FEDERAL RESERVE BANK OF MINNEAPOLIS
Research Division

WORKING PAPER
No. 783

The Welfare Costs of Inflation

September 2021

Luca Benati
University of Bern

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

DOI: https://doi.org/10.21034/wp.783
Keywords: Money demand; Lower bound on interest rates
JEL classification: E41, E43, E52

The Welfare Costs of Inflation*

Luca Benati
University of Bern†

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

Abstract

We revisit the estimation of the welfare costs of inflation originating from lack of liquidity satiation. We use data for the United States and several other developed countries. Our computations are heavily influenced by the recent experience of very low, even negative, short-term rates observed in the countries we study. We obtain estimates that range between 0.20% and 1.5% of lifetime consumption for the United States and find even higher values for some European countries.

*We wish to thank Peter Ireland for comments on a previous draft and for very helpful suggestions. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Minneapolis, or of the Federal Reserve System.

†Department of Economics, University of Bern, Schanzeneckstrasse 1, CH-3001, Bern, Switzerland. Email: luca.benati@vwi.unibe.ch

‡Federal Reserve Bank of Minneapolis, 90 Hennepin Avenue, Minneapolis, MN 55401, United States. Email: juanpa@minneapolisfed.org

1 Introduction

We provide new estimates of the welfare cost of inflation. We follow the tradition of Bailey (1956), Friedman (1969), Lucas (2000), and Ireland (2009) in that we estimate the welfare cost using the area under the real money demand curve. Specifically, to compute the welfare cost of a given value for the interest rate— say, i_0— we compute the integral of the real money demand curve between the lower bound for the interest rate and i_0. This strategy is justified by a large class of theoretical models.

There is a wide range of estimates in the literature. For a steady state interest rate of 5%, Lucas (2000) computes the cost to be around 1.1% of lifetime consumption, which is a sizeable amount. However, Ireland (2009) challenges Lucas's interpretation of the data and obtains an estimate of a mere 0.04% percent of consumption. There are two key aspects of the money demand relationship that affect the computation, as both Lucas and Ireland note. The first is the functional form adopted. The second is the values assigned to its parameters. Obviously, both authors use data to discipline their choices. And so will we.

Our main contribution is to bring more data to the debate. We do so in two ways. First, we use the additional decade of data available since Ireland's work. This is a particularly abnormal and, at the same time, very interesting decade, since it was characterized by several observations with very low interest rates. Thus, it helps identify the behavior of money demand in that range, which, as we will discuss, is very important to identify the functional form.

Second, we also study evidence from developed countries whose inflation histories are similar to that of the United States. Although one could certainly entertain differences across countries, this evidence is also useful for identifying the functional form and the parameter values, as we show below. But, more importantly, the exploration of other countries highlights a third key feature that we bring to the analysis: the assumption regarding the true lower bound on the short term nominal interest rate. This is very relevant, since it determines the lower limit of the integral under the real money demand curve. Both Lucas and Ireland assumed the lower bound to be zero as did most of the monetary economics literature till 2010! However, the negative interest rates that have since been observed in the euro area, Sweden, and Switzerland have challenged that notion. Addressing this question will be at the heart of our analysis.

We find that for the United States, the cost of a steady state nominal interest rate of 5% is between 0.20% and 1.50% of lifetime consumption, depending on the assumptions regarding the lower bound and the functional form. The costs for the United Kingdom, Canada, and Japan are within the same range. Estimates are larger for the euro area, Sweden, and Switzerland, where they can go as high as 2% of lifetime consumption.

Modern analysis of optimal monetary policy is typically performed with New Keynesian models. These consider only moneyless economies, so they ignore the welfare effect of lack of money satiation that we focus on. Two factors, we believe, support this strategy. The first is the widespread belief that money is disappearing in modern economies. The second is the result in Ireland, which computes those costs to be negligible.

We challenge both notions. Regarding the first, we present overwhelming evidence that there is no sense in which modern economies are becoming moneyless.[1] Regarding the second, Ireland bases his computations on U.S. data only. This is problematic, since there was already at the time substantial evidence that the standard measure of M1, which had maintained a stable relationship with prices, interest rates, and nominal output during most of the twentieth century, became unstable in the early 1980s. Fully aware of that problem, Ireland makes a very reasonable adjustment by adding to M1 the retail sweeps that became very popular from 1994 onwards. He shows, however, that even after this adjustment, the behavior of real money demand after 1980 is different from the one that prevailed between 1900 and 1980.

Armed with this new monetary aggregate and a sample that starts in 1980, Ireland argues that the welfare cost of inflation is substantially lower than the one obtained by Lucas (2000),

[1] A more detailed analysis with yearly data that includes more countries can be found in Benati et al. (2021).

for two different reasons. First, he shows that the log-log functional form preferred by Lucas performs much worse than the semi-log specification. Second, for the semi-log specification, he estimates a semi-elasticity of the real money demand with respect to the nominal interest rate that is much lower than the one calibrated by Lucas.

We depart from Ireland and adopt the proposal in Lucas and Nicolini (2015), who argue that regulatory changes between 1982 and 1984 altered the availability of transactional assets in the United States. Specifically, they propose adding the Money Market Demand Accounts, which were created in 1984, to M1.[2] Once these new deposits are taken into account, a remarkably stable real money demand is obtained, which behaves in the same way before and after 1980. The estimates of the real money demand using the Lucas and Nicolini aggregate, that they label "New M1", imply larger estimates of the welfare cost of inflation than those obtained by Ireland. As before, the reason is twofold. First, even if we use the semi-log specification chosen by Ireland, the estimated elasticity is substantially larger. Second, the evidence against the functional form used by Lucas is not as clear-cut, especially if one allows for a negative lower bound on the short-term interest rate.

Besides finding — for obvious reasons — the argument in Lucas and Nicolini (2015) very compelling, we also report results for several other countries, for which there is no evidence of instability for the entire sample, using M1 as the monetary aggregate. Overall, the analysis for the other countries strongly supports the results for the United States when using the New M1 aggregate, in terms of both the estimated semi-elasticity for the functional form preferred by Ireland and the comparative weakness of the evidence against the functional specification preferred by Lucas.

Our estimates of the welfare cost of lack of money satiation suggest that in analyzing optimal monetary policy, ignoring money can be seriously misleading. For instance, Coibion et al. (2012) make a compelling argument against increasing the inflation target in countries like the United States. They do so in the context of a model with frictions in the setting

[2] The retail sweeps that Ireland adds to M1 are a relatively low fraction of the Money Market Demand Accounts.

of prices and recurrent, though not very frequent, episodes with the nominal interest rate at the zero lower bound. They compute the welfare effect of an interest rate of 5% in their preferred specification to be close to 0.6% percent of lifetime consumption. That number, which combines the cost created by price frictions and the probability of being at the zero lower bound, is well within the range of estimates we obtain for the United States. Relative to this number, the 0.04% estimated by Ireland does appear negligible. But 0.2%, the lowest number we estimate, is certainly not.[3] As it turns out, taking into account the effect that we study would actually reinforce the argument of their paper.

On the theory side, we innovate by constructing upper and lower bounds for the estimate of the cost. The area under the money demand curve is an almost exact measure of the welfare cost for a very general class of monetary models in the neighborhood of zero, as Alvarez et al. (2019) show. We extend their results for a quite general sub-class of the models they analyze and compute exact lower and upper bounds for the estimates of the costs, using the area under the money demand curve, for any value of the interest rate. As we show, the difference between the upper and the lower bound is extremely small for the range of interest rates that have ever been observed in the United States. We believe the formulas we derive would be useful in future work.

In our analysis, we follow the tradition of considering the most liquid monetary assets, which include cash and transactional deposits. We abstract from a detailed discussion of the demand for each of the components, an issue recently addressed by Kurlat (2019).[4]

The paper proceeds as follows. In Section 2, we discuss a family of monetary models for which we derive very tight lower and upper bounds for the welfare cost of inflation using the area under the real money demand curve. In Section 3, we plot the data and argue that in our view, there is very solid evidence in favor of stable money demand relationships

[3]Coibion et al. (2012) explicitly acknowledge that they do not take into account the costs derived from lack of money satiation.

[4]He shows that addressing these considerations in a model with imperfect competition substantially increases the estimates of the welfare cost relative to those of models that ignore the creation of inside money.

for the countries we analyze. Section 4 formally makes this statement by analyzing unit root and cointegration properties of the series. It also discusses the estimation results for three different empirical specifications used in the literature, including the ones Lucas and Ireland explored. Section 5 presents our computations for the welfare cost functions. Section 6 discusses stability tests and the potential existence of non-linearities at low interest rates, as suggested by Mulligan and Sala-i-Martin (2000). Section 7 concludes.

2 The Model

We study a labor-only economy with uncertainty, in which making transactions is costly.[5] The economy is inhabited by a unit mass of identical agents with preferences given by

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

where U is differentiable, increasing, and concave.

Every period, the representative agent chooses a number of portfolio transactions n_t that allow her to exchange interest-bearing illiquid assets for money, which is needed to buy the consumption good. The total cost of those transactions, measured in units of times, is given by a function $\theta(n_t, \nu_t)$, where ν_t is an exogenous stochastic process. This formulation generalizes the linear function assumed by Baumol (1952) and Tobin (1956).

The production technology for the consumption good is given by

$$y_t = c_t = z_t l_t,$$

where l_t is time devoted to the production of the final consumption good and z_t is an exogenous stochastic process.

The representative agent is endowed, each period, with a unit of time that is used to

[5] The baseline model is discussed at length in Benati et al. (2021).

produce goods and make transactions. Thus, equilibrium in the labor market implies that

$$1 = l_t + \theta(n_t, \nu_t),$$

and feasibility is given by

$$c_t = z_t(1 - \theta(n_t, \nu_t)).$$

It follows that the real wage is equal to z_t.

Purchases are subject to a cash in advance constraint

$$P_t c_t \leq n_t M_t, \tag{2}$$

where M_t is average money balances.

Contrary to standard practice, we allow for money to pay a nominal return, which we denominate r_t^m, that is allowed to become negative. This is required, as we explain below, for the model to be able to account for negative values of the short-term policy rate in equilibrium, as experienced in recent years by some of the countries we analyze in the paper.

At the beginning of each period, the agent starts with nominal wealth W_t, which can be allocated to money or interest-bearing bonds, B_t according to

$$M_t + B_t \leq W_t. \tag{3}$$

Nominal wealth at the beginning of next period, in state s_{t+1}, will then be given by

$$\begin{aligned} W_{t+1} &\leq M_t(1 + r_t^m) + B_t(1 + r_t^b) + T_t \\ &\quad + [1 - \theta(n_t, \nu_t)]\, z_t P_t - P_t c_t, \end{aligned} \tag{4}$$

where r_t^b is the return on government bonds and T_t is a transfer made by the monetary authority.

Notice that the unconstrained efficient outcome is to allocate all the labor input to the production of the consumption good so as to set $c_t = z_t$. Thus, a measure of the welfare cost of making transactions, as a fraction of consumption, is given by the value of $\theta(n_t, \nu_t)$ in equilibrium.

In the Online Appendix A, we show that as long as the cost function $\theta(n_t, \nu_t)$ is differentiable, an interior solution for n_t must satisfy

$$n_t^2 \frac{\theta_n(n_t, \nu_t)}{(1 - \theta(n_t, \nu_t))} = r_t^b - r_t^m. \tag{5}$$

We also show that as long as $r_t^b - r_t^m > 0$, the cash in advance is binding, so real money demand, as a proportion of output, is equal to the inverse of n_t. Note that equation (5) is independent of z_t. Thus, the theory implies a unit income elasticity of real money demand.

The solution for real money demand depends on the interest rate *differential* between bonds and money. In what follows, we let the interest rate differential between bonds and money be $r_t \equiv r_t^b - r_t^m$.

For the maximum problem of the agent to be well defined, it has to be the case that

$$r_t = r_t^b - r_t^m \geq 0, \tag{6a}$$

which is the well-known lower bound on the interest rates in bonds.[6] The popular zero-bound restriction on policy rates is obtained from (6a) plus the standard assumption in the literature that $r_t^m = 0$. The analyses of both Lucas (2000) and Ireland (2009) make this assumption.

The recent experience of prolonged negative short-term interest rates in several countries severely challenges this notion. As condition (6a) must hold in equilibrium, the challenge can be solved only by allowing for negative values of the own return on money— namely $r_t^m < 0$— at least when the short term interest rate r_t^b becomes small. To allow for that possibility, we

[6] Intuitively, were $r(s^t) - r^m(s^t)$ to be negative, the representative agent would have incentives to borrow from the government in unbounded quantities and hold money.

proceed as follows. As we identify our measure of money with M1 in the data, it is natural to think of the return on money as an average of the return of the two components of M1, cash and demand deposits. For cash, a negative return can be rationalized by the risk of being lost or stolen, as Alvarez and Lippi (2009) compute using survey data.[7] For deposits, we use a linear relation between their nominal return and the interest rate on bonds. Kurlat (2019) provides very strong empirical support for such a relationship.

These assumptions, taken together, are consistent with the return on money satisfying

$$r_t^m = -a + br_t^b, \tag{7}$$

for $a > 0$ and $b < 1$.[8] This linear relationship implies that r_t^m will be negative for small enough values of r_t^b. In addition, together with $(6a)$, it implies that

$$r_t^b + a - br_t^b \geq 0, \text{ or } r_t^b \geq -\frac{a}{(1-b)},$$

so the lower bound on the short-term rate is negative. In our welfare cost computations below, we consider two combinations of values for these parameters, in addition to the standard benchmark of $a = b = 0$.

The functional form of the real money demand function depends on the functional form of the transactions technology $\theta(n_t, \nu_t)$, and at this level of generality, the model is consistent with many different possibilities. In what follows, to clarify the main difference between Lucas (2000) and Ireland (2009), we consider three well-known functional forms that have been used in previous empirical work. All three exhibit a unit income elasticity, as implied by the model. The first specification is the log-log one,

$$\ln \frac{M_t}{P_t y_t} = a^1 - \eta \ln r_t + u_t^1, \tag{8}$$

[7] Alvarez and Lippi (2009) calibrate this return at -0.02 using survey data from Italy.
[8] Further details are provided in the Online Appendix B.

which exhibits a constant interest rate elasticity equal to η. Notice that as $i_t \to 0$, real money demand goes to infinity. It is this asymptote at zero that Lucas used to argue that the welfare cost of inflation is sizeable, even at low values for the interest rate. The other two formulations that we explore, the semi-log

$$\ln \frac{M_t}{P_t y_t} = a^2 - \gamma r_t + u_t^2, \tag{9}$$

which exhibits a constant semi-elasticity, γ, and the Selden-Latané

$$\frac{M_t}{P_t y_t} = \frac{1}{a^3 + \phi r_t + u_t^3}, \tag{10}$$

imply a finite level of the demand for real money balances when the interest rate differential becomes zero. This feature is emphasized by Ireland, who uses (9) in his revision of Lucas's estimate.

By exploiting recent data that include several years of very low (or even negative) interest rates for a few countries, we can provide a sharper comparison of the empirical performance of the three alternative functional forms.

As we show below, the welfare cost implications of the last two functional forms are similar. We do, however, choose to include the Selden-Latané specification together with the others, since it does have an overall better performance than the other two, as our econometric analysis shows.

In computing the welfare cost of inflation, we consider these three functional forms and three alternative assumptions regarding the parameters a and b, which relate the return on money and the short term interest rate as described in the linear relationship (7). Besides being based on empirical evidence, such a linear relationship has the advantage that the relevant opportunity cost r_t becomes

$$r_t = r_t^b - r_t^m = a + (1 - b) r_t^b,$$

which is a linear transformation of the observable short term interest rate r_t^b. As the last two functional forms are either a linear function of r_t or the inverse of a linear function of r_t, one needs only to test and estimate those two specifications under the benchmark case of $a = b = 0$, then adjust the estimates by the corresponding linear transformation. However, for the log-log specification, this is not the case, and both the cointegration tests and the estimates will depend on the specific assumption regarding the lower bound. As we show below, both are quite sensitive to the assumed lower bound, particularly so for the case of the United States.

In the next section, we show how to build tight upper and lower bounds for the welfare cost of inflation, using the area under the estimated real money demand function.

2.1 The welfare cost of inflation and the area under the money demand curve

In this section, we apply the techniques developed in Alvarez et al. (2019) to a class of models that are more restrictive than the ones they used. Specifically, we consider only representative agent models in which the cost of transforming liquid assets into illiquid ones, is given by the differentiable function $\theta(n_t, \nu_t)$ described above. For this restricted class of models, we obtain upper and lower bounds for the welfare cost of inflation that can be directly computed based on estimated money demand functions.

Alvarez et al. (2019) show that the area under the money demand curve approximates the welfare cost of inflation arbitrarily well as the opportunity cost of money (in our model, the interest rate differential r_t) approaches zero.[9] Our bounds can be used for any value of the interest rate.

As we show below, for the countries we consider, the distance between the upper and lower bound is positive, but extremely small— so much so that in most of the figures, the

[9]They also show in numerical examples that the approximation is remarkably accurate for a wide range of positive values of the opportunity cost.

difference between the two is invisible to the eye.

To make progress and simplify the notation, we eliminate the shock and the time dependence, and we write (5) as

$$n^2 \frac{\theta_n(n)}{(1-\theta(n))} = r, \tag{11}$$

where $r \geq 0$. As previously discussed, the welfare cost of inflation, measured as a fraction of consumption, is given by

$$\omega^W(r) = \theta(n(r)), \text{ where } \omega^W(0) = \theta(n(0)) = 0.$$

It follows that

$$\frac{\partial \omega^W(r)}{\partial r} = \frac{\partial \theta(n)}{\partial n}\frac{\partial n}{\partial r}(r) > 0. \tag{12}$$

We now show how the function $\omega^W(r)$ can be bounded above and below using the integral under the money demand curve.[10]

The area under the demand curve is equal to

$$\omega^D(r) = \int_0^r m(z)dz - m(r)r, \tag{13}$$

so

$$\frac{\partial \omega^D(r)}{\partial r} = -\frac{\partial m}{\partial r}(r)r > 0.$$

As real money demand $m(r)$ is the inverse of velocity, $n(r)$, it follows that

$$\frac{\partial n}{\partial r}(r) = -\frac{\partial m}{\partial r}(r)n^2,$$

which, using (11), becomes

$$\frac{\partial n}{\partial r}(r) = -\frac{\partial m}{\partial r}(r)r\frac{[1-\theta(n)]}{\frac{\partial \theta(n)}{\partial n}}.$$

[10]The analysis below closely follows the ideas in Alvarez et al. (2019).

If we use the definition in (12),

$$\frac{\partial \omega^W(r)}{\partial r} = -\frac{\partial m}{\partial r}(r) r \left[1 - \theta(n)\right] = \frac{\partial \omega^D(r)}{\partial r} \left[1 - \omega^W(r)\right].$$

Recall that $\omega^W(0) = \omega^D(0) = 0$. Thus, we can recover the welfare cost of inflation for an interest rate differential r_0 by integrating $\partial\omega^W/\partial r$ from zero to r_0, or

$$\int_0^{r_0} \frac{\partial \omega^W(z)}{\partial r} dz = \int_0^{r_0} \frac{\partial \omega^D(r)}{\partial r} \left[1 - \omega^W(z)\right] dz.$$

For all $z \in [0, r_0]$, however,

$$1 \geq \left[1 - \omega^W(z)\right] \geq \left[1 - \omega^W(r_0)\right].$$

Therefore,

$$\int_0^{r_0} \frac{\partial \omega^W(z)}{\partial r} dz \leq \int_0^{r_0} \frac{\partial \omega^D(r)}{\partial r} dz$$

and

$$\int_0^{r_0} \frac{\partial \omega^W(z)}{\partial r} dz \geq \left[1 - \omega^W(r_0)\right] \int_0^{r_0} \frac{\partial \omega^D(r)}{\partial r} dz,$$

which imply

$$\left[1 - \omega^W(r_0)\right] \omega^D(r_0) \leq \omega^W(r_0) \leq \omega^D(r_0).$$

We thus obtain the following bounds.

$$\frac{\omega^D(R)}{(1 + \omega^D(R))} \leq \omega^W(R) \leq \omega^D(R).$$

It is straightforward to see that the bounds are extremely tight. For example, for an opportunity cost equal to 3% of consumption, which is very large, the difference between the upper and the lower bound is equal to about one-tenth of a percentage point.

Explicit closed form solutions for the function $\omega^D(R)$ can be obtained for the three empirical specifications described in (8) to (10), as we show below.

3 A Look at the Raw Data

For the empirical analysis, we work with quarterly post-WWII data. The series and their sources are described in detail in Appendix A. For all but one country, we consider M1 as the relevant monetary aggregate.[11] The exception is the United States, for which we follow Lucas and Nicolini (2015) and use "New M1", which is obtained by adding Money Market Deposit Accounts (MMDAs) to the standard M1 aggregate produced by the Federal Reserve.[12]

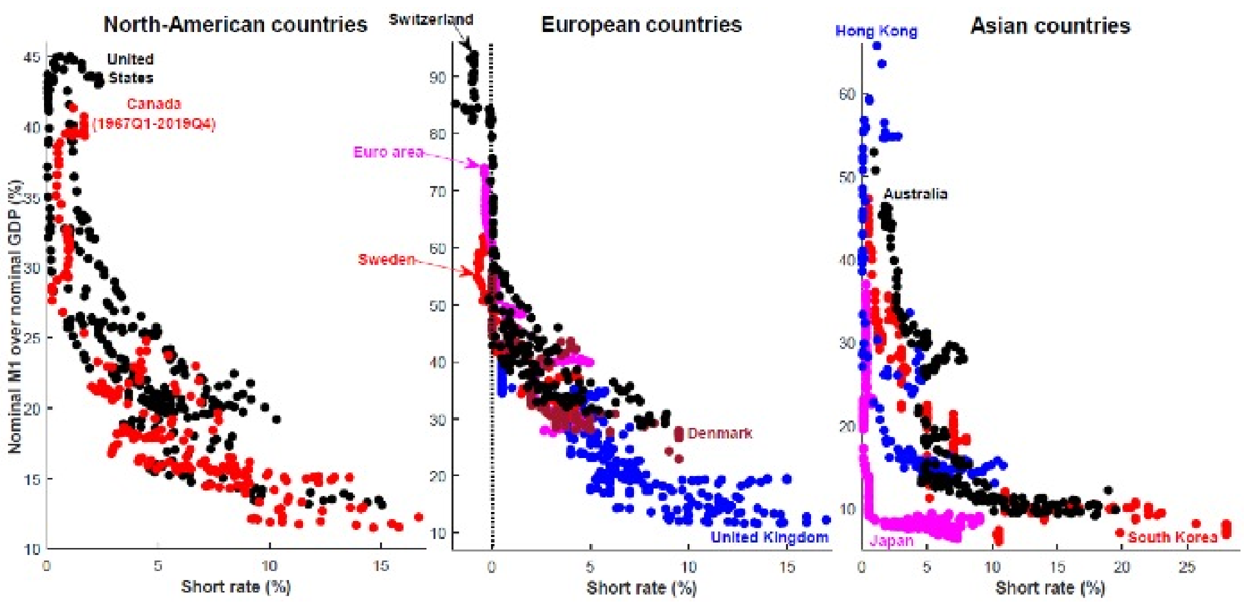

Figure 1: Scatterplots of nominal M1 over nominal GDP against the short rate

Figure 1 shows scatterplots of the ratio between nominal M1 and nominal GDP against a short-term nominal interest rate. We present three groups of countries organized by region. The three panels provide strong visual evidence of a negative relationship between the ratio of M1 to GDP and a short-term rate, a hallmark of the theory of real money demand. Comparing the three panels reveals several interesting features. The first is that there appear

[11]In Appendix C, we motivate our choice of working with "simple-sum" M1 aggregates, as opposed to their Divisia counterparts.

[12]Augmenting the standard M1 aggregate with MMDAs had originally been suggested by Goldfeld and Sichel (1990, pp. 314-315) in order to restore the stability of the long-run demand for M1.

to be clear and sizeable differences across (groups of) countries in terms of the level of the demand for real money balances. In particular, whereas the demand curves for the groups of North American and European countries exhibit a strong within-group similarity (this is especially apparent for the United States and Canada), those for the former group tend to be substantially lower than those for the latter one. This is especially clear at very low levels of the short rate. For our purposes this could be crucial, since it might affect the area under the demand curve. Asian countries exhibit even starker heterogeneity, with each individual country essentially having its own demand curve.[13] Finally, in three European countries (Switzerland, Sweden, and the euro area), short-term rates have consistently been negative over the most recent period, thus providing crucial and previously unavailable information about where the "true" lower bound for nominal short-term interest rates might lie.

4 Time-Series Properties of the Data

Figure 1 provides strong evidence, but it conceals the variables' behavior over time, thus failing to show how the persistent components of the two variables have co-moved along the sample. Figure 2 therefore shows the time series for M1 velocity and the short-term nominal rate in our sample. The data so displayed suggest that both series are I(1) and that they are cointegrated. As we now discuss, formal statistical tests strongly support this impression.

4.1 Evidence from unit root tests

Table A.1 in the Appendix reports results from Elliot et al.'s (1996) unit root tests for either the levels or the logarithms of M1 velocity and the short rate. In short, for nearly all countries and all series, the null hypothesis of a unit root cannot be rejected.[14] In searching

[13]Notice that since the short rate for Japan, Hong Kong, and, to a lesser extent South Korea has been at or around zero for a non-negligible portion of the sample, for these countries, the satiation level of real M1 balances is equal to the smallest level that has been observed with the short rate at zero. For example, for Japan, it is around 10%.

[14]For the short rate it can rejected only for Denmark (in levels) and Canada (1947Q3-2006Q4 in logarithms). For M1 velocity it can be rejected only for South Korea (in levels), whereas results for the euro

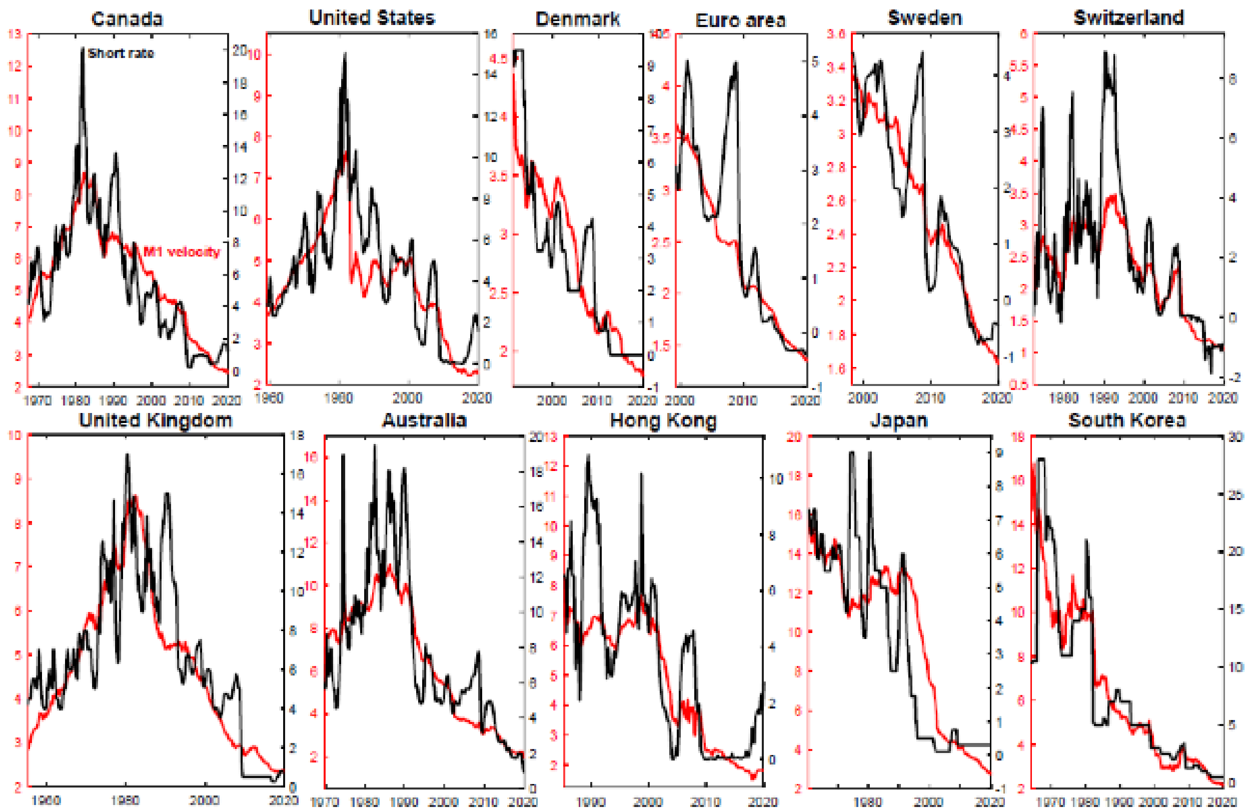

Figure 2: M1 velocity and the short rate

for a cointegration relationship between velocity and the short rate, we will therefore proceed as follows in the next section. First, taking the results from the unit root tests literally—that is, as indication that the series contain *exact* unit roots—we will test for cointegration based on Johansen's tests, which are predicated on the assumption that the series are indeed I(1). Since, however, a plausible alternative interpretation of the results in Table A.1 is that the series are *local-to-unity*—in which case, as shown by Elliot (1998), tests such as Johansen's tend to perform poorly—we will search for cointegration based on Wright's (2000) test, which is valid for both exact unit roots and roots that are local-to-unity. All of the technical details about the implementation of the tests are identical to those of Benati (2020) and Benati et

area (in levels) are ambiguous. In all of these cases, we will treat rejection of the null of a unit root as a fluke. There are two reasons for this. First, if the tests were perfectly sized (which, since we are here using Cavaliere et al.'s (2012) bootstrapping procedure, should be regarded as a good approximation), with 11 countries we should expect about one rejection for any of the four tests (two series, both either in levels or in logarithms). In fact, with three rejections, we obtain less than that. Second, visual inspection strongly suggests that the three series for which the null is rejected are in fact I(1).

al. (2021), which the reader is referred to.

4.2 Cointegration properties of the data

In studying a cointegration relationship between the demand for real money as a fraction of GDP and a short-term interest rate, we consider the functional forms (8), (9), and (10) for two alternative tests developed in the literature. In addition, for the log-log specification, we consider three different alternatives for the lower bound on the short-term interest rate.[15]

The first is the one in which $r^m = 0$, as is typically assumed in the literature, implying that $a = b = 0$. For the other two, we chose the values for a and b as follows. Kurlat (2019) very precisely estimates the slope parameter b to be 0.15, using microdata from the U.S. We adopt that value. For a, we explore two different values that allow us to accommodate the negative interest rate experiences of the countries in our sample. Specifically, we study the cases $a = 1$ and $a = 2$, which correspond to lower bounds of roughly -1.2% and -2.4%, respectively. The first lower bound can account for the observations on short-term rates in Denmark, the euro area, and Sweden, but it cannot account for those for Switzerland, for which the lowest value for the short-term interest rate was around -1.8%. The second lower bound can accommodate all cases.

The *true* lower bound on interest rates could be lower than the values we assumed: for decades it had been assumed that the lower bound on nominal interest rates was zero, and recent experiences have shown that this is not the case. For our purposes, a natural course of action is to consider the previously mentioned range of possibilities.

For any of the three money demand specifications discussed in Section 2, Table 1 reports bootstrapped p-values for Johansen's maximum eigenvalue test of 0 versus 1 cointegration

[15] As explained above, these different assumptions do not affect the cointegration tests for the semi-log and the Selden-Latané, since we assumed the interest rate differential to be linear in the short-term interest rate.

Table 1*a* Bootstrapped *p*-values[a] for Johansen's maximum eigenvalue[b] tests for (log) M1 velocity and (the log of) a short-term rate for *a*=*b*=0

		Money demand specification:		
Country	*Period*	*Selden-Latané*	*Semi-log*	*Log-log*
United States	1959Q1-2019Q4	0.0106	0.0302	0.4692
United Kingdom	1955Q1-2019Q4	0.0209	0.0507	0.5235
Canada	1947Q3-2006Q4	0.1758	0.0358	0.0730
	1967Q1-2019Q4	0.0388	0.4241	0.0035
Australia	1969Q3-2019Q4	0.0621	0.0367	0.3081
Switzerland	1972Q1-2019Q4	0.0166	0.0547	–[c]
Sweden	1998Q1-2019Q4	0.1142	0.1137	–[c]
Euro area	1999Q1-2019Q4	0.0877	0.1242	–[c]
Denmark	1991Q1-2019Q4	0.0501	0.1449	–[c]
South Korea	1964Q1-2019Q4	0.0000	0.0831	0.0399
Japan	1960Q1-2019Q4	0.0120	0.0066	0.0117
Hong Kong	1985Q1-2019Q4	0.0189	0.0189	0.0893

[a] Based on 10,000 bootstrap replications. [b] Null of 0 versus 1 cointegration vectors. [c] The last observations for the interest rate are either zero or negative.

vectors.[16] For Canada, we have two partially overlapping M1 series that cannot be linked, since they are slightly different. We present results based on each of them. Table 1*a* shows the results assuming a zero lower bound (i.e., $a = b = 0$), whereas Table 1*b* shows the corresponding results based on the other two assumptions for the log-log case, the only one for which the results are sensitive to the lower bound assumption. Table 2 reports the 90% confidence intervals for the second element of the normalized cointegration vector based on Wright's (2000) test. As was the case with Table 1, Table 2*a* shows the results for a zero lower bound, whereas Table 2*b* shows the results based on the other two assumptions for the log-log.

Based on Johansen's tests (Table 1) and the assumption of a zero lower bound, evidence of cointegration is strong based on the Selden-Latané specification, whereas it is slightly

[16]The corresponding results from the trace test are qualitatively the same, and they are available upon request.

Table 1*b* Bootstrapped *p*-values[a] for Johansen's maximum eigenvalue[b] tests for log M1 velocity and the log of a short-term rate				
		a=0	a=-1	a=-2
Country	*Period*	b=0	b=0.15	b=0.15
United States	1959Q1-2019Q4	0.4692	0.4397	0.0830
United Kingdom	1955Q1-2019Q4	0.5235	0.3566	0.2628
Canada	1947Q3-2006Q4	0.0730	0.0256	0.0230
	1967Q1-2019Q4	0.0035	0.0173	0.0377
Australia	1969Q3-2019Q4	0.3081	0.1733	0.1353
Switzerland	1972Q1-2019Q4	–[c]	–[c]	0.0509
Sweden	1998Q1-2019Q4	–[c]	0.2907	0.1189
Euro area	1999Q1-2019Q4	–[c]	0.1868	0.1455
Denmark	1991Q1-2019Q4	–[c]	0.5349	0.3988
South Korea	1964Q1-2019Q4	0.0399	0.0119	0.0097
Japan	1960Q1-2019Q4	0.0117	0.0056	0.0065
Hong Kong	1985Q1-2019Q4	0.0893	0.0492	0.0321

[a] Based on 10,000 bootstrap replications. [b] Null of 0 versus 1 cointegration vectors. [c] The last observations for the interest rate are either zero or negative.

weaker based on the semi-log, and it is materially weaker based on the log-log. In particular, based on Selden-Latané, cointegration is *not* detected at the 10% significance level *only* for the first period for Canada and almost marginally for Sweden (in this case, however, a likely explanation is that the sample period is quite short). Based on the semi-log, it is not detected for the euro area, Sweden, Denmark, and Canada's second sample.

For the log-log, cointegration is not detected for the United States, the United Kingdom, and Australia. The tests cannot be performed for the euro area, Sweden, Switzerland, and Denmark, since for these countries, the short rate has been either negative or exactly equal to zero for the most recent period. But these results are very sensitive to the assumption regarding the effective lower bound. In almost all cases, the p-values go down monotonically as the lower bound is reduced.[17] In a few cases, the p-values go down substantially, most notably in the U.S., for which the test does indeed detect cointegration at the 10% level,

[17]The only exceptions are Canada's second sample, for which the p-values increase, and Hong Kong, for which they do go down relative to the benchmark, but not monotonically. In both cases, however, the p-values are below 5% for all possible assumptions regarding the effective lower bound.

when the lower bound is -2.4%

Wright's (2000) tests (Table 2) detect cointegration based on the Selden-Latané specification for all countries except Canada's first sample and Denmark. Further, in all cases in which cointegration is detected, the upper bound of the 90% confidence interval is negative. Based on the semi-log, cointegration is detected for all countries except Canada's first sample. Finally, based on the log-log, cointegration is not detected for Canada's first sample for any value of the lower bound; for the euro area for $a = -1$ and $a = -2$; and for South Korea for $a = 0$ and $a = -1$.

4.3 Which specification do the data prefer?

Overall, the results from Johansen's and Wright's tests in Tables 1 and 2 suggest that the data tend to prefer the Selden-Latané specification to the semi-log or the log-log. In this sub-section, we perform a more systematic model comparison exercise. Since it is not possible to nest the three money demand specifications into a single encompassing one, we proceed as follows. We start from the comparison between the semi-log and the log-log. Intuitively, the comparison between (9) and (8) boils down to whether the dynamics of log M1 balances as a fraction of GDP (i.e., minus log velocity) are better explained by the level of the short rate or by its logarithm. For each country, we therefore regress $\ln(M_t/Y_t)$ on a constant, p lags of itself, and p lags of either the level of the short rate or its logarithm. A natural way of interpreting these regressions is the following. Under the assumption that cointegration is indeed there for all countries,[18] and based on either specification, both $Y_t^{SL} = [\ln(M_t/Y_t)\ R_t]'$ and $Y_t^{LL} = [\ln(M_t/Y_t)\ \ln(R_t)]'$ have a cointegrated VECM(p-1) representation, which maps into a restricted VAR(p) representation in levels (where the restrictions originate from the cointegration relationship). The equations we are estimating can therefore be thought of as the corresponding *unrestricted* form of the equations for $\ln(M_t/Y_t)$ in the VAR(p) representation in levels for either Y_t^{SL} or Y_t^{LL}. It

[18] If this assumption did not hold, the entire model comparison exercise would obviously be meaningless.

Table 2*a* Results from Wright's tests: 90% bootstrapped confidence interval for the second element of the normalized cointegration vector, based on systems for (log) M1 velocity and (the log of) a short-term rate[a] for *a*=*b*=0

		Money demand specification:		
Country	*Period*	*Selden-Latané*	*Semi-log*	*Log-log*
United States	1959Q1-2019Q4	[-0.5874 -0.3432]	[-0.1481 -0.0680]	[-0.3473 -0.0670]
United Kingdom	1955Q1-2019Q4	[-0.5323 -0.3441]	[-0.1150 -0.0790]	[-0.3931 -0.1969]
Canada	1947Q3-2006Q4	NCD	NCD	NCD
	1967Q1-2019Q4	[-0.4970 -0.3649]	[-0.1198 -0.0358]	[-0.4097 -0.2456]
Australia	1969Q3-2019Q4	[-0.9216 -0.7054]	[-0.1817 -0.0455]	[-1.2750 -0.9747]
Switzerland	1972Q1-2019Q4	[-0.4641 -0.2759]	[-0.2330 -0.1289]	–[b]
Sweden	1998Q1-2019Q4	[-0.3643 -0.3082]	[-0.1539 -0.1219]	–[b]
Euro area	1999Q1-2019Q4	[-0.6013 -0.3010]	[-0.2173 -0.1653]	–[b]
Denmark	1991Q1-2019Q4	NCD	[-0.1393 -0.0432]	–[b]
South Korea	1964Q1-2019Q4	[-0.5943 -0.5022]	[-0.1485 0.0276]	NCD
Japan	1960Q1-2019Q4	[-1.8658 -1.0569]	[-0.3175 -0.0333]	[-0.6108 -0.1223]
Hong Kong	1985Q1-2019Q4	[-1.0421 -0.5936]	[-0.2570 -0.1009]	[-0.4957 -0.0913]

[a] Based on 10,000 bootstrap replications. NCD = No cointegration detected.
[b] The last observations for the interest rate are either zero or negative.

Table 2*b* Results from Wright's tests: 90% bootstrapped confidence interval for the second element of the normalized cointegration vector, based on systems for log M1 velocity and the log of a short-term rate[a]

Country	*Period*	*a*=0, *b*=0	*a*=-1, *b*=0.15	*a*=-2, *b*=0.15
United States	1959Q1-2019Q4	[-0.3473 -0.0670]	[-0.6414 -0.1650]	[-0.7745 0.0784]
United Kingdom	1955Q1-2019Q4	[-0.3931 -0.1969]	[-0.5916 -0.2913]	[-1.6308 0.2791]
Canada	1947Q3-2006Q4	NCD	NCD	NCD
	1967Q1-2019Q4	[-0.4097 -0.2456]	[-0.5516 -0.4155]	[-0.7233 -0.2188]
Australia	1969Q3-2019Q4	[-1.2750 -0.9747]	[-1.8131 -1.1004]	[-1.6781 -1.3217]
Switzerland	1972Q1-2019Q4	–[b]	–[b]	[-2.0691 0.6696]
Sweden	1998Q1-2019Q4	–[b]	[-0.3191 -0.2631]	[-0.5678 -0.3516]
Euro area	1999Q1-2019Q4	–[b]	NCD	NCD
Denmark	1991Q1-2019Q4	–[b]	[-0.4595 -0.1952]	[-0.6246 0.3884]
South Korea	1964Q1-2019Q4	NCD	NCD	[-0.9116 -0.8115]
Japan	1960Q1-2019Q4	[-0.6108 -0.1223]	[-1.1097 -0.1487]	[-1.2568 0.1006]
Hong Kong	1985Q1-2019Q4	[-0.4957 -0.0913]	[-1.0071 -0.0942]	[-1.9103 -1.0414]

[a] Based on 10,000 bootstrap replications. NCD = No cointegration detected.
[b] The last observations for the interest rate are either zero or negative.

is important to stress that the two specifications we are estimating are in fact nested. The easiest way of seeing this is to think of them as two polar cases—corresponding to either $\theta = 1$ or $\theta = 0$—in the following representation based on the Box-Cox transformation of R_t:

$$\ln\left(\frac{M_t}{Y_t}\right) = \alpha + \sum_{j=1}^{p} \beta_j \ln\left(\frac{M_{t-j}}{Y_{t-j}}\right) + \sum_{j=1}^{p} \delta_j \left(\frac{R_{t-j}^{\theta} - 1}{\theta}\right) + \varepsilon_t. \quad (14)$$

We estimate (14) via maximum likelihood, stochastically mapping the likelihood surface via Random-Walk Metropolis (RWM). The *only* difference between the "standard" RWM algorithm, which is routinely used for Bayesian estimation, and what we are doing here is that the jump to the new position in the Markov chain is accepted or rejected based on a rule that does not involve any Bayesian priors, as it uniquely involves the likelihood of the data.[19] So one way of thinking of this is as Bayesian estimation via RWM with completely uninformative priors, so that the log-posterior collapses to the log-likelihood of the data. All the other estimation details are identical to those of Benati (2008), to which the reader is referred.

Table 3a reports, for either specification and for $p \in \{2, 4, 8\}$, the mode of the log-likelihood. The main result in the table is that whereas the semi-log appears as the preferred functional form for the U.S., the U.K., Canada, and Hong Kong, the log-log produces a larger value of the likelihood for Australia, South Korea, and Japan, so that neither of the two specifications clearly dominates the other one.[20]

[19] So, to be clear, the proposal draw for the parameter vector β, $\tilde{\beta}$, is accepted with probability min[1, $r(\beta_{s-1}, \tilde{\beta} \mid Y, X)$] and rejected otherwise, where β_{s-1} is the current position in the Markov chain and

$$r(\beta_{s-1}, \tilde{\beta} \mid Y, X) = \frac{L(\tilde{\beta} \mid Y, X)}{L(\beta_{s-1} \mid Y, X)},$$

which uniquely involves the likelihood. With Bayesian priors, it would be

$$r(\beta_{s-1}, \tilde{\beta} \mid Y, X) = \frac{L(\tilde{\beta} \mid Y, X) P(\tilde{\beta})}{L(\beta_{s-1} \mid Y, X) P(\beta_{s-1})},$$

where $P(\cdot)$ would encode the priors about β.

[20] This crucially hinges on the fact that we are here focusing exclusively on low-inflation countries. As shown by Benati et al. (2021) and Benati (2021), for high-inflation countries, and especially hyperinflationary episodes, the data's preference for the log-log is overwhelming.

Table 3*a* Model comparison exercise, semi-log versus log-log: Mode of the log-likelihood in regressions of log velocity on lags of itself and either the short rate or its logarithm							
		$p = 2$		$p = 4$		$p = 8$	
Country	*Period*	*Semi-log*	*Log log*	*Semi-log*	*Log-log*	*Semi-log*	*Log-log*
United States	1959Q1-2019Q4	766.1394	756.6280	763.2818	751.3266	765.1439	740.3543
United Kingdom	1955Q1-2019Q4	879.6821	877.9350	898.6224	893.7504	892.1970	887.1920
Canada	1947Q3-2006Q4	820.2401	807.8379	813.8001	804.9218	802.7001	794.7403
	1967Q1-2019Q4	775.0890	767.0845	775.9595	766.4531	771.9264	766.1943
Australia	1969Q3-2019Q4	650.7331	656.0624	649.9510	655.1057	642.6046	650.3903
South Korea	1964Q1-2019Q4	630.9515	633.8825	628.2222	634.6372	623.8333	628.0991
Japan	1960Q1-2019Q4	845.3632	850.7677	841.5156	848.6520	832.2577	840.2434
Hong Kong	1985Q1-2019Q4	328.0148	325.5701	326.1339	324.9236	319.8478	325.2641
For Switzerland, Sweden, Euro area, and Denmark, there is no comparison, because the last observations for the short rate are negative.							

As we showed above, both the cointegration tests and the estimation results for the log-log are very sensitive to the assumption of the lower bound. Thus, we repeated the test for the other two assumptions about the lower bound and also for the same three values for p. Table 3*b* presents the results for $p = 8$ (results for $p = 4$ and $p = 2$ are very similar, and they are reported in Tables A.2*a* and A.2*b* in the Online Appendix).

The first feature to highlight is that for the four countries with negative rates, for which we could not do the test before, the semi-log dominates the log-log. The second is that in line with the previous analysis, the likelihood of the log-log specification increases when the assumed lower bound is lower for most countries for which the semi-log is the preferred specification. This is particularly true for those countries for which the cointegration tests also improve substantially for lower values of the bound, like the United States or the United Kingdom. However, with the single exception of Canada, the increase is not enough for the log-log to dominate the semi-log.

Table 3b Model comparison exercise, semi-log versus log-log: Mode of the log-likelihood in regressions of log velocity on lags of itself and either the short rate or its logarithm (p=8)

			Log-log		
		Semi-	a=0	a=-1	a=-2
Country	*Period*	*log*	b=0	b=0.15	b=0.15
United States	1959Q1-2019Q4	765.1439	740.3543	749.5835	751.0460
United Kingdom	1955Q1-2019Q4	892.1970	887.1920	890.2113	891.0367
Canada	1947Q3-2006Q4	802.7001	794.7403	822.8528	823.7403
	1967Q1-2019Q4	771.9264	766.1943	770.6245	771.9067
Australia	1969Q3-2019Q4	642.6046	650.3903	649.4391	648.9755
Switzerland	1972Q1-2019Q4	567.8522	–[a]	–[a]	559.0320
Sweden	1998Q1-2019Q4	300.9987	–[a]	300.5980	300.5340
Euro area	1999Q1-2019Q4	316.0427	–[a]	315.9802	315.6161
Denmark	1991Q1-2019Q4	402.0518	–[a]	399.5286	400.2900
South Korea	1964Q1-2019Q4	623.8333	628.0991	629.3606	628.8658
Japan	1960Q1-2019Q4	832.2577	840.2434	838.1153	836.0811
Hong Kong	1985Q1-2019Q4	319.8478	325.2641	327.9993	327.0688

[a] The last observations for the short rate are negative.

Turning to the comparison between the semi-log and the Selden-Latané, we adopt the same logic as before, but this time we "flip" the specifications for velocity on their head by regressing the interest rate on lags of itself and of either the level or the logarithm of velocity. Once again, these two regressions can be thought of as particular cases of the nested regression

$$R_t = \alpha + \sum_{j=1}^{p} \varphi_j R_{t-j} + \sum_{j=1}^{p} \xi_j \left[\frac{\left(\frac{Y_{t-j}}{M_{t-j}}\right)^{\theta} - 1}{\theta} \right] + \varepsilon_t, \qquad (15)$$

with either $\theta = 1$ (corresponding to Selden-Latané) or $\theta = 0$ (corresponding to the semi-log).

At first sight, this approach might appear questionable: Since we are dealing with the demand for real M1 balances for a given level of the short-term nominal interest rate, why would it make sense to regress the short rate on M1 velocity? In fact, this approach is perfectly legitimate for the following reason. As shown by Benati (2020), M1 velocity is, to a first approximation (and up to a scale factor), the permanent component of the short-term rate,[21] so that if we focus, for example, on the Selden-Latané specification, $V_t = a^3 + \phi R_t^P$, where V_t is velocity; $a, b > 0$ are coefficients; and R_t^P is the unit-root component of the short rate (R_t), with $R_t = R_t^P + R_t^T$ and R_t^T being the transitory component.[22] Figure 2 shows this quite clearly for Australia, Canada, the euro area, Hong Kong, Sweden, Switzerland, and the U.K. Regressing R_t on V_t therefore amounts to regressing the short rate on its (rescaled) stochastic trend— that is, the dominant driver of its long-horizon variation— and it is therefore conceptually akin to (e.g.) regressing GDP on consumption.[23]

The results are reported in Table 3*c*. The evidence is much sharper than that for the previous comparison between the semi-log and the log-log: in particular, for p equal to 4 or 8, the Selden-Latané specification is preferred to the semi-log for all countries except the United Kingdom and the euro area.

Summing up, whereas the Selden-Latané functional form appears to be quite clearly preferred to the semi-log, the log-log and the semi-log seem to be, from an empirical standpoint, on a roughly equal footing.

These results are in line with the evidence in Figure 3, in which we compare the log-log with the Selden-Latané for the United States, the United Kingdom, and Japan. In all cases, we plot both velocity and the short-term interest rate. In the top panel, we plot both variables in levels, corresponding to the Selden-Latané specification. In the other two panels, we plot the series in logs, corresponding to the log-log specification. The middle panel shows the case of the zero lower bound, and the bottom panel the case in which the lower bound is -2.4%. As can be seen, the relationship is evident with the Selden-Latané specification in

[21] This expresses in the language of time-series analysis Lucas's (1988) point that real M1 balances are very smooth compared with the short rate.

[22] A simple rationalization of this fact is provided by a "preferred habitat" model (see Modigliani and Sutch, 1966, and Vayanos and Vila, 2021) in which 'long' investors such as pension funds play an important role in money demand. The intuition is that whereas permanent shocks to the short rate shift the entire term structure of interest rates, and therefore affect the demand for M1 coming from *all* investors, transitory shocks impact only the short end of the yield curve and therefore have a much smaller (and, in the limit, negligible) effect.

[23] See Cochrane (1994) on consumption being the permanent component of GDP.

Table 3c Model comparison exercise, Selden-Latané versus semi-log: Mode of the log-likelihood in regressions of the short rate on lags of itself and either velocity or its logarithm

		$p = 2$		$p = 4$		$p = 8$	
Country	*Period*	*Selden-Latané*	*Semi-log*	*Selden-Latané*	*Semi-log*	*Selden-Latané*	*Semi-log*
United States	1959Q1-2019Q4	-22.9102	-24.0809	-5.8335	-7.3440	12.9347	10.3522
United Kingdom	1955Q1-2019Q4	-85.7350	-84.1422	-85.4391	-83.9044	-83.6446	-82.1970
Canada	1947Q3-2006Q4	-72.0532	-71.7812	-64.4770	-66.2576	-62.2194	-64.2760
	1967Q1-2019Q4	-65.0057	-65.9778	-56.1253	-59.0260	-50.7916	-53.9112
Australia	1969Q3-2019Q4	-136.4591	-137.1389	-132.5116	-133.5144	-116.7487	-118.2407
Switzerland	1972Q1-2019Q4	-45.6989	-45.8396	-39.9744	-40.8984	-20.5636	-22.4888
Sweden	1998Q1-2019Q4	65.5876	65.4372	66.9821	66.9083	70.5126	68.9721
Euro area	1999Q1-2019Q4	63.8008	64.3157	64.5967	65.3372	74.7778	75.5777
Denmark	1991Q1-2019Q4	50.9544	50.7969	60.7088	60.1085	65.6600	64.4409
South Korea	1964Q1-2019Q4	-131.5950	-135.8924	-118.2770	-131.1253	-86.1317	-93.5032
Japan	1960Q1-2019Q4	-141.5147	-141.6026	-140.6631	-140.7865	-129.5219	-130.1270
Hong Kong	1985Q1-2019Q4	-65.6601	-65.7389	-60.9537	-61.4880	-50.8999	-51.6665

For Switzerland, Sweden, the euro area, and Denmark there is no comparison, because the last observations for the short rate are negative.

the top-panel and quite blurred for the log-log case with a zero bound in the middle panel. However, with the assumption of a -2.4% zero bound, the log-log bottom panel specification is as evident as the Selden-Latané specification.

We draw three main conclusions from the evidence so far. First, in line with the evidence in Figures 1 and 2, the data provide substantial support to the existence of a stable long-run demand for M1, as predicted by the theory. Second, the Selden-Latané specification appears to exhibit the best overall performance among the three. Third, the log-log specification substantially improves its performance when the lower bound on the short-term interest rate is assumed to be lower.

Based on these findings, we choose to use the Selden-Latané as our benchmark functional form. But we will also provide estimates for the other two specifications. As Lucas argued, the computed costs are substantially higher using the log-log specification, particularly when the lower bound is assumed to be lower than zero. These computations raise some caution for our benchmark estimates, since the log-log specification cannot always be clearly ruled out, particularly when the lower bound is assumed to be low. Tables with results for all the

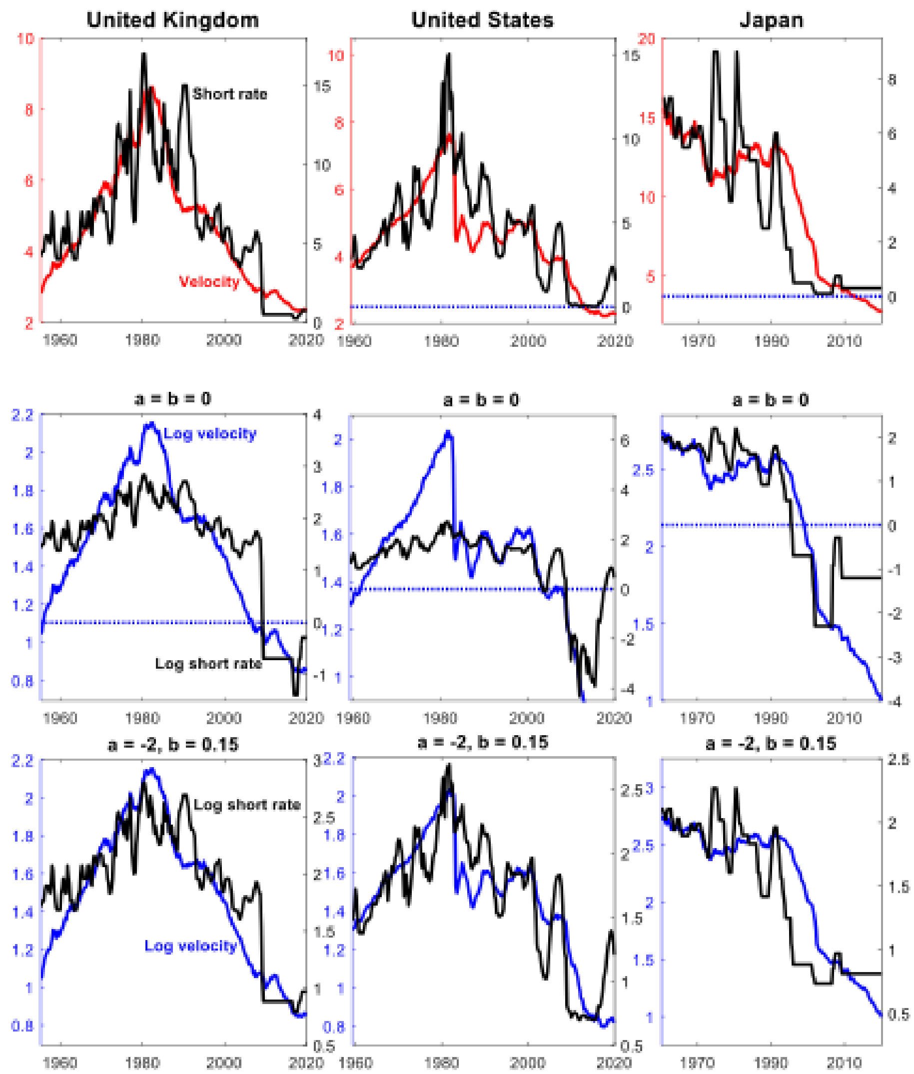

Figure 3: Comparing the Selden-Latané and log-log specifications

countries, the three functional forms and the three assumptions regarding the lower bound can be found in the Online Appendix.

4.4 Parameter estimates

Johansen's and Wright's tests provide strong evidence of a cointegrating relationship between real money holdings as a fraction of output and the short interest rate. However, neither of them directly provides point estimates for the parameters of the real money demand function.

Table 4*a* Point estimate and 90%-coverage bootstrapped[a] confidence interval for the coefficient on (the logarithm of) the short rate based on Stock and Watson's (1993) estimator for *a*=*b*=0

		Money demand specification:		
Country	*Period*	*Selden-Latané*	*Semi-log*	*Log-log*
United States	1959Q1-2019Q4	35.3 [26.3 41.8]	8.2 [5.7 10.0]	0.165 [0.087 0.235]
United Kingdom	1955Q1-2019Q4	36.5 [21.2 46.0]	7.9 [4.9 10.3]	0.284 [0.155 0.404]
Canada	1947Q3-2006Q4	39.5 [25.9 49.2]	7.9 [5.2 10.2]	0.373 [0.236 0.468]
	1967Q1-2019Q4	36.3 [23.5 45.0]	7.1 [4.1 9.3]	0.305 [0.200 0.382]
Australia	1969Q3-2019Q4	56.5 [33.5 70.2]	9.8 [6.1 12.4]	0.749 [0.518 0.892]
Switzerland	1972Q1-2019Q4	24.7 [18.0 29.5]	11.7 [7.8 14.4]	–[b]
Sweden	1998Q1-2019Q4	27.7 [22.7 32.2]	11.5 [9.5 13.4]	–[b]
Euro area	1999Q1-2019Q4	33.4 [27.0 40.2]	14.7 [12.2 17.6]	–[b]
Denmark	1991Q1-2019Q4	19.9 [13.9 26.7]	7.3 [4.3 10.0]	–[b]
South Korea	1964Q1-2019Q4	43.1 [35.3 46.9]	6.7 [4.8 7.9]	0.477 [0.401 0.539]
Japan	1960Q1-2019Q4	123.4 [71.3 155.5]	15.7 [10.5 20.5]	0.328 [0.172 0.440]
Hong Kong	1985Q1-2019Q4	60.4 [37.9 79.1]	14.6 [8.5 20.0]	0.171 [0.096 0.241]

[a] Based on 10,000 bootstrap replications. [b] The last observations for the interest rate are either zero or negative.

For the Johansen test, the corresponding money demand equation is estimated in its VECM form, from which the money demand parameters can be indirectly obtained. Wright's test, on the other hand, does not produce point estimates, but rather confidence intervals at the $x\%$ level for the parameters. To obtain point estimates for the welfare cost of inflation, we therefore choose to directly estimate the money demand equations using Stock and Watson's dynamic OLS procedure, which delivers point estimates for the parameters.

Table 4 shows the point estimates, as well as 90% confidence intervals, for the coefficients ϕ for the Selden-Latané specification, γ for the semi-log, and η for the log-log. Table 4a presents the results for the case of the zero lower bound. It is worth pointing out that the estimated value for the semi-elasticity parameter for the U.S., 8.2, is much closer to 7, the preferred value of Lucas (2000), than to 1.9, the one estimated by Ireland (2009). In fact the lower bound of the 90% confidence interval is 5.7. What explains this difference is the fact that Ireland used a different monetary aggregate, as we noted above.

Table 4*b* Point estimate and 90%-coverage bootstrapped[a] confidence interval for the coefficient on the logarithm of the short rate based on Stock and Watson's (1993) estimator

Country	*Period*	a=0, b=0	a=-1, b=0.15	a=-2, b=0.15
United States	1959Q1-2019Q4	0.165 [0.087 0.235]	0.406 [0.255 0.531]	0.543 [0.363 0.676]
United Kingdom	1955Q1-2019Q4	0.284 [0.155 0.404]	0.468 [0.259 0.630]	0.606 [0.355 0.801]
Canada	1947Q3-2006Q4	0.373 [0.236 0.468]	0.544 [0.357 0.676]	0.672 [0.439 0.838]
	1967Q1-2019Q4	0.305 [0.200 0.382]	0.467 [0.295 0.561]	0.585 [0.368 0.710]
Australia	1969Q3-2019Q4	0.749 [0.518 0.892]	0.916 [0.640 1.083]	1.064 [0.744 1.256]
Switzerland	1972Q1-2019Q4	–[b]	–[b]	0.552 [0.388 0.635]
Sweden	1998Q1-2019Q4	–[b]	0.250 [0.212 0.291]	0.429 [0.368 0.490]
Euro area	1999Q1-2019Q4	–[b]	0.398 [0.341 0.465]	0.610 [0.516 0.716]
Denmark	1991Q1-2019Q4	–[b]	0.298 [0.183 0.396]	0.417 [0.259 0.552]
South Korea	1964Q1-2019Q4	0.477 [0.401 0.539]	0.655 [0.565 0.722]	0.785 [0.674 0.863]
Japan	1960Q1-2019Q4	0.328 [0.172 0.440]	0.646 [0.281 0.917]	0.871 [0.419 1.208]
Hong Kong	1985Q1-2019Q4	0.171 [0.096 0.241]	0.587 [0.363 0.824]	0.816 [0.517 1.140]

[a] Based on 10,000 bootstrap replications. [b] The last observations for the interest rate are either zero or negative.

We did adopt the proposal in Lucas and Nicolini (2015), but that is a debatable choice. But we find it comforting that the estimate we obtain is very similar to the ones obtained for similar countries, like the U.K., Canada, and Australia, countries for which we use the measure of M1 reported by their corresponding central banks. Overall, there are point estimates close to 15 for Japan, Hong Kong, and the euro area, while the lowest point estimates, close to 7, are for Denmark and South Korea. Even for these two countries, the lowest bounds of the 90% confidence interval are above 4. All these estimates are much higher than the one used by Ireland (2009).

Table 4*b* presents the results for the log-log case, under the three assumptions regarding the zero bound. It is evident that the point estimates are very sensitive to the assumption regarding the true lower bound. This is consistent with the sensitivity of the cointegration tests, as discussed before, and is explained by the asymptote the log-log function implies. Notice that for the case of the lower bound of 2.4 percent, the estimate of the elasticity for the U.S., 0.54, is remarkably close to the value of 5 used by Lucas (2000). Overall, it is interesting that the estimates of the elasticity are quite close to 0.5, the value implied by the simplest version of the Baumol-Tobin model, when the lower bound is assumed to be 2.4.

5 The Estimated Welfare Cost Functions

The theoretical analysis implies that the parameters of the demand for real money balances and the lower bound we impose upon the short-term interest rate are the only relevant features to compute the welfare costs of inflation in any given country. In order to see this, it is useful to compute the integral under the money demand curve, as defined in (11), for the three specifications. The integrals are given by

$$\omega_{\log-\log}(r) = a^1 \frac{\eta}{1-\eta} r^{1-\eta}, \tag{16}$$

$$\omega_{semi-\log}(r) = \frac{a^2}{\gamma}\left(1 - \frac{1+\gamma r}{e^{\gamma r}}\right), \tag{17}$$

and

$$\omega_{Sel-Lat}(r) = \frac{1}{\phi} \ln\left(\frac{a^3+\phi r}{a^3}\right) - \frac{r}{a^3+\phi r}, \tag{18}$$

respectively, for the log-log, the semi-log and the Selden-Latané. As is apparent, each expression features a slope parameter and a level parameter. These two parameters, together with the assumption regarding the own return on money, fully summarize all of the information that is required for the computation of the welfare costs of inflation.

In what follows, we discuss in detail our results using the Stock and Watson estimates. We leave for the appendix the analysis with Wright's tests, in which we pick as "point estimates" the value that is most difficult to reject at the 10% level. Results in this case are very similar, except for a few cases in which the estimated values for the welfare cost are higher.[24] The methodology we use in order to compute the welfare costs of inflation follows

[24]These cases are Japan and the U.S. for the semi-log, the euro area for the Selden-Latané and Switzerland

Luetkepohl (1993, pp. 370-371). We first estimate via OLS the cointegrating regression corresponding to any of the three specifications—that is, to (8), (9), or (10). This gives us the point estimates of the parameters we need in order to compute the point estimates of the welfare cost functions. We then estimate the relevant VECM via OLS by imposing in estimation the previously estimated cointegration vector, and we characterize uncertainty about the point estimates of the welfare cost function by bootstrapping the VECM as in Cavaliere et al. (2012).

In line with the previous discussion in Section 4.1, this procedure is valid if the series contain *exact* unit roots. Under the alternative possible interpretation of the results from unit root tests—that is, that the series are *local-to-unity*— we proceed as in Benati et al. (2021, Section 4.2.1). Specifically, we compute, based on the just-mentioned VECM, the corresponding VAR in levels, which by construction features one, and only one, exact unit root. We turn it into its corresponding near unit root VAR by shrinking the unit root to λ=1-0.5(1/T), where T is the sample length.[25] 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. The two bootstrapping procedures deliver near-identical results, and in what follows, we will therefore report and discuss only those based on bootstrapping the VECM.

We start by focusing on the United States, Canada, the United Kingdom, and Japan, which in Figure 1 exhibit, for each level of the short rate, M1 balances as a fraction of GDP that are comparatively smaller than those of the remaining countries. Since for these countries the short rate has consistently been positive over the entire sample period, we first consider the case in which the own return of money is zero.

The results for these countries based on the Selden-Latané functional form (which, as discussed, we take as our benchmark) are reported in Figure 4.a. The point estimates of the upper and lower bounds are depicted as continuous black lines: as we previously anticipated,

for the log-log.

[25]For details see Benati et al.'s (2021) footnote 24.

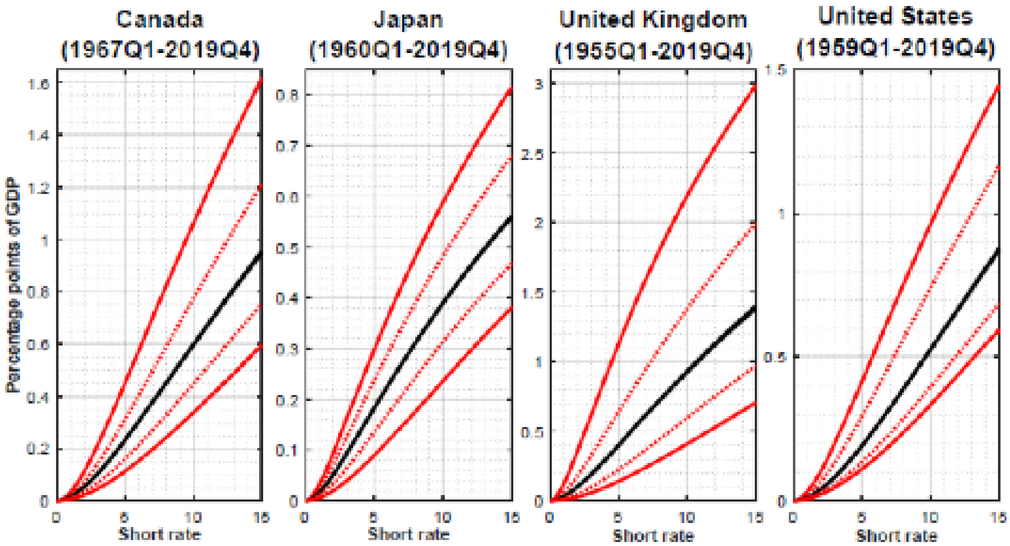

Figure 4.a: Estimated welfare cost functions based on the Selden-Latané specification with $a = b = 0$: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

in all cases the two lines are virtually indistinguishable, thus implying that the two bounds provide a very precise characterization of the welfare costs (the same holds for nearly all countries and all functional forms). The dotted and continuous red lines depict the 5th and 16th percentiles of the lower bounds and the 84th and 95th percentiles of the upper bounds of the bootstrapped distributions.

The point estimates for the welfare cost of a steady state interest rate equal to 5% are close to 0.2% of consumption for the United States and Japan and about 0.3% for Canada, while the estimate for the United Kingdom is about 0.4% of consumption. The estimate for the U.S. is the same as the one reported by Lucas (2000) when using the semi-log specification (9) and almost one order of magnitude above Ireland's (2009) estimate of 0.037%. As mentioned above, the main reason for the discrepancy is that our estimate of the semi-elasticity is substantially larger.

In Figure 4.b, we report the same point estimates using the Selden-Latané specification we report on Figure 4.a, together with the computations for the other two functional forms. The figure highlights the theoretical point made by Lucas (2000) that the log-log

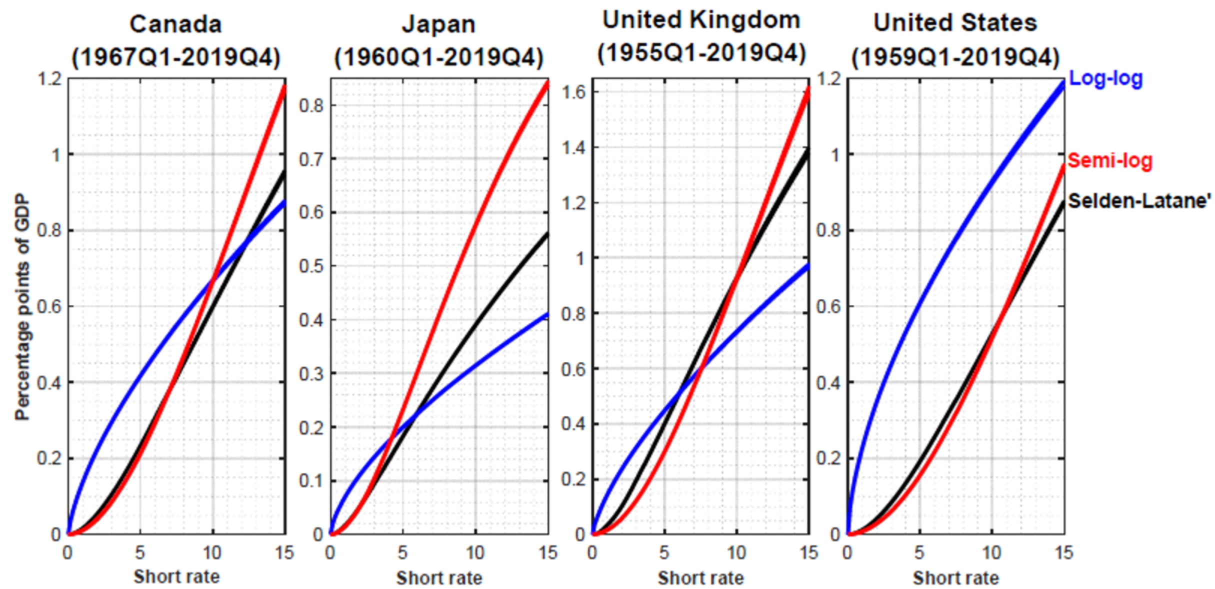

Figure 4.b: Point estimates of the welfare cost functions with $a = b = 0$ produced by alternative specifications

specification delivers substantially higher costs: about 0.6% compared with the 0.2% of the Selden-Latané.[26] However, the difference between the log-log and the other two specifications is much higher for the U.S. than for Japan and the U.K. and, to a lesser extent, Canada. This finding is explained by the fact that the point at which the cost implied by the log-log specification crosses the other two is quite sensitive to the estimates of particular countries. For Japan and the United Kingdom, they cross around 5% interest rate, for Canada at about 10%, and for the United States they cross at almost 20% (not depicted).

To explore for the possibility of a negative lower bound, in Figure 4.c, we report the results using the three specifications for the real money demand assuming a lower bound on short term interest rate of -2.4%, which corresponds to setting $b = 0.15$ and $a = 2$. We report the results only for the United States. We compare the results with the ones in Figure 4.b, which correspond the the case of a zero lower bound. As can be seen, the welfare costs of a 5% interest rate are twice as large as in the case of a zero lower bound: 0.4% of

[26] Our estimate of 0.6% of consumption is lower than the 1.2% reported by Lucas (2000). The difference lies in the fact that our estimate of the elasticity when the lower bound is zero is around 0.15. The elasticity used by Lucas is 0.5. We obtain the value of 0.5 when the lower bound is assumed to be -2.4%, a case we report below.

consumption for the Selden-Latané and 0.3% of consumption for the semi-log. However, for the log-log case, the increase is larger: it goes from 0.6% to 1.5% of consumption.

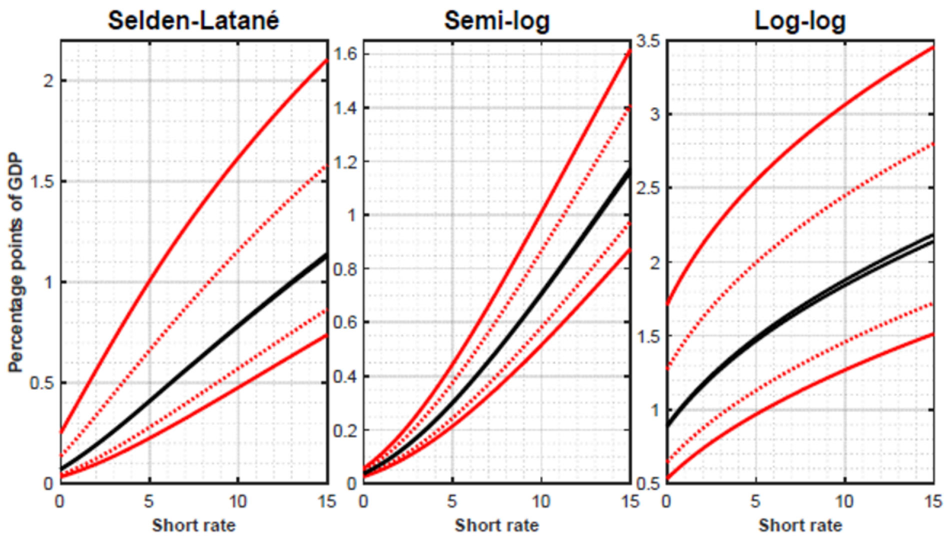

Figure 4.c: Estimated welfare cost functions for the United States (1959Q1-2019Q4) with $a = -2\%$ and $b = 0.15$: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

For our second set of results, we report the welfare cost computations for Switzerland, Sweden, and the euro area. There are two differences between these countries and the ones discussed above. The first is that they have, on average, higher money balances over output. The second is that they experienced negative short-term interest rates. As discussed above, this feature is consistent only with the notion that the own return on money is not zero and can become negative when the short-term interest rate becomes negative. In spite of that, in computing the cost of inflation, we consider three scenarios. The first is the benchmark case of a zero own return on money. We do this in order to compare the results with the ones reported in Figure 4.a. Any difference in the costs ought to be driven only by the different estimated parameters. For the other two scenarios, we follow the strategy adopted above.

Figure 5 presents the results for the estimated welfare cost functions, with one- and two-standard deviations bootstrapped confidence bands. In Figure 5.a, we report the point estimates based on the Selden-Latané. If we assume the lower bound to be zero, the welfare costs of a 5% interest rate are about 0.5% of consumption for the euro area and Switzerland and a bit smaller for Sweden, wich are more than double the ones for the United States, Canada, and Japan. These results are shown in the top panel. They are explained purely by differences in estimated parameters. In considering a lower bound that can accommodate the experiences of the euro area and Sweden, the cost increases to 0.75% of consumption for the euro area and Switzerland, and it is somewhat smaller for Sweden, as before. In the final case we consider, which is required to accommodate the experience of Switzerland, the cost can be almost 1.4% of consumption for the euro area, 1.2% for Switzerland and almost 1% for Sweden.

In Figure 5.b, we report a comparison for the point estimates for the three specifications (only two when the implied r is negative) for the three alternative assumptions regarding the effective lower bound. The message from this figure is similar to the one in Figure 4.b. First, the semi-log and the Selden-Latané provide very similar results for low interest rates. Second, the log-log implies much higher welfare costs at very low rates, with the exception of Sweden for the case in which the lower bound is -1.2%. However, the numbers are substantially larger than those in Figure 5.a, implying a larger interaction of the functional form with the assumption about the true lower bound on the short-term interest rates. For example, when the lower bound is assumed to be -2.4%, the log-log specification implies very large welfare costs of about 2.8% and 2.6% of consumption for the euro area and Switzerland and about 1.5% of consumption for Sweden.

To summarize, we now describe the two most extreme scenarios. In all cases, we report the welfare cost of a 5% nominal interest rate. Our lowest set of estimates corresponds to the first set of countries (Canada, Japan, the U.S. and the U.K.) for which the combined effect of the estimated parameters and the assumption of a zero bound on interest rates implies a range of estimates between 0.15% and 0.2% of consumption. However, for the United States, for example, one cannot reject the log-log specification together with the assumption of a lower bound of -2.4%. In this case, the estimated cost is 1.5% of consumption. Our largest

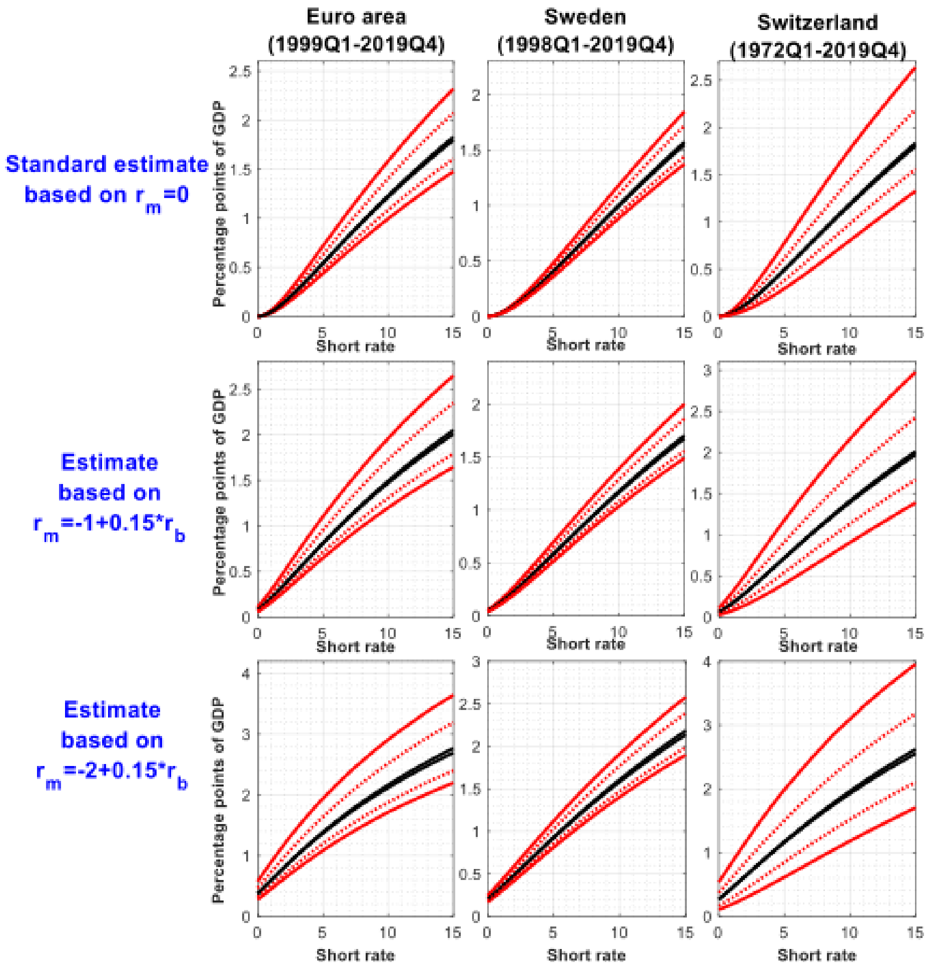

Figure 5.a: Estimated welfare cost functions based on the Selden-Latané specification for countries with negative interest rates (point estimates of the lower and upper bounds, 5-16 percentiles of the lower bounds, and 84-95 percentiles of the upper bounds of the bootstrapped distributions)

set of estimates corresponds to the second group of countries (the euro area, Sweden and Switzerland) for which the estimated cost when the lower bound is assumed to be zero is between 0.4% and 0.5% of consumption. However, these three countries experienced negative short-term rates. If we assume a zero lower bound of -2.4% and the log-log specification, the estimated costs now range between 1.5% and 2.8% of consumption.

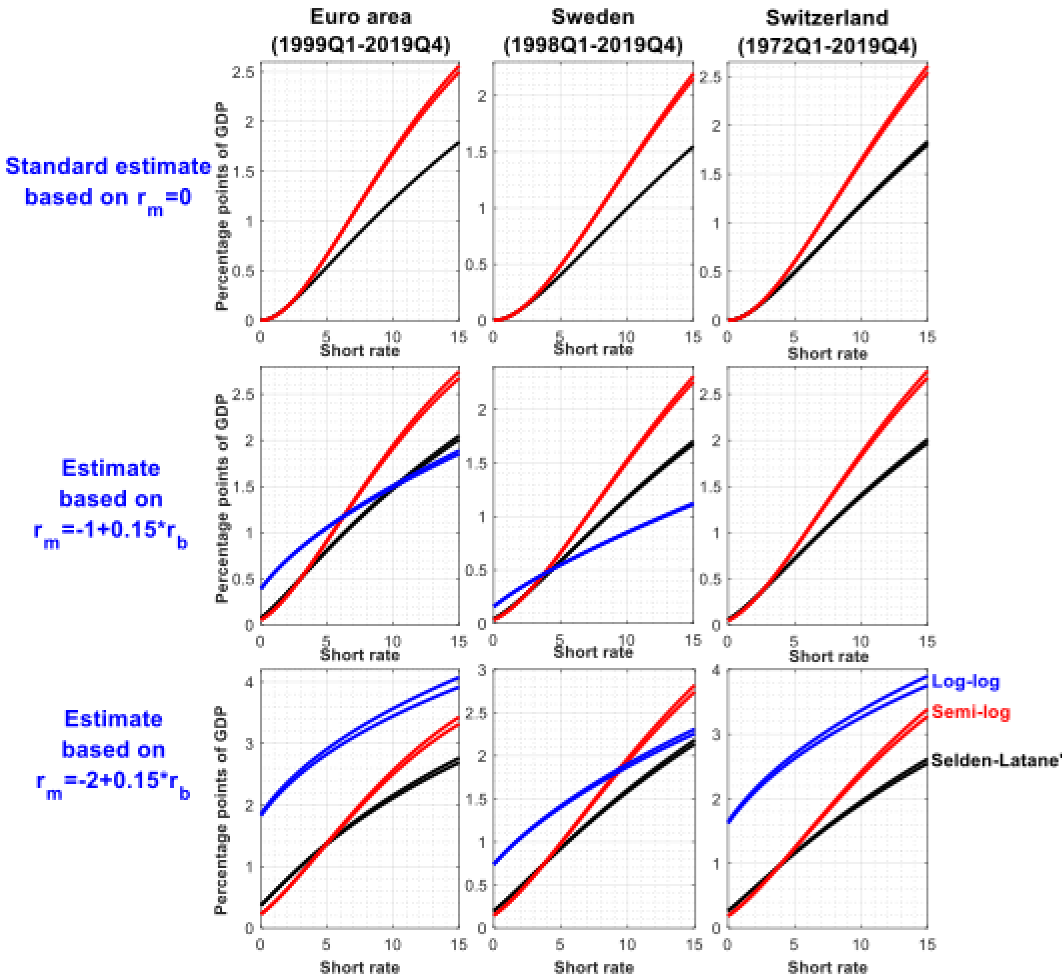

Figure 5.b: Point estimates of the welfare costs of inflation produced by alternative specifications

6 Exploring Stability and Non-linearities

A main concern in working with estimated money demand curves pertains to the stability of the *long-run* relationship over time. As previously mentioned, even without the econometric evidence produced, for example, by Friedman and Kuttner (1992), the simple visual evidence had been sufficient to discredit, long ago, any notion of stability of the U.S. demand for real M1 balances. As our results make clear, the solution proposed by Lucas and Nicolini (2015) has re-established stability of the U.S. demand for M1. However, since for all of the other countries in our dataset we work with the "standard" M1 aggregate, it is a legitimate question whether for (some of) these countries, too, some adjustment to the standard aggregate might be required in order to obtain stability of the long-run demand for M1.

6.1 Testing for stability in the cointegration vector

Table 5 reports evidence from Hansen and Johansen's (1999) tests for stability in the cointegration vector for our dataset,[27] based on any of the three money demand specifications. Only in two instances, Denmark and Japan based on the Selden-Latané specification, do the tests detect evidence of instability.[28]

Overall, there is very little evidence of a break in the real money demand relationship derived from the theory. This is reassuring in itself but also in reference to the issue raised by Ireland, which has prevailed in the discussion in the United States of a structural break in this relationship somewhere between the late '70s and the early '80s. It is the assumption of such a break that justifies focusing the analysis using only the recent data. These tests show, on one hand, that once we take into account United States-specific regulatory changes, there is no break in the money demand relationship. And on the other hand, they show that in other similarly developed countries that did not have the regulatory changes, the high inflation episode of the late '70s and early '80s is totally consistent with real money demand theory, using the standard M1 monetary aggregate.

[27] On the other hand, we do not test for stability of the loading coefficients, since they pertain to the short-term adjustment dynamics of the system towards its long-run equilibrium, and they are therefore irrelevant for the purpose of computing the welfare costs of inflation in the steady state. Finally, we eschew Hansen and Johansen's (1999) fluctuation tests, because as shown by Benati et al. (2021) via Monte Carlo, they exhibit overall a significantly inferior performance compared with the tests for stability in the cointegration vector and loading coefficients.

[28] This is in line with the evidence in Benati et al.'s (2021) Section 6.2. The main finding in that section was that evidence of breaks in either the cointegration vector or the loading coefficients vector is weak to non-existent. The estimated break dates for the cointegration vector are 2008Q1 for Denmark and 1979Q4 for Japan. The second element of the normalized cointegration vector for the first and second sub-periods is equal to -0.37 and -0.66 for Denmark and -0.41 and -0.74 for Japan.

Table 5 Bootstrapped p-values[a] for Hansen and Johansen's (1999) tests for stability in the cointegration vector for (log) M1 velocity and (the log of) a short-term rate				
		Money demand specification:		
Country	*Period*	*Selden-Latané*	*Semi-log*	*Log-log*
United States	1959Q1-2019Q4	0.5875	0.8030	0.9940
United Kingdom	1955Q1-2019Q4	0.5905	0.5480	0.9365
Canada	1947Q3-2006Q4	0.3535	0.6710	0.6910
	1967Q1-2019Q4	0.6900	0.7945	0.6070
Australia	1969Q3-2019Q4	0.7835	0.7880	0.6950
Switzerland	1972Q1-2019Q4	0.6378	0.8102	–[c]
Sweden	1998Q1-2019Q4	0.2335	0.1690	–[c]
Euro area	1999Q1-2019Q4	0.4880	0.2915	–[c]
Denmark	1991Q1-2019Q4	0.0085	0.2605	–[c]
South Korea	1964Q1-2019Q4	0.1460	0.5835	0.4485
Japan	1960Q1-2019Q4	0.0030	0.2600	0.4030
Hong Kong	1985Q1-2019Q2	0.5280	0.4510	0.8465

[a] Based on 10,000 bootstrap replications. [b] Null of 0 versus 1 cointegration vectors. [c] The last observations for the interest rate are either zero or negative.

6.2 Are there non-linearities in money demand at low interest rates?

A conceptually related issue pertains to the possibility that at low interest rates, money demand might exhibit sizeable non-linearities due to the presence of fixed costs associated with the decision to participate, or not to participate, in financial markets (see, e.g., Mulligan and Sala-i-Martin, 2000).[29] Based on this argument, at sufficiently low interest rates money demand (and therefore money velocity) should be largely *unresponsive* to changes in interest rates, since most (or all) households simply do not participate in financial markets.

[29] The intuition is straightforward. Suppose that the interest rate, R, is initially equal to zero, and consider a household with nominal assets A, which are entirely held in either cash or non-interest-bearing deposits. Crucially, suppose that if the household wants to switch a fraction of its assets into bonds B, it has to pay a fixed cost C. As R increases from zero to $R > 0$, unless $AR > C$, the household will keep all of its wealth in either cash or deposits, and only when the inequality is satisfied will it have an incentive to buy bonds. This implies that under the plausible assumption that C is heterogenous across the population, money demand should exhibit sizeable non-linearities (rather than a strict discountinuity) at low interest rates.

Although Hansen and Johansen's (1999) tests detect little evidence of instability in the cointegration vector, for the specific purpose of testing whether money demand curves might be flatter at low interest rates, these results should be discounted for (at least) two reasons.

First, as discussed by Bai and Perron (1998, 2003), when a coefficient experiences two breaks in opposite directions (e.g., first an increase, then a decrease), break tests that have not been explicitly designed to search for *multiple* breaks may have a hard time detecting the first break to begin with. Within the present context, this could be relevant for three countries: the U.S., the U.K., and Canada. In these cases, the short rate was below 5% (which, based on Mulligan and Sala-i-Martin, 2000, we take as the relevant threshold) at the beginning of the sample; it then significantly increased above 5% during the Great Inflation; and it has progressively decreased since the early 1980s. Under the assumption that money demand curves are comparatively flatter at low rates, this implies that the slope of the curve should have first increased and then decreased, which is precisely the kind of circumstance in which these tests may have problems in detecting a break.

Second, Hansen and Johansen's (1999) are tests for breaks at *unknown* points in the sample. In principle, it should be possible to perform more powerful tests if we had strong reasons for choosing a specific threshold for the short rate, which, as mentioned, we take to be 5%.

Before we delve into the econometric evidence, however, it is of interest to see what a simple visual inspection of the data suggests. Figure 6 shows informal evidence on the possible presence of non-linearities for five countries for which sub-samples with the short rate above and below 5% are both sufficiently long. In order to provide sharper evidence, for four countries (the U.S., the U.K., Canada, and Australia), we consider long samples of annual data that we do not further analyze.[30] The top row shows the raw data for M1 velocity and a short rate, whereas the bottom row shows the low-frequency components

[30] This is because, these being annual series, for all of them at least one of the sub-samples with the short rate either above or below 5% features too few observations to produce reliable results.

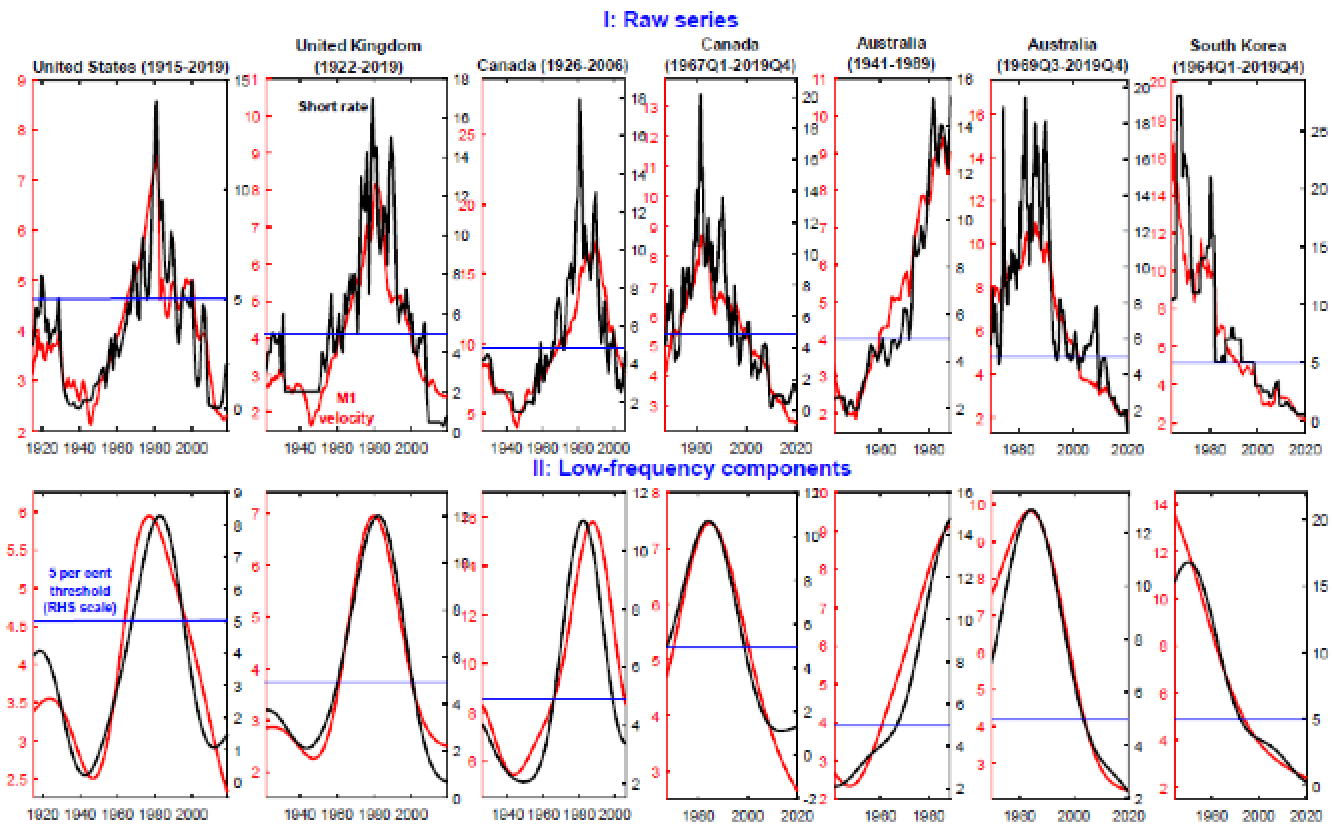

Figure 6: Informal evidence on the possible presence of non-linearities at low interest rates

of the two series.[31] The evidence in the figure speaks for itself, and it provides nearly *no support* to the notion that velocity—and therefore money demand—may be less responsive to interest rate changes at low interest rates. The only possible exceptions are the U.S. until WWII and the U.K. since the recent financial crisis. Counterexamples to these two cases, however, are provided by the U.S. and Canada since the financial crisis: for either country, the low-frequency component of velocity has plunged somewhat faster than the corresponding component of the short rate.

Overall, the "big picture" emerging from Figure 6 suggests that the relationship between M1 velocity and the short rate is virtually the same at all interest rate levels. Although we will shortly discuss the econometric results, in fact we regard this evidence, because of its simplicity, as the strongest argument against the notion that money demand curves may be

[31]The low-frequency components have been extracted via the methodology proposed by Müller and Watson (2018), setting the threshold to 30 years for the low frequencies.

flatter at low interest rates.

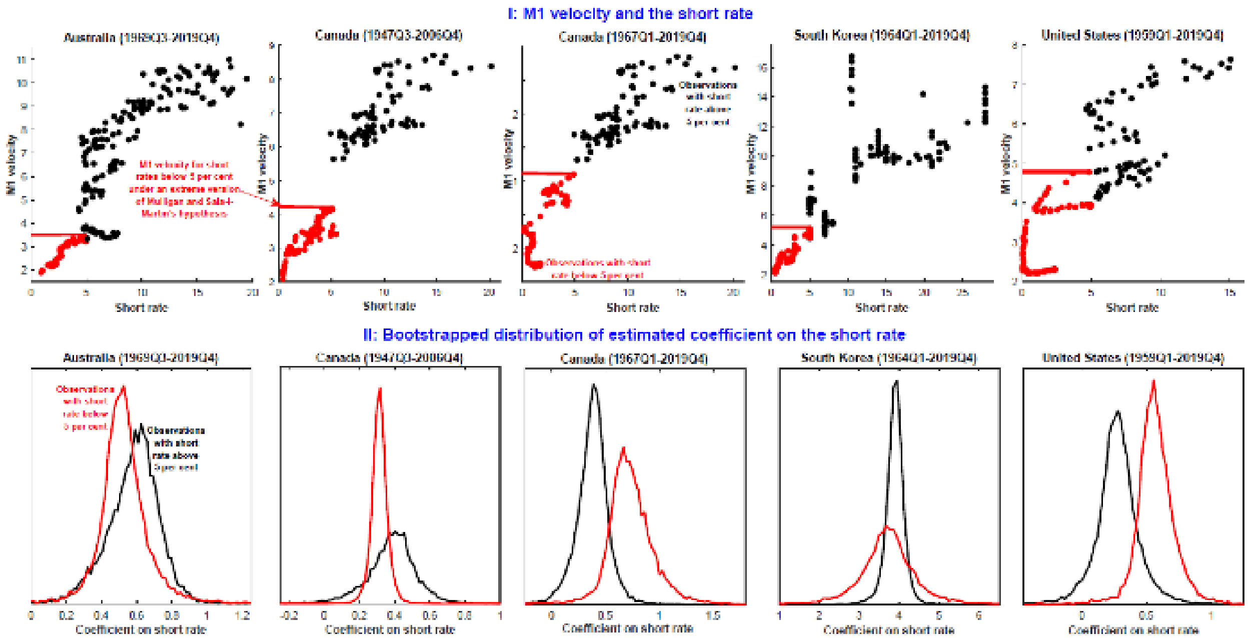

Figure 7: M1 velocity and short-term nominal interest rates: Observations with the short rate above and below 5 per cent (quarterly data)

Figure 7 shows evidence based on quarterly data for the four countries with sufficiently long continuous samples with the short rate both above and below the 5% threshold. The top row shows scatterplots of M1 velocity and the short rate, with the observations with the short rate above and below the threshold shown in black and red, respectively.[32] (The sub-samples with the short rate below and above 5% are reported in Table 6.) The panels also show a horizontal red line corresponding to an extreme version of the non-linearity hypothesis, in which when the short rate falls below 5% by an arbitrarily small quantity $\epsilon > 0$, velocity becomes completely insensitive to interest rate fluctuations (and therefore perfectly flat). The reason for reporting this extreme, and obviously implausible, case is that it provides a "reference benchmark": if the demand for M1 truly were to become flatter at low interest rates, the scatterplot with the red dots would also be flatter than the one with

[32]For Canada (1947Q3-2006Q4) it would seem that there is a discontinuity in the relationship between velocity and the short rate. In fact, this is not the case: rather, in order to obtain "clean" samples with the short rate almost entirely below or above 5%, we had to eliminate the period 1967Q4- 1973Q1, during which the short rate fluctuated around 5%. By the same token, for the U.S., we exclude the period 1991Q4-2000Q4.

the black dots, and compared with that, it would be rotated upwards and to the left towards the horizontal red line.

Table 6 Estimated coefficients on the short rate in Selden-Latané specifications for samples with the short rate above and below 5%[a]

		Based on samples with short rate:	
		below 5%	*above 5%*
Country	$P(\delta_{R<5}<\delta_{R>5})$	Estimate and 90% confidence interval	Median and 90% confidence interval
Australia	0.614	0.530 [0.321; 0.763]	0.604 [0.325; 0.802]
Canada, I	0.267	0.402 [0.138 0.612]	0.323 [0.248 0.399]
Canada, II	0.064	0.729 [0.451; 1.110]	0.399 [0.133; 0.612]
South Korea	0.584	0.351 [0.053; 0.651]	0.397 [0.305; 0.476]
United States	0.072	0.573 [0.369 0.826]	0.284 [0.008 0.561]

[a] Based on 10,000 bootstrap replications.

[a] *Samples with short rate below and above 5%:*
Australia: 2009Q1-2019Q4 and 1969Q3-2008Q4; Canada, I: 1947Q3-1967Q3 and 1973Q2-1993Q2; Canada, II: 2001Q1-2019Q4 and 1973Q2-1993Q2; South Korea: 1995Q3-2019Q4 and 1964Q1-1995Q2; United States: 2001Q1-2019Q4 and 1972Q4-1991Q3.

In fact, evidence that this might be the case is weak to non-existent. Specifically, for Australia, the visual evidence suggests that the slope is essentially the same at all interest rate levels, whereas the intercept appears to have been mostly different in the two sub-samples.[33] For Canada, the slope of the relationship between the two series appears to have been the same at all interest rate levels. For Korea, the fact that the observations with short rates above 5% are very spread out prevents us from making any strong statement. At the very least, however, evidence provides no support to the notion that the slope may have been flatter at low interest rates. Finally, evidence for the U.S. is idiosyncratic, with the observations below 5% clustered in two separate loops, but once again in no way does it suggest that the demand curve may be flatter at low interest rates.

For all sub-samples, the second row of Figure 7 reports the econometric evidence by

[33]The small cloud of black dots next to the red dots, however, suggests that the break in the intercept had nothing to do with the level of the interest rate.

showing the bootstrapped distribution of Stock and Watson's (1993) dynamic OLS (DOLS) estimator of the coefficient on the short rate in the Selden-Latané specification (10) which is our benchmark specification.[34]

Table 6 reports the point estimate of the coefficient, together with the 90% bootstrapped confidence interval and the p-value for testing the hypothesis that when the short rate is below 5%, the coefficient might be smaller than when it is above this threshold. The consistent message from Table 6, and from the bottom row of Figure 7, is that there is no econometric evidence in support of the notion that money demand curves may be flatter below 5%. First, the simple point estimates of ϕ are smaller for $R_t < 5\%$ only for Australia and South Korea, but in both cases, the p-values (at 0.614 and 0.584, respectively) are far from being significant even at the 10% level. Second, in two of the remaining cases (the U.S. and Canada, 1967Q1-2019Q4), the p-values (equal to 0.072 and 0.064, respectively) suggest that ϕ has been *larger*, rather than smaller, for short rates below 5% (this is also clearly apparent from the bottom row of Figure 5).

6.3 Spurious non-linearity from estimating log-log specifications

Suppose that the data have been generated by a Selden-Latané specification so that the relationship between the levels of velocity and the interest rate is *identical* at all interest rate levels. Since a given percentage change in the *level* of the interest rate (say, 1%) is associated with a larger change (in absolute value) in its *logarithm* at low interest rates than it is at higher interest rates,[35] this automatically maps into lower estimated elasticities (in absolute value) at low interest rates than at higher interest rates. This implies that if the true specification is the Selden-Latané specification, estimating a log-log specification automatically produces smaller elasticities (in absolute value) at lower rather than higher

[34]The methodology we use is the same as discussed in Section 5. Specifically, we estimate the cointegration vector via Stock and Watson's (1993) DOLS estimator; we then estimate the VECM for V_t and R_t via OLS, by imposing in estimation the previously estimated cointegration vector (which, as discussed in Luetkepohl, H., 1993, is correct in the presence of a single cointegration vector); and finally, we characterize uncertainty about the cointegration vector by bootstrapping the VECM as in Cavaliere et al. (2012).

[35]For example, ln(9)-ln(10)=-0.105, whereas ln(2)-ln(3)=-0.406.

interest rates.

This can be illustrated as follows. With the true money demand specification described by (10), estimating the log-log specification (8) produces the following theoretical value of the *estimated* elasticity:

$$\frac{d\ln\left(\frac{M_t}{P_t y_t}\right)}{d\ln r_t} = -\frac{\phi r_t}{a^3 + \phi r_t}, \tag{19}$$

which tends to -1 for $r_t \rightarrow \infty$, but tends to 0 for $r_t \rightarrow 0$ (in fact, for r_t=0, it is exactly equal to 0). The same argument obviously holds if the true specification is the semi-log.

In all the specifications estimated by Mulligan and Sala-i-Martin (2000) (as well as by Attanasio et al., 2002) the interest rate is entered in logarithms.[36] To be sure, this does not imply that Mulligan and Sala-i-Martin's finding, based on micro data, of a smaller elasticity at low interest rate levels is spurious. What it does imply, however, is that by entering the interest rate in logarithms, they would have automatically obtained this result even if the relationship between the levels of velocity and the short rate were identical at all interest rate levels.

7 Conclusions

How large is the cost of deviation from the Friedman rule if the nominal interest rate is set at 5% in the steady state? A well established tradition, started by Bailey (1956) and Friedman (1969), estimates those costs computing the area under the real money demand curve. Lucas (2000) follows this tradition and, arguing that a log-log specification is a good fit for the U.S. data during the 20th century, computes that cost to be 1.2% of lifetime consumption. A feature of the log-log specification is that it has an asymptote when the nominal interest rates go to zero. This feature makes the integral under the real money demand large.

However, Ireland (2009) argued that a semi-log specification provides a much better fit

[36]For Mulligan and Sala-i-Martin (2000), see equations (10), (11), (13) and (14). For Attanasio et al. (2002), see the estimates in Tables 3 and 7.

if one disregards the data until 1980. He also argues that the elasticity is much lower than the one used by Lucas. When both things are considered, Ireland estimates the welfare cost to be a mere 0.04% of consumption. A distinct feature of the semi-log specification is that it has a finite satiation point when the interest rate is zero, so the integral under the real money demand is not as large as with the log-log.

We use new data for the U.S., analyzed in detail in Lucas and Nicolini (2015), that provide a unified stable behavior for the U.S. from 1900 to 2020. In addition, we study the behavior of real money demand for several other developed countries. Finally, we also consider the functional form studied by Selden (1956) and Latané (1960). The Selden-Latané functional form shares with the semi-log the property that there is a finite satiation level of money balances when the interest rate is zero.

Many of these countries share with the United States the high nominal interest rates of the '70s and '80s and the very low interest rates of the later years of the data. In line with the analysis in Lucas and Nicolini (2015), the evidence of all these other countries is remarkably consistent with the notion of a stable real money demand when one includes the high interest rates period and also when interest rates are very low. This last feature is important, since it is when interest rates are very low that the log-log, the semi-log, and the Selden-Latané behave very differently. These are the observations that help identify the functional form that fits best.

The consideration of other countries also brings a new dimension to the analysis. It has been customary in the literature to assume that the own rate of money is zero. This assumption, together with utility maximization, implies that the lower bound on the short-term interest rate is zero. Thus, in computing the integral of real money demand, the lower bound on the interval has always been set to zero. However, the experience of the euro area, Sweden, and Switzerland, where the short-term interest rate has been negative for a substantial number of periods, makes evident that the true bound is lower than zero. As it turns out, this assumption is key in order to understand our results.

If we assume, as in the literature, that the lower bound is zero, then the Selden-Latané is the preferred specification overall. In all countries, this functional form performs better in the cointegration test, with the exceptions of one of the samples for Canada and one of the tests (Johansen's) that we run. In addition, in the tests that compare this functional form with the semi-log, it dominates in all cases. Thus, under the assumption that the lower bound is indeed zero, Ireland's argument that a real money demand that has a finite value at the lower bound dominates the log-log, which has an asymptote, is correct.

When using this functional form and assuming a zero lower bound, the welfare costs for the United States, Canada, Japan, and the United Kingdom are between 0.2% and 0.4% of consumption, substantially lower than the 1.2% in Lucas, but much higher than the 0.04% in Ireland. The reason why our estimate is much larger, in spite the fact that we use a functional form with finite money balances at the lower bound, is that the elasticity estimated with the monetary aggregate we use is five times larger than the one obtained by Ireland. Our estimate is consistent with the one Lucas used for the semi-log and is very similar to the ones we obtain for the other countries.

However, once we allow the lower bound to be negative, the log-log functional form, though not necessarily the specification with the best performance, cannot be rejected for several countries. In this case, for the United States, we detect cointegration, and our estimate of the elasticity parameter is 0.5, consistent with the squared-root formula in Baumol-Tobin models, and the same used by Lucas. The estimated cost for the United States of a 5% interest rate in a steady state is around 1.5% of consumption, higher than the one obtained by Lucas. The reason is that he integrated the curve starting at his assumed zero lower bound, while we start at a negative value. For some European countries we obtain even larger estimates, including up to 2.8% for Switzerland.

References

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

Alvarez, F, F. Lippi, and R. Robatto (2019): "Cost of Inflation in Inventory Theoretical Models,"*Review of Economic Dynamics*, 32, 206-226.

Attanasio, O. P., L. Guiso, and T. Jappelli (2002): "The Demand for Money, Financial Innovation, and the Welfare Cost of Inflation: An Analysis with Household Data", *Journal of Political Economy*, 110(2), 317-351.

Bai, J. and P. Perron (1998): "Estimating and Testing Linear Models with Multiple Structural Changes", *Econometrica*, 66(1), 47-78

Bai, J. and P. Perron (2003): "Computation and Analysis of Multiple Structural Change Models", *Journal of Applied Econometrics*, 18(1), 1-22

Bailey, M. J. (1956): "The Welfare Cost of Inflationary Finance", *Journal of Political Economy*, 64(2), 93-110.

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

Belongia, M. T., and P. N. Ireland (2019): "The Demand for Divisia Money: Theory and Evidence", *mimeo*, May 2019.

Benati, L. (2008): "Investigating Inflation Persistence across Monetary Regimes", *Quarterly Journal of Economics*, 123(3), 1005-1060.

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

Benati, L. (2020): "Money Velocity and the Natural Rate of Interest", *Journal of Monetary Economics*, 116, 117-134

Benati, L. (2021): "The Monetary Dynamics of Hyperinflations Reconsidered", *mimeo*

Benati, L., R. E. Lucas Jr., J. P. Nicolini, and W. Weber (2021): "International Evidence on Long-Run Money Demand", *Journal of Monetary Economics*, 117, 43-63.

Blanchard, O.J., G. Dell'Ariccia, and P. Mauro (2010): "Rethinking Macroeconomic Policy", IMF Staff Position Note 10/03.

Bresciani-Turroni, C. (1937): *The Economics of Inflation*, John Dickens and Co. Northampton.

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.

Christiano, L. J., and T. J. Fitzgerald (2003): "The Band Pass Filter", *International Economic Review*, 44(2), 435-465.

Cochrane, J.H. (1994): "Permanent and Transitory Components of GNP and Stock Prices", *Quarterly Journal of Economics*, 109(1), 241-265.

Coibion, O., Y. Gorodnichenko, and J. Wieland (2012): "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise their Inflation Targets in Light of the ZLB?", *Review of Economic Studies*, 79, 1371-1406.

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.

Dotsey, M., and P. Ireland (1996): "The Welfare Cost of Inflation in General Equilibrium", *Journal of Monetary Economics*, 37(1), 29-47.

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

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

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

Feldstein, Martin S. (1997): "The Costs and Benefits of Going from Low Inflation to Price Stability", in C. D. Romer and D. H. Romer, editors, *Reducing Inflation: Motivation and Strategy*, University of Chicago Press, 123–166.

Friedman, B. M., and K. N. Kuttner (1992): “Money, Income, Prices, and Interest Rates”, *American Economic Review*, 82(3), 472-492.

Friedman, M. (1969): “The Optimum Quantity of Money”, in M. Friedman, editor, *The Optimum Quantity of Money and Other Essays*, 1-50. Chicago: Aldine Publishing Company

Goldfeld, S. M. (1976): “The Case of the Missing Money”, *Brookings Papers on Economic Activity*, 3, 683-730.

Goldfeld, S. M., and D. M. Sichel (1990): “The Demand for Money”, in B. M. Friedman, and F.H. Hahn, editors, *Handbook of Monetary Economics*, Vol. I, Amsterdam, North Holland

Hamburger, M.J. (1977): “Behavior of the Money Stock: Is There a puzzle?”, *Journal of Monetary Economics*, 3(3), 265-288.

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

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

Ireland, P. (2009): “On the Welfare Cost of Inflation and the Recent Behavior of Money Demand”, *American Economic Review*, 99(3), 1040-1052.

Kurlat, P. (2019): “Deposit Spreads and The Welfare Cost of Inflation”, *Journal of Monetary Economics*, 106, 78-93.

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

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

Lucas, R.E., Jr. (2000): “Inflation and Welfare”, *Econometrica*, 68(2), 247-274.

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

Luetkepohl, H. (1993): *Introduction to Multiple Time Series Analysis*, 2nd edition.

Springer-Verlag.

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

Modigliani, F. and R. Sutch (1966): "Innovations in interest rate policy", *American Economic Review*, 56(1/2), 178-197.

Müller, U. and M. W. Watson (2018): "Long-Run Covariability", *Econometrica*, 86(3), 775-804

Mulligan, C. B., and X. Sala-i-Martin (2000): "Extensive Margins and the Demand for Money at Low Interest Rates", *Journal of Political Economy*, 108(5), 961-991.

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

Sidrauski, M. (1967*a*): "Inflation and Economic Growth", *Journal of Political Economy*, 75(6), 796-810.

Sidrauski, M. (1967*b*): "Rational Choice and Patterns of Growth in a Monetary Economy", *American Economic Review*, 57(2), 534-544.

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

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

Vayanos, D., and J.-L. Vila (2021): "A Preferred-Habitat Model of the Term Structure of Interest Rates", *Econometrica*, 89(1), 77-112.

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