Tilburg University

Preference Heterogeneity and Optimal Monetary Policy

Uras, Burak; van Buggenum, Hugo

PREFERENCE HETEROGENEITY AND OPTIMAL MONETARY POLICY

By

Burak R. Uras, Hugo van Buggenum

TILBURG ◆ UNIVERSITY

Preference Heterogeneity and Optimal Monetary Policy[*]

Burak R. Uras[†]　　　Hugo van Buggenum[‡]

Tilburg University　　　Tilburg University

October 20, 2020

Abstract

We study optimal policy design in a monetary model with heterogeneous preferences. In the model markets are incomplete and households are heterogeneous with respect to their current consumption preferences and discount factors. The government controls the supply of money (liquid) and nominal bonds (illiquid) based on which households make optimal portfolio choices. We uncover that the two types of preference heterogeneity have distinct implications for the optimal policy mix of the government. While the heterogeneity in current consumption preferences pushes the economy towards a zero-lower-bound (ZLB) associated with nominal interest rates, the heterogeneity in discount factors moves the economy away from the ZLB. We characterize the optimal policy design and quantify the welfare losses associated a binding ZLB - and thus potential welfare benefits from allowing negative interest rates on government bonds.

Keywords: Heterogeneous Consumption Preferences, Optimal Policy, Zero-Lower-Bound, Negative Interest Rates.
JEL Classification: E22, E41, E44, E52, O16.

1 Introduction

An extensive quantitative-macro literature studies the role of cross-sectional heterogeneity in explaining consumption and wealth inequality.[1] A branch of this literature highlights the potential relevance of heterogeneous consumption preferences (in the form of discount factor shocks) in accounting for inequality data (Krusell and Smith, 1998; Venti and Wise, 2000; Cagetti, 2003; Gelman, 2020). Preference heterogeneity may exacerbate inequality and cause welfare losses when financial markets are incomplete (Krusell and Smith, 1998; Hendricks, 2007) or when fiscal policy is set sub-optimally (Golosov et al., 2013).

In this paper, we approach the interaction between preference heterogeneity and welfare from a novel perspective: to what extent can monetary policy correct the allocation problem borne by heterogeneous preferences and improve welfare under incomplete markets? Our framework takes a general stance on modeling life-cycle consumption preferences to address this question. Specifically, in our model we separately incorporate intra-generational heterogeneity with respect to current marginal utility to consume and discount factors shocks into a monetary life-cycle model. Based on this framework, we argue that optimal monetary policy needs to influence both the price of liquid (money, i.e. inflation) and illiquid (bonds, i.e. nominal interest rates) assets to achieve an optimal cross-sectional allocation of consumption and wealth over the life-cycle. We also characterize the conditions for an endogenous zero lower bound (ZLB) on nominal interest rates to bind. We find that the ZLB limits optimal policy and results in welfare losses.

Our benchmark framework is a dynamic general equilibrium model of a monetary economy in spirit of Azariadis (1981) with three-period lived households who are connected over time through overlapping generations. At the beginning of the life-cycle, households produce and determine asset portfolios by taking expected preference shocks into account. During the remainder of their lifetime, the households consume and re-balance liquid and illiquid assets by transacting in a financial market. We define the notion of an asset's liquidity based on the exchangeability of that asset in return for consumption goods in the current period. As standard in monetary models, we assume that a special nominal object, called money, is the sole liquid asset available to the economy. Illiquid nominal assets, called bonds, serve as a saving instrument between periods. Important for our analysis is that the supplies of money and bonds are determined by the government.

At the core of our theoretical design are idiosyncratic preference shocks to which households are exposed in the second (early) and third (late) period of their life. We take a general approach in modeling this heterogeneity: there is a preference distribution to consume early as well as late, and individual preferences get realized at the beginning of the second period of

---

[1]See De Nardi and Fella (2017) for an extensive survey of the literature.

households' lifetime. Early consumption preference (current marginal utility to consume) of a household is not necessarily correlated with its preference to consume late (the discount factor). Moreover, preference shocks are uninsurable due to financial imperfections. Given this market incompleteness, the heterogeneity of preferences gives rise to an environment in which over- or under-supply of (il)liquidity can prevail when inflation and nominal interest rates follow the Friedman rule - equating the real return on all assets to the natural rate of interest. The key theoretical analysis of our paper concentrates on exploring optimal policy design that takes the distribution of consumption preferences as given.

As a primarily important result, we characterize a closed-form test statistic resulting from the distribution of preferences. We show that when the conditions governed by this statistic are met, an endogenous ZLB on nominal interest rates gets reached. This ZLB arises because besides bonds, households can use money as a savings instrument. Our test statistic has a simple economic interpretation, as it characterizes the average fraction of disposable resources that middle-aged households consume. The properties that we obtain with regard to this test statistic are straightforward. On the one hand, when households are relatively heterogeneous in terms of their preference to consume early (current MU) - and not that heterogeneous with respect to late consumption preferences (discount factors) - the test statistic is high and the economy reaches the ZLB. On the other hand, as heterogeneity in discount factors increases - while early consumption heterogeneity declines - the test statistic decreases and the economy moves away from the ZLB. In this respect, our analysis reveals that increasing heterogeneity of *impatience* (current MU) and *patience* (discount factors) have different consequences for government's ability to freely choose an optimal nominal interest rate in its policy mix.

We then study optimal policy design with respect to inflation and interest rates, and obtain an important property of the ZLB. That is, with a binding ZLB constraint a deviation from the Friedman rule becomes an optimal monetary policy choice. The intuition for this property is the following. When early consumption preferences are disperse (relative to the dispersion observed in discount factors) those agents with low MU to consume early like to increase savings which pushes the aggregate savings to an inefficiently high level. In order to reduce this saving inefficiency and stimulate early consumption of those with high MU, the government ideally reduces nominal interest rates and inflation. However, reductions in nominal rates are infeasible beyond the ZLB, at which point increasing the inflation becomes the optimal policy. The reason for this important property relates to compounded costs of inflation, which impacts late consumption more than the early consumption. On the flip-side of the same channel, when discount factor shocks are relatively disperse (compared to the heterogeneity in early consumption preferences), the aggregate savings rate tends to contract. In order to stimulate late consumption and savings of the middle-aged, the optimal policy increases both the nominal interest rate and the rate of inflation - moving the economy away from the ZLB.

3

Based on our theoretical insights, we conduct a series of quantitative macro experiments using a parameterized version of the model. A counterfactual analysis reveals that, on the one hand, when the relative standard deviation of first-best early consumption equals one (as a benchmark case), the optimal policy increases steady state welfare by 0.2% of the first-best level of welfare. On the other hand, when the relative standard deviation of first-best late consumption equals one (as another benchmark), the optimal policy increases steady state welfare by 1.2% of the first-best level of welfare. This quantitatively substantial difference is caused by the fact that the ZLB binds in the former case but not in the latter. Removing the ZLB and allowing for negative interest rates raises aggregate welfare by the 1.4% of first-best welfare when the relative standard deviation of first-best early consumption equals one.

As an important extension to the benchmark, we allow for private capital formation by households via a decreasing returns to scale investment technology. In the extended framework, we allow claims on capital to be traded as an illiquid asset to compete with government's nominal bonds. This extension reveals that endogenous capital formation does not change the theoretical property associated with the threshold to reach a ZLB. A high spread in early consumption shocks causes the ZLB to bind, also with endogenous capital formation, at which the government remains with only inflation as a policy instrument - leading to distorted allocations. A high spread in discount factors moves the economy away from the ZLB. The main difference compared to the benchmark model is that with private capital - even away from ZLB - optimal policies are closer to the Friedman rule, as the latter achieves first-best levels of capital formation.

Our results imply that understanding the mode of preference heterogeneity is important for optimal policy design. The findings reveal that uninsurable early consumption heterogeneity is likely to push an economy to the ZLB and constrain the capacity of monetary policy to maximize welfare. We also suggest the optimality of using money and bond pricing policies of the government in tandem and thus propose a new micro-foundation for the co-existence of money and bonds. Furthermore, since bond pricing policy becomes ineffective at ZLB, we argue the need to allow for negative interest rates.

Our paper speaks to several strands of literature. First, we contribute to the literature that models preference heterogeneity in quantitative macro models. Starting with the seminal work by Krusell and Smith (1998), there has been an interest to understand the role of discount factor shocks in explaining cross-sectional distributions of savings and wealth (Venti and Wise, 2000; Cagetti, 2003; Gelman, 2020) as well as asset prices (Garleanu and Panageas, 2015). Most papers in this line of research document that - especially in the presence of incomplete markets - a small dispersion in discount factor heterogeneity across households, can be key to explain the data. Our contribution to this literature is three-fold. (i) We develop a tractable life-cycle model that incorporates a general form of preference heterogeneity into a monetary framework.

(ii) Our analysis reveals that the implications of early and late consumption shocks are quite different from each other with respect to the welfare effects of optimally conducted monetary policy. (iii) Finally, we uncover the effects of preference heterogeneity on the ZLB and discuss optimality of negative interest rates.

Second, we also contribute to the recent monetary economics literature on general equilibrium models with preference heterogeneity. For instance, Curran and Dressler (2019) study the welfare implications of inflation in a model with heterogeneous coefficients of relative risk aversion. In two closely related papers, Boel and Camera (2006) and Boel and Waller (2015) embed heterogeneous discount factors in New Monetarist style models and - similar to our conclusions - argue for non-optimality of the Friedman rule. However, there are two essential differences between our paper and these two studies: (i) we take a more generalized approach in modeling preference heterogeneity (with respect to incorporating heterogeneity in both current consumption preferences and discount factors), and (ii) our approach allows to uncover the differential impact of the two different margins of preference shocks on co-existence of money and bonds, the ZLB, and optimality of negative interest rates.

Third we also contribute to a literature that studies optimal coexistence of money and bonds. Kocherlakota (2003) shows that in an infinite horizon economy, society can benefit from having both money and bonds when agents face shocks to marginal utility of current consumption. The mechanism identified in Kocherlakota is however transitory; welfare effects of introducing bonds last for one period only. Andolfatto (2011) demonstrates that the transitory mechanism of Kocherlakota (2003), persists in steady state when implementing the Friedman rule is infeasible. Van Buggenum (2020) considers an infinite horizon economy in which agents experience shocks to their subjective discount factors and shows that even when the Friedman rule is implementable, interest bearing bonds are essential to achieve an efficient distribution of savings. In an OLG framework with shocks to both the marginal utility from current consumption and subjective discount factors, the current paper demonstrates that bonds are essential when current consumption shocks dominate in characterizing preference heterogeneity and analyzes the welfare implications of a ZLB on bond prices.

Finally, several recent papers in the monetary macro literature have argued for the optimality of negative interest rates, such as Dong and Wen (2017) and Williamson (2019) - following the trends in interest rates in some European countries. Our paper is to first in this literature that uncovers a mechanism related to heterogeneity in preferences in justifying the conduct of optimal negative interest rate policies.

2 Benchmark Model

We study an overlapping generations model with three-period lived households. We refer to the households as *young* in the first period of their lifetime, and as *middle-aged* and *old* in the second and third period, respectively. There are three types of objects in the economy: a non-storable *consumption good*, and two nominal assets that we call as *money* and *bonds*. The supply of money and bonds are determined by the government. The consumption good is produced by the young, and consumed by the middle-aged and old.

During each time period, two markets open sequentially in two distinct sub-periods that allow trade of assets and consumption goods: the *financial market* (FM) and the *goods market* (GM). The consumption good is produced and sold in the GM by the young and it fully perishes within a time period. Money is the only recognizable object in the GM, so bonds cannot be traded in this sub-period. The unit-price of the consumption good in the GM is p_t - denoted in terms of money.

In the FM, money and bonds are supplied by the government and the two objects are traded among households. Money is storable across time, while bonds issued in the FM of period t mature in period $t+1$ and pay out money. The unit price of bonds is φ_t - denoted in terms of fiat money. In the FM, the government also levies lump-sum taxes on middle-aged households (or provides the middle-aged with lump-sum subsidies), which are paid with money. There is no enforcement problem with respect to the payment of taxes. The real consumption good equivalent of the lump-sum tax is denoted with τ_t.

Lifetime preferences of a household are given by

$$\mathcal{U}_\xi(h, q_\xi^m, q_\xi^o) = -h + \beta \varepsilon u(q_\xi^m) + \beta^2 \delta u(q_\xi^o). \tag{1}$$

In preference specification (1), h denotes production of goods when young, q_ξ^m denotes consumption of goods when middle-aged, i.e. *early consumption*, and q_ξ^o denotes consumption when old, i.e. *late consumption*. Parameter β is a component of the discount factor that is common to all households. Important for our model, are idiosyncratic preference shocks that households are subject to. Specifically, at the beginning of the second period of their lifetime, each middle-aged household draws a tuple $\xi = (\varepsilon, \delta)$ from a distribution with CDF $G(\varepsilon, \delta)$ and support $\Xi \in \mathbb{R}_+^2$. In this joint preference distribution, ε is an idiosyncratic preference shifter for early consumption and δ is an idiosyncratic preference shifter for late consumption which get realized simultaneously. We remain agnostic about G, which means ε could exhibit negative or positive correlation with δ. Without loss of generality, we make the following assumption.

Assumption 1. $\mathbb{E}[\varepsilon] = \mathbb{E}[\delta] = 1$.

The preference shocks serve to generate intra-generational heterogeneity across households,

the equilibrium implications of which constitute the core analysis of this paper.

2.1 Equilibrium Analysis

We denote gross inflation in between periods $t-1$ and t with $\pi_t = p_t/p_{t-1}$, which in steady state equals the growth rate of money supply. Because money is a store of value, and the nominal bonds pay out money, households do not hold bonds when the interest rate on bonds is negative ($\varphi > 1$). Hence, the effective gross real interest rate at which middle-aged households can save, is $1/[\min\{\varphi_{t+1}, 1\}\pi_{t+2}]$. Lifetime utility maximization of a household born in period t therefore implies

$$\mathcal{U}_t = \max_{h_t, \{q_{\xi,t}^m, q_{\xi,t}^o\}_{\xi \in \Xi}} \left\{ -h_t + \beta \mathbb{E}[\varepsilon u(q_{\xi,t}^m)] + \beta^2 \mathbb{E}[\delta u(q_{\xi,t}^o)] \right\} \tag{2a}$$

$$\text{s.t.} \quad \pi_{t+1}[q_{\xi,t}^m + \min\{\varphi_{t+1}, 1\}\pi_{t+2}q_{\xi,t}^o + \tau_{t+1}] \leq h_t, \quad \forall \, \xi \in \Xi. \tag{2b}$$

where (2a) is the expected life-time utility and (2b) is a state-contingent resource constraint. We define the concept of the equilibrium as follows.

Definition 1. *The equilibrium is defined by M_m and M_o units of money that initial middle-aged and old are endowed with and $\{h_t, \{q_{\xi,t}^m, q_{\xi,t}^o\}_{\xi \in \Xi}\}_{t=0}^\infty$ at which markets clear and households maximize life-time preferences* (2a) *subject to the state-contingent budget constraint* (2b) *by taking good prices, $\{p_t\}_{t=0}^\infty$, and the policy choice variables of the government, $\{\pi_t, \varphi_t, \tau_t\}_{t=0}^\infty$, as given.*

To obtain closed-form solutions, we assume logarithmic utility, such that the period utility function is given by $u(q) = \log(q)$ for both middle-age and old households. Based on this functional form, we first observe that labor supply by the young satisfies

$$h_t = \beta(1+\beta) + \pi_{t+1}\tau_{t+1},$$

where $\beta(1+\beta)$ hours of work are used to provide resources for early and late consumption, and $\pi_{t+1}\tau_{t+1}$ hours are provided to pay lump-sum taxes when middle-aged.

Furthermore, early and late consumption are given by:

$$q_{\xi,t}^m = \frac{\varepsilon}{\varepsilon + \beta\delta} \frac{\beta(1+\beta)}{\pi_{t+1}} \quad \text{and} \quad q_{\xi,t}^o = \frac{\beta\delta}{\varepsilon + \beta\delta} \frac{\beta(1+\beta)}{\min\{\varphi_{t+1}, 1\}\pi_{t+1}\pi_{t+2}}, \quad \text{for all } \xi \in \Xi. \tag{3}$$

The levels of early and late consumption at (3) are negatively related to the level of inflation. Late consumption is negatively related to the price of bonds and is more sensitive to inflation than early consumption, because of compounding.

As an important property of logarithmic utility, we obtain that only preference shifters and the common discount factor determine middle-aged households' desire to save. Specifically,

after taxation middle-aged households have $\beta(1+\beta)/\pi_{t+1}$ units of real money balances available in the FM. They consume a fraction $\theta = \varepsilon/(\varepsilon + \beta\delta)$ of this sum in the GM, and save the remaining fraction $1 - \theta$ for consumption in the next period, earning a real rate of return $1/[\min\{\varphi_{t+1}, 1\}\pi_{t+1}]$. We therefore define as a test statistic:

$$\Theta \equiv \iint \frac{\varepsilon}{\varepsilon + \beta\delta} \mathrm{d}G(\varepsilon, \delta). \tag{4}$$

The test-statistic defined at (4) is at the core of our policy analysis and measures the aggregate fraction of middle-aged households' disposable wealth devoted to early consumption. Analogously, $1 - \Theta$ captures the aggregate savings rate of middle-aged households.

Aggregate demand for money M_t and bonds B_t in FM of period t is given by:

$$\frac{M_t}{p_t} = \begin{cases} \beta(1+\beta)\left(\frac{\Theta}{\pi_t} + \frac{1-\Theta}{\min\{\varphi_{t-1},1\}\pi_{t-1}\pi_t}\right) & \text{if } \varphi_t < 1 \\ \in \left[\beta(1+\beta)\left(\frac{\Theta}{\pi_t} + \frac{1-\Theta}{\min\{\varphi_{t-1},1\}\pi_{t-1}\pi_t}\right), \beta(1+\beta)\left(\frac{1}{\pi_t} + \frac{1-\Theta}{\min\{\varphi_{t-1},1\}\pi_{t-1}\pi_t}\right)\right] & \text{if } \varphi_t = 1 \\ \beta(1+\beta)\left(\frac{1}{\pi_t} + \frac{1-\Theta}{\min\{\varphi_{t-1},1\}\pi_{t-1}\pi_t}\right) & \text{if } \varphi_t < 1 \end{cases} \tag{5}$$

$$\frac{\varphi_t B_t}{p_t} = \begin{cases} \beta(1+\beta)\left(\frac{1}{\pi_t} + \frac{1-\Theta}{\min\{\varphi_{t-1},1\}\pi_{t-1}\pi_t}\right) - \frac{M_t}{p_t} & \text{if } \varphi_t \leq 1 \\ 0 & \text{if } \varphi_t > 1 \end{cases}. \tag{6}$$

With a strictly positive supply of bonds, market clearing requires that $\varphi_t \leq 1$. An endogenous zero-lower-bound (ZLB) on nominal interest rates therefore arises. In what follows, without loss we assume $\varphi_t \leq 1$.

We close the equilibrium analysis, by solving for the lump-sum taxes and characterize the labor supply of the young as a function of inflation, bond prices, and preference parameters. For this purpose, we equate labor supply by young to total consumption by the middle-aged and old in a particular time period: $h_t = \mathbb{E}\{q^m_{\xi,t-1} + q^o_{\xi,t-2}\}$. The solution leads us to determine:

$$h_t = \beta(1+\beta)\left[\frac{\Theta}{\pi_t} + \frac{1-\Theta}{\varphi_{t-1}\pi_{t-1}\pi_t}\right]. \tag{7}$$

2.2 Optimal Policy

As is standard in OLG models, we let welfare across generations be aggregated according to the subjective discount factor β. A social planner with perfect commitment therefore maximizes

$$\mathcal{W}_t = \sum_{s=t}^{\infty} \beta^{s-t}\left[\mathbb{E}\left\{\varepsilon u(q^m_{\xi,s-1})\right\} + \mathbb{E}\left\{\delta u(q^o_{\xi,s-2})\right\} - h_s\right] \tag{8}$$

by choosing a sequence $\{\varphi_s, \pi_{s+1}\}_{s=t+1}^{\infty}$ subject to $\varphi_s \leq 1$ and taking as given $\{\varphi_s, \pi_{s+1}\}_{s=0}^{s=t}$.[2] Note the planner cannot choose π_t, π_{t+1}, and φ_t. These are variables that the young in periods

---

[2]The associated supply of money and bonds is given by Equations (5) and (6).

$t-1$ and $t-2$ need for their decisions, so what the planner communicated in those periods about π_t, π_{t+1}, and φ_t is what the planner has to commit to in period $t, t+1, ..., \infty$. Given the logarithmic utility specification, the welfare function at (8) can be expressed as:

$$\mathcal{W}_t = \sum_{s=t}^{\infty} \beta^{s-t} \left\{ \Delta - \log(\pi_s) - \frac{\Theta\beta(1+\beta)}{\pi_s} - \log(\varphi_{s-1}\pi_{s-1}\pi_s) - \frac{(1-\Theta)\beta(1+\beta)}{\varphi_{s-1}\pi_{s-1}\pi_s} \right\}, \quad (9)$$

where
$$\Delta = \iint \left[\varepsilon \log\left(\frac{\varepsilon\beta(1+\beta)}{\varepsilon + \beta\delta} \right) + \delta \log\left(\frac{\delta\beta^2(1+\beta)}{\varepsilon + \beta\delta} \right) \right] \mathrm{d}G(\varepsilon, \delta). \quad (10)$$

Let $\beta^{s+1-t}\chi_s$ denote the Lagrange multiplier on the ZLB constraint for φ_s, which allows to obtain the following first-order conditions that maximizes (9):

$$\pi_{s+2}: \quad 0 = -\frac{2+\beta}{\pi_{s+2}} + \frac{\Theta\beta(1+\beta)}{\pi_{s+2}^2} + \frac{\beta(1+\beta)(1-\Theta)}{\varphi_{s+1}\pi_{s+1}\pi_{s+2}^2} + \frac{\beta^2(1+\beta)(1-\Theta)}{\varphi_{s+2}\pi_{s+2}^2\pi_{s+3}}, \quad (11)$$

$$\varphi_{s+1}: \quad 0 = -\chi_{s+1} - \frac{1}{\varphi_{s+1}} + \frac{\beta(1+\beta)(1-\Theta)}{\varphi_{s+1}^2\pi_{s+1}\pi_{s+2}}. \quad (12)$$

Focusing on steady state policies, equations (11) and (12) reduce to:

$$\pi = \Theta\beta(1+\beta) + \chi\pi(1+\beta) \quad \text{and} \quad \varphi\pi^2 = \beta(1+\beta)(1-\Theta) - \chi\varphi\pi^2. \quad (13)$$

The closed-form test statistic Θ has a crucial implication for the ZLB constraint, which we summarize in the next lemma.

Lemma 1. *The ZLB constraint binds if and only if $\Theta(1+\beta) < 1$.*

Recall that $\theta = \varepsilon/(\varepsilon + \beta\delta)$ and $\Theta = \mathbb{E}[\theta]$, and then note that θ is concave in ε and convex in δ. On the one hand, as shocks to early consumption become more disperse, Θ increases. On the other hand, the dispersion in late consumption preferences decreases the value of Θ. These comparative statics with respect to Θ imply that enlarging heterogeneity in early and late consumption preferences have clear-cut and opposing effects with respect to the distance of an economy to the ZLB and the ability of the government to choose an unconstrained policy mix in equilibrium.

As benchmark for policy, we consider the *Friedman* rule.

Definition 2. *The Friedman rule sets $\pi = \beta$ and $\varphi = 1$ so that the real return on all assets equals the natural rate of interest, $1/\beta$.*

When early consumption preferences are disperse (relative to the dispersion observed in discount factors) those agents with low MU to consume early have a tendency to postpone consumption, which increases aggregate savings. *Ceterus paribus*, at the Friedman rule this results in an inefficiently low early consumption in the aggregate and an inefficiently high aggregate

late consumption. In order to stimulate the early consumption and suppress late consumption, relative to the allocations implied by the Friedman rule, inflation and nominal rates should be reduced in tandem. However, because of the ZLB, interest rates cannot be reduced to take values in the negative domain. Therefore, at ZLB, increasing the inflation becomes part of the (constrained) optimal policy mix of the government. The intuition for this surprising result relates to the compounded effects of inflation, which we observe in (3): inflation reduces late consumption more than early consumption. As we will quantitatively evaluate in the next section, constrained optimal policy results in output and welfare losses - relative to an unconstrained optimum with negative interest rates on nominal bonds.

When discount factor shocks are relatively disperse (compared to the heterogeneity of early consumption preferences) those agents with low MU to consume late have a tendency to increase early consumption, which results in inefficiently low aggregate savings at the Friedman rule. In order to correct for the inefficiently high early consumption implied by the Friedman rule, inflation and nominal interest rates should both increase in tandem in the optimal policy mix of the government. Hence, the economy moves away from the ZLB.

The qualitative properties that we obtained with respect to the effects of preference dispersion on tightness of the ZLB, yield the following steady-state characterization of optimal policy in equilibrium.

Proposition 1. *The optimal policy mix in a steady-state equilibrium satisfies:*

$$\varphi = \begin{cases} 1 & \text{if } \Theta(1+\beta) < 1 \\ \frac{1-\Theta}{\Theta^2\beta(1+\beta)} & \text{if } \Theta(1+\beta) \geq 1 \end{cases} \tag{14}$$

$$\pi = \begin{cases} (1+\beta)\frac{\Theta\beta+\sqrt{\beta}\sqrt{\beta\Theta^2+4(2+\beta)(1-\Theta)}}{2(2+\beta)} & \text{if } \Theta(1+\beta) < 1 \\ \Theta\beta(1+\beta) & \text{if } \Theta(1+\beta) \geq 1 \end{cases} \tag{15}$$

where the inflation rate when the ZLB binds, is the positive root of the following polynominal:

$$(2+\beta)\pi^2 - \Theta\beta(1+\beta)\pi - \beta(1+\beta)^2(1-\Theta) = 0. \tag{16}$$

These key results imply that while increasing the heterogeneity in current consumption preferences constrains the policy choice set of the government, raising the discount factor heterogeneity improves the freedom of the government in choosing a welfare maximizing optimal policy. We follow up on this proposition with a corollary to summarize the implications of preference heterogeneity on optimal inflation, which also follows from the basic intuition related to the effects of preference heterogeneity on tightness of the ZLB.

Corollary 1. *The optimal inflation rate satisfies $\beta \leq \pi$, with equality if and only if $\Theta(1+\beta) = 1$.*

With no idiosyncratic shocks to preferences, or when there are preference shocks but such that $\Theta(1+\beta) = 1$, the optimal inflation rate is at the Friedman rule. In the absence of the endogenous zero lower bound constraint, $\pi < \beta$ would be an optimal inflation policy when $\Theta(1+\beta) < 1$.

2.3 Qualitative Illustration of Optimal Policy

In order to conclude the qualitative discussions resulting from the benchmark model, we turn towards illustrating optimal policy design for some special cases with respect to the heterogeneity in preferences. First we discuss the case of *no preference shocks or policy neutral shocks.* Suppose there is no intra-generational heterogeneity of preferences, meaning there are no idiosyncratic preference shocks and $\Theta(1 + \beta) = 1$. In this case, like in standard micro-founded models of money, the Friedman rule is optimal. The same optimal policy prevails also when there is preference heterogeneity but with policy neutral implications. This happens with a non-degenerate distribution of shocks for which $\Theta(1 + \beta) = 1$. For instance, when $\varepsilon + \beta\delta = 1 + \beta$ for all $(\varepsilon, \delta) \in \Xi$.

Next, we consider a case with only *shocks to preferences for late consumption.* In this case $\Theta = \mathbb{E}\{1/(1 + \beta\delta)\}$ and because $1/(1 + \beta\delta)$ is convex in δ, Jensen's inequality implies that the ZLB does not bind. We find that optimal policy satisfies

$$\pi = \beta(1+\beta)\mathbb{E}\left\{\frac{1}{1+\beta\delta}\right\} \quad \text{and} \quad \varphi\pi = \frac{\beta(1+\beta)}{\pi}\mathbb{E}\left\{\frac{\beta\delta}{1+\beta\delta}\right\},$$

and because $\beta\delta/(1 + \beta\delta)$ is concave in δ, it follows from Jensen's inequality that $\pi > \beta$ and $\varphi\pi < \beta$. This means that optimal policy deviates from the Friedman rule in two dimensions. First, inflation is such that the real return on money falls short of the real natural interest rate. Second, the nominal interest rate is such that the real return on bonds exceeds the real natural interest rate.

Finally, we describe a case with only *shocks to preferences for early consumption.* Then $\Theta = \mathbb{E}\{\varepsilon/(\varepsilon + \beta)\}$ and because $\varepsilon/(\varepsilon + \beta)$ is concave in ε, Jensen's inequality implies that the ZLB constraint binds. Interestingly, the inflation rate satisfies $\pi > \beta$. The real return on money thus falls short of the real natural interest rate. Moreover, if the dispersion of shocks to preferences for early consumption is sufficiently large or the subjective discount factor is sufficiently high, we get a strictly positive rate of inflation as an optimal policy.

3 Quantitative Analysis

We are now ready to quantify the welfare, output, and consumption consequences of optimal policy and the endogenous ZLB across parameterized versions of the model. For this purpose, we start with a first-best economy. The first-best is characterized by a complete financial market structure, ensuring the most efficient distribution of resources. A complete financial

market structure implies that preference shocks are observable and insurable, so newborn agents maximize:

$$\mathcal{U}_t = \max_{h_t, \{q^m_{\xi,t}, q^o_{\xi,t}\}_{\xi \in \Xi}} \left\{ -h_t + \beta \mathbb{E}[\varepsilon u(q^m_{\xi,t})] + \beta^2 \mathbb{E}[\delta u(q^o_{\xi,t})] \right\}, \tag{17}$$

$$\text{s.t.} \quad \pi_{t+1}[\mathbb{E}[q^m_{\xi,t}] + \min\{\varphi_{t+1}, 1\}\pi_{t+2}\mathbb{E}[q^o_{\xi,t}] + \tau_{t+1}] \leq h_t. \tag{18}$$

This means, the state-contingent budget constraint only needs to hold in expectation. It is straightforward to show that the optimization problem at (17) and (18) implies:

$$h_t = \beta(1 + \beta) + \tau_{t+1}, \quad q^m_{\xi,t} = \varepsilon \beta / \pi_{t+1}, \quad \text{and} \quad q^0_{\xi,t} = \delta \beta^2 / [\min\{\varphi_{t+1}, 1\}\pi_{t+1}\pi_{t+2}]. \tag{19}$$

In turn, the first-best allocations at (19) yields the Friedman rule ($\pi = \beta$ and $\varphi = 1$) as the optimal monetary policy, and $h = 2$ as the optimal aggregate output.

The questions that we would like to address in our quantitative analysis are threefold: (i) How does an economy with preference heterogeneity and incomplete markets compare to a first-best economy? (ii) How large are quantitative welfare losses when policy would simply adhere to the Friedman rule? (iii) What are the welfare losses from being unable to let interest rates take values in the negative domain?

3.1 Parametrization

To parameterize the model, we set the length of a period to 20 years. In order to represent a real natural interest rate of 2% per annum, we choose $\beta = 1/(1.02)^{20} \approx 0.8195$. We assume that ε and δ are drawn from independent log-normal distributions with standard deviations σ_ε and σ_δ, respectively. The first-best levels of early and late consumption are then log-normally distributed with mean one, and (relative) standard deviations of σ_ε and σ_δ. Hence, σ_ε and σ_δ have intuitive economic interpretations. Moreover, choosing log-normal distributions for our preference shocks has straightforward implications too for the heterogeneity in agents' desire to save: the functional form implies that our key definition from Section 2, $\theta = \varepsilon/(\varepsilon + \beta\delta)$, follows a logit-normal distribution. Specifically, $\text{logit}(\theta)$ is normally distributed with mean $\mu_\varepsilon - \mu_\delta - \log(\beta)$ and standard deviation $\sqrt{\varsigma_\varepsilon^2 + \varsigma_\delta^2}$. The variables $\mu_x = -\log(1 + \sigma_x^2)/2$ and $\varsigma_x^2 = \log(1 + \sigma_x^2)$ are the mean and variance of $\log(x)$, respectively, with $x \in \{\varepsilon, \delta\}$.

The quantitative results from our parameterized model are presented in Figures 1-6. In these figures, we primarily study the consequences of varying the heterogeneity of preferences, σ_ε and σ_δ, on optimal policy design and macroeconomic outcomes. For this purpose, we let the standard deviations of preference shocks vary between zero and one.[3]

---

[3]The upper bound of one is motivated by the fact that numerically we approximate a continuous log-normal distribution with a discrete distribution. With high standard deviations, this approximation yields significant differences between the theoretical moments of a log-normal distribution and the actual moments of the distri-

3.2 Optimal Policy

Figure 1 illustrates the behaviour of the sufficient statistic Θ, which is the average consumption quote of the middle-aged. The plot in Figure 1 re-emphasizes the point that when the spread in late consumption shocks is small enough, even a very low standard deviation in early-consumption shocks induces ZLB to bind. This prevails when $\Theta(1 + \beta)$ is greater than 1 or equivalently, Θ is greater than $1/(1 + \beta) \approx 0.5596$.

Figures 2a and 2b present the implications of preference heterogeneity on the optimal policy design with respect inflation and nominal interest rates. Figure 2a shows that the heterogeneity in both early and late consumption shocks cause deviations of optimal inflation from the Friedman rule, $\pi = \beta$ and $\varphi = 1$, which requires a deflation of approximately 1.96% per annum. However, the quantitative policy effects of preference heterogeneity resulting from the spreads in ε and δ are substantially different from each other. For example, while keeping $\sigma_\varepsilon = 0$ and setting $\sigma_\delta = 1$ implies a +1.6%-point change from the Friedman rule as the optimal inflation policy, setting $\sigma_\varepsilon = 1$ and $\sigma_\delta = 0$ results in an optimal deviation of the inflation rate from the Friedman-rule by only +0.2%-points. Intuitively, this is because at the ZLB, high inflation discourages both early and late consumption. Ideally, a policymaker promotes early consumption and discourages late consumption when the spread in ε dominates the spread in δ. When ZLB is binding, it is impossible to stimulate early consumption and discourage late consumption simultaneously. Because of compounding, inflation affects late consumption more than the early consumption, and thus it becomes optimal to set the inflation rate higher than the Friedman rule. But, only a slight deviation results from the constrained optimum in order to also take into account the negative effects of inflation on early consumption.

Figure 2b reveals that when σ_δ is small, the ZLB binds for a wide range of σ_ε values. This range shrinks as σ_δ goes up. In Figure 2b we observe that when σ_ε is small, since the economy moves away from ZLB, an optimal annual nominal interest rate as high as 2% can be attained.

As a comparison, Figures 2c and 2d present the implications of preference heterogeneity on optimal policy absent a ZLB on nominal rates, i.e. policy rules that allow for nominal interest rates in the negative domain. It demonstrates that when the spread in early consumption preference shocks is large relative to discount factor shocks, a policy maker wants to shift resources towards early consumption. The policy maker achieves this by reducing inflation relative to the Friedman rule to promote early consumption. In order to undo the strong compounding effect of low inflation that would shift resources towards late consumption, nominal rates are moved into the negative territory and to take values as low as -1.8%.

---

bution we feed to the model. This makes it difficult to connect our numerical results with the properties of a log-normal distribution.

13

3.3 Welfare Effects

We then move on and explore the effects of preference heterogeneity on aggregate welfare. First, Figure 3a quantifies the welfare effects of not being able to implement first-best allocations in an environment where policy is set optimally and constrained by the ZLB. Since utility is linear in labor effort and in equilibrium labor effort is equal to output produced, we express welfare effects as a percentage of the first-best output. We observe that the welfare losses resulting from incomplete markets are substantial when preference heterogeneity arises in our framework. With heterogeneity in both early and late consumption preferences, losses close to 22.5% of first-best output prevail when the (relative) standard deviation of first-best early and late consumption both equal one.

Figure 3b quantifies steady-state welfare losses that arise when instead of being set optimally, policy adheres to the Friedman rule. In the absence of heterogeneity, i.e. when $\sigma_\varepsilon = \sigma_\delta = 0$, the Friedman rule is optimal, which we take as a special-case corner. In comparison to this special-case, implementing the Friedman rule at higher levels of preference heterogeneity results in welfare losses when heterogeneity either in early or late consumption preferences is present. These losses are especially high when heterogeneity in late consumption preferences dominate, in which case the ZLB constraint is slack. The welfare loss from sub-optimal Friedman rule implementation could be as high as 1.1% of first-best output when $\sigma_\delta = 1$.

Because it matters for optimal policy how an economy starts out, simply comparing steady-state welfare levels might be misleading. Therefore, in Figure 3c we consider what happens to an economy that starts with optimal policy but reverts back to the Friedman rule in period 1. That means, $\pi_{t>1} = \beta$ and $\varphi_{t\geq 1} = 1$. With allocations in period 0 and 1 being unaffected by such a policy change, we discount welfare losses back to period 2 and then express them as a percentage of the present value of current and future first-best output. Compared to Figure 3b, we see that equivalent effects arise for such a policy experiment.

In Figure 3d we evaluate the steady-state welfare losses borne by the ZLB. Specifically, we conduct a counterfactual analysis, in which we allow for negative interest rates and compare the resulting welfare against the welfare in our baseline model. This analysis shows that allowing for negative interest rates on nominal bonds imply welfare gains that are as high as 1.4% of the first-best output. Again, besides simply comparing steady-state welfare levels, Figure 3e considers what happens when we revert to standard optimal policy in an economy that starts with optimal negative interest rates. Compared to Figure 3d, we see that equivalent effects arise for such a policy experiment.

Finally, we also report welfare changes induced by implementing optimal policy rules. When the relative standard deviation of first-best early consumption equals one (as a benchmark case), the optimal policy increases steady state welfare by 0.2% of first-best level of welfare. However,

when the relative standard deviation of first-best late consumption equals one, the optimal policy increases steady state welfare by 1.2% of first-best level of welfare - a larger quantitative effect caused by a non-binding ZLB. Allowing for negative interest rates stimulates the aggregate welfare by 1.4% of first-best welfare when the relative standard deviation of first-best early consumption levels equals one.

3.4 Output Effects

Figure 4a shows the effect on output of not being able to implement first-best allocations. We see that only when the ZLB binds, output in an optimal policy regime changes relative to that in a first-best economy. Output losses of not-being able to implement negative interest rates are as large as 4% when $\sigma_\varepsilon = 1$ and $\sigma_\delta = 0$. Intuitively, these output losses are due to the fact that at the ZLB, suppressing late consumption comes at the cost of suppressing early consumption.

Figure 4a shows the effects on output when instead of being set optimally, policy adheres to the Friedman rule. When the ZLB does not bind, implementing the Friedman rule improves output but, as established earlier, at the expense of steady-state welfare.

3.5 Consumption Effects

We conclude the quantitative analysis by decomposing the output effects as consequences on early and late consumption. Figures 5a and 5c show that when the ZLB does not bind, aggregate early and late consumption are equal to their first-best levels. Since stimulating early consumption and suppressing late consumption simultaneously is difficult with an endogeous ZLB, we see that when early consumption shocks dominate and thus the ZLB binds, early consumption reduces and late consumption increases relative to first-best levels. With $\sigma_\varepsilon = 1$ and $\sigma_\delta = 0$, these effects are large; early consumption reduces by almost 20% while late consumption increases by more than 10% relative to first-best.

Figures 5b and 5d show that when implementing the Friedman rule instead of optimal policy, early consumption always reduces. When reverting to the Friedman rule, late consumption however increases if the ZLB is slack but reduces if the ZLB binds.

Finally, in Figure 6 we study the distributional effects of policy on consumption. We do so for two different cases; $\{\sigma_\varepsilon, \sigma_\delta\} = \{0.5, 1\}$ and $\{\sigma_\varepsilon, \sigma_\delta\} = \{1, 0.5\}$. For the first case, the ZLB does not bind, optimal inflation is -0.603%, and the optimal nominal rate is 1.310%. For the second case, the ZLB binds and optimal inflation is -0.847%. In the counterfactual of the ZLB absence, in the latter case optimal inflation would be -1.599% and the optimal nominal rate -1.879%. Figures 6a and 6b show how θ, i.e. middle aged agents' consumption quote, is distributed.

For the first case, Figures 6c and 6e show how optimal policy shifts the distribution of early

15

consumption to the left and the distribution of late consumption to the right, when compared to consumption distributions under the Friedman rule.

For the second case, Figures 6d and 6f show how optimal policy shifts the distribution of both early consumption and late consumption to the left when compared to consumption distributions under the Friedman rule. Absent a ZLB in the second case, optimal policy shifts the distribution of late consumption further to the left but the distribution of early consumption back to the right.

3.6 Summary

The quantitative analysis reveals that in the presence of heterogeneous preferences, optimal policy partially substitutes for the lack financial development. However, it is important to understand the nature of preference heterogeneity and in particular whether heterogeneity is resulting from differences in early consumption shocks or discount factors. Moreover, we also show that the presence of an endogenous ZLB is a quantitatively important constraining factor in maximizing the welfare gains from optimal policy implementation.

4 A Model with Real Investment and Private Assets

In our benchmark model households do not have the option to accumulate wealth through private means. In order to check the sensitivity of our qualitative findings to this benchmark specification, in this section, we introduce a private asset into the model and allow that asset to compete with nominal bonds. We continue to assume the same structure as before and also that in the GM money is the only medium of exchange. However, different from what we introduced in Section 2, households have access to a long-run investment technology which serves as an alternative form of saving instrument in addition to nominal bonds and yields a real return. We are interested in understanding whether the optimal policy mix that we derived in Section 2, remains qualitatively valid when households face an alternative (real) investment option.

Formally, we assume that in any time period t, all young households have access to a decreasing returns to scale investment technology during the GM. Specifically, for each k_t units of goods invested in GM t, the technology returns

$$y_{t+2} = \frac{A k_t^{1-\alpha}}{1-\alpha}, \quad \alpha \in (0,1) \tag{20}$$

units of goods in the GM of period $t+2$. Furthermore, in the FM of period $t+1$, a household can issue claims against the future output of its investment technology. Other households and also the government, can purchase these claims. When the government engages in asset purchases, it leads to the emergence of asset-backed money. That means, in FM of period $t+1$ the government

can buy privately issued financial claims by printing money. By selling off these claims in period $t+2$, the injected money can be withdrawn and the government may earn seignorage revenues. Importantly, the ability of the government to issue asset-backed money implies full control over inflation and nominal interest rates.

4.1 Optimizing Behavior and Equilibrium

We let ψ_t denote the FM t price of a claim against one unit of good to be delivered in the GM of period $t+1$ - measured in terms of the GM goods in period t. A no arbitrage condition implies $\psi_t = \varphi_t \pi_{t+1}$. The lifetime utility maximization problem of a household then becomes:

$$\mathcal{U}_t = \max_{h_t,k_t,\{q^m_{\xi,t},q^o_{\xi,t}\}_{\xi \in \Xi}} \left\{-h_t - k_t + \beta \mathbb{E}[\varepsilon u(q^m_{\xi,t})] + \beta^2 \mathbb{E}[\delta u(q^o_{\xi,t})]\right\}, \tag{21}$$

$$\text{s.t.} \quad \pi_{t+1}[q^m_{\xi,t} + \varphi_{t+1}\pi_{t+2}q^o_{\xi,t} + \tau_{t+1}] \leq h_t + \frac{\varphi_{t+1}\pi_{t+1}\pi_{t+2}Ak_t^{1-\alpha}}{1-\alpha}, \quad \forall \xi \in \Xi. \tag{22}$$

$$h_t \geq 0, \quad \text{and} \quad k_t \geq 0. \tag{23}$$

Based on the extended optimization problem at (21)-(23) we first establish the following lemma, which characterizes the equilibrium.

Lemma 2. *(i) Aggregate supply and demand in any time period t are characterized by:*

$$Q^s_t = \beta(1+\beta) + \pi_{t+1}\tau_{t+1} + A^{1/\alpha}\left([\varphi_{t-1}\pi_{t-1}\pi_t]^{\frac{1-\alpha}{\alpha}} - [\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{\frac{1}{\alpha}}\right)/(1-\alpha), \tag{24}$$

$$Q^d_t = \Theta\beta(1+\beta)/\pi_t + \beta(1+\beta)(1-\Theta)/(\varphi_{t-1}\pi_{t-1}\varphi_t). \tag{25}$$

(ii) The amount of labor devoted to acquiring cash and investment goods are given by:

$$h_t = \beta(1+\beta) + \pi_{t+1}\tau_{t+1} - [A\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{1/\alpha}/(1-\alpha) \quad \text{and} \quad k_t = [A\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{1/\alpha}. \tag{26}$$

(iii) Finally, in a steady-state equilibrium the following condition ensures that the labor-supply satisfies $h > 0$:

$$\frac{A^{\frac{1}{\alpha}}(\varphi\pi^2)^{\frac{1-\alpha}{\alpha}}}{1-\alpha} < \frac{\beta(1+\beta)(\varphi\pi\Theta+1-\Theta)}{\varphi\pi^2}. \tag{27}$$

We would like to first note that, in order to keep the analysis comparable to that of Section 2, we restrict our attention to cases in which $h > 0$ holds, because otherwise there is no inter-generational monetary exchange. Lemma 2 shows that $h > 0$ is the case if the productivity of the long-run investment technology is sufficiently low. The intuition for this property is straight-forward: when the long-run technology is too productive, agents can finance consumption needs in both periods using only the long-run technology. In addition to this property, Lemma 2 also characterizes the equilibrium quantities produced (and demanded), which different than the

benchmark specification contain terms related to labor- and capital-based output.

4.2 Optimal Policy

A government with perfect commitment maximizes

$$\mathcal{W}_t = \sum_{s=t}^{\infty} \beta^{s-t} \left[\mathbb{E} \left\{ \varepsilon u(q_{i,s-1}^m) \right\} + \mathbb{E} \left\{ \delta u(q_{i,s-2}^0) \right\} - h_s - k_s \right] \tag{28}$$

by choosing a sequence $\{\varphi_s, \pi_{s+1}\}_{s=t+1}^{\infty}$ s.t. $\varphi_s \leq 1$, and taking as given $\{\varphi_s, \pi_{s+1}\}_{s=0}^{s=t}$ and $\{k_s\}_{s=0}^{s=t}$. Using our logarithmic utility specification, the welfare function at (28) that the government maximizes can be written as:

$$\mathcal{W}_t = \sum_{s=t}^{\infty} \beta^{s-t} \left\{ \begin{array}{l} \Delta - \log(\pi_s) - \dfrac{\Theta\beta(1+\beta)}{\pi_s} - \log(\varphi_{s-1}\pi_{s-1}\pi_s) - \dfrac{(1-\Theta)\beta(1+\beta)}{\varphi_{s-1}\pi_{s-1}\pi_s} \\ + A^{\frac{1}{\alpha}} \dfrac{[\varphi_{s-1}\pi_{s-1}\pi_s]^{\frac{1-\alpha}{\alpha}} - (1-\alpha)[\varphi_{s+1}\pi_{s+1}\pi_{s+2}]^{\frac{1}{\alpha}}}{1-\alpha} \end{array} \right\}, \tag{29}$$

with Δ given by (10). We characterize the government's optimal policy mix in the following proposition.

Proposition 2. *The optimal policy mix in a steady-state equilibrium satisfies:*

$$\varphi = \begin{cases} 1 & \text{if } \Theta(1+\beta) < 1 \\ \dfrac{1-\Theta}{\Theta^2\beta(1+\beta)} + o_\varphi & \text{if } \Theta(1+\beta) \geq 1 \end{cases} \tag{30}$$

$$\pi = \begin{cases} (1+\beta)\dfrac{\Theta\beta+\sqrt{\beta}\sqrt{\beta\Theta^2+4(2+\beta)(1-\Theta)}}{2(2+\beta)} - o_\pi & \text{if } \Theta(1+\beta) \leq 1 \\ \Theta\beta(1+\beta) & \text{if } \Theta(1+\beta) > 1 \end{cases} \tag{31}$$

When the ZLB constraint does not bind, the difference between the optimal price of bonds with and without an investment technology is given by o_φ. This object is a positive solution of the following equation

$$o_\varphi[\Theta\beta(1+\beta)]^2 = \frac{\left\{ A\beta(1+\beta)[1-\Theta+o_\varphi\Theta^2\beta(1+\beta)] \right\}^{\frac{1}{\alpha}}}{\alpha} \left[1 - (1+\beta)\frac{1-\Theta+o_\varphi\Theta^2\beta(1+\beta)}{\beta} \right]. \tag{32}$$

When the ZLB constraint binds, the difference between optimal inflation with and without an investment technology is given by o_π. This object is the unique positive solution of the following equation

$$0 = o_\pi \left[o_\pi(2+\beta) - 2\Upsilon(2+\beta) + \Theta\beta(1+\beta) \right] - \left[A(\Upsilon - o_\pi)^2 \right]^{\frac{1}{\alpha}} \left[1 - (\Upsilon - o_\pi)^2/\beta^2 \right] (1+\beta)/\alpha, \tag{33}$$

where $\Upsilon = (1+\beta)[\Theta\beta + \sqrt{\beta}\sqrt{\beta\Theta^2 + 4(2+\beta)(1-\Theta)}]/[4+2\beta]$.

With A sufficiently large, o_φ may not be determined uniquely and additional second-order conditions need to be considered. In what follows, we consider A sufficiently small so that o_φ is pinned down uniquely by the associated equation in Proposition 2. With respect to the qualitative properties we discuss next, this simplification has been made without loss of generality.

Proposition 2 highlights an important robustness of our optimal policy characterization. Specifically, the threshold that determines the cut-off point for a ZLB constraint remains the same as in Section 2. To review from there, *ceteris paribus*, an increase in the heterogeneity of late-consumption preferences lowers the likelihood of a binding ZLB, while increasing the heterogeneity of early-consumption preferences raises the likelihood of a binding ZLB. Qualitatively, optimal policy also looks similar to what we characterized in Section 2, except either the optimal interest rate on nominal bonds (when the ZLB is slack) or the optimal inflation rate (when the ZLB binds) exhibits a distortionary wedge due to the presence of real asset formation by households. It also remains true as in Section 2 that $\pi \geq \beta$, with equality if and only if $\Theta(1 + \beta) = 1$, and $\varphi \leq 1$ with equality if and only if $\Theta(1 + \beta) \leq 1$.

With respect to the distortionary wedges on inflation and bond pricing policies, we obtain important comparative statics with an intuitive interpretation. Specifically, $\lim_{A \to 0} o_\varphi = 0$ and $\lim_{A \to 0} o_\pi = 0$. Also, $\partial o_\varphi / \partial A > 0$ and $\partial o_\pi / \partial A > 0$. This means, the higher the productivity of private investment, the larger become the distortionary wedges associated with optimal policy. The reason is that real investment is affected by the real return earned by government bonds through a no arbitrage condition. Since first-best investment prevails when this real rate of return equals the natural real interest rate, the presence of a real investment technology pulls optimal policy towards the Friedman rule.

4.3 Qualitative Illustration of Optimal Policy

When there are *no preference shocks or policy neutral shocks*, such that $\Theta(1 + \beta) = 1$, we conclude as in Section 2 with the optimality of the Friedman rule ($\pi = \beta$ and $\varphi = 1$). Observing how investment enters the first-order conditions, we see that the first-best level of investment requires $\varphi \pi^2 = \beta^2$. With no or policy-neutral idiosyncratic preference shocks, optimal levels of consumption are attained when $\pi = \beta$ and $\varphi = 1$. Thus, the introduction of a real investment opportunity and private asset formation has no effect on the optimality of the Friedman rule in this particular case.

In a specification with *shocks to only late consumption*, which implies $\Theta(1 + \beta) > 1$, the ZLB is slack. Early consumption is determined by π and late consumption by $\varphi \pi^2$. Because investment is determined by $\varphi \pi^2$, the investment technology only affects optimal nominal interest rates and not the optimal inflation rate. With $\Theta(1 + \beta) > 1$, optimal late consumption requires $\varphi \pi^2 = \beta(1 + \beta)(1 - \Theta) < \beta^2$ while optimal investment requires $\varphi \pi^2 = \beta^2$. Compared

to the Friedman rule, optimal policy balances these two conditions and therefore implies too little investment and too much late consumption - as in the benchmark specification.

When there are *shocks to early consumption only*, which implies $\Theta(1 + \beta) < 1$, the ZLB constraint binds. Optimal aggregate consumption requires $\pi > \beta$ while optimal investment requires $\pi = \beta$. Optimal policy trades off these two. The investment technology therefore affects the inflation, and compared to the Friedman rule, the optimal inflation implies $\pi > \beta$ with too much investment and too little aggregate consumption - as in the benchmark specification.

5 Conclusion

We studied preference heterogeneity and optimal monetary policy in a dynamic general equilibrium model of money with incomplete financial markets. Our benchmark theoretical model incorporates two dimensional consumption preference heterogeneity: idiosyncratic shocks to current consumption and discount factors. Although we model preference shocks in a general way, the model is tractable to yield clean theoretical insights and quantifiable effects. We show that the shocks to current (early) consumption could cause an endogenous ZLB to bind - an effect that could get mitigated by discount factor heterogeneity. We characterize optimal policy and study the behavior of optimal inflation and nominal interest rates. The optimal policy, in general, deviates from the Friedman rule, but the motives for deviation are different depending on the type of preference heterogeneity. Our qualitative results are robust to allowing for alternative savings instruments.

The quantitative analysis based on the theoretical framework reveals that - when implemented optimally - policy can partially substitute for the lack of complete financial markets. However, the ZLB is a constraining factor. We conduct a counterfactual analysis to evaluate the macroeconomic effects of optimal negative interest rates. This counterfactual analysis shows that allowing for negative interest rates on nominal bonds can result in welfare gains that are as high as 1.4% of first-best output.

References

[1] Andolfatto, D. (2011). A note on the societal benefits of illiquid bonds. *Canadian Journal of Economics/Revue canadienne d'économique*, 44(1), 133-147.

[2] Azariadis, C. (1981). Self-fulfilling prophecies. *Journal of Economic Theory*, 25(3), 380-396.

[3] Boel, P., and Camera, G. (2006). Efficient monetary allocations and the illiquidity of bonds. *Journal of Monetary Economics*, 53(7), 1693-1715.

[4] Boel, P., and Waller, C. (2018). On the Theoretical Efficacy of Quantitative Easing at the Zero Lower Bound. *Federal Reserve Bank of St. Louis Working Paper*, 2015-27.

[5] Curran, M., and Dressler, S. (2019). Preference Heterogeneity, Inflation, and Welfare (No. 40). *Villanova School of Business Department of Economics and Statistics, Working Paper*.

[6] Cagetti, M. (2003). Wealth accumulation over the life cycle and precautionary savings. *Journal of Business & Economic Statistics*, 21(3), 339-353.

[7] De Nardi, M., and Fella, G. (2017). Saving and wealth inequality. *Review of Economic Dynamics*, 26, 280-300.

[8] Garleanu, N., and Panageas, S. (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3), 670-685.

[9] Gelman, M. (2020). What drives heterogeneity in the marginal propensity to consume? temporary shocks vs persistent characteristics. *Journal of Monetary Economics*.

[10] Golosov, M., Troshkin, M., Tsyvinski, A., and Weinzierl, M. (2013). Preference heterogeneity and optimal capital income taxation. *Journal of Public Economics*, 97, 160-175.

[11] Hendricks, L. (2007). How important is discount rate heterogeneity for wealth inequality?. *Journal of Economic Dynamics and Control*, 31(9), 3042-3068.

[12] Kocherlakota, N. R. (2003). Societal benefits of illiquid bonds. *Journal of Economic Theory*, 108(2), 179-193.

[13] Krusell, P., and Smith, Jr, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5), 867-896.

[14] Van Buggenum, H. (2020). Coexistence of Money and Interest-Bearing Bonds. *Tilburg University CentER Working Paper*.

[15] Venti, S. F., and Wise, D. A. (2000). Choice, chance, and wealth dispersion at retirement (No. w7521). *National Bureau of Economic Research.*

[16] Wen, Y., and Dong, F. (2017). Optimal monetary policy under negative interest rate. *FRB St. Louis Working Paper*, (2017-19).

[17] Williamson, S. D. (2019). Low real interest rates and the zero lower bound. *Review of Economic Dynamics*, 31, 36-62.

Appendix

Derivations of real quantities in the benchmark model

We let $\beta\lambda_{\xi,t}$ denote the Lagrange multiplier associated with the (idiosyncratic) state-contingent budget constraint (2b). With log-utility, first-order conditions become:

$$h_t: \quad 0 = -1 + \beta\mathbb{E}[\lambda_{\xi,t}] \tag{A.1}$$

$$q_{\xi,t}^m: \quad 0 = -\pi_{t+1}\lambda_{\xi,t} + \varepsilon/q_{\xi,t}^m \tag{A.2}$$

$$q_{\xi,t}^o: \quad 0 = -\min\{\varphi_{t+1}\}\pi_{t+1}\pi_{t+2}\lambda_{\xi,t} + \beta\delta/q_{\xi,t}^o \tag{A.3}$$

Combining (A.2) and (A.3) results in

$$q_{\xi,t}^o = q_{\xi,t}^m \frac{\beta\delta}{\varepsilon} \frac{1}{\min\{\varphi_{t+1}, 1\}\pi_{t+2}}. \tag{A.4}$$

Using that the budget constraint (2b) must hold with equality, we find

$$h_t = \pi_{t+1}\left[q_{\xi,t}^m \frac{\varepsilon + \beta\delta}{\varepsilon} + \tau_{t+1}\right] \tag{A.5}$$

Rewriting (A.5) in terms of $q_{\xi,t}^m$ and substituting into (A.2), we obtain

$$\lambda_{\xi,t} = \frac{\varepsilon + \beta\delta}{h_t - \pi_{t+1}\tau_{t+1}} \tag{A.6}$$

Using this in (A.1) together with Assumption 1 ($\mathbb{E}[\varepsilon] = \mathbb{E}[\delta] = 1$) we find

$$h_t = \beta(1 + \beta) + \pi_{t+1}\tau_{t+1}. \tag{A.7}$$

Using (A.7) in (A.5) yields an expression for $q_{\xi,t}^m$, which we can then use in (A.4) to obtain an expression for $q_{\xi,t}^0$:

$$q_{\xi,t}^m = \frac{\varepsilon}{\varepsilon + \beta\delta}\frac{\beta(1 + \beta)}{\pi_{t+1}} \quad \text{and} \quad q_{\xi,t}^o = \frac{\beta\delta}{\varepsilon + \beta\delta}\frac{\beta(1 + \beta)}{\min\{\varphi_{t+1}, 1\}\pi_{t+1}\pi_{t+2}}. \tag{A.8}$$

$$\square$$

Proof of Lemma 1

When the ZLB does not bind, first-order conditions resulting from optimizing the welfare function at steady-state with respect to π (inflation rate) and φ (bond prices) yields

$$\pi^* = \Theta\beta(1 + \beta), \tag{A.9}$$

$$\varphi^*(\pi^*)^2 \quad = \quad \beta(1+\beta)(1-\Theta). \tag{A.10}$$

Clearly, the ZLB binds if and only if $\varphi^* > 1$. Using (A.9) and (A.10), this is the case when

$$\Theta^2(1+\beta)^2\beta^2 + \Theta\beta(1+\beta) - \beta(1+\beta) < 0. \tag{A.11}$$

Condition (A.11) is satisfied for $-1/\beta < \Theta < 1/(1+\beta)$. Because $\Theta \in (0,1)$, it follows that the ZLB binds if and only if $\Theta(1+\beta) < 1$. □

Proof of Proposition 1

There are two cases to consider: (i) binding ZLB (with $\varphi = 1$) and (ii) non-binding ZLB (with $\varphi < 1$). When $\varphi = 1$, optimal π is determined from:

$$\pi \quad = \quad \frac{\beta(1+\beta)\Theta}{1 - \chi(1+\beta)}, \tag{A.12}$$

$$1 + \chi \quad = \quad \frac{\beta(1+\beta)}{\pi^2}(1-\Theta), \tag{A.13}$$

using which optimal inflation rate π can be solved as the unique positive root of the polynomial:

$$\pi^2(2+\beta) - \pi\beta(1+\beta)\Theta - \beta(1-\beta)^2(1-\Theta) = 0, \tag{A.14}$$

as stated in Proposition 1.

In the case of non-binding ZLB with $\varphi < 1$ (and thus $\chi = 0$), the optimal inflation rate is determined as:

$$\pi = \Theta\beta(1+\beta), \tag{A.15}$$

which also yields:

$$\varphi = \frac{1-\Theta}{\Theta^2\beta(1+\beta)}. \tag{A.16}$$

□

Proof of Corollary 2

That the optimal inflation rate satisfies $\pi \geq \beta$ for the case of $\Theta(1+\beta) \geq 1$ (with $\pi = \beta$ when $\Theta(1+\beta) = 1$) follows immediately from the equation that pins down optimal policy mix Proposition 1. To observe that the optimal inflation is still greater than the Friedman rule, $\pi > \beta$, when $\Theta(1+\beta) < 1$, we first note that at this particular case the equation that determines the optimal inflation satisfies

$$(2+\beta)\pi^2 - \Theta\beta(1+\beta)\pi - \beta(1+\beta)^2(1-\Theta) = 0, \tag{A.17}$$

which has a unique positive root for π. Note that $\pi = \beta$ solves (A.17) if $\Theta(1 + \beta) = 1$. Then, starting at $\pi = \beta$ and $\Theta(1 + \beta) = 1$, if Θ declines (such that $\Theta(1 + \beta) < 1$) the left-hand-side of (A.17) which solves for the optimal inflation with $\Theta(1 + \beta) < 1$ goes down, since $\beta < 1$. Also because at $\pi = \beta$ an increase in π raises the left-hand-side of the polynomial, it yields the result that the optimal inflation rate satisfies $\pi > \beta$ for $\Theta(1 + \beta) < 1$. $\qquad\square$

Proof of Lemma 2

Let $\beta\lambda_{\xi,t}$, $\nu_{h,t}$, and $\nu_{k,t}$ denote the Lagrange multipliers of the associated constraints. Because the non-negativity constraint for capital never binds due to the DRS technology specification, we obtain the following first-order conditions:

$$h_t: \quad 0 = -1 + \nu_{h,t} + \beta\mathbb{E}[\lambda_{\xi,t}], \tag{A.18}$$

$$k_t: \quad 0 = -1 + \varphi_{t+1}\pi_{t+1}\pi_{t+2}Ak_t^{-\alpha}\beta\mathbb{E}[\lambda_{\xi,t}], \tag{A.19}$$

$$q_{\xi,t}^m: \quad 0 = \varepsilon/q_{\xi,t}^m - \pi_{t+2}\lambda_{\xi,t}, \tag{A.20}$$

$$q_{\xi,t}^o: \quad 0 = \delta\beta/q_{\xi,t}^0 - \varphi_{t+1}\pi_{t+1}\pi_{t+2}\lambda_{\xi,t}. \tag{A.21}$$

The solution clearly depends on whether the non-negativity constraint associated with young's labor supply, h, binds and thus whether there are intergenerational transfers in spirit of what we studied in Section 2. In what follows, we focus on equilibria that involve $h > 0$. The amount of labor devoted to acquiring cash and investment goods is then given by

$$h_t = \beta(1+\beta) + \pi_{t+1}\tau_{t+1} - [A\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{1/\alpha}/(1-\alpha) \quad\text{and}\quad k_t = [A\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{1/\alpha}, \tag{A.22}$$

respectively. When middle-aged, households' real financial wealth after taxation therefore equals $\beta(1 + \beta)/\pi_{t+1}$. As in the benchmark model, early and late consumption are then given by (3) and aggregate supply of and demand for goods in any time period t becomes:

$$Q_t^s = \beta(1+\beta) + \pi_{t+1}\tau_{t+1} + A^{1/\alpha}\left([\varphi_{t-1}\pi_{t-1}\pi_t]^{\frac{1-\alpha}{\alpha}} - [\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{\frac{1}{\alpha}}\right)/(1-\alpha), \tag{A.23}$$

$$Q_t^d = \Theta\beta(1+\beta)/\pi_t + \beta(1+\beta)(1-\Theta)/(\varphi_{t-1}\pi_{t-1}\varphi_t). \tag{A.24}$$

The level of lump-sum taxes that closes the model and equilibriates demand and supply is given by:

$$\pi_t\tau_t = \beta(1+\beta)\left\{\Theta[1/\pi_{t-1} - 1] + (1-\Theta)[1/(\varphi_{t-2}\pi_{t-2}\pi_{t-1}) - 1]\right\}$$
$$- \frac{A^{1/\alpha}\left([\varphi_{t-2}\pi_{t-2}\pi_{t-1}]^{\frac{1-\alpha}{\alpha}} - [\varphi_t\pi_t\pi_{t+1}]^{\frac{1}{\alpha}}\right)}{1-\alpha}. \tag{A.25}$$

It follows that total labor supply by young households satisfies:

$$h_t + k_t = \frac{\Theta\beta(1+\beta)}{\pi_t} + \frac{\beta(1+\beta)(1-\Theta)}{\varphi_{t-1}\pi_{t-1}\varphi_t} - A^{1/\alpha}\left(\frac{[\varphi_{t-1}\pi_{t-1}\pi_t]^{\frac{1-\alpha}{\alpha}}}{1-\alpha} - [\varphi_{t+1}\pi_{t+1}\pi_{t+2}]^{\frac{1}{\alpha}}\right).$$

(A.26)

Focusing on a steady state, to ensure h is strictly positive we need the following condition to hold in equilibrium:

$$\frac{A^{\frac{1}{\alpha}}(\varphi\pi^2)^{\frac{1-\alpha}{\alpha}}}{1-\alpha} < \frac{\beta(1+\beta)(\varphi\pi\Theta+1-\Theta)}{\varphi\pi^2}.$$

(A.27)

$\square$

Proof of Proposition 2

Let $\beta^{s+1-t}\chi_s$ denote the Lagrange multiplier on the ZLB constraint to obtain the following first order conditions

$$\pi_{s+2}: \quad 0 = \beta^2\left(-\frac{2+\beta}{\pi_{s+2}} + \frac{\Theta\beta(1+\beta)}{\pi_{s+2}^2} + \frac{\beta(1+\beta)(1-\Theta)}{\varphi_{s+1}\pi_{s+1}\pi_{s+2}^2} + \frac{\beta^2(1+\beta)(1-\Theta)}{\varphi_{s+2}\pi_{s+2}^2\pi_{s+3}}\right)$$

(A.28)

$$+ \frac{[A\varphi_{s+1}\pi_{s+1}\pi_{s+2}]^{\frac{1}{\alpha}}}{\alpha\pi_{s+2}}\left[\frac{\beta^2}{\varphi_{s+1}\pi_{s+1}\pi_{s+2}} - 1\right] + \beta\frac{[A\varphi_{s+2}\pi_{s+2}\pi_{s+3}]^{\frac{1}{\alpha}}}{\alpha\pi_{s+2}}\left[\frac{\beta^2}{\varphi_{s+2}\pi_{s+2}\pi_{s+3}} - 1\right],$$

$$\varphi_{s+1}: \quad 0 = \beta^2\left(-\chi_{s+1} - \frac{1}{\varphi_{s+1}} + \frac{\beta(1+\beta)(1-\Theta)}{\varphi_{s+1}^2\pi_{s+1}\pi_{s+2}}\right) + \frac{[A\varphi_{s+1}\pi_{s+1}\pi_{s+2}]^{\frac{1}{\alpha}}}{\alpha\varphi_{s+1}}\left[\frac{\beta^2}{\varphi_{s+1}\pi_{s+1}\pi_{s+2}} - 1\right].$$

(A.29)

Focusing on steady state, the first-order conditions reduce to:

$$\pi = \Theta\beta(1+\beta) + \chi\pi(1+\beta) \text{ and } \varphi\pi^2 = \beta(1+\beta)(1-\Theta) - \chi\varphi\pi^2 + [A\varphi\pi^2]^{\frac{1}{\alpha}}\left[1 - \varphi\pi^2/\beta^2\right]/\alpha.$$

(A.30)

We note that in steady state, the investment technology enters only in the first-order condition for φ. However, through the Lagrange multiplier for the ZLB, inflation can be affected by the investment technology.

Then, we show that the ZLB binds if and only if $\Theta(1+\beta) < 1$. Consider first $\Theta(1+\beta) \geq 1$ and conjecture that the ZLB is slack this case. Then equation (A.30) implies

$$\pi \geq \beta \quad \text{and} \quad \varphi\pi^2 \leq \beta^2 + [A\varphi\pi^2]^{\frac{1}{\alpha}}\left[1 - \varphi\pi^2/\beta^2\right]/\alpha \Rightarrow \varphi\pi^2 \leq \beta^2.$$

(A.31)

It follows that $\varphi \leq 1$, which verifies our conjecture. Consider second that $\Theta(1+\beta) < 1$ and suppose that the ZLB is also slack in this case. Then equation (A.30) implies

$$\pi < \beta \quad \text{and} \quad \varphi\pi^2 > \beta^2 + [A\varphi\pi^2]^{\frac{1}{\alpha}}\left[1 - \varphi\pi^2/\beta^2\right]/\alpha \Rightarrow \varphi\pi^2 > \beta^2.$$

(A.32)

It follows that $\varphi > 1$, which yields a contradiction. Hence the ZLB must bind when $\Theta(1+\beta) < 1$.

Next, we characterize optimal policy with a slack ZLB constraint. Clearly

$$\pi = \Theta\beta(1+\beta) \quad \text{and} \quad \varphi\pi^2 = \beta(1+\beta)(1-\Theta) + [A\varphi\pi^2]^{\frac{1}{\alpha}}\left[1 - \varphi\pi^2/\beta^2\right]/\alpha. \quad \text{(A.33)}$$

Letting $\varphi\pi^2 = \beta(1+\beta)(1-\Theta) + o'_\varphi$ we have that

$$o'_\varphi = \frac{A[\beta(1+\beta)(1-\Theta) + o'_\varphi]^{1/\alpha}}{\alpha}\left[1 - \frac{\beta(1+\beta)(1-\Theta) + o'_\varphi}{\beta^2}\right]. \quad \text{(A.34)}$$

With $\Theta(1 + \beta) \geq 1$, this implies that $o'_\varphi \geq 0$ with strict inequality if $A > 0$, as well as $\lim_{A\to 0} o'_\varphi = 0$. When A is sufficiently small, Equation (A.34) also pins down o'_φ uniquely. Then, $\partial o'_\varphi/\partial A > 0$. Finally,

$$\varphi = \frac{1-\Theta}{\Theta\beta(1+\beta)} + o_\varphi, \quad o_\varphi = \frac{o'_\varphi}{[\Theta\beta(1+\beta)]^2} \quad \text{(A.35)}$$

where o_φ has the same qualitative properties as o'_φ.

Finally, we characterize optimal policy with a tight ZLB constraint. Using $\varphi = 1$, Equation (A.30) implies that

$$0 = (2+\beta)\pi^2 - \Theta\beta(1+\beta)\pi - \beta(1+\beta)^2(1-\Theta) - [A\pi^2]^{1/\alpha}[1 - \pi^2/\beta](1+\beta)/\alpha. \quad \text{(A.36)}$$

With $\Theta(1+\beta) < 1$, the RHS of Equation (A.36) is negative for all $\pi \leq \beta$. The partial derivative of the RHS of Equation (A.36) w.r.t. π satisfies

$$2\pi(2+\beta) - \Theta\beta(1+\beta) - 2\pi[A\pi^2]^{1/\alpha}\left[\frac{1}{\alpha}\left(\frac{1}{\pi^2} - \frac{1}{\beta}\right) - \frac{1}{\beta}\right]\frac{1+\beta}{\alpha}, \quad \text{(A.37)}$$

which is strictly positive for $\pi > \beta$. Therefore, there exists a unique $\pi > \beta$ which solves Equation (A.36) and constitutes the optimal rate of inflation. Clearly, this solution is decreasing in A and can be written as

$$\pi = \Upsilon - o_\pi \quad \text{and} \quad \Upsilon = (1+\beta)\frac{\theta\beta + \sqrt{\beta}\sqrt{\beta\Theta^2 + 4(2+\beta)(1-\Theta)}}{2(2+\beta)}, \quad \text{(A.38)}$$

with $o_\pi = 0$ if $A = 0$, $o_\pi > 0$ if $A > 0$, $\partial o_\pi/\partial A > 0$, and $\lim_{A\to 0} o_\pi = 0$. Here o_π is the unique positive solution to:

$$0 = o_\pi\left[o_\pi(2+\beta) - 2\Upsilon(2+\beta) + \Theta\beta(1+\beta)\right] - \left[A(\Upsilon - o_\pi)^2\right]^{\frac{1}{\alpha}}\left[1 - (\Upsilon - o_\pi)^2/\beta^2\right](1+\beta)/\alpha. \quad \text{(A.39)}$$
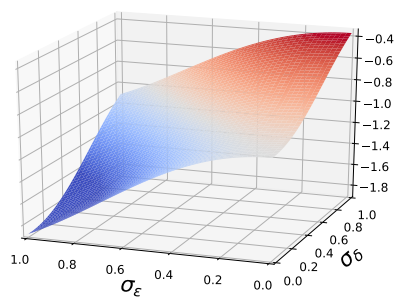
$\square$

27

Tables and Figures
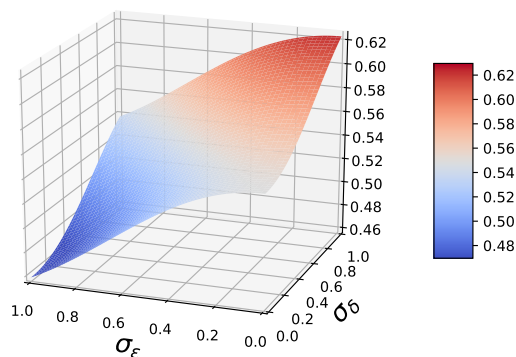
Figure 1: Sufficient statistic Θ for log-normally distributed preference shocks.
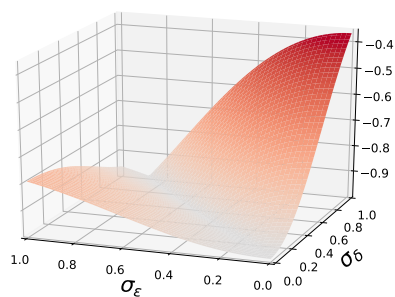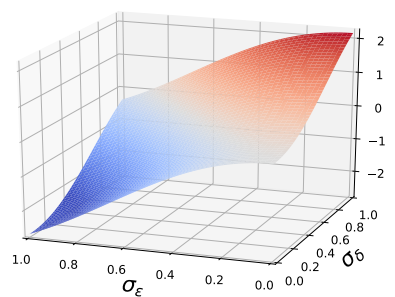

(a) Optimal inflation.
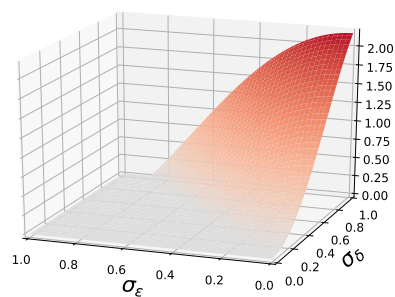
(b) Optimal nominal interest rate.

(c) Optimal inflation without a ZLB.

(d) Optimal nominal interest rate without a ZLB.

Figure 2: Sufficient statistic Θ and optimal policy with log-normal preference shocks. Inflation and nominal interest rates are expressed as annualized percentages. The colormap for Θ is normalized around $1/(1 + \beta)$. The colormap for inflation and nominal interest rates are normalized around the Friedman rule.
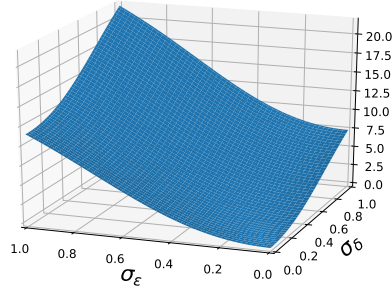
(a) Steady-state welfare loss of not being able to implement first-best allocations.
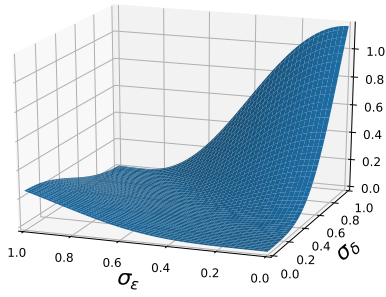
(b) Steady-state welfare loss of Friedman rule.
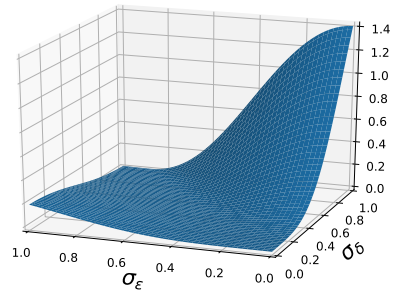
(c) Dynamic welfare loss of Friedman rule.
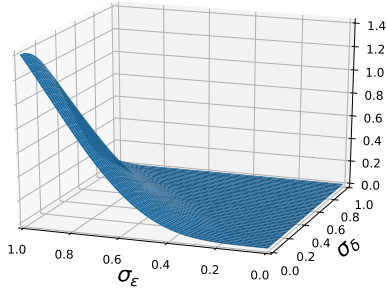
(d) Steady-state welfare loss of ZLB.
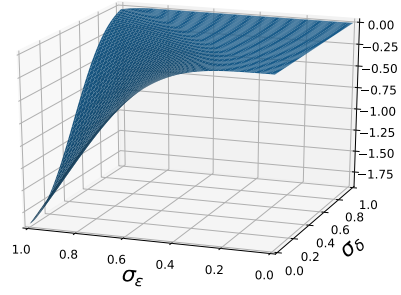
(e) Dynamic welfare loss of ZLB.

Figure 3: Welfare analysis for log-normally distributed preferences shocks. All effects are expressed as a percentage of steady-state output for first-best allocations.
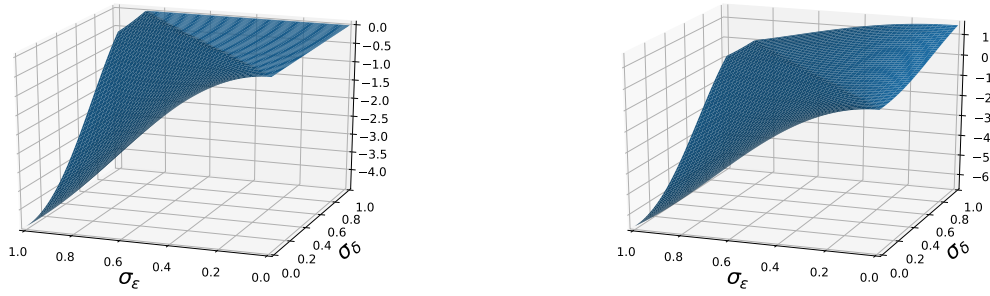
(a) Output with optimal policy relative to first-best.

(b) Output at Friedman rule relative to optimal policy.

Figure 4: Output effects for log-normally distributed preferences shocks. All effects are expressed as a percentage of steady-state output for first-best allocations.
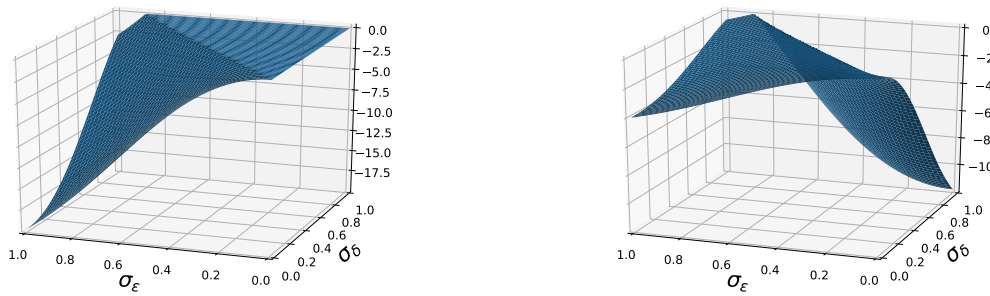
(a) Early cons. with opt. policy relative to first-best.

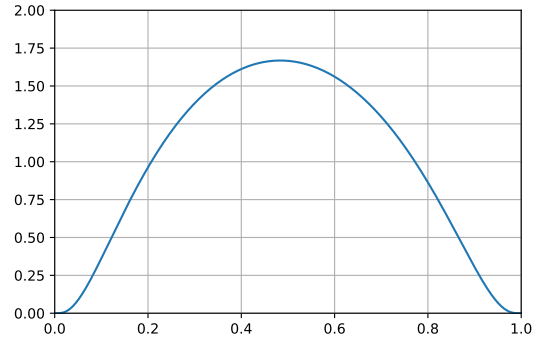(b) Early cons. at Friedman rule relative to opt. policy.
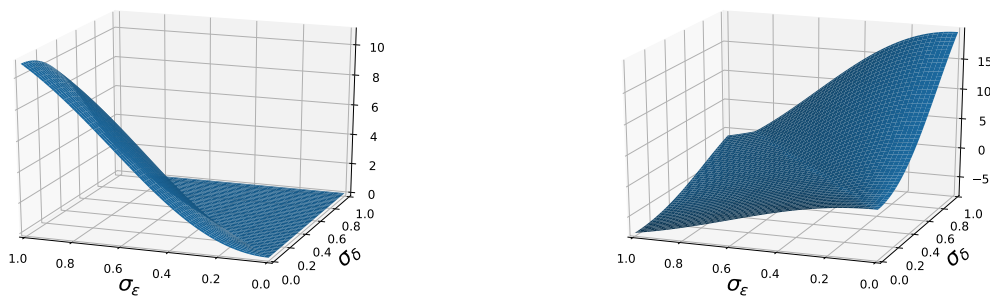
(c) Late cons. with opt. policy relative to first-best.

(d) Late cons. at Friedman rule relative to opt. policy.

Figure 5: Aggregate consumption effects for log-normally distributed preferences shocks. All effects are expressed a as a percentage.
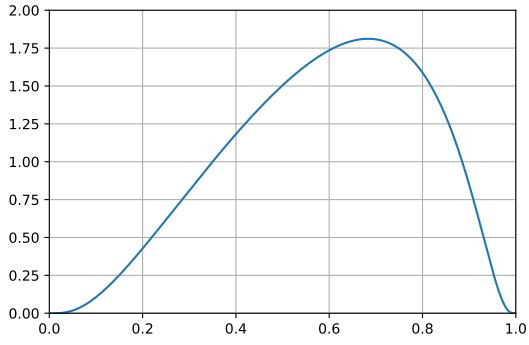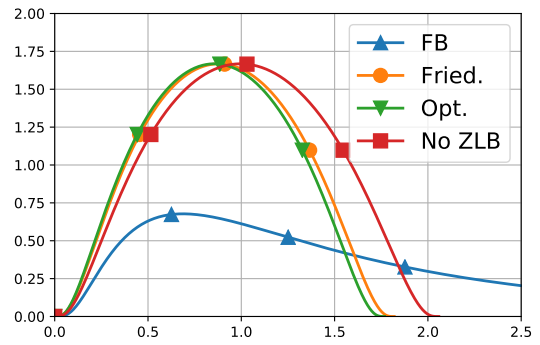
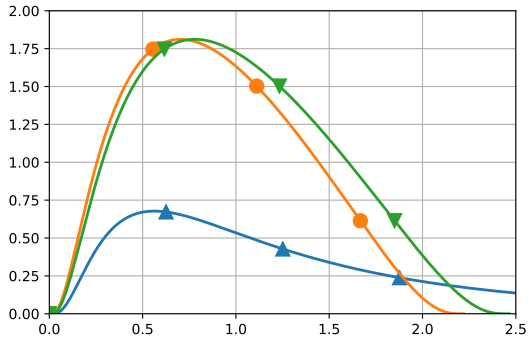(a) Pdf for θ with $\sigma_\varepsilon = 0.5$ and $\sigma_\delta = 1$.

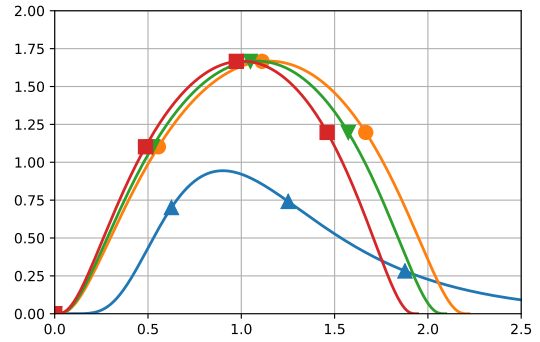(b) Pdf for θ with $\sigma_\varepsilon = 1$ and $\sigma_\delta = 0.5$.
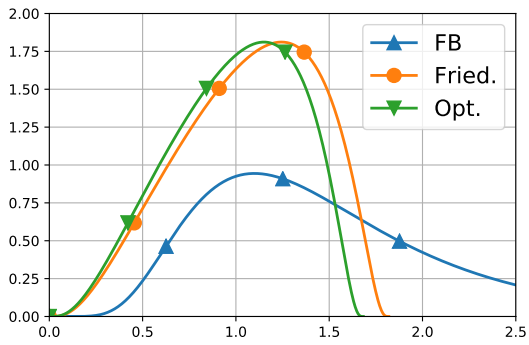
(c) Pdf for early cons. with $\sigma_\varepsilon = 0.5$ and $\sigma_\delta = 1$.

(d) Pdf for early cons. with $\sigma_\varepsilon = 1$ and $\sigma_\delta = 0.5$.

(e) Pdf for late cons. with $\sigma_\varepsilon = 0.5$ and $\sigma_\delta = 1$.

(f) Pdf for late cons. with $\sigma_\varepsilon = 1$ and $\sigma_\delta = 0.5$.

Figure 6: Distributional of effects for log-normally distributed preferences shocks. With $\sigma_\varepsilon = 0.5$ and $\sigma_\delta = 1$ in the left-column, optimal inflation is -0.603% and the optimal nominal rate is 1.310%. With $\sigma_\varepsilon = 0.5$ and $\sigma_\delta = 1$ in the right column, optimal inflation is -0.847% and the optimal nominal rate is 0% as the ZLB binds. Absent a ZLB, in the latter case optimal inflation would be -1.599% and the optimal rate -1.879%.

31