

Two Illustrations of the Quantity Theory of Money Reloaded¹

Han Gao (University of Minnesota)

Mariano Kulish (University of Sydney)

Juan Pablo Nicolini (Federal Reserve Bank of Minneapolis/Universidad Di Tella)

DOCUMENTO DE TRABAJO N° 162

Julio de 2022

¹ Publicado en Federal Reserve Bank of Minneapolis
<https://doi.org/10.21034/sr.633>

Los documentos de trabajo de la RedNIE se difunden con el propósito de generar comentarios y debate, no habiendo estado sujetos a revisión de pares. Las opiniones expresadas en este trabajo son de los autores y no necesariamente representan las opiniones de la RedNIE o su Comisión Directiva.

The RedNIE working papers are disseminated for the purpose of generating comments and debate, and have not been subjected to peer review. The opinions expressed in this paper are exclusively those of the authors and do not necessarily represent the opinions of the RedNIE or its Board of Directors.

Citar como:

Gao, Han, Mariano Kulish y Juan Pablo Nicolini (2022). Two Illustrations of the Quantity Theory of Money Reloaded. *Documento de trabajo RedNIE N°162*.

Two Illustrations of the Quantity Theory of Money Reloaded*

Han Gao[§] Mariano Kulish[†]
University of Minnesota University of Sydney

Juan Pablo Nicolini[‡]
Federal Reserve Bank of Minneapolis
and Universidad Di Tella

December 15, 2021

Abstract

In this paper, we review the relationship between inflation rates, nominal interest rates, and rates of growth of monetary aggregates for a large group of OECD countries. If persistent changes in the monetary policy regime are accounted for, the behavior of these series maintains the close relationship predicted by standard quantity theory models. With an estimated model, we show those relationships to be relatively invariant to alternative frictions that can deliver quite different high-frequency dynamics. We also show that the low-frequency component of the data derived from statistical filters does reasonably well in capturing these regime changes. We conclude that the quantity theory relationships are alive and well, and thus they are useful for policy design aimed at controlling inflation.

JEL Classification: E41, E51, E52.

Keywords: Money Demand, Monetary Aggregates, Monetary Policy.

*We thank V.V. Chari, Jonathan Heathcote, Andy Neumeyer, Pedro Teles, and Martin Uribe for helpful comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[§]1925 Fourth Street S, Minneapolis, MN, 55455, United States, gaoxx825@umn.edu

[†]Room 604, A02 Social Sciences, School of Economics, University of Sydney, Sydney, NSW, 2006, Australia, mariano.kulish@sydney.edu.au

[‡]90 Hennepin Avenue, Minneapolis, MN, 55401, United States, juanpa@minneapolisfed.org

“The four incidents we have studied are akin to laboratory experiments in which the elemental forces that cause and can be used to stop inflation are easiest to spot.”

“The Ends of Four Big Inflations” (p. 90), Thomas Sargent, 1982.

“These methods will yield clear results only if a good enough ‘experiment’ has been run by ‘nature’ over the sample period used.”

“Two Illustrations of the Quantity Theory of Money” (p. 1013), Robert E. Lucas Jr, 1980.

1 Introduction

In the mid '60s, the inflation rate in several OECD countries, which was relatively low at the end of the 50s, started to increase, reaching its highest value by the late '70s or early '80s, depending on the country. It then went down, and by the early 2000s, it reached the low levels that prevailed by the early '60s. [Figure 1](#) summarizes the rising inflation and its subsequent conquest. It depicts the average inflation rate for several OECD countries from 1960 to 2005, together with a one-standard-deviation band.




Figure 1: Average inflation for 13 OECD countries (1960–2005)
The 13 OECD countries are the USA, Australia, Canada, Germany, Denmark, Italy, Japan, Netherlands, New Zealand, Portugal, South Korea, Spain, and the UK.

In this paper, we argue that the two main laws of the quantity theory of money can explain the most relevant movements in the inflationary experiences depicted in [Figure 1](#). To do so with precision, we need to define what we mean by both “quantity theory of money” and “most relevant,” and we will do so in the paper. As a first approximation, we take the two main laws of the quantity theory to mean that: i) there is a one-to-one relationship between the nominal interest rates and the inflation rate, and ii) there is a one-to-one relationship between the growth rate of money and the inflation rate — once changes in the nominal interest rate are accounted for. As monetary policy can control the evolution of either the quantity of money or the short-term interest rate, the two main laws imply that the “most relevant” movements in inflation are explained by monetary forces.

As emphasized in the two opening quotes, in order to identify the quantity theory laws in the data, one needs natural experiments in which the underlying monetary policy regime changes over time. To put it differently, if the target for inflation set by the monetary authorities, either explicitly or implicitly, does not change over time, it will not be possible to detect the quantity theory implications. The reason is that above and beyond the monetary policy regime put in place, there are various shocks that affect the price level in any economy, among other variables. The properties of these shocks, together with the functioning of markets, determine the unconditional distribution of the inflation rate in equilibrium. The main tenet of the quantity theory is that the unconditional mean of that distribution can be uniquely pinned down by monetary policy.

This logic explains several exercises that have been performed in the literature. The first is Sargent’s (1982) study of four hyperinflations in Europe following WWI, which were abruptly ended through changes in policy regimes.¹ As with all successful attempts to end hyperinflations, there is ample independent evidence regarding the time and nature of the policy regime change, which explains why Sargent says they are “akin to a laboratory

¹Sargent went one step further, and explained the changes in the monetary policy regimes as the result of changes in the fiscal policy regimes in those countries. We will not attempt to explain why there were regime changes in the cases we analyze.

experiment” in the first opening quote of our paper.

A second exercise is the cross country analysis, which averages data over long periods of time, and includes in the sample countries with different monetary policy regimes. It can be found in [Vogel \(1974\)](#), [McCandless and Weber \(1995\)](#) and, more recently, [Teles et al. \(2016\)](#).

The third one is the strategy proposed in [Lucas \(1980\)](#), from which our second opening quote is taken. We heavily borrow from [Lucas \(1980\)](#), as attested to by the title of this paper. His strategy is to use a filter, with the hope that by removing from the data the effect of short-lived policy reactions to short-lived shocks, the effect of the more persistent regime change will dominate the fluctuations in inflation. Lucas defends this strategy as follows:

“One could in principle test the neoclassical laws by deriving their implications for the parameters of a structural econometric model. This course, while attractive in theory (since it broadens considerably the class of data which might shed light on the laws), is in practice a difficult one, since it involves nesting the two hypotheses in question within a complex maintained hypothesis, which must be accepted as valid in order to carry out the test. The virtue of relatively atheoretical tests ... is that they correspond to our theoretically based intuition that the quantity theoretic laws are consistent with a wide variety of possible structures.”

The evidence presented in [Figure 1](#) strongly suggests monetary policy regimes that evolved over time and is therefore fertile land for our exploration.

We go beyond the analysis in Lucas in several ways. First, we use a different filter, which gives us a precise definition of what we mean by “the most relevant” movements in inflation. We also significantly extend the sample by adding four more decades of data and several other countries that to some extent are comparable to the United States.

Second, and more importantly, we also move in the direction that Lucas disregards in the previous quote: we test the quantity theory implications in a particular structural

model. Specifically, we conduct a structural estimation of a New Keynesian monetary model on United States data, in which we allow, but do not impose, monetary policy regime changes. A regime change in monetary policy can be estimated quite precisely. Using the estimated model and estimated magnitude of the regime changes, we show that if simulated data are filtered as the real data are, the filter captures the estimated regime change reasonably well. We also show that frictions in the setting of prices, typically present in New Keynesian models, play a minimal role in explaining those fluctuations.

Our methodology identifies changes in the policy regime that resemble laboratory experiments in the sense of [Lucas \(1980\)](#) and [Sargent \(1982\)](#). More specifically, we argue that the regime change becomes visible after monetary policy cycles are removed from the data. In the United States, monetary policy cycles typically last between two and six years, with an average duration of three and a half years. Once those monetary policy cycles are filtered away, the quantity theory relationships emerge. But this, of course, need not be the case; if there were no regime change, removing monetary policy cycles would not reveal the quantity theory relationships.² Our results appear to be in contrast with Friedman’s famous “long and variable lags” description of the effects of monetary policy changes, since we apply the same filter uniformly to all countries and there is a strong resemblance between data and theory in most cases.

Overall, we believe that the joint evidence offered by the two complementary exercises provides support to the notion that the most relevant movements in inflation are caused by monetary forces.

Our results challenge prevailing narratives of the high inflation episode in the ’70s and ’80s in the United States, as well as in some other countries we analyze. These accounts base their explanations on the role of the Phillips curve in shaping inflation but do so with models in which the inflation target remains fixed for the whole sample. An example is the pioneering analysis in [Galí and Gertler \(1999\)](#), which opened up a large and influential

²This appears to be the case in Germany.

empirical literature — see, for instance, [Smets and Wouters \(2007\)](#). According to this narrative, the evolution of post-WWII inflation in developed economies is the result of real shocks interacting with frictions in the setting of prices, rather than the monetary forces of the quantity theory.³ And the substantial losses of output observed during the early '80s were key drivers of the fall in inflation.

By contrast, the results of our paper imply that inflation fell as a consequence of a regime switch in monetary policy, while price frictions played no substantive role. Our interpretation is therefore consistent with and makes an explicit case for the main argument in [Hazell et al. \(2020\)](#), in which time-variation in long-run inflation expectations is central in driving medium-term inflation. Our analysis, like the one in [Hazell et al. \(2020\)](#), is also consistent with the view that in the United States, there was no “missing inflation” during 2008 and 2009 and no “missing deflation” after 2012.

Our results are also in line with [Ireland \(2007\)](#), who estimates a cashless model with no regime change; however, it allows shocks to permanently affect the inflation target. They are also in line with the analysis of [Uribe \(2020\)](#), which makes a critical distinction between permanent and transitory shocks to monetary policy.

The paper proceeds as follows. In Section 2, we briefly discuss the theoretical model that guides our empirical analysis. In Section 3, we separate the data for a relatively large set of countries into a short-run and a long-run component, as pioneered by [Lucas \(1980\)](#) — which is alluded to in the title of this paper — and used by [Benati \(2009\)](#) and [Sargent and Surico \(2011\)](#), among others. Our second exercise is presented in Section 4, in which we estimate a New Keynesian model that is standard, except we allow for regime changes in the inflation target and explicitly model and estimate a real money demand. A concluding section provides a broad discussion of the policy implications of the evidence discussed in the paper.

³A key role is played by cost-push shocks in this narrative.

2 The model

We study a labor-only representative agent economy with uncertainty, in which making transactions is costly.⁴ The preferences of the representative agent are

$$E_0 \sum_{t=0}^{\infty} \beta^t U(x_t), \quad (1)$$

where x_t is consumption at date t and U is differentiable, increasing, and concave. The goods production technology is given by

$$y_t = x_t = z_t l_t,$$

where l_t is time devoted to the production of the final consumption good and z_t is an exogenous stochastic process. Each period, the representative agent is endowed with a unit of time, with l_t used to produce goods and $1 - l_t$ used to carry out transactions.

We assume that households choose the number n of “trips to the bank,” in the manner of the classic Baumol-Tobin model. Thus, purchases over a period are then subject to a cash-in-advance constraint

$$P_t x_t \leq M_t n_t, \quad (2)$$

where M_t is money and n_t is the velocity of money.

We assume that the cost of going to the bank is linear in the number of trips, as in the Baumol-Tobin case, according to

$$\theta n_t \nu_t,$$

where θ is a positive parameter and ν_t is an exogenous stationary stochastic process. Total time available for production is therefore

$$l_t = 1 - \theta n_t \nu_t,$$

⁴The model is a special case of the one developed in detail in [Benati et al. \(2021\)](#).

so consumption must satisfy

$$x_t = z_t(1 - \theta n_t \nu_t). \quad (3)$$

In Appendix A, we show that as long as the cash-in-advance constraint (2) is binding, the optimal solution for real money balances is given by

$$\frac{m_t}{x_t} = \sqrt{\frac{\theta \nu_t}{i_t}}. \quad (4)$$

This relationship between real money balances as a proportion of output and the nominal interest rate in bonds is the celebrated squared root formula derived by [Baumol \(1952\)](#) and [Tobin \(1956\)](#).

Assuming that the cash-in-advance constraint is binding is quite reasonable for the countries in the period considered in [Figure 1](#), with the possible exception of Japan, which since 1995 has had near-zero interest rates. We discuss the case of Japan post-1990 in a separate subsection in which we also study the period of very low interest rates in a few other countries following the financial crisis of 2008-09.

We also show in Appendix A that in equilibrium, it must be the case that

$$E_t \left[\frac{1}{1 + r_{t+1}} \frac{1}{1 + \pi_{t+1}} \right] = \frac{1}{1 + i_{t+1}}, \quad (5)$$

where r_{t+1} is a measure of the real interest rate.⁵ This last expression is the well known Fisher equation relating the nominal interest rate with the real interest rate and the inflation rate.

In summary, the theory delivers two equilibrium relationships, (4) and (5), which involve three endogenous variables: the inflation rate, the rate of money growth relative to output, and the nominal interest rate. These two conditions do not fully characterize the

⁵This real interest rate is measured in terms of marginal utilities of real wealth, using the indirect utility function. In Appendix A, we show how this relates to a real interest rate measured in units of consumption, rather than in wealth.

behavior of prices, money and interest rates. Conspicuous by its absence is a description of monetary policy. This was a conscious choice, since according to the theory, the two implications ought to hold independently of whether the central bank adopts a money rule or an interest rate rule. Thus, to validate the empirical performance of the two illustrations in the next section, we do not need to take a stand on how monetary policy is executed.

3 Empirical analysis

In order to obtain an explicit solution for inflation, we take log differences in (4) to obtain a relationship between inflation π , the growth rate of the nominal quantity of money μ , the growth rates of output and the nominal interest rates, denoted as g^x and g^i , plus the growth rate of the shock:

$$\pi_{t+1} = \mu_{t+1} - g_{t+1}^x + \frac{1}{2}g_{t+1}^i - \frac{1}{2}g_{t+1}^\nu. \quad (6)$$

The left-hand side and the first three terms of the right-hand side are observable. They correspond to our measures of inflation, money growth, output growth, and the growth rate of the short-term interest rate.

Our theoretical assumptions (Baumol and Tobin's assumptions, really) pin down the coefficients on the right-hand side, so there is no room for parameter estimation in this exercise. Equation (6) departs from most of the previous papers that focused on the low frequency, like [Lucas \(1980\)](#) and [Benati \(2009\)](#), which set the value of the interest rate elasticity to zero, rather than to 1/2, as the Baumol-Tobin model implies.⁶

Equation (5) requires some additional manipulation. First, we use a log-linear approximation to write it as

$$i_{t+1} = r_{t+1} + E_t \pi_{t+1},$$

⁶Had we followed their strategy, the fit of the model to the data would have worsened for most of the countries we analyze below.

which involves an expectation term. But we can write

$$\pi_{t+1} = E_t \pi_{t+1} + \xi_{t+1}^\pi,$$

where ξ_{t+1}^π is a zero-mean shock, independent from any of the variables in the information set at time t , since they are expectational errors. Thus, for the empirical implementation, we use

$$i_{t+1} = \pi_{t+1} + r_{t+1} - \xi_{t+1}^\pi, \tag{7}$$

and we treat the ξ_{t+1}^π as unobservable. The nominal interest rate on the left-hand side of equation (7) is observable. However, since the availability of index bonds is very limited in practice, we do not have direct observations on the real interest rate; this lack poses a problem in testing the empirical implications of this equation. In order to proceed, we will make the following assumption.

Assumption 1 (Integrated capital markets): During the period under consideration and for the countries analyzed, there were no restrictions on capital movements, so real interest rates should be the same across countries.

Assumption 1 is clearly problematic, since it requires, among other things, the risk of default to be the same for all countries. It also requires differences in the treatment of capital income taxes across all these countries not to create wedges between the return to capital across countries. It is also particularly incorrect for the period before the '80s, when capital controls were the norm around the world.⁷

In spite of its problems, Assumption 1 has a practical advantage: we can use data for the USA, assume that the Fisher equation holds, and use US data plus equation (7) to estimate a real interest rate. Our assumption implies that we can use that real interest rate to test the Fisher equation in all other countries. That will be our strategy. In fact,

⁷We are also assuming that real exchange rates are constant, which is known not to be the case. The high volatility and persistence of real exchange is one of the major puzzles in international economics, so we make no attempt to correct for those changes.

as we will focus on the medium-term component, we need only to assume that deviations from perfect capital market integration are short-lived, which is a weaker assumption.

In studying particular countries, it should clearly be possible to do better. For each country, one could try to estimate real interest rates using other data, like the return to capital from national income accounts. To a large extent, our conclusion will be that improvements in the fit of the theory, while worth making on a country-by-country basis, will bring modest progress to our ability to understand the medium- and long-run behavior of inflation for this group of countries as a whole.

Our model abstracts from all sorts of plausible frictions, so it has no hope of matching high-frequency data. We therefore follow [Lucas \(1980\)](#), abandon that specific quest at the outset and use a statistical filter to remove the high-frequency component. In any event, we present below both the low-frequency component and the original data. Our eyes see the original data in a different way after observing the low-frequency component — and yours also will, we hope.

By construction, whatever one may learn from this strategy is of little use for quarter-to-quarter or even year-to-year policy questions. However, as we argue at the end of the paper, the lessons derived from this exercise are still ignored in policy debates today, 40 years after the publication of Lucas's (1980) analysis.

There is a key degree of freedom that needs to be settled: the ability to split the data between the short-run and the long-run components. [Lucas \(1980\)](#) does not take a stand on that question. He chooses a family of filters that depend on a single parameter. The higher the value of that parameter, the lower is the frequency that is extracted from the data. He then shows that the two illustrations emerge beautifully in his figures as the parameter that extracts the low frequency is increased. Lucas's paper is like a mystery movie. If you stare at the data, chaos prevails. But as the viewer moves along the sequence of plots, each retaining lower frequencies, the patterns start to emerge. By the time the reader arrives at the last plot, the two illustrations shine and order prevails over chaos.

Just like the book of Genesis.

Our paper offers just a picture: we take a stand on a particular way to split the data. This, in turn, provides a specific definition of what we mean by medium run. This definition clarifies for which policy questions our framework will not be useful and for which questions it may be.

3.1 The filter

To decompose the data, we use the Hodrick-Prescott (HP) filter, popularized by the real business cycle (RBC) literature. The filter’s decomposition between the short-run and the long-run components is controlled by a single parameter, denoted by λ . By taking a stand on the value for λ , we take a stand on a particular way to decompose the data between the “cycle” and the “trend.”

So far, we have used, imprecisely, terms like “short-run” and “high-frequency” and treated them as interchangeable. Below, we will estimate a structural model subject to monetary policy regime changes that can shift the unconditional mean of nominal variables. Each regime is covariance stationary, and so oscillations of all frequencies are present in each regime. Although we label the extracted components from the HP filter as “cycle” and “trend”—terms that are commonplace—it will become evident from the analysis of the structural model that regime changes that shift the unconditional mean get picked up by the low-frequency component of the HP filter.⁸

In order to discipline the choice of λ , we use the recent history of monetary policy in the USA. Specifically, we base our choice of λ on a particular narrative regarding the behavior of the short-term interest rate in the USA. We believe it to be a widely accepted narrative among macroeconomists. To describe it, it is useful to refer to [Figure 2\(a\)](#), which depicts the time path for the federal funds rate, as well as two computations of the low-frequency component extracted using two alternative values for λ . The relative merits of the two

⁸See [Kulish and Pagan \(2021\)](#) for a discussion on the distinction between cycles and oscillations.

values for λ are discussed in detail below. In [Figure 2\(b\)](#), we plot the two corresponding measures of the high-frequency component, obtained by subtracting from the original data the two measures of the low-frequency component in [Figure 2\(a\)](#).

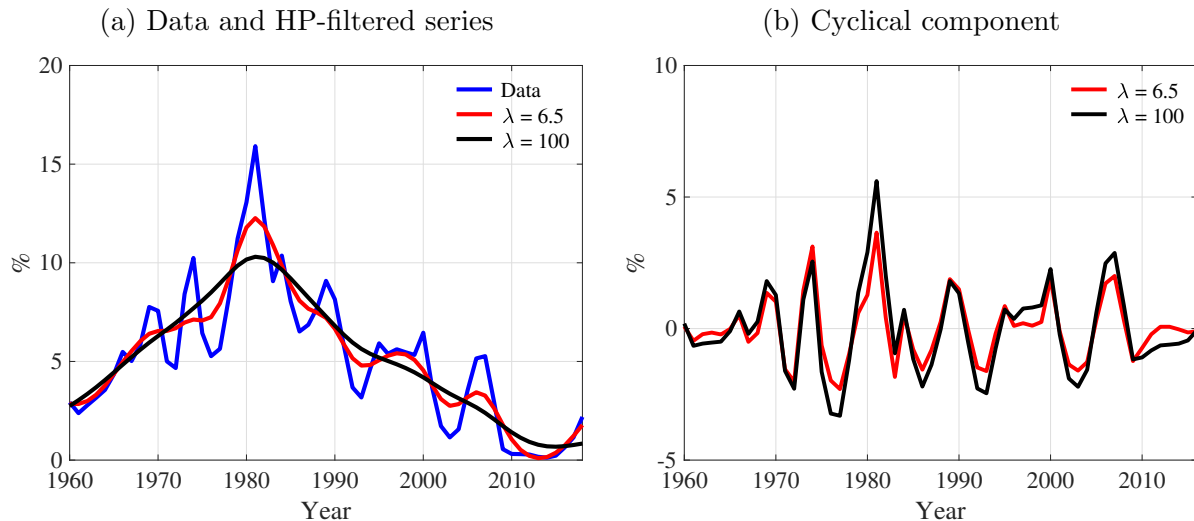


Figure 2: US nominal interest rates

The key historical element to build the narrative is the notion of a “tightening cycle.” Any such cycle is defined as a series of consecutive periods exhibiting increasing values for interest rates. These are clearly visible in [Figure 2](#), more obviously so in panel (b). Particularly famous tightening cycles are the ones known as the Volcker stabilization — starting at the end of the ’70s — and Greenspan’s conundrum — the one that starts in 2004.⁹

The narrative interprets these cycles in the interest rate as the policy response to temporary shocks, so as to stabilize the economy around certain desired values. This role of policy finds its strongest intellectual rationale in the New Keynesian literature, which emphasizes frictions in the setting of prices. In these models, price frictions generate only temporary effects on the equilibrium, which vanish “in the long run.”

Although we do not need to take a stand on how policy is executed, it is convenient to consider an interest rate policy that follows a rule as in [Taylor \(1993\)](#). Thus, let the policy

⁹The “tightening” cycles are followed by their corresponding “easing” period, in which the interest rate is decreasing.

rate be given by

$$i_t = i^* + \phi_\pi(\pi_t - \pi^*) + \phi_y(y_t - y^*) + \varepsilon_t^i, \quad (8)$$

where i_t , π_t , and y_t represent the policy interest rate, inflation, and output, respectively, and ε_t^i is a monetary policy shock. The triplet (i^*, π^*, y^*) is typically interpreted as the steady state values for the variables.

In the literature, the second and third terms on the right-hand side of the Taylor rule are meant to be the cause of the cycles described in [Figure 2\(b\)](#). They represent the attempt by the monetary authority to stabilize the equilibrium values of inflation and output around π^* and y^* . Most of the literature uses a variation of this Taylor rule, in which the triplet (i^*, π^*, y^*) is assumed to be time-invariant.¹⁰

Our interpretation of changes in the policy regime amounts to allowing for a target for inflation, denoted by π_t^* , that changes over time. These changes ought to be accompanied by the corresponding changes in the target for the interest rate, as implied by the Fisher equation. So the value for the interest rate target, i_t^* , must also be time-varying.

We prefer to adopt a policy rule as in equation (8), but in which the deviations of inflation and the interest rate are made relative to values that are changing over time.¹¹

The simple quantity theory model spelled out above, with all its simplifying assumptions, has no bearing on interest rate movements that correspond to the second and third terms in the Taylor rule.¹² This is so much so that the implied relationship between the nominal interest rate and inflation in our model — as described, for instance, in (7) — is positive and one to one. In contrast, the conventional wisdom in central banks, supported by the workings of New Keynesian models, is that increases in the nominal interest rate imply *reductions* of inflation.¹³ In this section, the quest to understand the fluctuations depicted

¹⁰For exceptions, see [Ireland \(2007\)](#), [Cogley and Sbordone \(2008\)](#), and [Ascari and Sbordone \(2014\)](#).

¹¹We make no attempt at explaining why those values change as they did during the period.

¹²In some formulations, such as [Woodford \(2003\)](#), the term $(y_t - y^*)$ in the Taylor rule is the difference between the equilibrium value for output and the one that would prevail under flexible prices — the output gap. In our simple quantity theory model, the output gap is by definition zero, so even that term disappears from the rule.

¹³See [Uribe \(2020\)](#) for a masterfull integration of these seemingly contradictory statements.

in Figure 2(b) is abandoned.

Therefore, in deciding the best value for λ , we aim to capture the slow-moving term i_t^* , while we expect the filter to remove the second and third terms in the rule, as well as temporary stochastic disturbances.

The distinction just made between deviations from a steady state, which imply a set of values that are constant over time, and deviations from a given trend is key. We address this issue in detail in the next section, in which we estimate a small-scale New Keynesian model and allow for shocks to the targets i_t^* and π_t^* . We make very precise in the model this distinction between movements that capture the tightening cycles around a trend and the ones that explain the trend, and we let the data separate the two. For the analysis of this section, we use our discussion above, plus the evidence in Figure 2, to justify our choice of λ .

This discussion explains our criterion for choosing the value for λ : the smallest value that eliminates from the data the tightening cycles. In Figure 2(a), we plot two alternative values for the low-frequency component, corresponding to values of 6.5 and 100. The first value, 6.5, is the one that the RBC literature suggests for yearly data. Its object of study is very different from ours (note the R in RBC), so there is no reason why what fits its objective should fit ours. And as can be seen in the figure, it does not: when using $\lambda = 6.5$, the tightening cycles are still visible. On the other hand, when using $\lambda = 100$, the cycles are completely removed from the policy rate.¹⁴ Therefore, in what follows, we set $\lambda = 100$. In Appendix B.3, we also show the results when using $\lambda = 6.5$.

By taking a stand on a particular value for λ , we take a stand on our definition of “medium run.” Conventional wisdom states that to see the mechanics of the quantity theory operating in the data for countries like the United States, one needs to look at averages over decades. This piece of conventional wisdom is consistent with Milton Friedman’s own

¹⁴The behavior of the low frequency obtained for values of lambda between 90 and 110 are indistinguishable to the eye, given the size of these figures. How could we resist the seductive power of a round number like 100?

view of the lags in monetary policy. For instance, in his 1970 Wincott Memorial Lecture, delivered at the University of London, he writes, “In the short run, which may be as much as five or ten years, monetary changes affect primarily output. Over decades, on the other hand, the rate of monetary growth affects primarily prices.”

Our choice of filtering implies a much tighter definition of “medium run.” We make this explicit in [Figure 2\(b\)](#), in which we plot the high-frequency component of the interest rate for the two values of the parameter in the HP filter. As expected, the fluctuations are higher when using our preferred parameter of 100. But the two series are very similar. Both identify the same number of cycles, defined as the time period contained between two consecutive crossings of the horizontal axis. Those correspond to a tightening cycle when the curve is increasing or an easing cycle when the curve is decreasing. For both measures, the average cycle is about three and a half years, with a maximum of six years and a minimum of one year, in 1967. One interpretation of the filter we use, which we adopt, is to leave out of the data all fluctuations that last about three and a half years on average, the average duration of the monetary policy cycles in the United States.

3.2 Preliminaries

We now take equations (6) and (7) to the data. We selected countries that are members of the OECD and for which we have complete data since 1960. These are the 13 countries included in [Figure 1](#), plus Colombia, Mexico and Turkey. These three countries experienced substantially higher inflation rates than the rest, so they have been left out of [Figure 1](#).

We use the short-term interest rate on government debt for i , gross domestic product for output, and the CPI for prices. For the monetary aggregate, we use $M1$, which is the sum of currency plus checkable deposits. For the United States, $M1$ provides a misleading measure of total assets available for transactions, owing to regulatory changes that occurred in the early '80s. [Lucas and Nicolini \(2015\)](#) discuss this issue in detail and propose a new measure called NewM1, which adds the Money Market Demand accounts created in 1982

to the standard measure of M1. Thus, for the USA only, we use NewM1 rather than M1. Doing so raises the issue of whether the simple model described above could account for a regime change due to the regulatory changes in the middle of the sample.¹⁵ Thus, for the USA only, we will also show the results using the currency component of M1, which according to [Lucas and Nicolini \(2015\)](#) should be relatively invariant to the regime change. In Appendix B.2, we discuss the data and their sources in detail.

The period we focus on is 1960-2005, consistent with the data in [Figure 1](#). There are a few exceptions. For the countries that joined the eurozone, accurate measures of M1 are not available after 1999, since currency in circulation cannot be properly measured. For those countries, we use data up to 1999 in evaluating the performance of the money demand equation (6).

The presence of very low interest rates presents additional theoretical considerations that are worth discussing separately. The reason is that to obtain equation (6), we assumed a binding cash-in-advance constraint. The validity of that assumption at very low rates is questionable. Thus, for Japan, we end the sample in 1990, before the country lowered its interest rate to almost zero. In a final subsection, we separately discuss our analysis's policy implications for Japan since 1990, as well as the evidence since 2005 for other countries that experienced very low interest rates. Finally, because of data availability, we start the analysis of Turkey only in 1970.

As we mentioned above, we have no independent estimate for the real interest rate in the USA. Therefore, in the case of that country, we simply plot the inflation rate and the nominal interest rate, so as to appreciate the positive correlation.

3.3 Filtering results

In the top panels of [Figure 3](#) to [Figure 6](#), we show the data corresponding to the money demand equation (6). We first plot the raw data for the inflation rate and for the growth

¹⁵The model in [Lucas and Nicolini \(2015\)](#) does imply, not surprisingly, that such a regime change should change the relationship between the nominal interest rate and the ratio of money to output.

rate of nominal money over real output. In the plots of raw data for Illustration 1, we do not make the adjustment for changes in the the nominal interest rate, as (6) implies. The reason is that this adjustment makes the theoretical prediction for inflation way more volatile than in the data, since the high-frequency movements in interest rates are very volatile and the value for the elasticity implied by the Baumol-Tobin formulation is high for high-frequency data.¹⁶ We then plot the low-frequency component for the theoretical inflation, as predicted by equation (6), together with the low-frequency component of inflation in the data. The bottom panels of Figure 3 to Figure 6 show the data corresponding to the Fisher equation (7). In all plots, we also report the correlation between the series.

The first column of Figure 3 presents the results for the United States. As mentioned above, in this case we use both Cash and NewM1 for Illustration 1. The yearly data do not make apparent the relation between money growth and inflation when using New M1. In fact the correlation is just 0.18. However, once the low-frequency movements are isolated and the effect of changes in the interest rate is taken into account, as equation (6) implies, the match between the theory and the data is quite notable, with a correlation coefficient of 0.86 when using NewM1. This is in spite of the regulatory changes in the early '80s. The match when using just cash is much better in the data and almost as good when using only the low frequency.

This reasonable match with the theory offers an alternative interpretation for the experience in the United States besides the one proposed by Sargent and Surico (2011). They replicate the analysis in Lucas (1980), using the same filter he uses. They extend the sample in Lucas (1980) to include data from 1980 till 2005. They use a monetary aggregate that is very close to M1 and show — as we do in Appendix B.1 — that the data do not align well with the theory.¹⁷ They propose a model with regime changes in the monetary

¹⁶This is consistent with the old empirical literature on money demand, which argued that the estimated “short-run” interest rate elasticity was much smaller than the “long-run” elasticity. See Lucas (1988) for a discussion.

¹⁷They also ignored — as Lucas (1980) did — the effect of the movements in the interest rate, which are important. But the main difference is the monetary aggregate they use.

Illustration 1

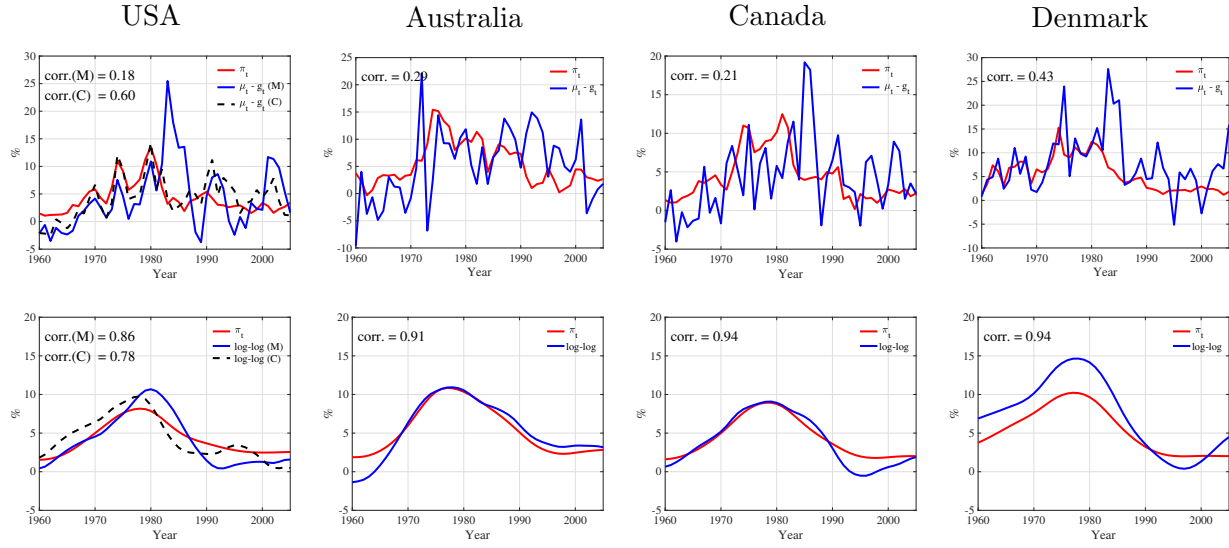


Illustration 2




Figure 3: Countries in Group 1 (a)

Illustration 1

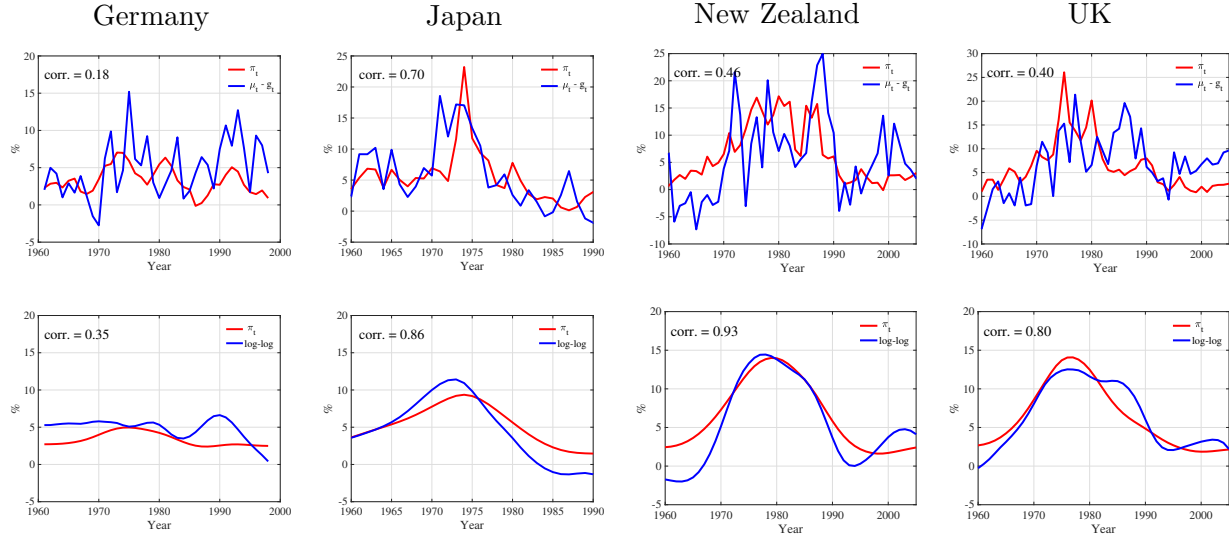


Illustration 2

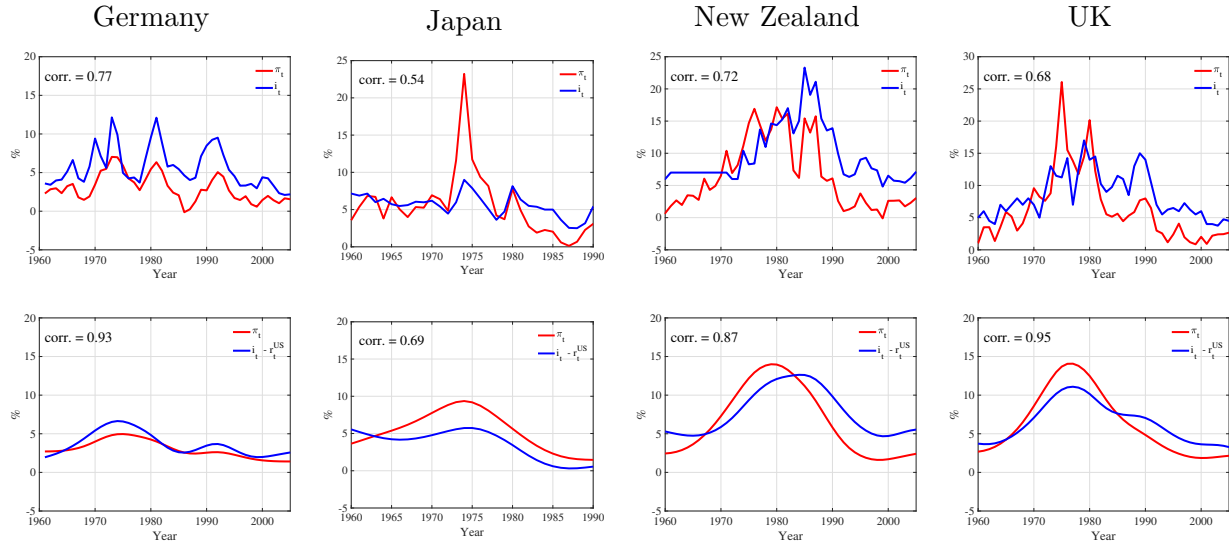


Figure 4: Countries in Group 1 (b)

Illustration 1

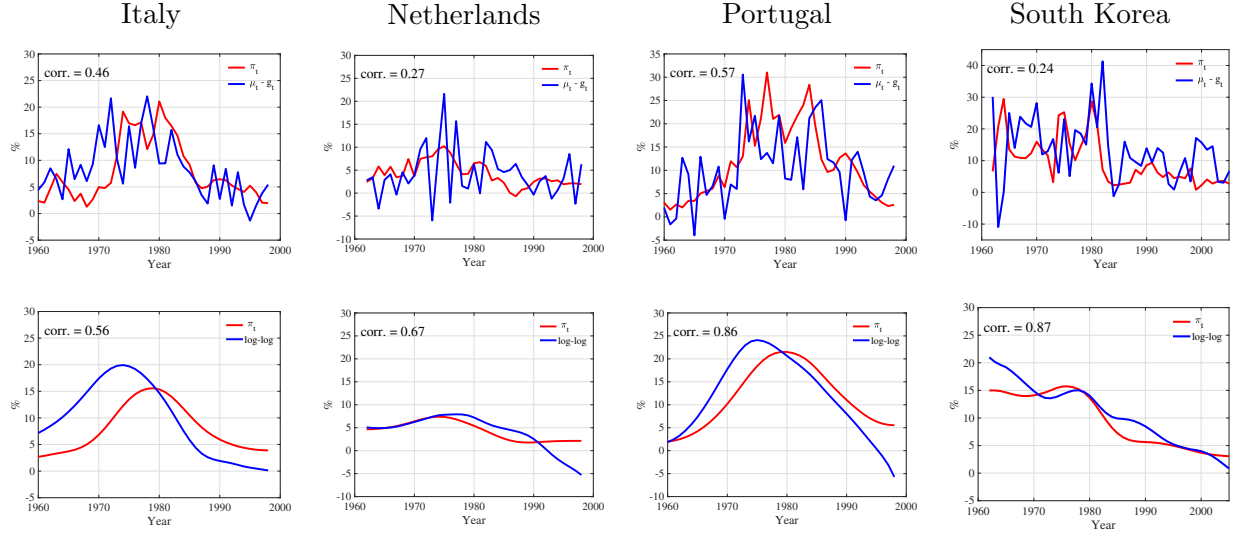


Illustration 2

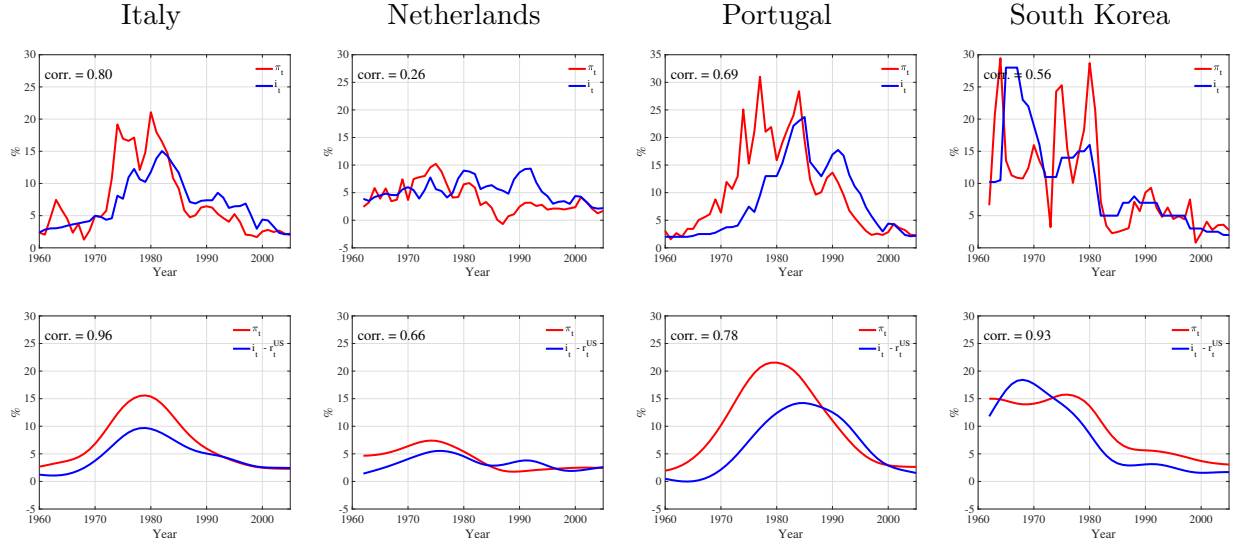


Figure 5: Countries in Group 2 (a)

Illustration 1

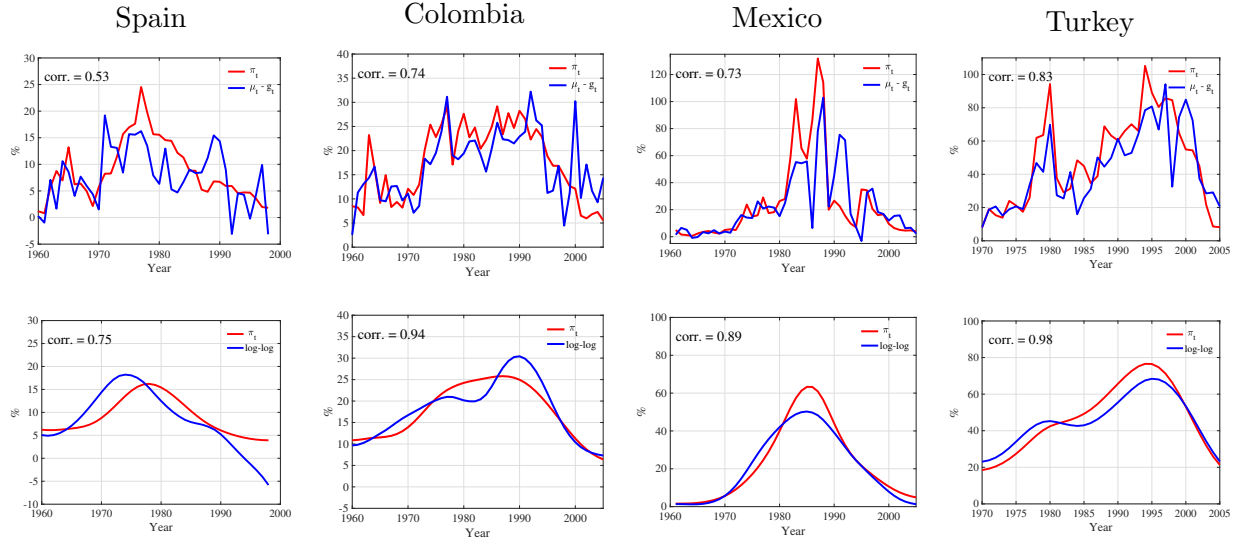


Illustration 2

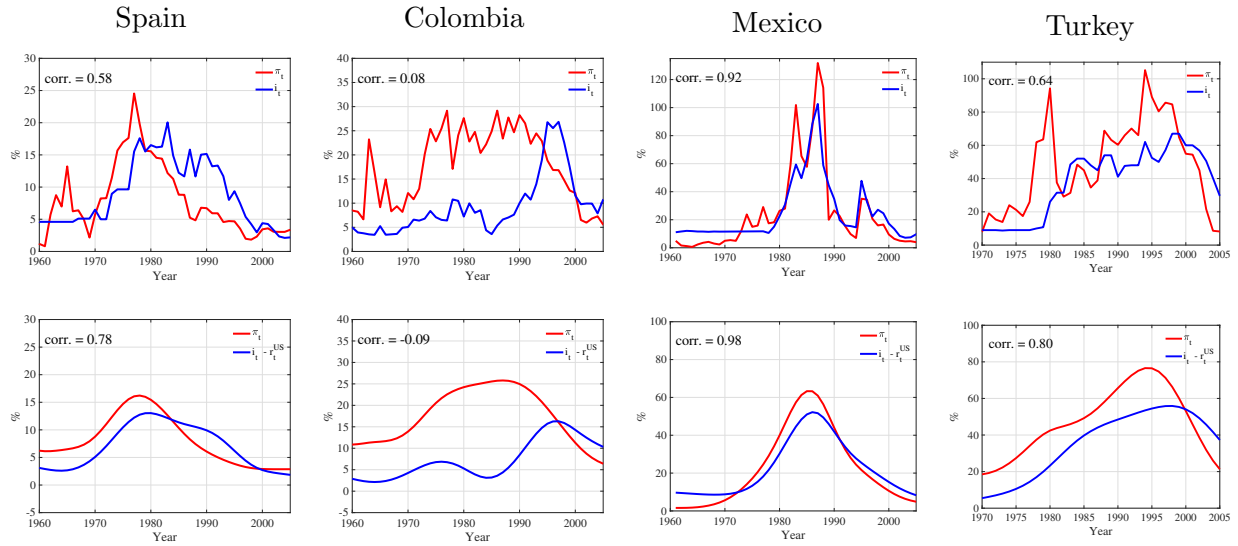


Figure 6: Countries in Group 2 (b)

policy rule to account for that failure. In using either Cash or NewM1, we show that no puzzle arises. As we show next, this phenomenon is specific to the United States. In the analysis for all the other countries that follows, we use M1 as the measure of money, as the results are notably good.

We separate the countries in two groups. The first group, shown in [Figure 3](#) and [Figure 4](#), includes the countries for which we do not find any particular behavior that makes our assumptions for the model especially suspicious. The worst fit in this group is Germany for Illustration 1, where the correlation is only 0.35. As we explain below, we interpret this failure as a lack of regime change for Germany, which transforms it into an outlier in light of the evidence of [Figure 1](#). The second group, shown in [Figure 5](#) and [Figure 6](#), includes a set of countries for which the nominal interest rate is lower than the inflation rate for several years in the first two decades of the period analyzed. To us, this behavior suggests government intervention in the credit market, relatively common in the '60s and '70s and called “financial repression” at the time. Under this condition, the observed interest rate is not a market-determined price, so it does not represent the true opportunity cost of money. This imposes a bias in our two theoretical predictions. This group also includes cases with higher inflation rates.

Each picture is worth a thousand words. As we provide plenty of pictures, words will be kept to a minimum. We read the sequence of plots as an affirmation of our simple theory, particularly when compared with other theories in the social sciences. We mostly let the readers evaluate the pictures themselves and emphasize just a few features of the plots.

First, while the correlation between the data and the theory is very high, in some cases there are sizable differences, of up to a few percentage points; these do matter for policy. A 2% or larger difference between observed inflation and the theoretical counterpart, as observed in many cases, is an important difference that can and should be further studied on a case-by-case basis. It is most likely that in order to understand those differences, country-specific features ought to be brought to the policy debate table.

Second, for the group of countries in [Figure 5](#) and [Figure 6](#), for which we guess that financial repression was prevalent in the first decades of the sample, the evidence is worse, particularly when evaluating the second implication (the Fisher equation). A poster child of this issue is Colombia, where financial repression was the norm till the reforms of the early '90s.

3.4 The near-zero nominal interest rates periods

A positive interest rate is required for the cash-in-advance constraint to be binding. When this is not the case, real money demand is not uniquely determined. In our simple representative agent economy, the result is stark: as long as the interest rate is positive, the constraint is binding and the equilibrium of the model is uniquely pinned down. However, sensitive modifications, like agent-specific borrowing limits or heterogeneous returns on nominal assets due to heterogeneous access to credit markets, would affect the implications of the model when the nominal interest rate is positive but very close to zero.¹⁸

To further clarify this discussion, consider the solution of our simple model, given by [\(4\)](#). Notice that the solution for real money balances as a fraction of output goes to infinity when the nominal interest rate goes to zero. How can that be a solution for agents that have finite wealth? The answer is that in equilibrium, the private sector's borrowing from the government is also going to infinity, keeping the wealth of the private sector bounded. While this is mathematically correct for any positive interest rate, it is of little, if any, applied interest.

To illustrate the difficulties in using real money demand theory at very low interest rates, we now analyze those countries that experienced them for a prolonged period of time. Besides the solution in [\(4\)](#), we use an alternative functional form proposed by [Selden \(1956\)](#) and [Latané \(1960\)](#) and explored in detail in [Benati et al. \(2021\)](#). The specific functional

¹⁸For an analysis with heterogeneous borrowing constraints, see [Benati et al. \(2021\)](#). An alternative model that delivers similar results is analyzed in [Alvarez and Lippi \(2009\)](#).

Illustration 1




Illustration 2




Figure 7: Illustrations for countries with periods of low interest rates (1960–2018)

form is given by

$$\frac{M_t}{P_t x_t} = \frac{A}{1 + b i_t}. \quad (9)$$

Notice that when $i_t = 0$, real money demand - as a fraction of output - is finite. Thus, it departs from the Baumol-Tobin specification at very low interest rates. On the other hand, the parameters A, b can be chosen so that (4) and (9) are very close to each other for interest rates above 2% and all the way up to 30%, a range that includes all the experiences we now discuss.

In Figure 7, we extend the analysis presented in Figure 3 and Figure 4 to several countries in Group 1 that maintained an interest rate very close to zero for several periods. As mentioned above, in computing theoretical inflation, we present both the log-log case used before and the Selden-Latané case, in which the parameters have been chosen to match as much as possible the solution in (4).¹⁹

¹⁹Specifically, we choose parameter $b = 0.14$ for the Selden-Latané specification. The value for A is irrelevant for computing growth rates, as we do in this section.

As the figure shows, when using the log-log specification, the implications of money demand theory become really off the mark during the periods of very low interest rates. The Selden-Latané alternative specification does better but still fails to perform as it did in the previous years. [Figure 7](#) suggests that inference about the behavior of real money demand at very low rates using evidence of periods with relatively higher rates, when the cash-in-advance constraint can be safely assumed to be uniformly binding, could be misleading. Thus, monetary aggregates may be uninformative at very low rates.

Note, on the other hand, that the evidence regarding the second illustration is as good for the low interest rate period as it is for the rest of the sample.

Two policy implications follow. First, when interest rates are very low, the effect of expansions of the the central bank’s balance sheet on the real side of the economy — the so called “unconventional policies” — is hard to predict, since it is hard to estimate the demand for those assets. Second, changing the low-frequency component of the the policy rate in a way that resembles a positive shock to the target — as the Federal Reserve did between 2015 and 2018 — can act as an effective tool to fight persistently low inflation, even at the lower bound.

4 An estimated model with monetary regime changes

The filter used above allowed us to stay away from taking a stand on what drives short-run fluctuations in inflation. The advantage, of course, is that this analysis can be applied to any theory in which the effect of frictions does not last more than three years on average. But doing so prevents us from having a full description of the data and from understanding our filter’s ability to capture underlying regime changes.

Therefore, in this section, we move in the direction that [Lucas \(1980\)](#) discussed but did not pursue and estimate a small-scale New Keynesian model. In order to evaluate Illustration 1, we do not consider the money-less limit and use data on money in estimation.

Our main point of departure from the literature is that we allow for policy regime changes that originate from a time-varying inflation target. We let the data speak about the timing and magnitude of these regime changes; in other words, we let the data inform us about how important these regime changes are in explaining movements in inflation. For reasons of space and data availability, we focus only on the USA. As we show below, we clearly identify a regime change that started with a gradual increase of the inflation target in the late '60s. A new regime change occurred with a gradual decrease of the inflation target in the early '80s.

Looking through the lens of the estimated model, we think of the data as being generated from different regimes; in our case, from regimes with different unconditional means for the inflation target. Any one regime is covariance-stationary, and so oscillations of all frequencies are present. A change of the unconditional mean is a change of the zero frequency of the process. Applying the HP filter, as we did in Section 3, to time series data generated by two or more regimes with different means gives rise to a low-frequency component that does not really belong to any one regime. So one should not interpret the low-frequency component as literally belonging to a single regime; rather, our aim is to assess quantitatively the extent to which the low-frequency component of our filter manages to capture the estimated monetary regime changes. This section integrates the quantity theory, the low frequency components of Section 3, and the estimated magnitude of regime changes.

We use the estimated model to simulate data and filter them the same way we filtered the data in the previous section. We then compare the results of the same model but with the regime change shut down and again filter the simulated data. The comparison between the two exercises makes clear that the policy regime change is essential in explaining what appear as low-frequency movements in inflation, interest rates and money growth.

We repeat the exercise but vary the degree of price frictions. We show that the price frictions barely change the implications regarding the low-frequency behavior of inflation,

interest rates, and money growth, which in all cases is explained by the policy regime shocks. We interpret these exercises as evidence that the strength of the price frictions in the model does not change the medium-run implications of the simulated data. We see all this evidence as a validation of our filtering choice, since it is the case that when applying the same filter to simulated data, we always remove the tightening cycles.

Our analysis is consistent and complementary to the one in [Uribe \(2020\)](#), which integrates the two effects within a single theoretical model, considering both temporary and permanent shocks to the policy rate. In his estimated model, permanent changes in monetary policy have an almost immediate effect on inflation, with almost no effect on output. Our analysis is also complementary to that of [Ireland \(2007\)](#), which allows shocks to have a permanent effect on the inflation target.

We now briefly describe the model, discuss the estimation strategy and present the results. A full description of the estimation and a detailed analysis of the simulation exercises is in Appendix C.

4.1 The model

For the analysis that follows, we extend the simple model in Section 2 to include sticky prices, endogenous leisure, and preference, markup and monetary policy shocks – essentially the workhorse New Keynesian model used by [Ireland \(2004\)](#) and [Sargent and Surico \(2011\)](#), complemented with a real money demand function.²⁰ Our only point of departure is to allow for shocks to the policy rule target.

In contrast with the previous section, to estimate the model, we need to take a stand on how monetary policy is conducted. In theory, the evolution of the most relevant part of inflation depends on the monetary policy regime, but not on the instrument used to implement such a regime. This is also the case in practice, since we estimate the model using an interest rate rule first and a money rule second. The estimated policy regime

²⁰Details of the non-linear model can be found in [Ireland \(2004\)](#).

change is essentially the same one, independent of the policy instrument used. In the Online Appendix D, we show the detailed results when we specify policy by means of a money growth rule. In what follows, we present results when following the New Keynesian tradition of specifying policy by means of an interest rate rule. For purposes of comparison, we also show some results under the money rule.

The model under interest rate rule is composed of the familiar Euler equation (which is the Fisher equation of our simple model of Section 2), the New Keynesian Phillips curve, and the Taylor rule shown below:

$$x_t = (z - \ln \beta) - (i_t - \mathbb{E}_t \pi_{t+1}) + \mathbb{E}_t x_{t+1} + (1 - \omega)(1 - \rho_a)a_t \quad (10)$$

$$\pi_t = (1 - \beta)\pi^s + \beta \mathbb{E}_t \pi_{t+1} + \psi x_t - e_t \quad (11)$$

$$i_t = i_t^* + \rho_i (i_{t-1} - i_{t-1}^*) + \phi_\pi (\pi_t - \pi_t^*) + \phi_x x_t + \varepsilon_{i,t}. \quad (12)$$

In the equations above, x_t is the output gap, π_t is the log of the gross rate of inflation, and i_t is the log of the gross nominal interest rate.

We differ from [Ireland \(2004\)](#) in that we allow for the inflation target, π_t^* , and the corresponding target for the nominal interest rate i_t^* to depend on time, as shown in equation (12). Specifically, we assume that

$$\pi_t^* = (1 - \rho_\pi)\pi^s + \rho_\pi \pi_{t-1}^* + \mathbb{I}^s \varepsilon_{\pi^*,t} \quad (13)$$

$$i_t^* = z - \ln \beta + \pi_t^*. \quad (14)$$

According to equation (14), the implied target for the nominal interest rate, i_t^* , is determined by the steady state real interest rate, $z - \ln \beta$, and the inflation target, π_t^* . The variable z is the steady state growth rate of labor-augmenting productivity, which follows in logs a unit root with drift z , and β is the household's discount factor. The inflation target, π_t^* , follows a regime-dependent AR(1) process in which \mathbb{I}^s is an indicator variable that is

turned on at T^{on} and then turned off at T^{off} ; that is,

$$\mathbb{I}^s = \begin{cases} 1 & \text{for } t \in [T^{\text{on}}, T^{\text{off}}) \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

At the start of the sample (1960Q1), we set $\mathbb{I}^s = 0$ and $\pi^s = 0.005$, which is equivalent to an inflation target of 2% in annualized terms. Before T^{on} , shocks to the inflation target are turned off, and the model is a standard New Keynesian model with a constant inflation target.

At T^{on} , the inflation target changes in two ways. First, $\mathbb{I}^s = 1$, so $\varepsilon_{\pi^*,t}$ now affect the inflation target, π_t^* . Second, we allow in estimation, but do not require, the long-run inflation target — that is, π^s — to change from $\pi^s = 0.005$ to $\pi^s = 0.005 + \Delta_\pi$, and Δ_π is to be estimated. Thus, at time T^{on} , the inflation target is subject to a permanent shock, Δ_π , and to persistent but temporary shocks, $\varepsilon_{\pi^*,t}$, until T^{off} . Notice that because of the persistence of the inflation target process, ρ_π in equation (13), the long-run inflation target is reached gradually. Finally, at time T^{off} policy reverts to its original regime; that is, $\mathbb{I}^s = 0$ and $\pi^s = 0.005$.

This choice allows for a potentially slow-moving component that pushes up inflation during the first two decades, capturing the rise in inflation in the '70s, with a reversion to the original 2% per year inflation rate observed since the '80s. In estimating the model, we let the data choose the values for the key five parameters, $\{\Delta_\pi, \rho_\pi, T^{\text{on}}, T^{\text{off}}, \sigma_{\varepsilon_\pi}\}$, where σ_{ε_π} is the standard deviation of the shock $\varepsilon_{\pi t}$ in (13).

The economy is subject to the following non-policy shocks: a preference shock, a_t ; a markup shock, e_t ; a money demand shock, ξ_t ; and a technology shock, z_t . These are

governed by the equations below:

$$a_t = \rho_a a_{t-1} + \varepsilon_{a,t} \quad (16)$$

$$e_t = \rho_e e_{t-1} + \varepsilon_{e,t} \quad (17)$$

$$\xi_t = \rho_\xi \xi_{t-1} + \varepsilon_{m,t} \quad (18)$$

$$z_t = z + \varepsilon_{z,t}. \quad (19)$$

Real money demand m_t follows

$$m_t = \bar{m} + \rho_m m_{t-1} - (1 - \rho_m) \eta \left(\frac{1 + i^s}{i^s} \right) i_t + \xi_t. \quad (20)$$

4.1.1 The model with a money rule

Note that the model with interest rate rule is block recursive in the sense that conditional on the shocks, inflation, interest rate, and output gap are determined by equations (10), (11), and (12). Money is thus determined by equation (20) outside the equations block.

Under a money rule, the model is not block recursive anymore. But we can combine the Euler equation (10) with the money demand equation (20) to eliminate the interest rate. The rest of the system can then be used to solve for inflation, money growth and the output gap. As we did when we considered interest rate rules, we allow for regime changes in the target for money growth. Thus, μ_t^* evolves according to

$$\mu_t = \mu_t^* + \rho_\mu (\mu_{t-1} - \mu_{t-1}^*) + \theta_\pi (\pi_t - \pi_t^*) + \theta_x x_t + \varepsilon_{\mu,t} \quad (21)$$

$$\mu_t^* = (1 - \rho^*) \mu^s + \rho^* \mu_{t-1}^* + \mathbb{I}^s \varepsilon_{\mu^*,t}, \quad (22)$$

where ρ_μ is the response to deviation of money growth, and θ_π and θ_x capture the responses of policy to inflation deviations and output gaps. To solve the model, we use equation (22) as a replacement of equation (13). The regime indicator function \mathbb{I}^s is defined in same way

as in equation (15), and the implied target inflation is given by

$$\pi_t^* = \mu_t^* - z. \quad (23)$$

4.2 Estimation

4.2.1 Estimation strategy

We estimate the dates of regime change, T^{on} and T^{off} , alongside the structural parameters, following the method outlined by [Kulish and Pagan \(2017\)](#). The model is estimated with five observable series: real GDP per capita growth; the federal funds rate; core inflation as measured by the CPI, excluding food and energy; the Michigan survey measure of inflation expectations; and money growth. For the United States, as discussed above, we use NewM1, the monetary aggregate proposed in [Lucas and Nicolini \(2015\)](#).²¹

The equations linking the observable variables; output growth, g_t ; and money growth, μ_t , to the endogenous variables are given by

$$g_t = \hat{y}_t - \hat{y}_{t-1} + z_t \quad (24)$$

$$x_t = \hat{y}_t - \omega a_t \quad (25)$$

$$\mu_t = m_t - m_{t-1} + \pi_t + g_t \quad (26)$$

$$\mathbb{E}_t^{\text{obs}} \pi_{t+1} = \frac{1}{4} \left(\sum_{j=1}^4 \mathbb{E}_t \pi_{t+j} \right) + v_t, \quad (27)$$

where \hat{y}_t is the percentage deviation of stochastically detrended output, Y_t/Z_t , from its steady state; μ_t is money growth; and $m_t = \ln(M_t/P_t Y_t)$ is the log of real money balances to output. The constant \bar{m} pins down real money balances to output in steady state. We use the Surveys of Consumers from the University of Michigan as the measure of inflation expectations, $\mathbb{E}_t^{\text{obs}} \pi_{t+1}$, and allow for measurement error, v_t .

²¹The same results are obtained if we use the cash component of M1, which was not much affected by the regulatory changes on the '80s, as can be seen in the top-left panel of [Figure 3](#).

In steady state, $\pi_t = \pi_t^* = \pi^s$, $i_t = i^s$, $g_t = z$, and $i^s = \pi^s + z - \ln \beta$; all other variables (including the output gap) settle on zero. The reason nominal variables are left in levels, as opposed to percentage deviations from steady state, is that in estimation we allow for changes in the steady states of these variables.

We estimated the model treating the regime changes as unanticipated. This seems to us a reasonable choice, particularly for the shock T^{on} : it is conceivable that the breakdown of the Bretton Woods system and the inflation that ensued took most by surprise. It is less plausible that the disinflation shock, T^{off} , was a surprise. However, we allow the change in target to be very slow-moving by allowing for the autoregressive component in (13). A high value for ρ_π implies that the economy slowly approaches the new long-run target. The estimation does deliver a very high value for ρ_π , so although Δ_π is unanticipated, the transition path that it triggers for π_t^* towards its new steady state is anticipated.

To guard against the possibility that our proposed policy regime change captures the higher macroeconomic volatility before the Great Moderation, we use a parsimonious specification and introduce the parameter κ , which multiplies the standard deviations of all structural shocks, except that of money demand, before T_κ . That is, the standard deviation of structural shock i is given by $\kappa\sigma_i$ before T_κ and shifts to σ_i at T_κ . The standard deviations of all shocks change by the same proportions. By adding this “great moderation” shock, the estimation is free to rely on shocks other than the inflation target shock to account for the increased volatility in the earlier part of the sample.

The parameters that determine the steady state of output growth, the nominal interest rate, inflation and the ratio of money to output are set prior to estimation. In particular, we set $\beta = 0.9975$, $z = 0.0044$, $\pi^s = 0.005$ and $\bar{m} = 1$. Jointly, they imply a mean growth rate of real GDP per capita of 1.8% in annual terms, a mean nominal interest rate of 4.75%, an inflation rate of 2% in annual terms, and a ratio of money to output of about 25% in annual terms. We set the slope of the NK Phillips curve to $\psi = 0.3$, which in our model depends on a quadratic price adjustment cost. One way to interpret this slope is by

considering a version of the NK model with a Calvo price friction. With log-utility and linearity in hours worked, a value of $\psi = 0.3$ in the standard NK model would correspond to a parameter ζ of 0.6, consistent with the findings of [Fitzgerald et al. \(2020\)](#).²²

At the mode, we estimate a value for κ of 2, implying that the volatility of structural shocks halved after T_κ , which at the mode is estimated precisely around 1985Q1.

4.2.2 The regime change

The two parameters characterizing the policy regime change are the persistence parameter ρ_π and change in inflation target between regimes Δ_π . The prior on ρ_π is a beta distribution with a mean of 0.5 and a standard deviation of 0.2. The estimated posterior mode for ρ_π is 0.98, implying a very slow adjustment of the target to its newer, higher value. We use a wide uniform prior for Δ_π that ranges from -8% to 24% in annual terms. At the mode, Δ_π is estimated at roughly 0.01, which in annual terms amounts to a jump in the target of about 4% per year.

For the date breaks, T^{on} and T^{off} , we use uniform priors but restrict T^{off} to lie between 1979Q4 and 1983Q4, the quarters corresponding to the Volcker disinflation. In turn, T^{on} is restricted to take place simply before 1979Q4. Importantly, while the estimation allows for changes in the policy regime, these changes are not imposed. The estimation is free to choose $\Delta_\pi = 0$ and $\sigma_\pi = 0$, if it so desires.

The data strongly favor a specification in which the increase in inflation in the '70s is in large part interpreted as permanent, with π^s smoothly increasing from 2% to roughly 6% at the mode and with negligible mass for $\Delta_\pi < 0$. Most of the remaining variation is explained by temporary shocks to the inflation target. The date breaks are precisely estimated, with the inflationary regime beginning in the late '60s and ending in the early '80s. The estimates of the policy rule parameters are in line with those found in the literature. In the interest of space, the full set of estimates of the structural parameters and date breaks

²² ζ refers to the fraction of sticky firms.

appears in Appendix C.

The main difference across regimes is that when $\mathbb{I}^s = 0$, the inflation target shocks are shut down, and once $\mathbb{I}^s = 1$, shocks to the inflation target can have an impact on endogenous variables. To gauge how the contribution of structural shocks changes with the policy regime to inflation fluctuations, we conduct the following two exercises.

In our first exercise, we simulate the estimated model, setting the policy target shocks equal to their estimated values while setting the value for all other shocks to zero, as shown in Figure 8. In our second exercise, we repeat the exercise but set the policy target shocks to zero and set all other shocks to their estimated values, as shown in Figure 9. We conduct the two exercises for the model with an interest rate rule and also the model with a money rule. We show the corresponding counterfactuals as solid black and dashed blue lines in Figure 8 and Figure 9, respectively. For the ease of comparison, we also show the values for inflation in the United States during the period. Figure 8 makes clear that, independent of the monetary policy instrument we specify, the shocks to the target alone can do a very good job at tracking the evolution of the low-frequency component of inflation in the data. All other shocks only add to high-frequency inflation fluctuations but fail in capturing the high inflation from mid '60s to early '90s, as shown in Figure 9.

We further compare variance decompositions for Regime 1 and Regime 2 to understand the contribution of structural shocks changes with the policy regime to other variables.²³ Table 1(A) and 1(B) present variance decomposition results for interest rate rule and money rule, respectively. Table 1(A) shows that shocks to the inflation target, ε_{π^*} , account for the bulk of fluctuations in inflation and the nominal interest rate in Regime 2. Interestingly, the variance decomposition for real GDP growth, g_t , is essentially the same for the two regimes, with productivity shocks accounting for around three-fourths of its variance across regimes and the target shocks accounting for just 0.5% of its volatility.

²³This decomposition of the unconditional variance is due to the structural shocks alone, capturing what the unconditional variance would be if the regime were to prevail indefinitely. It does not account for the fraction of the variance in the data that results from permanent changes of the inflation target from $\Delta\pi$.




Figure 8: Policy target shocks only




Figure 9: All shocks but policy target shocks

Thus, not accounting for these monetary policy regime changes will wrongly assign these fluctuations to the other shocks and will therefore give rise to biases in the estimates. Table 1(B) further shows that the same pattern emerges in the model with a money rule as well and that policy shocks explain most of the variation in inflation and interest rates

Table 1: Variance decomposition

	Regime 1: $\mathbb{I}^s = 0$					Regime 2: $\mathbb{I}^s = 1$				
	i_t	π_t	g_t	μ_t	x_t	i_t	π_t	g_t	μ_t	x_t
<i>Shocks</i>										
ε_i	2.2	6.2	4.0	3.7	22.6	0.4	0.8	4.0	2.5	21.2
ε_a	94.5	15.6	23.2	19.4	70.8	18.1	1.9	23.1	11.1	66.5
ε_e	3.3	78.2	0.4	1.2	6.6	0.6	9.7	0.4	1.6	6.2
ε_z	0.0	0.0	72.4	20.7	0.0	0.0	0.0	72.0	18.6	0.0
ε_{π^*}	0.0	0.0	0.0	0.0	0.0	80.8	87.6	0.5	16.6	6.0
ε_m	0.0	0.0	0.0	55.0	0.0	0.0	0.0	0.0	49.5	0.0

(A) Interest rate rule

	Regime 1: $\mathbb{I}^s = 0$					Regime 2: $\mathbb{I}^s = 1$				
	i_t	π_t	g_t	μ_t	x_t	i_t	π_t	g_t	μ_t	x_t
<i>Shocks</i>										
ε_μ	0.2	3.0	1.2	7.2	10.6	0.2	0.2	1.9	4.7	14.5
ε_a	97.3	16.8	20.3	55.6	54.2	33.9	0.2	11.1	9.0	10.4
ε_e	0.8	72.2	0.2	3.5	4.5	0.6	6.0	0.3	1.8	5.9
ε_z	0.5	2.1	75.8	10.0	9.3	0.4	0.2	80.3	12.3	14.5
ε_{μ^*}	0.0	0.0	0.0	0.0	0.0	64.3	92.9	2.0	43.8	23.2
ε_m	1.1	5.9	2.5	23.7	21.3	0.7	0.5	4.5	28.4	31.4

(B) Money rule

in Regime 2.

These results are in contrast with findings of [Galí and Gertler \(1999\)](#) and [Smets and Wouters \(2007\)](#), who assume a fixed inflation target and find that wage and price markup shocks play dominant roles in explaining inflation fluctuations. Without a regime change of the inflation target, they find that monetary policy shocks explain a small part of federal funds rate's fluctuations and account for only a little part of inflation fluctuations.

4.3 Simulation analysis

We now use the estimated model to run several simulations that show the importance of the estimated regime changes in the inflation target.²⁴ We use the simulations to evaluate the performance of the filter we applied to the data by treating the simulated data with the same filter.

First, we simulate the model 40 times, setting the estimated regime change and shocks to the target at the mode of the estimated posterior distribution and drawing all other shocks randomly from their posteriors. In the left panel of [Figure 10\(a\)](#), we plot the 40 simulations, together with the inflation data. The right panel of the same figure shows the filtered version of the series in the left panel. In both cases, the data is represented with a wider black line. We next re-estimate the model but set all the shocks to the target to be zero, which is the standard procedure in the New Keynesian literature. We then simulate again the model 40 times, drawing all shocks from their posterior distributions. The left panel of [Figure 10\(b\)](#) shows the 40 simulations plus the data for the period, and the right panel shows the filtered version of the series in the left panel. As the figures make clear, the model without the shocks to the target cannot reproduce the low-frequency movements detected in the data.

The preceding exercises all point in the same direction: the low-frequency movements in the data that we discussed in Section 3 are well captured by the shocks to the target, while all the other shocks typically used in the literature have a very hard time accounting for them, even if we do not allow for the shocks to the target in the estimation.

Sensitivity to price-setting frictions In order to evaluate the role of the price frictions, we simulated the model by setting all shocks equal to their estimated values, but varied the value for the Calvo parameter ζ . We used the values 0.9, 0.6, and 0.1. Recall that we calibrated the Calvo parameter to be 0.6 in the estimation, so that case corresponds to the

²⁴In the text we focus the analysis on the evolution of inflation. In Appendix C, we show the results for the nominal interest rates and for money growth.

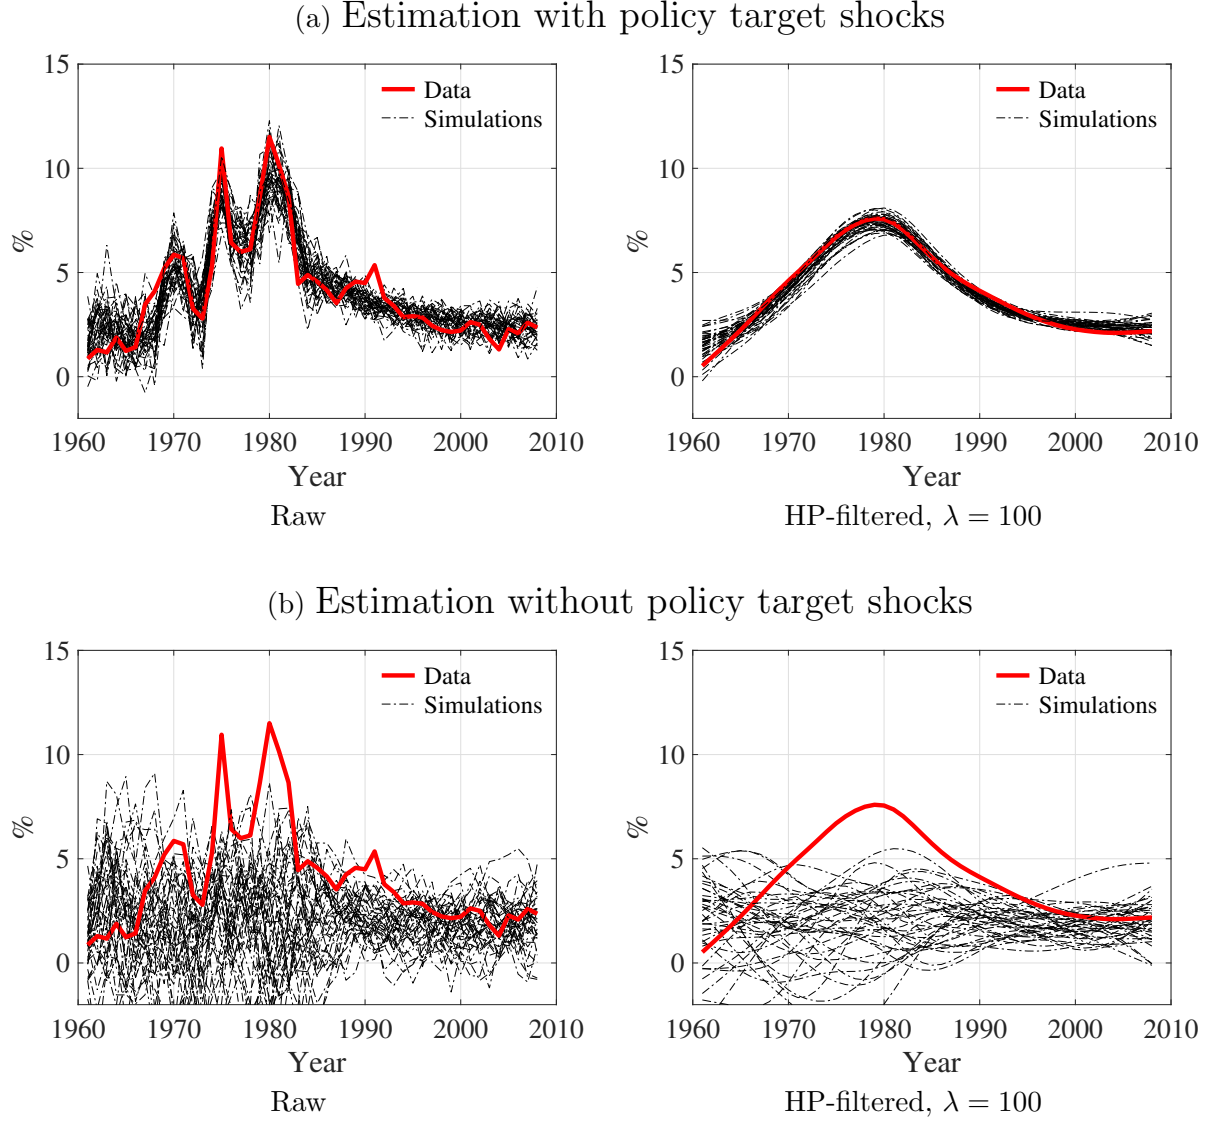


Figure 10: Model fitness of inflation rates – interest rate rule

true data. Then, we filtered the simulated data. We present the results for inflation and the interest rate in [Figure 11\(a\)](#).

As the figure shows, while the specific value for that parameter does change both the maximum inflation attained and the date at which it occurs, the differences are relatively small, even though the variations on the Calvo parameter are very large. In [Figure 11\(b\)](#), we report the results of the same exercise, with the values of the estimated shocks to the target equal to zero. As in [Figure 11\(a\)](#), the differences across regimes with different price

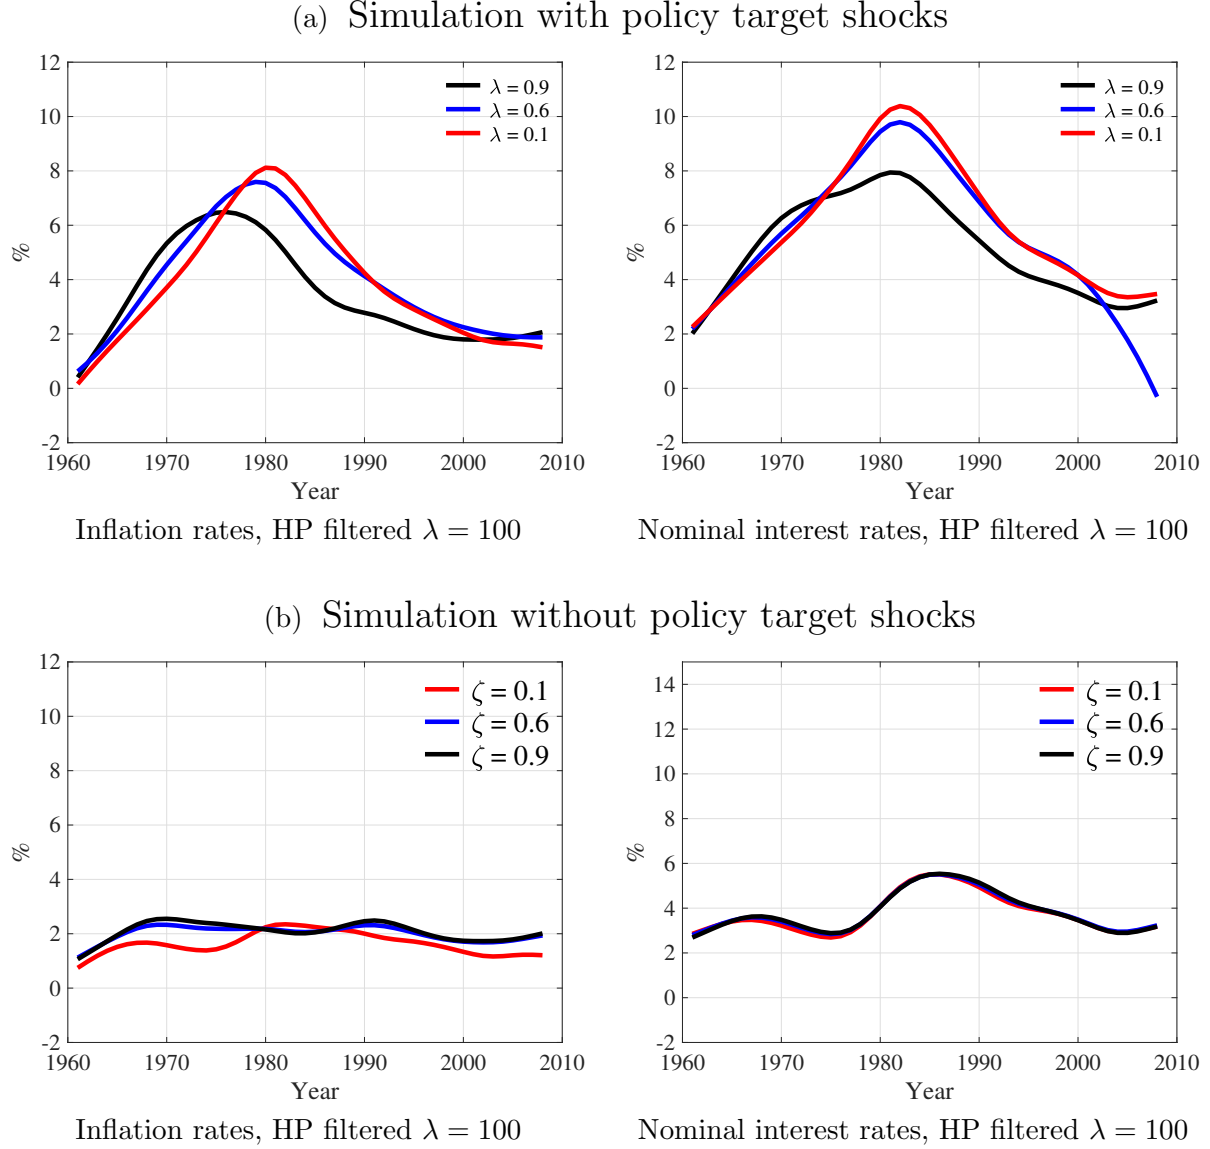


Figure 11: Sensitivity to price stickiness – interest rate rule

frictions is negligible, and in no case is it possible to reproduce the rise and subsequent fall in inflation that characterized the data in the United States between 1960 and 1990. These results reinforce the notion that the strength of the monetary transmission mechanism is not crucial for understanding the main trend observed in the inflation rates of the OECD countries presented in [Figure 1](#).


Quantity theory correlations In [Figure 3](#) through [Figure 6](#), we report the correlation between the variables involved in the two illustrations, which in the absence of real shocks ought to all be equal to 1. A common feature in all cases is that the correlation increases substantially when using filtered data.

In order to evaluate the role of the shocks to the target in those correlations, we simulate the estimated model 10,000 times, drawing all shocks from their posterior distributions, and treat the data the same way that we treat the true data in [Section 3](#). As in that section, we compute the correlation between the inflation rate and the two theoretically computed inflation rates — our two illustrations. We do so both for the simulated data and for the filtered version. We then compute the distribution of the correlation in all cases. The distributions obtained through this process are depicted in the left panels in [Figure 12](#). As a comparison, we repeat the exercise but set the shocks to the target equal to zero and draw all other shocks from their posterior distributions. We depict the distribution in right panels. The distributions of the correlation coefficients before and after the filtering in each case are depicted in blue and orange, respectively. Figures in panel (a) show the first illustration, while figures in panel (b) show the second. We also report means for all distributions.

Even though the two quantity theory predictions hold by construction in the model, the lack of shocks to the target implies that the correlations are lower and that filtering the data actually worsens the fit, in line with the two opening quotes of the paper.

This may be the reason why the country with the worst fit in our empirical analysis of [Section 3](#) was Germany, for which there is very little evidence of an important regime change in the period analyzed.

(a) Illustration 1



(b) Illustration 2




Figure 12: Correlations of series in simulated data – interest rate rule

5 Conclusions: Policy implications

Good day-to-day central banking is a complicated task: it amounts to monitoring and assessing massive amounts of data, simulating alternative scenarios, studying the robustness of policies in each scenario, and deciding the right judgment in each policy decision. These decisions affect the actions of many different members of society, none of whom know exactly how the economy functions. It is very tempting, given the complicated nature of economic relationships, to disregard the lesson of very simple, almost naive theoretical constructions like the quantity theory of money. We build a case for not falling into that temptation. The immediate effect of a monetary policy change depends on details of the environment, and relatively minor changes can sometimes substantially affect the theoretical conclusions. But to understand medium-term inflation, we argue that the quantity theory, though a simple and utterly unrealistic abstraction, suffices.

The combined analysis of simple filtering techniques and the estimation of a specific structural model suggest that a reasonable definition of medium run is between three and four years. This being so, are there direct policy implications that come out of our analysis? We believe so, but clearly it depends on the question at hand. We illustrate this by addressing several very topical policy issues.

We start with one debate to which we have nothing to contribute. By the second half of 2016, the yearly inflation rate in the USA, as measured by the core personal consumption expenditures index, was gradually going up, to the point of getting very close to its target of 2%. However, at the beginning of 2017, the behavior reverted and inflation started falling below 1.3% by August of that year, raising concerns regarding the optimal future path for the policy rate. The analysis of this paper is helpless in trying to understand and amend that two-quarter event. No useful policy advice derives from our analysis.

A longer-run issue has also been a source of ample debate. The Federal Reserve announced an official target of 2% in January 2012, a time in which inflation, again measured with the core personal consumption index, had finally exceeded 2% for the first

time since the eruption of the 2008 financial crisis. Inflation then remained above the 2% target till April of the same year, fell to 1.5% by the end of 2012, and was below the target till the early months of 2020. Our analysis implies that had the nominal interest rate been 50 basis points higher than it was during those years, inflation would have been closer to its target, on average, during the last seven years.

A more dramatic case is that of Japan. For over two decades, the Bank of Japan has been concerned about the low inflation rates. This can be seen in panel (a) of Figure 13, which plots the low-frequency component of inflation in Japan — the solid red line — together with the equivalent measure for the other seven countries in Group 1. Japan appears as the clear outlier, with substantially lower inflation all the way till the end of the sample, at which point its inflation rate seems to converge with the group. Does the policy followed by the Bank of Japan explain this fact? We believe so. Panel (b) of Figure 13 plots the low-frequency components of the policy rates. Japan is again the clear outlier, with interest rates systematically lower than the rest, except at the end of the sample. In the natural counterfactual in which Japan had maintained permanently higher interest rates — say, at the average value for the other countries — the inflation rate in Japan would have also been higher — say, at the average value of the other countries.




Figure 13: Low-frequency movements of Group 1 countries since 1990

The figure also hints at the reason why, over a decade ago, inflation in Japan started to increase somewhat. Notice that the negative trend in nominal interest rates of all the other countries in panel (b) is not followed by the inflation rates in those same countries since the year 2000 or so in panel (a). According to the model, this is possible only if the real interest rate also falls during the period by a magnitude similar to that of the fall in the nominal interest rates. Under this interpretation, inflation went up marginally in the second decade of this century in Japan, owing solely to the lower real rates that have been observed globally, since there is no movement in the low-frequency component of the policy rate in Japan. This implies that if global real rates start returning to positive values once the pandemic is over, even lower inflation readings in Japan become more likely — unless the nominal interest rate in Japan goes up in the medium term.²⁵

This analysis also sheds light on possible future scenarios in the United States, following the policy decisions made after the onset of the 2020 pandemic.²⁶ We make reference to two different decisions. The first was the one to set the interest rate at its effective zero lower bound. Given the current negative real rates exhibited by indexed bonds in the United States, that situation is compatible with positive inflation rates. However, once the economic effect of the pandemic recedes, one could reasonably expect the real rate to return to positive values. Once that happens, and to the extent that the policy rate remains at zero, our theory and our data imply a decreasing trend for inflation, possibly to negative territory. Is it reasonable to expect the policy rate to remain at zero for a long period? This brings us to the second important policy changes made in 2020: the new monetary policy framework announced by the Fed in August 2020 and the FOMC statement that followed in September. These decisions make a prolonged period of policy rates very close to zero more likely. In a nutshell, they imply that the Fed will refrain from increasing the nominal interest rate until inflation is on track to moderately exceed 2% for some time.

Sure enough, high-frequency movements of the kind we ignored here may generate

²⁵See [Uribe \(2020\)](#) for a complementary analysis that points towards similar conclusions.

²⁶The analysis also applies to several other central banks of developed economies.

inflation rates above 2% for “some time,” as has actually happened in the first quarters of 2021. If these high inflation rates continue for a few more quarters, the policy rate may increase. Our paper has nothing to contribute to that debate. However, keeping the interest rate at zero for long periods may become a “target” shock with deflationary pressures like the ones described above. If this happens and inflation tends to negative values, the new framework implies that the nominal interest rate will remain at zero, forcing the trend inflation to remain in negative territory if the real rates become positive again. In this case, low inflation and low interest rates reinforce each other. This low-frequency component of policy risks bringing about a convergence of the United States inflation rate and the Japanese experience since the mid-'90s.

References

- Alvarez, F. and Lippi, F. (2009). Financial Innovation and the Transactions Demand for Cash. *Econometrica*, 77(2):363–402.
- Ascari, G. and Sbordone, A. M. (2014). The Macroeconomics of Trend Inflation. *Journal of Economic Literature*, 52(3):679–739.
- Baumol, W. J. (1952). The Transactions Demand for Cash: An Inventory Theoretic Approach. *Quarterly Journal of Economics*, 66(4):545–556.
- Benati, L. (2009). Long Run Evidence on Money Growth and Inflation. ECB working paper.
- Benati, L., Lucas, Jr., R. E., Nicolini, J. P., and Weber, W. (2019). Online Appendix for: International Evidence on Long-Run Money Demand. Staff report 588, Federal Reserve Bank of Minneapolis.
- Benati, L., Lucas Jr, R. E., Nicolini, J. P., and Weber, W. (2021). International Evidence on Long-Run Money Demand. *Journal of Monetary Economics*, 117:43–63.
- Cogley, T. and Sbordone, A. M. (2008). Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve. *American Economic Review*, 98(5):2101–2126.
- Fitzgerald, T. J., Jones, C., Kulish, M., and Nicolini, J. P. (2020). Is There a Stable Relationship between Unemployment and Future Inflation? Staff report 614, Federal Reserve Bank of Minneapolis.
- Galí, J. and Gertler, M. (1999). Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics*, 44(2):195–222.
- Hazell, J., Herreno, J., Nakamura, E., and Steinsson, J. (2020). The Slope of the Phillips Curve: Evidence from US States. NBER working paper no. W28005.

- Ireland, P. (2004). Technology Shocks in the New Keynesian Model. *Review of Economics and Statistics*, 86(4):923–936.
- Ireland, P. (2007). Changes in the Federal Reserve’s Inflation Target: Causes and Consequences. *Journal of Money, Credit and Banking*, 39(8):1851–1882.
- Kulish, M. and Pagan, A. (2017). Estimation and Solution of Models with Expectations and Structural Changes. *Journal of Applied Econometrics*, 32(2):255–274.
- Kulish, M. and Pagan, A. (2021). Turning Point and Oscillatory Cycles: Concepts, Measurement, and Use. *Journal of Economic Surveys*, 35(4):977–1006.
- Kulish, M. and Rees, D. M. (2017). Unprecedented Changes in the Terms of Trade. *Journal of International Economics*, 108:351–367.
- Latané, H. A. (1960). Income Velocity and Interest Rates: A Pragmatic Approach. *Review of Economics and Statistics*, 42(4):445–449.
- Lucas, R. E. (1980). Two Illustrations of the Quantity Theory of Money. *American Economic Review*, 70(5):1005–1014.
- Lucas, R. E. (1988). Money Demand in the United States: A Quantitative Review. *Carnegie-Rochester Conference Series on Public Policy*, 29:137–167.
- Lucas, R. E. and Nicolini, J. P. (2015). On the Stability of Money Demand. *Journal of Monetary Economics*, 73:48–65.
- McCandless, G. T. and Weber, W. E. (1995). Some Monetary Facts. *Federal Reserve Bank of Minneapolis Quarterly Review*, 19(3):2–11.
- Sargent, T. J. (1982). The Ends of Four Big Inflations. In Hall, R. E., editor, *Inflation: Causes and Effects*, pages 41–98. University of Chicago Press.

- Sargent, T. J. and Surico, P. (2011). Two Illustrations of the Quantity Theory of Money: Breakdowns and Revivals. *American Economic Review*, 101(1):109–128.
- Selden, R. (1956). Monetary Velocity in the United States. In Friedman, M., editor, *Studies in the Quantity Theory of Money*, pages 405–454. University of Chicago Press.
- Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3):586–606.
- Taylor, J. B. (1993). Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214.
- Teles, P., Uhlig, H., and Valle e Azevedo, J. (2016). Is Quantity Theory Still Alive? *Economic Journal*, 126(591):442–464.
- Tobin, J. (1956). The Interest-Elasticity of Transactions Demand for Cash. *Review of Economics and Statistics*, 38(3):241–247.
- Uribe, M. (2020). The Neo-Fisher Effect: Econometric Evidence from Empirical and Optimizing Models. Working paper, Columbia University.
- Vogel, R. C. (1974). The Dynamics of Inflation in Latin America, 1950-1969. *American Economic Review*, 64(1):102–114.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.