The Welfare Costs of Inflation Reconsidered*

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 for 12 low-inflation and 5 high-inflation countries, and for Weimar Republic's hyperinflation. Our evidence suggests that, contrary to the implicit assumption in much of the literature, these costs are far from negligible. For the U.S. our point estimates are equal to about one-third of those computed by Lucas (2000), and an order of magnitude larger than those obtained by Ireland (2009). Crucially, the most empirically plausible money-demand functional form points towards sizeable 'upward risks' for these costs, with the 90% confidence interval associated with a 4% nominal interest rate stretching to 0.9 per cent of GDP. The welfare costs of inflation in the Euro area are nearly twice as large as in the U.S., thus suggesting that, *ceteris paribus*, the inflation target should be materially lower. At the peak of the inflation episodes, welfare costs had ranged between 0.6 and 2.3 per cent of GDP for low-inflation countries; between about 4 and 14 per cent for high-inflation ones; and between 25 and 33 per cent for Weimar's hyperinflation.

*We wish to thank participants to the EEA 2024 meeting in Rotterdam for useful comments and discussions. 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

The welfare costs of inflation originating from lack of liquidity satiation—as discussed in the classic work of Bailey (1956), Friedman (1969), Lucas (2000), and Ireland (2009)—tend to be consistently disregarded in the current policy debate. A case in point is the proposal to increase inflation targets in order to decrease the probability that the Zero Lower Bound (ZLB) on monetary policy rates may become binding. Following Blanchard, Dell'Ariccia and Mauro (2010), a vast literature has explored the costs of a binding ZLB on monetary policy rates within cashless economies. Although this literature has documented the benefits that increasing the inflation target delivers to society, such increase also comes at a cost, originating from moving the economy further away from liquidity satiation.

A likely reason for the literature's disregard of these costs is the widespread belief in the absence of stable money-demand relationships, which are a necessary condition for the computation of such costs. The recent work of Lucas and Nicolini (2015) for the U.S., and Benati, Lucas, Nicolini, and Weber (2021) for 38 countries since World War I, has shown however that the notion of instability of the long-run demand for M1 is, in fact, incorrect. On the contrary, the existence of a stable long-run demand for M1 appears to be one of the most robust stylized facts in the entire field of macroeconomics. As discussed by Benati *et al.* (2021), and as we discuss below in detail, crucial issues in being able to identify a stable long-run demand for M1 are (*i*) imposing unitary income elasticity, which is implied by theory and it is in fact supported by the data; and (*ii*) adopting the correct functional form for the demand for real money balances.

A second reason for the literature's consistent disregard of the welfare costs of inflation originating from lack of liquidity satiation is the (by now) long-standing tradition of evaluating monetary policies within cashless economies, with monetary aggregates playing no role whatsoever. As we show below, however, there is no sense in which modern advanced economies are becoming cashless. For countries such as the U.S., the U.K., Sweden, and Switzerland, or for the Euro area, the ratio of M1 to GDP has fluctuated in recent years between 40 and 100 per cent. As we show, this implies that for empirically plausible parameterizations of the demand for money balances, the point estimates of the welfare costs of inflation are in general non-negligible. Further, we show that once taking into account of statistical uncertainty, in several cases—notably, the U.S.—the most plausible functional form for the demand for real M1 balances points towards material 'upward risks' for the welfare costs of inflation, in the sense that (e.g.) 90 per cent confidence intervals for these costs include values that are definitely non-negligible, and in fact often sizeable.

A final likely reason for the literature's disregard for these costs is that they are often thought to be negligible. For example, Feldstein (1997, p. 145), in his review of the costs of inflation, characterized the welfare costs originating from lack of liquidity satiation as *'small relative to the other effects that have been discussed in this paper'*.

As mentioned, we show that this presumption is in general incorrect.

The number of studies presenting empirical estimates of the welfare costs of inflation originating from lack of liquidity satiation is quite surprisingly limited. Further, previous studies are sub-optimal along several dimensions. First, several of them (first and foremost Lucas, 2000) are based on calibrated, as opposed to estimated models. Second, previous estimates are uniformly based on functional forms for the demand for real money balances that Benati *et al.* (2021) have shown to be empirically highly implausible.[1] As we show, in general the money-demand specification does matter for the computations of welfare costs along several dimensions. Third, to the best of our knowledge the only previous analysis of the welfare costs for hyperinflations is Barro (1972). Last, but not least, the period following the outbreak of the financial crisis has provided previously unavailable information on the behavior of money demand at very low interest rates, which as discussed by Lucas (2000) should play a crucial role in the determination of the welfare costs of inflation. In principle this should allow us to obtain better and more precise estimates of these costs.

In this paper we revisit the estimation of the welfare costs of inflation for twelve low-inflation and five high-inflation countries, and for Weimar Republic's hyperinflation. In our analysis we follow the tradition of considering the most liquid monetary assets, which include cash and transactional deposits. We abstract from a detailed discussion of the demand for each of the components, an issue recently addressed by Kurlat (2019).[2]

Our evidence suggests that, contrary to the explicit or implicit assumption in much of the literature, these costs are often far from negligible. For the U.S. our point estimates are equal to about one-third of those computed by Lucas (2000), and an order of magnitude larger than those obtained by Ireland (2009).[3] Crucially, the most empirically plausible money-demand functional form points towards sizeable 'upward risks' for these costs, with the 90% confidence interval associated with a 4% nominal interest rate stretching to 0.9 per cent of GDP. At the peak of the inflation episodes, welfare costs had ranged between 0.6 and 2.3 per cent of GDP for low-inflation countries; between 4 and about 14 per cent for high-inflation ones; and between 25 and 33 per cent for Weimar's hyperinflation.

Our evidence suggests that ignoring money in analyzing optimal monetary policies can be seriously misleading. For instance, Coibion, Gorodnichenko and Wieland (2012)[4] make a compelling argument against increasing the inflation target in coun-

[1] Notable cases in point are Lucas (2000) and Ireland (2009).

[2] 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.

[3] For a steady-state interest rate of 5 per cent, Lucas (2000) computed the cost to be around 1.1 per cent of lifetime consumption. Ireland (2009) challenged Lucas' interpretation of the data, and estimated a mere 0.04 per cent of lifetime consumption.

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

tries like the U.S., based on a model with frictions in price-setting and with recurrent, though not very frequent, episodes with the nominal interest rate at the ZLB. Based on their preferred specification they compute the welfare effect of an interest rate of 5 per cent to be close to 0.6 per cent of lifetime consumption. Such an estimate, combining the cost created by price frictions and the probability to be at the ZLB, is of the same order of magnitude of the estimates we obtain for the U.S., and it is in fact smaller than the upper bound of our 90 per cent confidence interval.

In any of the low-inflation countries the demand for M1 as a fraction of GDP at very low, or even negative interest rates has (so far) exhibited no obvious difference compared to its behavior at higher interest rates. These results contrast with those of Mulligan and Sala-i-Martin (2000), who, based on U.S. households micro data, provided evidence that money demand becomes comparatively flatter at low interest rates. We provide a straightforward explanation for Mulligan and Sala-i-Martin's (2000) finding, by showing mathematically that if the true money demand specification is the one that, as shown by Benati *et al.* (2021), is the most empirically plausible at low inflation rates, Mulligan and Sala-i-Martin's (2000) approach automatically produces spurious evidence of a flatter demand curve at low interest rates.

From a theoretical standpoint, following Alvarez, Lippi and Robatto (2019) we construct upper and lower bounds for the welfare costs of inflation. As they show, 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. We extend their results for a quite general sub-class of the models they analyze and compute exact lower and upper bounds for 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 low-inflation countries such as the U.S.

For policy purposes, our main finding is that the welfare costs inflation in the Euro area are nearly twice as large as in the U.S.. This suggests that, *ceteris paribus*, the inflation target should be materially lower.

The paper proceeds as follows. In Section 2 we discuss a family of monetary models for which we derive very tight lower and upper bounds for the welfare cost of inflation using the area under the real money demand curve. In Section 3 we 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. Section 5 presents our 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 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, in 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. We allow for money to pay a nominal return that we denominate r_t^m, which in what follows, in line with the literature, we will set to zero.

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

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

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

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

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

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, since we assume that $r_t^m = 0$, real money demand only depends on the interest rate on 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.

[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.

2.1 The functional form for the demand for real money balances

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,

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

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

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

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

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

both imply a finite level of the demand for real money balances when the interest rate becomes zero. This feature is emphasized by Ireland, who uses (9) 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.[7]

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

2.2 The welfare costs 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. Specifi-

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

cally, 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, r_t) approaches zero.[8] Our bounds can be used for any value of the interest rate.

As we show below, for low-inflation countries the distance between the upper and lower bound is positive, but negligible, so that in the figures the difference between the two is invisible to the naked eye.

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

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

As previously discussed, the welfare cost of inflation, measured as a fraction of consumption, is given by

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

It follows that

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

We now show how the function $\omega^W(r)$ can be bounded above and below using the integral under the money demand curve.[9]

The area under the demand curve is equal to

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

so

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

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

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

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

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

which, using (11), becomes

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

Using the definition in (12),

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

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

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

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

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

Therefore

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

and

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

which imply

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

We 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 (8) to (10), as we show below.

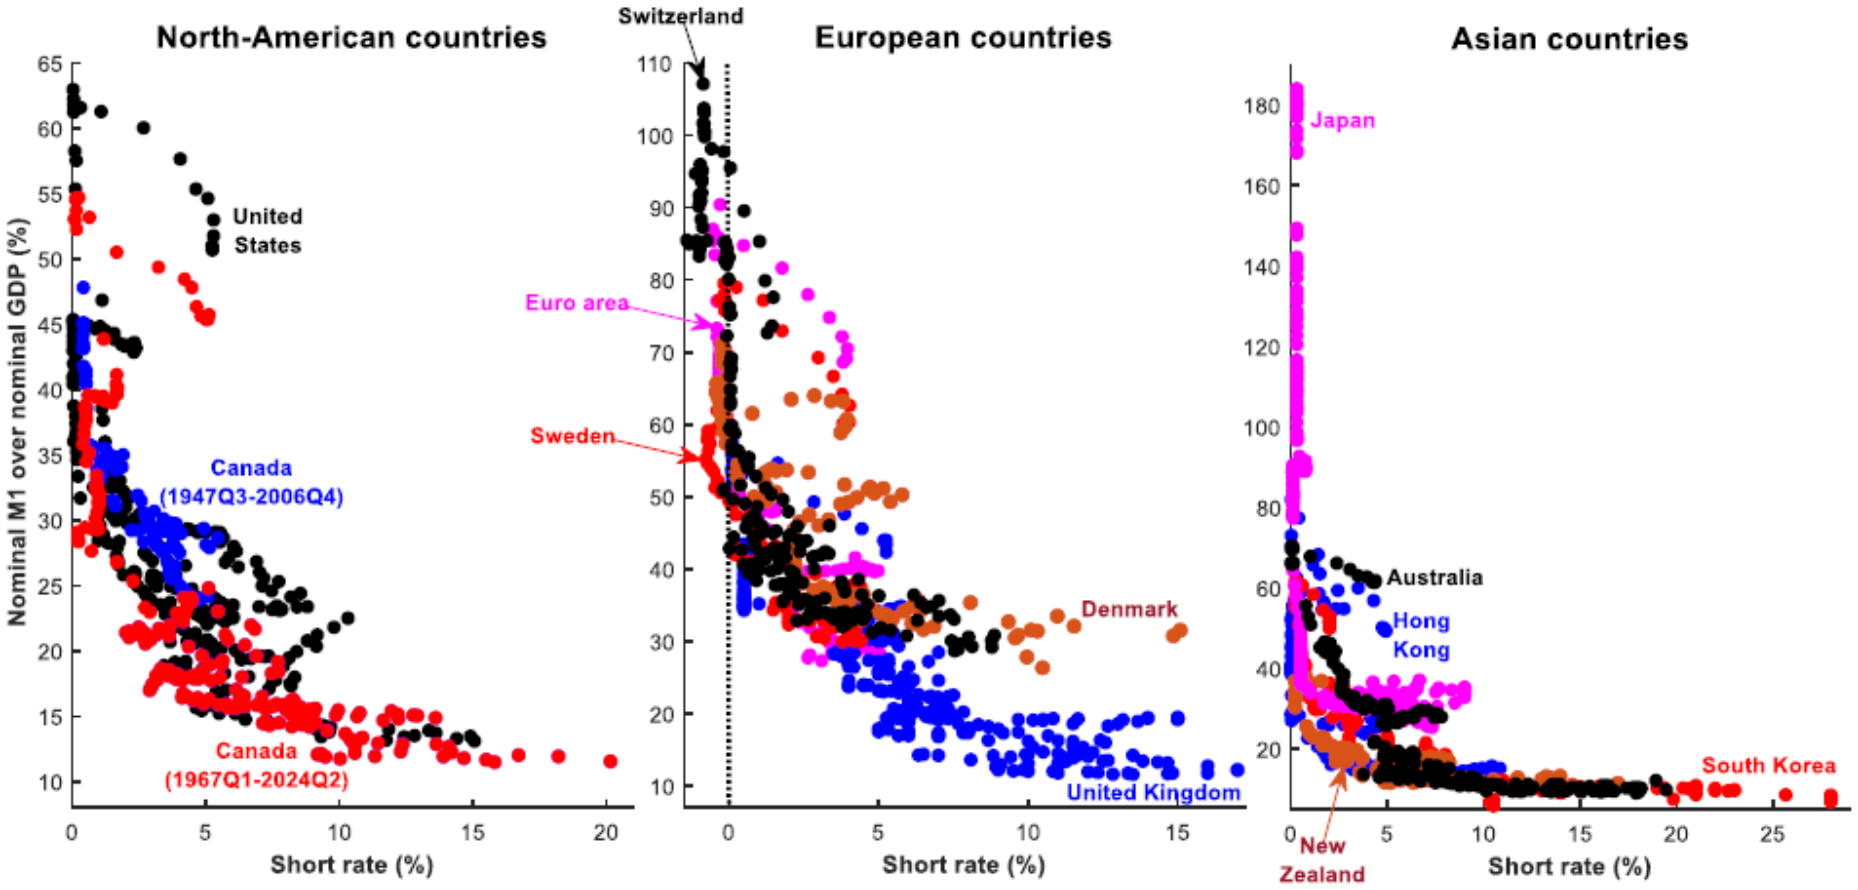


Figure 1 Scatterplots of nominal M1 over nominal GDP against the short rate

3 A Look at the Raw Data

For low-inflation countries we work with quarterly post-WWII data. The series and their sources are described in detail in Appendix B. For all but one country we consider M1 as the relevant monetary aggregate.[10] 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.[11]

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.[12] 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 the behavior of money demand at very low interest rates.

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 extremely useful, and it provides a powerful visual

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

[11] 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.

[12] 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 40 per cent.

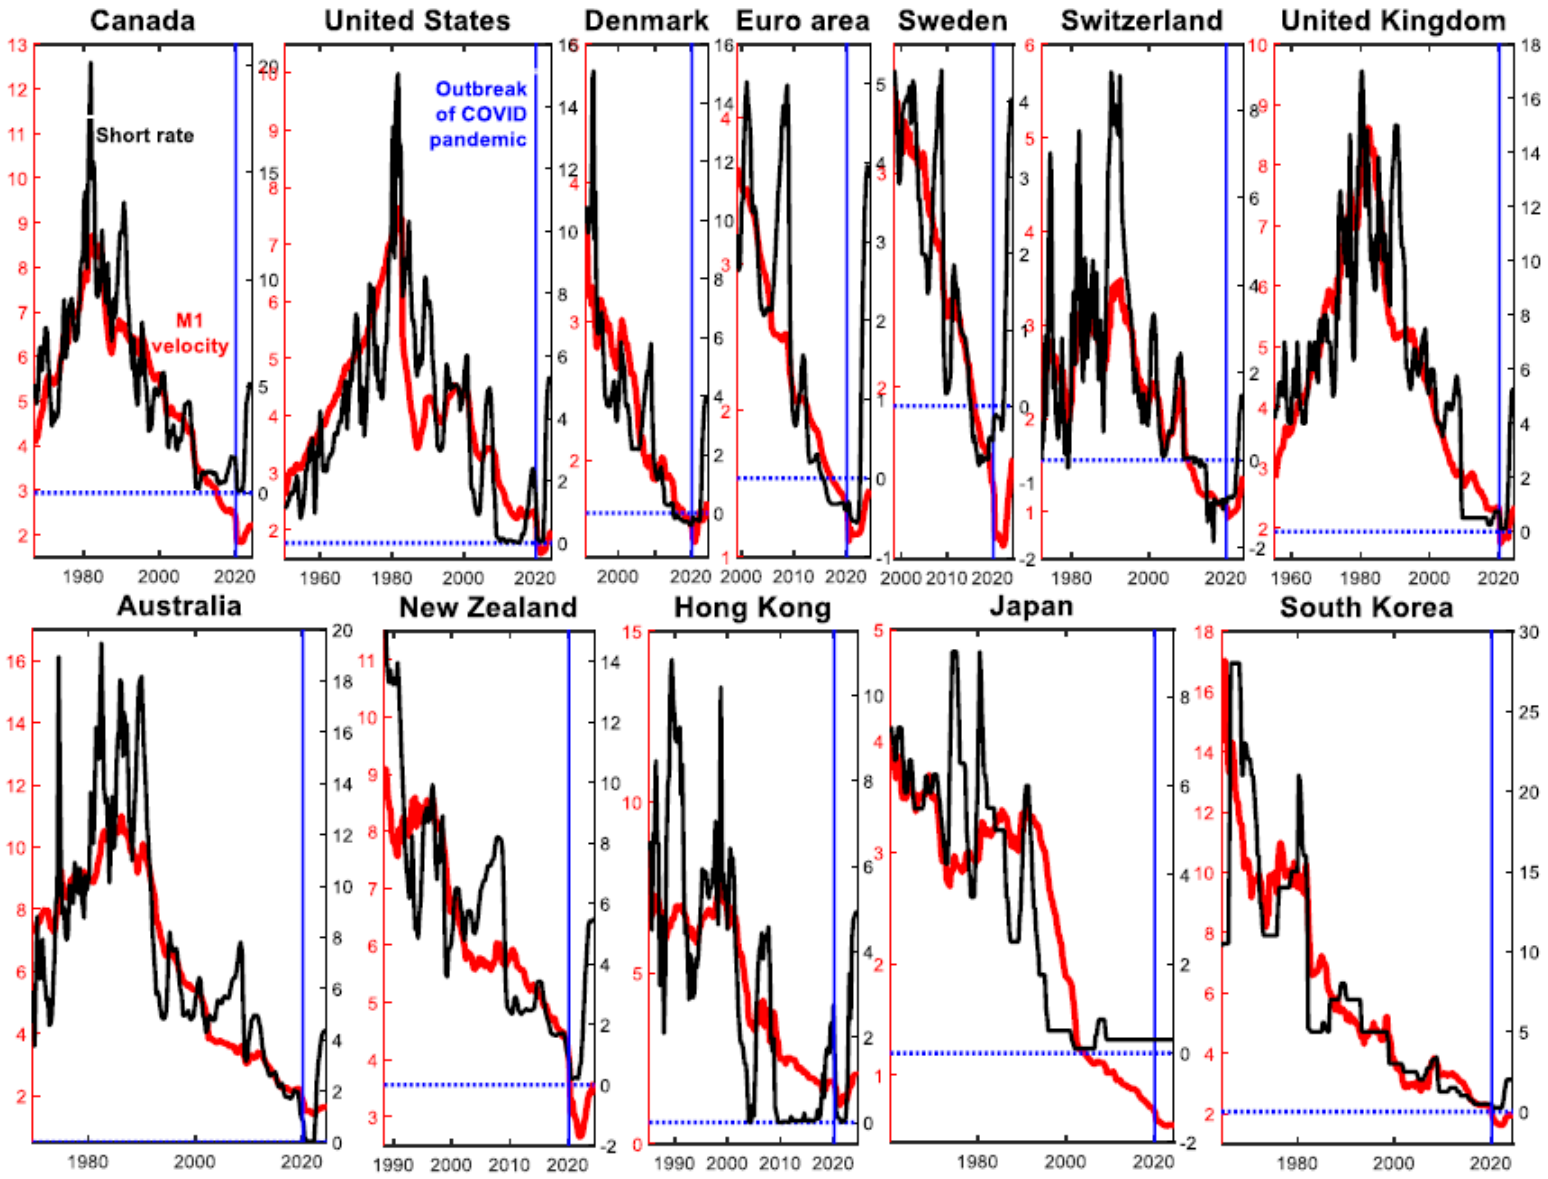


Figure 2 M1 velocity and the short rate

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 Unit root and cointegration properties of the data

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.[13] In short, the null hypothesis of a unit root cannot be rejected for nearly all countries and all series.[14] One possible interpretation of this evidence is that for all countries both series feature *exact* unit roots. A more plausible interpretation is that they are *near* unit roots, which in small samples are statistically indistinguishable from exact unit root processes. This is the case in particular for interest rates, for which a direct implication of them featuring an exact unit root—i.e., that they could literally take *any* value between minus and plus infinity—appears as manifestly absurd. Assuming that the series are near unit root processes, we therefore proceed to test for cointegration based on Wright's (2000) test, which was designed to be equally valid for both exact and near unit root processes. 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.

Table 1 reports, for any of the three money demand specifications discussed in Section 2, the 90% bootstrapped confidence intervals for the second element of the normalized cointegration vector based on Wright's (2000) test. For Canada we have two partially overlapping M1 series that cannot be linked, since they are slightly different. Cointegration is detected near-uniformly across the board: the only two exceptions are the first sample for Canada and Japan based on the semi-log. An important point to stress is that, as we show in the next sub-section, evidence strongly suggests that for low-inflation countries the semi-log is *not* the most plausible money-demand specification.

[13] As discussed in the Appendix, for high-inflation countries and Weimar's hyperinflation we compute the opportunity cost of money as the *maximum* between a short-term nominal interest rate and inflation. The reason is that in all of these cases inflation is highly volatile, and for non-negligible fractions of the samples it is persistently higher than the short rate.

[14] 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).

Table 1 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*	Money demand specification: Selden-Latané	Semi-log	Log-log
		Low-inflation countries:		
United States	1950Q1-2024Q2	[-0.672 -0.435]	[-0.180 -0.036]	[-0.513 0.031]
United Kingdom	1955Q1-2024Q2	[-0.563 -0.366]	[-0.135 -0.079]	[-0.469 -0.137]
Canada	1947Q3-2006Q4	[-0.713 -0.543]	NCD	[-0.512 -0.440]
	1967Q1-2024Q2	[-0.607 -0.410]	[-0.177 0.023]	[-0.570 -0.134]
Australia	1969Q3-2024Q2	[-0.933 -0.704]	[-0.184 -0.128]	[-0.830 -0.630]
Switzerland	1972Q1-2024Q2	[-0.494 -0.314]	[-0.253 -0.172]	–[b]
New Zealand	1988Q2-2024Q2	[-0.913 -0.189]	[-0.173 0.039]	[-0.568 -0.168]
Sweden	1998Q1-2024Q2	[-0.730 -0.246]	[-0.281 -0.097]	–[b]
Euro area	1999Q1-2024Q2	[-0.743 -0.383]	[-0.308 -0.216]	–[b]
Denmark	1991Q1-2024Q2	[-0.261 -0.189]	[-0.125 -0.061]	–[b]
South Korea	1964Q1-2024Q2	[-0.664 -0.424]	[-0.121 -0.053]	[-0.552 -0.480]
Japan	1960Q1-2024Q2	[-0.461 -0.297]	NCD	[-0.553 -0.280]
Hong Kong	1985Q1-2024Q2	[-1.238 -0.649]	[-0.298 -0.158]	[-0.580 -0.160]
		High-inflation countries:		
Argentina	1946-2023	–	–	[-1.056 -0.171]
Chile	1946-2023	–	–	[-0.952 0.357]
Bolivia	1980-2022	–	–	[-0.483 -0.282]
Israel	1982Q1-2024Q2	–	–	[-0.551 -0.290]
Mexico	1985Q4-2024Q2	–	–	[-0.674 -0.394]
		Hyperinflations:		
Weimar Republic	Sep. 1920-Oct. 1923	–	–	[-0.526 -0.349]

[a] Based on 10,000 bootstrap replications. NCD = No cointegration detected.

[b] The last observations for the short rate are either zero or negative.

For high-inflation countries, Benati *et al.* (2021) show that the log-log specification provides the most plausible description of the data. By the same token, Benati (2024) shows that for 20 hyperinflations evidence in favor of the log-log is overwhelming. As for low-inflation countries, the evidence in Benati *et al.* (2021) suggests that the Selden-Latané specification is the most plausible one, but the evidence is less clear-cut. For this group of countries we therefore now turn to a systematic model comparison exercise.

4.2 Which specification do the data prefer?

Since it is not possible to nest the three money demand specifications into a single encompassing one, we proceed as follows. We start from the comparison between the semi-log and the log-log. Intuitively, the comparison between (9) and (8) boils down to whether the dynamics of log M1 balances as a fraction of GDP (i.e., minus log velocity) is better explained by the level of the short rate, or by its logarithm. For each low-inflation country we therefore regress $\ln(M_t/Y_t)$ on a constant, p lags of itself, and p lags of either the level of the short rate or its logarithm. A natural way of interpreting these regressions is the following. Under the assumption that cointegration is indeed there for all countries,[15] 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:

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

We estimate (14) *via* maximum likelihood, stochastically mapping the likelihood surface *via* Random-Walk Metropolis (RWM). The *only* difference between the 'standard' RWM algorithm which is routinely used for Bayesian estimation and what we are doing here is that the jump to the new position in the Markov chain is accepted or rejected based on a rule which does not involve any Bayesian priors, as it uniquely involves the likelihood of the data.[16] So one way of thinking of this is as Bayesian

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

[16] 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

Table 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

		$p = 2$		$p = 4$		$p = 8$	
Country	*Period*	*Semi-log*	*Log log*	*Semi-log*	*Log-log*	*Semi-log*	*Log-log*
United States	1950Q1-2024Q2	932.0394	911.9087	933.4777	912.2283	931.3384	896.3368
United Kingdom	1955Q1-2024Q2	872.4031	869.1347	878.5606	875.5319	864.6897	861.8448
Canada	1947Q3-2006Q4	820.2401	807.8379	813.8001	804.9218	802.7001	794.7403
	1967Q1-2024Q2	742.6498	741.2765	736.8099	735.6754	722.9794	724.0862
Australia	1969Q3-2024Q2	687.7033	674.8213	685.5811	672.2297	674.6775	662.1495
New Zealand	1988Q2-2024Q2	417.2078	424.8559	416.6719	418.7911	406.0082	407.0024
South Korea	1964Q1-2024Q2	652.3785	656.8891	651.4499	660.2585	639.5034	646.4666
Japan	1960Q1-2024Q2	883.7449	891.6378	880.9667	886.2008	867.3169	871.7063
Hong Kong	1985Q1-2024Q2	351.4688	348.0478	346.8881	347.2495	337.1892	340.9246

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

Table 2*b* 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	1950Q1-2024Q2	-13.6244	-15.3535	8.3996	3.3564	16.6472	13.1504
United Kingdom	1955Q1-2024Q2	-87.6333	-88.1512	-89.6394	-89.9796	-89.2024	-91.7205
Canada	1947Q3-2006Q4	-72.0532	-71.7812	-64.4770	-66.2576	-62.2194	-64.2760
	1967Q1-2024Q2	-64.9929	-66.2571	-61.9475	-63.1689	-50.2364	-58.9810
Australia	1969Q3-2024Q2	-144.2463	-146.9065	-140.3467	-142.2191	-126.7830	-130.1300
New Zealand	1988Q2-2024Q2	1.4811	1.8822	5.6223	3.2538	4.2795	2.8030
Switzerland	1972Q1-2024Q2	-43.4559	-44.1075	-39.0693	-40.4712	-21.0416	-25.0099
Sweden	1998Q1-2024Q2	72.2438	72.3716	69.4531	69.4531	67.4080	66.4256
Euro area	1999Q1-2024Q2	75.7278	75.5544	74.8249	74.4669	73.9300	70.9283
Denmark	1991Q1-2024Q2	-30.2668	-30.6901	-27.2982	-28.0958	13.9556	3.1774
South Korea	1964Q1-2024Q2	-135.0929	-142.1783	-119.2569	-137.8390	-91.4879	-99.8509
Japan	1960Q1-2024Q2	209.8299	209.6898	205.6815	205.6815	210.2267	210.0041
Hong Kong	1985Q1-2024Q2	-74.2537	-74.3769	-70.9778	-71.1801	-64.2284	-64.4633

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 2*a* 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, New Zealand, South Korea, and Japan, so that neither of the two specifications clearly dominates the other one.

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

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

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,[17] 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.[18] This can be seen quite clearly in Figure 2 for Australia, Canada, the Euro area, Hong Kong, Sweden, Switzerland, and the U.K..

---

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.

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

[18] 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

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 therefore conceptually akin to (e.g.) regressing GDP on consumption.[19]

The results are reported in Table 2*b*. The evidence is much sharper than for the 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, although in two cases (Japan, and Hong Kong) the difference between the modes of the log-likelihood is very small.

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

We draw two 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, for low-inflation countries the Selden-Latané specification appears to exhibit the best overall performance among the three.

Based on this, for low-inflation countries we will use the Selden-Latane as our benchmark functional form, but we will also provide estimates for the other two specifications.

4.3 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.

---

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.

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

Table 3 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*
		Low-inflation countries:		
United States	1950Q1-2024Q2	0.713	0.823	0.873
United Kingdom	1955Q1-2024Q2	0.672	0.819	0.958
Canada	1947Q3-2006Q4	0.354	0.671	0.691
	1967Q1-2024Q2	0.797	0.871	0.823
Australia	1969Q3-2024Q2	0.791	0.780	0.847
Switzerland	1980Q1-2024Q2	0.632	0.819	–[c]
New Zealand	1988Q2-2024Q2	0.814	0.457	0.561
Sweden	1998Q1-2024Q2	0.865	0.598	–[c]
Euro area	1999Q1-2024Q2	0.843	0.636	–[c]
Denmark	1991Q1-2024Q2	0.678	0.152	–[c]
South Korea	1964Q1-2024Q2	0.101	0.662	0.928
Japan	1960Q1-2024Q2	0.031	0.031	0.684
Hong Kong	1985Q1-2024Q2	0.666	0.557	0.817
		High-inflation countries:		
Argentina	1946-2023	–	–	0.590
Chile	1946-2023	–	–	0.908
Bolivia	1980-2022	–	–	0.721
Israel	1982Q1-2024Q2	–	–	0.681
Mexico	1985Q4-2024Q2	–	–	0.422
		Hyperinflations:		
Weimar Republic	Sep. 1920-Oct. 1923	–	–	0.210

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

4.3.1 Testing for stability in the cointegration vector

Table 3 reports evidence from Hansen and Johansen's (1999) tests for stability in the cointegration vector[20] for our dataset, based on any of the three money demand

[20] 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.

specifications. Only for Japan based on either the Selden-Latané or the semi-log specifications the tests detect evidence of instability.Overall, there is virtually no evidence of a break in the real money demand relationship derived from theory. This is reassuring in itself, but also in reference to the issue raised by Ireland, and that has prevailed in the discussion in the United States, related to an alleged structural break in money-demand relationships sometimes between the late 70s and the early 80s. It was the alleged presence of such a break that justified Ireland's focus on the most recent data. By contrast, the results in Table 3 show, on the one hand, that once we take into account U.S.-specific regulatory changes there is no evidence of a break in U.S. money-demand relationships over the post-WWII period. On the other hand, they show that in other developed countries that did not experience regulatory changes, the high inflation episode of the late 70s and early 80s is consistent with a stable demand for real M1 balances just based on the standard M1 monetary aggregate.

4.3.2 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).[21] Based on this argument, at sufficiently low interest rates money demand (and therefore money velocity) should be largely *unresponsive* to changes in interest rates, since most (or all) households simply do not participate in financial markets. 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, since 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

[21] 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 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.

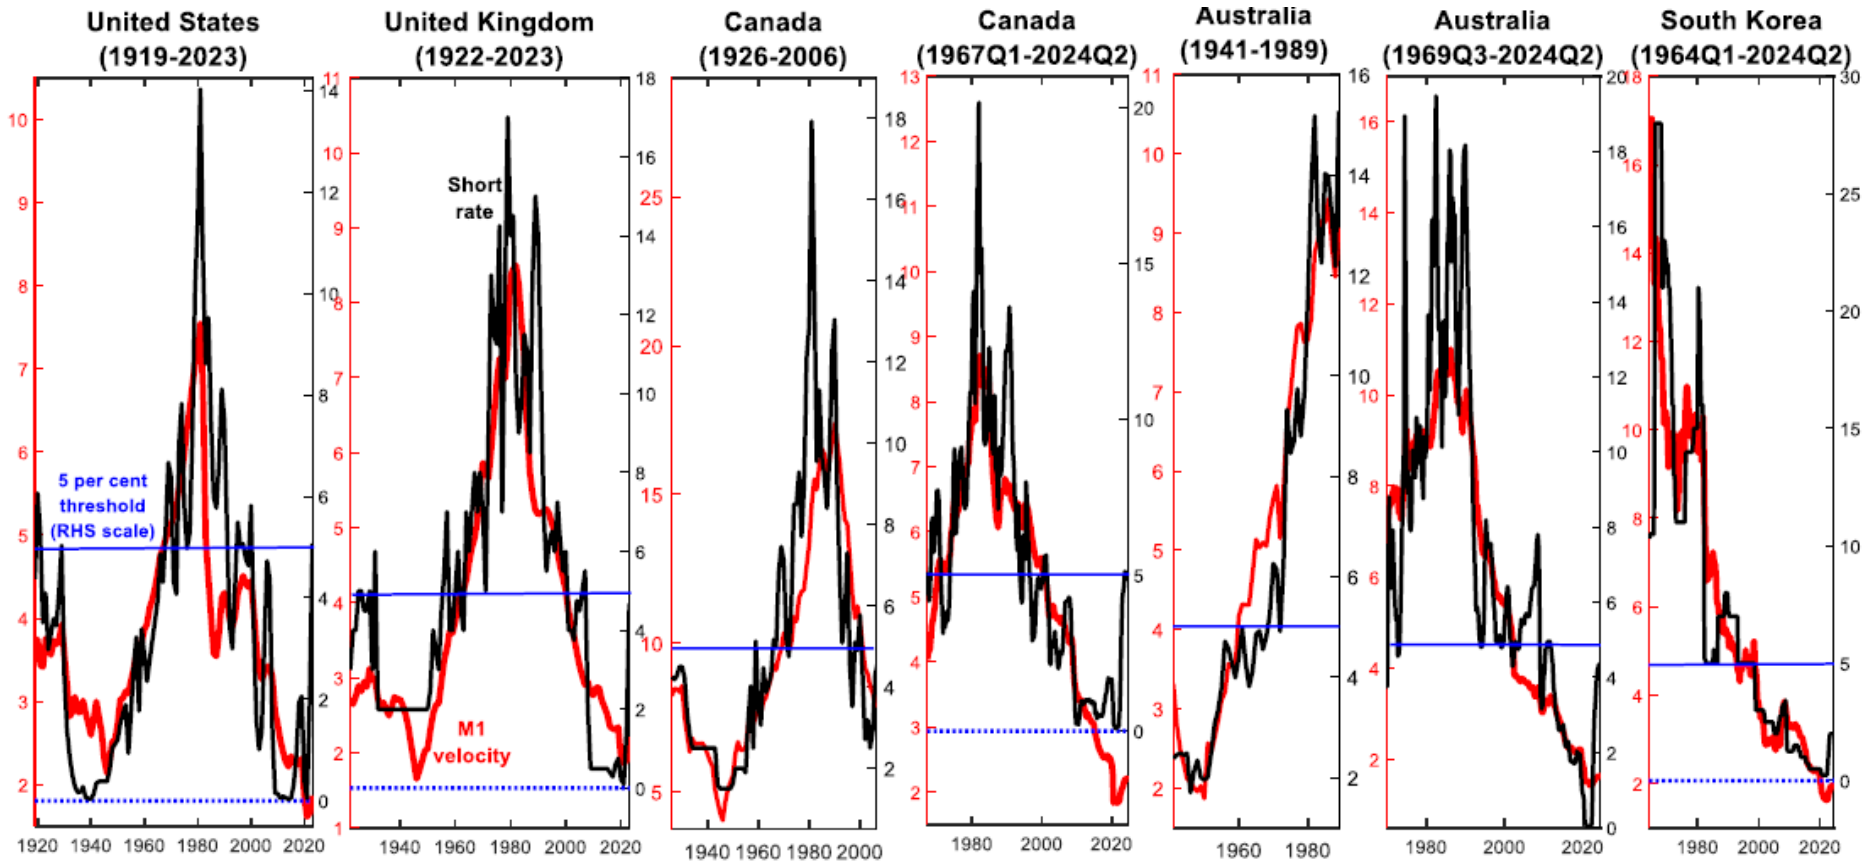


Figure 3 Informal evidence on the possible presence of non-linearities at low interest rates

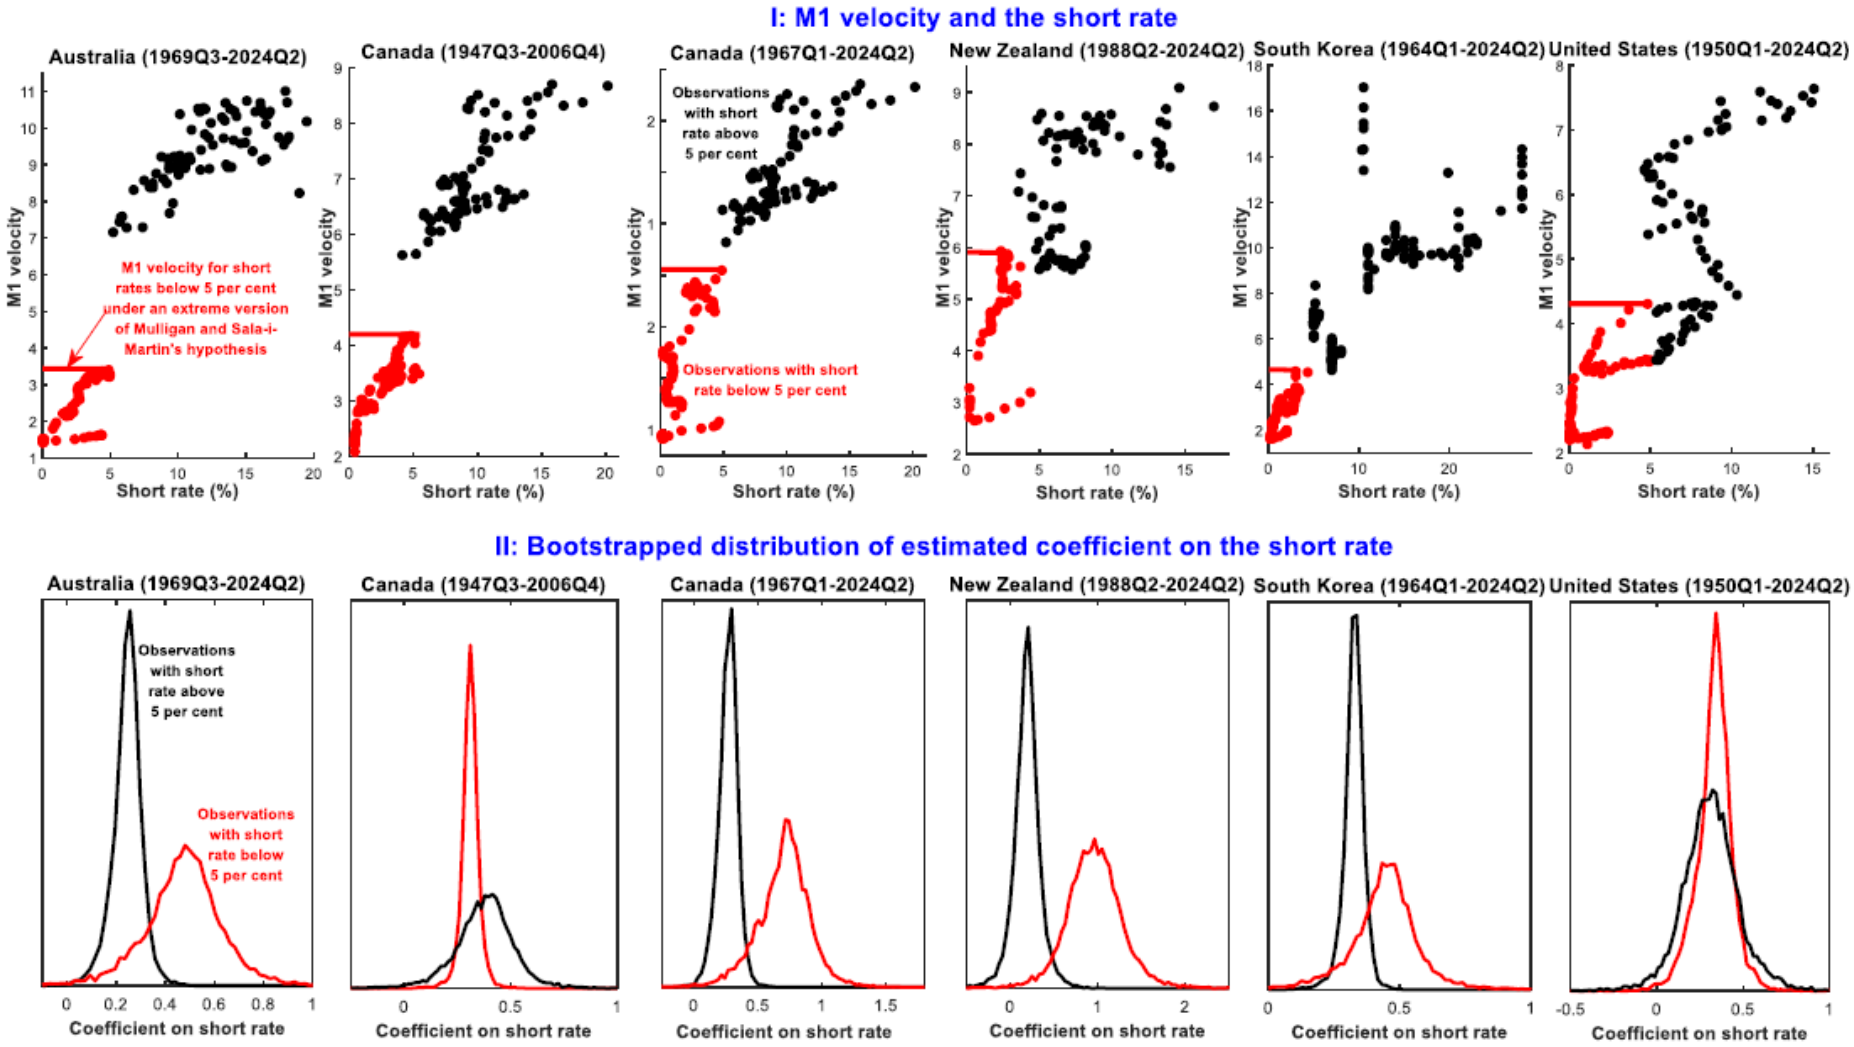


Figure 4 M1 velocity and short-term nominal interest rates: observations with the short rate above and below 5 per cent (quarterly data)

a hard time in 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, following 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 3 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.[22] The figure shows the raw data for M1 velocity and a short rate. The evidence speaks for itself, and it provides *no support* to the notion that velocity—and therefore money demand—may be less responsive to interest rate changes at low interest rates. The only possible exception is the U.S. until WWII.

Overall, the 'big picture' emerging from Figure 3 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.[23]

Figure 4 shows evidence based on quarterly data for five 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.[24] (The sub-samples with the short rate below and above 5%

[22] 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.

[23] 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.

[24] 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.

are reported in Table 4.) 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. For Canada, in line with Figure 3, 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 does not allow to make any strong statement. At the very least, however, evidence provides no support to the notion that the slope may have been flatter at low interest rates. Finally, evidence for the U.S. is idiosyncratic, with the observations below 5% clustered in two separate loops, but once again, in no way does it suggest that the demand curve may be flatter at low interest rates.

Table 4 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	$P(\phi_{R<5}<\phi_{R>5})$	Estimate and 90% confidence interval	Median and 90% confidence interval
Australia	0.072	0.481 [0.269; 0.637]	0.288 [0.152; 0.409]
Canada, I	0.703	0.318 [0.256; 0.380]	0.392 [0.151; 0.594]
Canada, II	0.098	0.788 [0.286; 1.169]	0.415 [0.142; 0.634]
South Korea	0.341	0.413 [0.246; 0.558]	0.383 [0.307; 0.451]
United States	0.172	0.586 [0.198; 0.865]	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: 2009Q1-2023Q1 and 1973Q1-1993Q2; Canada, I: 1947Q3-1967Q3 and 1973Q2-1993Q2; Canada, II: 2001Q1-2023Q2 and 1973Q3-1993Q3; South Korea: 1995Q3-2023Q2 and 1964Q1-1995Q2; United States: 2001Q1-2023Q1 and 1972Q4-1991Q3.

The second row of Figure 4 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 (10) , which is our benchmark specification[25]

Table 4 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 4, and from the bottom row of Figure 4, is that there is *no 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 Canada's first sample, but the p-value, at 0.703, is far from being significant even at the 10% level. Second, in two of the remaining cases (Australia, and Canada's second sample) the p-values (equal to 0.072 and 0.098, 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 4).

But then, how can we rationalize Mulligan and Sala-i-Martin's (2000) evidence? The next sub-section provides a straightforward explanation.

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 is associated with a larger change (in absolute value) in its *logarithm* at low interest rates than it is at higher interest rates,[26] 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 being described by (10), estimating the log-log specification (8) produces the following theoretical value of the *estimated* elasticity

$$\frac{d\ln\left(\frac{M_t}{P_t y_t}\right)}{d\ln r_t} = -\frac{\phi r_t}{a^3 + \phi r_t}, \tag{16}$$

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

[25] 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).

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

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.[27] 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.

We are now ready to discuss our welfare cost computations.

5 Evidence

5.1 Low-inflation countries

The previous theoretical analysis implies that the parameters of the demand for real money balances are the only features that are relevant in order to compute the welfare costs of inflation. In order to see this, it is useful to compute the integral under the money demand curve, as defined in (11), for the three specifications. The integrals are given by

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

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

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 fully summarize all of the information that is required for the computation of the welfare costs of inflation.

[27] 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.

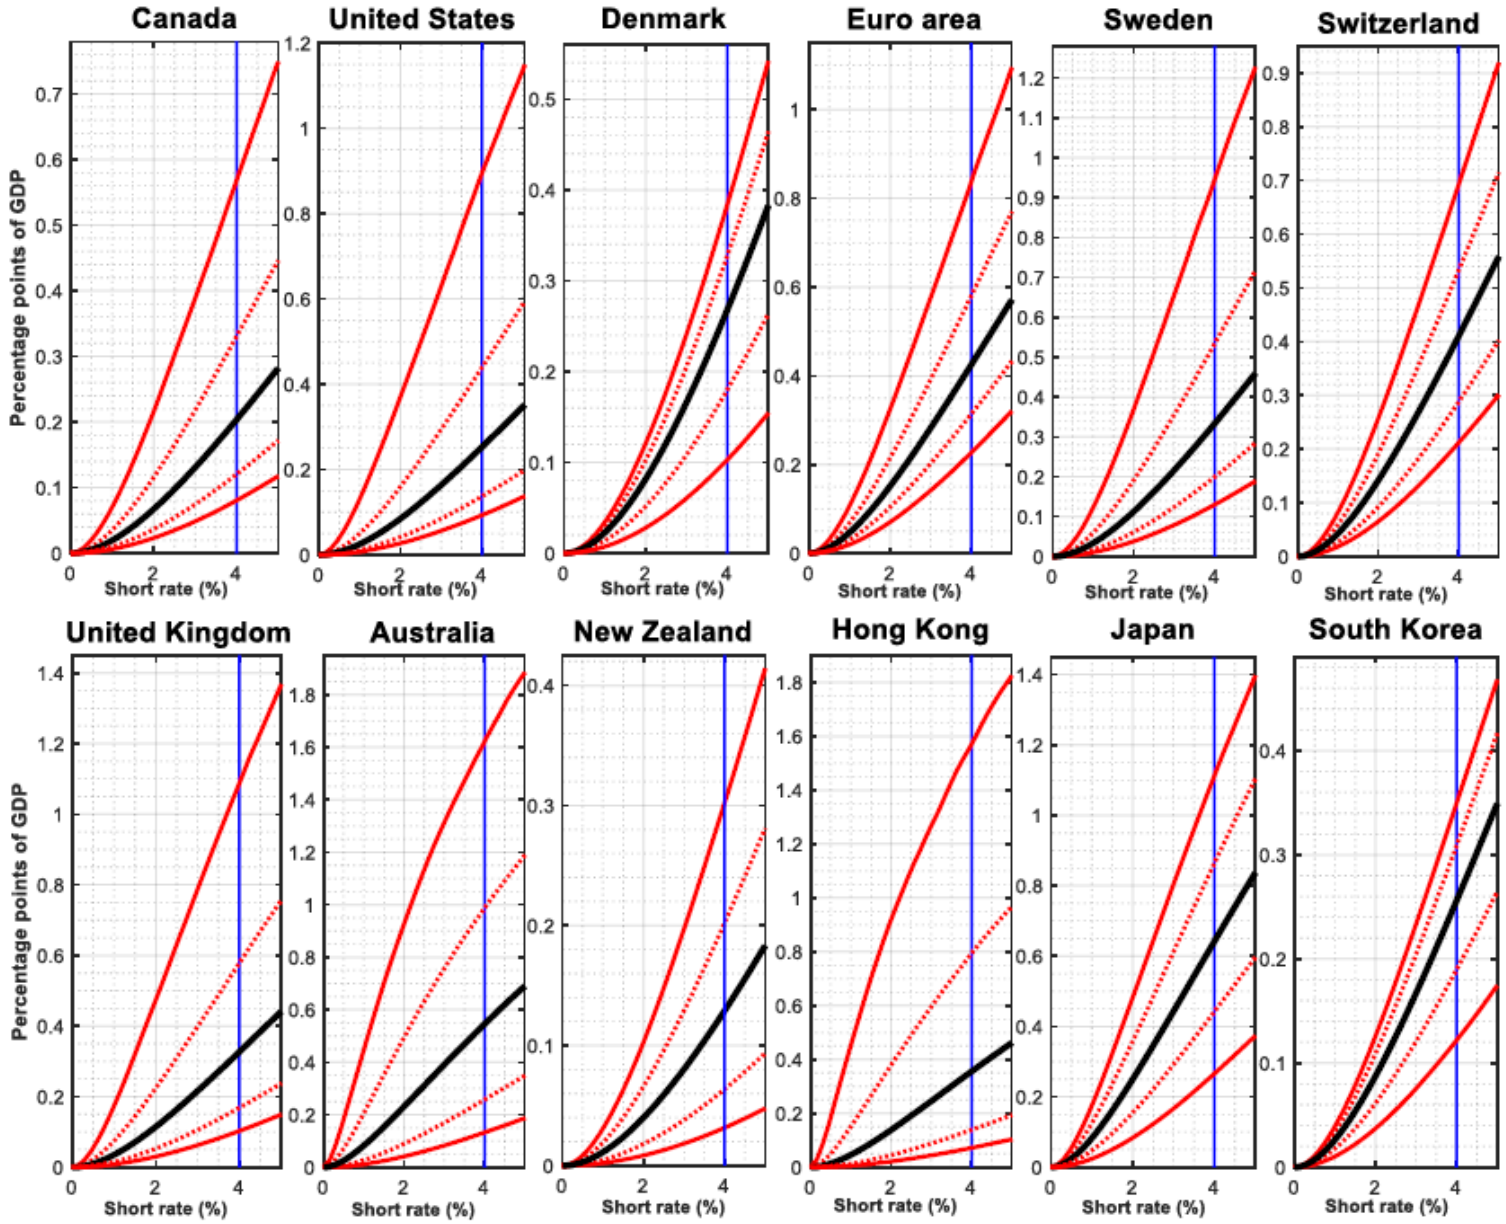


Figure 5 Estimated welfare cost functions based on the Selden-Latané specification: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

In what follows, we discuss in detail our results using the Stock and Watson DOLS estimates, and leave for the appendix the analysis with Wright's (2000) tests, where, as 'point estimates', we pick the value that is most difficult to reject at the 10% level. Results based on Wright's (2000) approach are very similar to those produced by Stock and Watson's estimator, except for a few cases in which the estimated welfare cost are slightly higher.[28] The methodology we use in order to estimate the demand for real M1 balances follows Luetkepohl (1991, pp. 370-371). Specifically, we start by estimating *via* OLS the cointegrating regression corresponding to any of the three specifications, i.e. to either (8), (9), or the inverse of (10).[29] 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, 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.[30] 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).

Figure 5 shows, for any of the twelve countries, the estimated welfare cost functions based on the Selden-Latané specification, which based on the previous discussion we take as our benchmark functional form for this group of countries. We plot the point estimates of the lower and upper bounds—which, in line with the discussion in Section 2, are nearly indistinguishable—the 5th and 16th percentiles of the lower bounds, and the 84th and 95th percentiles of the upper bounds of the bootstrapped distributions. Assuming that the inflation target and the natural rate of interest are both equal to 2 per cent, the monetary policy rate will be equal, on average, to 4 per cent. In turn, by arbitrage the same will approximately hold for short-term nominal rate. On the other hand, the thought experiment considered by both Lucas (2000) and Ireland (2009) involved a steady-state nominal rate of 5 per cent. In order to be able to draw

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

[29] So, to be clear, when we work with the Selden-Latane specification we run the cointegrating regression for M1 velocity, rather than its inverse, M1 as fraction of GDP.

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

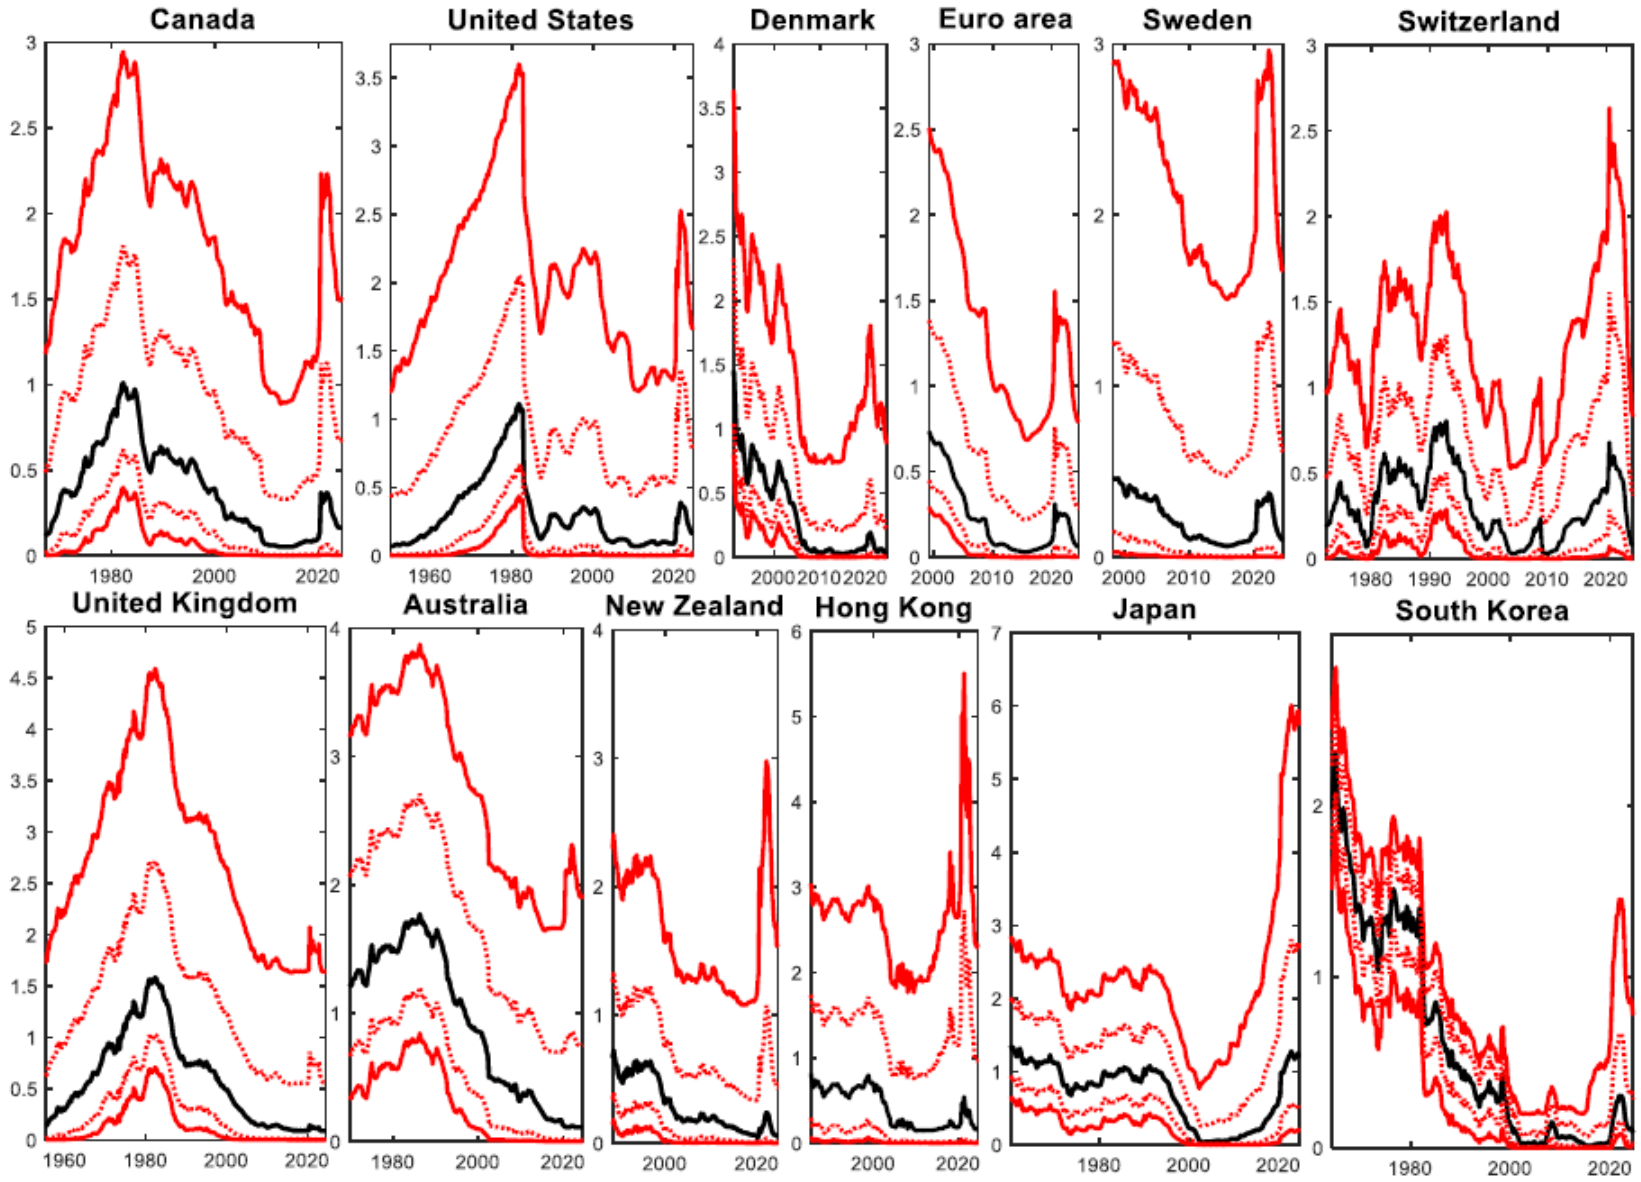


Figure 6 Estimated welfare cost at each point in time based on the Selden-Latané specification: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

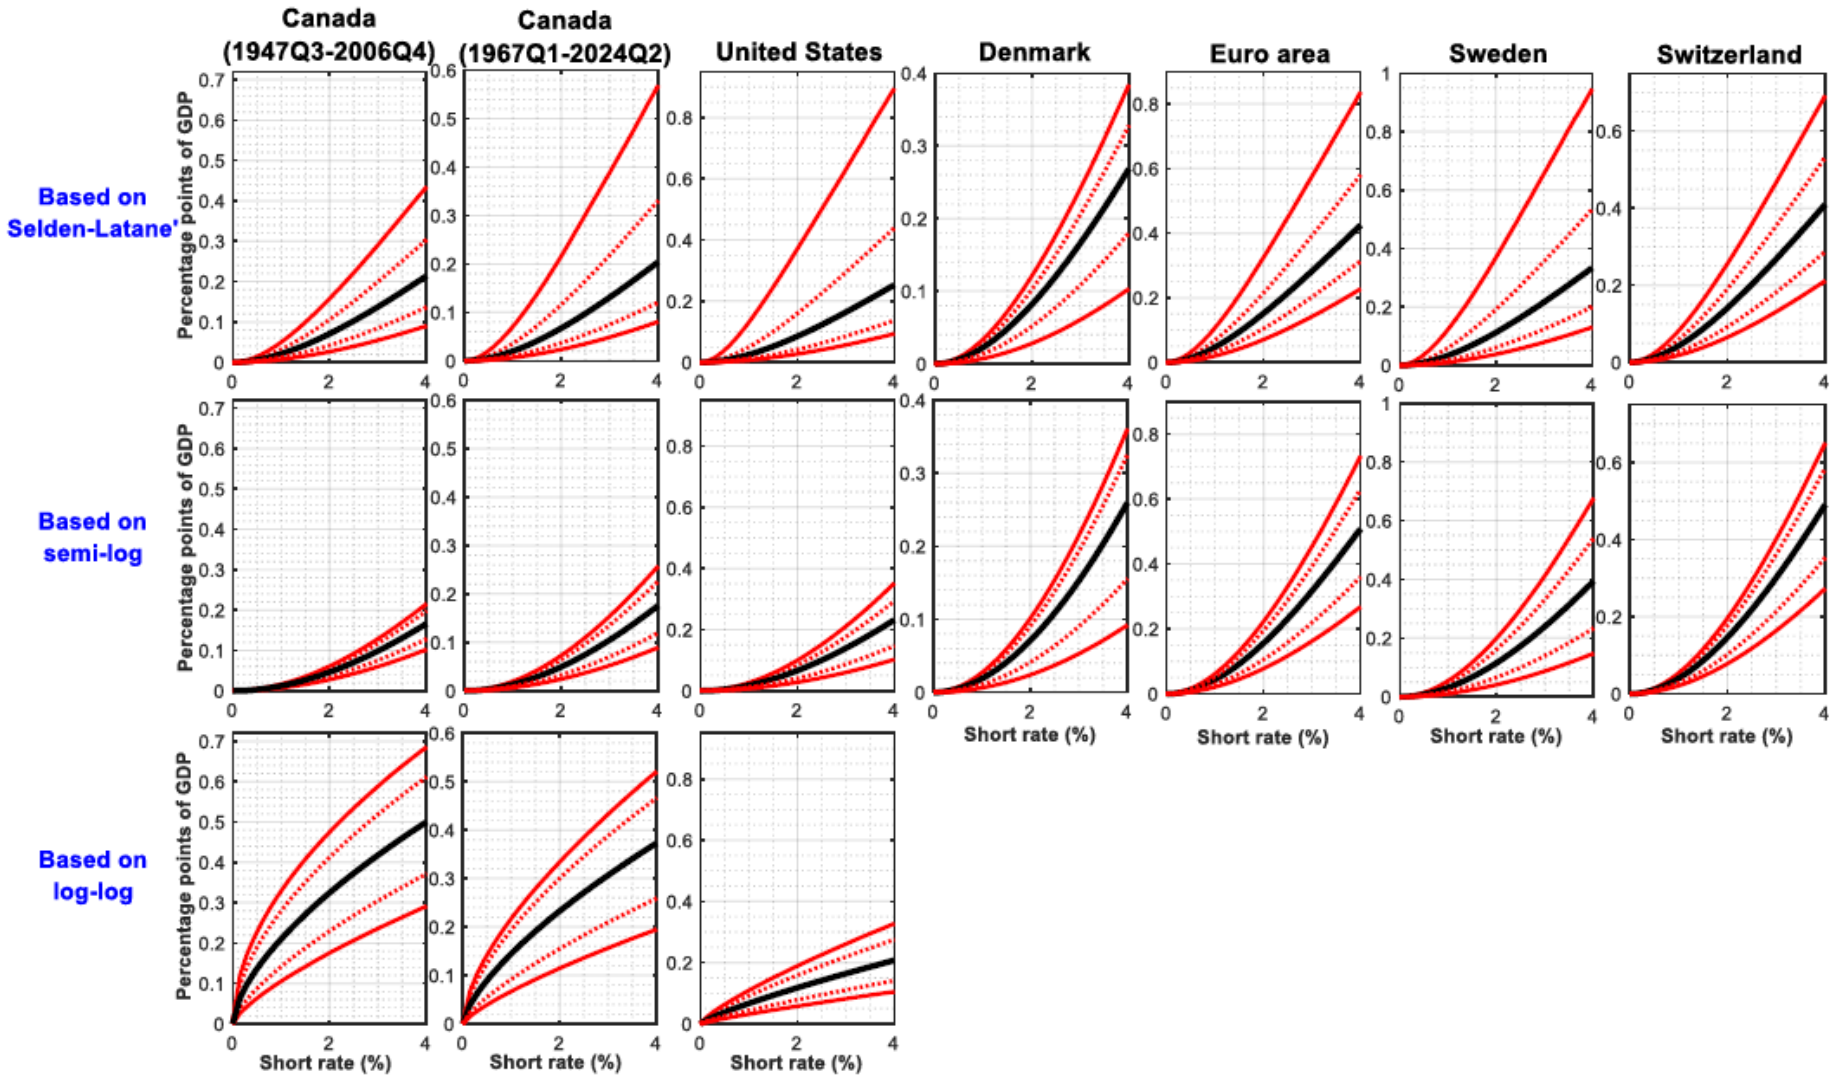


Figure 7a Comparing the estimated welfare cost functions produced by alternative specifications: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

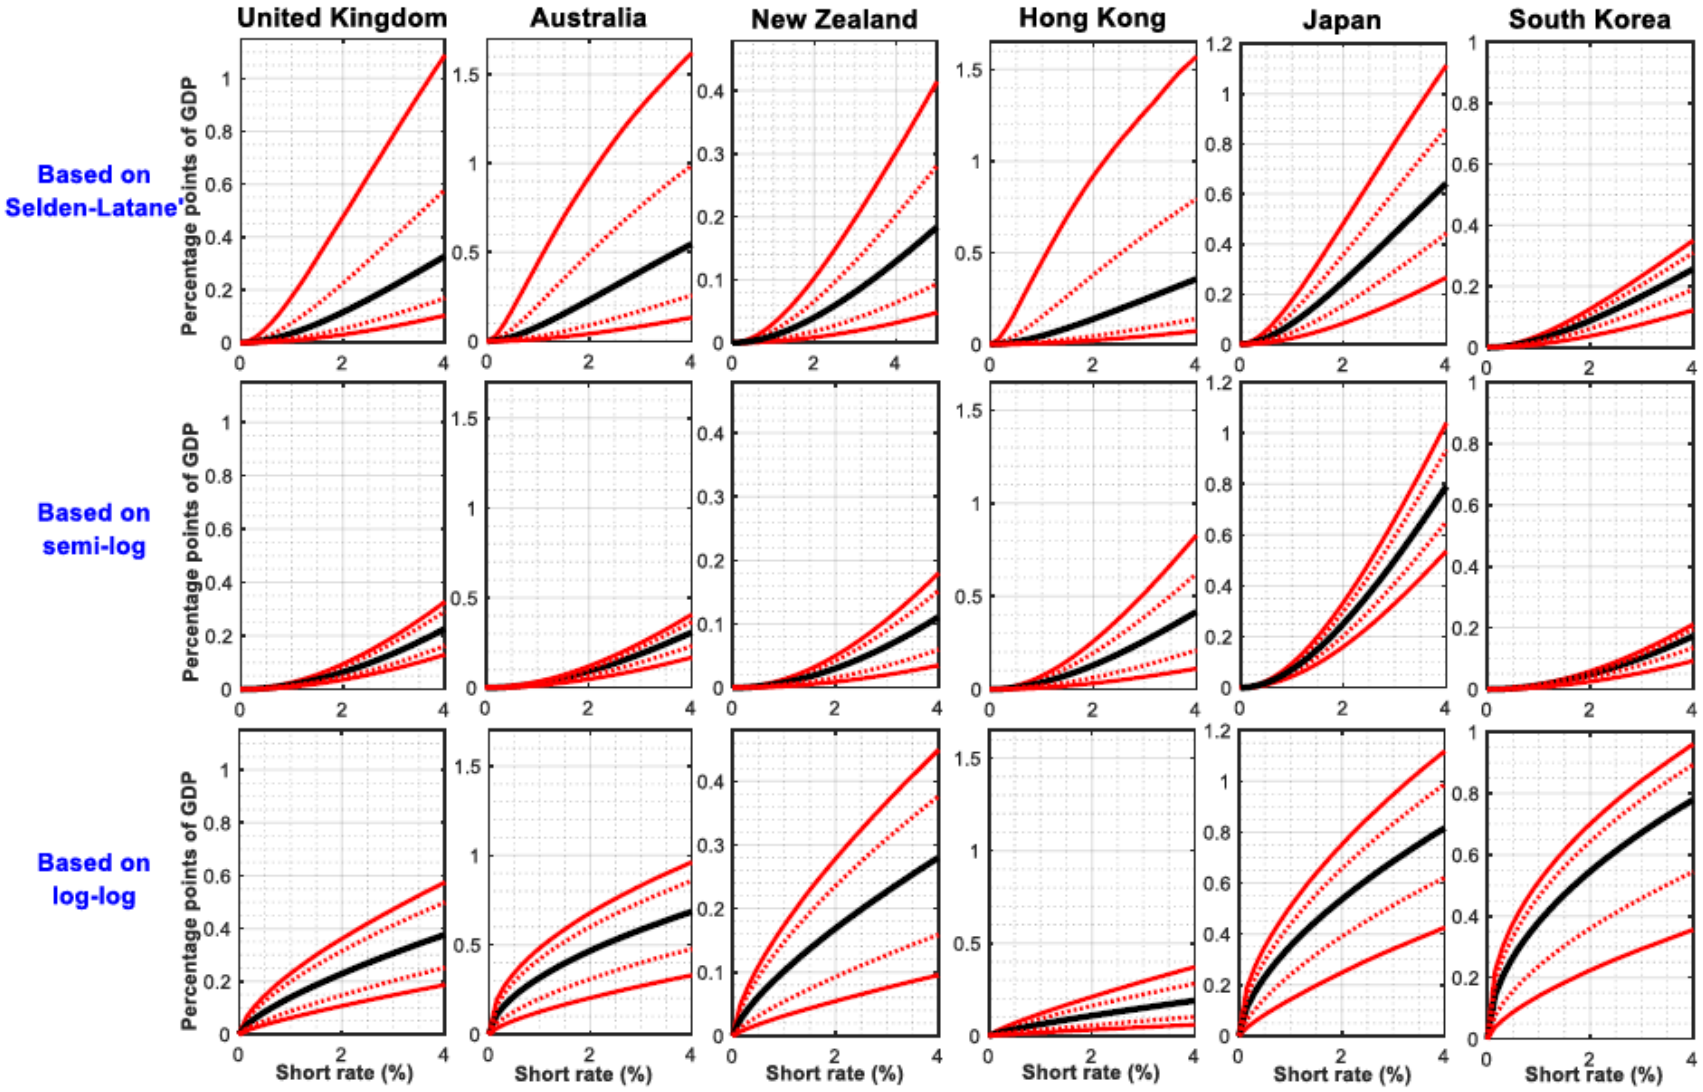


Figure 7b Comparing the estimated welfare cost functions produced by alternative specifications: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

a comparison with their results we therefore plot the welfare cost functions over the domain from 0 to 5 per cent for the short rate.

Starting from the U.S., which has been the focus of most previous research, based on a calibrated log-log Lucas (2000) computed a welfare loss of a 5 per cent nominal rate equal to 1.1% of permanent consumption. Based on a semi-log estimated over the period 1980Q1-2006Q4, on the other hand, Ireland (2009) estimated a welfare loss of just 0.04 per cent of lifetime consumption. Our point estimate in Figure 5 is equal to 0.35 per cent of GDP, about one-third of the loss computed by Lucas (2000), and an order of magnitude larger than that estimated by Ireland (2009). So, even just focusing on point estimates, our evidence for the U.S. suggests that these losses are not negligible, since—it is important to recall—they are suffered *every single year*. Further, and crucially, the 90 per cent confidence interval in Figure 5 stretches between 0.15 and 1.15 per cent. This points towards the existence of sizeable 'upward risks' potentially associated with the welfare costs of inflation, which policymakers obviously should take into account when choosing inflation objectives. At our reference benchmark of a 4 per cent nominal interest rate, the upper bound of the 90 per cent confidence interval is still equal to 0.9 per cent of lifetime consumption, which would definitely be a sizeable annual loss. The argument about 'upward risks' holds for *all* countries. In particular, for Hong Kong, Japan, Australia, the U.K., the Euro area, Sweden, are quantitatively either close to the U.S.'s, or even greater. A policymaker aiming at implementing *robust policies* would take such 'upward risks' into account, and *ceteris paribus* would set the inflation target *lower* than the value just implied by the point estimate.

A second finding emerging from Figure 5 is a non-negligible extent of heterogeneity across countries. For example, focusing on point estimates, and our benchmark thought experiment of a natural rate of interest and an inflation target both equal to 2 per cent, welfare costs in the Euro area, equal to 0.43 per cent, are materially greater than those in the U.S., equal to 0.25 per cent, thus suggesting that *ceteris paribus* the inflation target in the Euro area should be lower. More generally, point estimates range betweeen 0.2 per cent for Canada and 0.65 per cent for Japan, whereas the 95th percentiles of the bootstrapped distribution ranges betweeen 0.38 per cent for Denmark and 1.6 per cent for Australia. Again, these are sizeable figures.

Figure 6 brings this home in the starkest possible way. The figure reports the point estimates of the lower and upper bounds of these costs at each point in time along the sample, together with the 5th percentiles of the lower bounds and 95th percentiles of the upper bounds of the bootstrapped distributions. Once again, the point estimates of the lower and upper bounds are nearly indistinguishable. A first finding emerging from the figure is a dramatic extent of variation in these costs, due to a corresponding large variation in the level of short-term nominal interest rates over the sample periods. The main finding in Figure 6, however, is that if we uniquely focus on countries whose samples include the Great Inflation, at the peak of the inflation episodes point estimates of the welfare costs had ranged between 0.6 per cent of GDP

for Switzerland and 2.3 per cent for South Korea. Again, these numbers are far from negligble, and in fact they are uniformly sizeable, reiterating once again one of our main points: the assumption, explicit or implicit in the previous literature, that the welfare costs of inflation can be ignored is unwarrated.

A key implication of our results is that the common practice in the literature of ignoring money in the analyses of optimal monetary policies is likewise unwarranted, and it can be seriously misleading. For example, Coibion, Gorodnichenko and Wieland (2012)[31] advance a compelling argument against increasing the inflation target in countries like the U.S., based on a model with frictions in price-setting and with recurrent, although not very frequent episodes with the nominal interest rate constrained by the ZLB. Based on their preferred specification they compute the welfare effect of an interest rate of 5 per cent to be close to 0.6 per cent of lifetime consumption. Such an estimate, combining the cost created by price frictions and the probability to be at the ZLB, is of the same order of magnitude of the estimates we obtain for the U.S., and it is in fact *less than half* of the upper bound of our 90 per cent confidence interval in Figure 5.

Finally, Figures 7*a*-7*b* explore whether the three functional forms for the demand for real money balances do, or do not produce materially different welfare cost functions. Overall, evidence is mixed. Starting from the U.S., at a 4 per cent short rate point estimates range between 0.2 (based on the log-log) and 0.25 per cent (based on the Selden-Latane) of GDP. This shows that the conventional-wisdom notion that, *ceteris paribus* (e.g., for given observations for velocity and the interest rate), a log-log specification should be expected to produce comparatively higher welfare costs is, in general, incorrect. The reason for this is straightforward: Although the log-log specification does not have a finite satiation level of money balances at R_t=0 , empirical estimates of the intercepts and the coefficient on the (logarithm of the) interest rate do in fact matter. Therefore, focusing e.g. on the semi-log and the log-log, it is perfectly possible that the parameters estimates are such that

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

In several instances, however, the standard assumption appears to be validated: this is the case for Canada, the U.K., Australia, New Zealand, and South Korea. We now turn to the welfare costs of high inflations and Weimar Republic's hyperinflation

5.2 High inflation countries and Weimar Republic's hyperinflation

As it has been extensively documented, very high inflations have uniformly been associated with macroeconomic mayhem and the destruction of wealth held in nom-

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

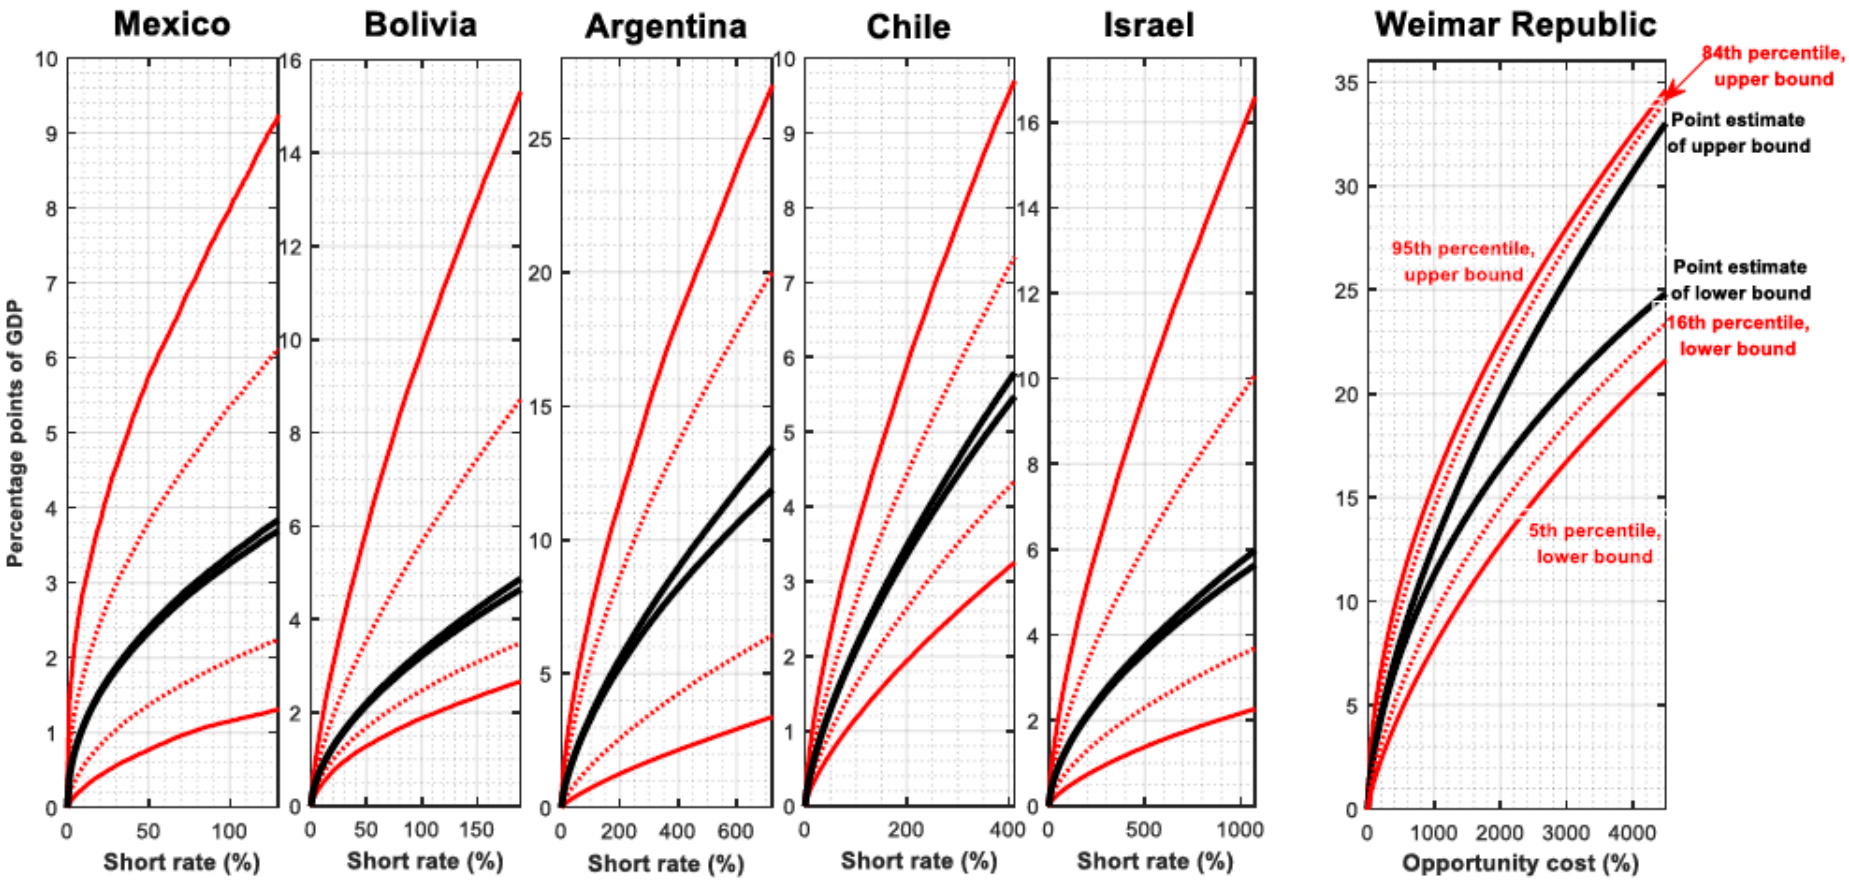


Figure 8 Estimated welfare cost functions for high-inflation countries and Weimar Republic's hyperinflation: point estimates of the lower and upper bounds, 5th and 16th percentiles of the lower bounds, and 84th and 95th percentiles of the upper bounds of the bootstrapped distributions

inal assets. Evidence is especially stark for hyperinflations. For Weimar Republic's episode, for example, the data reported in Table XL of Graham (1930, p. 317) show that the unemployment rate among trade union members, which in 1922 had oscillated between 0.6 and 3.3 per cent, increased rapidly following the invasion of the Ruhr on the part of France in January 1923, which as pointed out by Bresciani-Turroni (1937) '*gave the coup de grâce to the national finances and the German mark*', thus inaugurating the most extreme phase of the hyperinflation. Unemployment reached 6.2 per cent in May, 9.9 in September, and it further increased to a remarkable 28.2 per cent in December, the last month of the hyperinflation.

In this section we explore the welfare costs of these episodes originating from lack of liquidity satiation, which, quite surprisingly, have been near-uniformly overlooked by the previous literature. The only exception we are aware of is Barro (1972), which reports evidence for the Weimar Republic's episode and four other hyperinflations that had been studied by Cagan (1956). For all countries we estimate Meltzer's (1963) log-log specification. Our main finding is that for very high inflations and hyperinflations these costs are very far from negligible, as at the inflation peaks of the respective episodes they range from nearly 4 per cent of output for Mexico to between 25 and 33 per cent for Weimar's hyperinflation.

Figure 8 show, for either country/episode, the estimated welfare cost functions (in percentage points of GDP) for values of the opportunity cost of money from zero to the maximum value that it had taken over the sample period. Several facts clearly emerge from the figure. In particular,

first, in line with the previous evidence for low-inflation countries, the point estimates of the upper and lower bounds of the welfare cost functions are very tight. The only exception is Weimar's hyperinflation, for which they can be clearly distinguished.

Second, focusing again on the point estimates, the evidence in Figure 8 uniformly suggests that for all of these episodes the welfare costs had been sizeable-to-large. In particular, at the peaks of the inflation episodes these costs had been equal to nearly 4 per cent of GDP for Mexico; between 4 and 12-14 per cent for Bolivia, Argentina, Chile, and Israel; and between 25 and 33 per cent of income for the Weimar Republic. This provides a stark illustration of how, beyond the already well known and widely documented costs of very high inflations and hyperinflations in terms of economic mayhem and the destruction of wealth held in nominal assets, these episodes have consistently imposed non-negligible, and sometimes large costs *uniquely* in terms of lack of liquidity satiation, by compelling agents to hold comparatively low levels of real money balances.

Third, uncertainty is near-uniformly substantial, sometimes remarkably so. E.g., for Argentina the upper bound of the 90 per cent-coverage bootstrapped confidence interval at the peak of the inflation episode is equal to 27 per cent of GDP.

6 Conclusions

In this paper we revisit the estimation of the welfare costs of inflation for twelve low-inflation and five high-inflation countries, and for Weimar Republic's hyperinflation. In our analysis we follow the tradition of considering the most liquid monetary assets, which include cash and transactional deposits. We abstract from a detailed discussion of the demand for each of the components, an issue recently addressed by Kurlat (2019).[32]

Our evidence suggests that, contrary to the explicit or implicit assumption in much of the literature, these costs are often far from negligible. For the U.S. our point estimates are equal to about one-third of those computed by Lucas (2000), and an order of magnitude larger than those obtained by Ireland (2009).[33] Crucially, the most empirically plausible money-demand functional form points towards sizeable 'upward risks' for these costs, with the 90% confidence interval associated with a 4% nominal interest rate stretching to 0.9 per cent of GDP. At the peak of the inflation episodes, welfare costs had ranged between 0.6 and 2.3 per cent of GDP for low-inflation countries; between 4 and nearly 14 per cent for high-inflation ones; and between 25 and 33 per cent for Weimar's hyperinflation.

Our evidence suggests that ignoring money in analyzing optimal monetary policies can be seriously misleading. For instance, Coibion, Gorodnichenko and Wieland (2012)[34] make a compelling argument against increasing the inflation target in countries like the U.S., based on a model with frictions in price-setting and with recurrent, though not very frequent, episodes with the nominal interest rate at the ZLB. Based on their preferred specification they compute the welfare effect of an interest rate of 5 per cent to be close to 0.6 per cent of lifetime consumption. Such an estimate, combining the cost created by price frictions and the probability to be at the ZLB, is of the same order of magnitude of the estimates we obtain for the U.S., and it is in fact *less than half* of the upper bound of our 90 per cent confidence interval.

In any of the low-inflation countries the demand for M1 as a fraction of GDP at very low, or even negative interest rates has (so far) exhibited no obvious difference compared to its behavior at higher interest rates. These results contrast with those of Mulligan and Sala-i-Martin (2000), who, based on U.S. households micro data, provided evidence that money demand becomes comparatively flatter at low interest rates. We provide a straightforward explanation for Mulligan and Sala-i-Martin's (2000) finding, by showing mathematically that if the true money demand

[32] 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.

[33] For a steady-state interest rate of 5 per cent, Lucas (2000) computed the cost to be around 1.1 per cent of lifetime consumption. Ireland (2009) challenged Lucas' interpretation of the data, and estimated a mere 0.04 per cent of lifetime consumption.

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

specification is the one that, as shown by Benati *et al.* (2021), is the most empirically plausible at low inflation rates, Mulligan and Sala-i-Martin's (2000) approach automatically produces spurious evidence of a flatter demand curve at low interest rates.

From a theoretical standpoint, following Alvarez, Lippi and Robatto (2019) we construct upper and lower bounds for the welfare costs of inflation. As they show, 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. We extend their results for a quite general sub-class of the models they analyze and compute exact lower and upper bounds for 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 low-inflation countries such as the U.S.

For policy purposes, our main finding is that the welfare costs inflation in the Euro area are materially larger than in the U.S.. This suggests that, *ceteris paribus*, the inflation target should be lower.

7 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.

Barro, R.J. (1972): "Inflationary Finance and the Welfare Cost of Inflation", *Journal of Political Economy*, 80(5), 978-1001.

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.(2024): "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. (1973): "The Demand for Money Revisited", *Brookings Papers on Economic Activity*, 3, 577-646.

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

Graham, F. D. (1930): *Exchange, Prices, and Production in Hyperinflation*, Princeton University Press.

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.

Holtfrerich, C.L. (1980): *Die Deutsche Inflation 1914-1923: Ursachen und Folgen in Internationaler Perspektive.* Berlin/New York: Walter de Gruyter, S. 70. Primärquelle: Die Wirtschaftskurve 1922, 1923 passim.

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.

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

$$\beta^t U(c_t) = P_t\lambda_t + P_t\delta_t \tag{20}$$

$$P_t\lambda_t\theta_n(n_t, \nu_t)z_t = M_t\delta_t \tag{21}$$

$$\xi_t = \lambda_t(1 + r_t^m) + \delta_t n_t \tag{22}$$

$$\xi_t = \lambda_t(1 + r_t^b) \tag{23}$$

$$\lambda_t = E_t\xi_{t+1} \tag{24}$$

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

$$\frac{n_t}{r_t^b - r_t^m}\theta_n(n_t, \nu_t)z_t = \frac{M_t}{P_t} \tag{25}$$

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 that

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

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 (25) 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 Data

Here follows a detailed description of the dataset.

B.1 Low-inflation countries

B.1.1 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.

B.1.2 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.

B.1.3 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').

B.1.4 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).

B.1.5 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).

B.1.6 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).

B.1.7 Euro area

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

B.1.8 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.

B.1.9 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.

B.1.10 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.

B.1.11 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.

B.2 High-inflation countries

B.2.1 Bolivia

Series for nominal GDP, M1, and a short-term nominal interest rate are from the Unidad de Analisis de Politicas Sociales y Economicas (Bolivia's national statistical agency, known as UDAPE for short).

B.2.2 Chile

Annual series for nominal GDP, the GDP deflator, and M1 are from Braun-Llona et al. (1998) for the period 1940-1995. As for the period 1996-2012, they are from the Banco Central de Chile, Chile's central bank (specifically, nominal GDP and the GDP deflator are from the Banco Central 's Anuarios de Cuentas Nacionales, whereas M1 is from Banco Central 's Base Monetaria y Agregados Monetarios Privados). A short-term nominal interest rate ("1-day interbank interest rate, financial system average (annual percentage)") from Banco Central is available for the period 1940-1995. In order to extend our analysis to the present as much as possible, we therefore also consider, as an alternative measure of the opportunity cost of money, GDP deflator inflation.

B.2.3 Ecuador

All of the data are from the website of Banco Central del Ecuador (henceforth, BCE), Ecuador's central bank. Most of them are from "85 Años, 1927-2012: Series Estadísticas Históricas," a special publication celebrating BCE's 85th anniversary. Specifically, a series for annual CPI inflation ("Variación Anual del Indice Ponderado de Precios al Cunsumidor por Ciudades y por Categorias de Divisiones de Concumo, Nacional"), available for the period 1940-2011, is from Chapter 4 of "85 Años." An annual series for a nominal interest rate has been constructed by linking the series "Tasas, Máxima Convencional, En porcentajes," available for the period 1948-1999, "Tasas de Interés Referenciales Nominales en Dólares, Máxima Convencional," available for the period 2000-2007; and "Tasas de Interés Referenciales Efectivas en Dólares, Máxima Convencional," available for the period 2007-2011. All of them are from from Chapter 1 of "85 Años." An annual series for nominal M1 in US dollars has been constructed by linking the M1 aggregate ("Oferta Monetaria M1, En millones de dólares al final del período"), available for the period 2000-2011, which is expressed in US dollars, and the M1 aggregate ("Medio Circulante (M1), Saldos en millones de sucres"), available for the period 1927-1999, which is expressed in Ecuador's national currency, the sucre (both series are from Chapter 1 of "85 Años"). The latter M1 aggregate has been converted in US dollars based on the series for the sucre/dollar nominal exchange rate found in Chapter 2 of "85 Años," which is available for the period 1947-1999. Specifically, the exchange rate series (sucre per dollar) has been computed as the average between the "Compra" (i.e., buy) and the "Venta" (i.e., sell) series. An annual series

for nominal GDP in U.S. dollars ("Producto interno bruto (PIB), Miles de dólares"), available for the period 1965-2011, is from Chapter 4 of "85 Años". An important point to stress is that since we are working with M1 velocity—defined as the ratio between nominal GDP and nominal M1—the specific unit in which the two series are expressed (US dollars, or Ecuadorian sucres) is irrelevant.

B.2.4 Israel

Quarterly seasonally adjusted data on nominal GDP and the CPI are from the Central Bureau of Statistics, whereas a series for M1 is from Israel's central bank. A series for the Treasury bill rate is from the International Monetary Fund's International Financial Statistics.

B.2.5 Mexico

Quarterly seasonally adjusted data on nominal GDP are from Mexico's statistical agency, INEGI. Quarterly seasonally adjusted data for the CPI, M1, and a 3-month government bond yield are all from the Banco de México.

For all countries, in what follows we work with money velocity (i.e., the inverse of money balances as a fraction of GDP), and a series for the opportunity cost of money, which we compute as the maximum, at each point in time, between inflation and the series for the nominal short-term interest rate (for Argentina we were not able to find an interest rate series, and we therefore work with inflation).

B.3 Weimar's Republic

Monthly data on the velocity of circulation of money based on wholesale prices are from Table XXII of Bresciani-Turroni (1937). The series had been normalized by 1913 (i.e. for the year 1913 it took a value of one). Based on Benati, Lucas, Nicolini and Weber's (2021) data, however, in 1913 German money velocity had been equal to 7.49. Consequently, we have rescaled Bresciani-Turroni's money velocity series by multiplying it by 7.49.[35] A series for the inflation rate is from Cagan (1956). A series for the money market rate ('Tägliches Geld') is from Table 23 of Holtfrerich (1980). The sample period is September 1920-October 1923.

C 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

[35] In fact, working with Bresciani-Turroni's original series produces manifestly absurd results, with the welfare costs of inflation even taking values in excess of 100 per cent of GDP.

for the United States (from the *Center for Financial Stability*, henceforth *CFS*) and for the United Kingdom (from the Bank of England). A second problem is that, for the United States, 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 C.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 C.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.

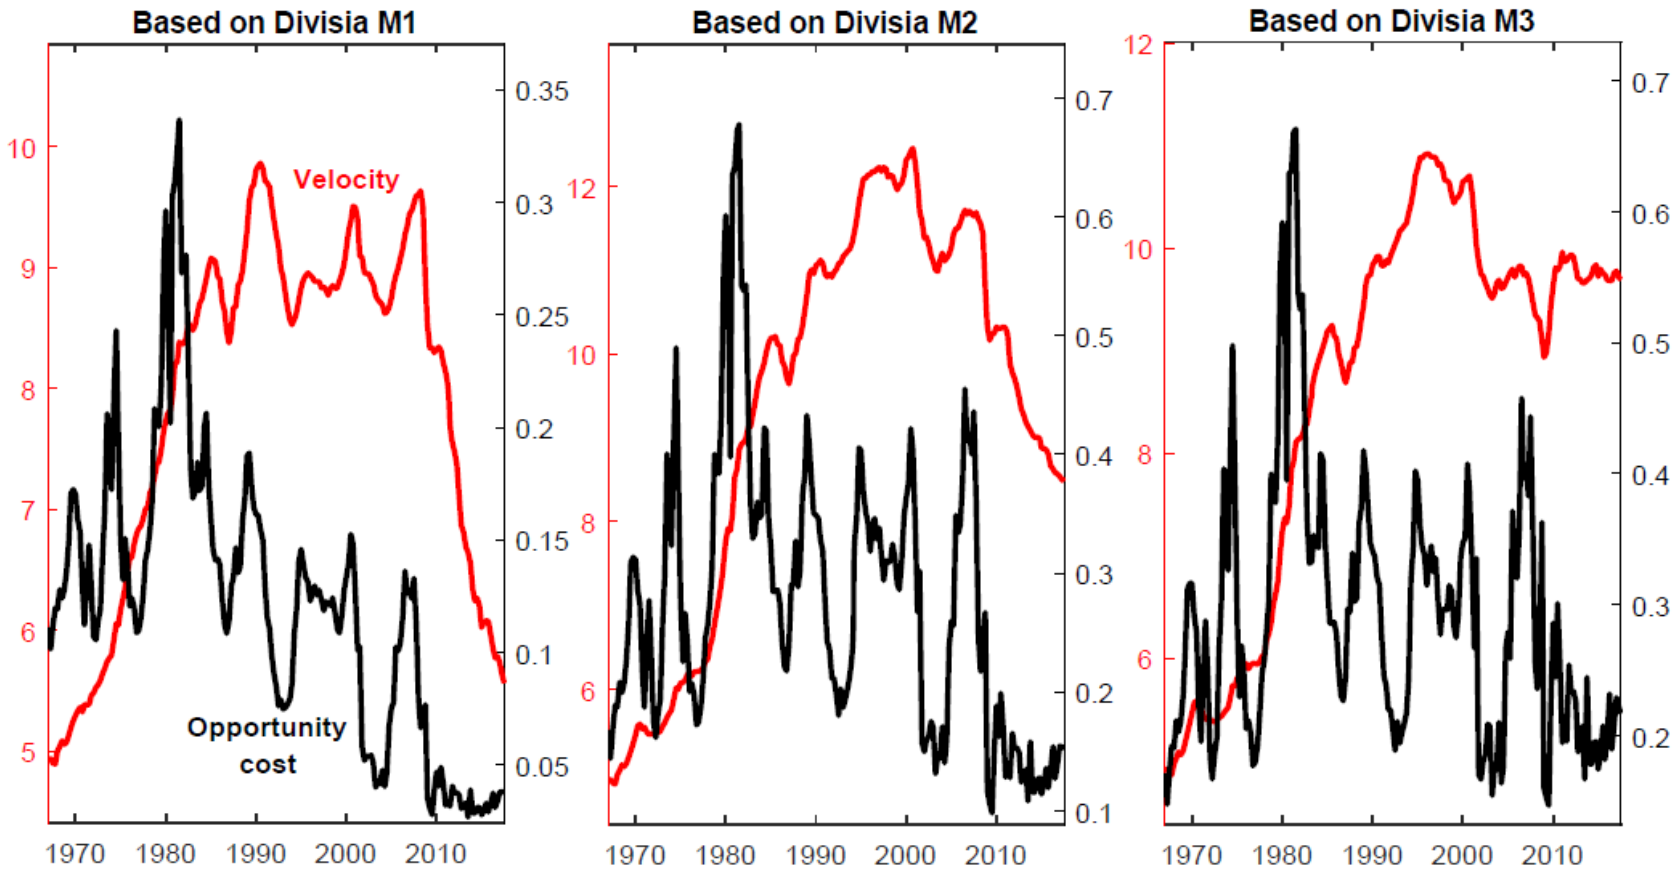


Figure C.1 United States: money velocity based on Divisia aggregates, and the corresponding opportunity costs

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
		Low-inflation countries:							
United States	1959Q1-2023Q2	0.9809	0.9323	0.9236	0.9207	0.3122	0.1794	0.1053	0.2918
	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	1955Q1-2023Q2	0.9257	0.8719	0.8042	0.8490	0.2896	0.3225	0.4118	0.4673
	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
	1967Q1-2023Q1	0.9772	0.9682	0.9557	0.9554	0.3996	0.3831	0.3605	0.6097
Australia	1969Q3-2023Q1	0.9665	0.9600	0.9580	0.9213	0.4143	0.3218	0.5442	0.7099
	1969Q3-2008Q4	0.9883	0.9823	0.9834	0.9679	0.3078	0.2719	0.4116	0.5810
Switzerland	1980Q1-2023Q1	0.8459	0.7684	0.7566	0.6917	0.3639	0.4453	0.1953	0.2024
Sweden	1998Q1-2023Q1	0.5217	0.4343	0.5006	0.6716	0.3254	0.4637	0.5665	0.5411
Euro area	1999Q1-2023Q1	0.1096	0.0847	0.0515	0.0075	0.4403	0.2511	0.2468	0.3625
Denmark	1991Q1-2023Q1	0.1027	0.2238	0.2289	0.1194	0.0999	0.0433	0.0212	0.0090
South Korea	1964Q1-2023Q1	0.0195	0.0000	0.0150	0.0535	0.5079	0.4863	0.3828	0.0867
Japan	1960Q1-2023Q2	0.7305	0.7323	0.8163	0.5827	0.2394	0.3962	0.4375	0.4192
Hong Kong	1985Q1-2023Q2	0.6293	0.6881	0.5964	0.6065	0.3103	0.1820	0.1338	0.3085
[a] Based on 10,000 bootstrap replications of estimated ARIMA processes. Tests are with an intercept and no time trend.									

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
		Low-inflation countries:							
United States	1959Q1-2023Q2	0.9850	0.9688	0.9712	0.9759	0.1579	0.0883	0.2274	0.2411
	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	1955Q1-2023Q2	0.9793	0.9496	0.9284	0.9224	0.2758	0.2210	0.2902	0.7804
	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
	1967Q1-2023Q1	0.9951	0.9898	0.9803	0.9902	0.1001	0.1006	0.3021	0.7613
Australia	1969Q3-2023Q1	0.9898	0.9863	0.9770	0.9609	0.1160	0.0509	0.4413	0.8191
	1969Q3-2008Q4	0.9978	0.9946	0.9953	0.9923	0.3915	0.3312	0.5893	0.6785
Switzerland	1980Q1-2023Q1	0.8647	0.8136	0.7686	0.7438	_[b]	_[b]	_[b]	_[b]
Sweden	1998Q1-2023Q1	0.7177	0.5494	0.5106	0.6041	0.2756	0.2584	0.3197	0.3893
Euro area	1999Q1-2023Q1	0.4266	0.3080	0.2638	0.0954	0.4084	0.6636	0.6686	0.6779
Denmark	1991Q1-2023Q1	0.2648	0.3825	0.4415	0.2066	_[b]	_[b]	_[b]	_[b]
South Korea	1964Q1-2023Q1	0.3815	0.1906	0.4337	0.4606	0.6510	0.6106	0.5806	0.3996
Japan	1960Q1-2023Q2	0.9694	0.9607	0.9528	0.8539	0.5903	0.6472	0.6553	0.6345
Hong Kong	1985Q1-2023Q2	0.7613	0.7888	0.7563	0.7412	0.1532	0.1217	0.1317	0.2013
		High-inflation countries:							
Israel	1982Q1-2019Q4	0.8997	0.8346	0.7721	0.5814	0.5826	0.5698	0.5049	0.4309
Mexico	1982Q1-2019Q4	0.4687	0.4159	0.2106	0.1188	0.2463	0.1085	0.0400	0.0147
		p=1	p=2			p=1	p=2		
Bolivia	1980-2019	0.7971	0.8018			0.7588	0.7822		
Chile	1946-2019	0.6862	0.5399			0.7214	0.6946		
Ecuador	1980-2019	0.9309	0.8715			0.8018	0.8928		
		Hyperinflations:							
		p=1	p=2	p=3	p=4	p=1	p=2	p=3	p=4
Weimar Republic	Sep. 1920-Oct. 1923	0.9309	0.9601	0.8592	0.8569	0.5006	0.5459	0.7271	0.5317

[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.