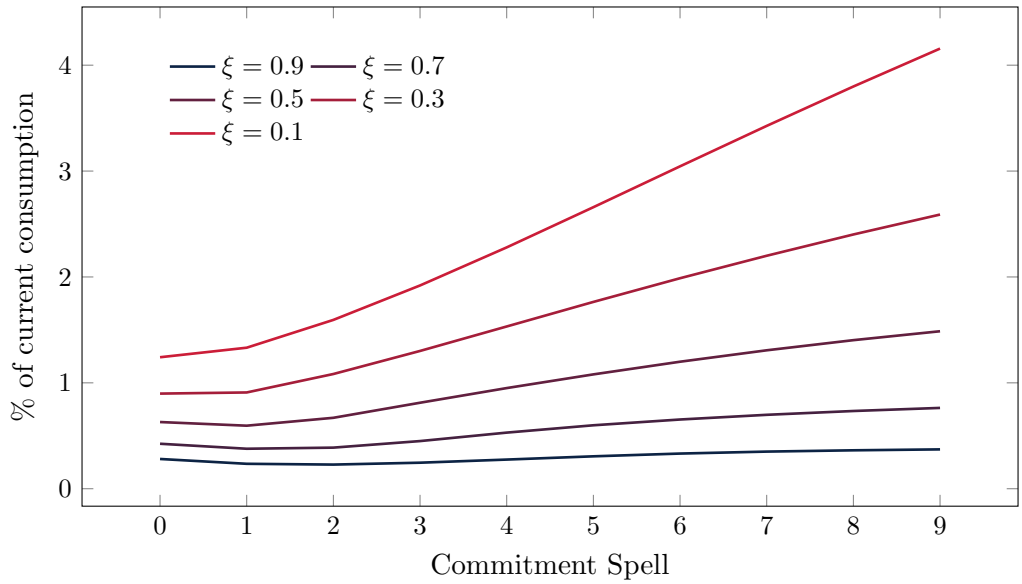

Note: the figure shows the average welfare loss in terms of current consumption for the first 10 commitment spells and selected levels of commitment.

Figure 21: Probability of commitment spell

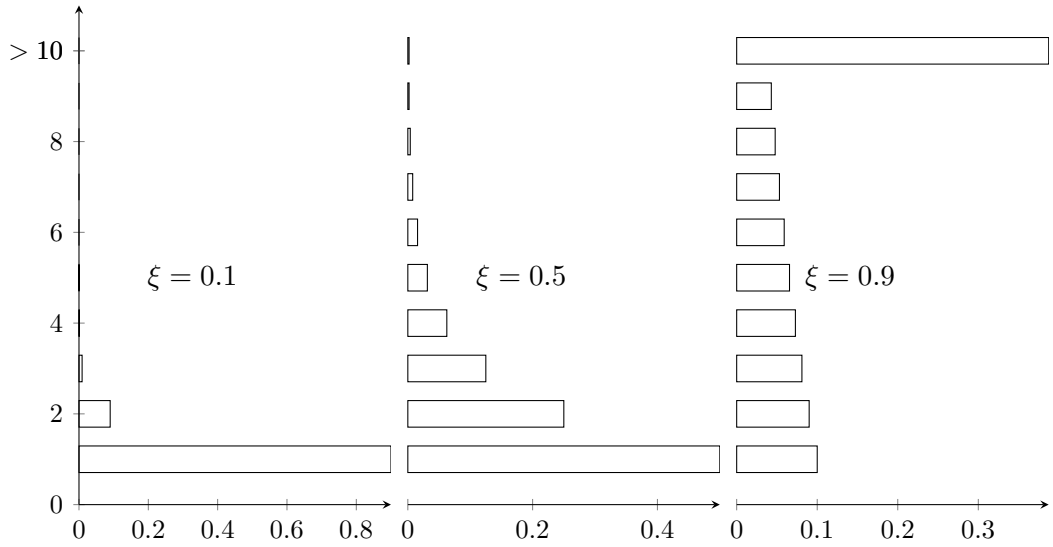

Note: the figure shows the probability of commitment spell for selected levels of commitment. The figure points out that the larger probability of each spell is maximum at spell zero. Nonetheless, higher commitment levels accumulate more probability mass at higher spells.

Figure 22: Probability of hitting the ZLB

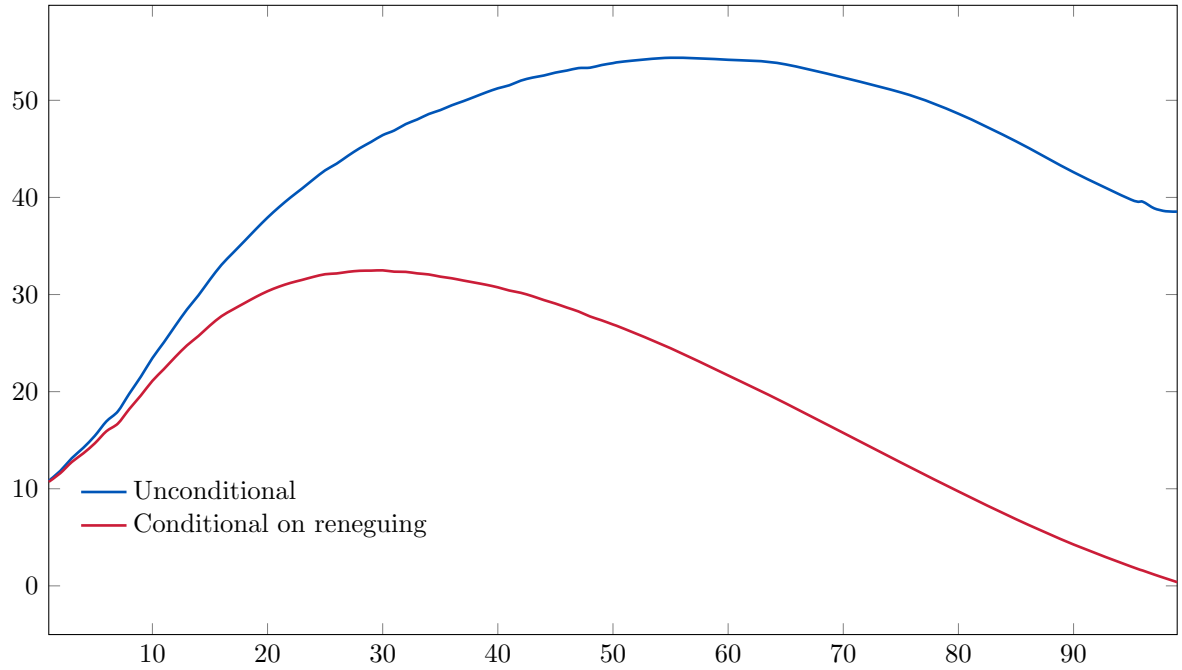

Note: The figure shows how the unconditional ZLB probability and conditional on reneging vary with commitment level.

Figure 23 shows a simulation of the model with demand shocks for several levels of commitment. The simulations have a burnout sample of 900 observations. Inflation is in annual percent points while output is in percent deviations from the efficient allocation.

Figure 23: Simulation of the model with loose commitment

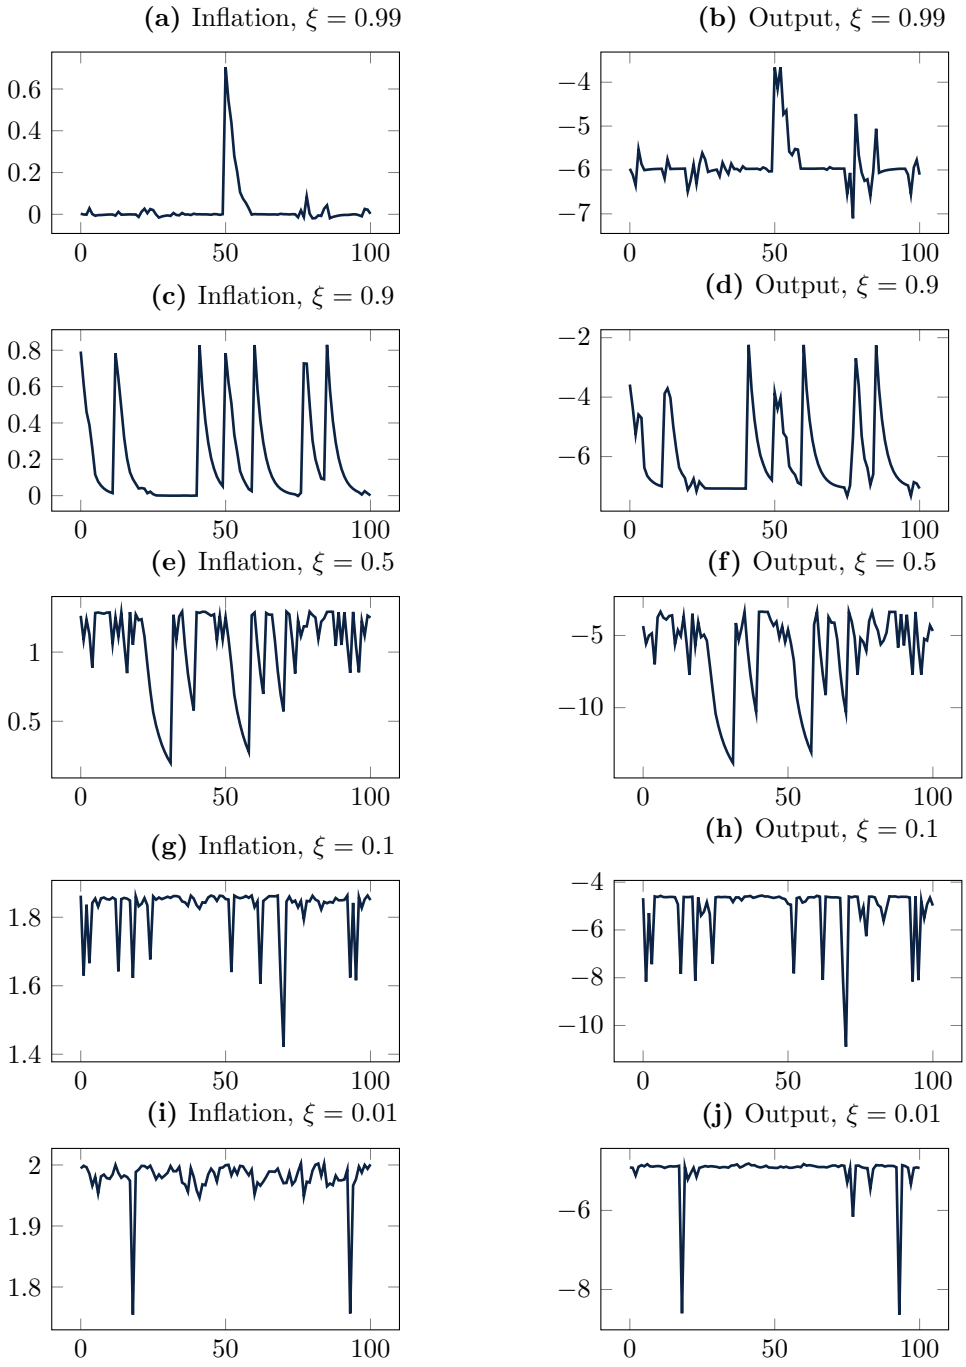

B Model First-Order Conditions and Steady State

In this section, I describe the first-order conditions from which the solution algorithms will iterate. In each subsection, there is a description of the non-stochastic steady-state from which the algorithms will start.

B.1 Discretion

The problem of the central bank in the discretionary world is to maximize the current utility of the household

$$\max_{X_t} \frac{c_t^{1-\sigma}-1}{1-\sigma} - \chi \frac{n_t^{1+\eta}}{1+\eta} \tag{42}$$

Here the vector X_t denote the policy variables R_t, c_t, π_t, and n_t. The previous maximization is subject to the attainable equilibrium:

$$c_t^{-\sigma} R_t^{-1} = \delta_t \beta \mathbb{E}_t \left[\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right] \tag{43.1}$$

$$c_t^{-\sigma} z_t n_t \left[1 - \nu + \phi \nu \pi_t \left(\pi_t - \pi_f\right)\right] - \frac{\chi n_t^{1+\eta}}{1+\tau} = \phi \nu \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} z_{t+1} n_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1} \right] \tag{43.2}$$

$$c_t = z_t n_t \left[1 - \frac{\phi}{2} \left(\pi_t - \pi_f\right)^2 \right] \tag{43.3}$$

$$R_t \geq 1 \tag{43.4}$$

Define as λ_i, with $i \in \{1,2,3,4\}$ as the Lagrange multipliers of the Euler equation, Phillips curve, resource constraint, and ZLB, respectively. The first-order conditions of this problem are:

$$\partial \lambda_t^1 : \quad 0 = R_t^{-1} c_t^{-\sigma} - \beta \delta_t \mathbb{E}_t \left[\frac{c_{t+1}^{-\sigma}}{\pi_{t+1}} \right] \tag{44.1}$$

$$\partial \lambda_t^2 : \quad 0 = c_t^{-\sigma} n_t z_t \left(1 - \nu + \phi \nu \pi_t \left(\pi_t - \pi_f\right)\right) - \frac{\chi n_t^{1+\eta}}{1+\tau} - \phi \nu \beta \delta_t \mathbb{E}_t \left[c_{t+1}^{-\sigma} z_{t+1} n_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1} \right] \tag{44.2}$$

$$\partial \lambda_t^3 : \quad 0 = c_t - z_t n_t \left[1 - \tau_1 \frac{\phi}{2} \left(\pi_t - \pi_f\right)^2 \right] \tag{44.3}$$

$$\partial \lambda_t^4 : \quad 0 = \lambda_t^4 (R_t - 1) \tag{44.4}$$

$$\partial c_t : \quad 0 = c_t^{-\sigma} - \sigma \lambda_t^1 R_t^{-1} c_t^{-1-\sigma} - \sigma \lambda_t^2 c_t^{-1-\sigma} z_t n_t \left[1 - \nu + \phi \nu \pi_t \left(\pi_t - \pi_f\right)\right] - \lambda_t^3 \tag{44.5}$$

$$\partial n_t : \quad 0 = -\chi n_t^{\eta} + \lambda_t^2 c_t^{-\sigma} z_t \left(1 - \nu + \phi \nu \pi_t \left(\pi_t - \pi_f\right)\right) - \lambda_t^2 (1+\eta) \chi n_t^{\eta} + \lambda_t^3 z_t \left[1 - \tau_1 \frac{\phi}{2} \left(\pi_t - \pi_f\right)^2 \right] \tag{44.6}$$

$$\partial \pi_t : \quad 0 = \lambda_t^2 c_t^{-\sigma} z_t n_t \phi \nu (2\pi_t - \pi_f) - \lambda_t^3 z_t n_t \tau_1 \phi \left(\pi_t - \pi_f\right) \tag{44.7}$$

$$\partial R_t : \quad 0 = \lambda_t^1 R_t^{-2} c_t^{-\sigma} - \lambda_t^4 \tag{44.8}$$

Note that when the economy is not in the ZLB, λ_t^4 equals zero. Under this assumption, and since the interest rate and consumption have to be positive, we can see from (44.8) that λ^1 has to be zero as

well. This relates to the "control" that the monetary authority has in normal times (i.e. non-ZLB). In normal times welfare is not affected by the expectations in the Euler equation since the central bank can always move the interest rate accordingly (i.e. it controls the aggregate demand). In the zero lower bound though, there could be improvements if the central bank was able to conduct forward guidance.Define x_{ss} to the variable x in the non-stochastic steady state where the ZLB does not bind. With this, the previous system becomes:

$$R_{ss} = \frac{\pi_{ss}}{\beta} \tag{45.1}$$

$$\frac{\chi n_{ss}^{\eta} c_{ss}^{\sigma}}{1+\tau} = 1 - \nu + \phi\nu(1-\beta)(\pi_{ss} - \pi_f)\pi_{ss} \tag{45.2}$$

$$c_{ss} = n_{ss}\left[1 - \tau_1 \frac{\phi}{2}\left(\pi_{ss} - \pi_f\right)^2\right] \tag{45.3}$$

$$\lambda_{ss}^3 = c_{ss}^{-\sigma} - \sigma\lambda_{ss}^2 c_{ss}^{-1-\sigma} n_{ss}\left[1 - \nu + \phi\nu\pi_{ss}\left(\pi_{ss} - \pi_f\right)\right] \tag{45.4}$$

$$\begin{aligned}\chi n_{ss}^{\eta} = {} & \lambda_{ss}^2 c_{ss}^{-\sigma}\left(1 - \nu + \phi\nu\pi_{ss}\left(\pi_{ss} - \pi_f\right)\right) \\ & - \lambda_{ss}^2(1+\eta)\chi n_t^{\eta} + \lambda_{ss}^3\left[1 - \tau_1\frac{\phi}{2}\left(\pi_{ss} - \pi_f\right)^2\right]\end{aligned} \tag{45.5}$$

$$\lambda_{ss}^3 = \frac{\lambda_{ss}^2 c_{ss}^{-\sigma}\phi\nu(2\pi_{ss} - \pi_f)}{\tau_1\phi\left(\pi_{ss} - \pi_f\right)} \tag{45.6}$$

The previous system is solved by a unique positive value of π_{ss}. This is done numerically. Once π_{ss} is known equation (a) in the system above pins down the steady state nominal interest rate, equations (b) and (c) identify consumption and labor, and (d) and (e) the Lagrange multipliers.

B.2 Commitment

I assume a New Keynesian economy where the intermediate good firms face a quadratic cost of adjusting prices as in Rotemberg (1982b). The monetary authority chooses the interest rate in order to maximize the discounted utility of the household and it is constrained from below by the Zero Lower Bound (ZLB). The Central Bank chooses its policy under perfect commitment. i.e., it chooses a plan at the beginning of time and follows it with probability 1. The problem at time zero can be resumed as follows:

$$\max_{\{R_t,\pi_t,w_t,y_t,c_t,n_t\}} \mathbb{E}_0 \sum_{t=0}^{\infty}\left[\Pi_{k=0}^{t-1}\delta_k\right]\beta^t\left\{\frac{c_t^{1-\sigma}-1}{1-\sigma} - \chi\frac{n_t^{1+\eta}}{1+\eta}\right\} \tag{46.1}$$

Subject to the attainable equilibrium

$$
\begin{aligned}
c_t^{-\sigma} R_t^{-1} &= \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} \pi_{t+1}^{-1} \right] && (46.2)\\
c_t^{-\sigma} y_t &= \frac{\phi \nu \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} y_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1} \right]}{1 - \nu + \phi \nu \pi_t (\pi_t - \pi_f) - \dfrac{w_t}{z_t}} && (46.3)\\
c_t &= y_t \left[1 - \frac{\phi}{2} (\pi_t - \pi_f)^2 \right] && (46.4)\\
R_t &\geq 1 && (46.5)\\
y_t &= z_t n_t && (46.6)\\
w_t &= \chi n_t^{\eta} c_t^{\sigma} && (46.7)
\end{aligned}
$$

Using the results in Marcet and Marimon (2019), we can re-write the previous problem as a min max problem. The first-order conditions of such a problem are:

$$
\begin{aligned}
\partial c \quad &\Longrightarrow \quad \left(1 - \lambda_t^3 c_t^{\sigma}\right) c_t = \sigma \left\{ \lambda_t^1 R_t^{-1} - \mu_t^1 \pi_t^{-1} \right. \\
&\qquad \left. + z_t n_t \left[\lambda_t^2 (1 - \nu) + \left(\lambda_t^2 - \mu_t^2 \right) \phi \nu \pi_t (\pi_t - \pi_f) \right] \right\} && (47.1)\\
\partial n \quad &\Longrightarrow \quad \chi n_t^{\eta} \left[1 + \lambda_t^2 (1 + \eta) \right] = \lambda_t^3 z_t \left[1 - \frac{\phi}{2} (\pi_t - \pi_f)^2 \right] \\
&\qquad + c_t^{-\sigma} z_t \left[\lambda_t^2 (1 - \nu) + \left(\lambda_t^2 - \mu_t^2 \right) \phi \nu \pi_t (\pi_t - \pi_f) \right] && (47.2)\\
\partial \pi \quad &\Longrightarrow \quad \lambda_t^3 z_t n_t \phi (\pi_t - \pi_f) = \mu_t^1 c_t^{-\sigma} \pi_t^{-2} \\
&\qquad + c_t^{-\sigma} z_t n_t \left(\lambda_t^2 - \mu_t^2 \right) \phi \nu (2\pi_t - \pi_f) && (47.3)\\
\partial R \quad &\Longrightarrow \quad \lambda_t^4 = \lambda_t^1 c_t^{-\sigma} R_t^{-2} && (47.4)\\
\partial \lambda^1 \quad &\Longrightarrow \quad c_t^{-\sigma} R_t^{-1} = \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} \pi_{t+1}^{-1} \right] && (47.5)\\
\partial \lambda^2 \quad &\Longrightarrow \quad c_t^{-\sigma} z_t n_t = \frac{\phi \nu \delta_t \beta \mathbb{E}_t \left[c_{t+1}^{-\sigma} z_{t+1} n_{t+1} (\pi_{t+1} - \pi_f) \pi_{t+1} \right]}{1 - \nu + \phi \nu \pi_t (\pi_t - \pi_f) - \dfrac{\chi n_t^{\eta} c_t^{\sigma}}{z_t}} && (47.6)\\
\partial \lambda^3 \quad &\Longrightarrow \quad c_t = z_t n_t \left[1 - \frac{\phi}{2} (\pi_t - \pi_f)^2 \right] && (47.7)\\
\partial \lambda^4 \quad &\Longrightarrow \quad \lambda_t^4 (R_t - 1) = 0 && (47.8)
\end{aligned}
$$

Where $\mu_t^i = \lambda_{t-1}^i$ and $\mu_0^i = 0$, $i = \{1, 2\}$, represents the burden of previous commitments. The preference shock does not appear in the conditions with respect to c_t, n_t, π_t and R_t: the forward-looking term of a constraint at $t-1$ is multiplied by $\delta_{t-1}\beta$, which is exactly the discount factor between $t-1$ and t, so it enters period t as $-\mu_t$ times the corresponding current term. In the non-stochastic steady state below $\delta_{ss} = 1$.Define x_{ss} to the variable x in the non-stochastic steady

state where the ZLB does not bind. With this, the previous system becomes:

$$\begin{aligned}
\left(1-\lambda_{\text{ss}}^{3}c_{\text{ss}}^{\sigma}\right)c_{\text{ss}} &= \sigma n_{\text{ss}}\lambda_{\text{ss}}^{2}\left(1-\nu\right) && (48.1)\\
\chi n_{\text{ss}}^{\eta}\left[1+\lambda_{\text{ss}}^{2}(1+\eta)\right] &= \lambda_{\text{ss}}^{3}\left[1-\frac{\phi}{2}\left(\pi_{\text{ss}}-\pi_{f}\right)^{2}\right]+c_{\text{ss}}^{-\sigma}\lambda_{\text{ss}}^{2}\left(1-\nu\right) && (48.2)\\
\lambda_{\text{ss}}^{3}n_{\text{ss}}\phi\left(\pi_{\text{ss}}-\pi_{f}\right) &= 0 && (48.3)\\
\chi n_{\text{ss}}^{\eta}c_{\text{ss}}^{\sigma} &= 1-\nu+(1-\beta)\phi\nu\pi_{\text{ss}}\left(\pi_{\text{ss}}-\pi_{f}\right) && (48.4)\\
c_{\text{ss}} &= n_{\text{ss}}\left[1-\frac{\phi}{2}\left(\pi_{\text{ss}}-\pi_{f}\right)^{2}\right] && (48.5)
\end{aligned}$$

Note that (48.3) implies that $\pi_{\text{ss}}=\pi_{f}$ or $\lambda_{\text{ss}}^{3}=0$. The latter is not possible since it would be feasible to increase the household's welfare by decreasing labor and increasing consumption until the resource constraint binds. Therefore it must be that $\pi_{\text{ss}}=\pi_{f}$, which pins down the rest of the variables.

B.3 Loose commitment

I assume a New Keynesian economy where the intermediate good firms face a quadratic cost of adjusting prices as in Rotemberg (1982b). The monetary authority chooses the interest rate in order to maximize the discounted utility of the household and it is constrained from below by the Zero Lower Bound (ZLB). The Central Bank chooses its policy under loose commitment, meaning that it chooses a plan at the beginning of time and is allowed to renege on it with probability $1-\xi$. Denote by x_{t}^{c} the allocation on the committed branch and by x_{t}^{r} the allocation on the reneged branch. Conditional on $s_{0}=\text{c}$, the problem at time zero can be resumed as follows:

$$\max_{\{R_{t}^{\text{c}},\pi_{t}^{\text{c}},w_{t}^{\text{c}},y_{t}^{\text{c}},c_{t}^{\text{c}},n_{t}^{\text{c}}\}}\mathbb{E}_{0}\sum_{t=0}^{\infty}\left[\Pi_{k=0}^{t-1}\delta_{k}\right](\beta\xi)^{t}\left\{\begin{array}{l}\dfrac{(c_{t}^{\text{c}})^{1-\sigma}-1}{1-\sigma}-\chi\dfrac{(n_{t}^{\text{c}})^{1+\eta}}{1+\eta}\\+\delta_{t}\beta(1-\xi)\mathbb{E}_{t}V_{t+1}^{\text{r}}\end{array}\right\} \qquad (49.1)$$

Subject to the attainable equilibrium

$$
\begin{aligned}
\left(c_t^{\mathrm{c}}\right)^{-\sigma}\left(R_t^{\mathrm{c}}\right)^{-1} &= \delta_t\beta\mathbb{E}_t\begin{bmatrix}\xi\left(c_{t+1}^{\mathrm{c}}\right)^{-\sigma}\left(\pi_{t+1}^{\mathrm{c}}\right)^{-1}\\ +\left(1-\xi\right)\left(c_{t+1}^{\mathrm{r}}\right)^{-\sigma}\left(\pi_{t+1}^{\mathrm{r}}\right)^{-1}\end{bmatrix} && (49.2)\\
\left(c_t^{\mathrm{c}}\right)^{-\sigma} y_t^{\mathrm{c}} &= \frac{\phi\nu\delta_t\beta\mathbb{E}_t\begin{bmatrix}\xi\left(c_{t+1}^{\mathrm{c}}\right)^{-\sigma} y_{t+1}^{\mathrm{c}}\left(\pi_{t+1}^{\mathrm{c}}-\pi_f\right)\pi_{t+1}^{\mathrm{c}}\\ +\left(1-\xi\right)\left(c_{t+1}^{\mathrm{r}}\right)^{-\sigma} y_{t+1}^{\mathrm{r}}\left(\pi_{t+1}^{\mathrm{r}}-\pi_f\right)\pi_{t+1}^{\mathrm{r}}\end{bmatrix}}{1-\nu+\phi\nu\pi_t^{\mathrm{c}}\left(\pi_t^{\mathrm{c}}-\pi_f\right)} && (49.3)\\
&\quad -\frac{w_t^{\mathrm{c}}}{z_t}\\
c_t^{\mathrm{c}} &= y_t^{\mathrm{c}}\left[1-\frac{\phi}{2}\left(\pi_t^{\mathrm{c}}-\pi_f\right)^2\right] && (49.4)\\
R_t^{\mathrm{c}} &\geq 1 && (49.5)\\
y_t^{\mathrm{c}} &= z_t n_t^{\mathrm{c}} && (49.6)\\
w_t^{\mathrm{c}} &= \chi\left(n_t^{\mathrm{c}}\right)^{\eta}\left(c_t^{\mathrm{c}}\right)^{\sigma} && (49.7)
\end{aligned}
$$

The two forward-looking constraints already reflect that, from today's perspective, tomorrow is the committed branch with probability ξ and the reneged branch with probability $1-\xi$ — this is equations (14) and (15) averaged across the two possible continuations. Following the same min-max reformulation used above (main text, equation (34)), and denoting $\mu_t^1 = \lambda_{t-1}^1$, $\mu_t^2 = \lambda_{t-1}^2$ the

burden of previous commitments (with $\mu_0^1 = \mu_0^2 = 0$), the first-order conditions are

$$\begin{aligned}
\partial c^{\mathrm{c}} \quad &\Longrightarrow \quad \left(1 - \lambda_t^3 \left(c_t^{\mathrm{c}}\right)^{\sigma}\right) c_t^{\mathrm{c}} = \sigma \Big\{ \lambda_t^1 \left(R_t^{\mathrm{c}}\right)^{-1} - \mu_t^1 \left(\pi_t^{\mathrm{c}}\right)^{-1} \\
&\qquad + z_t n_t^{\mathrm{c}} \left[\lambda_t^2 (1-\nu) + \left(\lambda_t^2 - \mu_t^2\right) \phi \nu \pi_t^{\mathrm{c}} \left(\pi_t^{\mathrm{c}} - \pi_f\right) \right] \Big\} \qquad (50.1) \\
\partial n^{\mathrm{c}} \quad &\Longrightarrow \quad \chi \left(n_t^{\mathrm{c}}\right)^{\eta} \left[1 + \lambda_t^2 (1+\eta)\right] = \lambda_t^3 z_t \left[1 - \frac{\phi}{2} \left(\pi_t^{\mathrm{c}} - \pi_f\right)^2\right] \\
&\qquad + \left(c_t^{\mathrm{c}}\right)^{-\sigma} z_t \left[\lambda_t^2 (1-\nu) + \left(\lambda_t^2 - \mu_t^2\right) \phi \nu \pi_t^{\mathrm{c}} \left(\pi_t^{\mathrm{c}} - \pi_f\right) \right] \qquad (50.2) \\
\partial \pi^{\mathrm{c}} \quad &\Longrightarrow \quad \lambda_t^3 z_t n_t^{\mathrm{c}} \phi \left(\pi_t^{\mathrm{c}} - \pi_f\right) = \mu_t^1 \left(c_t^{\mathrm{c}}\right)^{-\sigma} \left(\pi_t^{\mathrm{c}}\right)^{-2} \\
&\qquad + \left(c_t^{\mathrm{c}}\right)^{-\sigma} z_t n_t^{\mathrm{c}} \left(\lambda_t^2 - \mu_t^2\right) \phi \nu \left(2\pi_t^{\mathrm{c}} - \pi_f\right) \qquad (50.3) \\
\partial R^{\mathrm{c}} \quad &\Longrightarrow \quad \lambda_t^4 = \lambda_t^1 \left(c_t^{\mathrm{c}}\right)^{-\sigma} \left(R_t^{\mathrm{c}}\right)^{-2} \qquad (50.4) \\
\partial \lambda^1 \quad &\Longrightarrow \quad \left(c_t^{\mathrm{c}}\right)^{-\sigma} \left(R_t^{\mathrm{c}}\right)^{-1} = \delta_t \beta \mathbb{E}_t \begin{bmatrix} \xi \left(c_{t+1}^{\mathrm{c}}\right)^{-\sigma} \left(\pi_{t+1}^{\mathrm{c}}\right)^{-1} \\ + (1-\xi) \left(c_{t+1}^{\mathrm{r}}\right)^{-\sigma} \left(\pi_{t+1}^{\mathrm{r}}\right)^{-1} \end{bmatrix} \qquad (50.5) \\
\partial \lambda^2 \quad &\Longrightarrow \quad \left(c_t^{\mathrm{c}}\right)^{-\sigma} z_t n_t^{\mathrm{c}} = \frac{\phi \nu \delta_t \beta \mathbb{E}_t \begin{bmatrix} \xi \left(c_{t+1}^{\mathrm{c}}\right)^{-\sigma} z_{t+1} n_{t+1}^{\mathrm{c}} \left(\pi_{t+1}^{\mathrm{c}} - \pi_f\right) \pi_{t+1}^{\mathrm{c}} \\ + (1-\xi) \left(c_{t+1}^{\mathrm{r}}\right)^{-\sigma} z_{t+1} n_{t+1}^{\mathrm{r}} \left(\pi_{t+1}^{\mathrm{r}} - \pi_f\right) \pi_{t+1}^{\mathrm{r}} \end{bmatrix}}{1 - \nu + \phi \nu \pi_t^{\mathrm{c}} \left(\pi_t^{\mathrm{c}} - \pi_f\right) - \dfrac{\chi \left(n_t^{\mathrm{c}}\right)^{\eta} \left(c_t^{\mathrm{c}}\right)^{\sigma}}{z_t}} \qquad (50.6) \\
\partial \lambda^3 \quad &\Longrightarrow \quad c_t^{\mathrm{c}} = z_t n_t^{\mathrm{c}} \left[1 - \frac{\phi}{2} \left(\pi_t^{\mathrm{c}} - \pi_f\right)^2\right] \qquad (50.7) \\
\partial \lambda^4 \quad &\Longrightarrow \quad \lambda_t^4 \left(R_t^{\mathrm{c}} - 1\right) = 0 \qquad (50.8)
\end{aligned}$$

Equations (50.5) and (50.6) are the loose-commitment Euler equation and Phillips curve: each is simply the natural expectation of the corresponding constraint across the continuing (ξ) and reneging ($1-\xi$) branches. Every other first-order condition above is identical in form to the fully-committed case, equations (47.4)–(47.8) above (with the "c" superscript added), since the explicit reneging term is dated at $t+1$ and therefore has zero derivative with respect to the period-t controls $c_t^{\mathrm{c}}, n_t^{\mathrm{c}}, \pi_t^{\mathrm{c}}, R_t^{\mathrm{c}}$; the multiplier updating rule $\mu_t^i = \lambda_{t-1}^i$ is likewise unchanged, since the possibility of reneging is captured entirely by the two-branch average in $\partial\lambda^1$ and $\partial\lambda^2$ above, not by the multiplier's own law of motion. Setting $\xi = 1$ collapses (50.5)–(50.6) to (47.5)–(47.6); setting $\xi = 0$ removes the forward-looking piece of the plan entirely, consistent with the discretionary problem of Appendix D.The imperfect commitment is created by a stochastic variable driving the economy to a reset, therefore there is no clear definition for what the deterministic steady state should be. In this paper I use the weighted average between commitment and discretion, described above, to initialize the solution finding algorithm.Define the variable X_{SS}^{c} and X_{SS}^{r} as the values of X in the deterministic steady state when the economy is in the committed world and the discretionary world, respectively.

Then the "steady state" of X when the level of commitment is ξ is defined as

$$X_{SS}^{\xi} = \xi X_{SS}^{c} + (1-\xi) X_{SS}^{r} \tag{51}$$

C Solution method for loose commitment

In this section, I describe the method used to solve the limited commitment model presented in the paper. The main challenge in solving it lies in an ergodic distribution broader than pure commitment or discretion. When the monetary authority is discretionary, the only state variable is the total factor productivity process. When the policy is under commitment, the endogenous state– the previous value of the Phillips and Euler equations Lagrange multipliers— begins at the origin but it will never return to it. This implies that the ergodic distribution will be dense in an area far from the origin. When the monetary policy is highly but not perfectly committed, the highest probability mass will be concentrated at the origin and the path to a higher commitment spells.In the universe of solution procedures, I stick with a "global method" since approximations around an arbitrary point might introduce numeric error biases due to the existence of the Zero Lower Bound kink and the presence of inflation bias. This decision is also induced by the fact that there is no steady state, and the probability of being at the origin in the endogenous state —another candidate to linearize around— is not the same for different levels of commitment. Work in the literature has been done linearizing the model around the zero-inflation steady state and leaving the ZLB non-linear restriction. Their results accuracy depends on how far the zero lower bound is from the zero-inflation state. In this paper we show results with inflation bias that depends on different commitment levels and compare the welfare loss induced by it, therefore this technique would not be appropriate.The main feature in the solution method is the time iteration proposed by Coleman (1991). I rely on an algorithm similar to the one in Maliar and Maliar (2015) to find an approximation of the ergodic distribution, and form a non-regular and non-squared set of nodes where I solve the model. Finally, I use Delaunay triangulation —Matlab built-in routines— to get spline interpolation on the irregular grid.

Policy iteration algorithm

The objective of this algorithm is to find a solution for the model in a given set of scattered state nodes. I leave the discussion on how to select those nodes for the next subsection.Define the scalar PI_u as an update dampener, and PI_{tol} as a convergence tolerance level. The algorithm starts with an initial guess of the policies $x_j(z,\mu)$ with $j=0$ and iterates over j as follows[20]

1. For each node in the grid, compute the expectation of the Euler and Phillips equations using the state evolution and policy function guess.

[20]Note that the policies are invariant over time, therefore the subscript j in x_j denotes an iteration in the solution finding algorithm

2. Assume the Zero Lower Bound does not bind, and find the roots of the first-order conditions. Call them $x_j^*(z,\mu)$

3. If the interest rate found is lower than zero, solve the system assuming the ZLB binds, if not proceed to the next step

4. Compute the percent difference between x_j and x_j^*. If the difference is lower than the tolerance level PI_{tol}, then stop. Otherwise, proceed to the next step.

5. Update the policy function guess using $x_{j+1}(z,\mu) = \lambda_u x_j^*(z,\mu) + (1-\lambda_u)x_j(z,\mu)$ and go back to step 1.

Granted the convergence of the algorithm, the numerical quality of the solution found on each node depends on how precise the estimated expectations are. This estimation uses the evolution of the state and a linear interpolation between nodes.Given the policy function guesses, we can compute the next period endogenous state as $\mu_{t+1} = \lambda_t$ with probability ξ, and $\mu_{t+1} = 0$ with probability $(1-\xi)$. This is independent of the future total factor productivity or time preferences shock, which will evolve according to (19) and (20). I use a Gauss-Hermite quadrature with N_ε nodes on ε, which will lead to N_ε possible values of z_{t+1} and its correspondent probability weights. This procedure leads to an exact solution of $\mathbb{E}(h(x))$ when $h(x)$ is a polynomial of degree $2N_\varepsilon - 1$.Since the next period state is most likely to be out of the scattered nodes, I need to interpolate the policies in those states. To do that I use the Matlab routine *scatteredInterpolant* which uses the Delaunay triangulation to find the interpolated function value in a scattered set of nodes.[21] In particular, I use the linear interpolation and extrapolation methods.A computational solution of the model is a set of functions that receive as input any value of the model's state (z_t or δ_t, μ_t^1, and μ_t^2) and returns the —potentially interpolated— policies.

Ergodic distribution

The policy iteration algorithm described above tends to break when a dense squared state grid is provided. To avoid including points that will never be part of the solution —which happens with a dense regular grid— I select nodes using their probability mass. Nonetheless, the distribution of the endogenous state is apriori unknown. In this section, I describe how to get a discrete approximation.Given a computational solution, the first step in the Ergodic distribution algorithm is to select minimum and maximum values for the endogenous state. The minimum value is set in the origin, which is where the monetary authority goes every time it is allowed to. The maximum value for μ_1 is related to the probability of reaching the zero lower bound, while the maximum value of μ_2 is to the probability of a long commitment spell. I perform deterministic simulations using the solution of the model, holding the exogenous shock at the bottom, the mean, and the top of its grid during a commitment spell of length L_E. The spell length L_E is set to match the probability

[21] More on irregularly spaced data interpolation methods Amidror (2002)

of reaching higher commitment spells to 0.001%. The maximum values for μ are selected from this simulation.With the minimum and maximum in hand, the state is divided into a homogeneous grid with $NE_z \times NE_{\mu_1} \times NE_{\mu_2}$ nodes. Noting that each commitment spell starts at the origin with probability one and that the long-run probabilities of the total factor productivity process are known, each node's probability mass is assigned in the following way:

1. Generate a vector with the probability of reaching each commitment history $p_t = (1-\xi)\xi^t$ with $t \in \{0, \cdots, L_E\}$
2. Set t = 0 and generate a $(NE_{\mu_1} \times NE_{\mu_2})$ matrix and call it P_t. Set the first element in this matrix to one, and the rest to zero ($P_0(1,1) = 1$ and $P_0(i,j) = 0 \forall (i,j) \neq (1,1)$). P_0 is the distribution of the endogenous state conditional on being in a zero commitment spell.
3. For each non-zero elements $P_t(i,j)$ of P_t (indexed by h) and its associated endogenous grid point:
   (a) For each node in the exogenous state, compute the future endogenous state assuming the commitment history will continue.
   (b) Distribute $P_t(i,j)$ among states using the ergodic distribution of the exogenous process. To relocate them in the endogenous grid, divide the probability mass in the four closest nodes using the distance as weight. Denote this relocation as $\tilde{P}_h$
4. Set $P_{t+1} = \sum_h \tilde{P}_h$
5. If $P_{t+1} = P_t$, then set $p_t = \sum_{j=t}^{L_E} p_j$ and break the algorithm. Else go back to step 3.

Initial guess and coarse solution

The initial guess used might delay the solution and ergodic state-finding algorithms. To give them a better start I use a coarse and squared endogenous and exogenous grid to get a sense of where the ergodic distribution might be.To start things out I compute the steady state of the commitment and discretion model and set the initial policy guess for each commitment level as the weighted combination of both. Then, the following iteration is performed

1. Given the initial guess and coarse state grid, compute the solution of the model.
2. Find the ergodic distribution using the algorithm above and compare the maximum values of the grid
3. If the values are similar, break. If not, replace the new 'max' in the grid with the ones found in the ergodic distribution

The big loop

The final solution is computed by iterating at most N times the following

1. Find an initial guess and coarse solution
2. Compute the ergodic distribution
3. Select the nodes with probability higher than p_l and set them as new endogenous nodes.
4. If the number of exogenous nodes is lower than desired, increase it.
5. Find the solution of the model in these new nodes.
6. If the solution is immediately found, break. Otherwise, go back to 2.

Numerical accuracy

Because the main results compare economies along the whole commitment line, the approximation error must be small and, just as importantly, of similar size at every level of commitment, so that differences across ξ are not artifacts of the solution method. I evaluate the accuracy of the solutions with unit-free residuals of the equilibrium and optimality conditions, in the spirit of Judd (1992), computed away from the solution grid. For each level of commitment I simulate 100 histories of 500 quarters with the computed policy functions. Every history starts at a plan reset and the plan continues with probability ξ each period, so the sample covers short and long commitment spells in their ergodic proportions. At each of the 50,000 simulated states I take all variables from the interpolated solution —inflation, consumption, labor, the nominal interest rate and the Lagrange multipliers— and nothing is re-solved, so the residuals measure the error of the approximated policy functions themselves. I report the residual of the Euler equation, $1 - \beta\delta_t R_t c_t^{\sigma} \mathbb{E}_t \left[c_{t+1}^{-\sigma} / \pi_{t+1} \right]$, the residual of the non-linear Phillips curve and the residuals of the three first-order conditions of the policy problem, each divided by the sum of the absolute value of its terms. Expectations are computed with 20 Gauss–Hermite nodes, instead of the 6 used by the solver, so the residuals also capture the integration error.Table 3 reports the results for selected levels of commitment. The average Euler-equation error lies between $10^{-4.7}$ and $10^{-4.0}$ across the commitment line, and its 95th percentile does not exceed $10^{-3.5}$; the Phillips curve is slightly more accurate, and the first-order conditions of the policy problem hold to $10^{-5.9}$ or better in 95% of the states. Errors are of similar size at and away from the ZLB, except close to discretion, where the Euler residual at the ZLB is about half an order of magnitude larger; the kink in the policy functions does not otherwise concentrate the approximation error. These magnitudes are in the range usually reported in the global-solution literature (Judd, 1992; Maliar and Maliar, 2015), and they are at least two orders of magnitude smaller than the welfare losses reported in the main text, which range between 0.3% and 1.4% of consumption. Errors also do not increase around the commitment levels where the probability of hitting the ZLB peaks, so the non-monotonic relation between commitment and the

ZLB is not driven by approximation error.Table 4 documents the grids and the convergence of each solution. Most of the integration error comes from the quadrature: evaluated with the solver's own 6 nodes, the average Euler residual falls to between $10^{-6.7}$ and $10^{-5.9}$, while moving from 6 to 20 nodes changes the Euler-equation expectation by at most 0.12% across all commitment levels. The adaptive grid covers the simulated states well: less than 0.5% of them fall outside the convex hull of the grid, where the policies are extrapolated, and the multiplier on the Euler equation has the right sign in all but 0.002% of the ZLB periods.

Table 3: Accuracy of the loose-commitment solutions

		Euler		Phillips		Planner FOCs		Euler, mean	
ξ (%)	Pr(ZLB) (%)	Mean	P95	Mean	P95	Mean	P95	ZLB	No ZLB
1	9.3	−4.37	−3.74	−4.66	−4.04	−7.79	−6.22	−3.79	−4.43
10	22.1	−4.06	−3.65	−4.40	−4.02	−7.57	−6.09	−4.13	−4.04
30	43.9	−4.24	−3.85	−4.65	−4.24	−7.07	−6.11	−4.33	−4.17
50	50.9	−4.47	−3.96	−4.96	−4.44	−6.77	−5.95	−4.42	−4.52
65	49.7	−4.62	−4.03	−5.16	−4.59	−6.66	−5.92	−4.57	−4.66
75	47.2	−4.72	−4.03	−5.29	−4.64	−6.68	−5.90	−4.58	−4.84
90	41.9	−4.49	−3.91	−5.13	−4.56	−6.69	−5.94	−4.53	−4.47
99	35.5	−4.19	−3.71	−4.85	−4.41	−6.68	−5.91	−4.21	−4.17

Note: Unit-free residuals in $\log_{10}$ units (e.g. -4 is an error of 0.01%), evaluated at 50,000 states simulated with the computed policy functions (100 histories of 500 periods, each starting at a plan reset, fixed seed). All variables, including the interest rate and the Lagrange multipliers, are taken from the interpolated solution, and expectations are computed with 20 Gauss–Hermite nodes instead of the 6 used by the solver. *Euler* is $1 - \beta\delta_t R_t c_t^{\sigma} \mathbb{E}_t[c_{t+1}^{-\sigma}/\pi_{t+1}]$; *Phillips* is the residual of the non-linear Phillips curve divided by the sum of the absolute value of its terms; *Planner FOCs* is the largest of the three first-order conditions of the policy problem, normalized in the same way. *Mean* is the average and *P95* the 95th percentile of the $\log_{10}$ absolute residual. The last two columns split the Euler residual between periods at and away from the ZLB.

D Solution Method for Discretion

In this section, I describe the method used to solve the discretionary model presented above. The discretionary model has only one exogenous state variable δ_t. The algorithm is standard in the literature: time iteration as in Coleman (1991).The method starts by discretizing the state. Since the ergodic distribution of the exogenous state is known, the task of finding grid points is trivial. I add N nodes between plus and minus 4 standard deviations from the mean, using the distribution to space them apart. A linear spline is used between nodes.Since one of the most relevant points in the policy functions is where the kink where the zero lower bound starts binding, I add a bisection search algorithm to look for this point and add a fixed number of nodes between the states identified the ZLB would live within.

Table 4: Grids, quadrature and convergence of the loose-commitment solutions

ξ (%)	Nodes δ	Nodes total	Quadrature nodes	Tolerance ($\log_{10}$)	Outside grid (%)	Δ Euler exp. 6 vs 20 (%)	Longest spell
1	51	2,091	6	−6	0.47	0.03	2
10	51	4,947	6	−6	0.03	0.03	4
30	51	8,058	6	−6	0.05	0.03	10
50	51	13,158	6	−6	0.00	0.05	13
65	51	19,125	6	−5	0.01	0.03	20
75	51	24,684	6	−5	0.00	0.03	28
90	51	60,945	6	−5	0.00	0.02	76
99	51	24,582	6	−6	0.00	0.03	499

Note: *Nodes* δ is the number of nodes of the preference shock and *Nodes total* the number of points of the final (irregular) state grid. *Tolerance* is the convergence criterion on the policy functions in the last iteration (maximum percent change). *Outside grid* is the share of simulated states outside the convex hull of the grid, where the policies are extrapolated. Δ *Euler exp.* is the largest percent change in the Euler-equation expectation when moving from 6 to 20 Gauss–Hermite nodes. *Longest spell* is the longest commitment spell in the simulation, in quarters.