b
UNIVERSITÄT
BERN

Faculty of Business, Economics
and Social Sciences

Department of Economics

The Welfare Costs of Inflation

Luca Benati, Juan-Pablo Nicolini

21-13

August, 2021

DISCUSSION PAPERS

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

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 are close to those obtained by Lucas (2000), and an order of magnitude higher than those in Ireland (2009).

*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. One such class is discussed below.

There is a wide range of estimates in the literature. For a steady state interest rate of 5 percent, Lucas (2000) computes the cost to be around 1.1 percent of lifetime consumption, which is a sizeable amount. However, Ireland (2009) challenges Lucas' 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 the 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, data is used by both authors to discipline their choices. And so will we.

Our main contribution is to bring more data to the debate. We do so in two dimensions. 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 rate. Thus, it helps identify the behavior of money demand in that range that, as we will discuss, is very important to identify the functional form.

Second, we also use data from several other developed countries, that had similar inflation histories as the United States. Although one could certainly entertain differences across countries, this evidence is also useful to identify 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 - as most of the monetary economics literature till 2010! - assumed the lower bound to be zero. However, the negative interest rates observed in the Euro area, Sweden and Switzerland 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 percent is between 0.20 to 1.50 percent of lifetime consumption, depending on the functional form assumed and the assumption regarding the lower bound. 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 percent of lifetime consumption.

Modern analysis of optimal monetary policy is typically performed with models belonging to the New Keynesian paradigm. These models consider money-less economies only, so they ignore the welfare effect of lack of money satiation that we focus on. Two reasons, 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, that computes those costs to be negligible.

We challenge both notions. Regarding the first reason, we present overwhelming evidence that there is no sense in which modern economies are becoming money less.[1] Regarding the second, Ireland bases his computations on USA data only. This is problematic, since there was already at the time substantial evidence that the standard measure of M1, that had maintained a stable relationship with interest rates and nominal output by 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 since 1994. He shows, however, that even after this adjustment, the behavior of real money demand 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), for two different reasons. First, he shows that the log-log functional form preferred by Lucas performs much worse that the semi-log specification. Second, for the semi-log specification, he estimates a much lower semi-elasticity of the real money demand with respect to the nominal interest rate 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 changed the availability of transactional assets in the United States. Specifically, they propose to add the Money Market Demand accounts, created in 1984, to M1.[2] Once these new deposits are taken into account, a remarkably stable real money demand is obtained, that behaves the same way before and after 1980. The estimates of the real money demand using the Lucas and Nicolini aggregate, that they label NewM1, imply larger estimates of the welfare cost of inflation than those obtained by Ireland. As before, the reason is two-fold. First, even using 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 support the results for the United States when using the NewM1 aggregate, in terms of both the estimated semi-

[1] A more detailed analysis with yearly data that includes more countries can be found in Benati, Lucas, Nicolini and Weber (2021).

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

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 ignoring money in analyzing optimal monetary policy can be seriously misleading. For instance, Coibion, Gorodnichenko and Wieland (2012) make a compelling argument against increasing the inflation target in countries like the United States, in a model with frictions in the setting of prices and with 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 percent in their preferred specification to be close to 0.6 percent of life-time consumption. That number, that combines the cost created by price frictions and the probability to be 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 reinforce the argument of their paper.

On the theory side, we innovate in that we construct 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, Lippi and Robatto (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 ever 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, that 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 momey demand curve. In Section 3 we discuss the data and several Figures that, in our view, present very solid evidence in favor of stable money demand relationships for the countries we analyze. Section 4 makes formally 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

[3] Coibion, Gorodnichenko and Wieland (2012) explicitly aknowledge 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 models that ignore the creation of inside money.

computations for the welfare cost functions. For each interest rate level, we compute the bootstrapped distribution of the welfare costs (and therefore median estimates, and confidence intervals) expressed in percentage points of GDP. 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, that 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 produce goods and to 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 are average money balances and n_t is the number of portfolio adjustments within each period. The variable n_t is the only economically relevant decision to be made by the representative agent.

[5] The baseline model is discussed at length in Benati et. al. (2020).

We allow for money to pay a nominal return, that we denominate r_t^m, that is allowed to become negative. This implies a point of departure from most of the literature that sets $r_t^m = 0$. This is important, as we explain below, for the model 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, that can be allocated to money or interest bearing bonds, B_t so a restriction to the optimal problem of the agent is

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

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

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

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 1, 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, which implies that

$$\frac{M_t}{P_t y_t} = \frac{1}{n_t}, \tag{6}$$

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, secular increases in productivity do not affect the optimal solution for n_t, so the theory implies a unit income elasticity of real money demand.

Note that the solution for n_t, and therefore the solution for real money demand, depends on the interest rate *differential* between bonds and money. As mentioned above, in most of the money demand literature, it is customary to assume that $r_t^m = 0$. in which case real money demand depends on the interest rate in bonds. For further references, we let the interest rate differential between bonds and money to 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

$$r_t = r_t^b - r_t^m \geq 0. \tag{7a}$$

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 ($7a$) plus the standard assumption in the literature that $r_t^m = 0$. The analyses of both Lucas (2000) and Ireland (2009) are done under this standard assumption.

The recent experience of prolonged negative short term interest rates in several countries severely challenges this notion. As condition ($7a$) must hold in equilibrium, the challenge within the confines of the model we use can only be solved 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 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. As for cash, a negative return can be rationalized by the risk of being lost or stolen, as Alvarez and Lippi (2013) compute using survey data.[7] As 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

$$r_t^m = -a + br_t^b, \tag{8}$$

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 ($7a$) , 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, and 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 of the three functional forms exhibit a unit income elasticity, as implied by the model. The first specification is the log-log one,

$$\ln \frac{M_t}{P_t y_t} = a^1 - \eta \ln r_t + u_t^1, \tag{9}$$

that 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

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

[7] Alvarez and Lippi (2013) calibrate this return at -0.02 using survey data from Italy.

[8] Further details are provided in the Online Appendix 2.

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

$$\ln \frac{M_t}{P_t y_t} = a^2 - \gamma r_t + u_t^2, \tag{10}$$

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

$$\frac{M_t}{P_t y_t} = \frac{1}{a^3 + \phi r_t + u_t^3}, \tag{11}$$

both 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 (10) in his revision of Lucas's estimate.

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

As we show below, the welfare costs 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.[9]

In computing the welfare cost of inflation we consider these three functional forms and three alternative assumptions regarding the parameters a and b, that relate the return on money and the short term interest rate, as described in the linear relationship (8). Besides being based on empirical evidence, such a linear relationship has the additional 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 or the inverse of r_t, one only needs to test and estimate those two specifications under the benchmark case of $a = b = 0$. and 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 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.

[9] This is in line with the evidence in Benati, Lucas, Nicolini, and Weber (2021).

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, Lippi and Robatto (2019) to a class of models that is more restrictive than the ones they used. Specifically, we only consider representative agent models in which the cost of transforming liquid into illiquid assets 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, Lippi and Robatto (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.[10] 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. This is so much so that in most of the figures the difference between the two is invisible to the eye.

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

$$n^2 \frac{\theta_n(n)}{(1-\theta(n))} = r, \tag{12}$$

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

$$\frac{\partial \omega^W(r)}{\partial r} = \frac{\partial \theta(n)}{\partial n} \frac{\partial n}{\partial r}(r) > 0. \tag{13}$$

We now show how the function $\omega^W(r)$ can be bounded above and below using the integral under the money demand curve.[11]

The area under the demand curve is equal to

$$\omega^D(r) = \int_0^r m(z)dz - m(r)r, \tag{14}$$

so

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

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

[11] The analysis below follows closely the ideas in Alvarez, Lippi and Robatto (2019).

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 (12), becomes

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

Using the definition in (13),

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

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 therefore obtain our bounds as

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

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 (9) to (11), as we show below.

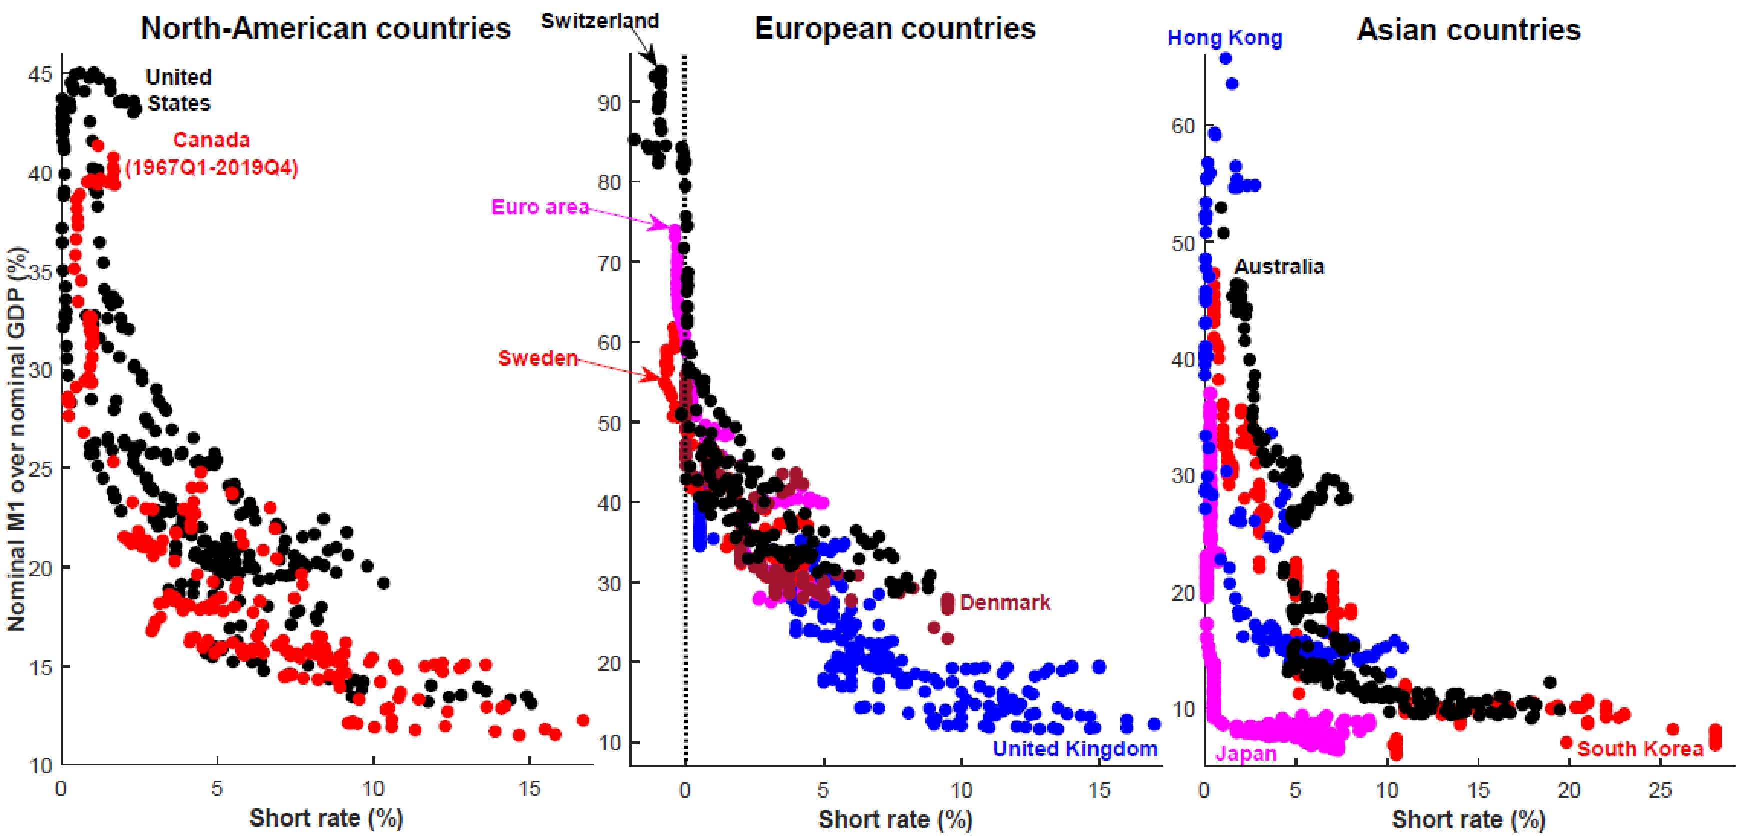

Figure 1 Scatterplots of nominal M1 over nominal GDP against the short rate

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.[12] The single exception is the United States, in which case 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.[13]

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, which is the hallmark of the theory of real money demand. A comparison between the three panels highlights several interesting features. The first is that there appear 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 an even starker extent of heterogeneity, with each individual country essentially having its own demand curve.[14] 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 shows the raw data in the way that has become standard in empirical studies of money demand. Depicted in this way, however, the plots conceal the variables' behavior over time, thus failing to show the persistence exhibited by both series, and in particular how the persistent components of the two variables have co-moved along the sample. This information is also very useful, and it provides a powerful visual

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

[13] 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.

[14] Notice that since for Japan, Hong Kong, and to a lesser extent South Korea the short rate 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. E.g., for Japan it is around 10 per cent.

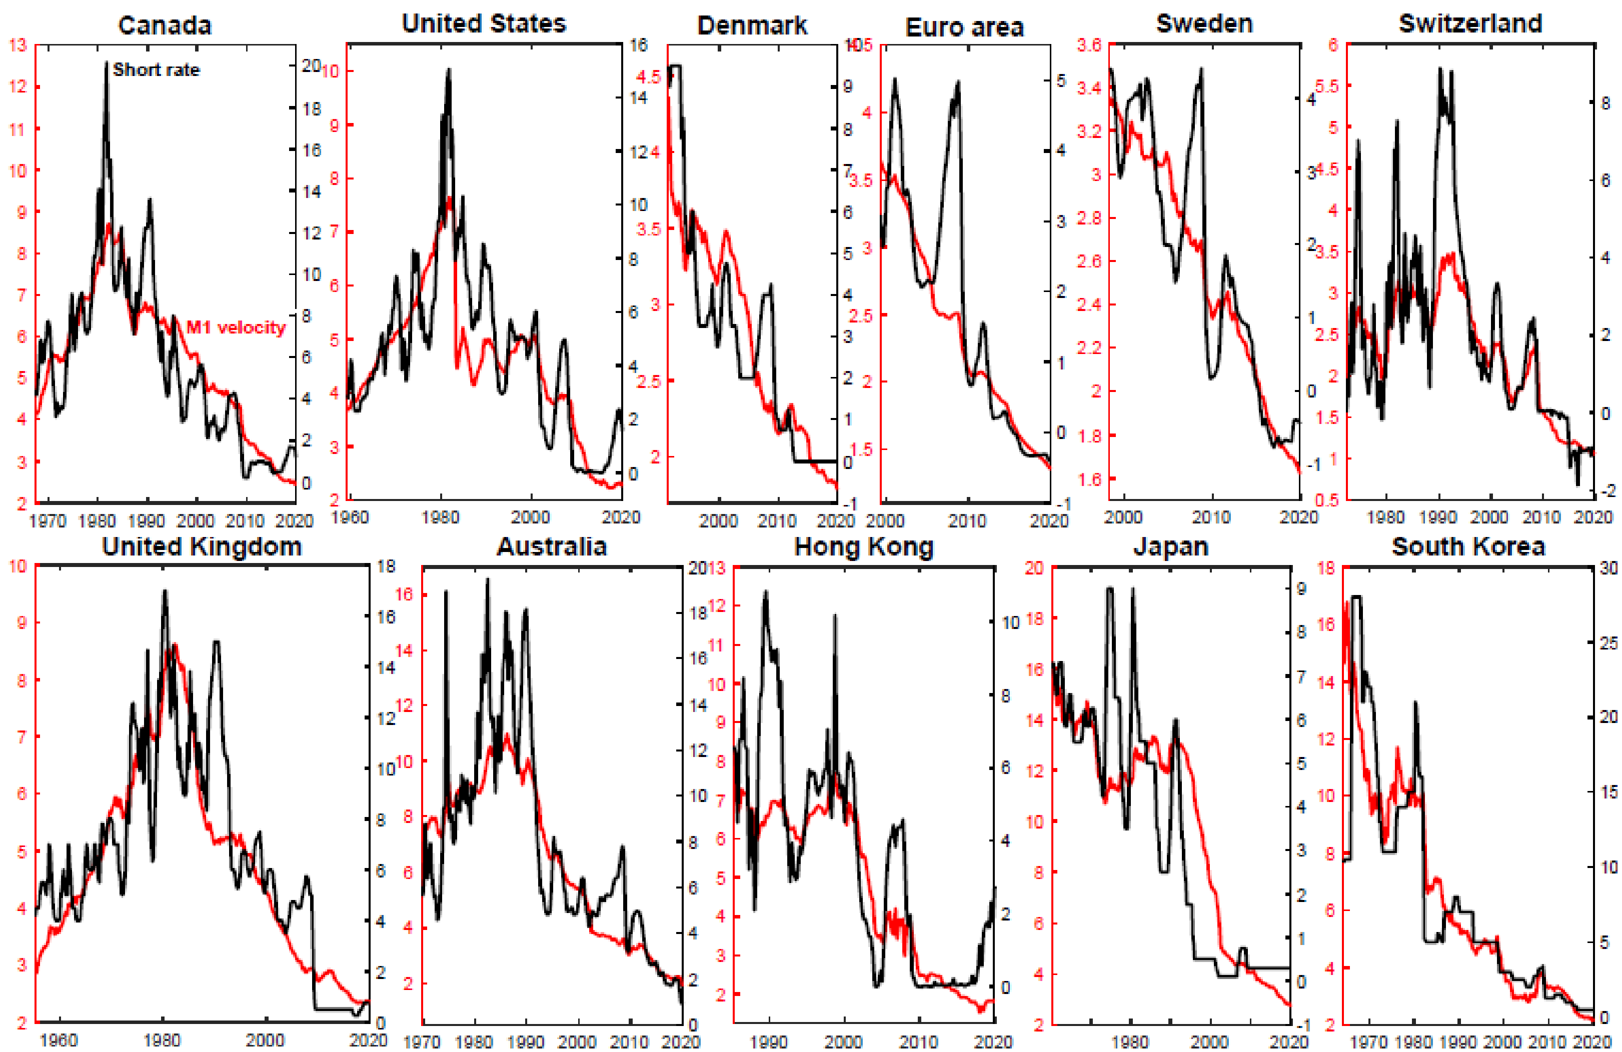

Figure 2 M1 velocity and the short rate

motivation for the cointegration methods that we use in the rest of the paper. Figure 2 therefore shows the time series for M1 velocity and the short-term nominal rate in our sample. The data so displayed suggests 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, the null hypothesis of a unit root cannot be rejected for nearly all countries and all series.[15] In searching for a cointegration relationship between velocity and the short rate, in the next section we will therefore proceed as follows. First, taking the results from the unit root tests literally—i.e, as indication that the series contain *exact* unit roots—we will test for cointegration based on Johansen's tests, which are predicated in 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 Benati (2020) and Benati *et 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 ought to specify a functional form, and to define a lower bound for the interest rate. In what follows we present results for the functional forms (9), (10) and (11), 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.[16]

[15] For the short rate it can rejected only for Denmark (in levels) and Canada (1947Q3-2006Q4) in logarithms. For M1 velocity it can only be rejected for South Korea (in levels), whereas results for the Euro 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 2014 bootstrapping procedure, should be regarded as a good approximation), with eleven 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).

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

To compute welfare costs, we consider three alternative values for the lower bound on interest rates.[17] The first case we explore is the one in which $r^m = 0$, as it is typically assumed in the literature, which implies $a = b = 0$. This assumption is inconsistent with the evidence for several countries in our sample. Indeed, in the Euro area, Switzerland, and Sweden, short term interest rates have been consistently negative for the last several years (see Figure 1). In order for the theory to account for negative short-term rates, in the other two cases we assume that $r^m = -a + br^b$.

Kurlat (2019) very precisely estimates the slope parameter b to be 0.15, using micro-data from the US. We adopt that value. On the other hand, the constant a depends on fixed costs of holding deposits, as well as the negative return in cash, related to the probability of being lost of stolen. For this parameter, we explore with two different values, that allows 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 percent and -2.4 percent, 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, where 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 the 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.

Table 1 reports, for any of the three money demand specifications discussed in Section 2, bootstrapped p-values for Johansen's maximum eigenvalue test of 0 versus 1 cointegration vectors.[18] For Canada we have two partially overlapping M1 series that cannot be linked, since they are slightly different. We present results based on either 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 before, 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 assuming a zero lower bound, evidence of cointegration is strong based on the Selden-Latané specification, whereas it is slightly weaker based on the semi-log, and it is materially weaker based on the log-

[17] While the cointegration tests and the parameter estimates are invariant to the assumption of the lower bound on the short rate for two of the specifications, the estimates of the welfare cost are not, since the lower bound of the integral depends on it.

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

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.

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.

For the log-log, cointegration is not detected for the United States, the United Kingdom, and Australia, whereas 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.[19] In a few cases, the p-values go down substantially, most notably in the USA where the test does indeed detect cointegration at the 10% level, for the lower bound assumed to be -2.4 percent, the number consistent with the experience in Switzerland.

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

[19] The only exceptions are Canada for the second sample, where the p-values increase, and Hong Kong, wher 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.

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 either $a = 0$ or $a = -1$.

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

Country	*Period*	a=0, b=0	a=-1, b=0.15	a=-2, 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.

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 either 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 (10) and (9) boils down to whether the dynamics of log M1 balances as a fraction of GDP (i.e., minus log velocity) is 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

Table 2a 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.

there for all countries,[20] 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 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:

$$\ln\left(\frac{M_t}{Y_t}\right) = \alpha + \sum_{j=1}^{p} \beta_j \ln\left(\frac{M_{t-j}}{Y_{t-j}}\right) + \sum_{j=1}^{p} \delta_j \left(\frac{R_{t-j}^{\theta} - 1}{\theta}\right) + \varepsilon_t \tag{15}$$

We estimate (15) *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 which does not involve any Bayesian priors, as it uniquely involves the likelihood of the data.[21] 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 of the other estimation details are identical to Benati (2008), to which the reader is referred to.

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.[22]

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

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

[21] 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 encodes the priors about β.

[22] This crucially hinges on the fact that we are here exclusively focusing 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.

Table 3*b* 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		
Country	*Period*	*Semi-log*	a=0 b=0	a=-1 b=0.15	a=-2 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.

Table 3*c* 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, Euro area, and Denmark there is no comparison because the last observations for the short rate are negative.

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.1a-A.1b 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 where the semi-log is the preferred specification. This is particularly so for those countries where the cointegration tests also improve substantially for lower values of the bound, like the United States of the United Kingdom. However, with the single exception of Canada, the increase is not enough for the log-log to dominate the semi-log.

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

$$R_t = \alpha + \sum_{j=1}^{p} \varphi_j R_{t-j} + \sum_{j=1}^{p} \xi_j \left[\frac{\left(\frac{Y_{t-j}}{M_{t-j}} \right)^{\theta} - 1}{\theta} \right] + \varepsilon_t \tag{16}$$

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

At first sight this approach might appear as questionable: since we are here 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,[23] so that focusing (e.g.) on the Selden-Latané specification, $V_t = a + bR_t^P$, 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.[24] This can be seen quite clearly in Figure 2 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, i.e. the dominant driver of its long-horizon variation, and it is

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

[24] 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 only impact the short end of the yield curve, and therefore have a much smaller (and in the limit negligible) effect.

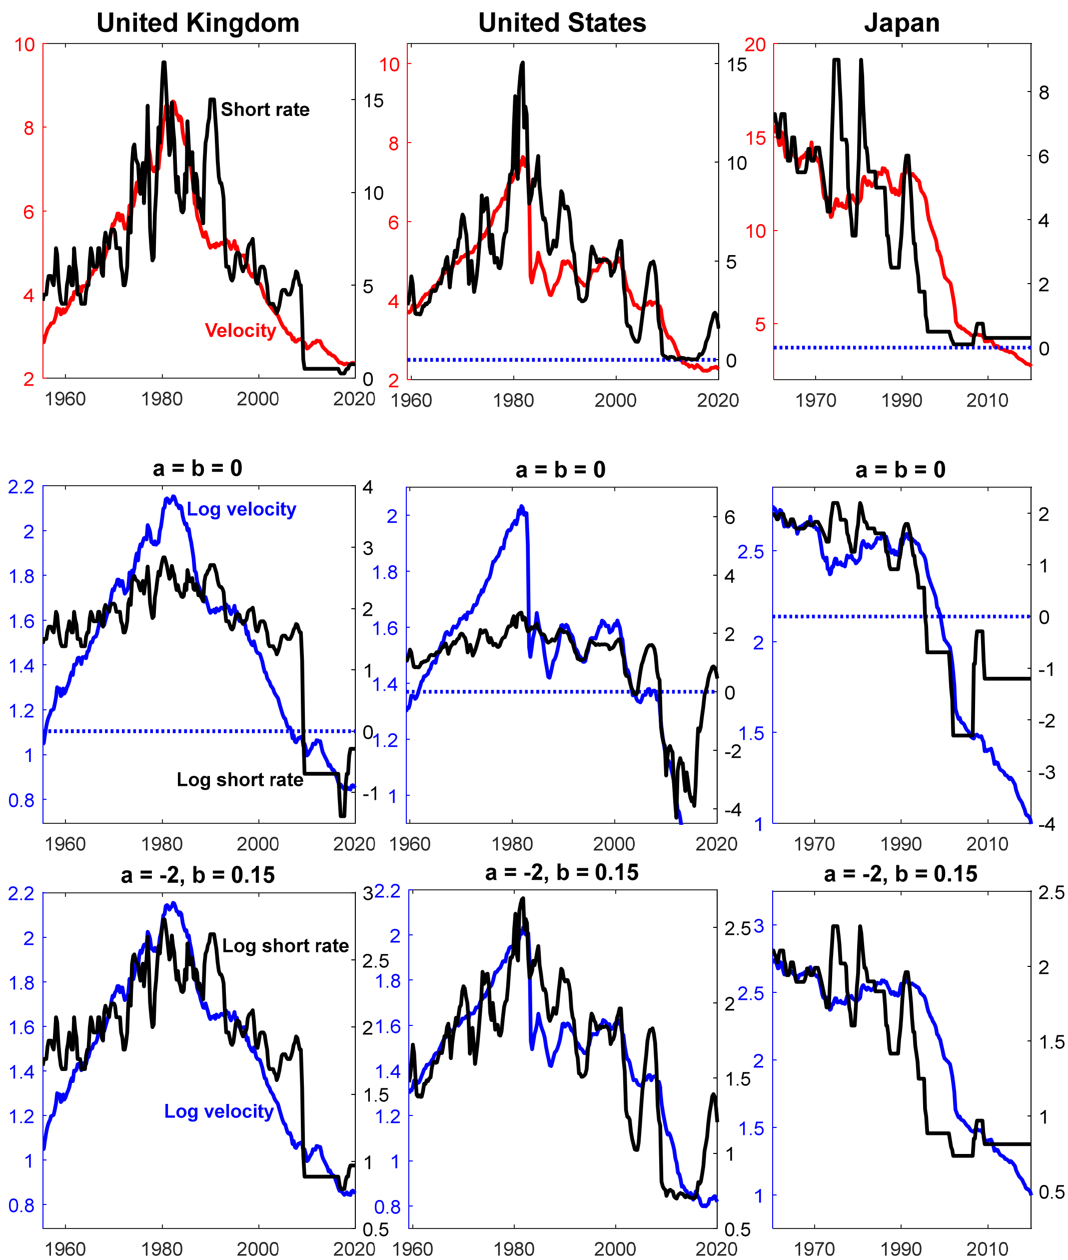

Figure 3 Comparing the Selden-Latané and log-log specifications

therefore conceptually akin to (e.g.) regressing GDP on consumption.[25]

The results are reported in Table 3*c*. The evidence is much sharper than for the previous comparison between the semi-log and the log-log: in particular, for p equal to either 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, where we visually compare the log-log with the Selden-Latene 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-Latane 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. The bottom panel, the case where the lower bound to be -2.4 percent. As it can be seen, the relationship is evident with the Selden-Latane top-panel specification and quite blurred for the log-log case middle panel with a zero bound. However, with the assumption of a -2.4 percent zero bound, the log-log bottom panel specification is as evident as with the Selden-Latane 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 this, we choose to use the Selden-Latane as our benchmark functional form. But we will also provide estimates for the other two specifications. As argued by Lucas, the computed costs is substantially higher using the log-log specification, particularly when the lower bound is assumed to be lower than zero. These computations raise some caution on our benchmark estimates since the log-log specification cannot always be clearly ruled out, particularly so when the lower bound is assumed to be low. Tables with results for all the countries, the three functional forms and the three assumptions regarding the lower bound can be found in the Online Appendix.

4.4 Parameter Estimates

The 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. To do the Johansen test, the corresponding money demand equation is estimated in its VECM form, from which the money demand

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

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 chose to directly estimate the money demand equations using Stock and Watson's dynamic OLS procedure, that 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-Latane specification, γ for the semi-log and η for the log-log. Table 4.a 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 US, of 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 the Ireland used a different monetary aggregate, as we explained above.

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 UK, 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 way 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. As it can be seen, 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 USA, of 0.54, is remarkable 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.

We are now ready to discuss our welfare cost computations.

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 (12), for the three specifications. The integrals are given by

$$\omega_{\log-\log}(r) = a^1 \frac{\eta}{1-\eta} r^{1-\eta}, \tag{17}$$

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.

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.

$$\omega_{semi-\log}(r) = \frac{a^2}{\gamma}\left(1 - \frac{1+\gamma r}{e^{\gamma r}}\right) \tag{18}$$

and

$$\omega_{Sel-Lat}(r) = \frac{1}{\phi}\ln\left(\frac{a^3+\phi r}{a^3}\right) - \frac{r}{a^3+\phi r}, \tag{19}$$

respectively, for the log-log, the semi-log and the Selden-Latané. As it 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, and leave for the appendix the analysis with Wright's tests, where, as 'point estimates', we pick 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.[26] The methodology we use in order to compute the welfare costs of inflation follows Luetkepohl (1991, pp. 370-371). We first estimating *via* OLS the cointegrating regression corresponding to any of the three specifications, i.e. to either (9), (10), or (11). 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, i.e. 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, and we turn it into its corresponding near unit root VAR by shrinking the unit root to λ=1-0.5$\cdot$(1/T), where T is the sample length.[27] 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. In short, the two bootstrapping procedures produce numerically near-identical results, and in what follows we will therefore exclusively report and discuss those based on bootstrapping the VECM (the alternative set of results is however available upon request).

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, comparatively smaller M1 balances as a fraction of GDP than the remaining countries. Since for

[26]These cases are Japan and the US for the semi-log, the Euro area for the Selden-Latane and Switzerland for the log-log.

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

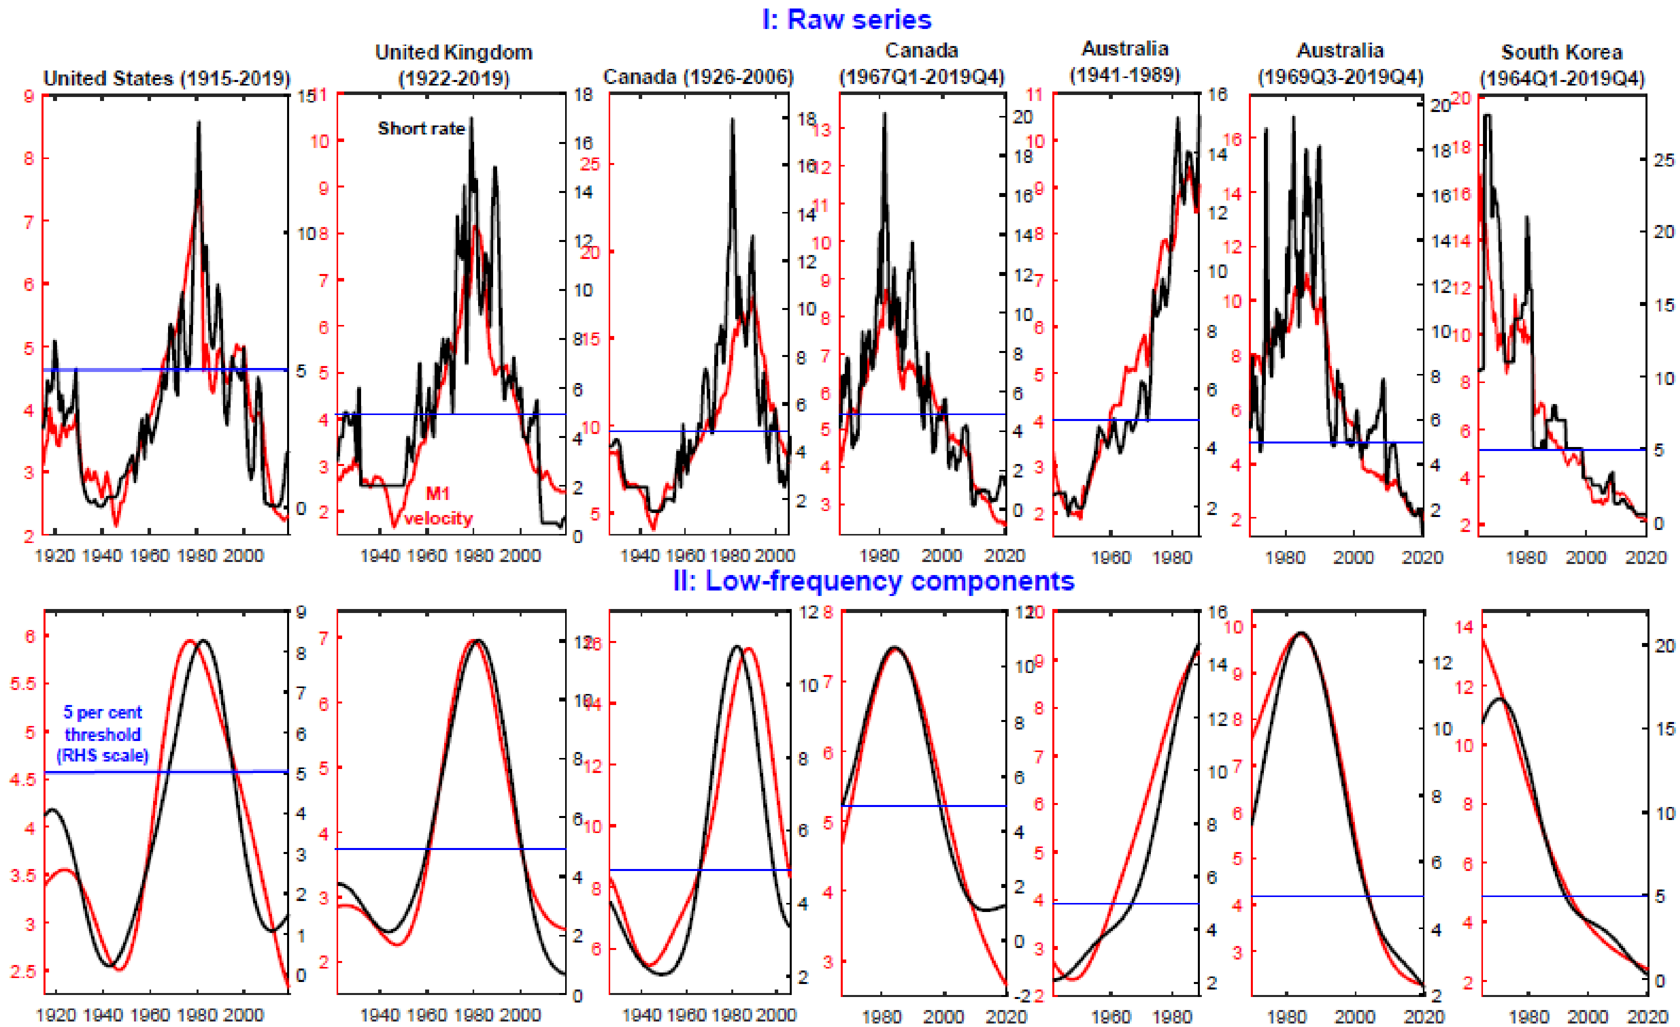

Figure 4 Informal evidence on the possible presence of non-linearities at low interest rates

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, 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. It is clear that the econometric uncertainty (captured, e.g., by the distance between the 5th and the 95th percentiles) is at least one order of magnitude larger (and likely more) than the theoretical imprecision captured by the point estimates of the lower and upper bounds.

The point estimates for the welfare cost of a steady state interest rate equal to 5 percent are close to 0.2 percent of consumption for the United States and Japan, about 0.3 percent for Canada while it is about 0.4 percent of consumption for the United Kingdom. The estimate for the US is the same as the one reported by Lucas (2000) when using the semi-log specification (10), and almost one order of magnitude above the estimate of Ireland (2009) of 0.037 percent. As mentioned above, the main reason for the discrepancy is the small value for the semi-elasticity obtained by Ireland, based on a different monetary aggregate. That explains most of the difference: notice from the Figure that the semi-log, used by Ireland, does imply a cost of only 0.15. The Selden-Latane implies an additional 0.05 percent.

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), in that the log-log specification delivers substantially higher costs, about 0.6 percent, than the 0.2 percent of the Selden-Latane.[28] However, the difference between the log-log and the other two specifications is much higher for the USA than for Japan and the UK and, to a lesser extent, than Canada. This 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 about 5 percent interest rate, for Canada at about 10 percent and for the United States they cross at almost 20 percent (not depicted).

Although the interest rate in the countries discussed so far was always positive, the experience of some European countries suggests that the standard assumption of a zero lower bound may not be appropriate. Thus, we believe that allowing for

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

a negative lower bound cannot be completely ruled out. To explore that possibility, 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 percent, that corresponds to setting $b = 0.15$ and $a = 2$. We only report the results 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 it can be seen, the welfare cost of a 5 percent interest rate are twice as large as in the case of a zero lower bound: 0.4 percent of consumption for the Selden-Latane and 0.3 percent of consumption for the semi-log. However, for the log-log case, the increase is larger, it goes from 0.6 to 1.5 percent of consumption.

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 only consistent 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 by the different estimated parameters only. For the other two scenarios, we follow the strategy adopted for estimation, and let $r_t^m = -a + br_t^b$, where we set $b = 0.15$ and consider two different values for the constant, $a = 1$ and $a = 2$.

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 - top panel - the welfare costs of a 5 percent interest rate is about 0.5% of consumption for the Euro area and Switzerland, and a bit smaller for Sweden, more than double the ones for the United States, Canada and Japan. This is 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 percent of consumption for the Euro area and Switzerland, and somewhat smaller for Sweden, as before. In the final case we consider, that is required to accommodate the experience of Switzerland, the cost can be almost $1,4$ percent of consumption for the Euro area, 1.2 percent for Switzerland and almost 1 percent 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 percent. However, the numbers are substantially larger than in Figure 5.a, im-

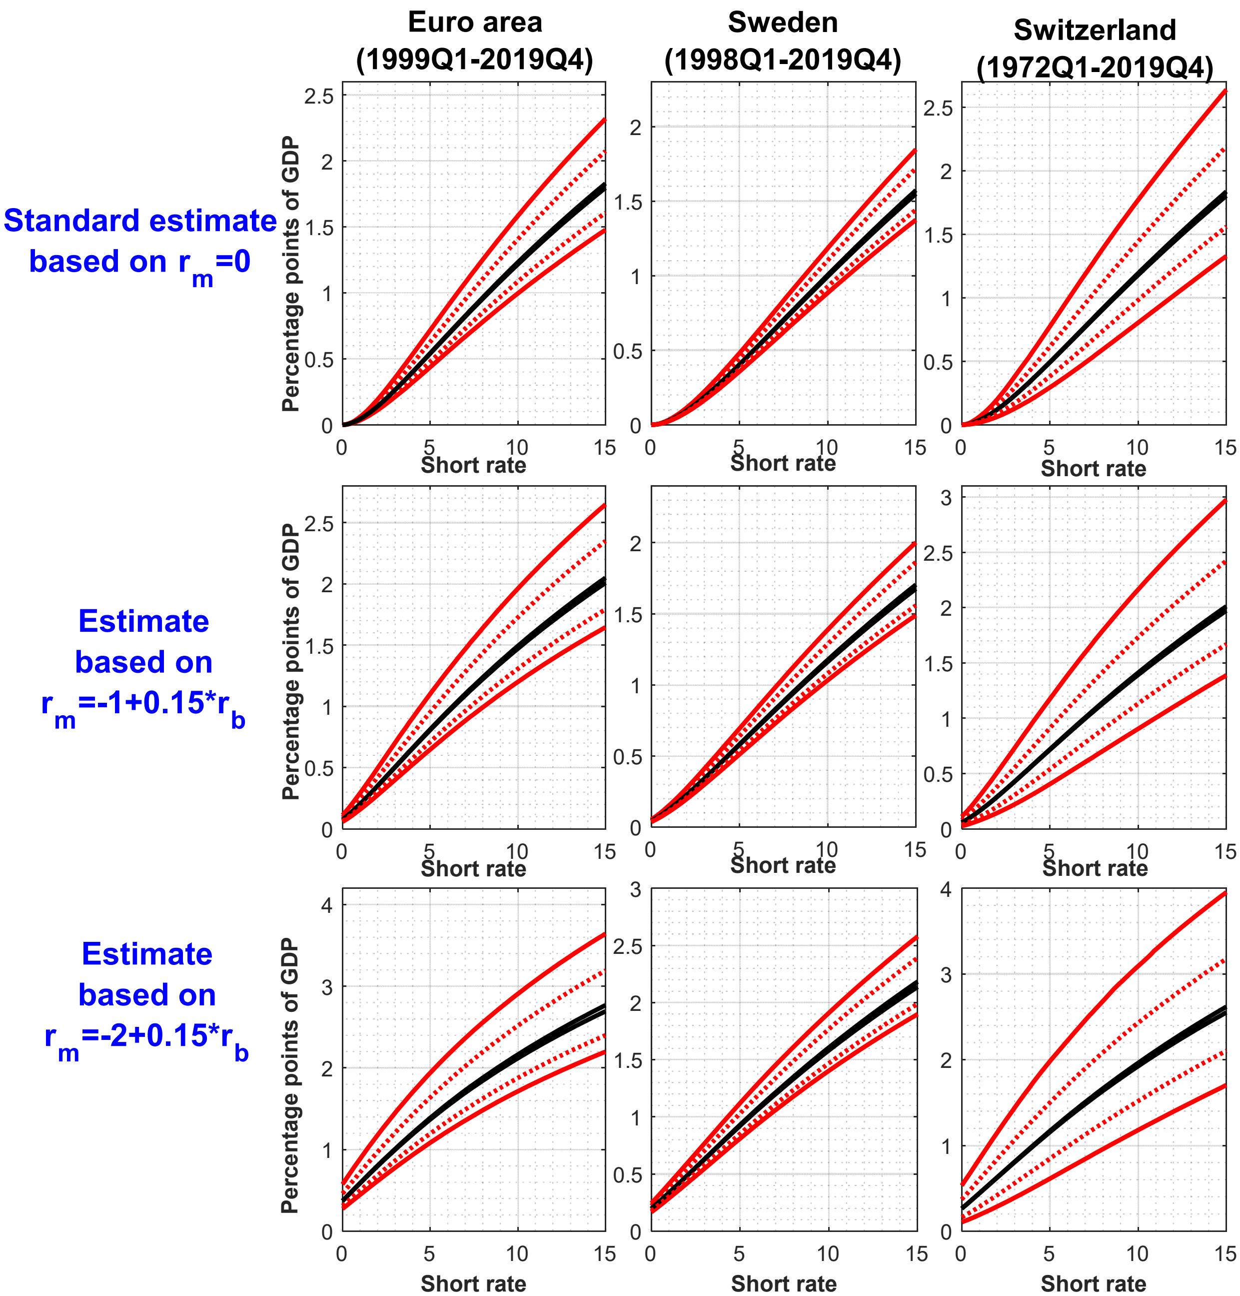

Figure 5.a Estimated welfare cost functions based on the Selden-Latane' specification for countries with negative interest rates (point estimates of the lower and upper bounds, 5-16 percentiles of the lower bounds, and 84-95 percentiles of the upper bounds of the bootstrapped distributions)

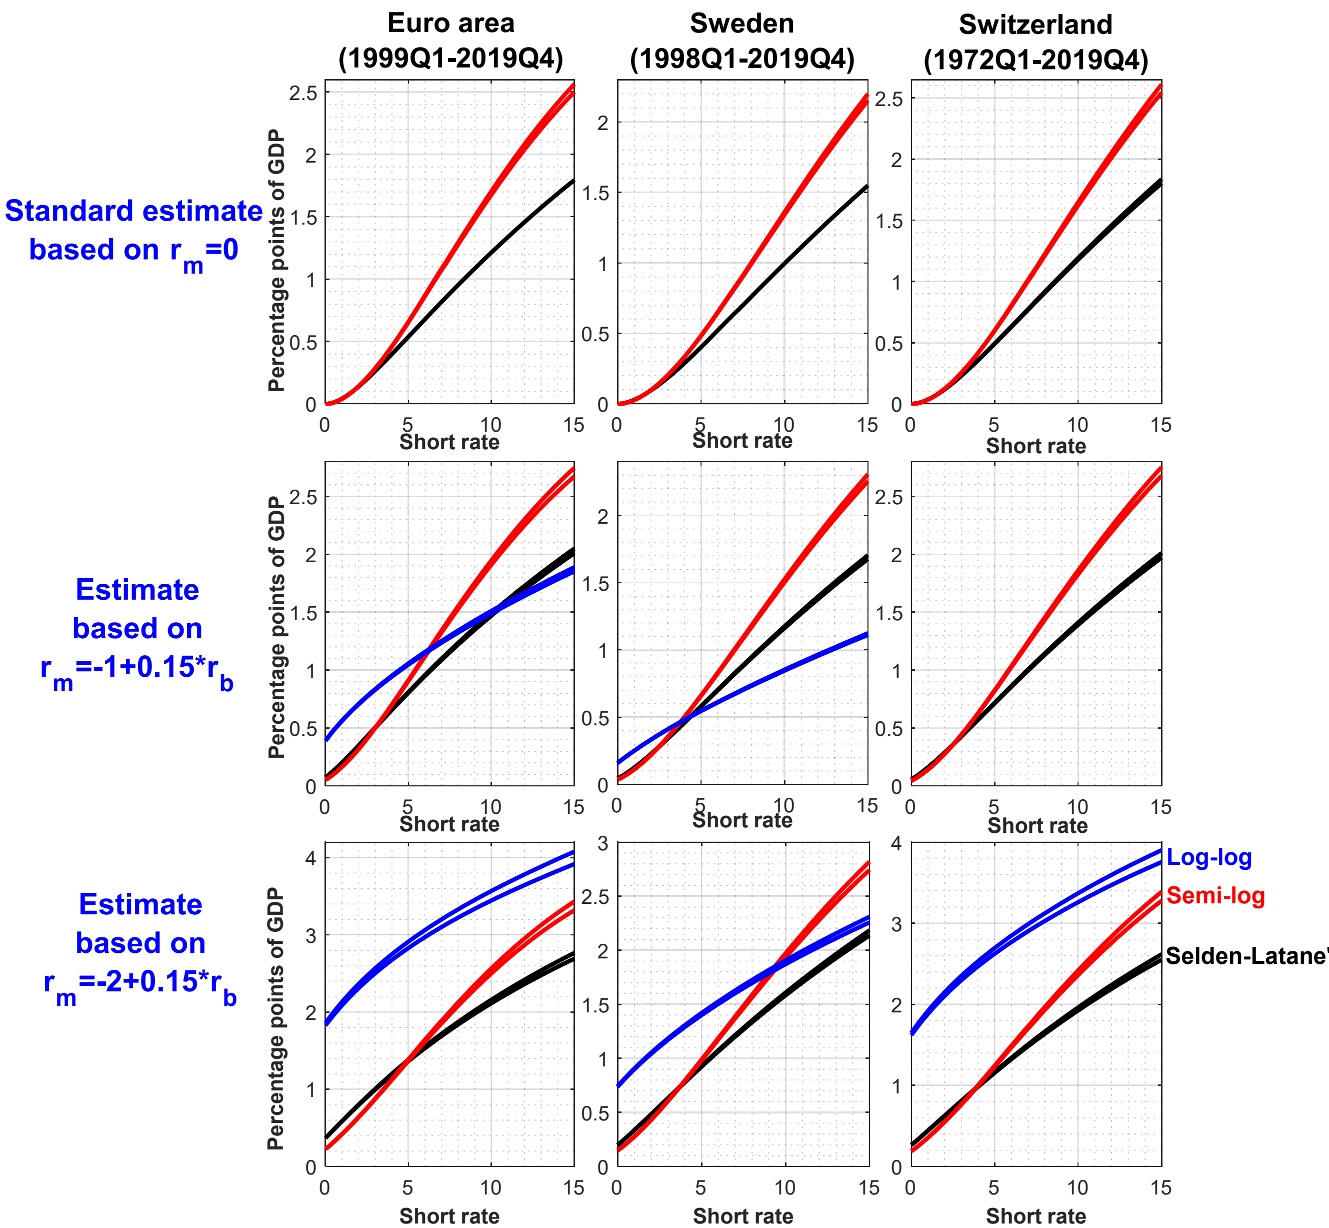

Figure 5b Point estimates of the welfare costs of inflation produced by alternative specifications

plying a larger interaction of the functional form with the assumption about the true lower bound on the short-term interest rates. For example, the log-log specification implies very large welfare costs, of about 2.8 and 2.6 percent of consumption for the Euro area and Switzerland and about 1.5 percent of consumption for Sweden when the lower bound is assumed to be -2.4 percent.

To summarize, we now describe the two most extremes scenarios. In all cases, we report the welfare cost of a 5 percent nominal interest rate. Our lowest set of estimates correspond to the first set of countries (Canada, Japan, USA and UK) for which the combined effect of the estimated parameters and the assumption of a zero bound on interest rates imply a range of estimates between 0.15 and 0.2 percent 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 percent. In this case, the estimated cost is 1.5 percent of consumption. Our largest set of estimates correspond the 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 are between 0.4 and 0.5 percent 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 percent of consumption.

Our estimates show that under the benchmark scenario that assumes the Selden-Latané and a zero lower bound, the welfare cost of inflation evaluated at a 5% average nominal interest rate are between 0.2% and 0.5% of consumption, depending on the country, which is not a negligible number. However, they also show that they depend critically on the functional form and the underlying assumption regarding the true lower bound on the short term interest rates. The log-log specification, that cannot be clearly rejected in some countries with specific assumptions about the true lower bound, delivers much higher estimates, particularly when the true lower bound is assumed to be lower. For example, the log-log specification cannot be rejected for the USA using both Johansen's and Wright's tests when $a = -2$. In this case, the welfare cost at 5% is estimated to be 1.5 percent. The same happens for Switzerland, where the cost is estimated to be 2.5 percent.

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 (e.g.) 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[29] for our dataset, based on any of the three money demand specifications. Only in two instances, Denmark, and Japan based on the Selden-Latané specification, the tests detect evidence of instability.[30]

Overall, there is very little evidence of a break in the real money demand relationship derived from the theory. This is reassuring on itself, but also in reference to the issue raised by Ireland and that has prevailed the discussion in the United States, related to 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 in 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.

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).[31] Based on this argument, at

[29] 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 to the tests for stability in the cointegration vector and loading coefficients.

[30] This is in line with the evidence in Benati *et al.*'s (2021) Section 6.2. The main finding there 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 to -0.41 and -0.74 for Japan.

[31] 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 form, and only when the inequality is

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.

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

		Based on samples with short rate:	
		below 5 per cent	*above 5 per cent*
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 per cent:*
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.

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. The implication is that it should not be possible to reliably estimate money demand functions (and therefore the welfare costs of inflation) based on aggregate time series data, as only the use of micro data allows to meaningfully capture the non-linearities associated with the cost of participating in financial markets.

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, and then a decrease), break tests which 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 any of these cases the short rate had been 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 it to be 5%.

Before delving 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 nonlinearities for five countries for which both sub-samples with the short rate above, and respectively below 5% are 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.[32] The top row shows the raw data for M1 velocity and a short rate, whereas the bottom row shows the low-frequency components of the two series.[33] The evidence in the figure speaks for itself, and it provides nearly *no support* to the notion that velocity—

satisfied it will 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.

[32]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.

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

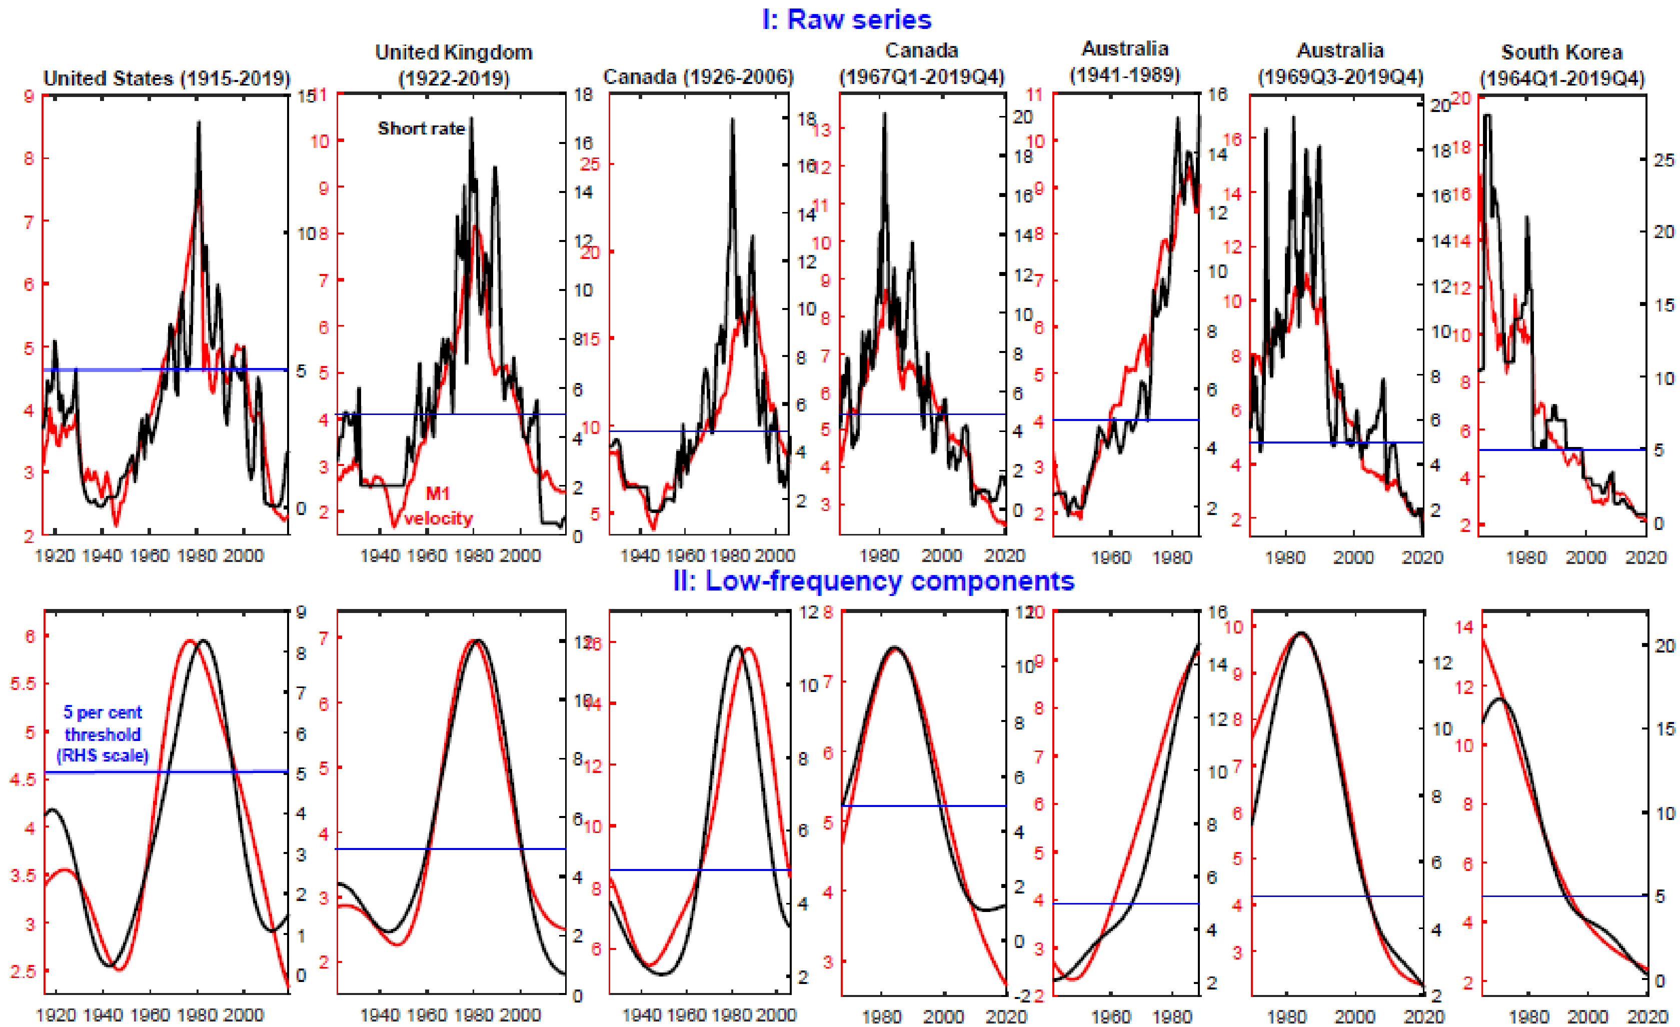

Figure 6 Informal evidence on the possible presence of non-linearities at low interest rates

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. Counter-examples 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 flatter at low interest rates.[34]

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 being shown in black and red, respectively.[35] (The sub-samples with the short rate below and above 5% are reported in Table 6.) The panels also show an 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 should also be flatter than the one with the black dots, and compared to that it should be rotated upwards and to the left towards the horizontal red line.

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.[36] For Canada, in line with Figure 6, 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 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

[34] This is in line with Summers' (1991) point that the most convincing type of evidence, and the one that, historically, had the most impact in terms of changing the profession's views, is simple evidence based on either raw data, or data that have been subjected to very simple manipulations.

[35] 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.

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

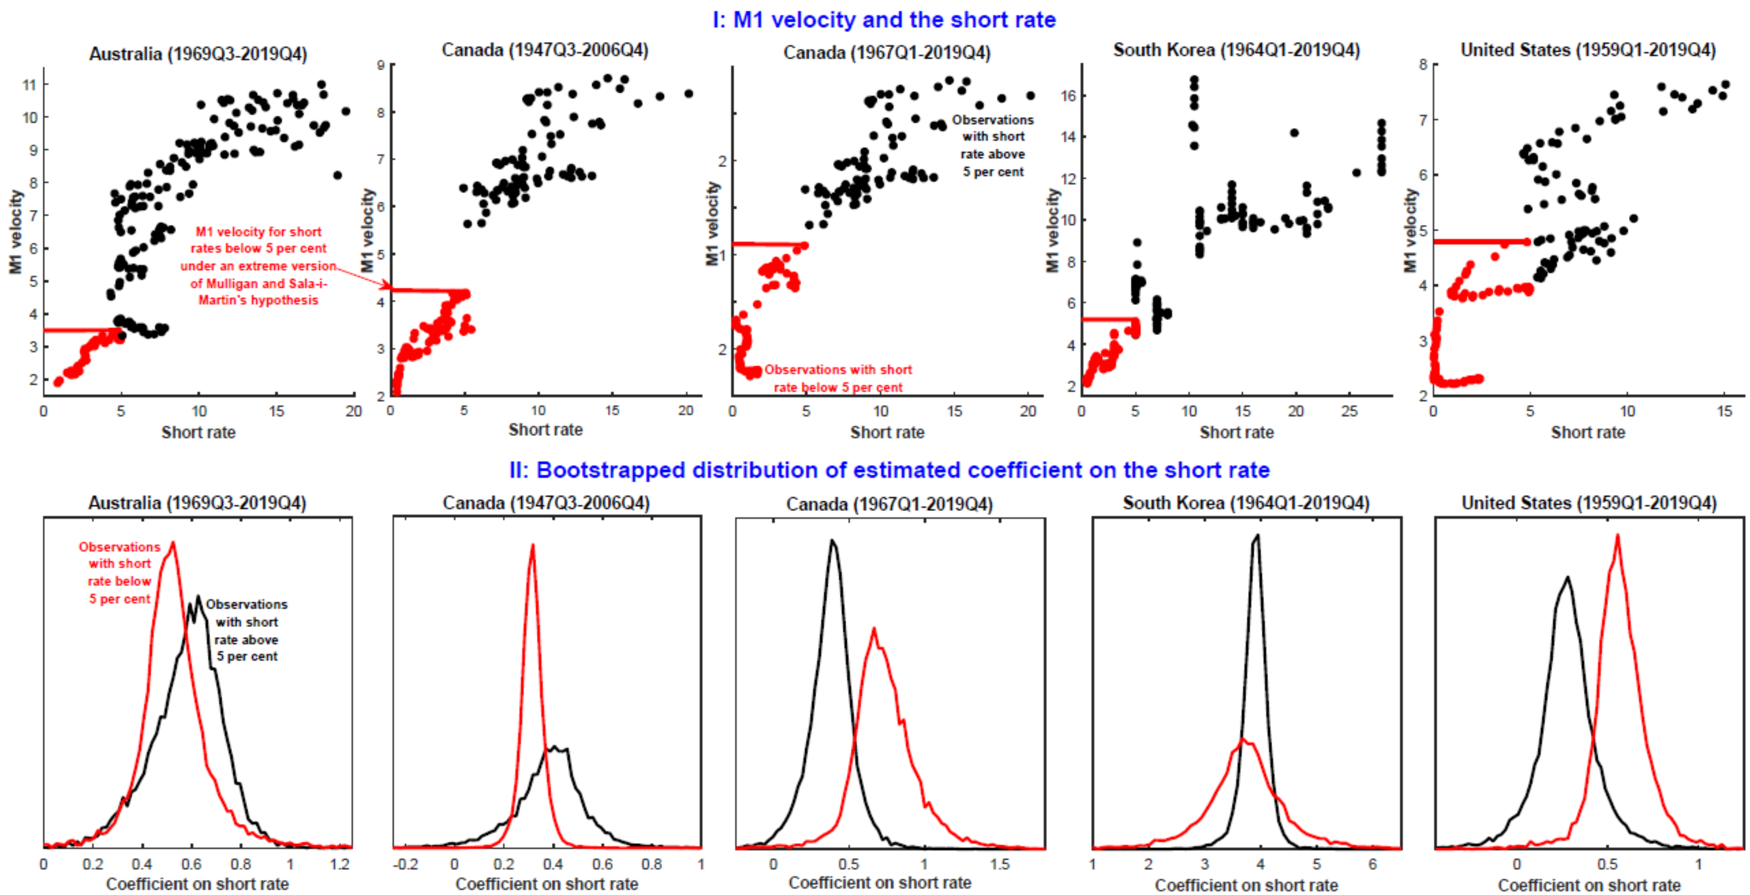

Figure 7 M1 velocity and short-term nominal interest rates: observations with the short rate above and below 5 per cent (quarterly data)

rates. Finally, evidence for the U.S. is idiosyncratic, with the observations below 5% clustered in two separate loops,[37] but once again, in no way does it suggest that the demand curve may be flatter at low interest rates.

The second row of Figure 7 reports the econometric evidence, by showing, for any of the sub-samples, 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 (11) , which is our benchmark specification[38]

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, below 5%, money demand curves may be flatter. 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 7).

6.3 Spurious nonlinearity 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,[39] 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 interest rates. The same argument obviously holds if the true specification is the semi-log.

This can be illustrated as follows. With the true money demand specification

[37] This is partly due to the fact that, as mentioned in footnote 36, we had to eliminate the period 1991Q4-2000Q4.

[38] The methodology we use is standard. 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., 1991, 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).

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

being described by (11), estimating the log-log specification (9) produces the following theoretical value of the *estimated* elasticity

$$\frac{d\ln\left(\frac{M_t}{P_t y_t}\right)}{d\ln r_t} = -\frac{\phi r_t}{a^3+\phi r_t}, \tag{20}$$

which tends to -1 for $r_t \to \infty$, but tends to 0 for $r_t \to 0$ (in fact, for r_t=0, it is exactly equal to 0). By the same token, if the true specification is of the semi-log type, estimating a log-log specification produces the following theoretical value of the estimated semi-elasticity

$$\frac{d\ln\left(\frac{M_t}{P_t y_t}\right)}{d\ln r_t} = -\gamma r_t.$$

which tends to -∞ for $r_t \to \infty$, tends to 0 for $r_t \to 0$, and it is exactly equal to 0 for r_t=0. The implication is that in either case, estimating a log-log specification produces entirely spurious evidence of a lower (semi) elasticity at interest rates approaching zero.

In fact, in each single one of the specifications estimated by Mulligan and Sala-i-Martin (2000) (as well as by Attanasio, Guiso, and Jappelli (2002)) the interest rate entered in logarithms.[40] 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 by setting the nominal interest rate 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 of the US data during the XX century, computed 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 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 argues, the estimated cost is only 0.04% of consumption. A distinct feature of the

[40]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.

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 US, analyzed in detail in Lucas and Nicolini (2015), that provides a unified stable behavior for the US 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 Latane (1960). The Selden-Latane 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 last years. 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-Latane 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, imply 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-Latane is the preferred specification overall. In all countries this functional form performs better in the cointegration test, with the single exception of one of the samples for Canada and for one of the tests (Johansen's) that we run. In addition, in the tests that compares this functional form with the semi-log, it dominates in all cases. Thus, under the assumption that the lower bound is indeed zero, Ireland argument that a real money demand that has a finite value at the lower bound dominates the log-log, that has an asymptote, is correct.

When using this functional form and assuming a zero lower bound, the welfare cost for the Unites 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 consistent with the experiences of the European countries, 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, of 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, Elsevier for the Society for Economic Dynamics, vol. 32, pages 206-226, April.

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(April)), 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 Model", *Quarterly Journal of Economics*, 66, 545-372.

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(November), 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 Hyperinflation 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 SPN/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 Cointegration Rank in Vector Autoregressive Models", *Econometrica*, 80(4), 1721-1740.

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

Coibion, Olivier. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise their Inflation Targets in Light of the ZLB?", with Yuriy Gorodnichenko and Johannes Wieland, 2012, 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(29), 29-47.

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): “Cointegration 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 Christina D. Romer and David 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.

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

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

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

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, October.

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

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

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

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

Luetkepohl, H. (1991): *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*, 75(4), 569-589.

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, ed., *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.

Summers, L.H. (1991): "The Scientific Illusion in Empirical Macroeconomics", *Scandinavian Journal of Economics*, 93(2), 129-148.

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

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

Online Appendix for: The Welfare Costs of Inflation

Luca Benati
University of Bern*

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

*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

A Model Solution

The problem of the agent is to maximize (1) in the main text by choosing c_t, n_t, M_t, B_t, and W_{t+1} subject to (3) (4) and (2) . Assume that the function $\theta(n_t, \nu_t)$ is differentiable.

If we let ξ_t, λ_t and δ_t be the corresponding Lagrange multipliers, the first order conditions are given by

$$\beta^t U\left(c_t\right) = P_t\lambda_t + P_t\delta_t \tag{21}$$

$$P_t\lambda_t\theta_n(n_t, \nu_t)z_t = M_t\delta_t \tag{22}$$

$$\xi_t = \lambda_t(1 + r_t^m) + \delta_t n_t \tag{23}$$

$$\xi_t = \lambda_t(1 + r_t^b) \tag{24}$$

$$\lambda_t = E_t\xi_{t+1} \tag{25}$$

The first-order conditions imply that, as long as $r_t^b - r_t^m > 0$,

$$\lambda_t = \frac{\delta_t n_t}{r_t^b - r_t^m}$$

and from this we obtain

$$P_t\frac{\delta_t n_t}{r_t^b - r_t^m}\theta_n(n_t, \nu_t)z_t = M_t\delta_t$$

or

$$\frac{n_t}{r_t^b - r_t^m}\theta_n(n_t, \nu_t)z_t = \frac{M_t}{P_t} \tag{26}$$

Note also that , as long as $r_t^b - r_t^m > 0$, it ought to be the case that $\delta_t > 0$, which means that the cash-in-advance constraint is binding, so

$$c_t = n_t\frac{M_t}{P_t}$$

which together with feasibility

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

implies

$$\frac{z_t(1 - \theta(n_t, \nu_t))}{n_t} = \frac{M_t}{P_t}$$

Replacing on (26) above

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

Thus, the solution for n_t depends only on the two stochastic processes $r_t^b - r_t^m$ and ν_t. Note, in particular, that it does not depend on z_t, so the theory implies a unit income elasticity.

B The Return on Money

As mentioned above, we assume that the monetary aggregate is the sum of cash and deposits, in fixed proportions. Thus, if we let D_t and C_t be the stock of deposits and cash, then

$$D_t = (1-\gamma)M_t \text{ and } C_t = \gamma M_t,$$

where $\gamma \in (0,1)$. If we let r_t^d and r_t^c be the nominal returns on deposits and cash, then

$$r_t^m = (1-\gamma)r_t^d + \gamma r_t^c$$

As for cash, we assume that $r_t^c = -r^c$, for some non-negative constant r^c. This reflects the chances of cash being lost or stolen, which for simplicity we assume independent of the state of the economy s^t.[43] In discussing the interest rate on deposits, we refer the reader to models similar to the one we discussed above, but enlarged by modeling banks, that create deposits backed with goverenment bonds, as the one in Freeman and Kydland (2000) for example. In those models, due to costs of creating deposits, the interest rate on deposits is a function of the bond rate, so $r_t^d = R(r_t^b)$, with the following properties

$$R(r_t^b) \leq r_t^b,\ 1 \geq R_r(r_t^b) \geq 0, \tag{27}$$

which implies that the spread between the interest rate in bonds and the interst rate in deposits is positive, and that the interest rate on deposits is a non-decreasing function of the interest rate in bonds and that it does not change more that one to one with the interst rate in bonds.[44]

With these two assumptions, the diffference between the interest rate in bonds and the interest rate on money is given by

$$r_t = (1-\gamma)\left[r_t^b - R(r_t^b)\right] + \gamma\left[r_t^b - r^c\right] \tag{28}$$

The maximum problem is well defined for values of r_t such that $r_t \geq 0$. The properties specified in (27) imply that r_t is non-decreasing on r_t^b. Thus, feasibility implies that $r_t \geq r^{\min}$, where $r^{\min}$ satisfies $r_t = 0$. It immediately follows that as long as $r^c > 0$ or $R(0) < 0$, then $r^{\min} < 0$, which implies that r_t^b may indeed be negative.

Kurlat convincingly argues that the function $R(r_t^b)$ is linear, using micro data from the United States, where the slope is very precisely estimated, while the constant is essentially zero. On the other hand, Alvarez and Lippi estimate the return on cash $r^c = -0.02$, using survey data. As a consequence, we use in the paper a liner return for money of the form

$$r_t^m = -a + br_t^b.$$

Our benchmark case sets $a = b = 0$, which is the standard assumption in the literature. To account for the countries that experienced negative rates, we also consider

[43] See Alvarez and Lippi (2009) for some survey-based evidence on this cost.

[44] Kurlat (2019) presents ample evidence for these two properties using USA data.

two more cases. In both we set $b = 0.15$, following the findings in Kurlat. We then explore two alternative values for the constant $a \in \{1, 2\}$. The first value implies that the lower bound on r_t^b is given by

$$r_t^b - r_t^m = r_t^b + a - br_t^b \geq 0$$

or

$$r_t^b \geq -1,17\%,$$

which is lower than the negative rates in the Euro area and Sweden, where rates were always above -1%. However, rates in Switzerland went all the way down to -1.85%, that is why we also explore the case of $b = 2$, which implies a lower bound of -2.35%.

C The Data

Here follows a detailed description of the dataset. All data are from official sources, i.e., either central banks or national statistical agencies.

C.1 United States

For the United States, seasonally adjusted series for nominal GDP and the standard M1 aggregate, and series for the 3-month Treasury bill rate and the 10-year government bond yield, are all from the St. Louis FED's internet data portal, FRED II (their acronyms are GDP, M1SL, TB3MS, and GS10, respectively). The standard M1 aggregate starts in 1959Q1. Before that, the series has been linked to the series M173Q4 in the spreadsheet m1QvMd.xlsx from the Federal Reserve Bank of Philadelphia's real-time data portal, which starts in 1947Q1. Over the period of overlapping the two M1 series are virtually identical, which justifies the linking. The series for Money Market Deposits Accounts (MMDAs), starting in 1982Q4, is from the *Federal Reserve*'s mainframe. A series for currency is from the *Federal Reserve*'s website.

C.2 United Kingdom

For the United Kingdom, a seasonally adjusted series for nominal GDP ('YBHA, Gross Domestic Product at market prices: Current price, Seasonally adjusted £m') is from the *Office for National Statistics*. A seasonally adjusted and break-adjusted stock of M1 is from 'A millennium of macroeconomic data for the UK, The Bank of England's collection of historical macroeconomic and financial statistics, Version 3 - finalised 30 April 2017', which is from the Bank of England's website. Likewise, series for a 10-year bond yield and a Treasury bill rate are all from the same spreadsheet.

C.3 Canada

For Canada, a seasonally adjusted series for nominal GDP ('Gross domestic product (GDP) at market prices, Seasonally adjusted at annual rates, Current prices') is from *Statistics Canada*. Series for the 3-month Treasury bill auction average rate and the benchmark 10-year bond yield for the government of Canada, are from *Statistics Canada*. M1 ('v41552787, Table 176-0020: currency outside banks, chartered bank chequable deposits, less inter-bank chequable deposits, monthly average') is from Statistics Canada. Data on currency are from Statistics Canada ('Table 176-0020 Currency outside banks and chartered bank deposits, monthly average, Bank of Canada, monthly').

C.4 Australia

Nominal GDP ('Gross domestic product: Current prices, $ Millions, Seasonally Adjusted, A2304418T') is from the *Australian Bureau of Statistics*. The short rate ('3-month BABs/NCDs, Bank Accepted Bills/Negotiable Certificates of Deposit-3 months; monthly average, Quarterly average, Per cent, ASX, 42767, FIRMMBAB90') is from the *Reserve Bank of Australia* (henceforth, *RBA*). M1 ('M1: Seasonally adjusted, $ Millions') is from the *Reserve Bank of Australia* since 1975Q2, and from FRED II (at the St. Louis FED's website) for the period 1972Q1-1975Q1 (over the period of overlapping, i.e. since 1975Q2, the two series are identical, which justifies their linking). 5-and 10-year government bond yields are from the *RBA*. Specifically, they are from the *RBA*'s spreadsheet 'F2.1 Capital Market Yields – Government Bond', which is available at the *RBA*'s website. A quarterly seasonally adjusted series for the 'Unemployment rate, Unemployed persons as percentage of labour force' has been computed by taking averages within the quarter of the corresponding monthly series from the *Australian Bureau of Statistics* (the series' code is GLFSURSA).

C.5 Switzerland

For Switzerland, both M1 and the short rate ('Monetary aggregate M1, Level' and 'Switzerland - CHF - Call money rate (Tomorrow next)', respectively) are from the *Swiss National Bank*'s internet data portal. A seasonally adjusted series for nominal GDP ('Gross domestic product, ESA 2010, Quarterly aggregates of Gross Domestic Product, expenditure approach, seasonally and calendar adjusted data, In Mio. Swiss Francs, at current prices') is from the *State Secretariat for Economic Affairs* (SECO) at https://www.seco.admin.ch/seco/en/home. A series for the 10-year government bond yield is from the St. Louis FED's internet data portal, FRED II (the acronym is IRLTLT01CHM156N).

C.6 Sweden

For Sweden, a seasonally adjusted series for nominal GDP ('BNPM - GDP at market prices, expenditure approach (ESA2010) by type of use, seasonally adjusted current prices, SEK million.') is from *Statistics Sweden*. Series for M1 and the 3-month Treasury bill rate ('Money supply, notes and coins held by Swedish non-bank public, M1 (SEK millions)' and 'Treasury Bills, SE 3M', respectively) are from *Statistics Sweden*. A series for the 10-year government bond yield is from the St. Louis FED's internet data portal, FRED II (the acronym is IRLTLT01SEM156N).

C.7 Euro area

For the Euro area, all of the data are from the European Central Bank.

C.8 Denmark

For Denmark, M1 ('Money stock M1, end of period, Units: DKK bn.') is from Denmark's central bank. Nominal GDP ('B.1GF Gross domestic product at factor cost, Seasonally adjusted, Current prices, 1-2.1.1 Production, GDP and generation of income (summary table) by seasonal adjustment, price unit, transaction and time, Units: DKK mio.') and rwal GDP ('B.1*g Gross domestic product, real, Seasonally adjusted, 2010-prices, real value, Units: DKK mio.') are from Statistics Denmark. The central bank's discounrt rate is from the central bank's website.

C.9 South Korea

For South Korea, all of the data are from the central bank: nominal and real GDP ('10.2.1.1 GDP and GNI by Economic Activities (seasonally adjusted, current prices, quarterly), Gross domestic product at market prices(GDP), Bil.Won' and '10.2.2.2 Expenditures on GDP (seasonally adjusted, chained 2010 year prices, quarterly), Expenditure on GDP, Bil.Won' respectively); M1 (''1.1.Money & Banking (Monetary Aggregates, Deposits, Loans & Discounts etc.), Seasonally Ajusted M1(End of), Bil.Won since 1969Q4; Before that: 1.1.Money & Banking (Monetary Aggregates, Deposits, Loans & Discounts etc.), M1(Narrow Money, End Of), Bil.Won, adjusted via ARIMA X-12); and the central bank's discount rate.

C.10 Japan

A series for the discount rate is from the Bank of Japan. A seasonally adjusted series for nominal GDP is from the Economic and Social Research Institute, Cabinet Office, Government of Japan. A seasonally adjusted series for M1 has been constructed based on MA'MAM1NAM3M1MO ('M1/Average amount outstanding/money stock') and

MA'MAM1YAM3M1MO ('M1/Percent changes from the previous year in average amounts outstanding/Money Stock'), both from the Bank of Japan.

C.11 Hong Kong

For Hong Kong, the HIBOR (Hong Kong Inter-Bank Offered Rate) is from the Hong Kong Monetary Authority (HKMA). M1 ('M1, Total, (HK$ million)') is from HKMA, and it has been seasonally adjusted via ARIMA X-12. Nominal GDP ('GDP, HK$ million, From: Table031: GDP and its main expenditure components at current market prices, National Income Section (1)1,') is from Hong Kong's Census and Statistics Department. It has been seasonally adjusted *via* ARIMA X-12.

D Why We Do Not Use Divisia Aggregates

Throughout the entire paper we work with 'simple-sum' M1 aggregates. In this appendix we briefly discuss why we have chosen to ignore Divisia indices. A first problem is that, to the very best of our knowledge, such indices are only available for the U.S. (from the *Center for Financial Stability*, henceforth *CFS*) and for the U.K. (from the Bank of England). A second problem is that, for the U.S., the Divisia M1 series constructed by the *CFS* does not feature MMDAs (which are instead included in Divisia M2). This means that although the resulting index of monetary services has been constructed by optimally weighting the underlying individual assets, it suffers from the crucial shortcoming that it is not including a key component of the transaction technology. As a result, although Divisia M1 is in principle superior to the standard simple-sum M1 aggregate, it ultimately suffers from the same shortcoming of not including MMDAs.

So the key question is: What is more important? Including MMDAs, or optimally weighting the underlying assets? Figure D.1 provides evidence on this, by showing the same evidence shown in Figure 2 in the main text of the paper, but this time with velocity being computed based on Divisia aggregates. The figure speaks for itself, and provides *no evidence* of a stable relationship between the velocity of any Divisia aggregate and its opportunity cost (computed based on the user cost series from the *CFS*). In particular, a comparison between the first panel of Figure D.1, and the second panel in Figure 2, clearly shows that, for the purpose of detecting a stable long-run demand for M1 in the United States, the crucial issue is including MMDAs in the definition of M1, rather than computing the aggregate by optimally weighting the underlying assets. So although, in theory, Divisia M1 possesses optimal properties, because of the specific way in which is has been constructed, within the present context such optimal properties are trumped by the fact that, exactly as its simple-sum counterpart, it does not include MMDAs.

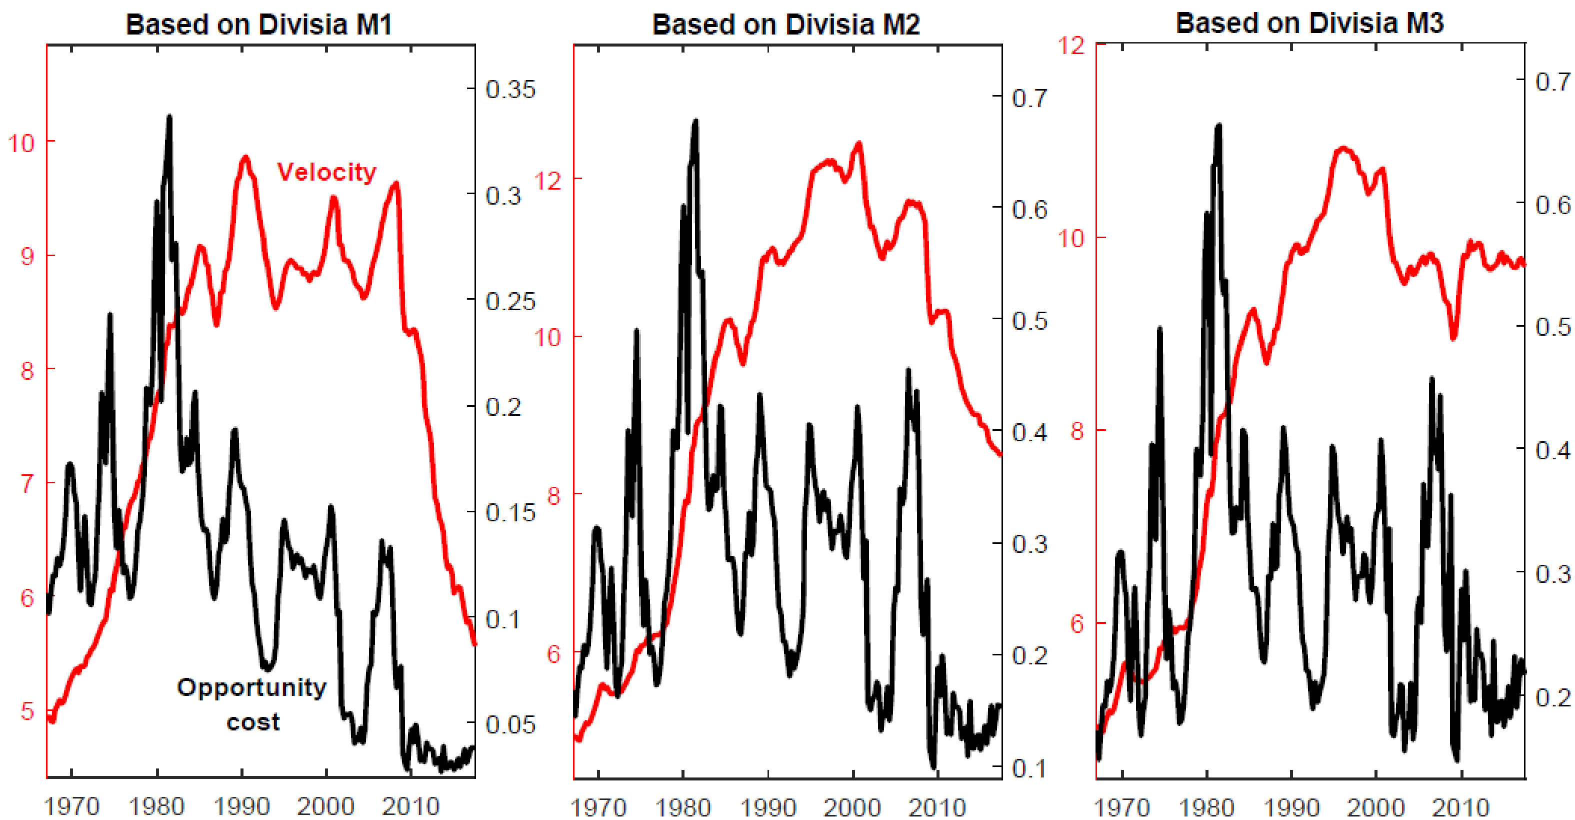

Figure D.1 United States: money velocity based on Divisia aggregates, and the corresponding opportunity costs

Tables for the online appendix

Table A.1*a* Bootstrapped *p*-values for Elliot, Rothenberg, and Stock unit root tests[a]

		M1 velocity				short rate			
		p=2	p=4	p=6	p=8	p=2	p=4	p=6	p=8
United States	1959Q1-2019Q4	0.8633	0.8362	0.9048	0.8764	0.4382	0.2861	0.1903	0.4334
	1959Q1-2001Q4	0.3529	0.2989	0.4112	0.3768	0.3238	0.2756	0.1428	0.2979
United Kingdom	1955Q1-2019Q4	0.9129	0.8806	0.8170	0.8609	0.4010	0.4436	0.5495	0.5951
	1955Q1-2008Q3	0.8187	0.8012	0.7262	0.7740	0.1416	0.1640	0.2386	0.2600
Canada	1947Q3-2006Q4	0.4641	0.6307	0.3987	0.5405	0.2298	0.2466	0.2224	0.3600
	1967Q1-2019Q4	0.9780	0.9725	0.9636	0.9641	0.5020	0.5018	0.4775	0.7151
Australia	1969Q3-2019Q4	0.9766	0.9766	0.9695	0.9418	0.4745	0.3996	0.6217	0.7796
	1969Q3-2008Q4	0.9883	0.9823	0.9834	0.9679	0.3078	0.2719	0.4116	0.5810
Switzerland	1980Q1-2019Q4	0.8918	0.7756	0.8488	0.8296	0.1631	0.2756	0.1845	0.2488
Sweden	1998Q1-2019Q4	0.7086	0.6588	0.7850	0.9379	0.3931	0.5558	0.6600	0.5850
Euro area	1999Q1-2019Q4	0.2303	0.1188	0.0399	0.0074	0.6275	0.3936	0.3316	0.3453
Denmark	1991Q1-2019Q4	0.3557	0.6095	0.4426	0.5297	0.1394	0.0431	0.0288	0.0102
South Korea	1964Q1-2019Q4	0.0205	0.0001	0.0346	0.0289	0.5296	0.5266	0.4200	0.1089
Japan	1960Q1-2019Q4	0.7583	0.7323	0.8419	0.6102	0.2549	0.4236	0.4468	0.4413
Hong Kong	1985Q1-2019Q4	0.6960	0.7432	0.6730	0.6678	0.3566	0.2372	0.1851	0.3773

[a] Based on 10,000 bootstrap replications of estimated ARIMA processes. Tests are with an intercept and no time trend. [b] The short rate has a few negative observations at the end of the sample.

Table A.1*b* Bootstrapped *p*-values for Elliot, Rothenberg, and Stock unit root tests[a]

		Logarithm of:							
		M1 velocity				short rate[b]			
		p=2	p=4	p=6	p=8	p=2	p=4	p=6	p=8
United States	1959Q1-2019Q4	0.9749	0.9550	0.9784	0.9611	0.4835	0.3839	0.4577	0.2198
	1959Q1-2001Q4	0.3054	0.2694	0.3626	0.3169	0.4250	0.3985	0.3023	0.3665
United Kingdom	1955Q1-2019Q4	0.9728	0.9709	0.9468	0.9628	0.7598	0.7917	0.8321	0.8835
	1955Q1-2008Q3	0.8090	0.8838	0.8299	0.8706	0.1484	0.2162	0.3100	0.4094
Canada	1947Q3-2006Q4	0.1103	0.2339	0.1159	0.2931	0.0590	0.0474	0.0229	0.0275
	1967Q1-2019Q4	0.9986	0.9979	0.9960	0.9974	0.4761	0.4924	0.6513	0.7351
Australia	1969Q3-2019Q4	0.9967	0.9988	0.9957	0.9933	0.9297	0.8868	0.9786	0.9822
	1969Q3-2008Q4	0.9978	0.9946	0.9953	0.9923	0.3915	0.3312	0.5893	0.6785
Switzerland	1980Q1-2019Q4	0.9391	0.8697	0.9167	0.9019	[b]	[b]	[b]	[b]
Sweden	1998Q1-2019Q4	0.9332	0.8848	0.9314	0.9712	0.7487	0.7071	0.8188	0.8947
Euro area	1999Q1-2019Q4	0.6377	0.4587	0.2992	0.0952	0.8988	0.9366	0.9539	0.9768
Denmark	1991Q1-2019Q4	0.5492	0.7791	0.6813	0.6186	[b]	[b]	[b]	[b]
South Korea	1964Q1-2019Q4	0.3406	0.2072	0.4245	0.3604	0.9125	0.8993	0.8952	0.7023
Japan	1960Q1-2019Q4	0.9718	0.9639	0.9767	0.9090	0.6177	0.6604	0.6732	0.6732
Hong Kong	1985Q1-2019Q4	0.8316	0.8502	0.8396	0.8061	0.2699	0.2891	0.2719	0.3173

[a] Based on 10,000 bootstrap replications of estimated ARIMA processes. Tests are with an intercept and no time trend. [b] The short rate has a few negative observations at the end of the sample.

Table A.2*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)

			Log-log		
		Semi-	a=0	a=-1	a=-2
Country	*Period*	*log*	b=0	b=0.15	b=0.15
United States	1959Q1-2019Q4	766.1394	756.6280	763.3872	764.4667
United Kingdom	1955Q1-2019Q4	879.6821	877.9350	878.9924	879.3578
Canada	1947Q3-2006Q4	820.2401	807.8379	835.9188	836.6934
	1967Q1-2019Q4	775.0890	767.0845	772.1874	773.4180
Australia	1969Q3-2019Q4	650.7331	656.0624	655.5800	654.8957
Switzerland	1972Q1-2019Q4	547.7844	–[a]	–[a]	542.5159
Sweden	1998Q1-2019Q4	317.3933	–[a]	316.9838	317.3565
Euro area	1999Q1-2019Q4	333.1620	–[a]	333.0533	333.1005
Denmark	1991Q1-2019Q4	404.1328	–[a]	404.0745	404.0309
South Korea	1964Q1-2019Q4	630.9515	633.8825	633.7839	633.4732
Japan	1960Q1-2019Q4	845.3632	850.7677	850.8201	849.3694
Hong Kong	1985Q1-2019Q4	328.0148	325.5701	334.2308	332.9143

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

Table A.2*b* 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=4)

			Log-log		
		Semi-	a=0	a=-1	a=-2
Country	*Period*	*log*	b=0	b=0.15	b=0.15
United States	1959Q1-2019Q4	763.2818	751.3266	758.5811	759.4505
United Kingdom	1955Q1-2019Q4	898.6224	893.7504	895.9872	897.0579
Canada	1947Q3-2006Q4	813.8001	804.9218	832.4689	833.5199
	1967Q1-2019Q4	775.9595	766.4531	772.1997	773.9634
Australia	1969Q3-2019Q4	649.9510	655.1057	654.3638	653.9727
Switzerland	1972Q1-2019Q4	544.4609	–[a]	–[a]	538.9411
Sweden	1998Q1-2019Q4	311.4090	–[a]	312.9552	312.7534
Euro area	1999Q1-2019Q4	326.4400	–[a]	326.5178	326.6725
Denmark	1991Q1-2019Q4	405.6384	–[a]	406.5285	406.1349
South Korea	1964Q1-2019Q4	628.2222	634.6372	635.1622	634.5937
Japan	1960Q1-2019Q4	841.5156	848.6520	846.0627	844.8858
Hong Kong	1985Q1-2019Q4	326.1339	324.9236	335.0398	333.2458

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

Figures for the online appendix

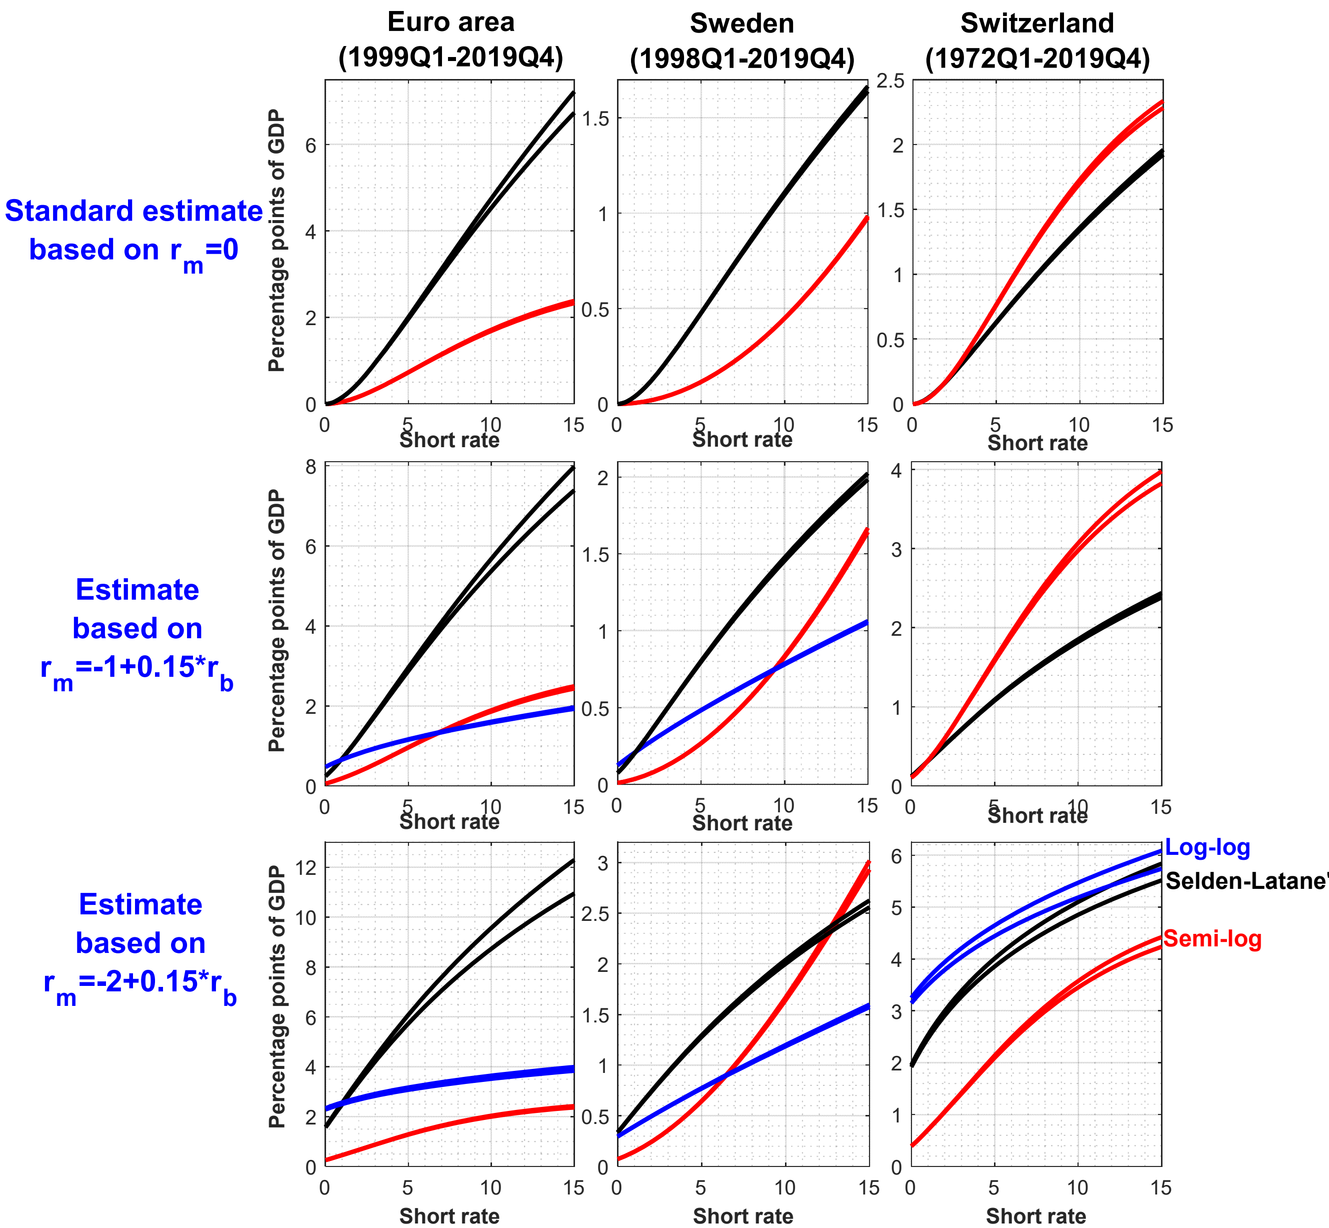

Figure A.1 Estimates of the welfare costs of inflation produced by alternative specifications, based on Wright's (2000) estimator