

What Lower Bound? Monetary Policy with Negative Interest Rates*

Matthew Rognlie[†]

Job Market Paper

November 23, 2015

Click [here](#) for most recent version

Abstract

Policymakers and academics have long maintained that nominal interest rates face a zero lower bound (ZLB), which can only be breached through major institutional changes like the elimination or taxation of paper currency. Recently, several central banks have set interest rates as low as -0.75% without any such changes, suggesting that, in practice, money demand remains finite even at negative nominal rates. I study optimal monetary policy in this new environment, exploring the central trade-off: negative rates help stabilize aggregate demand, but at the cost of an inefficient subsidy to paper currency. Near 0%, the first side of this tradeoff dominates, and negative rates are generically optimal whenever output averages below its efficient level. In a benchmark scenario, breaking the ZLB with negative rates is sufficient to undo most welfare losses relative to the first best. More generally, the gains from negative rates depend inversely on the level and elasticity of currency demand. Credible commitment by the central bank is essential to implementing optimal policy, which backloads the most negative rates. My results imply that the option to set negative nominal rates lowers the optimal long-run inflation target, and that abolishing paper currency is only optimal when currency demand is highly elastic.

*I am deeply grateful to my advisors Iván Werning, Daron Acemoglu, and Alp Simsek for their continual guidance and support, and to Adrien Auclert for invaluable advice on all aspects of this project. I also thank Alex Bartik, Vivek Bhattacharya, Nicolas Caramp, Yan Ji, Ernest Liu, Miles Kimball, Ben Moll, Emi Nakamura, Christina Patterson, Jón Steinsson, and Ludwig Straub for helpful comments. Thanks to the NSF Graduate Research Fellowship for financial support. All errors are my own.

[†]MIT Department of Economics. Email: mrognlie@mit.edu

1 Introduction

Can nominal interest rates go below zero? In the past two decades, the zero lower bound (ZLB) on nominal rates has emerged as one of the great challenges of macroeconomic policy. First encountered by Japan in the mid-1990s it has, since 2008, become a constraint for central banks around the world, including the Federal Reserve and the European Central Bank. These central banks' perceived inability to push short-term nominal rates below zero has led them to experiment with unconventional policies—including large-scale asset purchases and forward guidance—in order to try to achieve their targets for inflation and economic activity, with incomplete success.

Events in the past year, however, have called into question whether zero really is a meaningful barrier. Central banks in Switzerland, Denmark, and Sweden have targeted negative nominal rates with apparent success, and without any major changes to their monetary frameworks. Policymakers at other major central banks, including the Federal Reserve and the ECB, have recently alluded to the possibility of following suit.¹²

In this paper, I consider policy in this new environment, where negative nominal rates are a viable option. I argue that these negative rates, though feasible, are not costless: they effectively subsidize paper currency, which now receives a nominal return (zero) that exceeds the return on other short-term assets. Policymakers face a tradeoff between the burden from this subsidy and the benefits from greater downward flexibility in setting rates. This paper studies the tradeoff in depth, exploring the optimal timing and magnitude of negative rates, as well as their interaction with other policy tools.

The traditional rationale behind the zero lower bound is that the existence of money, paying a zero nominal return, rules out negative interest rates in equilibrium: it would be preferable to hoard money rather than lend at a lower rate. This view was famously articulated by Hicks (1937):

If the costs of holding money can be neglected, it will always be profitable to hold money rather than lend it out, if the rate of interest is not greater than

¹In response to a question while testifying before Congress on November 4, 2015, Federal Reserve Chair Janet Yellen stated that if more stimulative policy were needed, “then potentially anything, including negative interest rates, would be on the table.” (Yellen 2015.) In a press conference on October 22, 2015, ECB President Mario Draghi stated: “We’ve decided a year ago that [the negative rate on the deposit facility] would be the lower bound, then we’ve seen the experience of countries and now we are thinking about [lowering the deposit rate further].” (Draghi 2015.)

²By some measures, the ECB has already implemented negative rates, since the Eurosystem deposit facility (to which Draghi 2015 alluded) pays -0.20%. Excess reserves earn this rate, which has been transmitted to bond markets: as of November 20, 2015, short-term government bond yields are negative in a majority of Euro Area countries. Since the ECB’s benchmark rate officially remains 0.05%, however, I am not classifying it with Switzerland, Denmark, and Sweden.

zero. Consequently the rate of interest must always be positive.

Of course, this discussion presumes that money pays a zero nominal return, which is not true of all assets that are sometimes labeled “money”. Bank deposits can pay positive interest or charge the equivalent of negative interest through fees; similarly, central banks are free to set the interest rate on the reserves that banks hold with them. The one form of money that *is* constrained to pay a zero nominal return is paper currency—which in this paper I will abbreviate as “cash”. The traditional argument for a zero lower bound, therefore, boils down to the claim that cash yielding zero is preferable to a bond or deposit yielding less—and that any attempt to push interest rates below zero will lead to an explosion in the demand for cash.

In light of recent experience, I argue that this claim is false: contrary to Hicks’s assumption, the costs of holding cash cannot be neglected. I write a simple model of cash use in which these costs make it possible for interest rates to become negative. These very same costs, however, make negative rates an imperfect policy tool: since cash pays a higher return, households hold it even when the marginal costs exceed the benefits. The distortionary subsidy to cash creates a deadweight loss. This is the other side of a mainstay of monetary economics, the Friedman rule, which states that nominal rates should be optimally set at zero, and that any deviation from zero creates a welfare loss. The Friedman rule has traditionally been used to argue that positive nominal rates are suboptimal, but I argue the same logic captures the loss from setting negative rates—and this loss may be of far greater magnitude, since cash demand and the resulting distortion can grow unboundedly as rates become more negative.

I integrate this specification for cash demand into a continuous-time New Keynesian model. With perfectly sticky prices, nominal interest rates determine real interest rates, which in turn shape the path of consumption and aggregate output. The challenge for policy is to trade off two competing objectives—first, the need to set the nominal interest rate to avoid departing too far from the equilibrium or “natural” real interest rate, determined by the fundamentals of the economy; and second, the desire to limit losses in departing from the Friedman rule. Optimal policy navigates these two objectives by smoothing interest rates relative to the natural rate, to an extent determined by the level and elasticity of cash demand. These results echo earlier results featuring money in a New Keynesian model, particularly [Woodford \(2003b\)](#), though my continuous-time framework provides a fresh look at several of these previous insights, in addition to a number of novel findings.

I then provide a reinterpretation of the ZLB in this new framework. Under my standard specification of cash demand, motivated by the evidence from countries setting neg-

ative rates, the ZLB is not a true constraint on policy, though it is possible to consider optimal policy when it is imposed as an exogenous additional constraint. I argue that this optimal ZLB-constrained policy is equivalent to optimal policy in a counterfactual environment, where the net marginal utility from cash is equal to zero for any amount of cash above a satiation point. Central banks that act as if constrained by a ZLB, therefore, could be motivated by this counterfactual view of cash demand.

In the baseline case where cash demand does not explode at zero, I show that it is generally optimal to use negative rates. The key observation is that the *zero bound is also the optimal level of interest rates prescribed by the Friedman rule*. In the neighborhood of this optimum, any deviation leads to only second-order welfare losses, which are overwhelmed by any first-order gains from shaping aggregate demand. These first-order gains exist if, over any interval that begins at the start of the planning horizon, the economy will on average (in a sense that I will make precise) be in recession. Far from being a hard constraint on rates, therefore, zero is a threshold that a central bank *should* go beyond whenever needed to boost economic activity.

With this in mind, I revisit the standard “liquidity trap” scenario that has been used in the literature to study the ZLB. As in [Eggertsson and Woodford \(2003\)](#) and [Werning \(2011\)](#), I suppose that the natural interest rate is temporarily below zero, making it impossible for a ZLB-constrained central bank to match with its usual inflation target of zero. With negative rates as a tool, it is possible to come much closer to the optimal level of output, but this response is mitigated by the desire to avoid a large deadweight loss from subsidizing cash.

In the simplest case, I assume that the natural rate reverts to zero after the “trap” is over, and that it is impossible to commit to time-inconsistent policies following the trap. Solving the model for optimal policy with negative rates, the key insight that emerges is that the most negative rates should be *backloaded*. Relative to the cost of violating the Friedman rule, which does not vary over time, negative rates have the greatest power to lift consumption near the end of the trap. The optimal path of rates during the trap, in fact, starts at zero and monotonically declines, always staying above the natural rate. If full commitment to time-inconsistent policies is allowed, it becomes optimal to keep rates negative even after the trap has ended and the natural rate is no longer below zero—taking backloading one step further, and effectively employing forward guidance with negative rates.

Quantitatively, I compare the outcomes of ZLB-constrained and unconstrained policy using my benchmark calibration. Freeing the policymaker to set negative rates closes over 94% of the gap between equilibrium utility and the first best. A second-order ap-

proximation to utility, which is extremely accurate for the benchmark calibration, offers insight into the forces governing the welfare improvement: negative rates offer greater gains when the trap is long and the welfare costs of recession are high, but they are less potent when the level and elasticity of cash demand are large.

I also consider the case where, following the trap, the natural rate reverts to a positive level. This allows a ZLB-constrained central bank to engage in forward guidance, continuing to set rates at zero after the trap. In this environment, I show that the optimal ZLB-constrained and unconstrained policies produce qualitatively similar results: they both use forward guidance to create a boom after the trap, which limits the size of the recession during the trap. ZLB-constrained policy, however, produces far larger swings in output relative to the first-best level, in both the positive and negative directions. With negative rates, it is possible to smooth these fluctuations by more closely matching the swings in the natural rate.

I next relax the assumption of absolute price stickiness, assuming instead that prices are rigid around some trend inflation rate, which can be chosen by the central bank. This allows me to evaluate the common argument that higher trend inflation is optimal because it allows monetary policy to achieve negative *real* rates despite the zero lower bound (see, for instance, [Blanchard, Dell’Ariccia and Mauro 2010](#)). I show that once negative *nominal* rates are available as a policy tool, the optimal trend inflation rate falls, as inflation becomes less important for this purpose. The ability to act as a substitute for inflation may add to negative nominal rates’ popular appeal.

Finally, I consider supplemental policies that limit the availability of cash. The most extreme such policy is the abolition of cash, frequently discussed in conjunction with the zero lower bound (see, for instance, [Rogoff 2014](#)). This policy is equivalent of imposing an infinite tax on cash, and in that light can be evaluated using my framework: the crucial question is whether the distortion from subsidizing cash when rates are negative is large enough to exceed the cost from eliminating cash altogether. I argue that this depends on the extent of asymmetry in the demand for cash with respect to interest rates, and I describe a simple sufficient condition that makes it optimal for policymakers to retain cash. As an empirical matter, I conclude that it is probably not optimal to abolish cash—but this does depend on facts that are not yet settled, including the extent to which cash demand rises when rates fall below levels that have thus far been encountered. One possible intermediate step is the abolition of larger cash denominations, which have lesser holding costs and are demanded more elastically than small denominations. In an extension of my cash demand framework to multiple denominations, I show that it is always optimal to eliminate these large denominations first.

Related literature. This paper relates closely to several literatures.

The literature on negative nominal interest rates has seen considerable growth in the past decade. In contrast to my paper, this literature generally makes the same presumption as [Hicks \(1937\)](#): it assumes that cash demand becomes infinite once cash offers a higher pecuniary return than other assets. When this is true, major institutional changes are required before negative rates are possible. [Buiter \(2009\)](#) summarizes the options available: cash can either be abolished or made to pay a negative nominal return. The former option, the abolition of cash, has been explored in detail by [Rogoff \(2014\)](#). The latter option, a negative nominal return, can be implemented either by finding some way to directly tax cash holdings, or by decoupling cash from the economy's numeraire.

The idea of taxing cash originated with [Gesell \(1916\)](#), who proposed physically stamping cash as proof that tax has been paid. At the time, this proposal was influential enough to be cited by [Keynes \(1936\)](#). More recently, similar ideas have been explored by [Goodfriend \(2000\)](#), who proposes including a magnetic strip in each bill to keep track of taxes due; by [Buiter and Panigirtzoglou \(2001, 2003\)](#), who integrate a tax on cash into a dynamic New Keynesian model; and more whimsically by [Mankiw \(2009\)](#), who suggests that central banks hold a lottery to invalidate cash with serial numbers containing certain digits.

The idea of decoupling cash from the numeraire originated with [Eisler \(1932\)](#), who envisioned a floating exchange rate between cash and money in the banking system, with the latter as the numeraire. This floating rate makes it possible to implement negative nominal interest rates in terms of the numeraire, even as cash continues to pay a zero nominal rate in cash terms, by engineering a gradual relative depreciation of cash. More recently, [Buiter \(2007\)](#) has resurrected this approach, and [Agarwal and Kimball \(2015\)](#) provide a detailed guide to its implementation and possible advantages.

Each of these approaches makes negative rates unambiguously feasible, but at the cost of major changes to the monetary system: either abolishing cash, taxing it via a tracking technology, or removing its status as numeraire. My paper, by contrast, primarily focuses on the consequences of negative rates within the existing system, as they are currently being implemented in Switzerland, Denmark, and Sweden. For policymakers who are not yet ready or politically able to make major reforms to the monetary system, the paper provides a framework for understanding negative rates; by clarifying the costs of negative rates within the existing system, it also provides a basis for comparison to the costs of additional reforms.

Some very recent work explores the practical side of the negative rate policies now in effect. [Jackson \(2015\)](#) provides an overview of recent international experience with

negative policy rates, and [Jensen and Spange \(2015\)](#) discuss the pass-through to financial markets and impact on cash demand from negative rates in Denmark. [Humphrey \(2015\)](#) evaluates ways to limit cash demand in response to negative rates.

This paper is also closely related to the modern zero lower bound literature, which began with [Fuhrer and Madigan \(1997\)](#) and [Krugman \(1998\)](#) and subsequently produced a flurry of papers. I revisit the “trap” scenario contemplated in much of this work— notably [Eggertsson and Woodford \(2003\)](#) and [Werning \(2011\)](#)—in which the natural rate of interest is temporarily negative and cannot be matched by a central bank subject to the zero lower bound. One particularly important theme—both in the ZLB literature and in this paper—is forward guidance, which is the focus of a large emerging body of work that includes [Levin, López-Salido, Nelson and Yun \(2010\)](#), [Campbell, Evans, Fisher and Justiniano \(2012\)](#), [Del Negro, Giannoni and Patterson \(2012\)](#), and [McKay, Nakamura and Steinsson \(2015\)](#). I also consider the interaction of the ZLB, negative rates, and the optimal rate of trend inflation, which has been covered by [Coibion, Gorodnichenko and Wieland \(2012\)](#), [Williams \(2009\)](#), [Blanchard et al. \(2010\)](#), and [Ball \(2013\)](#), among others.

At its core, this paper uses the canonical New Keynesian framework laid out by [Woodford \(2003a\)](#) and [Galí \(2008\)](#), but since price dynamics are not a focus, for simplicity I replace pricesetting à la [Calvo \(1983\)](#) with the assumption of fully rigid prices. I follow [Werning \(2011\)](#) by using a continuous-time version of the model, which permits a sharper characterization of both cash demand and the liquidity trap. In adding cash to the model, the paper is reminiscent of much of the New Keynesian literature with money, including [Khan, King and Wolman \(2003\)](#), [Schmitt-Grohé and Uribe \(2004b\)](#), and [Siu \(2004\)](#). It perhaps comes closest to [Woodford \(1999\)](#) and [Woodford \(2003b\)](#), which also find that smoothing interest rates is optimal in the model with money—though this smoothing takes a particularly stark form in the continuous-time framework I provide.

This paper is deeply connected with the literature on the Friedman rule, since it emphasizes deviation from the Friedman rule—in a novel direction—as the reason why negative rates are costly. This literature began eponymously with [Friedman \(1969\)](#), and was exhaustively surveyed by [Woodford \(1990\)](#). The seminal piece opposing the Friedman rule was [Phelps \(1973\)](#), which argued that a government minimizing the overall distortionary burden of taxation should rely in part on the inflation tax as a source of revenue; much subsequent work has investigated this claim. The key intuition for why the Friedman rule may be optimal, even when alternative sources of government revenue are distortionary, is that money is effectively an intermediate good, facilitating transactions: versions of this idea are in [Kimbrough \(1986\)](#), [Chari, Christiano and Kehoe \(1996\)](#), and [Correia and Teles \(1996\)](#).

As [Schmitt-Grohé and Uribe \(2004a\)](#) and others point out, however, positive nominal interest rates may be optimal as an indirect tax on monopoly profits. Inversely, [da Costa and Werning \(2008\)](#) find that negative rates may be preferable due to the complementarity of money and work effort, although they interpret this finding as showing that the Friedman rule is optimal as a corner solution, under the presumption that negative rates are not feasible. In this paper I sidestep much of the complexity in the literature by taking a simple model where the government has a lump-sum tax available, and the Friedman rule is therefore unambiguously optimal absent nominal rigidities. If, in a richer model, the optimum nominal rate is positive or negative instead, much of the analysis in the paper still holds, except that zero no longer has the same special status as a benchmark.

2 Model and assumptions on cash

2.1 Zero lower bound and cash demand


Why should zero be a lower bound on nominal interest rates? Traditionally, the literature has held that negative rates imply *infinite money demand*, which is inconsistent with equilibrium.

For instance, the influential early contribution by [Krugman \(1998\)](#) models money demand using a cash-in-advance constraint. Once this constraint no longer binds, the nominal interest rate falls to zero—but it cannot fall any further, because individuals prefer holding money that pays zero to lending at a lower rate. Similarly, [Eggertsson and Woodford \(2003\)](#) posit that real money balances enter into the utility function, and that marginal utility from money is exactly zero once balances exceed some satiation level. Again, rates can fall to zero, but no further: once the marginal utility from money is zero, holding wealth in the form of money is indistinguishable from holding it in the form of bonds, and if bonds pay a lower rate there will be an unbounded shift to money.

Many traditional models of money demand similarly embed this zero lower bound. In the Baumol-Tobin model ([Baumol 1952](#) and [Tobin 1956](#)), for instance, the interest elasticity of real money demand is $-1/2$. As the nominal interest rate i approaches 0, money demand $M/P \propto i^{-1/2}$ approaches infinity. The same happens in any model where the interest elasticity of money demand is bounded away from zero in the neighborhood of $i = 0$, including many of the specifications in the traditional empirical money demand literature, which assume a constant interest elasticity—see for instance, [Meltzer \(1963\)](#).³

³This feature has played a prominent role in welfare calculations: under specifications assuming a constant interest elasticity, [Lucas \(2000\)](#) finds that the costs of moderate departures from the Friedman rule are

Figure 1: Different views of money demand



In contrast, other empirical studies of money demand, dating back to [Cagan \(1956\)](#), assume a constant interest *semi*elasticity—see, for instance, [Ball \(2001\)](#) and [Ireland \(2009\)](#). With this specification, money demand does not explode as $i \rightarrow 0$; indeed, if extended to cover negative i , the specification continues to imply finite money demand.

Figure 1 displays three possible shapes for the demand curve for money with respect to interest rates. The first is a curve featuring a constant elasticity of demand, such that money demand explodes as $i \rightarrow 0$. The second is a curve featuring a constant semielasticity, such that money demand remains finite even as i becomes negative. The third is a modification of the second curve along the lines of [Eggertsson and Woodford \(2003\)](#) and much of the other zero lower bound literature, where money demand is unbounded at $i = 0$ even though it remains finite in the limit $i \rightarrow 0$. The first and third cases feature a zero lower bound, while the second does not.

As argued by [Ireland \(2009\)](#), modern experience with low nominal interest rates contradicts the first case in figure 1: money demand does not explode in inverse proportion to rates near zero. It has been an open question, however, whether money demand more closely resembles the second or third case: does it smoothly expand as rates dip below zero, or does it abruptly become infinite at zero? The modern zero lower bound literature has generally assumed the latter, either implicitly (when the bound is imposed as an ad-hoc constraint) or explicitly (when the bound is microfounded using money demand).

Cash vs. other forms of money. At this point, it is useful to distinguish between different forms of “money”. Inside money, consisting of bank deposits and other liquid liabilities of private intermediaries, is not subject in principle to any zero lower bound:

significant, while under specifications assuming a constant interest *semi*elasticity, the costs are much smaller. Roughly speaking, when assuming a constant elasticity, the explosion in money demand as $i \rightarrow 0$ means that the deadweight loss from setting $i > 0$ is much larger.

it can pay negative interest as well as positive interest, and sometimes does so implicitly through account fees. There may be frictions in adjusting to negative rates, but these are highly specific to the institution and regulatory regime, and are not central to the zero lower bound as a general notion.⁴

Most central bank liabilities can also pay negative interest: for instance, a central bank can charge banks who hold reserve balances with it. In fact, this is exactly what central banks that implement negative rates do. In a world where all liabilities of the central bank could pay negative interest, there would be no hint of a lower bound.⁵

The difficulty is that one central bank liability, paper currency, has a nominal return that is technologically constrained to be zero.⁶ If nominal interest rates on other assets are negative, the concern is that demand for paper currency—which I abbreviate as *cash*—will become infinite. If this is true, zero does serve as an effective lower bound on interest rates. Interpreting figure 1 as depicting alternative possible shapes for the *cash* demand function, the crucial question is therefore whether the second or third possibility is more accurate.⁷

New evidence: successful implementation of negative rates. In the last year, three central banks have set their primary rate targets at unprecedentedly negative levels: both Switzerland and Denmark at -0.75%, and Sweden at -0.30%. This is depicted in figure 2.

Implementation has been successful: in line with the targets, market short-term nominal interest rates have fallen well into negative territory.⁸ Indeed, expectations that the negative rate policy will be continued in Switzerland are sufficiently strong that even the 10-year Swiss government bond yield has been negative for much of 2015.

This novel policy experiment provides a useful test of whether negative market interest rates are consistent with bounded cash demand. Thus far, the verdict has been clear: not only has cash demand remained finite, but its response to negative rates has been quite mild. For Switzerland, where monthly data on banknotes outstanding is publicly

⁴For instance, [McAndrews \(2015\)](#) mentions the dilemma of retail and Treasury-only money market mutual funds in the US, which as currently structured would “break the buck” and be forced to disband in an environment with negative rates. Money market mutual funds elsewhere, however, have successfully adapted to negative rates.


⁵In fact, in the canonical treatment of the New Keynesian model in [Woodford \(2003a, p. 68\)](#), the lower bound on interest rates i_t is derived to be the interest i_t^m paid on money by the central bank.

⁶As discussed in section 1, there have been proposals to remove this constraint through changes in technology: for instance, the idea of [Goodfriend \(2000\)](#) to embed a magnetic strip in paper currency that tracks taxes paid on it.

⁷As before, the argument in [Ireland \(2009\)](#) rules out the first possibility: cash demand appears not to explode as nominal interest rates asymptote to zero.

⁸For example, as of November 20, 2015, one-month government bond yields are -0.89% in Switzerland, -0.70% in Denmark, -0.39% in Sweden.

Figure 2: Target interest rates in Switzerland, Denmark, and Sweden



(Target rates are 3-month Libor CHF for Switzerland, Danmarks Nationalbank certificates of deposit rate for Denmark, and Riksbank repo rate for Sweden.)

available, figure 3 shows the total value of cash in circulation against the path of the Swiss target rate. Following the decline to -0.75% at the beginning of 2015, there has been little perceptible break in the trend. For Denmark, [Jensen and Spange \(2015\)](#) have similarly noted little increase in cash demand.


Among the possibilities depicted in figure 1, therefore, the empirical cash demand schedule appears to most closely resemble the middle case, with no discontinuity at $i = 0$. My analysis will build upon this observation.

Cost of negative rates: a subsidy to cash. If negative rates do not lead to infinite cash demand, and are therefore feasible, is there potentially any reason to avoid them? Yes.

To build intuition, it is useful to consider an extreme case: suppose that setting $i = -1\%$ leads cash demand to increase by a factor of 100. Since this is not quite an infinite increase, it is still feasible in equilibrium, but there is a considerable cost. If the central bank holds short-term bonds on the asset side of its balance sheet, for instance, then its cash liabilities will pay 0% while its assets earn -1%. The effective 1% subsidy to cash, relative to the market interest rate, will cost the central bank greatly—and on a massively expanded base of cash, leading to annual losses equal in magnitude to the entire prior level of cash in circulation.

The public will both benefit from this subsidy and ultimately pay the cost of providing it, via a larger tax burden. Under certain assumptions, which I will use in this paper, this

Figure 3: Target interest rate vs. cash in circulation, Switzerland



cost and benefit will cancel to first order in i .⁹ But there will be a second-order net cost—which, if cash demand increases by a factor of 100, will be quite large—since the subsidy leads the public to demand more cash than is socially optimal.


The intuition for this second-order cost is similar to that for any subsidy. If the public decides to hold more cash when the subsidy is 1% than when it is 0%, there must be a net nonpecuniary cost to the marginal unit of cash: the inconvenience of holding wealth in the form of cash exceeds, at the margin, the liquidity benefits. Once the public pays for the subsidy through taxes, all that remains is this inefficiently high level of cash demand—with, perhaps, a large drain of resources going to the manufacturers of safes.

This is the inverse of the traditional story, in which *positive* nominal interest rates act as a tax, leading the public to demand inefficiently little cash. The idea that the optimal level of nominal interest rates is zero—with neither a tax nor a subsidy on cash—is called the *Friedman rule*, in recognition of [Friedman \(1969\)](#). Generally, only one side of the Friedman rule has been discussed: prior to recent events, negative rates were not viewed as a feasible option, and it made little sense to talk about the inefficiency from too much cash demand.

But this inefficiency, in fact, is at the center of the policy tradeoff with negative rates. The prior consensus that negative rates were infeasible, due to an explosion in cash demand at zero, can be interpreted as just an extreme form of the same point: as cash demand becomes more and more elastic with respect to negative interest rates, the inefficiency increases until negative rates become infinitely costly in the limit. More generally,

⁹This first-order cancellation arises because 0% is the Friedman rule optimum. More generally, whether or not the Friedman rule holds depends on assumptions about distributive effects, fiscal instruments available to the government, and so on. See the discussion of the literature in section 1.

Figure 4: Utility from cash



it is plausible that the cost of deviating from the Friedman rule is much more severe on the negative side than on the traditional, positive one, because the rise in cash demand is potentially unbounded.

Interpretation in a simple model of cash demand. In the section 2.2, I will integrate cash demand into a simple infinite-horizon New Keynesian model by including concave flow utility $v(m(t))$ from real cash balances $m(t)$ into household preferences (2). The opportunity cost of holding wealth in the form of cash rather than bonds is the nominal interest rate i on bonds, and real cash demand $M^d(i, c)$ as a function of nominal interest rates i and consumption c is given by the optimality condition

$$v'(M^d(i, c)) = iu'(c) \quad (1)$$

where $u'(c)$ is marginal utility from consumption.

If there is finite cash demand when $i = 0$, its level $m^* = M^d(0, c)$ is given by $v'(m^*) = 0$; and if cash demand continues to be finite for negative i as well, then v' must be strictly declining at m^* . It follows that $v(m^*)$ is a global maximum of v . This is depicted in figure 4.

Positive i corresponds to $v' > 0$ and to an inefficiently low level of cash demand $m < m^*$, while negative i corresponds to $v' < 0$ and an inefficiently high level of cash demand $m > m^*$. The utility shortfall relative to $v(m^*)$ can be obtained in consumption terms by integrating marginal utility v' , which according to (1) is proportional to the cash

Figure 5: Loss from violating Friedman rule: integrating under the demand curve





Figure 6: Approximate cost of deviating from Friedman rule: Harberger triangle



demand curve times $u'(c^*)$. This is visualized in figure 5, which shows as shaded areas the loss from setting positive i (the standard case) and the loss from setting negative i (the new case). Figure 6 shows the second-order Harberger triangle approximation to the loss from negative i , which depends on the level of cash demand at $i = 0$, $M^d(0, c) = m^*$, and crucially the local semielasticity of cash demand $\partial \log M^d(0^-, c) / \partial i$. When the semielasticity is higher and cash demand grows more rapidly as i falls below 0, the loss is more severe—and in the limit as the semielasticity becomes infinite, the cost becomes infinite as well, leading in effect to a zero lower bound.

As figures 5 and 6 illustrate, therefore, the cost of setting negative rates fits squarely into the standard microeconomic analysis of distortions. With this view in mind, I now turn to a dynamic framework, studying how this cost trades off against the other objec-

tives of monetary policy.

2.2 Benchmark model

In this section, I describe the basic infinite-horizon continuous-time model that will be used for the analysis, with a particular focus on the specification for cash demand.

Households. Households have the objective

$$U(\{c(t), n(t), M(t), P(t)\}) = \int_0^\infty e^{-\int_0^s \rho(u) du} \cdot \left(u(c(t)) - \chi(n(t)) + v\left(\frac{M(t)}{P(t)}\right) \right) dt \quad (2)$$

where $c(t)$ is consumption, $n(t)$ is labor supplied, $M(t)$ is the level of cash held by the household, $P(t)$ is the price of the consumption good, and $\rho(t)$ is the time-varying rate of time preference. In general, I will assume that both utility from consumption u and disutility from labor χ are isoelastic, denoting the elasticity of intertemporal substitution in consumption by σ and the Frisch elasticity of intertemporal substitution in labor by ψ :

$$u(c) = \frac{c^{1-\sigma^{-1}} - 1}{1 - \sigma^{-1}} \quad \chi(n) = \gamma \frac{n^{1+\psi^{-1}} - 1}{1 + \psi^{-1}} \quad (3)$$

The assumption that v is separable from the rest of the utility function is in line with much of the New Keynesian literature featuring money in the utility function. Here, the assumption is made primarily for analytical convenience, but calibrated studies have generally found that (for instance) ignoring the possible complementarity between consumption and money does not have significant quantitative ramifications.

Households have access to two stores of value, cash M and bonds B , and face the nominal flow budget constraint

$$\dot{M}(t) + \dot{B}(t) + P(t)c(t) = i(t)B(t) + W(t)n(t) + \Pi(t) + T(t) \quad (4)$$

where $i(t)$ is the nominal interest rate paid on bonds and $W(t)$ is the nominal wage paid for labor by firms. $\Pi(t)$ is firms' profit, and $T(t)$ is net lump-sum transfers by the government, both of which will be specified later. Cash is assumed to pay zero interest in (4). As discussed in section 2.1, other liabilities of the central bank—such as electronic reserves—can pay nonzero interest. Here am abstracting away from the difference between these liabilities and bonds $B(t)$, since both are short-term interest-paying liabilities of the government.

Dividing by $P(t)$, the real flow budget constraint becomes

$$\dot{m}(t) + \dot{b}(t) + c(t) = r(t)b(t) + w(t)n(t) + \frac{\Pi(t)}{P(t)} + \frac{T(t)}{P(t)} \quad (5)$$

where $m(t)$ and $b(t)$ are real cash and bonds, respectively, $w(t)$ is the real wage rate, and $r(t) \equiv i(t) - \dot{P}(t)/P(t)$ is the real interest rate, and $\Pi(t)/P(t)$ and $T(t)/P(t)$ are real transfers. Integrating (5) and imposing a no-Ponzi condition gives the infinite-horizon version of the budget constraint:

$$\int_0^\infty e^{-\int_0^t r(s)ds} (c(t) + i(t)m(t))dt = \int_0^\infty e^{-\int_0^t r(s)ds} \left(w(t)n(t) + \frac{\Pi(t)}{P(t)} + \frac{T(t)}{P(t)} \right) dt \quad (6)$$

Given paths $\{i(t), r(t), w(t)\}$ for prices and $\{\Pi(t)/P(t), T(t)/P(t)\}$ for transfers, the household's problem is to choose $\{c(t), n(t), m(t)\}$ to maximize (2) subject to (6).

Firms. A continuum of monopolistically competitive firms $j \in [0, 1]$ produce intermediate goods using labor as the only input, subject to a potentially time-varying productivity parameter $A(t)$:

$$y_j(t) = A(t)f(n_j(t)) \quad (7)$$

I will also generally assume that f is isoelastic, with $1 - \alpha$ as the constant elasticity of output with respect to labor n :

$$f(n) = \frac{n^{1-\alpha}}{1-\alpha} \quad (8)$$

These firms' output is aggregated into production $y(t)$ of the final consumption good by a perfectly competitive final good sector, which operates a final constant elasticity of substitution production technology $y(t) = \left(\int_0^1 y_j(t)^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}}$. Demand by this sector for firm j 's output is $y_j(t) = (P_j(t)/P(t))^{-\varepsilon} y(t)$, where $P(t) = \left(\int_0^1 P_j(t)^{1-\varepsilon} dj \right)^{1/(1-\varepsilon)}$ is the aggregate price index. Market clearing for labor requires that firms' total demand for labor equals household labor supply: $n(t) = \int_0^1 n_j(y) dj$.

I consider two possible specifications of firms' pricesetting. In the benchmark *flexible price* case, they choose prices at each t to maximize profits

$$\Pi_j(t) = \max_{P_j(t)} P_j(t) \left(\frac{P_j(t)}{P(t)} \right)^{-\varepsilon} y(t) - C(y_j(t); t) \quad (9)$$

where $C(y; t) \equiv f^{-1}(y/A(t))W(t)$ is the nominal cost of producing y at time t . Profits

(9) are maximized when $P_j(t)$ is set at a markup of $\epsilon/(\epsilon - 1)$ over marginal cost:

$$P_j(t) = \frac{\epsilon}{\epsilon - 1} C_y(y; t) = \frac{\epsilon}{\epsilon - 1} \frac{W(t)}{f'(f^{-1}(y/A(t)))}$$

It follows that all firms j set the same price at time t and produce the same output, and that real wages are given by

$$w(t) = \frac{\epsilon - 1}{\epsilon} A(t) f'(n(t)) \quad (10)$$

In the *sticky price* case, by contrast, prices are rigid at $P(t) \equiv \bar{P}$ for all t . This simple assumption will create an aggregate demand management role for the monetary authority, generating the tradeoff at the heart of this paper: the distortionary costs from setting interest rates below zero, versus the benefits of bringing output closer to its optimal level.

In both cases, I assume that aggregate profits $\Pi(t) = \int_0^1 \Pi_j(t) dj$ are immediately rebated to the household, as seen earlier in (4).

Government. The government, representing both the fiscal and monetary authorities, has two liabilities, bonds $B(t)$ and cash $M(t)$. Nominal interest $i(t)B(t)$ is earned on bonds, while the nominal interest rate on cash is fixed at zero. A lump-sum transfer $T(t)$ to households, which can be positive or negative, is also available.

The government's nominal flow budget constraint is then

$$\dot{M}(t) + \dot{B}(t) = i(t)B(t) - T(t) \quad (11)$$

which, when normalized by $P(t)$ and integrated subject to a no-Ponzi condition, becomes

$$\int_0^\infty e^{-\int_0^t r(s) ds} i(t) m(t) dt = \int_0^\infty e^{-\int_0^t r(s) ds} \frac{T(t)}{P(t)} dt \quad (12)$$

which states that the net present value of real seignorage $i(t)m(t)$ must equal that of real transfers $T(t)/P(t)$ to the public.

Equilibrium. With these ingredients in place, I am now ready to define equilibrium.¹⁰

¹⁰Note that for economy of notation, this definition of flexible-price equilibrium assumes that all firms set the same price, so that there is no need to carry around the distribution of individual prices as an equilibrium object. This is true given my assumptions on firms.

Definition 2.1. A *flexible-price equilibrium* consists of quantities

$$\{c(t), n(t), y(t), M(t), \Pi(t), T(t)\}_{t=0}^{\infty}$$

and prices

$$\{i(t), W(t), P(t)\}_{t=0}^{\infty}$$

such that households optimize intertemporal utility (2) subject to (4), firms optimize profits (9), the government satisfies its budget constraint (11), and goods, factor, and asset markets all clear. In a *sticky-price equilibrium*, profit optimization is replaced by a sticky-price constraint $P_j(t) = \bar{P}$.

Natural rate. The real interest rate achieved in flexible-price equilibrium—which is uniquely pinned down by fundamentals $\{A(t), \rho(t)\}$ —will prove useful as a benchmark for sticky-price equilibrium as well. Following common usage, I call it the *natural rate*.

Lemma 2.2. In flexible-price equilibrium, $c(t)$, $y(t)$, $n(t)$, and $w(t)$ are uniquely determined by the two equations

$$\begin{aligned} \frac{v'(n(t))}{u'(c(t))} &= w(t) = \frac{\epsilon - 1}{\epsilon} A(t) f'(n(t)) \\ c(t) &= y(t) = A(t) f(n(t)) \end{aligned}$$

Assuming isoelastic preferences (3) and technology (8), the equilibrium real interest rate, which I denote by $r^n(t)$, is then given by

$$r^n(t) = \rho(t) + \frac{1 + \psi}{\sigma + \psi + (\sigma - 1)\psi\alpha} \frac{\dot{A}(t)}{A(t)} \quad (13)$$

Definition 2.3. The *natural rate* $r^n(t)$ is the flexible-price equilibrium real interest rate in (13).

Note that the natural rate reflects both the rate of pure time preference $\rho(t)$ and the rate of productivity growth $\dot{A}(t)/A(t)$.

3 Optimal policy and negative rates

In this section, I set up the optimal policy problem and discuss the implications for negative rates.

Characterizing equilibria. The equilibrium concept in definition 2.1 is such that the paths for real quantities $\{c(t), n(t), y(t), m(t)\}$ and prices $\{w(t), i(t)\}$ are uniquely characterized by a much smaller set of paths.

For flexible-price equilibrium, lemma 2.2 already shows that $c(t)$, $n(t)$, $y(t)$, and $m(t)$ are determined by exogenous fundamentals. Given the nominal interest rate $i(t)$, the quantity of cash is then given by $m(t) = M^d(i(t), c(t))$.

In contrast, the real quantities and prices in sticky-price equilibrium are not pinned down by nominal interest rates alone. Instead, conditional on nominal interest rates $\{i(t)\}$ there is a single degree of indeterminacy in the consumption path. This indeterminacy can be indexed by the level of consumption at some selected time, which I choose to be $t = 0$ for simplicity.

With this in mind, given any path $\{i(t)\}$ for the nominal interest rate and the time-0 level of consumption $c(0)$, consumption at any time t can be obtained by integrating the household's consumption Euler equation $\dot{c}(t)/c(t) = \sigma(i(t) - \rho(t))$

$$\log c(t) = \log c(0) + \int_0^t \sigma(i(s) - \rho(s)) ds \quad (14)$$

With $c(t)$ known, output $y(t) = c(t)$ and labor input $n(t) = f^{-1}(y(t)/A(t))$ are given by market clearing and the production function. The quantity of cash is given by $m(t) = M^d(i(t), c(t))$.

The following proposition summarizes these observations.

Proposition 3.1. *Given any path $\{i(t)\}_{t=0}^{\infty}$ for nominal interest rates, real quantities*

$$\{c(t), n(t), y(t), m(t)\}_{t=0}^{\infty}$$

and prices

$$\{w(t), i(t)\}_{t=0}^{\infty}$$

are uniquely determined in flexible-price equilibrium. Additionally, given the level $c(0)$ of consumption at time 0, these real quantities and prices are uniquely determined in sticky-price equilibrium as well.

By offering a straightforward characterization of equilibria, proposition 3.1 simplifies the search for equilibria that are optimal from a household welfare (2) standpoint.

Optimal policy: definition and solution under flexible prices. I assume that the policymaker can freely choose between equilibria, as characterized by proposition 3.1. For

flexible-price equilibria, this is natural, since the nominal interest rate path $\{i(t)\}$ chosen by the government is sufficient to characterize the equilibrium.

For sticky-price equilibria, this is slightly less natural, since the time-0 level $c(0)$ of consumption must also be specified. To pin down a particular level for $c(0)$ —and, by extension, the entire path $\{c(t)\}$ —the government requires some additional policy tool, which I show in the Online Appendix can be a Taylor-style rule for $i(t)$ off the equilibrium path. Here, I simply assume that the policymaker is capable of choosing $c(0)$.

Definition 3.2. *Optimal policy for flexible-price equilibrium* is the choice of path $\{i(t)\}_{t=0}^{\infty}$ for nominal interest rates such that the flexible-price equilibrium characterized by proposition 3.1 maximizes household utility (2).

Optimal policy for sticky-price equilibrium is the choice of $\{i(t)\}_{t=0}^{\infty}$, along with time-0 consumption $c(0)$, such that the sticky-price equilibrium characterized by proposition 3.1 maximizes household utility (2).

Note that optimal policy by this definition is not necessarily time consistent, and that I am therefore assuming full commitment by the policymaker. I will relax this assumption in section 4.2.

The flexible-price case turns out to be extremely simple. Since consumption $c(t)$ and labor supply $n(t)$ are already pinned down by fundamentals as per lemma 2.2, the only quantity entering into household utility (2) that can be affected by policy is real cash $m(t)$. The $v(m(t))$ term is maximized under the Friedman rule $i(t) = 0$.

Proposition 3.3. *Optimal policy for flexible price equilibrium is given by $i(t) = 0$ for all t .*

With optimal policy for flexible price equilibrium characterized, I will focus on sticky price equilibrium for the remainder of the paper.

Optimal policy under sticky prices. The sticky-price case, by contrast, involves a non-trivial tradeoff: as before, the nominal interest rate affects the level of cash, but it also directly affects the path of consumption in (14). Optimal policy now requires balancing the first force against the second.

This can be formulated as an optimal control problem with state $c(t)$ and control $i(t)$. Letting $\mu(t)$ be the costate on $\log c(t)$, the current-value Hamiltonian is (dropping dependence on t for economy of notation):

$$H \equiv g(c; A) + v(M^d(i, c)) + \mu\sigma(i - \rho) \quad (15)$$

where $g(c; A) \equiv u(c) - \chi(f^{-1}(c/A))$ is defined to be the *net utility* from consumption c minus the disutility from the labor required to produce that consumption.

It follows from the maximum principle that i must maximize (15), and therefore that

$$v'(M^d(i, c)) \cdot \frac{\partial M^d(i, c)}{\partial i} + \mu\sigma = 0 \quad (16)$$

The law of motion for the costate μ is

$$\frac{\dot{\mu}}{\mu} = -c\mu^{-1} \left(g'(c; A) + v'(M^d(i, c)) \cdot \frac{\partial M^d(i, c)}{\partial c} \right) + \rho \quad (17)$$

Since I assume that the policymaker can optimally choose c at time 0, $c(0)$ is free and the corresponding costate is zero:

$$\mu(0) = 0 \quad (18)$$

Together with the Euler equation $\dot{c}/c = \sigma(i - \rho)$, conditions (16), (17), and (18) characterize optimal policy.

Simplifying optimal policy. Define $\hat{\mu} \equiv \mu/u'(c)$, which is the costate in consumption-equivalent terms. Dividing (16) by $u'(c)$, and using $v'(M^d(i, c))/u'(c) = i$, I obtain

$$\hat{\mu}\sigma = i \cdot m \cdot -\frac{\partial \log M^d}{\partial i} \quad (19)$$

Also note that $\hat{\mu}/\hat{\mu} = \dot{\mu}/\mu + \sigma^{-1}\dot{c}/c = \dot{\mu}/\mu + i - \rho$, which allows (17) to be rewritten as

$$\frac{\dot{\hat{\mu}}}{\hat{\mu}} = -\hat{\mu}^{-1} \left(c \frac{g'(c; A)}{u'(c)} + \frac{v'(M^d(i, c))}{u'(c)} \cdot \frac{\partial M^d}{\partial \log c} \right) + i \quad (20)$$

Now, let $\tau(c; A) \equiv 1 - \frac{\chi'(f^{-1}(c/A))/u'(c)}{f'(f^{-1}(c/A))}$ denote the *labor wedge*, defined as one minus the ratio of the marginal rate of substitution between leisure and consumption χ'/u' to the marginal product of labor f' . Since $g'(c; A) = u'(c) - \frac{\chi'(f^{-1}(c/A))}{f'(f^{-1}(c/A))}$, it follows that $\tau(c; A) = g'(c; A)/u'(c)$. Using this result and again $v'(M^d(i, c))/u'(c) = i$, and rearranging:

$$i\hat{\mu} - \dot{\hat{\mu}} = c\tau + i \cdot m \cdot \frac{\partial \log M^d}{\partial \log c} \quad (21)$$

It is useful to pause and interpret the terms in the above expression. The costate $\hat{\mu}$ gives the present discounted value, in terms of current consumption, from proportionately in-

creasing consumption at all future dates. This value includes two terms, visible on the right side of (21).

The first term, $c\tau$, captures the effect on net utility g from increasing consumption. If, for instance, the labor wedge τ is positive—meaning that consumption is low relative to the first best—this value is positive, because increasing consumption is beneficial. The second term, $i \cdot m \cdot \frac{\partial \log M^d}{\partial \log c}$, captures the effect on utility from cash. For instance, if i is positive—meaning that cash is low relative to the first best—then this term is positive, because the increase in cash demand induced by a rise in consumption brings the household closer to the first best.

Under the assumption in (2) of separable utility from cash, an additional simplification of (21) is possible. Differentiating $v'(M^d(i, c)) = iu'(c)$ with respect to i and $\log c$ gives

$$v''(M^d(i, c)) \cdot \frac{\partial M^d}{\partial i} = u'(c) \text{ and } v''(M^d(i, c)) \cdot \frac{\partial M^d}{\partial \log c} = -i\sigma^{-1}u'(c)$$

respectively. It follows that

$$\frac{\partial \log M^d}{\partial \log c} = -i\sigma^{-1} \frac{\partial \log M^d}{\partial i}$$

Substituting this identity into (21) and applying (19) gives

$$i\hat{\mu} - \hat{\mu} = c\tau - i\sigma^{-1} \left(i \cdot m \cdot \frac{\partial \log M^d}{\partial i} \right) = c\tau + i\hat{\mu}$$

and cancelling the $i\hat{\mu}$ on both sides, the law of motion (21) simplifies to just

$$\hat{\mu} = -c\tau$$

This cancellation reflects the equality of two forces in the optimal policy problem: discounting in the law of motion for μ , and the interaction of $\log c$ with the inefficiency in cash demand. Without separable utility from cash, this equality no longer holds, but the results are most likely robust to the presence of complementarities of plausible magnitude. Full cancellation also depends on the assumption of perfectly sticky prices: since discounting depends on the real interest rate while cash demand depends on the nominal interest rate, nonzero inflation would lead to another term, which I will derive once inflation is introduced in section 5.1.

To sum up, sticky price equilibrium under optimal policy is characterized by the fol-

lowing system:

$$\frac{\dot{c}}{c} = \sigma(i - \rho) \quad (22)$$

$$\hat{\mu}\sigma = i \cdot m \cdot -\frac{\partial \log M^d}{\partial i} \quad (23)$$

$$\dot{\hat{\mu}} = -c\tau \quad (24)$$

$$\hat{\mu}(0) = 0 \quad (25)$$

The basic tradeoff: demand management versus the Friedman rule. The great advantage of (24) is that it permits an especially simple characterization of the optimal policy tradeoff. Integrating (24) forward using the initial condition $\hat{\mu}(0) = 0$ from (18) gives

$$\hat{\mu}(t) = -\int_0^t c(s)\tau(s)ds$$

Substituting this into (23) gives

$$\sigma \int_0^t c(s)\tau(s)ds = i \cdot m \cdot \frac{\partial \log M^d}{\partial i} \quad (26)$$

which characterizes the basic optimal policy tradeoff.


(26) can be interpreted as equating the benefits and costs of an decrease in interest rates at time t . Holding consumption from time t onward constant, decreasing $i(t)$ raises the path of consumption prior to t , providing benefits of $\sigma \int_0^t c(s)\tau(s)$. If this integral is positive, which (loosely speaking) means that consumption is on average too low over the interval $[0, t]$, then the right side of (26) must be positive as well; since the interest semielasticity $\partial \log M^d / \partial i$ of cash demand is negative, this means that the nominal rate must be negative.

Smoothing and the natural rate. Optimal interest rate policy is characterized here by *smoothing*. One striking manifestation of this feature is the continuity of optimal $\{i(t)\}$.

Proposition 3.4. *Under optimal policy, $i(t)$ is continuous.*

This continuity holds regardless of any discontinuities in the fundamentals ρ or A . It emerges as a feature of the optimum because (26) trades off the benefit from reshaping the overall path of consumption—which changes continuously—against the cost of departing from the Friedman rule.

Figure 7: Optimal i relative to r^n , for varying cash demands



The costs of departing from the Friedman rule, however, depend on cash's importance in preferences (2). As cash becomes less important, the right side of (26) diminishes in magnitude, allowing interest rate policy to more closely match the natural rate.

This can be formalized by introducing the parameter α , and writing

$$v(m; \alpha) \equiv \alpha Y(\alpha^{-1} m)$$

Here, cash demand is proportional to α : $M^d(i, c; \alpha) = \alpha M^d(i, c; 1)$.

Proposition 3.5. *Under optimal policy, $i(t) \rightarrow r^n(t)$ for all t as $\alpha \rightarrow 0$.*

Together, these two propositions reflect the two sides of optimal policy: proposition 3.4 capturing the tendency toward smoothing, and proposition 3.5 showing how this tendency weakens as cash demand shrinks.

Figure 7 illustrates the contest between these two forces, by taking a simple example where the natural rate is -1% prior to $t = 0$ and 1% afterward, and considering optimal policy over several different levels of cash demand. In all cases, proposition 3.4 holds: despite the discontinuous natural rate, the optimal policy rate varies continuously. Yet the smoothing is much stronger in the $M^d(0) = 1.0$ case than the $M^d(0) = 0.01$ case—and in the latter, policy comes much closer to matching the natural rate.

ZLB-constrained optimal policy. As already discussed, zero has a special role as a benchmark for nominal interest rates: it is the optimal level of rates prescribed by the *Friedman rule*. Proposition 3.3 shows that zero rates are, in fact, optimal in the flexible-price case, where the path of interest rates only affects welfare by changing the level of cash demand. This does not carry over to the sticky-price case, and indeed figure 7 pro-

vides an example where optimal policy involves both a path for nominal rates with both strictly negative and strictly positive values.

Until recently, however, zero was significant for a different reason: it was the perceived lower bound on nominal interest rates, and central banks did not attempt to target rates beneath it. To consider the effects of this perceived bound, I will define the concept of *ZLB-constrained optimal policy*. This is identical to the original notion of optimal policy from definition 20, except that the constraint $i(t) \geq 0$ is exogenously imposed.

Definition 3.6. *ZLB-constrained optimal policy under sticky prices* is the choice of $\{i(t)\}_{t=0}^{\infty}$, along with time-0 consumption $c(0)$, such that the sticky-price equilibrium characterized by proposition 19 maximizes household utility (15), subject to the constraint that $i(t) \geq 0$ for all t .

Proposition 3.7. *ZLB-constrained optimal policy, given v , is identical to (unconstrained) optimal policy under the alternative utility function from cash*

$$\tilde{v}(m) = \begin{cases} v(m) & m \leq m^* \\ v(m^*) & m \geq m^* \end{cases} \quad (27)$$

where m^* is given by $v'(m^*) = 0$.

Proposition 25 provides one way that ZLB-constrained optimal policy can be interpreted: as optimal policy under an *alternative hypothesis about the utility from cash*. Figure 8 depicts the difference between the original v and the \tilde{v} defined in (27). The modified utility \tilde{v} flattens out at m^* , which corresponds to a zero nominal interest rate; since \tilde{v}' never becomes strictly negative, a strictly negative nominal interest rate is not possible in equilibrium.


It is natural to ask when this implicit misapprehension *matters*: when is ZLB-constrained optimal policy different from unconstrained optimal policy—or, equivalently, when does unconstrained optimal policy feature negative nominal interest rates? The next proposition provides a simple characterization in terms of the ZLB-constrained optimal policy.

Proposition 3.8 (Optimality of negative rates). *Unconstrained optimal policy features negative nominal rates if and only if under ZLB-constrained optimal policy, there is some t for which*

$$\int_0^t c^{ZLB}(s) \tau^{ZLB}(s) ds > 0 \quad (28)$$

According to proposition 3.8, negative rates are optimal when the ZLB-constrained solution features, at *any* time t , a positive (consumption-weighted) average labor wedge

Figure 8: Utility from cash: v versus \tilde{v}



between 0 and t . Loosely speaking, this means that negative rates are optimal if there is any t at which the economy has on average, to date, been in a slump rather than a boom.

To build intuition for this result, take some t where (28) holds, and consider a small downward perturbation $-\Delta i$ to the interest rate over the small interval $[t, t + \Delta t]$. The welfare impact of this perturbation working through the path of consumption is approximately

$$\left(\int_0^t c^{\text{ZLB}}(s) \tau^{\text{ZLB}}(s) ds \right) \Delta i \Delta t > 0, \quad (29)$$

which is positive and first order in Δi .

If $i(t) > 0$, this perturbation also brings us closer to the Friedman rule and is therefore unambiguously optimal—contradicting the assumption that we start at the ZLB-constrained optimal policy.

Consider alternatively the case where $i(\cdot) = 0$ on the interval $[t, t + \Delta t]$. Here we start at the Friedman rule, and the downward perturbation to interest rates moves us away from it—but the cost of this deviation is second order in Δi , at approximately

$$- \left(\int_t^{t+\Delta t} m \frac{\partial \log M^d(0, c^{\text{ZLB}}(s))}{\partial i} \cdot ds \right) \frac{1}{2} (\Delta i)^2 \Delta t < 0 \quad (30)$$

For sufficiently small Δi , the first-order benefit in (29) from increasing consumption over the interval $[0, t]$ dominates the second-order cost in (30) from deviating from the Friedman rule over the interval $[t, t + \Delta t]$. Hence a perturbation toward negative rates offers a welfare gain, and ZLB-constrained optimal policy does not coincide with unconstrained optimal policy.

The foundation of this argument is the fact that *the zero lower bound coincides with the*

Friedman rule. Pushing interest rates below zero, assuming that cash demand does not become infinite, creates a distortion—but since zero is the Friedman rule optimum, the resulting welfare loss is second-order. As long as negative rates bring the economy closer to an optimal level of output, creating first-order benefits, they are warranted.

Although proposition 3.8 is conceptually important, the condition (28) may be unwieldy to verify. The following corollary offers a much simpler test for when negative rates are optimal.

Corollary 3.9. *Unconstrained optimal policy features negative nominal rates if and only if under ZLB-constrained optimal policy, there is some t for which $i^{\text{ZLB}}(t) = 0$ and $\tau^{\text{ZLB}}(t) \neq 0$.*

This corollary demonstrates that negative rates are *generically* optimal whenever ZLB-constrained optimal policy features zero interest rates. The only exception is when consumption is at precisely its first best level, $\tau^{\text{ZLB}}(t) = 0$, for all t where $i^{\text{ZLB}}(t) = 0$.

The following proposition expands upon proposition 3.8 and 3.9 by offering a striking, novel characterization of how ZLB-constrained policy and unconstrained policy differ.

Proposition 3.10. *If unconstrained optimal policy features negative nominal rates, then $i^{\text{ZLB}}(t) \leq \max(i(t), 0)$, with strict inequality whenever $i(t) > 0$.*

In short, if the zero lower bound constraint is *ever* binding, then a ZLB-constrained policymaker optimally sets interest rates lower at *every* t where it is feasible to do so. These lower rates are used to compensate for the higher-than-optimal rates during periods when the zero lower bound binds.

4 Revisiting ZLB traps

Following the general results in section 3, in this section I consider a more specific scenario: a zero lower bound “trap”, featuring a negative natural real rate over some period of time.

Specifically, I suppose that in the interval $[0, T]$ —the “trap”—the natural rate takes some strictly negative value $-\bar{r}$, followed by a return to a nonnegative steady state value $r^{\text{ss}} \geq 0$.

$$r^n(t) = \begin{cases} -\bar{r} & 0 \leq t < T \\ r^{\text{ss}} & T \leq t \end{cases}$$

To start, I assume that $r^{\text{ss}} = 0$, which greatly simplifies characterization of the solution and facilitates some useful analytical results. This path for the natural rate is depicted in

Figure 9: Trajectory of the natural rate during the trap episode




figure 9. I will later consider the case where $r^{ss} > 0$ in section 4.4.¹¹

Exercises of this form are ubiquitous in the literature on the zero lower bound—for instance, a stochastic trap is in Eggertsson and Woodford (2003), and a deterministic trap similar to my own is in Werning (2011). The negative natural rate trap is popular because it epitomizes the problems created by the zero lower bound: when interest rates cannot be set low enough to match the natural rate, the level of output during the trap—relative to the first-best level—must fall below the level expected after the trap. If the central bank is expected to target first-best output after the trap, then this means that output during the trap is inefficiently low: there is a zero lower bound recession. If, on the other hand, the central bank can commit at the beginning of the trap to policy after the trap, then it optimally engages in “forward guidance”—using low interest rates to generate a boom once the trap is over, lifting up the level of economic activity during the trap as well.

Once negative rates are available as a policy tool, however, this standard analysis of the trap no longer applies. The central bank can now, in principle, set rates to match the natural rate at every point—but given the costs of setting negative rates, of course, this policy is not optimal. My goal in this section is to study the structure of optimal policy under negative rates in detail and contrast its outcomes with the traditional ZLB-constrained policy, with a particular focus on the extent to which negative rates can close the welfare gap relative to the first best.

¹¹For simplicity, I will assume that productivity is constant, so that $r^n(t) = \rho(t)$ according to (13). Variation in the time preference $\rho(t)$ of the representative household can be interpreted as reflecting variation in the effects of idiosyncratic uncertainty and incomplete markets in an underlying heterogenous-agents model; see, for instance, Werning (2015). For the dynamics of interest rates and the output gap, to a first-order approximation it does not matter whether variation in $r^n(t)$ is driven by time preference $\rho(t)$ or productivity growth $\dot{A}(t)/A(t)$, but the assumption that $r^n(t) = \rho(t)$ is needed for the fully nonlinear solution here.

4.1 Calibration

Now that I am interested in a more quantitative analysis, I need to specify a calibration of the model. Aside from cash, there are three parameters: the elasticity of intertemporal substitution σ and Frisch elasticity ψ in (3), and the elasticity of output $1 - \alpha$ with respect to labor input in (8).

Since $1 - \alpha$ is also the labor share in the model, I calibrate it on this basis at $1 - \alpha = 0.56$, to match the labor share of factor income in the United States in 2014.¹² I calibrate the Frisch elasticity of labor supply to be $\psi = 0.86$, reflecting the Frisch elasticity for aggregate hours obtained from studies with micro identification in the [Chetty, Guren, Manoli and Weber \(2013\)](#) meta-analysis. Finally, I calibrate the elasticity of intertemporal substitution to be $\sigma = 0.50$, which the meta-analysis in [Havránek \(2013\)](#) identifies as the mean value in the literature, and [Hall \(2009\)](#) describes as the “most reasonable” choice for the parameter.¹³

My calibrated functional form for the utility from cash is given by

$$v'(m) = -\frac{u'(c^*)}{b} \log\left(\frac{m}{m^*}\right) \quad (31)$$

This function implies a roughly constant interest semielasticity $\partial \log M^d / \partial i$ of cash demand, which is exactly constant and equal to $-b$ when consumption c is at its first-best level c^* . Cash demand at $i = 0$ is given by m^* , which I calibrate to match the current ratio of cash in circulation to GDP in the United States (where $i \approx 0$), which is 0.075.¹⁴

The literature on the interest semielasticity has produced varied results, and it generally looks at demand for M1—including both cash and demand deposits—rather than isolating cash. I tentatively adopt the estimate from [Ball \(2001\)](#), drawn from the postwar United States, of an interest semielasticity equal to -5 . Although this estimate does not cover a period with negative interest rates, it nevertheless appears consistent with the behavior of Swiss cash demand in response to negative rates as displayed in figure 3: with a semielasticity of -5 , the recent drop in Swiss target rates from 0% to -0.75% would be ex-

¹²This is taken from NIPA Table 1.10, as the ratio of compensation of employees to gross domestic income minus production taxes net of subsidies.

¹³[Havránek \(2013\)](#) argues that this mean value is inflated somewhat due to publication bias. On the other hand, since I am interpreting consumption c in this model as a measure of the overall level of economic activity—which also includes the much more interest-sensitive category of fixed investment—the relevant EIS here should be higher than the estimates obtained in the literature for private consumption alone. There are also aggregate redistributive effects that boost the consumption response to interest rates, as identified in [Auclert \(2015\)](#). Altogether, for my benchmark calibration, I assume that these biases roughly offset each other.

¹⁴I am normalizing first-best aggregate output, c^* , to 1.

Table 1: Calibration

Parameter		Source
Elasticity of intertemporal substitution	$\sigma = 0.5$	Havránek (2013), Hall (2009)
Frisch elasticity of labor supply	$\psi = 0.86$	Chetty et al. (2013)
Elasticity of output to labor input	$1 - \alpha = 0.56$	Labor share in US, 2014
Cash demand at 0% interest rates	$m^* = 0.075$	Cash/GDP in US, 2014
Interest semielasticity of cash demand	$\partial \log M^d / \partial i = 5$	Ball (2001)

pected to produce just shy of a 4% increase in cash demand, similar to the slight increase in cash demand actually observed in Switzerland relative to trend.

Table 1 summarizes this calibration. My benchmark scenario features a natural rate of $-\bar{r} = -2\%$ during the trap, and a trap length of $T = 4$. This is intended to generate a moderately severe recession under ZLB-constrained policy, with output starting the trap at 4% below its first-best level, as seen in the next section.

4.2 Partial commitment case


To facilitate comparison with the zero lower bound literature—which often emphasizes the case where the monetary authority lacks commitment—I will start by modifying the assumption of full commitment from section 3. Dropping commitment entirely, however, is not a viable option in this environment: in the limit where policy is continually reoptimized, nominal interest rates are simply set to 0% at all times to satisfy the Friedman rule.

Instead, I consider a simple case with *partial commitment*, where optimal policy is reoptimized at $t = T$, the end of the trap. Figure 10 shows the results. Under both ZLB-constrained and unconstrained policy, the nominal rate is set to zero—equal to the natural rate—following $T = 4$, and consumption is stabilized at its natural level, thereby achieving the first best from $T = 4$ onward.

Prior to $T = 4$, a recession ensues in both cases, and $i(t)$ exceeds the natural rate for all t . Unsurprisingly, however, this recession is far more severe in the ZLB-constrained case; with negative rates, the gap between $i(t)$ and the natural rate $-\bar{r}$ shrinks substantially.

The path $i(t)$ of interest rates in the unconstrained case in figure 10 exhibits three distinctive features that are, in fact, general—as summarized by the following proposition.

Figure 10: Optimal policy under partial commitment: with and without ZLB



Proposition 4.1. *In the trap under partial commitment, $i(0) = 0$, $i(T^-) > -\bar{r}$, and $i(t)$ is strictly decreasing on the interval $[0, T)$.*

All three features of $i(t)$ result from the central tradeoff in the model—the tradeoff between the costs of departing from the Friedman rule by setting negative rates and the benefits of increasing consumption during the trap.

A lower rate at time t raises consumption over the interval $[0, t]$, and the overall welfare gains naturally depend on the length of this interval. At $t = 0$, the length is zero, implying that the tradeoff is resolved entirely in favor of the Friedman rule: $i(0) = 0$. As t increases, the benefits grow relative to the costs, implying that optimal $i(t)$ decreases. But $i(t)$ does not decrease so steeply that it falls below \bar{r} : if it did, there would be a boom in the latter part of the trade episode $[0, T]$, and the policymaker could increase welfare by smoothing the path of $i(t)$ and eliminating this boom.

This *backloading* of the most negative rates is an important feature of optimal policy. This differs sharply from the path of $i(t)$ implied by, for instance, a Taylor rule where i depends on the output gap—which would be strictly *increasing*. Hence, although the case for negative rates is often illustrated by an appeal to the negative rate implied by a Taylor rule¹⁵, this view misses a distinctive feature of the optimal policy.

To what extent do negative rates close the utility gap? Let V^{ZLB} denote household utility (2) under ZLB-constrained optimal policy, V^* denote household utility under un-

¹⁵See, for instance, Rudebusch (2010), contrasting the path of nominal interest rates implied by a simple linear policy rule—which falls to nearly 6% in 2009–10—with the actual ZLB-constrained path.

constrained optimal policy, and V^{FB} denote household utility in the first best.

One natural measure of how unconstrained optimal policy improves upon ZLB constrained optimal policy is the extent to which it shrinks the gap in utility relative to the first best: the ratio $(V^* - V^{FB}) / (V^{ZLB} - V^{FB})$. As displayed in figure 10, in the benchmark scenario under partial commitment this ratio is extremely small, at 5.6%: the ability to set negative rates eliminates the vast majority of the utility shortfall.

This ratio can be analytically characterized as a function of primitives, up to a second-order approximation, as revealed by the following proposition.

Proposition 4.2. *The following is a second-order approximation for the decline in the welfare gap:*

$$\frac{V^* - V^{FB}}{V^{ZLB} - V^{FB}} = 3 \left(\frac{1}{T\sqrt{A}} \right)^2 + O \left(\left(\frac{1}{T\sqrt{A}} \right)^3 \right) \quad (32)$$

where

$$A \equiv \frac{\sigma^2 \times \frac{\partial \log \tau(c)}{\partial \log c}}{\frac{\partial \log M^d(0, c^*)}{\partial i} \times m^*} \quad (33)$$

(32) implies that the decline in the welfare gap is more dramatic when T is large. This is no surprise: when the trap is longer, negative rates can lift output over a longer period, making them a more useful tool.

The decline is also increasing in the composite parameter A . As (33) reveals, A is increasing in the elasticity of intertemporal substitution σ (which determines the influence of negative rates on consumption) and the elasticity $\partial \log \tau(c) / \partial \log c$ of the labor wedge with respect to consumption (which determines the magnitude of welfare loss from the output gap). It is decreasing in the interest semielasticity of cash demand and the level of cash demand m^* at $i = 0$, which as depicted in figure 6 determine the costs from deviating from the Friedman rule.


In addition to these qualitative insights, (32) is remarkably accurate from a quantitative standpoint, as long as $T\sqrt{A}$ is relatively large. In the benchmark parameterization above, for instance, the approximation is

$$3 \left(\frac{1}{T\sqrt{A}} \right)^2 = 3 \times \frac{\frac{\partial \log M^d(0, c^*)}{\partial i} \times m^*}{T^2 \times \sigma^2 \times \frac{\partial \log \tau(c)}{\partial \log c}} = 3 \times \frac{5 \times 0.075}{4^2 \times 0.5^2 \times 4.86} = 0.058 \quad (34)$$

which is very close to the actual ratio 0.056 obtained in the simulation.

It is clear from (34) how various features of the calibration in section 4.1 contribute to closing the utility gap. A higher interest semielasticity of cash demand, for instance,

Figure 11: Optimal policy under full commitment: with and without ZLB



would result in a larger utility gap under unconstrained policy—but even if the semielasticity of 5 was replaced by 25, negative rates would still cut the utility gap relative to ZLB-constrained policy by over two-thirds.

4.3 Full commitment case

Now I revert to the original assumption of full commitment. Figure 10 then becomes figure 11.


The path of ZLB-constrained optimal policy does not change. Since the natural rate after the trap is zero, it is not possible to create a boom by committing to hold rates below the natural rate after the trap, as in Eggertsson and Woodford (2003) or Werning (2011).

The path of unconstrained optimal policy, meanwhile, involves negative rates even *after* the trap ends at T . This leads to a boom in consumption in the neighborhood of time T , which pulls up the entire path of consumption over the interval $[0, T]$ and brings output closer to its first-best level. The relevant features of the solution are general, and summarized in the following proposition.

Proposition 4.3. *The optimal solution under full commitment features a path:*

- for $i(t)$ that is decreasing from 0 to T' (where $T' < T$), reaches a minimum at T' , and is increasing from T' onward, such that $i(0) = 0$, $i(t) \rightarrow 0$ as $t \rightarrow \infty$, and $i(t) > -\bar{r}$ for all t .
- for $c(t)$ that is increasing from 0 to T and decreasing from T onward, with $c(0) < c^*$, $c(T) > c^*$, and $c(t) \rightarrow c^*$ as $t \rightarrow \infty$.

Figure 12: Optimal policy under commitment: positive natural rate after trap



The key insight from the full commitment case is that negative rates are not just an alternative to forward guidance. Instead, the two are complements, in the sense that it is optimal to do forward guidance *with* negative rates. Quantitatively, however, this is not of great importance: as figure 11 reveals, with full commitment the utility gap is reduced to 5.1% of its ZLB-constrained level, not much better than the 5.6% achieved with partial commitment in figure 10.

4.4 Full commitment case, with a positive natural rate after the trap

The assumption that $r^{ss} = 0$, and therefore that $r^n(t) = 0$ for all $t \geq T$, simplified characterization of optimal policy, but it made forward guidance in the ZLB-constrained case impossible. I now relax this assumption, considering $r^{ss} > 0$ instead. Figure 12 shows the paths that result when $r^{ss} = 1.5\%$.¹⁶

Letting $i(t)$ and $c(t)$ denote interest rates and consumption with unconstrained optimal policy, and $i^{ZLB}(t)$ and $c^{ZLB}(t)$ denote these with ZLB-constrained optimal policy, the following proposition summarizes the general features of the solution.

Proposition 4.4. *In the optimal solution under full commitment, $i(t)$ starts below $i^{ZLB}(t)$ and crosses it once. $c(t)$ starts above $c^{ZLB}(t)$ and crosses it once.*

¹⁶The utility gap measure reported in previous cases is no longer meaningful, since the assumption that the steady-state natural rate is above zero implies that there is inevitably a departure from first-best utility in the steady state, which creates a large wedge between actual intertemporal utility (2), over the interval $[0, \infty)$, and the first best.

As figure 12 depicts, outcomes under ZLB-constrained policy and unconstrained policy are now qualitatively similar. Interest rates are set below the natural rate after the trap, generating a boom (in which output exceeds its first-best level) in the neighborhood of $T = 4$. Interest rates exceed the natural rate during the trap, leading to a bust (in which output falls short of its first-best level) for the majority of the trap.

The key difference between the two cases is *quantitative*: the output gap is much closer to zero, and less variable, when negative rates are available. At the same time, interest rates are much more volatile. Effectively, once the zero lower bound constraint is lifted, optimal policy accepts more interest rate volatility in order to stabilize output. In contrast, under the zero lower bound, rates are artificially stable, as policymakers are forced to maintain $i = 0$ for a prolonged period after the trap in order to lift economic activity during the trap.

5 Interaction with other policies

5.1 Trend inflation


One limitation of the framework in the model thus far is the simplifying assumption of perfectly sticky prices. With inflation fixed at zero by assumption, it is impossible to evaluate the common idea—proposed by Blanchard et al. (2010), and evaluated formally by Coibion et al. (2012) and others—that higher trend inflation alleviates the limitations on policy imposed by the zero lower bound.

In this section, I relax the assumption of perfect price stickiness, allowing for nonzero inflation. To preserve the parsimony of the model, however, I continue to make a strong assumption on prices: I replace sticky prices with sticky inflation, where the path of prices is constrained to take the form $P(t) = e^{\bar{\pi}t}P(0)$ for some trend inflation rate $\bar{\pi}$. This embeds my earlier case, which corresponds to $\bar{\pi} = 0$. It is intended to capture the role of trend inflation in the simplest way possible, and it can also be understood as a stylized representation of well-anchored inflation expectations under an inflation targeting regime.

The Euler equation characterizing the path of consumption becomes

$$\frac{\dot{c}}{c} = \sigma(i - \bar{\pi} - \rho) \tag{35}$$

Figure 13: Optimal policy under full commitment: with and without trend inflation



Optimal policy is now characterized by the system

$$\hat{\rho}\sigma = i \cdot m \cdot -\frac{\partial \log M^d}{\partial i} \quad (23)$$

$$\hat{\rho} = -c\tau - \bar{\pi}\hat{\rho} \quad (36)$$

$$\hat{\rho}(0) = 0 \quad (25)$$

where (23) and (25) are unchanged from the formulation in section 2.2, while (36) is a modified version of (24) with the additional term $\bar{\pi}\hat{\rho}$. When comparing optimal policy under distinct trend inflation rates $\bar{\pi}$, I will make the dependence on $\bar{\pi}$ explicit by introducing it as an argument: for instance, $i(t; \bar{\pi})$.

As trend inflation $\bar{\pi}$ increases, the average level of nominal interest rates over time has to increase along with it: otherwise, the Euler equation implies that consumption will diverge unboundedly to either 0 or ∞ . But this does not imply a parallel shift in the path of nominal interest rates i : optimal policy accommodates inflation in a way that minimizes the cost of deviating from the Friedman rule.

For instance, figure 13 shows how the optimal path of nominal interest rates $i(t)$ from the scenario in section 4.3 changes as trend inflation is set to $\bar{\pi} = 1\%$.¹⁷ In the ZLB-constrained case, rates do not change at all during and immediately after the trap, staying at zero; instead, the planner takes advantage of trend inflation in order to implement a lower real interest rate and lessen the severity of the recession. In the unconstrained case,

¹⁷Note that the dashed line in figure 13 shows the *nominal* natural rate $r^n(t) + \bar{\pi}$, which is shifted up when $\bar{\pi} > 0$.

$i(0)$ remains at 0%, although rates shift up by approximately $\bar{\pi}$ at later t .

Given the subtlety of the adjustment in figure 13, is it possible to make any general statements about how optimal policy responds to trend inflation? The following proposition shows that it is.

Proposition 5.1. *The path of optimal nominal interest rates under optimal policy is, for all $t > 0$, strictly increasing in trend inflation:*

$$\bar{\pi}' > \bar{\pi} \text{ implies } i(t; \bar{\pi}') > i(t; \bar{\pi})$$

In short, although the response of optimal nominal interest rates to rising trend inflation is uneven—with the Fisher equation holding only on average—it is unambiguously true that interest rates will rise at each t .

Optimal trend inflation. How do different levels of $\bar{\pi}$ affect household utility under optimal policy—and, in particular, how does this differ depending on whether the zero lower bound is imposed? Let $W(\bar{\pi})$ denote household utility (2) under unconstrained optimal policy, and $W_{ZLB}(\bar{\pi})$ denote the same under ZLB-constrained optimal policy.

Proposition 5.2. *There exists some $\bar{\pi}_u$ such that $W_{ZLB}(\bar{\pi}) - W(\bar{\pi})$ is strictly increasing in $\bar{\pi}$ for all $\bar{\pi} < \bar{\pi}_u$, and $W_{ZLB}(\bar{\pi}) = W(\bar{\pi})$ for all $\bar{\pi} \geq \bar{\pi}_u$.*

This proposition states that the zero lower bound constraint is *complementary* to higher inflation: as long as the constraint is binding, the utility cost of imposing the constraint becomes smaller as trend inflation rises. Eventually, a high enough level of inflation π_u is reached that the zero lower bound is no longer binding at all, and $W_{ZLB}(\pi_u) = W(\pi_u)$. This follows from the usual intuition: higher inflation allows more negative *real* rates to be realized without setting negative *nominal* rates.

Let $\pi^* = \operatorname{argmax} W(\pi)$ and $\pi_{ZLB}^* = \operatorname{argmax} W_{ZLB}(\pi)$ be the levels of inflation that maximize household utility. If we interpret the sticky inflation rate $\bar{\pi}$ as the result of a long-term policy of anchoring inflation expectations, then these are the policymaker's optimal choices of trend inflation, as characterized by the following corollary to proposition 5.2.

Corollary 5.3. *Optimal inflation with unconstrained policy is below optimal inflation with the zero lower bound: $\pi^* \leq \pi_{ZLB}^*$. This inequality is strict, and $\pi^* < \pi_{ZLB}^* < \bar{\pi}_u$, except in the special case $r^n(t) = \bar{r}$ where the natural rate is constant over all t , in which case $\pi^* = \pi_{ZLB}^* = \bar{\pi}_u = -\bar{r}$.*

In general, optimal inflation under unconstrained policy is lower than optimal inflation under ZLB-constrained policy ($\bar{\pi}^* < \bar{\pi}_{ZLB}^*$), because inflation is necessary to achieve negative real rates when there is a zero lower bound.

The corollary also states that $\bar{\pi}_{ZLB}^* < \bar{\pi}_u$; in other words, that inflation should not be set so high as to negate the zero lower bound constraint altogether. For levels of $\bar{\pi}$ immediately below $\bar{\pi}_u$, the zero lower bound barely binds, and the cost imposed by the constraint is second-order—which is overridden by the first-order benefits from lowering the trend rate of inflation and thereby bringing the average level of nominal interest rates slightly closer to the Friedman rule optimum.

The special case $r^n(t) = \bar{r}$ where the natural rate never varies is an exception, and it features particularly simple optimal policy. By setting trend inflation equal to minus the natural rate, the policymaker can match the natural rate while setting the nominal interest rate equal to zero at all times. This achieves the first best: it implements a first-best level of output while also satisfying the Friedman rule.

The crucial message of corollary 5.3 is that once negative nominal interest rates become available as a tool, trend inflation should be brought down. This weakens a longstanding argument—dating back at least to Summers (1991) and Fischer (1994)—that low inflation is dangerous due to the limitations of interest rate policy. It also hints at a source of political appeal: although negative rates may not seem attractive on their own, they can facilitate a low inflation target, which is a broadly popular idea.

5.2 Abolishing cash


As discussed in section 1, one policy that has been commonly suggested as a response to the zero lower bound is the abolition of cash—see, for instance, Rogoff (2014) for details. Once negative rates are permissible, this original rationale no longer holds in the same form.

Since negative rates are not costless, however, it remains possible that abolishing cash is optimal, if the cost from the subsidy to cash under negative rates exceeds the benefits from having cash. Figure 14 depicts this possibility. There is some level of cash \bar{m} at which $v(\bar{m}) = v(0) = 0$, beyond which it would be optimal to replace $m > \bar{m}$ with $m = 0$.

Figure 14 additionally shows the interest rate $\underline{i} < 0$ such that $\underline{i}u'(c^*) = v'(\bar{m})$, and therefore $\bar{m} = M^d(\underline{i}, c^*)$, where c^* is the first-best level of consumption. If consumption is at its first-best level, and the nominal interest rate is kept at or above \underline{i} , flow utility is higher when cash is kept rather than abolished.

This observation gives rise to the following proposition, which offers a simple suf-

Figure 14: Threshold \bar{m} for cash being inefficient



efficient condition for keeping cash (rather than abolishing it) to be optimal given some sequence of fundamentals.

Proposition 5.4. *Let $c^* = \min_t c^*(t)$, where $c^*(t)$ is the first-best level of consumption at time t , and write $\underline{i} \equiv v'(\bar{m})/u'(c^*)$. Then it is optimal to keep cash if the natural nominal rate is bounded from below by \bar{i} : $\bar{\pi} + r^n(t) \geq \bar{i}$ for all t .*

Proof. One feasible policy is to set $i(t) = \bar{\pi} + r^n(t)$ and $c(t) = c^*(t)$ for all t . Then since $i(t) = \bar{\pi} + r^n(t) \geq \underline{i}$,

$$m(t) = M^d(i(t), c^*(t)) \leq M^d(\bar{i}, c^*) = \bar{m}$$

and therefore $v(m(t)) \geq v(\bar{m}) = v(0)$. It follows that this policy achieves weakly better utility than that from the no-cash $m = 0$ case. \square

The idea behind proposition 5.4 is straightforward. Suppose that the natural rate is always above the interest rate threshold $\underline{i} < 0$ at which the cumulative utility from cash might become negative. Then one option for the policymaker is to set the nominal interest rate to equal the natural nominal rate at all t , achieving the first-best level of economic activity at all t while never exceeding the level \bar{m} at which $v(\bar{m}) = v(0)$. With this option available, the policymaker would never opt to abolish cash.

What is a plausible level of \underline{i} ? This depends on the shape of v , which can be mapped onto measurable objects like the interest semielasticity.

Lemma 5.5. *Suppose that $M^d(i, c^*)$ has a constant semielasticity of $-b$ with respect to i . Then $\underline{i} = -1/b$.*

Lemma 5.6. *Let v and \tilde{v} be two alternative utility functions for cash, with corresponding cash demand schedules M^d and \tilde{M}^d . Suppose that the interest semielasticity of cash demand for M^d is always smaller than that of \tilde{M}^d :*

$$-\frac{\partial \log M^d(i, c)}{\partial i} < -\frac{\partial \log \tilde{M}^d(i, c)}{\partial i} \text{ for all } i, c$$

Then $\underline{i} < \tilde{\underline{i}}$.

Lemma 5.5 shows how a constant semielasticity maps onto the lower bound \underline{i} in proposition 5.4. Note that \underline{i} is more negative when the semielasticity is smaller. Intuitively, this is because when cash demand is less interest-elastic, the inframarginal benefits of cash are larger compared to the marginal losses, and the level of cash demand has to rise higher—corresponding to a much more negative interest rate—until these inframarginal benefits are wiped out. Lemma 5.6 formalizes this point in the general case, showing that when cash demand is less interest elastic, \underline{i} is lower.

Calibrating lemma 5.5 using the values for the interest semielasticity from section 4.1, the implied \underline{i} is extremely low. Given $b = 5$, for instance, $\underline{i} = -20.0\%$. In practice, the semielasticity probably becomes much higher for such low nominal interest rates: indeed, at $i = -20.0\%$, cash demand may be virtually infinite.

But even with a much higher semielasticity—intended to reflect the more elastic response of cash demand in the relevant range of negative interest rates— \underline{i} is still quite low. At $b = 50$, for instance, lemma 5.5 implies $\underline{i} = -2\%$; and since the empirical semielasticity is much lower than this when nominal interest rates are positive or mildly negative, a more precise calculation would obtain a lower \underline{i} . Regardless, based on estimates in the literature, it seems plausible that the natural nominal rate $\bar{\pi} + r^n(t)$ has always, or almost always, been above -2% in the United States.¹⁸ Proposition 5.4 would then indicate that it is not optimal to abolish cash.

More generally, proposition 5.4—in conjunction with lemmas 5.5 and 5.6—tells us that the path of the natural rate and the semielasticity of cash demand are key considerations in calculating whether or not to keep cash. Since these are both matters of some empirical controversy, further work will be needed to obtain a more definitive answer.

¹⁸Cúrdia (2015) estimates that the $r^n(t)$ reached a minimum of slightly above -4% in the aftermath of the Great Recession, which corresponds to a natural nominal rate $\bar{\pi} + r^n(t)$ of slightly above -2% under trend inflation of $\bar{\pi} = 2\%$. Del Negro, Giannoni, Cocci, Shahanaghi and Smith (2015) directly provide figures for the natural nominal rate and show that it fell to slightly above -2% for much of the Great Recession, though it dipped below -2% for a brief period around the beginning of 2013. Laubach and Williams (2015)—using a statistical model rather than the structural models in the other two papers—find much higher natural rates, with a real natural rate of roughly 0% with the benchmark methodology and a real natural rate of slightly below -2% using an alternative methodology for estimating the output gap.

5.3 Multiple denominations

One alternative to abolishing cash altogether is to selectively eliminate certain denominations of cash. A number of observers have proposed eliminating high-value denominations (for instance, the \$100 bill) to curtail cash demand under negative rates, with the reasoning that these denominations have lower holding costs per unit value and are thus particularly likely to be hoarded. I evaluate this idea by extending the model such that cash comes in more than one denomination.

For simplicity, I consider a case with only two denominations, letting $z(m_h, m_l)$ denote the utility from holding m_h in high-denomination cash and m_l in low-denomination cash.¹⁹ I then define

$$\begin{aligned} v_{bd}(m) &\equiv \max_{m_h+m_l=m} z(m_h, m_l) \\ v_{hd}(m) &\equiv z(m, 0) \\ v_{ld}(m) &\equiv z(0, m) \end{aligned}$$

Here v_{bd} denotes the utility from having cash m that can be spread between *both denominations*, while v_{hd} denotes the utility from cash only in the *high denomination* and v_{ld} denotes the utility from cash only in the *low denomination*. I take these as representing alternative policy choices: when setting up the monetary system, the government can choose to provide both denominations or only one. Whatever choice is adopted, the household utility function (2) then uses the corresponding v .

Let M_{bd}^d , M_{hd}^d , and M_{ld}^d be the corresponding demand functions for cash. I then make the following assumption.

Assumption 5.7. *When $i = 0$, cash demand is highest when both denominations are available and lowest when only the low denomination is available:*

$$M_{ld}^d(0, \cdot) < M_{hd}^d(0, \cdot) < M_{bd}^d(0, \cdot)$$

Furthermore, for any c ,

$$\frac{M_{bd}^d(i, c)}{M_{hd}^d(i, c)} \text{ and } \frac{M_{ld}^d(i, c)}{M_{hd}^d(i, c)} \tag{37}$$

are strictly increasing functions in i .

Both parts of assumption 5.7 are natural for a model with multiple denominations. The second part, in particular, captures a key distinction between high and low denom-

¹⁹I show in the Online Appendix how a utility function from cash of this form can be microfounded.

inations: since the holding costs for high denominations are lower, the cost of foregone nominal interest i looms larger, making demand for these denominations more elastic with respect to the nominal interest rate. (For instance, the nominal interest rate is much more important when deciding how many \$100 bills to carry than when deciding how many \$1 bills to carry.)

The following example shows how a simple functional form for z can produce cash demand functions that satisfy assumption 5.7. I show in the Online Appendix how such a functional form can be microfounded in a model of cash transactions.

Example 5.8. Suppose that

$$z(m_h, m_l) = \frac{A_h m_h^{1-\zeta}}{1-\zeta} + \frac{A_l m_l^{1-\zeta}}{1-\zeta} - \alpha_h m_h - \alpha_l m_l$$

where $\alpha_h < \alpha_l$ and $A_h/\alpha_h > A_l/\alpha_l$. Then the conditions of assumption 5.7 are satisfied.

Now define W_{bd} , W_{hd} , W_{ld} , and W_{nd} to be household utility (2) under optimal policy with both denominations, the high denomination only, the low denomination only, and no cash, respectively.

Proposition 5.9. *If eliminating the low denomination increases utility under optimal policy, then utility is increased further by either eliminating the high denomination instead, or abolishing cash altogether:*

$$\text{if } W_{hd} \geq W_{bd}, \text{ then } W_{ld} > W_{hd} \geq W_{bd} \text{ and } W_{nd} > W_{hd} \geq W_{bd}$$

This proposition states that it is never optimal to eliminate only the low denomination: if this leads to an improvement, then it is even better to either eliminate the high denomination or to abolish cash.

The key force driving proposition 5.9 is the assumption of increasing ratios (37). Demand for the high denomination is *relatively* greater when interest rates are lower. The only reason to eliminate denominations is to cut down on excessive cash demand when interest rates are negative; hence, if anything, the high denomination should be axed, since its demand increases most disproportionately at these negative rates.

6 Conclusion

This paper studies, for the first time, the use of negative nominal interest rates as part of optimal monetary policy—without any major changes to the monetary system. I show

that negative rates are costly in this environment because they imply an inefficient subsidy to cash, violating the Friedman rule in the opposite of the traditional direction. I replace the zero lower bound, which is imposed as an ad-hoc constraint in many New Keynesian models, with a more flexible, microfounded tradeoff between the distortionary costs of negative rates and their benefits in raising aggregate demand.

The first insight that emerges from this framework is that when the economy is, on average, in a slump, negative rates are always optimal to some degree: the first-order benefits of boosting aggregate output outweigh the second-order costs of deviating from the Friedman rule optimum. An effective zero lower bound only emerges in a limit case, where cash demand becomes infinitely elastic at zero—a case that is not consistent with recent evidence from currencies with negative rates.

Revisiting the liquidity trap scenarios that are studied in the zero lower bound literature, I show that negative rates bring significant improvements. In a benchmark scenario, introducing negative rates brings utility 94% closer to the first best. Optimal policy dictates that the most negative rates should be backloaded, and indeed that negative rates should continue even after the trap has ended. Although ZLB-constrained policy can mitigate the worst consequences of the trap through forward guidance, negative rates facilitate a path for output that is much closer to the first best.

Negative rates are most useful as a tool when cash demand is relatively inelastic, because this is when the distortion from violating the Friedman rule is least severe. Policies that can constrain cash demand in this region are therefore important complements to negative rates. In the most extreme case, cash can be abolished; but more limited measures, such as the elimination of larger denominations, may also be worthwhile. An important task of future research will be to study the properties of cash demand at very low interest rates, and to search for new measures that can contain this demand—ensuring that negative rates realize their great potential as a policy tool.

References

- Agarwal, Ruchir and Miles Kimball**, “Breaking Through the Zero Lower Bound,” IMF Working Paper 15/224 2015.
- Auclert, Adrien**, “Monetary Policy and the Redistribution Channel,” Working Paper 2015.
- Ball, Laurence**, “Another Look at Long-Run Money Demand,” *Journal of Monetary Economics*, February 2001, 47 (1), 31–44.

- , “The Case for Four Percent Inflation,” *Central Bank Review*, 2013, 13 (2), 17–31.
- Baumol, William J.**, “The Transactions Demand for Cash: An Inventory Theoretic Approach,” *The Quarterly Journal of Economics*, 1952, 66 (4), 545–556.
- Blanchard, Olivier, Giovanni Dell’Ariccia, and Paolo Mauro**, “Rethinking Macroeconomic Policy,” *Journal of Money, Credit and Banking*, September 2010, 42, 199–215.
- Buiter, Willem H.**, “Is Numeraiology the Future of Monetary Economics?,” *Open Economies Review*, May 2007, 18 (2), 127–156.
- , “Negative Nominal Interest Rates: Three Ways to Overcome the Zero Lower Bound,” *The North American Journal of Economics and Finance*, December 2009, 20 (3), 213–238.
- **and Nikolaos Panigirtzoglou**, “Liquidity Traps: How to Avoid Them and How to Escape Them,” in Wim F.V. Vanthoor and Joke Mooij, eds., *Reflections on Economics and Econometrics, Essays in Honour of Martin Fase*, Amsterdam: De Nederlandsche Bank NV, 2001, pp. 13–58.
- **and** – , “Overcoming the Zero Bound on Nominal Interest Rates with Negative Interest on Currency: Gesell’s Solution,” *The Economic Journal*, October 2003, 113 (490), 723–746.
- Cagan, Phillip**, “The Monetary Dynamics of Hyperinflation,” in “Studies in the Quantity Theory of Money,” Chicago: University of Chicago Press, 1956.
- Calvo, Guillermo A.**, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, September 1983, 12 (3), 383–398.
- Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano**, “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, 2012, 2012 (1), 1–80.
- Chari, V. V., Lawrence J Christiano, and Patrick J Kehoe**, “Optimality of the Friedman Rule in Economies with Distorting Taxes,” *Journal of Monetary Economics*, April 1996, 37 (2), 203–223.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber**, “Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities,” *NBER Macroeconomics Annual*, 2013, 27 (1), 1–56.

- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland**, “The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?,” *Review of Economic Studies*, 2012, 79 (4), 1371–1406.
- Correia, Isabel and Pedro Teles**, “Is the Friedman Rule Optimal When Money Is an Intermediate Good?,” *Journal of Monetary Economics*, October 1996, 38 (2), 223–244.
- Cúrdia, Vasco**, “Why so Slow? A Gradual Return for Interest Rates,” *FRBSF Economic Letter*, 2015, 2015 (32).
- da Costa, Carlos E. and Iván Werning**, “On the Optimality of the Friedman Rule with Heterogeneous Agents and Nonlinear Income Taxation,” *Journal of Political Economy*, 2008, 116 (1), 82–112.
- Draghi, Mario**, “Introductory Statement to the Press Conference (with Q&A),” October 2015.
- Eggertsson, Gauti B. and Michael Woodford**, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, January 2003, 2003 (1), 139–211.
- Eisler, Robert**, *Stable Money: The Remedy for the Economic World Crisis: A Programme of Financial Reconstruction for the International Conference 1933*, London: The Search Publishing Co, 1932.
- Fischer, Stanley**, “The Costs and Benefits of Disinflation,” in J. Onno De Beaufort Wijnholds, Sylvester C. W. Eijffinger, and Lex H. Hoogduin, eds., *A Framework for Monetary Stability*, number 27. In ‘Financial and Monetary Policy Studies.’, Springer Netherlands, 1994, pp. 31–42.
- Friedman, Milton**, “The Optimum Quantity of Money,” in “The Optimum Quantity of Money and Other Essays,” Chicago: Aldine, 1969.
- Fuhrer, Jeffrey C. and Brian F. Madigan**, “Monetary Policy when Interest Rates Are Bounded at Zero,” *Review of Economics and Statistics*, November 1997, 79 (4), 573–585.
- Galí, Jordi**, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton: Princeton University Press, March 2008.
- Gesell, Silvio**, *Die natürliche Wirtschaftsordnung durch Freiland und Freigeld*, Berlin-Lichterfelde: G. Blumenthal, 1916.

- Goodfriend, Marvin**, "Overcoming the Zero Bound on Interest Rate Policy," *Journal of Money, Credit and Banking*, November 2000, 32 (4), 1007–1035.
- Hall, Robert E.**, "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, 2009, 117 (2), 281–323.
- Havránek, Tomáš**, "Publication Bias in Measuring Intertemporal Substitution," IES Working Paper 15/2013 2013.
- Hicks, J. R.**, "Mr. Keynes and the "Classics"; A Suggested Interpretation," *Econometrica*, 1937, 5 (2), 147–159.
- Humphrey, David B.**, "Negative Interest Rates and the Demand for Cash," Working Paper 2015.
- Ireland, Peter N.**, "On the Welfare Cost of Inflation and the Recent Behavior of Money Demand," *The American Economic Review*, 2009, 99 (3), 1040–1052.
- Jackson, Harriet**, "The International Experience with Negative Policy Rates," Staff Discussion Paper 2015-13, Bank of Canada 2015.
- Jensen, Carina Moselund and Morten Spange**, "Interest Rate Pass-Through and the Demand for Cash at Negative Interest Rates," *Danmarks Nationalbank Monetary Review*, 2015, 2nd Quarter.
- Keynes, John Maynard**, *The General Theory of Employment, Interest, and Money*, London: Macmillan, 1936.
- Khan, Aubhik, Robert G. King, and Alexander L. Wolman**, "Optimal Monetary Policy," *The Review of Economic Studies*, October 2003, 70 (4), 825–860.
- Kimbrough, Kent P.**, "The Optimum Quantity of Money Rule in the Theory of Public Finance," *Journal of Monetary Economics*, November 1986, 18 (3), 277–284.
- Krugman, Paul R.**, "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, January 1998, 1998 (2), 137–205.
- Laubach, Thomas and John C. Williams**, "Measuring the Natural Rate of Interest Redux," Federal Reserve Bank of San Francisco Working Paper 2015-16 2015.
- Levin, Andrew, David López-Salido, Edward Nelson, and Yack Yun**, "Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound," *International Journal of Central Banking*, 2010, 6 (1), 143–189.

- Lucas, Robert E.**, "Inflation and Welfare," *Econometrica*, 2000, 68 (2), 247–274.
- Mankiw, N. Gregory**, "It May Be Time for the Fed to Go Negative," *New York Times*, April 2009.
- McAndrews, James**, "Negative Nominal Central Bank Policy Rates: Where Is the Lower Bound?," May 2015.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson**, "The Power of Forward Guidance Revisited," Working Paper 20882, National Bureau of Economic Research January 2015.
- Meltzer, Allan H.**, "The Demand for Money: The Evidence from the Time Series," *Journal of Political Economy*, 1963, 71 (3), 219–246.
- Negro, Marco Del, Marc P. Giannoni, and Christina Patterson**, "The Forward Guidance Puzzle," *FRB of New York Staff Report*, 2012, (574).
- , – , **Matthew Cocci, Sara Shahanaghi, and Micah Smith**, "Why Are Interest Rates So Low?," May 2015.
- Phelps, Edmund S.**, "Inflation in the Theory of Public Finance," *The Swedish Journal of Economics*, 1973, 75 (1), 67–82.
- Rogoff, Kenneth**, "Costs and Benefits to Phasing Out Paper Currency," in "NBER Macroeconomics Annual 2014, Volume 29," University of Chicago Press, 2014.
- Rudebusch, Glenn D.**, "The Fed's exit strategy for monetary policy," *FRBSF Economic Letter*, 2010, 18 (2010).
- Schmitt-Grohé, Stephanie and Martín Uribe**, "Optimal Fiscal and Monetary Policy Under Imperfect Competition," *Journal of Macroeconomics*, June 2004, 26 (2), 183–209.
- and – , "Optimal Fiscal and Monetary Policy Under Sticky Prices," *Journal of Economic Theory*, February 2004, 114 (2), 198–230.
- Siu, Henry E.**, "Optimal Fiscal and Monetary Policy with Sticky Prices," *Journal of Monetary Economics*, April 2004, 51 (3), 575–607.
- Summers, Lawrence**, "Panel Discussion: Price Stability: How Should Long-Term Monetary Policy Be Determined?," *Journal of Money, Credit and Banking*, 1991, 23 (3), 625–631.
- Tobin, James**, "The Interest-Elasticity of Transactions Demand For Cash," *The Review of Economics and Statistics*, 1956, 38 (3), 241–247.

- Werning, Iván**, “Managing a Liquidity Trap: Monetary and Fiscal Policy,” Working Paper 17344, National Bureau of Economic Research August 2011.
- , “Incomplete Markets and Aggregate Demand,” Working Paper 21448, National Bureau of Economic Research August 2015.
- Williams, John C.**, “Heeding Daedalus: Optimal Inflation and the Zero Lower Bound,” *Brookings Papers on Economic Activity*, 2009, 2009 (2), 1–37.
- Woodford, Michael**, “Chapter 20 the Optimum Quantity of Money,” in *BT Handbook of Monetary Economics*, ed., *BT Handbook of Monetary Economics*, ed., Vol. 2, Elsevier, 1990, pp. 1067–1152.
- , “Optimal Monetary Policy Inertia,” *Manchester School*, 1999, 67 (0), 1–35.
- , *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, September 2003.
- , “Optimal Interest-Rate Smoothing,” *The Review of Economic Studies*, October 2003, 70 (4), 861–886.
- Yellen, Janet**, “Testimony Before the Committee on Financial Services, U.S. House of Representatives,” November 2015.

A Proofs

Proof of proposition 5.1

Proof. Since i is strictly increasing in $\hat{\mu}$ by (23), it suffices to show that $\hat{\mu}(t; \bar{\pi}') > \hat{\mu}(t; \bar{\pi})$ for all $t > 0$.

Suppose to the contrary that the set $A = \{\hat{\mu}(t; \bar{\pi}') \leq \hat{\mu}(t; \bar{\pi}); t > 0\}$ is nonempty, and let $t' \equiv \inf A$. We have $\hat{\mu}(t'; \bar{\pi}') = \hat{\mu}(t'; \bar{\pi})$ and $\hat{\mu}'(t'; \bar{\pi}') \leq \hat{\mu}'(t'; \bar{\pi})$:

- If $t' = 0$, then $\hat{\mu}(t'; \bar{\pi}') = \hat{\mu}(t'; \bar{\pi}) = 0$, and since A contains t arbitrarily close to 0 such that $\hat{\mu}(t; \bar{\pi}') \leq \hat{\mu}(t; \bar{\pi})$, we must have $\hat{\mu}'(0; \bar{\pi}') \leq \hat{\mu}'(0; \bar{\pi})$.
- If $t' > 0$, then since $\hat{\mu}(t; \bar{\pi}') > \hat{\mu}(t; \bar{\pi})$ for all $t < t'$ by construction, it must be that $\hat{\mu}(t'; \bar{\pi}') = \hat{\mu}(t'; \bar{\pi})$ and $\hat{\mu}'(t'; \bar{\pi}') \leq \hat{\mu}'(t'; \bar{\pi})$.

From (36) it follows that $c(t'; \bar{\pi}') \leq c(t'; \bar{\pi})$.

I argue now that we must have $A = [t', \infty)$. Suppose to the contrary that $\inf([t', \infty) \setminus A) = t''$. Then $i(t; \bar{\pi}') \leq i(t; \bar{\pi})$ for all $t \in [t', t'']$, and it follows from the Euler equation and $c(t'; \bar{\pi}') \leq c(t'; \bar{\pi})$ that $c(t; \bar{\pi}') < c(t; \bar{\pi})$ for all $t \in [t', t'']$. Integrating (36) from t' to t'' then implies that $\hat{\mu}(t''; \bar{\pi}') < \hat{\mu}(t''; \bar{\pi})$, which is a contradiction.

But if $A = [t', \infty)$, then $i(t; \bar{\pi}') \leq i(t; \bar{\pi})$ for all $t > t'$, implying that $\dot{c}(t; \bar{\pi}')/c(t; \bar{\pi}') \leq \dot{c}(t; \bar{\pi})/c(t; \bar{\pi}) + (\bar{\pi}' - \bar{\pi})$, and hence that $\log c(t; \bar{\pi}') - \log c(t; \bar{\pi}) \leq -(\bar{\pi}' - \bar{\pi})t$, leading consumption under the two rates of trend inflation to diverge without bound. This violates the transversality condition.

I conclude that $A = \emptyset$ and hence that $\hat{\mu}(t; \bar{\pi}') > \hat{\mu}(t; \bar{\pi})$ for all $t > 0$. \square

Proof of lemma 5.5

Proof. If $M^d(i, c^*)$ has a constant semielasticity of $-b$ with respect to i , i.e. $\frac{\partial \log M^d(i, c^*)}{\partial i} = -b$, then integrating

$$\log M^d(i, c^*) = \log M^d(0, c^*) + \int_0^i \frac{\partial \log M^d(i', c^*)}{\partial i'} di' = \log m^* - bi$$

Rearranging:

$$\begin{aligned} i &= -\frac{\log M^d(i, c^*) - \log m^*}{b} \\ v'(m) = iu'(c^*) &= -\frac{\log m - \log m^*}{b} u'(c^*) \end{aligned} \quad (38)$$

Now, integrating $v'(m)$ starting from the initial condition $v(0) = 0$ gives

$$\begin{aligned} v(m) &= \int_0^m v'(m') dm' = -\frac{u'(c^*)}{b} \int_0^m \log m' - \log m^* dm' \\ &= -\frac{u'(c^*)m}{b} (\log m - \log m^* - 1) \end{aligned}$$

and \bar{m} is given by $\log \bar{m} - \log m^* = 1$. Plugging this into (38) implies $\underline{i} = -\frac{\log \bar{m} - \log m^*}{b} = -\frac{1}{b}$ as desired. \square

Verification of example 5.8

Solving for cash demand:

$$\begin{aligned} M_{hd}^d(i, c) &= \left(\frac{A_h}{\alpha_h + iu'(c)} \right)^{1/\zeta} \\ M_{ld}^d(i, c) &= \left(\frac{A_l}{\alpha_l + iu'(c)} \right)^{1/\zeta} \\ M_{bd}^d(i, c) &= M_{hd}^d(i, c) + M_{ld}^d(i, c) \end{aligned}$$

It follows that $M_{hd}^d(0, \cdot) / M_{ld}^d(0, \cdot) = \left(\frac{A_h}{\alpha_h} / \frac{A_l}{\alpha_l} \right)^{1/\zeta} > 1$, and also

$$\frac{M_{ld}(i, c)}{M_{hd}(i, c)} = \left(\frac{A_l}{A_h} \right)^{1/\zeta} \left(\frac{\alpha_h + iu'(c)}{\alpha_l + iu'(c)} \right)^{1/\zeta}$$

which is strictly increasing in i since $\alpha_h < \alpha_l$. Hence $M_{bd}(i, c) / M_{hd}(i, c) = (M_{hd}(i, c) + M_{ld}(i, c)) / M_{hd}(i, c) = 1 + M_{ld}(i, c) / M_{hd}(i, c)$ is strictly increasing as well, and the conditions of assumption 5.7 are satisfied.