b
UNIVERSITÄT
BERN

Faculty of Business, Economics
and Social Sciences

Department of Economics

International Evidence on Long-Run Money Demand

Luca Benati, Robert E. Lucas Jr.,
Juan Pablo Nicolini, Warren Weber

20-21

December, 2020

DISCUSSION PAPERS

Schanzeneckstrasse 1
CH-3012 Bern, Switzerland
http://www.vwi.unibe.ch

International Evidence on Long-Run Money Demand

Luca Benati[a], Robert E. Lucas Jr.[b], Juan Pablo Nicolini[c], Warren Weber[d]

[a] *University of Bern*
[b] *University of Chicago*
[c] *Federal Reserve Bank of Minneapolis and Universidad Di Tella*
[d] *University of South Carolina, Bank of Canada, and Federal Reserve Bank of Atlanta*

Journal of Monetary Economics, forthcoming

Abstract

We explore the long-run demand for M1 based on a dataset comprising 38 countries and relatively long sample periods, extending in some cases to over a century. The evidence supports the existence of a stable long-run relationship between the ratio of M1 to GDP and a short-term interest rate for a large majority of the countries. The log-log specification provides a good characterization of the data, with the exception of periods featuring very low interest rates. An extension of the theory that imposes limits on the amount households can borrow results in a truncated log-log specification, which is in line with what we observe in the data. We estimate the interest rate elasticity to be between 0.3 and 0.6.

Keywords: Long-run money demand, Cointegration

*We thank Francisco Buera, Giuseppe Cavaliere, V.V. Chari, Jesus Fernandez-Villaverde, Patrick Kehoe, Tim Kehoe, Helmut Lüetkepohl, Albert Marcet, Rodolfo Manuelli, Edward Nelson, B. Ravikumar, Pedro Teles, and Nancy Stokey. Special thanks to Joao Ayres, Han Gao, Nicole Gorton, Ruth Judson, Eva Werner, and Junji Yano for invaluable help with the data. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Banks of Minneapolis or Atlanta or the Federal Reserve System.

1. Introduction

This paper describes and analyzes a new dataset containing annual measurements of money supplies, both real and nominal output (GDP) and thus price levels, and short-term nominal interest rates for 38 countries, for periods that go from three decades to over a century. The framework we use for organizing these data is a money demand function that relates the money that the public and private sectors of the economy choose to hold to the rate of production of goods and the short-term interest rate,

$$M_t = P_t y_t \phi(r_t), \tag{1}$$

where M_t is the monetary aggregate, P_t is the price level, y_t is total real production, r_t is a short-term nominal interest rate, and ϕ is a decreasing function of r_t.

The formula (1) contains some strong implications. One is that, if r_t is stationary, M_t and $P_t y_t$ should grow at a common rate in the long run, for any continuous function ϕ. If, on the other hand, r_t has a unit root—possibly because inflation is driven in part by permanent shocks—then M_t and $P_t y_t$ should grow at a common rate in the long run, once we control for the impact of permanent shocks to r_t. Another implication is that it should be possible to use both cross-country and within-country time series to trace out the function ϕ. This is the agenda carried out in this paper.

In recent years, many economists and central bankers have come to doubt the usefulness of measures of the money supply (such as M1) in the conduct of monetary policy. What was thought to be a central pillar of the monetary policies of the newly established European Central Bank in 1999 has come to be seen as too unreliable to be of any use. These concerns were not without an empirical basis.

First, since shocks to real money demand have historically been volatile, money-based rules may induce substantial volatility in policy outcomes. In contrast, as interest rate rules are immune to those shocks, they isolate the economy from them. The remarkable success that all central banks in the developed world (and many in the developing world) have had in maintaining inflation rates very close to target over the last decades, using interest rate rules, provides strong empirical support for this notion. More recently, however, with real interest rates at historically low values, the threat of the zero lower bound (ZLB) became a serious constraint on monetary policy. With the policy rate constrained by the ZLB, it is conceivable that a money-based rule may perform better, as recently shown by Belongia and Ireland (2019b). A maintained assumption in Belongia and Ireland (2019b) is that a stable relationship such as the one described in (1) exists. We document this is indeed the case for a large set of countries.

Second, shocks to real money demand exhibit very high persistence. This means that deviations of the data from the theoretical expression in (1) are long-lived. A large literature has developed to understand this fact, assuming that agents make transactions in physically segmented markets.[1] A recent example is provided in Alvarez and Lippi (2014). Heterogeneous access to markets implies that monetary policy affects agent's decisions with lags, so that the aggregate real money demand responds with lags to disturbances. Alvarez and Lippi (2014) show how this modification to the standard model can successfully account for several of the short-run facts. On the other hand, their theory cannot quantitatively reproduce the long-run movements exhibited by the data. Our empirical exploration will abstract from these short run fluctuations. In order to to so, we use the methods of cointegration and apply them uniformly to a wide variety of countries. The virtue of these methods for our purposes is that they make precise what we mean by long run relations. These methods also characterize the short run behavior, by estimating the moments of the cointegration errors - which as expected in light of the previous discussion, are quite persistent. This quantification is useful to discipline future attempts to integrate the long run component we focus in this paper with the short run evidence that the work of Alvarez and Lippi (2014) can rationalize.

Finally, the argument has been made that as the formula (1), which was shown to perform very well empirically by Meltzer (1963) and Lucas (1988), has broken down in the last decades, so it cannot be a reliable guide for policy. This has indeed been the case in the United States for M1, which

[1]See Grossman and Weiss (1983) and Rotemberg (1984) for early contributions.

adds cash to demand deposits. But several papers (e.g., Carlson, Hoffman, Keen and Rasche (2000), Teles and Zhou (2005), Serletis and Gogas (2014), Judson, Schlusche and Wong (2014), Barnett (2016), Belongia and Ireland (2019a), Anderson, Bordo and Duca (2017)) have argued that it has not been the case for other measures, such as MZM (or Money Zero Maturity), or by aggregating the components of money using Divisia indexes. Others (e.g., Lucas and Nicolini (2015)) have shown that accounting for regulatory changes that allowed for newly created deposits that are very close substitutes to checking accounts, and that occurred precisely at the time of the breakdown can restore a stable money demand function.[2] This literature has forcefully argued that the apparent breakdown of real money demand in the United States is just the result of regulatory changes that made the measure of M1 reported by the Federal Reserve an unreliable measure of means of transactions.

With 38 different datasets, some covering more than a century, we can expect surprises and there will indeed be a few. But first we want to explain why in 37 countries we apply M_t as M1 and only for the United States we also apply "New M1," a short hand for M1 plus money market demand accounts, or more briefly NewM1 = M1 + MMDA. The addition of MMDA appears only in the U.S. and only because the 1933 Glass-Steagall Act imposed the Regulation Q that prohibited interest payments on checking accounts. For

[2]See also Ireland (2008) for a related discussion.

the early years of Glass-Steagall checking accounts remained, with free checking and other services, fairly close to what competitive banking would have done without Regulation Q. But by the late 1970s and early 1980s returns to bankers were on the order of 8 percent or more on deposits. Thus, in 1982, banks were allowed to issue the newly created created MMDA. Large cash holders substituted away from checking accounts and into MMDAs, so M1 continued to fall. Here we chose NewM1 = M1 + MMDA, the monetary aggregate proposed by Lucas and Nicolini (2015), to capture this.[3] For the 37 others without Regulation Q no such NewM1 was needed.

Central banks in the developed world are revising their policy frameworks to adapt them to a world with very low real interest rates. None of them are seriously considering money-based rules, and that is understandable: money demand models are not yet be ready for prime time. We need to better understand the behavior of monetary aggregates in the countries that lack solid evidence. We also need to integrate into a common theoretical framework the low-frequency components analyzed in this paper with the high frequency behavior that Alvarez and Lippi (2014) successfully analyzed. Such an integrated framework, assuming that it passes the quantitative tests, would provide support for policy evaluation exercises such as those in Belongia and Ireland (2019b).[4] The evidence described in this paper provides sizable em-

[3]Results using other aggregates as MZM or adding Mutal Money Market Funds to NewM1 are very similar to the ones obtained using NewM1.

[4]Belongia and Ireland (2019b) model the lagged response of real money demand as adjustment costs in the utility function.

pirical support for the notion of a stable real money demand. This is a required first input into the analysis of money rules as a policy option when real interest rates become very low.

We take pains to ensure that terms such as "short-term interest rate" and "money" are measures of the same thing (or almost!) in different countries and over time within countries. The set of countries is highly heterogeneous in terms of size, income per capita, and world region. More importantly, the countries' respective monetary histories also differ substantially : our sample includes countries that experienced hyperinflation as well as countries in which inflation has been almost always within a single digit. The periods covered include very different growth experiences, different monetary arrangements, and different degrees of integration within the world markets. We will explicitly ignore all those differences and will look at this diverse set of countries through the lens of an extremely simple model. The high degree of variation in nominal interest rates across countries and over time within each country is what we will exploit in building our case that the basic features of the demand function for money are in general quite solid for a large set of countries.

A particular expression for the function $\phi(r_t)$ is the well-known squared root formula that Baumol (1952) and Tobin (1956) derived over half a century ago. In this paper, we use a theory that generalizes this Baumol-Tobin expression along a couple of dimensions. The first generalization allows for a technology to transform bonds into money that encompasses the linear technology assumed by Baumol and Tobin, but which also allows for nonlinear relationships. The second generalization is the consideration of borrowing constraints, which affects the behavior of money demand at very low interest rates.

We address the elements of agents' decision problems in detail in Section 2, where we derive an equation like (1) that generalizes the familiar Baumol-Tobin specification. In Section 3, we plot the implied predictions of a particular case of the model against the data and let the graphics speak for themselves. The methods of cointegration are described in Sections 4 and 5, where we discuss the methodology that we use throughout the paper and discuss the results from cointegration analysis. In Section 6, we discuss some extensions, and Section 7 concludes.

2. A Model of Money Demand

We study a labor-only, representative agent economy in which making transactions is costly. Preferences are given by

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

where $\beta < 1$, x_t is consumption at date t, and the function U is differentiable, increasing, and concave. The agent is endowed in each period with a unit of time, with l_t allocated to goods production and $1 - l_t$ used to carry out transactions. The goods production technology is given by $y_t = x_t = z_t l_t$, where z_t is an exogenous stochastic process.

We assume that households choose the number n of "trips to the bank" in the manner of the classic Baumol-Tobin model. At the beginning of a period, a household begins with some nominal wealth that can be allocated to money M_t or to risk-free government bonds B_t. During the first of the n subperiods, one member of the household uses money to buy consumption goods. During this same initial subperiod, another member of the household produces and sells goods in exchange for money. At the end of the subperiod, producers transfer to the bank the proceeds from their transactions. Thus, the situation at the beginning of the second subperiod exactly replicates the situation at the beginning of the first. This process is repeated n times during the period. The choice of this variable n will be the only economically relevant decision made by households. Purchases over a period are then subject to a cash-in-advance constraint, $P_t x_t \leq M_t n_t$.

Notice that n is the velocity of money, and its inverse in equilibrium is the money-to-output ratio, or the demand for real money, which is the concept that we care about. Baumol-Tobin assumed that the cost of carrying out these transactions increases linearly in the number of trips. We consider a more general specification in which the total cost of making transactions, measured in units of time, is given by

$$\theta(n_t, \nu_t) = \gamma n_t^{\sigma} \nu_t, \tag{2}$$

where γ and σ are positive constants and ν_t is an exogenous stochastic process. The natural interpretation of the stochastic shock ν_t is aggregate disturbances in intermediation technologies. This random component is important for motivating the econometric analysis at the core of the paper. The expression in (2) becomes the Baumol-Tobin linear case when we set the curvature parameter σ equal to 1.

Equilibrium in the labor and goods markets implies

$$x_t = z_t l_t = z_t(1 - \gamma n_t^\sigma \nu_t),$$

so the equilibrium real wage is equal to z_t.

At the beginning of each period, the agent starts with wealth in real terms w_t, which can be allocated to money m_t or interest-bearing bonds b_t, both also measured in real terms. We can then write this constraint as

$$m_t + b_t \leq w_t. \tag{3}$$

In addition, we impose a productivity-adjusted borrowing constraint for the agent, in the sense that its bond holdings, b_t, cannot be too negative. Specifically, we impose

$$b_t \geq z_t b^* \tag{4}$$

for some arbitrary value of b^*.[5] Below we discuss how the equilibrium money

[5] Our dataset contains long samples of over a century for a few countries. All of the countries had substantial increases in productivity over the length of the sample. We find

demand relationship is affected by this constraint.

The agent takes the nominal return on short-term bonds, r_t, as exogenously given. We do not need to take a stand on how monetary policy is executed. Our framework allows for r_t to be a process determined by monetary policy, in which case the behavior of the growth rate of the money supply is restricted by other equilibrium conditions. But it also allows for policy to be described as money rules, in which case the nominal interest rate will be given by a Fisher equation.[6]

So far, we have been silent with respect to what our measure of money accounts for. For the theoretical analysis, one can allow for money to pay interest.[7] In what follows, we consider the case in which the interest on money is zero. We discuss this choice below in choosing the empirical counterpart of the monetary aggregate.

The agent's wealth next period, contingent on the actions taken in the current period and the realization of the exogenous shock, is given by

$$w_t^{t+1} \leq \frac{m_t + b_t(1+r_t) + \left[1 - \gamma n_t^{\sigma} \nu_t\right] z_t - x_t}{\pi_t^{t+1}} + \tau_t^{t+1}, \tag{5}$$

where π_t^{t+1} denotes the gross inflation rate between period t and period $t+1$ in that particular state, and τ^{t+1} is the real value of the monetary transfer the

---

it natural to assume that the borrowing constraint depends on the level of technology.

[6] Details are provided in Online Appendix B.

[7] That more general case and its details are discussed in the working paper version of this paper (Benati et al. (2017)).

government makes to the representative agent. Finally, the cash-in-advance constraint can be written in real terms as

$$x_t \leq m_t n_t. \tag{6}$$

We now consider the decision problem of a single atomistic agent who takes as given the price level, the inflation rates π_t^{t+1}, the interest rate r_t, the real wage z_t, and the shock ν_t. Given the initial wealth w_t, the agent chooses consumption x_t, the number of trips n_t, and the assets m_t, b. These choices, restricted by equations (3), (4), (5), and (6), determine the wealth w_t^{t+1} carried into the next period.

In Online Appendix B, we show that as long as the borrowing constraint (4) does not bind, the equilibrium number of portfolio adjustments n depends on the interest rate r_t according to

$$r_t = \sigma\gamma n_t^{\sigma+1} \frac{\nu_t}{1 - \gamma n_t^{\sigma} \nu_t}. \tag{7}$$

The solution involves an extended quadratic formula for the equilibrium value of n. Using the cash-in-advance constraint (6), the following relationship must hold in equilibrium:

$$\frac{m_t}{x_t} = \frac{1}{n_t(r_t, \nu_t)}, \tag{8}$$

which is a particular case of (1).

This solution has several empirical implications. First, notice that the solution for the money-to-output ratio does not depend on the technology level, z_t. In the model, income growth must be associated with a positive trend in technology. This implies an income elasticity of the real demand for money that is equal to one, which is the specification we will study.[8] Second, the stochastic properties of the money-to-output ratio, m_t/x_t, are inherited from the stochastic properties of r_t and ν_t. This has testable implications as long as ν_t is stationary, as we will assume throughout the paper. Specifically, if r_t is stationary, m_t/x_t should be too, whereas if r_t has a unit root, m_t/x_t should have a unit root as well. Our cointegration analysis below will address these issues.

2.1. Characterization of the solution

We first characterize the solution for the case in which the borrowing constraint (4) is not binding. Then, we provide a general characterization.

2.1.1. When the borrowing constraint does not bind

A simple inspection of equation (7) reveals that velocity is an increasing and continuous function of the interest rate $r_{t.}$ As (8) makes clear, the ratio of money to income is then negatively related to r_t.

To obtain a simple parametric form that we can take to the data, we discuss one approximation. Note that $\gamma\nu_t n_t^{\sigma}$ represents the welfare cost of

[8] In Online Appendix B, we allow for a more general specification that does not restrict the income elasticity to be one, and where we are able to test this unitary income elasticity implication.

inflation as a ratio of maximum potential output, which is arbitrarily close to zero when the interest rate r_t is small. For moderate interest rates, computations of the welfare cost are negligible. Even for interest rates as high as 20%, estimates of the welfare cost of inflation are barely above 4%, so the denominator in the expression above would range from 1 to 0.96.[9] We then use the approximation $1-\gamma\nu_t n_t^\sigma \simeq 1$ and write the solution as $n_t^{\sigma+1}\sigma\gamma\nu_t \simeq r_t$. Taking logs and using the fact that $m/x = 1/n$, we then obtain

$$\ln\left(m_t/x_t\right) = \frac{1}{(\sigma+1)}\left[\ln\sigma\gamma - \ln r_t + \ln\nu_t\right], \tag{9}$$

which is the log-log function typically used in the literature, with an interest rate elasticity of $1/(1+\sigma)$. The Baumol-Tobin case is the one obtained when $\sigma = 1$, so the elasticity is equal to $1/2$.

Notice that a property of this specification is that real money balances over output, m_t/x_t, go to infinity when the nominal interest rate differential, r_t, goes to zero. How can this be a solution for a representative agent with finite wealth? Inspecting the budget constraint (3) suggests that bond holdings must therefore be approaching negative infinite. In this closed economy model, this means that agents are borrowing unbounded amounts from the government. Imagine, then, that agents could run away with the cash they hold and keep some fraction of it. The borrowing constraints (4) would

[9] The approximation error in a model calibrated to match US data is very small, below 2%, even for interest rates as high as 50% per year.

naturally arise from an optimal contracting problem with enforceability constraints.

We next turn to considering the case of a binding borrowing constraint.

2.1.2. When the borrowing constraint binds

When the borrowing constraint binds, $b = z_t b^*$, the solution is trivial, since there is really no economic problem to be decided by the agent. Note that the budget constraint (3) implies

$$m_t = w_t - z_t b^*,$$

which fully determines the real quantity of money. The values for x_t and n_t are then determined by the equilibrium conditions $x_t = z_t(1 - \theta(n_t))$ and $x_t = n_t m_t$.

2.1.3. A full characterization

To provide a full characterization of the relationship between money balances to output and the interest rate differential, given any value for the state variable w, it is useful to state the following lemma, proved in the Online Appendix.

Lemma 1. *Given values for w and b^*, if constraint (4) binds for an arbitrary interest rate r_0, then it also binds for any $r < r_0$. In addition, if constraint (4) is not binding for an arbitrary interest rate r_1, then neither does it bind for any $r > r_1$.*

The previous lemma implies that for any value of real initial wealth, w, there exists a value for the net interest rate $\widetilde{r}$ such that the money demand equation is given by the solution to (7), for all $r \geq \widetilde{r}$, whereas it is equal to $\frac{m}{x}(\widetilde{r})$ for all $r \leq \widetilde{r}$. Such a money demand is depicted as the dotted blue line, labeled A in Figure 1.

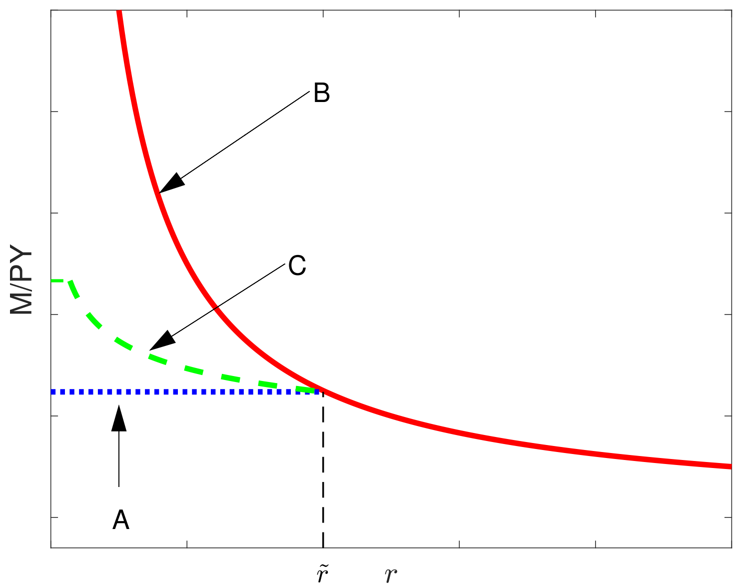

Figure 1 Borrowing constraint and money demand when interest rate is near 0

The location of the kink in the money demand depends on the value for $w - zb^*$. The larger that value, the lower will be the value for the interest rate $\widetilde{r}$. Note also that if we let b^* go to negative infinity, real money balances also go to infinity as the interest rate goes to zero. Such a money demand is depicted as the solid red line, labeled B in Figure 1.

This sharp characterization at near-zero interest rates depends on the representative agent assumption, and it is not robust to sensible generalizations. Thus, we may not want to take the flat portion of the money demand in schedule A of Figure 1 literally when pursuing our empirical analysis. We discuss this issue next.

2.1.4. Connecting the theory to the data

Consider a model like the one above, with a continuum of agents that are alike in all respects except that they differ in their productivity. To be more specific, assume that idiosyncratic productivity for agent j is equal to $\xi^j z_t$, where $\xi^j \in [\xi_l, \xi^h]$, and where the mean of ξ^j is equal to one. It would be natural in this environment to impose agent-specific borrowing constraints, since agents' ability to pay would vary across types.

We explore such an economy in Online Appendix B. There we show that under certain conditions, there will be a threshold interest rate $\widehat{r}$ such that for interest rates higher than $\widehat{r}$, no agent is constrained, so all individual money demand functions are well approximated by the log-log specification. It follows that the aggregate money demand function is also log-log. For interest rates lower than $\widehat{r}$, the aggregate money demand is a combination of two types of agents. For the first type, the constraint binds, so their aggregate demand is insensitive with respect to the interest rate. For the second type, the constraint does not bind, and their elasticity is given by the log-log specification. As the interest rate keeps going down, the fraction of agents for which the elasticity is positive goes down, so the aggregate elasticity also goes down. The aggregate money demand is decreasing for this range, but with an interest rate elasticity that is lower than the log-log specification. Such a money demand is the dashed green line, labeled C in

Figure 1.[10]

In light of this discussion, in our empirical strategy, we will follow two complementary approaches. First, we will ignore the borrowing constraints and use the log-log functional form implied by the theory. We expect this strategy to work well in countries that did not experience low interest rates. Second, we will use a parametric form that is observationally very similar to the log-log specification for interest rates that are not too small, and which differs from that specification at very low levels of interest rates in a way that closely resembles the behavior of the real money demand with the borrowing constraints described above. This parametric form, used by Selden (1956) and Latané (1960), is given by

$$\frac{m_t}{x_t} = \frac{1}{a + br_t}. \tag{10}$$

In Figure 2, we plot two curves relating real money balances to the interest rate. The range of short-term interest rates is the relevant one for the United States in the last century: between 0% and 15%. The blue circles correspond to a log-log specification, with an elasticity equal to 1/2, as implied by the Baumol-Tobin linear technology discussed above. The constant in that equation has been chosen so that the ratio of money to ouput is close to 25% when the interst rate is 6%, which matches the US data reasonably well. The solid

[10] Similar results for the aggregate money demand arise in a model in which agents get the few first portfolio transactions for free. Such a model is developed and estimated in Alvarez and Lippi (2009).

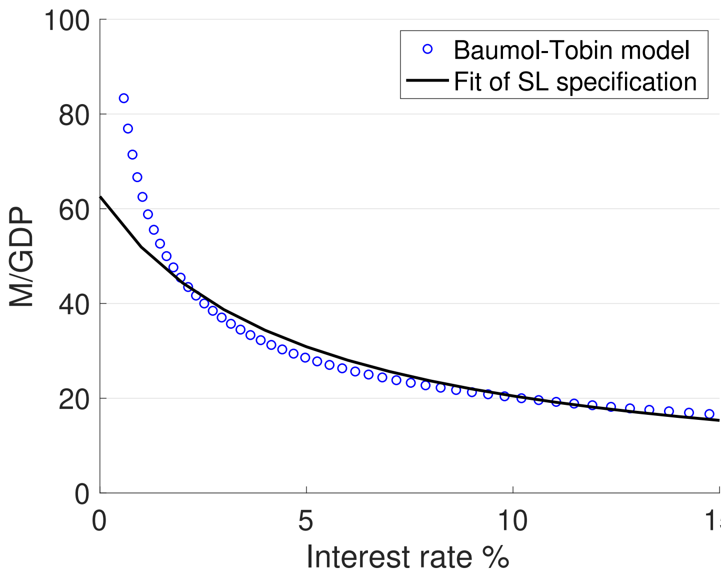

Figure 2 Fitting Baumol-Tobin model with Selden-Latané specification

black line corresponds to the best fit of equation (10) to the blue circles in the figure.[11]

As the figure makes clear, both functional forms behave in a remarkably similar way for interest rates between 15% and 2%. They are so similar, in fact, that it appears that the ability to identify one functional form from the other would require a gigantic amount of data for interest rates within that range. On the other hand, the two functional forms do behave very differently at interest rates between 2% and 0%. To the extent that borrowing constraints are relevant, this formulation ought to work better in countries that did experience low interest rates in several periods. By comparing the empirical performance of the two specifications, we will be able to draw some

[11]The two parameters of the solid black line were calibrated using OLS.

conclusions regarding the quantitative relevance of the borrowing constraints.

2.1.5. The role of regulation and technology

The literature has long recognized that changes in regulation or technology can change the equilibrium relationship between interest rates and real money balances. For instance, Lucas and Nicolini (2015) argue that regulatory changes introduced in the United States in the early 1980s can explain the apparent instability of real money demand in the United States. Alvarez and Lippi (2009) show that advances in banking technology are important in explaining their household level data on cash holdings. The theoretical implications of such changes can be analyzed with this model. In the working paper version of this paper (Benati et al. (2017)), we show that when money pays an interest rate $r_t^m > 0$, the solution for the number of trips to the bank, n_t, is an equation equal to (7) , except that on the right-hand side, r_t must be replaced by $r_t - r_t^m$. Thus, if banks are allowed to compete by paying interest on deposits, the optimal choices of n_t would change even if the interest rate r_t is unchanged. In addition, a change in the level parameter γ changes the optimal value for money balances, again keeping r_t constant.

As mentioned in the Introduction, we will ignore both regulatory and technological changes in this paper. As it turns out, the data analysis shows that by and large, these theoretical considerations have little empirical relevance: the relationship in (8) derived from the model is confirmed by the data for a large set of countries, even though we cover samples that are several decades long. As we document below, we found just a handful of cases for which M1 real money demand appears to have breaks that are suggestive of further analysis. Our general conclusion is that the apparent breakdown in the real money demand relationship observed in the United States, requiring a detailed and country-specific analysis, is an exception rather than the rule.

3. A First Look at the Data

In this section, we present the data and provide a visual comparison with the theory. To begin, we discuss how we map our theoretical construct M_t to the data. This choice is associated with the discussion of its nominal return. We have no data on the interest rate paid by deposits, so we choose to work with M1, which includes cash and only checking accounts. In deciding to set the interest rate on money to zero in the theory, we implicitly assumed that checking accounts pay no interest. This is a questionable assumption, but it is certainly more appropriate for M1 than for broader aggregates, which typically include interest-paying deposits.[12] Accordingly, we identify money in the model with M1.[13]

[12] Deposits did pay interest in the United States after Regulation Q was modified in the early 1980s. Also, some deposits included in M1 did pay interest in very high-inflation countries such as Argentina and Brazil.

[13] We have data on interest rates paid on deposits for the United States. As a robustness check, we use it to compute the opportunity cost of New M1 as the difference between the 3-month Treasury bill rate and money's own rate(details on the computation can be found in Online Appendix A.1.31). Based on Johansen's tests, evidence is qualitatively the same as that based on the Treasury bill rate, with p-values equal to 0.066 for Selden-Latane', and 0.394 based on log-log. Based on Wright's tests, results are the opposite, with the null of cointegration not being rejected based on the log-log, and being rejected for the

Online Appendix A describes the data and the data sources in detail. All of the series are standard, with the single exception of the United States, where we also consider the NewM1 monetary aggregate proposed in Lucas and Nicolini (2015). Specifically, we add to the standard M1 aggregate the money market deposit accounts (MMDAs) that were created in 1982. We call this aggregate New M1.[14] Our simple theory abstracts from investment, so output and consumption are the same. It also abstracts from money demand by firms, which raises the question of wether total output is a better measure than consumption. We chose to use output as our measure of economic activity, as it is customary in the literature.[15]

We first present the raw data in the form of cross plots between the short-term interest rate and the ratio of money to nominal income. The data were not manipulated in any way. Figures 3a to 3c are scatterplots of the short rate and the ratio between nominal M1 and nominal GDP (i.e., the inverse of M1 velocity). We also plot the theoretical curve that corresponds to equation (9), specialized to the case in which the elasticity is equal to 1/2. We allow the per-unit cost, γ^j, to be different across countries. Thus, for each country

Selden-Latane' specification.

[14]The results are the same with an alternative aggregate in which currency has been adjusted along the lines of Judson (2017) to take into account the sizable expansion in the fraction of US currency held by foreigners since the early 1990s. See Benati (2019a) for details.

[15]There is ample evidence that output and consumption are cointegrated, so this choice is likely to be of little relevance. We we have checked this for all countries for which we could find consumption data, and results are qualitatively the same. These results are reported in Online Appendix K

j , we plot the curve

$$\frac{M_t^j}{Y_t^j} = \frac{\gamma^j}{\sqrt{r_t^j}}, \tag{11}$$

where Y_t^j is nominal income at time t in country j and M_t^j is M1, except for the United States, where we use New M1, as mentioned above. We calibrate the single free parameter for each country to be the one that minimizes the mean squared errors between the curve and the data, but imposing the elasticity to be 1/2.

The criteria used to group the countries follow the results of the tests performed below. Details will be discussed then, but a rough approximation is that we start with the countries for which the evidence of comovement between velocity and the interest rate is very strong and show at the end the countries for which the evidence is weak or nonexistent. In our view, it is surprising how well this simple theory, which allows for a single free parameter per country, performs in this first inspection. Our own summary is the following. For the 12 countries in Figure 3a, the evidence is very strong, with two important caveats: for both Brazil and Argentina, a few points clearly lie below the theoretical line. In both cases, those points correspond to the years following the successful stabilization of hyperinflations.[16] The model in the previous section is too simple to be used to address these cases,

[16]Neither Bolivia nor Israel, which also experienced hyperinflations, exhibits the same puzzling behavior. But for those cases, we do not have data for many years before the hyperinflation, so they are not completely comparable.

so we leave that topic for future research. The following 14 countries in Figure 3b also exhibit clear evidence of a negative relationship, though in some cases (such as Bahrain, Barbados, and Thailand), the slope seems to be different from the 1/2 implied by the linear technology. The 6 countries in the top row of Figure 3c also provide good evidence, though in some cases (such as Finland and West Germany), the picture is not as clear as before. Finally, the 6 countries in the bottom row of Figure 3c are the most problematic. Portugal and to a lesser extent Paraguay and Norway, also seem to display evidence of a negative relationship, but it is far from the theoretical curve. The last three cases are, from the viewpoint of the cross plots, blunt failures.

The plots in Figures 3a through 3c depict the data in a way that became traditional in empirical studies of money demand. In addition, the data can be visually compared with the theoretical curve indicated by the theory (11). But depicted in this way, the plots conceal the behavior of the variables over time, and they thus fail to show the persistence exhibited by both series, particularly how the persistent components of the two series have moved together over time. We find that information also very valuable as a visual motivation for the cointegration methods that we use in the rest of the paper. The figures also add the time dimension, which helps to explain some of the apparent failures discussed above.

Thus, in Figures 4a to 4c, we present the time series for both the short

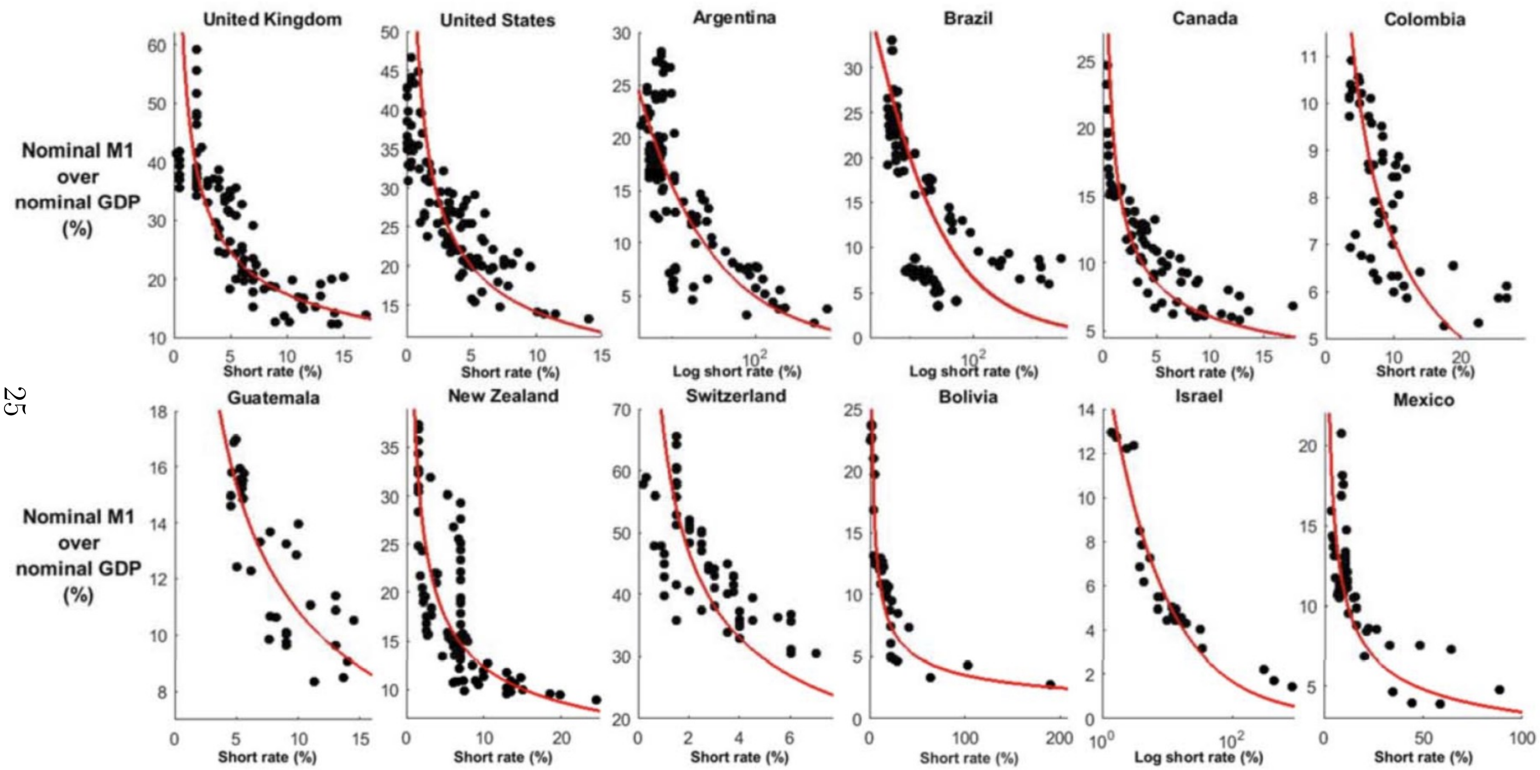

Figure 3a The raw data: scatterplots of the short rate and M1 as a fraction of nominal GDP

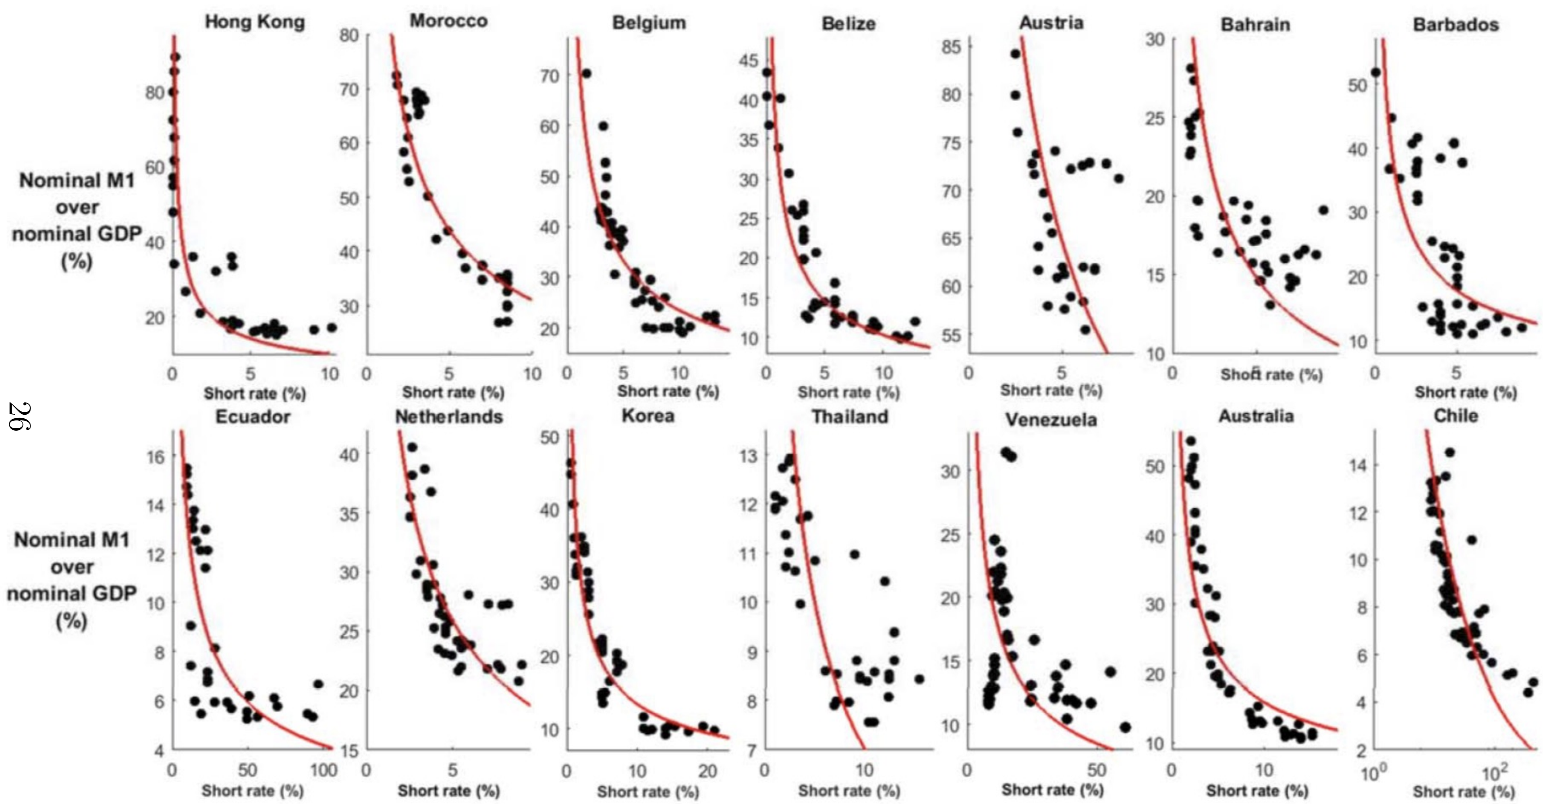

Figure 3b The raw data: scatterplots of the short rate and M1 as a fraction of nominal GDP

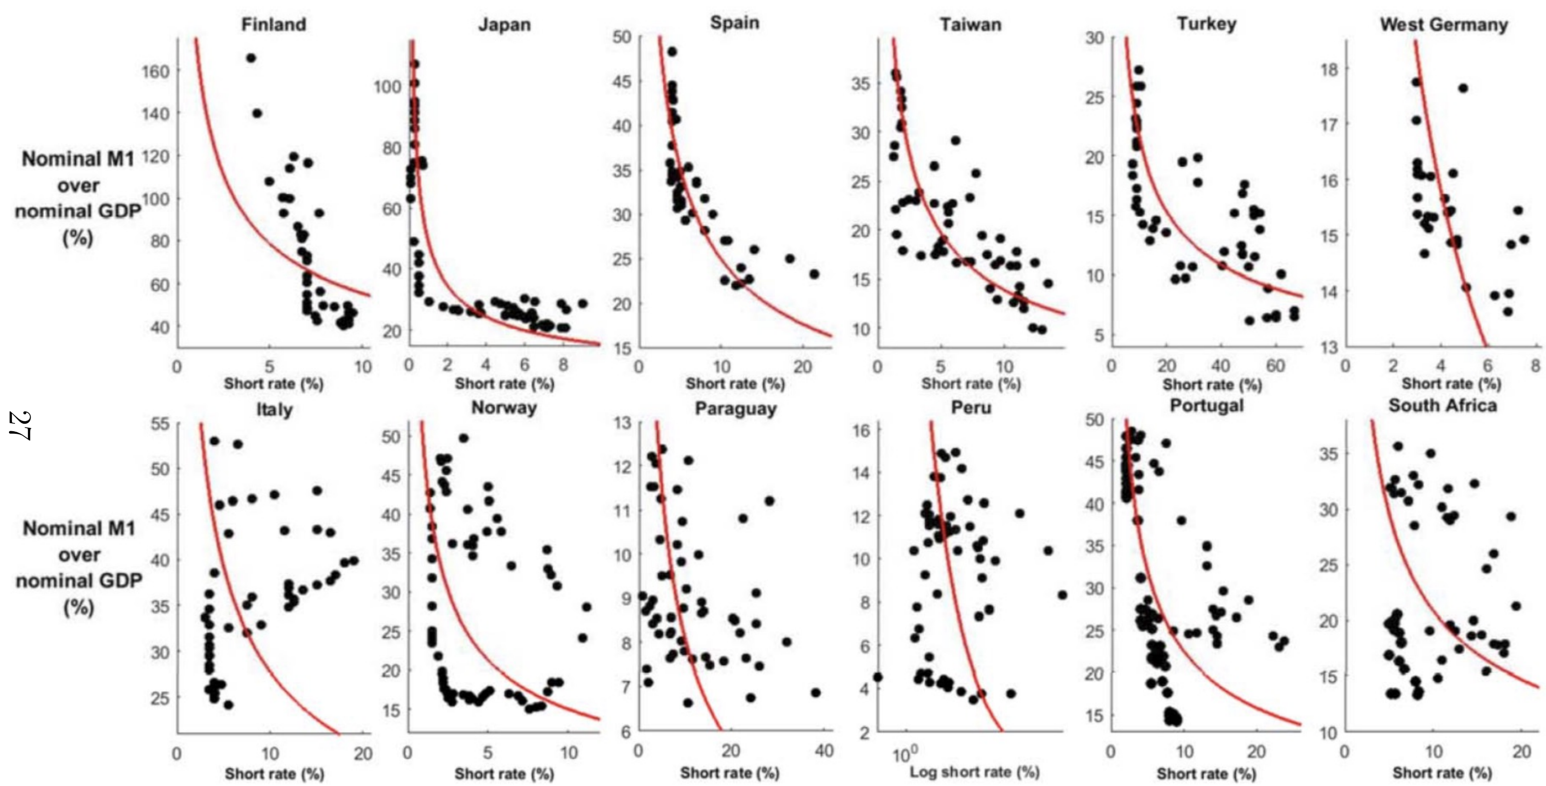

Figure 3c The raw data: scatterplots of the short rate and M1 as a fraction of nominal GDP

term interest rate and velocity.[17] In this case, we did some –minimal– manipulation by plotting the two variables along different axes. This manipulation amounts to making a linear transformation of one of the variables, which is consistent with the theoretical constructs in (9) and (10). We also go beyond the previous comparison with the theory, where we used the log-log specification for all countries. The theory suggests that the formulation (10) is more likely to be a better description of the data when the borrowing constraints are quantitatively relevant. Thus, we classify countries into two groups. For the first group of countries, we plot the log of the interest rate and the log of velocity, as specified in equation (9). For the second group of countries, we plot only the levels of the same variables, as specified in equation (10). This second group of countries comprises those that never had their interest rates too high and even had several years of interest rates close to zero, as the theory suggests.[18]

For the 26 countries in Figures 4a and 4b, the comovement between velocity and the interest rate is very strong. The same caveats regarding Argentina and Brazil apply, and notice that in both Israel and Bolivia, our data mostly cover the years following the stabilization. Three more puzzles appear clearly in the time series: Venezuela, and to a lesser extent Ecuador, seem to exhibit

[17]We find it convenient to show velocity – the inverse of real money demand according to the model – since it ought to be positively related to the interest rate.

[18]We use the following criterion: we present the data in levels for all countries with some observations below 5% and no observations beyond 100%, and we present the data in logs for all other countries.

different behavior during the first half of the sample relative to the second. In addition, the data for Chile show, as do the data for Argentina and Brazil mentioned above, that the behavior right after the stabilization of a very high inflation does not conform to the theory. The countries in the top row of Figure 4c also exhibit solid evidence of comovement between velocity and the interest rate, though it is less clear-cut than the previous cases. Finally, among the group of countries in Figure 4c, 4 of them suggest blunt failure, whereas both Norway and Portugal seem to conform to the theory in the last few decades, but not before.

To summarize, the first data inspection suggests the following: 26 countries exhibit remarkable evidence of comovement between velocity and the interest rate, 6 countries offer good evidence, and for the final group of 6 countries, the evidence is either weak or nonexistent. In addition, an inspection of the countries that experienced extremely high inflation suggests that in some cases, it takes several years after nominal interest rates have returned to normal for real money demand to recover to its previous levels. Finally, we identify only three candidates (Venezuela, Norway and Portugal) that seem to exhibit a breakdown in the money demand relationship that could justify further analysis of whether regulation may have played a role—an analysis along the same lines discussed in Lucas and Nicolini (2015) for the United States.

Despite the attractiveness of looking at simple plots, however, the previous analysis has several limitations. We would like to formally test whether,

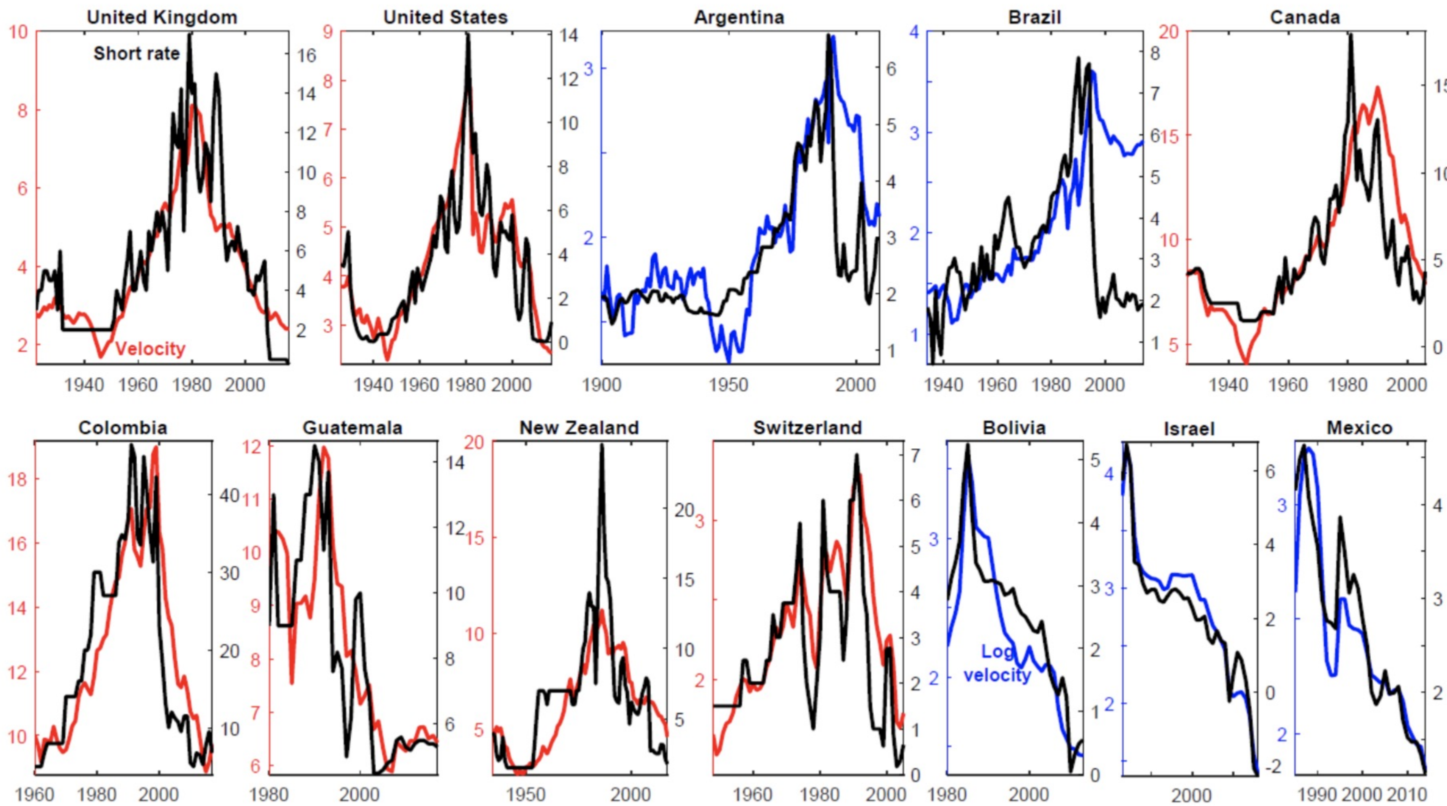

Figure 4a The raw data: M1 velocity and the short rate

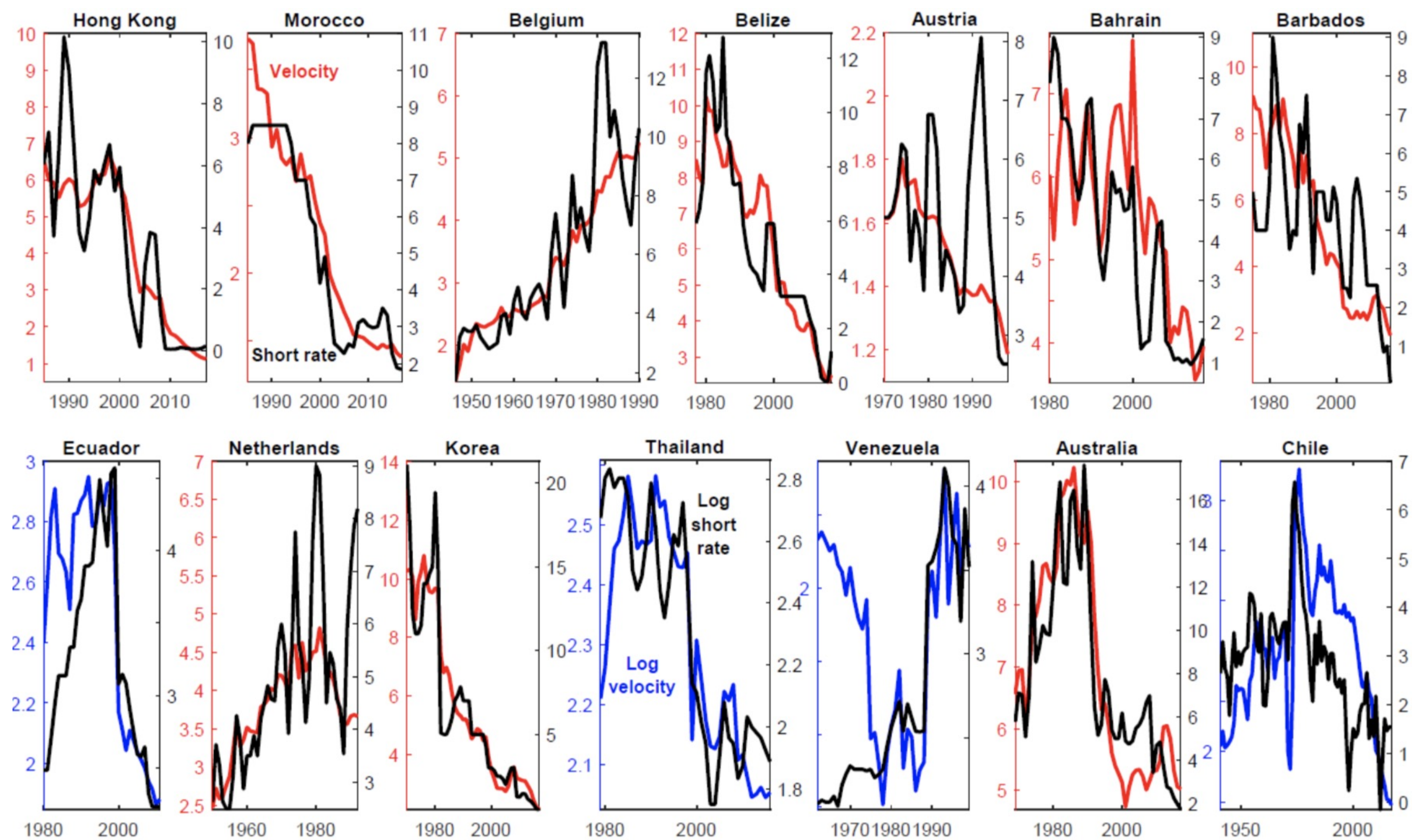

Figure 4b The raw data: M1 velocity and the short rate

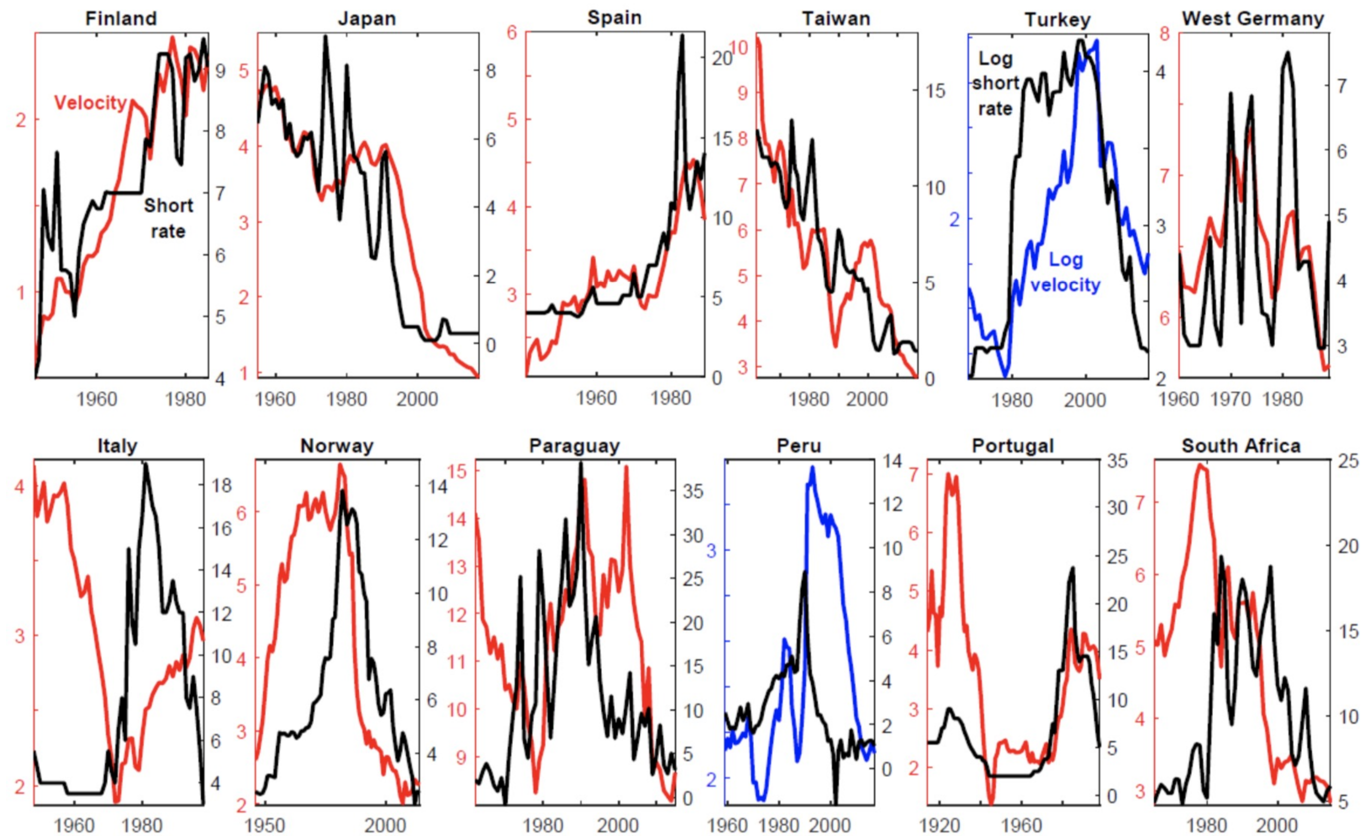

Figure 4c The raw data: M1 velocity and the short rate

as the simple model above implies, the ratio between real money balances and output inherits a unit root when the short-term interest rate exhibits a unit root. We also want to formally estimate the interest rate elasticities and see how good $1/2$ is as an approximation, the value implied by the linear technology in the original Baumol-Tobin specification. We would also like to let the data indicate the quantitative effect of the borrowing constraints when interest rates are very low, as it has been in countries such as Japan, the United States, and the United Kingdom. If this were the case, the Selden-Latané specification ought to deliver better results.

The plots in Figures 4a through 4c show how persistent the series are and provide support for the use of the cointegration methods that we use below. In fact, the statistical tests overwhelmingly identify unit roots both in the ratio of money to output and in the interest rates (see Online Appendix C). In spite of the results of the unit root tests, one may have theoretical reasons to believe that the interest rate, being a policy variable, ought to be stationary. One such reason is that in most monetary models used to justify inflation targeting policies, the policies that stabilize the economy around the inflation target deliver a stationary series for the equilibrium interest rate, even if the economy is subject to unit root shocks, as long as the real interest rate is stationary. And these models approximate very well the behavior of inflation rates in the data in countries that have successfully managed to keep inflation very close to its target. Clearly, temporary but persistent deviations from that policy, as one may interpret the US experience from 1965 to 1985, may imply very persistent movements in the interest rate. And it is indeed quite difficult to distinguish that behavior from a unit root, statistically speaking, given the our sample size.

The good news is that those speculations are of little quantitative relevance: a crucial feature of one of the methods we use in Section 6.1, owing to Wright (2000), is that they perform very well for highly persistent series, even if they are not exactly unit roots. We illustrate this property in the specific context of the model described above. There we run Monte Carlo experiments with model-generated data. We simulate stationary but very persistent series and show that these methods identify the true parameters very well. It is because we are handling very persistent series that we are fully comfortable embracing the cointegration techniques that follow.

We now turn to a brief discussion of the main features of our approach and several methodological issues.

4. Main Features of Our Approach

The cointegration techniques we use were justified above on statistical grounds: the series we work with are highly persistent, to the point where, in nearly all cases, it is not possible to reject the null of a unit root. The notion of cointegration boils down to the existence of a long-run relationship between series driven by permanent shocks: those shocks are the source of identification of the relationship between the short rate and velocity. The existence of the cointegration relationship implies that, in the long run, any permanent increase in the interest rate maps into a corresponding permanent increase in velocity and therefore a decrease in real money balances; the exact amount will be captured by the cointegration vector. Further, any deviation of the two series from their long-run relationship—that is, the cointegration residual—is transitory and bound to disappear in the long run. The persistence of the residual is therefore a measure of how long-lived these short-run deviations are. As we will show, these estimated residuals are indeed very persistent themselves. Our analysis therefore leaves unexplained a substantial fraction of the dynamic interactions between the short-term interest rate and the money-to-output ratio. As mentioned in the introduction, models with segmented markets, such as the one in Alvarez and Lippi (2014) have been developed to account for these persistent short-run deviations. The statistical properties of the cointegration errors we obtain can be used to discipline those models. Developing models that can successfully integrate long-run with short-run behavior are left for future research.

We present the analysis in two steps. In the first step, discussed in this section, we take the results of the unit root test literally—that is, we interpret the lack of rejection of the null hypothesis as evidence that the series contain *exact* unit roots—and present evidence from Johansen's cointegration tests, which take no cointegration as the null hypothesis. Then, in a second step, we present the results from Wright's (2000) tests, which take cointegration as the null hypothesis.[19] There are (at least) two reasons for also considering the Wright test, in addition to the Johansen test. First, although the overwhelming majority of the papers in the money demand literature have been based on Johansen's procedure, there is no reason to regard no cointegration as the "natural null hypothesis." Rather, it might be argued that, since we are searching here for a long-run money demand for *transaction* purposes, cointegration should be the natural null.[20] Second, as discussed by Wright (2000), Wright's test works equally well both when the data contain exact unit roots, and when they are local-to-unity. On the other hand, as shown by Elliot (1998), when the data are local-to-unity, tests (such as Johansen's) that are predicated on the assumption that the data contain exact unit roots can perform poorly.

Once cointegration is detected, we can use standard methods to estimate the parameter that governs the elasticity of the money demand relationship. For reasons of robustness, we consider Johansen's just-mentioned procedure, as well as Stock and Watson's (1993). We also compare the results of using the log-log specification and the Selden-Latané one, and use those results to

[19]The Wright (2000) test searches across the parameter space for all the values of β in the normalized cointegration vector $[1 - \beta]'$ such that the null hypothesis that $[1 - \beta]'.y_t$ is $I(0)$ cannot be rejected, where y_t features the levels of velocity and the short rate for the Selden-Latané specification, and their logarithms for the log-log. A $(1-\alpha)$%-coverage confidence interval for β is computed as the set of all values of β for which the null hypothesis that $[1 - \beta]'.y_t$ is $I(0)$ cannot be rejected at the α% level.

[20]Basic economics logic suggests that, up to fluctuations in the opportunity cost of money, the nominal quantity of money demanded should be proportional to the nominal volume of transactions (proxied by nominal GDP).

discuss the behavior of real money demand at very low interest rates.

4.1. Integration properties of the data

Online Appendix C reports evidence from our extensive investigation of the integration properties of the data based on the unit root tests in Elliot, Rothenberg and Stock (1996)Ȯur main results can be summarized as follows.

First, there is overwhelming evidence of unit roots in the vast majority of the series, with the bootstrapped p-values being near-uniformly greater than the 10% threshold, which, throughout the entire paper, we take as our benchmark significance level, and in most cases markedly so. In the very few instances in which this is not the case, we eschew the relevant specifications (e.g., if we can reject the null of a unit root for the logarithm of the short rate but not for the level, we eschew the log-log specification, and we uniquely focus on the Selden-Latané specification).

Second, for both the first difference and the log-difference of either velocity or the short rate, the null of a unit root can be rejected almost uniformly. In the few instances in which this is not the case—so that the relevant series should be regarded, according to the tests in Elliot, Rothenberg and Stock (1996), as $I(2)$—we do not run cointegration tests.[21] As for nominal M1 and especially nominal GDP, on the other hand, the opposite is true, with the null of a unit root not being rejected most of the time. In all of these

[21] Both Johansen's and Wright's tests are predicated on the assumption that the series contain (near) unit roots, but that their order of integration is at most one.

cases, we will therefore eschew unrestricted specifications for the logarithms of nominal M1, nominal GDP, and a short rate.

4.2. Methodological issues pertaining to cointegration tests

In this section, we briefly discuss three issues. First, we describe the bootstrapping procedures we use. Second, we describe a series of Monte Carlo experiments that help to interpret the test results. Finally, we discuss evidence related to the persistency of the cointegration residual.

4.2.1. Issues pertaining to bootstrapping

All model statistics in this paper are bootstrapped. In this section, we briefly discuss details of the bootstrapping procedures we use and how such procedures perform. In our discussion, we extensively refer to online Appendix D, which contains the Monte Carlo evidence motivating both our choices and the way in which we will interpret the evidence based on the actual data.

As for Johansen's tests, we bootstrap trace and maximum eigenvalue statistics via the procedure proposed by Cavaliere, Rahbek and Taylor (2012; henceforth, CRT). In a nutshell, for tests of the null of no cointegration against the alternative of one or more cointegrating vectors, the model that is being bootstrapped is a simple, noncointegrated VAR in differences (for the maximum eigenvalue tests of h versus $h+1$ cointegrating vectors, on the other hand, the model that ought to be bootstrapped is the VECM estimated under the null of h cointegrating vectors). All of the technical details can be found in CRT, to which the reader is referred. We select the VAR lag order as the maximum[22] between the lag orders chosen by the Schwartz and the Hannan-Quinn criteria[23] for the VAR in levels.

As for the Wright (2000) test, since it has been designed to be equally valid for data-generation processes (DGPs) featuring either exact or near unit roots, we consider two alternative bootstrapping procedures, corresponding to either of the two possible cases. (In practice, as a comparison between the results reported in Table 2 in the text and Table E.1 in online Appendix E makes clear, the two procedures produce near-identical results.) The former procedure involves bootstrapping—as detailed in CRT, and briefly recounted in the previous paragraph—the cointegrated VECM estimated (based on Johansen's procedure) under the null of one cointegration vector. This bootstrapping procedure is the correct one if the data feature exact unit roots. For the alternative possible case in which velocity and the short rate are near unit root processes, we proceed as follows. Based on the just-mentioned cointegrated VECM estimated under the null of one cointegration vector, we compute the implied VAR in levels, which, by construction, features one, and only one eigenvalue equal to 1. Bootstrapping this VAR would obviously be

[22] We consider the maximum between the lag orders chosen by the SIC and HQ criteria because the risk associated with selecting a lag order smaller than the true one (model misspecification) is more serious than the one resulting from choosing a lag order greater than the true one (overfitting).

[23] On the other hand, we do not consider the Akaike Information Criterion since, as discussed by Luetkepohl (1991), for example, for systems featuring $I(1)$ series, the AIC is an inconsistent lag selection criterion, in the sense of not choosing the correct lag order asymptotically.

equivalent to bootstrapping the underlying cointegrated VECM, that is, it would be correct if the data featured exact unit roots. Since, on the other hand, here we want to bootstrap under the null of a near unit root DGP, we turn such an exact unit root VAR in levels into its corresponding near unit root, by shrinking down the single unitary eigenvalue to λ=1-0.5$\cdot(1/T)$, where T is the sample length.[24] The bootstrapping procedure we implement for the second possible case, in which the processes feature near unit roots, is based on bootstrapping such a near unit root VAR.

We now turn to discussing Monte Carlo evidence on the performance of the two bootstrapping procedures.

4.2.2. Monte Carlo evidence

Tables D.1 and D.2 in online Appendix D report extensive Monte Carlo evidence on the performance of the bootstrapping procedures, which is discussed in detail in Sections D.1.1 and D.1.2 of online Appendix D. We perform the Monte Carlo experiments based on two types of DGPs, featuring no cointegration and cointegration, respectively. For either DGP, we consider several alternative sample lengths, from T = 50 to T = 1,000. For the DGP

[24] In particular, we do that via a small perturbation of the parameters of the VAR matrices B_j's in the cointegrated VECM representation $\Delta Y_t = A + B_1 \Delta Y_{t-1} + ... + B_p \Delta Y_{t-p} + G Y_{t-1} + u_t$, where Y_t collects (the logarithms of) M1 velocity and the short rate, and the rest of the notation is obvious. By only perturbating the elements of the VAR matrices B_j's—leaving unchanged the elements of the matrix G (and therefore both the cointegration vector and the loading coefficients)—we make sure that both the long-run equilibrium relationship between velocity and the short rate, and the way in which disequilibria in such a relationship map into subsequent adjustments in the two series, remain unchanged.

featuring cointegration, we also consider several alternative values for the persistence of the cointegration residual, which we model as an AR(1). Finally, for the experiments pertaining to Johansen's tests, we only consider DGPs with exact unit roots, but for those pertaining to Wright tests we also consider the corresponding DGPs with roots local-to-unity, which we obtain by replacing, in the former DGPs, the exact unit root with λ=1-0.5$\cdot$(1/T). In the case of cointegrated DGPs featuring exact unit roots, we bootstrap Wright's test statistics based on the first procedure discussed in the previous subsection (that is, based on bootstrapping the VECM estimated conditional on one cointegration vector, as in CRT). In the case of cointegrated DGPs featuring near unit roots, on the other hand, we bootstrap the tests via the alternative procedure, based on bootstrapping the corresponding near unit root VAR in levels.

Our main results can be summarized as follows.

As for the Johansen test, if the true DGP features no cointegration, CRT's procedure performs remarkably well irrespective of sample size, with empirical rejection frequencies (ERFs) very close to the nominal size. This is in line with the Monte Carlo evidence reported in CRT's Table I, p. 1731, and with the analogous evidence reported in Benati (2015). If, however, the true DGP features cointegration, the tests perform well only if the persistence of the cointegration residual is sufficiently low, the sample size is sufficiently large, or both: if the residual is persistent, the sample is short, or both, the tests fail to detect cointegration a nonnegligible fraction of the time. This is in line with some of Engle and Granger's (1987) evidence, and it has a simple explanation: as the residual becomes more and more persistent, it gets closer and closer to a random walk (in which case there would be no cointegration), so that the procedure needs larger and larger samples to detect the truth (i.e., that the residual is highly persistent but ultimately stationary).

As for the Wright test, evidence is qualitatively the same, and quantitatively very close, in the case of either exact or near unit root DGPs. Specifically, if the true DGP features cointegration, the procedure works remarkably well if the sample size is sufficiently long, the persistence of the cointegration residual is sufficiently low, or both, with ERFs very close to the tests' nominal size. As the sample size decreases and/or the persistence of the cointegration residual increases, however, the ERFs increase systematically, to the point where, for example, for $T = 50$ and the autoregressive parameter of the cointegration residual equal to 0.95, the test incorrectly rejects the null at about twice the nominal size. The explanation for this is straightforward, and it has to do, once again, with Engle and Granger's (1987) previously mentioned point: when the cointegration residual is highly persistent, only sufficiently long samples allow the test to detect the truth (i.e., that the deviation between the two series is ultimately transitory, so that they are in fact cointegrated). But, under these circumstances, the shorter the sample period, the more likely it will be to mistakenly infer that the deviation between the series is permanent, so that they are not, in fact, cointegrated. If, however, the true DGP features no cointegration, the test tends to reject the null at about twice the nominal size, essentially irrespective of sample length.

These results can be summarized as follows. If the Johansen test detects cointegration, we should have a reasonable presumption that cointegration is there. If, however, it does not detect it, a possible explanation is that the sample is too short, the cointegration residual is highly persistent, or both. As for the Wright test, lack of rejection of cointegration does not represent very strong evidence that cointegration truly is there, as this also happens with a comparatively high frequency for DGPs featuring no cointegration.

We now turn to the issue of how persistent cointegration residuals in fact are.

4.2.3. Evidence on the persistence of cointegration residuals

Tables H.1 and H.2 in Online Appendix H report Hansen's (1999) "grid bootstrap" median-unbiased (henceforth, MUB) estimates of the sum of the AR coefficients in AR(2) representations for the "candidate cointegration residuals" in our dataset.[25] By "candidate cointegration residual" (henceforth, CCR), we mean the linear combination of the variables in the system that will indeed be regarded as a cointegration residual *if* cointegration is detected.[26] For reasons of robustness, for either the Selden-Latané specification

[25]Results are based on 2,000 bootstrap replications for each possible value of the sum of the AR coefficients in the grid. Bootstrapping has been performed as in Diebold and Chen (1996). For reasons of robustness, we report results based on two alternative estimators of the cointegration vector, Johansen's and Stock and Watson's (1993).

[26]We label it as a candidate cointegration residual because, as the Monte Carlo evidence in the previous section has shown, if a residual is highly persistent, cointegration might

(Table H.1) or the log-log specification (Table H.2) we consider two alternative estimators of the cointegration residual: either Johansen's or Stock and Watson's (1993).

Evidence points toward both a nonnegligible extent of persistence of the CCRs and a wide extent of heterogeneity across countries. Focusing on results based on the log-log specification, the MUB estimate based on Johansen's estimator of the cointegration vector—let's label it as $\hat{\rho}^J_{MUB}$—ranges from a minimum of 0.27 for Belize to a maximum of 1.17 for the Barbados islands. By classifying the $\hat{\rho}^J_{MUB}$'s, in an admittedly arbitrary fashion, as "highly persistent" ($\hat{\rho}^J_{MUB} \geq 0.8$); "moderately persistent" ($0.4 < \hat{\rho}^J_{MUB} < 0.8$); and "not very persistent" ($\hat{\rho}^J_{MUB} \leq 0.4$), we end up with 22 $\hat{\rho}^J_{MUB}$'s in the first group, 14 in the second, and 4 in the third. Results based on Stock and Watson's estimator are qualitatively the same, with the three groups comprising, respectively, 25, 13, and 2 countries.

Under these circumstances, statistical tests will often have a hard time detecting cointegration even if it truly is present, especially when $\hat{\rho}_{MUB}$ is high and the sample period is comparatively short. This implies that results from cointegration tests should not be taken strictly at face value, but rather should be interpreted in the light of the previously mentioned Monte Carlo evidence in Tables D.1 and D.2 in Online Appendix D.

well not be detected even if it is present, which would prevent the candidate from being identified as a true cointegration residual.

5. Results

In presenting the results, we first discuss the cointegration tests and then show the parameter estimates. We finish with a comparison between the log-log and Selden-Latané specifications.

5.1. Cointegration tests

In this section, we discuss the results from bivariate systems for velocity and the short rate, as implied by equation (7).[27] Table 1 reports results from Johansen's maximum eigenvalue test of 0 versus 1 cointegration vectors, together with the Monte Carlo-based ERFs computed conditional on the null of one cointegration vector. We highlight in yellow all p-values for maximum eigenvalue tests smaller than 10% and all ERFs smaller than 50%, corresponding to a less-than-even chance of detecting cointegration if it is truly present in the data.

The table reports the cointegration test results for only one of the functional forms when the unit root tests for either the level or the log of one of the variables were rejected. For two of the countries, Morocco and Hong Kong, the series were identified as $I(2)$ for both the level and the log of the variables, so no cointegration tests are reported. Thus, Table 1 reports results for only 36 countries.

We ordered the countries according to the test results. Within each cate-

[27] In Online Appendix F, we discuss test results for unrestricted specifications between the log of the interest rate, the log of nominal output, and the log of M1.

Table 1 Bootstrapped p-values[a] for Johansen's maximum eigenvalue[b] test and empirical rejection frequencies of the tests under the null

		I: Bootstrapped p-values		II: Empirical rejection frequencies	
Country	*Period*	*Selden-Latané*	*Log-log*	*Selden-Latané*	*Log-log*
United Kingdom	1922-2016	**0.003**	0.793	0.975	0.661
US – M1 + MMDAs	1915-2017	**0.063**	0.212	0.859	0.625
US – M1	1915-2017	0.869	0.218	**0.099**	**0.275**
Argentina	1914-2009	–	**0.031**	–	0.789
Brazil	1934-2014	–	**0.093**	–	**0.341**
Canada	1967-2017	**0.015**	**0.028**	0.965	0.939
	1926-2006	**0.007**	0.361	0.968	0.628
Colombia	1960-2017	**0.032**	**0.027**	0.648	0.593
Guatemala	1980-2017	**0.007**	**0.038**	0.536	**0.448**
New Zealand	1934-2017	**0.099**	**0.030**	0.690	0.819
Switzerland	1948-2005	**0.000**	**0.017**	0.923	0.769
Bolivia	1980-2013	**0.053**	0.681	0.686	**0.125**
Israel	1983-2016	**0.000**	0.252	0.767	**0.197**
Mexico	1985-2014	**0.007**	0.505	0.537	**0.200**
Belgium	1946-1990	0.361	**0.062**	0.699	0.721
Belize	1977-2017	0.704	**0.007**	**0.107**	**0.394**
Austria	1970-1998	0.203	0.180	**0.220**	**0.265**
Bahrain	1980-2017	0.401	0.335	**0.082**	**0.085**
Barbados	1975-2016	0.542	0.677	**0.348**	**0.335**
Ecuador	1980-2011	0.838	0.686	**0.043**	**0.053**
Netherlands	1950-1992	0.349	0.568	**0.463**	**0.325**
South Korea	1970-2017	0.364	0.955	**0.086**	**0.169**
Thailand	1979-2016	0.101	0.212	**0.101**	**0.091**
Venezuela	1962-1999	0.776	0.922	**0.087**	**0.059**
Australia	1969-2017	0.134	0.960	0.720	**0.425**
	1941-1989	0.642	0.722	**0.168**	**0.198**
Chile	1941-2017	0.442	0.307	0.824	0.628
	1940-1995	0.133	0.175	**0.111**	0.859
Finland	1946-1985	0.246	–	**0.286**	–
Japan	1955-2017	0.567	0.142	**0.363**	0.590
Spain	1941-1989	0.120	0.196	0.636	**0.205**
Taiwan	1962-2017	–	0.909	–	**0.139**
Turkey	1968-2017	0.460	–	**0.229**	–
West Germany	1960-1989	–	0.352	–	**0.219**
Italy	1949-1996	0.171	–	0.629	–
Norway	1946-2014	**0.035**	**0.043**	0.749	0.789
Paraguay	1962-2015	**0.074**	0.168	**0.443**	**0.366**
Peru	1959-2017	**0.003**	0.171	0.992	**0.389**
Portugal	1914-1998	0.857	**0.047**	**0.048**	**0.059**
South Africa	1965-2015	0.116	0.157	0.563	**0.329**

[a] Based on 10,000 bootstrap replications. [b] Null of 0 versus 1 cointegration vectors.

gory, we ordered the countries alphabetically. For each country, we first mention the time period for which we have consistent data. In some cases (Australia, Canada, and Chile), we have two different datasets, for long enough periods, but they do not completely overlap. The series are not exactly the same, so they cannot be used to construct a single series that can suitably be analyzed using cointegration. We report the results using both series. The third and fourth columns report the p-values of the tests for both the Selden-Latané and the log-log specifications. Finally, we show the ERF of the Monte Carlo exercises for both the Selden-Latané and the log-log specifications.

We first report the results for the United Kingdom and the United States, for which we have close to a century of data. For the case of the United States, we use both M1 and New M1 (the monetary aggregate proposed in Lucas and Nicolini 2015). The second group of countries contains the ones for which both p-values (or the only one that we could run) are below 10%. The next two groups include countries for which one and only one of the p-values is below 10%. The fifth and sixth groups contain countries for which both p-values are above 10%, but either the two ERFs are below 50% (fifth group) or only one is below 50% (sixth group).[28] Finally, the last group includes the six countries for which we believe the evidence is weak or nonexistent based on the visual evidence, in spite of the test results.

We first discuss how to interpret the United States and United Kingdom

[28]In classifying countries for which we have more than one set of series, we chose the one that contains the most recent data.

results in detail. The other numbers in the table are to be interpreted in a similar way. Recall that for both countries, the evidence displayed in the simple graphs in Figures 3 and 4 is quite remarkable. The results of the tests confirm that notion. In using M1 for the United Kingdom and New M1 for the United States, the p-values for the Selden-Latané specification are below 10%, but the ones for the log-log specification are both above 10%. For both countries, the ERFs are substantially larger than 50%. This strong preference of the data for the Selden-Latané specification is consistent with the fact that both countries had several periods with interest rates very close to zero. Taken together, these results provide evidence of a satiation point at zero in the aggregate real money demand. Finally, when using the standard M1 aggregate for the United States, both p-values are higher than 10%, although the ERFs are below 50% for both specifications, indicating that the power of the test is low.

With the exception of Hong Kong and Morocco—for which we could not run the tests—the order of the countries in the table is the same as the order of the countries in the figures in Section 3. The first four groups contain 14 countries in total. For all of them, the tests detect cointegration in at least one of the specifications, even though in several cases the ERFs are low. The next two groups contain a total of 16 countries for which cointegration is not detected, but the ERFs are low in the two tests (8 countries in group 5) or in one test (8 countries in group 6). For these 30 countries, the visual evidence is very good – with the caveat of the behavior right after the stabilization of very high inflations, as in Argentina and Brazil.

The final group contains 6 countries for which the visual evidence was problematic or nonexistent. Two problematic cases are Norway and Portugal. In both cases, the tests do detect cointegration in at least one of the specifications. However, the visual evidence suggests a different behavior over time, somewhat similar to what occurred with M1 in the United States. Exploring whether regulation could explain those 2 cases seems to be worth an avenue pursuing, but not in this paper. For the remaining 4 countries, the visual evidence does not suggest such a pattern (or any other pattern!). Even though in 2 of those 4 cases the tests do detect cointegration, we can only classify those 4 countries as failing to behave as the theory implies.

To summarize: we find supporting evidence for 32 out of the 38 countries analyzed (the 30 countries in groups 1 to 6, plus Hong Kong and Morocco). Of the remaining 6 countries, regulatory changes may explain the behavior of real money demand in 2, whereas the other 4 are blunt failures.

In Table 2, we present the results for the Wright test. We report 90% confidence intervals for the second element of the normalized cointegration vector $(1-\beta)$. As mentioned above, they represent the set of all values of β for which the null hypothesis that $(1-\beta)\prime.y_t$ is $I(0)$ cannot be rejected at the $\alpha\%$ level, where y_t is a vector that contains either the levels or the logs of the short rate and velocity. The order of the countries is the same as in Table 1. In those cases in which cointegration is not detected, the entry in the table is NCD. We highlight in yellow the cases in which the confidence

Table 2 Results from the Wright (2000) test: 90% coverage confidence intervals for the second element of the normalized cointegration vector

Country	*Period*	Bootstrapped process: cointegrated VECM *Selden-Latané*	*Log-log*
United Kingdom	1922-2016	**[-0.529; -0.417]**	NCD
US - M1 + MMDAs	1915-2017	**[-0.613; -0.393]**	**[-0.352; -0.108]**
US - M1	1915-2017	NCD	**[-0.506; -0.029]**
Argentina	1914-2009	**[-0.107; -0.087]**	**[-0.513; -0.245]**
Brazil	1934-2014	**[-0.065; -0.009]**	[-1.366; 0.276]
Canada	1926-2006	**[-1.490; -1.053]**	**[-0.719; -0.607]**
	1967-2017	**[-0.578; -0.494]**	**[-0.389; -0.345]**
Colombia	1960-2017	NCD	NCD
Guatemala	1980-2017	**[-0.752; -0.448]**	**[-0.678; -0.414]**
New Zealand	1934-2017	NCD	**[-0.589; -0.312]**
Switzerland	1948-2005	NCD	NCD
Bolivia	1980-2013	**[-0.369; -0.193]**	**[-0.520; -0.388]**
Israel	1983-2016	NCD	**[-0.388; -0.320]**
Mexico	1985-2014	**[-0.260; -0.184]**	**[-0.422; -0.314]**
Belgium	1946-1990	**[-0.465; -0.289]**	**[-1.146; -0.710]**
Belize	1977-2017	**[-0.840; -0.692]**	**[-2.567; 1.433]**
Austria	1970-1998	[-0.601; 0.080]	[-1.040; 0.618]
Bahrain	1980-2017	NCD	**[-0.254; -0.194]**
Barbados	1975-2016	**[-2.006; -0.748]**	**[-2.899; 0.101]**
Ecuador	1980-2011	NCD	NCD
Netherlands	1950-1992	**[-0.394; -0.290]**	**[-0.483; -0.331]**
South Korea	1970-2017	**[-0.617; -0.521]**	**[-0.639; -0.338]**
Thailand	1979-2016	NCD	**[-0.498; -0.386]**
Venezuela	1962-1999	NCD	[-0.249; 0.287]
Australia	1941-1989	**[-0.691; -0.526]**	**[-0.808; -0.704]**
	1969-2017	**[-0.484; -0.404]**	**[-0.506; -0.314]**
Chile	1940-1995	**[-0.140; -0.028]**	**[-0.382; -0.278]**
Finland	1946-1985	**[-0.530; -0.414]**	**[-2.693; -1.780]**
Japan	1955-2017	**[-0.520; -0.312]**	**[-0.513; -0.125]**
Spain	1941-1989	**[-0.163; -0.159]**	NCD
Taiwan	1962-2017	**[-0.449; -0.341]**	**[-0.453; -0.253]**
Turkey	1968-2017	NCD	NCD
West Germany	1960-1989	[-0.963; 0.931]	[-0.489; 0.692]
Italy	1949-1996	[0.032; 0.204]	[0.159; 0.511]
Norway	1946-2014	[-0.961; 0.985]	[-0.227; 1.043]
Paraguay	1962-2015	[-0.328; 0.125]	**[-0.200; -0.024]**
Peru	1959-2017	[-0.042; 0.026]	[-0.493; 0.692]
Portugal	1914-1998	[-0.340; 0.433]	[-0.018; 0.210]
South Africa	1965-2015	[-0.170; 0.427]	[-0.052; 1.065]

NCD = No cointegration detected.

interval lies entirely in the range of negative numbers, so that cointegration is detected in the data, and furthermore the relationship is negative, as it is in the theory. The results in Table 2 are in general even stronger than the ones in Table 1, but the results are consistent. For the 6 countries we had identified as having weak or nonexistent evidence, the results are also weak in this case. On the other hand, for the 16 countries in groups 5 and 6 of Table 1, where the Johansen test failed to detect cointegration (the p-values were above 10%) but where at least one of the ERFs were low, the Wright test identifies cointegration in 15 of them (West Germany being the only exception). For the 12 countries for which we did identify cointegration using the Johansen test, there is conflicting evidence for only one country: results for Switzerland are very strong in Table 1 but not in Table 2.[29]

5.2. The estimated coefficients on the short rate

The coefficients can be estimated using both Johansen's and Stock and Watson's procedures. In addition, the Wright test also delivers confidence intervals for the coefficient on the short rate. The full set of results is presented in Online Appendix I, where we show the estimation using the three procedures for the two specifications. We will focus the discussion in this section on the estimates of the elasticity in the log-log specification using

[29]In light of the low power of the tests, one could repeat the analysis using quarterly data. For nine of the countries in our sample, the tests can be found in Benati (2020), using data since WWII. We could find quarterly (also since WWII) for an additional set of seven countries. Appendix J contains the result. In all the sixteen cases the results are in line with the ones we report in Tables 1 and 2.

Stock and Watson's procedure. The reason for focusing on the log-log specification is twofold. First, the coefficient on the log of the interest rate has the natural interpretation of an elasticity. Second, and most importantly, we can directly relate it to the transactions technology that is the key component of the theory. In that respect, our reference value has an elasticity equal to 0.5, which corresponds to the case of a linear technology, as in Baumol and Tobin. Higher values for the elasticity imply that the exponent σ in equation (2) is lower than 1, which implies that the marginal cost of making transactions is decreasing with the number of transactions. The opposite is true when the elasticity is lower than 0.5.

The reason for focusing on Stock and Watson's estimates is that they are based on a single equation, whereas Johansen's is based on estimating an entire cointegrated VAR, so there are many more coefficients. In small samples this approach may produce less precise results. In almost all cases, Johansen's estimates are broadly in line with Stock and Watson's but typically deliver larger standard errors. In 4 cases the results are different, and we conjecture that this result might be a small-sample issue. For details, see Online Appendix I. We are aware that as long as the data-generating process corresponds to the model with the borrowing constraint, the estimate of the elasticity will be biased downward in countries with several observations of interest rates near zero. The reason is that the procedure will try to match those observations with low interest rates that all lie below the log-log curve that has a good fit with the observations for moderate and high interest rates.

However, given that the number of observations at very low rates is not a very large fraction of the total sample, we expect this bias to be small.

In Figure 5, we present the results of the estimation for 33 countries using the Stock and Watson procedure. As explained above, for 2 countries (Hong Kong and Morocco), we could not run the tests, and for 3 of them (Finland, Italy, and Turkey), the test did not detect cointegration for the log-log specification. The horizontal axis represents the value for the estimator of the elasticity for the corresponding country, ranging from -1 to 1. For each country, we report the point estimator with a black rhombus and the 90% confidence interval with a dotted red line. We order the countries according to the point estimate, starting with the lowest one. Finally, in the figure we plot two vertical lines: one at zero, which corresponds to the null of no relationship between the log of the interest rate and the log of real money demand, and one at the value 0.5, which corresponds to the linear technology assumed by Baumol and Tobin.[30]

As a summary of the figure, the confidence interval includes zero for only 4 out of the 33 countries. Two of them (Norway and South Africa) belong to the group with either weak or nonexistent evidence. In no case is the estimate statistically greater than zero. Finally, for around 20 countries, the confidence interval includes 0.5 or is remarkably close to it. Table 2, together with the country plots in Figures 3a to 3c, gives very strong support to a

[30]For the 3 countries for which we had two different sets of data, here we report the set that includes the latest observations.

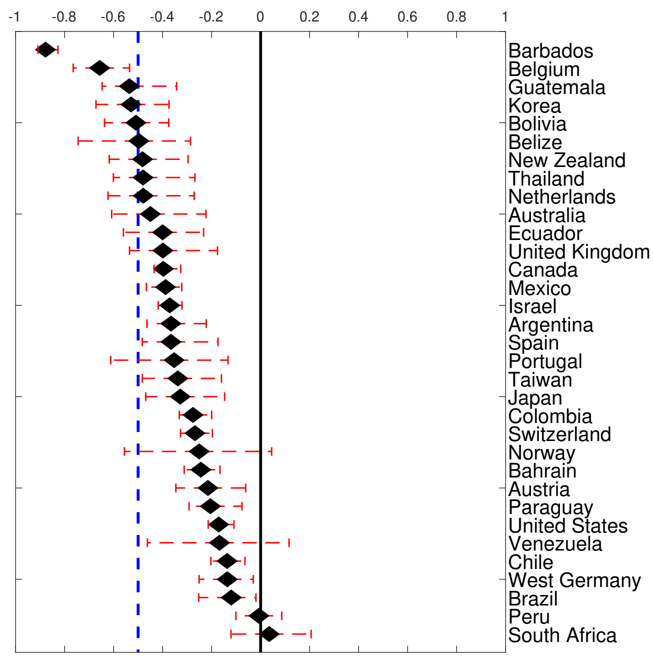

Figure 5 Estimation results using the procedure from Stock and Watson (1993)

rather simple theory that, in essence, was developed over half a century ago.

5.3. Evidence on the functional form

Figure 6 provides simple, informal evidence on which specification – the Selden-Latané or the log-log – provides the most plausible description of the data at both low and high interest rates. For selected low-inflation and high-inflation countries, the top row shows the levels of M1 velocity and the short rate, and the bottom row shows the logarithms of the two series. By plotting the series with different axes, we search for a linear relationship between either the levels or the logs. The evidence in the top row therefore corresponds to a Selden-Latané specification and the bottom row to the log-log specification.

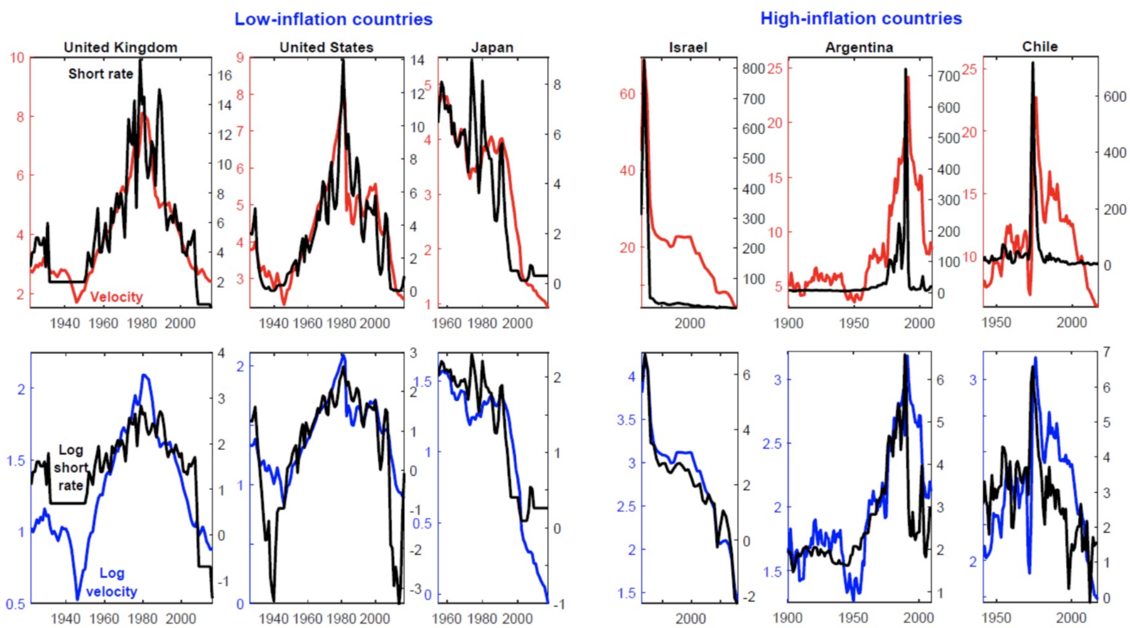

Figure 6 Comparing the Selden-Latané and log-log specifications: selected evidence for low-inflation and high-inflation countries

Two broad patterns emerge from Figure 6. First, for the low-inflation countries, both formulations do a very good job at capturing the rise and fall of both velocity and interest rates in the United Kingdom and the United States, and the persistent fall of both in Japan. The figure clearly shows, however, that the log-log specification is substantially worse when interest rates are close to zero for the three countries. This result is in line with our discussion of the borrowing constraints in the theoretical section. Second, for the high-inflation countries, the opposite is true: the specification in logs appears to deliver a linear relationship, whereas the specification in levels does not.[31] This overall pattern is consistent with the theory, where borrowing constraints play an important role in low interest rates.

We did not select countries in Figure 6 randomly; rather, they are the ones that either had interest rates close to zero for many periods or had very high inflation rates. The full comparison is presented in Online Appendix I, and while the general message is similar, the conclusion there is not as striking as the examples shown here.

6. Two Additional Issues

We now discuss two additional issues. First, we show that the estimators of the elasticity of money demand we used in Section 5.1—which are

[31] These results are in line with the evidence produced by Benati (2019b) based on either monthly or weekly data for 20 cases of hyperinflation, from the French Revolution to Venezuela's episode: in nearly all cases, econometric evidence shows a clear and often overwhelming preference for the log-log specification.

predicated on the assumption that the series feature exact unit roots—work equally well for series that are local-to-unity. Then, we discuss tests for the stability of the money demand cointegration relationship.

6.1. Robustness of the estimates of the elasticity of money demand to near unit roots

Table 3 reports results from the following Monte Carlo experiment. We simulate the following DGP for the logarithms of the short rate, R_t, and of M1 velocity, V_t:

$$\ln R_t = \lambda \ln R_{t-1} + \epsilon_t, \text{ with } \lambda\text{=1-0.5} \cdot (1/T),\ \epsilon_t \sim i.i.d.\ N(0,1) \tag{12}$$

$$u_t = \rho u_{t-1} + v_t, \text{ with } 0 \leq \rho < 1,\ v_t \sim i.i.d.\ N(0,1). \tag{13}$$

$$\ln V_t = \alpha_0 - \alpha_1 \ln R_t + u_t \tag{14}$$

We set α_0=1 and α_1 equal to Baumol and Tobin's benchmark value of 0.5. As for ρ, we consider six possible values ranging from 0 to 0.95, corresponding to alternative extents of persistence of the cointegration residual. Finally, we consider four possible values for the sample length, T, ranging from 50 to 1,000. For each possible combination of values for T and ρ, we simulate (12)-(14) 10,000 times, and based on each artificial sample, we estimate the elasticity of money demand as we did in Section 5.1, based on either Johansen's or Stock and Watson's (1993) procedures. Table 3 reports the mean of the Monte Carlo distribution for the estimates of α_1 based on Stock

ρ	0	0.25	0.5	0.75	0.9	0.95
T = 50	0.5002	0.4983	0.5018	0.4978	0.4889	0.4992
T = 100	0.5007	0.4996	0.4995	0.5002	0.5025	0.5010
T = 200	0.5000	0.4999	0.4997	0.4989	0.4990	0.4982
T = 1000	0.5000	0.5000	0.5000	0.4998	0.5002	0.5004

Table 3: Mean of Monte Carlo distribution for alternative values of T and ρ

and Watson's procedure (results based on Johansen's procedure are qualitatively the same, and they are available upon request). The evidence in the table speaks for itself and shows that the estimates of the elasticity of money demand we discussed in Section 5.1 are in fact robust to the series being local-to-unity, rather than featuring exact unit roots.

6.2. Testing for stability in cointegration relationships

We test for stability in cointegration relationships based on the three tests discussed by Hansen and Johansen (1999): two Nyblom-type tests for stability in the cointegration vector and the vector of loading coefficients, respectively; and a fluctuation test, which is essentially a joint test for time variation in the cointegration vector and the loadings. In either case, we bootstrap the test statistics via CRT's procedure, based on the VECM estimated conditional on one cointegration vector, and not featuring any break or time variation of any kind.

Table H.1 in the Online Appendix reports Monte Carlo evidence on the performance of the tests conditional on bivariate cointegrated DGPs, for alternative sample lengths and alternative degrees of persistence of the cointegration residual, which is modeled as an AR(1). The main results can be summarized as follows. The two Nyblom-type tests exhibit an overall reasonable performance, incorrectly rejecting the null of no time variation, most of the time, at roughly the nominal size. Crucially, this is the case irrespective of the sample length and of the persistence of the cointegration residual. The fluctuation test, on the other hand, exhibits good performance only if the persistence of the cointegration residual is low. The higher the residual's persistence, however, the worse the performance, so that for example, when the AR root of the residual is equal to 0.95, for a sample length $T = 50$, the test rejects at twice the nominal size. This result is clearly problematic since, as previously discussed, residuals are typically moderately to highly persistent. In what follows, we therefore focus on the results from the two Nyblom-type tests, but we eschew instead results from the fluctuation test (these results are reported in Tables H.2 and H.5 in the Online Appendix).

Tables H.2 and H.3 in the Online Appendix report results from Hansen and Johansen's (1999) Nyblom-type tests for stability in either the cointegration vector or the vector of loading coefficients. The key finding in the two tables is that evidence of breaks in either the cointegration vector or the loading coefficients vector is weak to nonexistent. Specifically, for the United States, based on the Selden-Latané specification, the null of no breaks in either feature is never rejected. Based on the log-log specification, stability in the cointegration vector is also never rejected, whereas breaks in the loadings are detected. Evidence for other countries is qualitatively the same. For instance, based on the Selden-Latané specification, stability in the cointegration vector is rejected in three cases, whereas stability in the loadings is rejected in six cases. Results for the log-log specification are along the same lines.

7. Conclusions

Our review of real money demand behavior leads us to reach the following conclusions. First, for about 26 of the countries, there is substantial evidence of a long-run relationship between the ratio of money to nominal income and the short-term interest rate; the evidence is weaker for 6 countries and nonexistent for the remaining 6. Second, for the set of countries for which the evidence is strong, the log-log specification that implies a constant elasticity is a very good representation of the data, except when interest rates are very close to zero. Third, there is evidence of a satiation point at zero interest rates, implying that the elasticity of real balances with respect to the interest rate is lower in that range, approaching zero as interest rates go to zero.

[1] Alvarez, F., and F. Lippi (2009): "Financial Innovation and the Transactions Demand for Cash," *Econometrica*, 77(2), 363-402.

[2] Alvarez, F., and F. Lippi (2014): "Persistent Liquidity Effects and Long-Run Money Demand," *American Economic Journal: Macroeconomics*, 6(2), 71-107.

[3] Anderson, R. G., M. Bordo, and J. V. Duca (2017): "Money and Velocity During Financial Crises: From the Great Depression to the Great Recession," *Journal of Economic Dynamics and Control*, 81(C), 32-49.

[4] Barnett, W. A. (2016) : "Friedman and Divisia Monetary Measures." In: Cord, R. A., Hammond, J.D.(Eds), *Milton Friedman: Contributions to Economics and Public Policy*, Oxford University Press, Oxford, 265-291.

[5] Baumol, W. J. (1952), "The Transactions Demand for Cash: An Inventory Theoretic Approach, "*Quarterly Journal of Economics*, 66(4), 545-556.

[6] Belongia, M. T., and P. N. Ireland (2019a): "The Demand for Divisia Money: Theory and evidence," *Journal of Macroeconomics*, 61, Article 103128.

[7] Belongia, M. T., and P. N. Ireland (2019b): "A Reconsideration of Money Growth Rules," Boston College Working Papers in Economics 976, Boston College Department of Economics.

[8] Benati, L. (2015): "The Long-Run Phillips Curve: A Structural VAR Investigation, " *Journal of Monetary Economics*, 76, 15-28.

[9] Benati, L. (2019a): "Money Velocity and the Natural Rate of Interest," mimeo, University of Bern.

[10] Benati, L. (2019b): "Cagan's Paradox Revisited,"mimeo, University of Bern.

[11] Benati, L., R. E. Lucas Jr., J.-P. Nicolini, and W. Weber (2017): "International Evidence on Long-Run Money Demand," Working Paper 737, Federal Reserve Bank of Minneapolis.

[12] Carlson, John B., Dennis L. Hoffman, Benjamin D. Keen, and Robert H. Rasche (2000): "Results of a study of the stability of cointegrating relations comprised of broad monetary aggregates," *Journal of Monetary Economics*, 46(2), 345 - 383.

[13] Cavaliere, G., A. Rahbek, and A.M.R. Taylor (2012): "Bootstrap Determination of the Co-integration Rank in Vector Autoregressive Models," *Econometrica*, 80(4), 1721-1740.

[14] Diebold, F. X., and C. Chen (1996): "Testing Structural Stability with Endogenous Breakpoint: A Size Comparison of Analytic and Bootstrap Procedures," *Journal of Econometrics*, 70(1), 221-241.

[15] Elliot, G. (1998): "On the Robustness of Cointegration Methods When Regressors Almost Have Unit Roots," *Econometrica*, 66(1), 149-158.

[16] Elliot, G., T. J. Rothenberg, and J. H. Stock (1996): "Efficient Tests for an Autoregressive Unit Root," *Econometrica*, 64(4), 813-836.

[17] Engle, R. F., and C. W. J. Granger (1987): "Co-integration and Error Correction: Representation, Estimation, and Testing," *Econometrica*, 55(2), 251-276.

[18] Grossman, S., and L. Weiss (1983): "A Transactions-Based Model of the Monetary Transmission Mechanism," *American Economic Review*, 73(5), 871-880.

[19] Goldfeld, S. M., and D. Sichel (1990), "The Demand for Money," in Benjamin M. Friedman and Frank H. Hahn, ed., *Handbook of Monetary Economics Volume 1*, Amsterdam: North Holland, pp. 299 - 356.

[20] Hansen, B. E. (1999): "The Grid Bootstrap and the Autoregressive Model," *Review of Economics and Statistics*, 81(4), 594-607.

[21] Hansen, H., and S. Johansen (1999): "Some Tests for Parameter Constancy in Cointegrated VAR-Models," *Econometrics Journal*, 2(2), 306-333.

[22] Ireland, P. N. (2008): "On the Welfare Cost of Inflation and the Recent Behavior of Money Demand," *American Economic Review*, 99(3), 1040-1052.

[23] Judson, R. (2017): "The Death of Cash? Not So Fast: Demand for U.S. Currency at Home and Abroad, 1990-2016," paper presented at the

International Cash Conference 2017 - War on Cash: Is there a Future for Cash? April 25-27, 2017, Island of Mainau, Germany, Deutsche Bundesbank, Frankfurt a. M.

[24] Judson R. A., B. Schlusche, and V. Wong (2014): "Demand for M2 at the Zero Lower Bound: The Recent U.S. Experience, " FEDS Working Paper No. 2014-22.

[25] Latané, H.A. (1960): "Income Velocity and Interest Rates: A Pragmatic Approach," *Review of Economics and Statistics*, 42(4), 445-449.

[26] Lucas, R. E. Jr. (1988): "Money Demand in the United States: A Quantitative Review," *Carnegie-Rochester Conference Series on Public Policy*, 29, 137-168.

[27] Lucas, R. E. Jr., and J.-P. Nicolini (2015): "On the Stability of Money Demand," *Journal of Monetary Economics*, 73, 48-65.

[28] Luetkepohl, H. (1991): *Introduction to Multiple Time Series Analysis*, 2nd ed. Springer-Verlag.

[29] Meltzer, A. H. (1963): "The Demand for Money: The Evidence from the Time Series," *Journal of Political Economy*, 71(3), 219-246.

[30] Rotemberg, J. (1984): " A Monetary Equilibrium Model with Transactions Costs." *Journal of Political Economics*, 92(1), 40-58.

[31] Selden, R. T. (1956): "Monetary Velocity in the United States,"in M. Friedman, ed., *Studies in the Quantity Theory of Money*, University of Chicago Press, pp. 405-454.

[32] Serletis, A., and P. Gogas (2014): "Divisia Monetary Aggregates, the Great Recessions, and Classical Money Demand Functions," *Journal of Money, Credit and Banking*, 46(1), 229-241.

[33] Stock, J.H., and M.W. Watson (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems,"*Econometrica*, 61(4), 783-820.

[34] Teles, P., and R. Zhou (2005): "A Stable Money Demand: Looking for the Right Monetary Aggregate," Economic Perspectives, 29(1), 50-63.

[35] Tobin, J. (1956): "The Interest Elasticity of Transactions Demand for Cash, "*Review of Economics and Statistics*, 38(3), 241-247.

[36] Wright, J. H. (2000): "Confidence Sets for Cointegrating Coefficients Based on Stationarity Tests,"*Journal of Business and Economic Statistics*, 18(2), 211-222.