



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics
and Social Sciences

Department of Economics

The Welfare Costs of Inflation

Luca Benati, Juan-Pablo Nicolini

19-10

December 2019

DISCUSSION PAPERS

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

The Welfare Costs of Inflation*

Luca Benati
University of Bern[†]

Juan-Pablo Nicolini
Federal Reserve Bank
of Minneapolis[‡]

Abstract

We estimate the welfare costs of inflation originating from lack of liquidity satiation—as in Bailey (1956), Friedman (1969), Lucas (2000), and Ireland (2009)—for the U.S., U.K., Canada, and three countries/economic areas (Switzerland, Sweden, and the Euro area) in which interest rates have recently plunged below zero. We pay special attention to *(i)* the fact that, as shown by recent experience, zero cannot be taken as the effective lower bound (ELB); *(ii)* the possibility that, as discussed by Mulligan and Sala-i-Martin (2000), the money demand curve may become flatter at low interest rates; *(iii)* the functional form for money demand.; and *(iv)* what the most relevant *proxy* for the opportunity cost is.

We report three main findings: (1) allowing for an empirically plausible ELB (e.g., -1%) materially increases the welfare costs compared to the standard benchmark of zero; (2) there is nearly *no evidence* that at low interest rates money demand curves may become flatter: rather, evidence for the U.S. (the country studied by Mulligan and Sala-i-Martin (2000)) clearly points towards a *steeper* curve at low rates; and (3) welfare costs are, in general, non-negligible: this is especially the case for the Euro area, Switzerland, and Sweden, which, for any level of interest rates, demand larger amounts of M1 as a fraction of GDP. For policy purposes the implication is that, *ceteris paribus*, inflation targets for these countries should be set at a comparatively lower level.

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

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

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

1 Introduction

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)—should in principle play an important role in the design of monetary policy regimes (specifically, in the choice of the numerical target for inflation). In fact, to the very best of our knowledge, these costs played a uniformly small-to-negligible role both in the design of inflation-targeting frameworks, and in the choice of the numerical inflation objective on the part of the Federal Reserve and the European Central Bank. More generally, the number of studies presenting empirical estimates of these costs is quite surprisingly limited, with Ireland (2009) being the only author to have ever presented estimates (for the U.S.) based on aggregate time-series data.¹

A likely reason for the limited interest in these costs is that they are uniformly thought to be negligible.² As we show in this paper, this presumption is, in general, incorrect. In particular, this is the case

(I) when accounting for the fact that, as shown by recent experience, zero cannot be taken as the effective lower bound (ELB) for interest rates: as we show, allowing for an empirically plausible ELB materially increases the welfare costs compared to the standard benchmark of zero; and especially

(II) when considering countries such as Sweden, Switzerland, and the Euro area, which, for any level of interest rates, hold significantly larger amounts of M1 (as fractions of GDP) than (e.g.) the U.S., or Canada. E.g., in recent years M1 has fluctuated between 40 and 45 per cent of GDP in the U.S.,³ and between 30 and 40 per cent in Canada, whereas it has been around 70 per cent in the Euro area, between 50 and 60 per cent in Sweden, and between 80 and 100 per cent in Switzerland. These larger M1 balances as fractions of GDP automatically map into greater welfare costs for any level of interest rates.

We estimate the welfare costs of inflation based on aggregate time series data for the U.S., the Euro area, Canada, and three European countries, the U.K., Switzerland, and Sweden. The Euro area, Swiss and Swedish experience is especially interesting because interest rates there have recently plunged below zero. This shows that the standard practice in the literature of taking zero as the ELB for the computation

¹Lucas (2000) computed the welfare costs of inflation for the U.S. based on a calibrated model for the demand for M1. Alvarez and Lippi (2009) estimated a model of the transaction demand for *cash* based on Italian micro data, and then recovered the welfare costs from the estimated money demand schedule. Being based on the demand for cash (rather than M1) these costs are however not comparable to either those computed by Lucas (2000), or to Ireland’s (2009) estimates (as well as those in the present work). The same holds for Attanasio *et al.* (2002), who estimated money demand schedules based on micro data for Italian households, and then used their estimates to back out the costs of inflation.

²E.g., Feldstein (1997, p. 145) characterized these costs as ‘small relative to the other effects that have been discussed in this paper’.

³As we discuss below, the M1 aggregate we use for the U.S. is Lucas and Nicolini’s (2015) ‘New M1’, which also includes Money Market Deposit Accounts.

of welfare costs is unwarranted, and raises the issue of whether allowing for a plausible negative ELB might make a material difference. As we show, this is indeed the case. Further, in any of these countries the demand for M1 at extremely low, or even negative interest rates has (so far) exhibited no obvious difference compared to its behaviour at higher interest rates. The same holds for the U.K. and Canada: in either country the demand for M1 below a ‘low interest rate threshold’ (which we take to be 6%) exhibits no appreciable difference compared to its behavior for samples above the threshold. 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. For the U.S. we do indeed detect strong evidence of non-linearities at low interest rates: our results, however (and, in fact, even the simple visual evidence), suggest that at low interest rates the demand for M1 is comparatively *steeper*, rather than flatter. We provide a simple explanation for Mulligan and Sala-i-Martin’s (2000) finding, by showing mathematically that if the true money demand specification is the one we work with, their methodology automatically produces spurious evidence of a flatter demand curve at low interest rates.

For policy purposes, our main finding that—especially in Europe—the welfare costs of inflation originating from lack of liquidity satiation are not negligible has two main implications:

first, these costs should be taken into account for the purpose of choosing an appropriate numerical target for inflation.

Second, ceteris paribus, inflation targets for the Euro area, and for countries such as Switzerland and Sweden should be set at a comparatively lower level.

The paper is organized as follows. The next section discusses the money demand specifications we use, and the associated welfare cost functions. Section 3 briefly discusses the data, whose properties and sources are outlined in detail in Appendix A, whereas Section 4 analyzes their unit root and cointegration properties, and it explores the issue of whether money demand curves may exhibit evidence of instability (in particular, non-linearities at low interest rates). Section 5 estimates money demand curves, and from these backs out the welfare cost functions, thus obtaining, for each interest rate level, the bootstrapped distribution of the welfare costs (and therefore median estimates, and confidence intervals) expressed in percentage points of GDP. Section 6 discusses the policy implications of our results. Section 7 concludes, and outlines possible directions for future research.

2 Money Demand Specifications and Welfare Cost Functions

2.1 Alternative money demand specifications

In what follows we consider three alternative money demand specifications. The first two, which have so far dominated the literature on money demand, are Cagan’s (1956) *semi-log* and Meltzer’s (1963) *log-log*, which are given by

$$\ln\left(\frac{M_t}{Y_t}\right) = \ln(A) - \xi R_t \quad (1)$$

and

$$\ln\left(\frac{M_t}{Y_t}\right) = \ln(B) - \eta \ln(R_t) \quad (2)$$

respectively, where M_t , R_t , and Y_t are the nominal money stock—which throughout the entire paper we take to be M1—the nominal interest rate, and nominal GDP, respectively; η and ξ are the elasticity and semi-elasticity of money demand, respectively; and A and B are constants.

The third specification is the one Benati *et al.* (2019) label as the ‘Selden-Latané specification’, from Selden (1956) and Latané (1960), who, to the very best of our knowledge, are the only ones to ever estimate specifications in which the level of money velocity, Y_t/M_t , is postulated to evolve as a linear function of the nominal interest rate,

$$V_t \equiv \frac{Y_t}{M_t} = C + \delta R_t, \quad (3)$$

thus implying that the ratio between nominal money balances and nominal GDP is given by

$$\frac{M_t}{Y_t} = \frac{1}{C + \delta R_t} \quad (4)$$

As shown by Benati *et al.* (2019, Section 2), at low interest rates (such as those analyzed herein) the Selden-Latané specification provides a very good approximation to a log-log specification with borrowing constraints. Intuitively, as the interest rate falls towards zero, the log-log specification implies that the demand for money balances tends to infinity, which is only possible if, in the limit, agents can borrow infinite amounts of money. On the other hand, within the more plausible scenario in which agents are subject to constraints on the maximum amount they can borrow, the log-log specification morphs into a functional form which is closely approximated by the linear relationship (3) between velocity and the nominal interest rate. A key feature of the Selden-Latané specification is that (like the semi-log) it implies a positive ‘satiation level’ for money balances as a fraction of GDP at $R_t=0$, which, in equation (4), is given by $1/C$. On the other hand, as R_t tends to infinity, money balances as a fraction of GDP shrinks to zero.

As discussed by Benati *et al.* (2019, Section 2), the three money demand specifications (1), (2), and (4) arise as particular cases of a theoretical framework in which a representative agent, subject to a cash-in-advance constraint, makes transactions in each period *via* either cash or demand deposits.⁴ Depending on the characteristics of the technology used to make the portfolio adjustments (i.e., the technology describing the ‘trips to the bank’), the framework can generate several alternative money demand specifications (e.g., one particular case is the classic Baumol (1952) and Tobin (1956) ‘square root’ formula).

2.2 The associated welfare cost functions

The welfare costs of inflation associated with a specific level of the nominal interest rate R_t , and for a given ELB R_0 , are given by⁵

$$WC(R_t, R_0) = \int_{R_0}^{R_t} m(r)dr - [R_t - R_0]m(R_t) \quad (5)$$

where $m(R_t)$ is the demand for real money balances as a function of the nominal interest rate R_t , and R_0 is the ELB for the nominal rate. Working as in Lucas (2000), it can be easily shown that for the semi-log and log-log specifications the welfare cost of inflation associated with a specific level of R_t , for a given R_0 , are given by

$$WC_{SL}(R_t) = \frac{A}{\xi}e^{-\xi R_0} + R_0 A e^{-\xi R_t} - \frac{A}{\xi}(1 + \xi R_t)e^{-\xi R_t} \quad (6)$$

and

$$WC_{LL}(R_t) = \frac{B\eta}{1-\eta}R_t^{1-\eta} + BR_0R_t^{-\eta} - \frac{B}{1-\eta}R_0^{1-\eta} \quad (7)$$

respectively. For $R_0=0$, the two expressions take the values found in Lucas (2000, p. 251), i.e. $A\frac{\eta}{1-\eta}R_t^{1-\eta}$ and $\frac{B}{\xi}[1 - (1 + \xi R_t)e^{-\xi R_t}]$, respectively. By the same token, it can be easily shown that the welfare cost function associated with the Selden-Latané specification (4) is given by


$$WC_{SE}(R_t) = \frac{1}{\delta} \ln\left(\frac{C + \delta R_t}{C + \delta R_0}\right) - \frac{R_t - R_0}{C + \delta R_t} \quad (8)$$

We now turn to discussing the data.

⁴Likewise, Belongia and Ireland (2019) derive the three specifications as special cases of Sidrauski’s (1967a, 1967b) ‘money in the utility function’ framework.

⁵This expression is the same as Lucas’ (2000, p. 250) equation (2.1), with the only difference that he assumed $R_0=0$, whereas we allow for $R_0 \leq 0$.

I: M1 velocity and the long rate



II: M1 velocity and the short rate




Figure 1a. M1 velocity and long- and short-term nominal interest rates




Figure 1b United States: M1 velocity and long- and short-term nominal interest rates





Figure 2 Scatterplots of nominal M1 over nominal GDP against the long rate

3 A Look at the Raw Data

Throughout the entire paper we will work with quarterly post-WWII data,⁶ and we will consider M1 as the relevant monetary aggregate.⁷ The series and their sources are described in detail in Appendix A. All of the series are standard, with one exception: for the U.S., instead of working with the standard M1 aggregate produced by the Federal Reserve—for which, as it has been repeatedly documented from Friedman and Kuttner (1992) to Benati *et al.* (2019), no stable money demand relationship exists—we work with Lucas and Nicolini’s (2015) ‘New M1’ aggregate, which is obtained by adding Money Market Deposit Accounts (MMDAs) to the standard M1 aggregate.⁸ As shown by Benati *et al.* (2019, Section 2), based on New M1 a stable long-run demand is indeed detected.

Figures 1*a* and 1*b* plot M1 velocity, computed as the ratio between nominal GDP and nominal M1, together with either a long- or a short-term nominal interest rate, whereas Figure 2 presents the data in a more standard fashion, as scatterplots of the ratio between M1 and GDP against a nominal rate. The evidence in Figures 1*a-1b* clearly suggests the following four facts:

- (i) M1 velocity and the two interest rate series are all I(1);
- (ii) velocity is cointegrated with either of the two interest rates;
- (iii) as documented by Benati (2020), up to a linear transformation, velocity is, essentially, the stochastic trend of the short rate; and
- (iv) velocity’s relationship with the long rate is significantly closer and stronger than that with the short rate.

Econometric evidence on (i) and (ii) will be discussed in Section 2. Evidence on (iii) for the six countries analyzed herein, as well as for additional ?? countries since World War I, can be found in Benati (2020).

The fact that M1 velocity is, to a first approximation, the permanent component of the short rate explains why—as it is so starkly apparent from Figures 1*a-1b*—its relationship with the long rate is significantly closer and stronger than that with the short rate. The intuition is straightforward. Empirically, as it is well known, short- and long-term nominal interest rates are cointegrated, so that their I(1) components are driven by a common permanent shock. Since the long rate is much closer to the common stochastic trend than the short rate—in the sense that the latter contains a

⁶In a previous version of the paper we worked with the annual long-run series from Benati *et al.*’s (2019) dataset. Our preference, within the present context, for working with quarterly data is motivated by the fact that the estimates of the welfare costs of inflation are significantly more precise.

⁷In Appendix C we motivate our choice of working with ‘simple-sum’ M1 aggregates, as opposed to their Divisia counterparts.

⁸Augmenting 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. The rationale for doing so is that MMDAs perform an economic function very similar to that of the standard demand deposits which are included in the official M1 series.

sizeable transitory component, whereas the former typically has a (near-) negligible one—it is also much closer to M1 velocity. Another way of putting this is that *both* M1 velocity *and* the long rate are very good *proxies* for the common stochastic trend in the system, whereas the short rate is not. It should therefore come as no surprise that the relationship between velocity and the long rate is significantly stronger than that with the short rate.

For the U.S., focusing on the relationship between velocity and the long rate also highlights three distinct periods over the post-WWII era, which are instead much less apparent if we concentrate on the relationship with the short rate:

(i) a former period up until the introduction of MMDAs, in which velocity and the long rate had exhibited a very close co-movement at the low frequencies.⁹

(ii) A ‘transition period’ between the introduction of MMDAs, in 1983Q1, and the mid-1990s. As it is apparent from the first panel of Figure 1*b*, during this period the introduction of MMDAs had temporarily and significantly disrupted the previous close relationship between the two series, with the long rate falling from about 14 to about 6 per cent, and velocity oscillating instead between 4 and 5 per cent.

(iii) A last, and still ongoing period, starting in the mid-1990s, in which a remarkably strong relationship between velocity and the long rate has reasserted itself, with the two series moving largely in lockstep.¹⁰

It is important to stress how neither (ii) nor (iii) are clearly apparent from the second panel of Figure 1*b*: the fact that the short rate contains a dominant transitory component¹¹ blurs the distinction between the second and the third periods (henceforth, periods II and III). The evidence in the first panel of Figure 1*b*, on the other hand, suggests that working with either the entire post-WWII sample, or with the sample following the introduction of MMDAs, will distort inference. In fact, as we will show in Section 2.2, whereas Johansen’s test produces very strong evidence of cointegration between velocity and the long rate for the period since the mid-1990s, it does *not* detect it for the full period following the introduction of MMDAs. This confirms what the evidence in the first panel of Figure 1*b* suggests: periods II and III are indeed very different, and they should not be mixed. For the purposes of this paper (in particular, the computation of the welfare costs of inflation) we will therefore mostly focus the period since the mid-1990s.

Going forward, although in the Appendix we will also present results based on the short rate, in the main body of the paper, based on the previous discussion, we will exclusively work with the long rate. Although it is standard practice, in the

⁹In fact, as we will discuss in Section 2.2, for this period both Johansen’s and Wright’s tests detect cointegration between the two series.

¹⁰In passing, it is worth stressing how this provides strong confirmation of the meaningfulness of working with Lucas and Nicolini’s (2015) New M1 aggregate. What the first panel of Figure 1*b* suggests is that, in fact, New M1 is the equivalent, for the period after the introduction of MMDAs, of the standard M1 aggregate for the previous period. In Section 4.4 we will provide further confirmation of this, based on the estimated coefficients on the long rate for the two aggregates.

¹¹See Benati (2020, Figure 3*a*).

literature on money demand, to take a short rate as the opportunity cost, in fact there is no strong rationale for doing so. As discussed (e.g.) by Goldfeld and Sichel (1990, Section 3.2.3, pp. 320-322) different authors have proposed several alternative *proxies* for the opportunity cost, with some (e.g.) proposing to use the entire term structure of interest rates, or stock prices' rate of return. The main proponent of using the long rate was Michael Hamburger,¹² who claimed¹³ that the instability in U.S. short-run money demand documented by Goldfeld (1976) disappeared when (a) working with velocity (i.e., imposing unitary income elasticity), and (b) taking the long rate as the opportunity cost. Goldfeld and Sichel (1990, p. 321) were skeptical of Hamburger's (1977) results, stating that 'Hamburger's model contains a number of constraints (i.e., a unitary income elasticity) that are not warranted by the data'. In fact, as the evidence in the top row of Figure 1a and the first panel of Figure 1b suggests, and as the econometric results in Section 2.2 will confirm, working with velocity uncovers an extraordinarily strong and stable low-frequency relationship with both the short and especially the long rate, and suggests that the latter exhibits the closest relationship with velocity.

It is important to stress once again that, from a time-series perspective, we do indeed have a strong rationale for using the long rate. As previously mentioned, both M1 velocity and the long rate are very close to the common stochastic trend in the system, whereas the short rate—whose fluctuations are dominated by transitory shocks—is not. As a result, since our objective is to estimate a strong and stable relationship between velocity and a measure of the opportunity cost, the long rate is the natural choice.

We now turn to discussing the time-series properties of the data.

4 Time-Series Properties of the Data

4.1 Evidence from unit root tests

Table A.1 in the Appendix reports results from Elliot *et al.* (1996) unit root tests for either the levels or the logarithms of M1 velocity, the long rate, and the short rate. In short, for all countries the null of a unit root cannot be rejected for any of the series. In searching for a long-run cointegration relationship between velocity and interest rates, in the next section we will therefore proceed as follows. First, taking the results from unit root tests literally—i.e, as indication that the series do contain exact unit roots—we will test for cointegration based on Johansen's tests, which are predicated in the assumption that the series are indeed I(1). Since, however, a plausible alternative interpretation of the results in Table A.1 is that the series are local-to-unity—in which case, as shown by Elliot (1998), tests such as Johansen's

¹²We wish to thank Peter Ireland for alerting us to Hamburger's largely forgotten work.

¹³See Hamburger (1977).

tend to perform poorly—we will search for cointegration based on Wright’s (2000) test, which is valid for both exact unit roots, and roots which are local-to-unity. All of the technical details about the implementation of the tests are identical to Benati *et al.* (2019) and Benati (2020), which the reader is referred to.

Table 1 Bootstrapped p-values^a for Johansen’s maximum eigenvalue^b tests for (log) M1 velocity and (the log of) a long-term rate				
<i>Country</i>	<i>Period</i>	Money demand specification:		
		<i>Selden-Latané</i>	<i>Semi-log</i>	<i>Log-log</i>
United States	1953Q2-1982Q4	0.085	0.500	0.177
	1983Q1-2018Q3	0.330	0.201	0.179
	1996Q1-2018Q3	0.002	9.0e-4	0.042
United Kingdom	1955Q1-2016Q4	0.090	0.108	0.465
Canada	1967Q1-2019Q2	0.002	0.078	0.013
Switzerland	1980Q1-2019Q2	0.250	0.130	– ^c
Sweden	1998Q1-2019Q2	0.047	0.050	0.640
Euro area	1999Q1-2019Q2	0.453	0.356	0.777

^a Based on 10,000 bootstrap replications. ^b Null of 0 *versus* 1 cointegration vectors. ^c The last observations for the interest rate are negative.

Table 2 Results from Wright’s tests: 90% confidence interval for the second element of the normalized cointegration vector, based on systems for (log) M1 velocity and (the log of) a long-term rate^a				
<i>Country</i>	<i>Period</i>	Money demand specification:		
		<i>Selden-Latané</i>	<i>Semi-log</i>	<i>Log-log</i>
United States	1953Q2-1982Q4	[-0.653; -0.365]	[-0.141; -0.090]	[-0.733; -0.338]
	1983Q1-2018Q3	NCD	[-0.135; -0.113]	[-0.567; -0.301]
	1996Q1-2018Q3	[-0.757; -0.670]	[-0.223; -0.189]	[-0.804; -0.731]
United Kingdom	1955Q1-2016Q4	[-0.548; -0.412]	[-0.131; -0.095]	[-0.841; -0.543]
Canada	1967Q1-2019Q2	NCD	[-0.124; -0.100]	[-0.639; -0.534]
Switzerland	1980Q1-2019Q2	[-0.408; -0.360]	[-0.212; -0.185]	– ^b
Sweden	1998Q1-2019Q2	[-0.356; -0.264]	[-0.172; -0.102]	[-0.607; -0.177]
Euro area	1999Q1-2019Q2	[-0.175; -0.103]	[-0.396; -0.157]	[-1.277; -0.347]

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

4.2 Cointegration properties of the data

Table 1 reports, for any of the three money demand specifications discussed in Section 2, bootstrapped p -values for Johansen’s maximum eigenvalue test of 0 versus 1 cointegration vectors,¹⁴ whereas Table 2 reports the 90% confidence intervals for the second element of the normalized cointegration vector based on Wright’s (2000) test. Table A.2 and A.3 in the Appendix report the corresponding results based on the short rate. Two main findings emerge from the Tables:

(i) Wright’s test detects cointegration near-uniformly across the board based on any of the three money demand specifications.

(ii) As for Johansen’s tests, evidence of cointegration is very strong based on the Selden-Latané specification, whereas it is weaker based on the other two functional forms. In particular, based on Selden-Latané cointegration is not detected for the U.S. based on the entire period following the introduction of MMDAs (whereas it is strongly detected for the period since the mid-1990s), and for Switzerland and the Euro area. As for the U.S., this just confirms what the visual evidence in the first panel of Figure 1*b* suggests: as we previously discussed, what we labelled as periods II and III are indeed very different, and they should not be mixed. The former is just a transition period, whereas in the latter the relationship between velocity and the long rate has fully reasserted itself. As for the Euro area, and especially Switzerland (in the light of the visual evidence in Figure 1*a*), we regard these failures to detect cointegration as flukes, possibly due to the issues originally discussed by Engle and Granger (1987), i.e. a combination of (i) small sample size (this is especially the case for the Euro area), and (ii) a high persistence of the cointegration residual.¹⁵

Based on the previous discussion, in what follows (1) we will therefore proceed under the assumption that cointegration is there in all samples (with the exception, for the U.S., of the sub-sample 1983Q1-2018Q3, which we will not further analyze), and (2) we will largely focus on the Selden-Latané specification, for which evidence of cointegration is stronger.

We now turn to the issue of stability of the cointegration relationships.

4.3 Testing for stability of the cointegration relationships

Table 3 reports bootstrapped p -values for Hansen and Johansen’s (1999) tests for stability in the cointegration vector.¹⁶ At the 10 per level stability is rejected in a

¹⁴The corresponding results from the trace test are qualitatively the same, and they are available upon request.

¹⁵Benati *et al.* (2019) present extensive Monte Carlo evidence on the empirical relevance of Engle and Granger’s point, by showing that a combination of (i) and (ii) can cause Johansen’s tests to fail to detect cointegration, when it is there, a high or very high fraction of the time.

¹⁶On 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 a steady-state

single case, i.e. for the U.S. based on the most recent sub-sample and the log-log specification, whereas based on the Selden-Latané specification the lack of rejection is borderline. In the light of the evidence in Table 1, a plausible explanation for the former result is simply that the log-log specification is not the correct one.¹⁷ In Section 4.6 we will explore the issue of which specification appears as the most plausible, and we will argue that the functional form preferred by the data is the Selden-Latané.

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 long-term rate				
Country	Period	<i>Selden-Latané</i>	<i>Semi-log</i>	<i>Log-log</i>
United States	1953Q2-1982Q4	0.661	0.525	0.978
	1996Q1-2018Q3	0.106	0.748	0.055
United Kingdom	1955Q1-2016Q4	0.574	0.498	0.372
Canada	1967Q1-2019Q2	0.432	0.619	0.717
Switzerland	1980Q1-2019Q2	0.629	0.457	— ^b
Sweden	1998Q1-2019Q2	0.272	0.321	0.324
Euro area	1999Q1-2019Q2	0.805	0.843	0.942

^a Based on 10,000 bootstrap replications.

Although Hansen and Johansen's (1999) tests detect essentially no evidence of instability in the cointegration vector, for the specific purpose of testing whether money demand curves might be flatter at low interest rates, these results should be discounted for (at least) two reasons.

First, as discussed by Bai and Perron (1998, 2003), when a coefficient experiences two breaks in opposite directions (e.g., first an increase, and then a decrease), break tests which have not been explicitly designed to search for *multiple* breaks may have a hard time detecting the first break to begin with. Within the present context this could be relevant for two countries, the U.S. and the U.K.: in both cases the long rate had been below 6% (which, as we will discuss below, we will take as the relevant threshold) at the beginning of the sample; it then significantly increased above 6% 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

characterized by a specific value of the long rate. Finally, we eschew Hansen and Johansen's (1999) fluctuation tests because, as shown by Benati *et al.* (2019) *via* Monte Carlo, they exhibit, overall, a significantly inferior performance compared to the Nyblom tests for stability in the cointegration vector and loading coefficients.

¹⁷Based on very long samples of annual data since World War I, indeed, Benati *et al.* (2019, Figure 6) present simple but powerful visual evidence that for low-inflation countries such as those analyzed herein the log-log specification is highly implausible, whereas the Selden-Latané is the preferred one.

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 long rate. In fact, as we discuss in the next sub-section, this is indeed the case.

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

A strand of literature—see, first and foremost, Mulligan and Sala-i-Martin (2000)—has argued that, at low interest rates, money demand exhibits sizeable nonlinearities due to the presence of fixed costs associated with the decision to participate, or not to participate, in financial markets.¹⁸ This implies that 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 do not participate in financial markets. The implication is that it should not be possible to reliably estimate money demand functions (and therefore the welfare costs of inflation) based on aggregate time series data, as only the use of micro data allows to meaningfully capture the nonlinearities associated with the cost of participating in financial markets.

Figure 3 shows evidence on this for the three countries (U.S., U.K., and Canada) with sufficiently long continuous samples both above and below a ‘low interest rate’ threshold, which in what follows we take to be 6%.¹⁹ The sub-samples with the long rate consistently above 6% are 1969Q1-1982Q3 for the U.S., 1965Q1-1998Q1 for the U.K., and 1967Q4-1997Q2 for Canada. The corresponding sub-samples with the long rate consistently below 6% are 1998Q2-2016Q4 for the U.K. and 2000Q3-2019Q2 for Canada, whereas for the U.S. there are two, 1947Q1-1968Q4 and 1996Q1-2018Q3. As previously mentioned, we exclude from the analysis for the U.S. the ‘transition period’ 1982Q4-1995Q4, during which the relationship between M1 velocity and the long rate had been thrown temporarily out of kilter by the introduction of MMDAs.

¹⁸The 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 discontinuity) at low interest rates.

¹⁹Mulligan and Sala-i-Martin (2000), working with *short*-term interest rates, take 5% as the threshold. Since in our dataset the average long-short spread has been equal to about 1% (or, in the case of the U.S., just slightly higher), we consider 6% as the corresponding threshold for long-term rates.

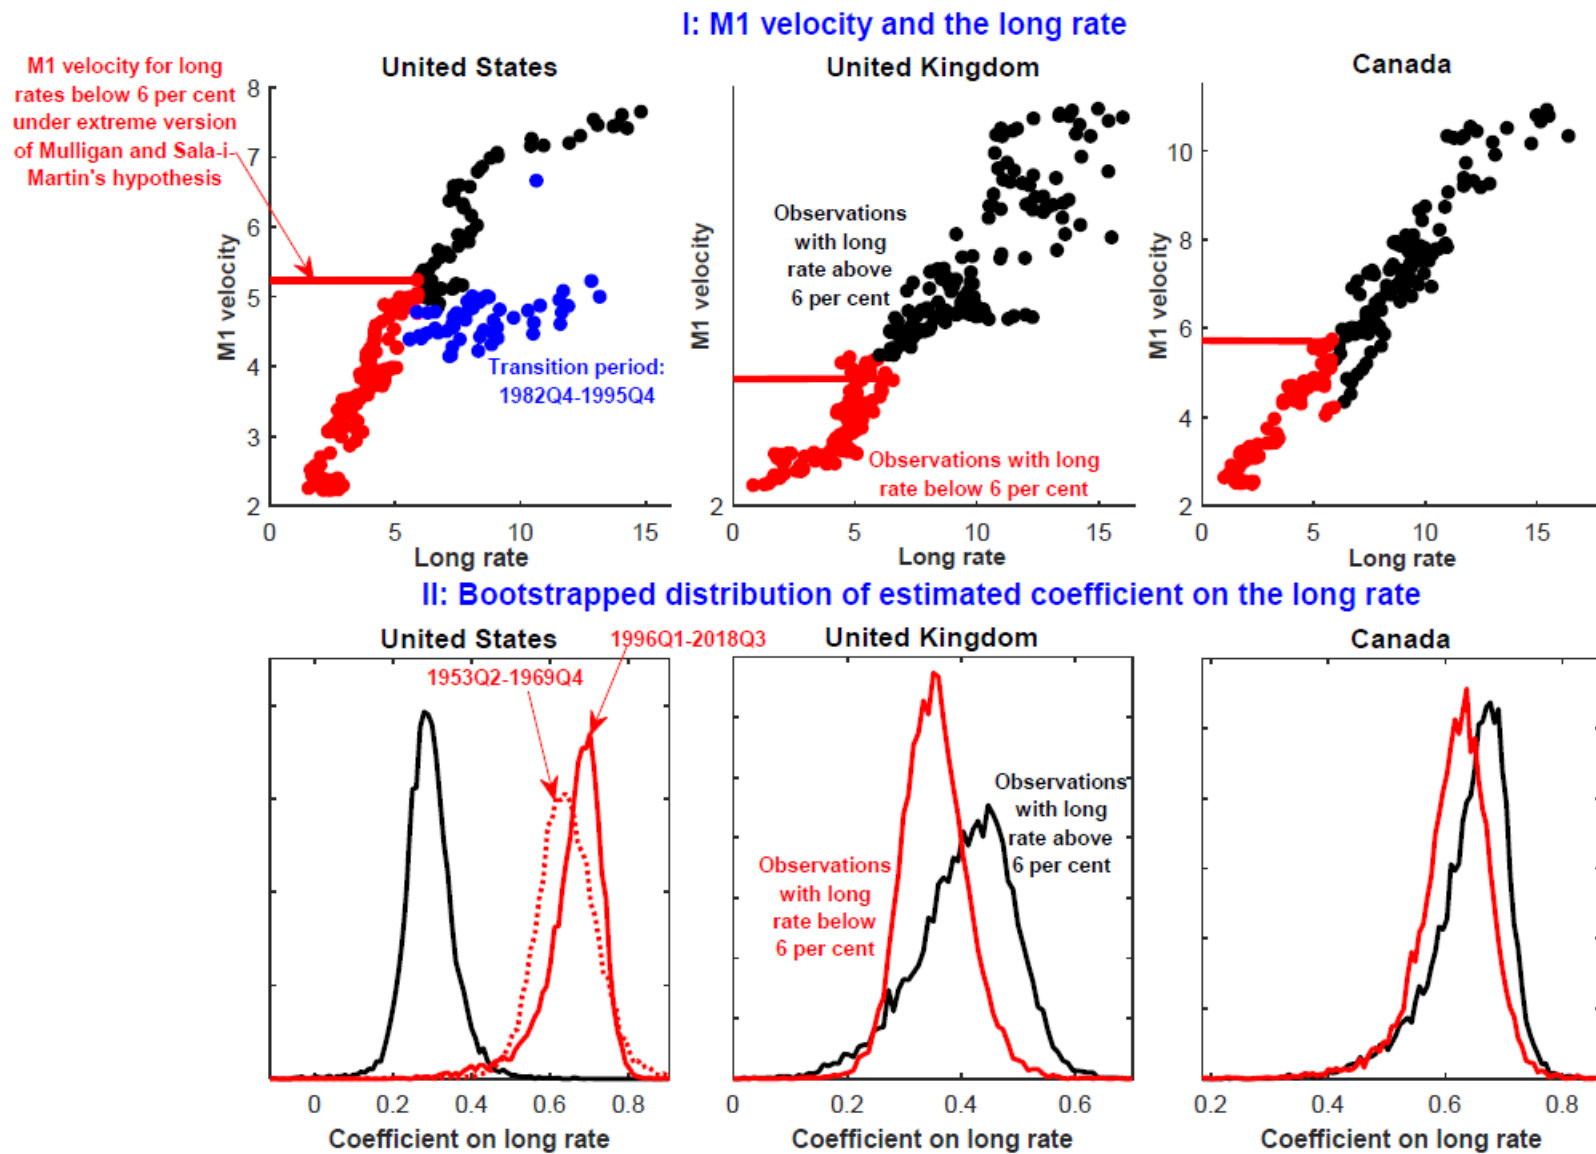


Figure 3 M1 velocity and long-term nominal interest rates: observations with the long rate above and below 6 per cent

The top row shows, for any of the three countries, scatterplots of M1 velocity and the long rate, with the observations with the long rate above and below the 6% threshold being shown in black and red, respectively. The three panels also show an horizontal red line corresponding to an extreme version of the non-linearity hypothesis, in which when the long rate falls below the 6% threshold 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 that 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,

(i) for Canada visual evidence strongly suggests that the relationship between velocity and the long rate is the same at all interest rate levels.

(ii) For the U.S. the *only* evidence that money demand might be flatter at low interest rates comes from the transition period 1982Q4-1995Q4, which, as we argued, should be excluded from the analysis. The evidence from the rest of the sample, on the other hand, seems to suggest that, if anything, for long rates below 6% the relationship had, and has in fact been *steeper*, rather than flatter. This evidence of non-linearity should however be regarded with some skepticism, since it crucially hinges on the highly volatile period between mid-1979 and 1982Q3: excluding those three years, the relationship appears to be the same at all interest rate levels.

(iii) Finally, for the U.K. there seems to be some evidence that the relationship might indeed have been flatter for long rates below about 5%.

Table 4 Estimated coefficients on the long rate in Selden-Latane⁷ specifications for samples with the long rate above and below 6 per cent^a			
<i>Country</i>	$P(\delta_{R<6} < \delta_{R>6})$	Based on samples with long rate:	
		<i>below 6 per cent</i>	<i>above 6 per cent</i>
		Median and 90% confidence interval	Median and 90% confidence interval
United States, I	0.001	0.646 [0.529; 0.775]	0.293 [0.205; 0.399]
United States, II	0.004	0.680 [0.513; 0.758]	0.293 [0.205; 0.399]
United Kingdom	0.699	0.356 [0.272; 0.456]	0.420 [0.252; 0.540]
Canada	0.626	0.625 [0.510; 0.706]	0.658 [0.513; 0.729]
^a Based on 10,000 bootstrap replications.			
Samples with long rate below and above 6 per cent: United States, I: 1953Q2-1968Q4 and 1969Q1-1982Q3; United States, II: 1996Q1-2018Q3 and 1969Q1-1982Q3; United Kingdom: 1998Q2-2016Q4 and 1965Q1-1998Q1; Canada: 2000Q3-2019Q2 and 1967Q4-1997Q2			

The second row of Figure 3 reports 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 long rate in the Selden-Latané specification $V_t = C + \delta R_t + \varepsilon_t$.²⁰ The evidence, reported in Table ??, provides support to the visual impression from the scatterplots in the top row. For Canada and the U.K. the median estimates of δ for the two sub-samples are very close, and the p -values for testing the null hypothesis that δ might have been smaller at low interest rates are equal to 0.625 and 0.699. For the U.S., on the other hand, equality of the slopes at high and low interest rates is strongly rejected based on either of the two sub-samples with the long rate below 6%, with p -values equal to 0.001 and 0.004. In line with the visual evidence in the top row, however, the estimates of δ based on either of the two sub-samples with the long rate below 6% are *larger* than the corresponding estimate based on the sub-sample with the long rate above 6%, thus rejecting the notion that the demand for M1 might be flatter at low interest rates. Interestingly, both the median estimates of δ for the two sub-samples with low interest rates, and, in fact, their entire bootstrapped distributions are quite remarkably close. This confirms the previously discussed visual impression from the first panel of Figure 1b: after the transition period 1982Q4-1995Q4, in which the relationship between velocity and the long rate had been disturbed by the introduction of MMDAs, the *very same relationship* which had prevailed up until the end of the 1960s has in fact fully reasserted itself.

Based on this evidence, in Section 5 we will proceed as follows. For both Canada and the U.K. we will estimate money demand curves, and then extract from them the corresponding welfare cost functions, based on the full sample periods. As for the U.S., on the other hand, we will exclusively work with either of the two sub-samples with the long rate below 6%, whereas we will eschew the sub-sample with the long rate above 6%. The obvious reason is that our objective is to characterize the welfare costs of inflation for *low-inflation* monetary regimes. Under this respect, for the U.S., its experience at high levels of the interest rate is, in the light of the previously discussed evidence, uninformative.

Before estimating money demand curves and welfare costs, however, it is worth addressing the issue of how to rationalize Mulligan and Sala-i-Martin’s finding of a smaller elasticity at low interest rates.

4.4.1 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

²⁰The methodology is the same we described before. 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; finally, we characterize uncertainty about the cointegration vector by bootstrapping the VECM as in Cavaliere *et al.* (2012).

rate (say, 1%) is associated with a larger change in its *logarithm* at low interest rates than it is at higher interest rates,²¹ 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 will automatically produce 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 (4), estimating the log-log specification (2) produces the following theoretical value of the *estimated* elasticity

$$\frac{d \ln \left(\frac{M_t}{Y_t} \right)}{d \ln R_t} = -\frac{\delta R_t}{C + \delta R_t}, \quad (9)$$

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

$$\frac{d \ln \left(\frac{M_t}{Y_t} \right)}{d \ln R_t} = -\gamma R_t.$$

which tends to $-\infty$ for $R_t \rightarrow \infty$, tends to 0 for $R_t \rightarrow 0$, and 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.

4.5 Which specification do the data prefer?

As previously discussed, the results from Johansen's tests in Table 1 suggest that the data tend to somehow prefer the Selden-Latané specification to either the semi-log or the log-log. In this sub-section we perform a more systematic model comparison exercise. Since it is not possible to nest the three money demand specifications into a single encompassing one, we proceed as follows. We start from the comparison between the semi-log and the log-log. Intuitively, the comparison between (1) and (2) 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 long rate, or by its logarithm. For each country we therefore regress $\ln(M_t/Y_t)$ on a constant, p lags of itself, and p lags of either the level of the long rate or its logarithm. A natural way of interpreting these equations is the following. Under the assumption that cointegration is indeed there for all countries, and based on any specification,²² both $Y_t^{SL} = [\ln(M_t/Y_t)$

²¹For example, $\ln(9)-\ln(10)=-0.105$, whereas $\ln(2)-\ln(3)=-0.406$.

²²If this assumption did not hold, the entire model comparison exercise would obviously be meaningless.

$R_t]$ ' 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). So the equations we are estimating can be thought of as the 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 :

$$\ln\left(\frac{M_t}{Y_t}\right) = \alpha + \sum_{j=1}^p \beta_j \ln\left(\frac{M_{t-j}}{Y_{t-j}}\right) + \sum_{j=1}^p \beta_j \left(\frac{R_{t-j}^\theta - 1}{\theta}\right) + \varepsilon_t \quad (10)$$

We estimate (10) *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.²³ All other estimation details are identical to Benati (2008), to which the reader is referred to.

Table 5a Model comparison exercise, semi-log <i>versus</i> log-log: mode of the log-likelihood in regressions of log velocity on lags of itself and either the long rate or its logarithm						
<i>Country</i>	$p = 2$		$p = 4$		$p = 8$	
	<i>Semi-log</i>	<i>log-log</i>	<i>Semi-log</i>	<i>Log-log</i>	<i>Semi-log</i>	<i>Log-log</i>
United States	328.4	326.2	322.8	320.3	304.7	303.1
United Kingdom	839.3	834.6	854.4	848.5	840.2	835.2
Canada	739.2	736.2	739.2	734.9	735.9	730.9
Sweden	307.6	305.0	299.3	296.1	284.0	281.6
Euro area	322.9	321.1	315.6	313.1	299.4	297.4

For Switzerland there is no comparison because the last observations for the long rate are negative.

²³So, to be clear, the proposal draw for the parameter vector $\beta, \tilde{\beta}$, is accepted with probability $\min[1, r(\beta_{s-1}, \tilde{\beta} | Y, X)]$, and rejected otherwise, where β_{s-1} is the current position in the Markov chain, and

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

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

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

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

Table 5a reports, for either specification, and for p equal to either 2, 4, or 8, the mode of the log-likelihood. The key result from the table is that for all countries, and for any lag order, regression (10) with $\theta=1$ (corresponding to the semi-log) uniformly dominates the corresponding regression with $\theta=0$ (corresponding to the log-log). This suggests that, for the low-inflation countries we are dealing here, the semi-log is preferred to the log-log. This is important because, as we will see in Section 5, the log-log tends to produce welfare cost functions materially different from those produced by the semi-log and the Selden-Latané, which are instead quite close.

Turning to the comparison between the semi-log and the Selden-Latané, we adopt the same logic as before, but this time we ‘flip’ specifications for velocity on their head, by regressing the *long 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

$$R_t = \alpha + \sum_{j=1}^p \beta_j R_{t-j} + \sum_{j=1}^p \beta_j \left[\frac{\left(\frac{Y_{t-j}}{M_{t-j}} \right)^\theta - 1}{\theta} \right] + \varepsilon_t \quad (11)$$

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

Table 5b Model comparison exercise, Selden-Latané <i>versus</i> semi-log: mode of the log-likelihood in regressions of the long rate on lags of itself and either the velocity or its logarithm						
<i>Country</i>	$p = 2$		$p = 4$		$p = 8$	
	<i>Selden-Latané</i>	<i>Semi-log</i>	<i>Selden-Latané</i>	<i>Semi-log</i>	<i>Selden-Latané</i>	<i>Semi-log</i>
United States	59.5	57.7	61.0	60.3	59.4	56.4
United Kingdom	26.3	27.7	27.9	28.7	26.3	28.0
Canada	47.6	45.9	49.5	47.3	50.8	46.6
Sweden	64.2	64.5	62.9	62.2	66.7	63.7
Switzerland	34.7	37.4	32.9	34.7	36.4	37.6
Euro area	62.4	62.3	60.7	61.3	56.3	56.9

Results are reported in Table 5b. Here evidence is not as clear-cut as in Table 5a. The Selden-Latané specification is preferred for any lag order for the U.S. and Canada, whereas the semi-log is preferred for the U.K. and Switzerland. Finally, results are not clear-cut for the Euro area and Sweden.

We now turn to discussing the estimated welfare cost functions.




Figure 4 Estimated welfare cost functions based on the Selden-Latané specification

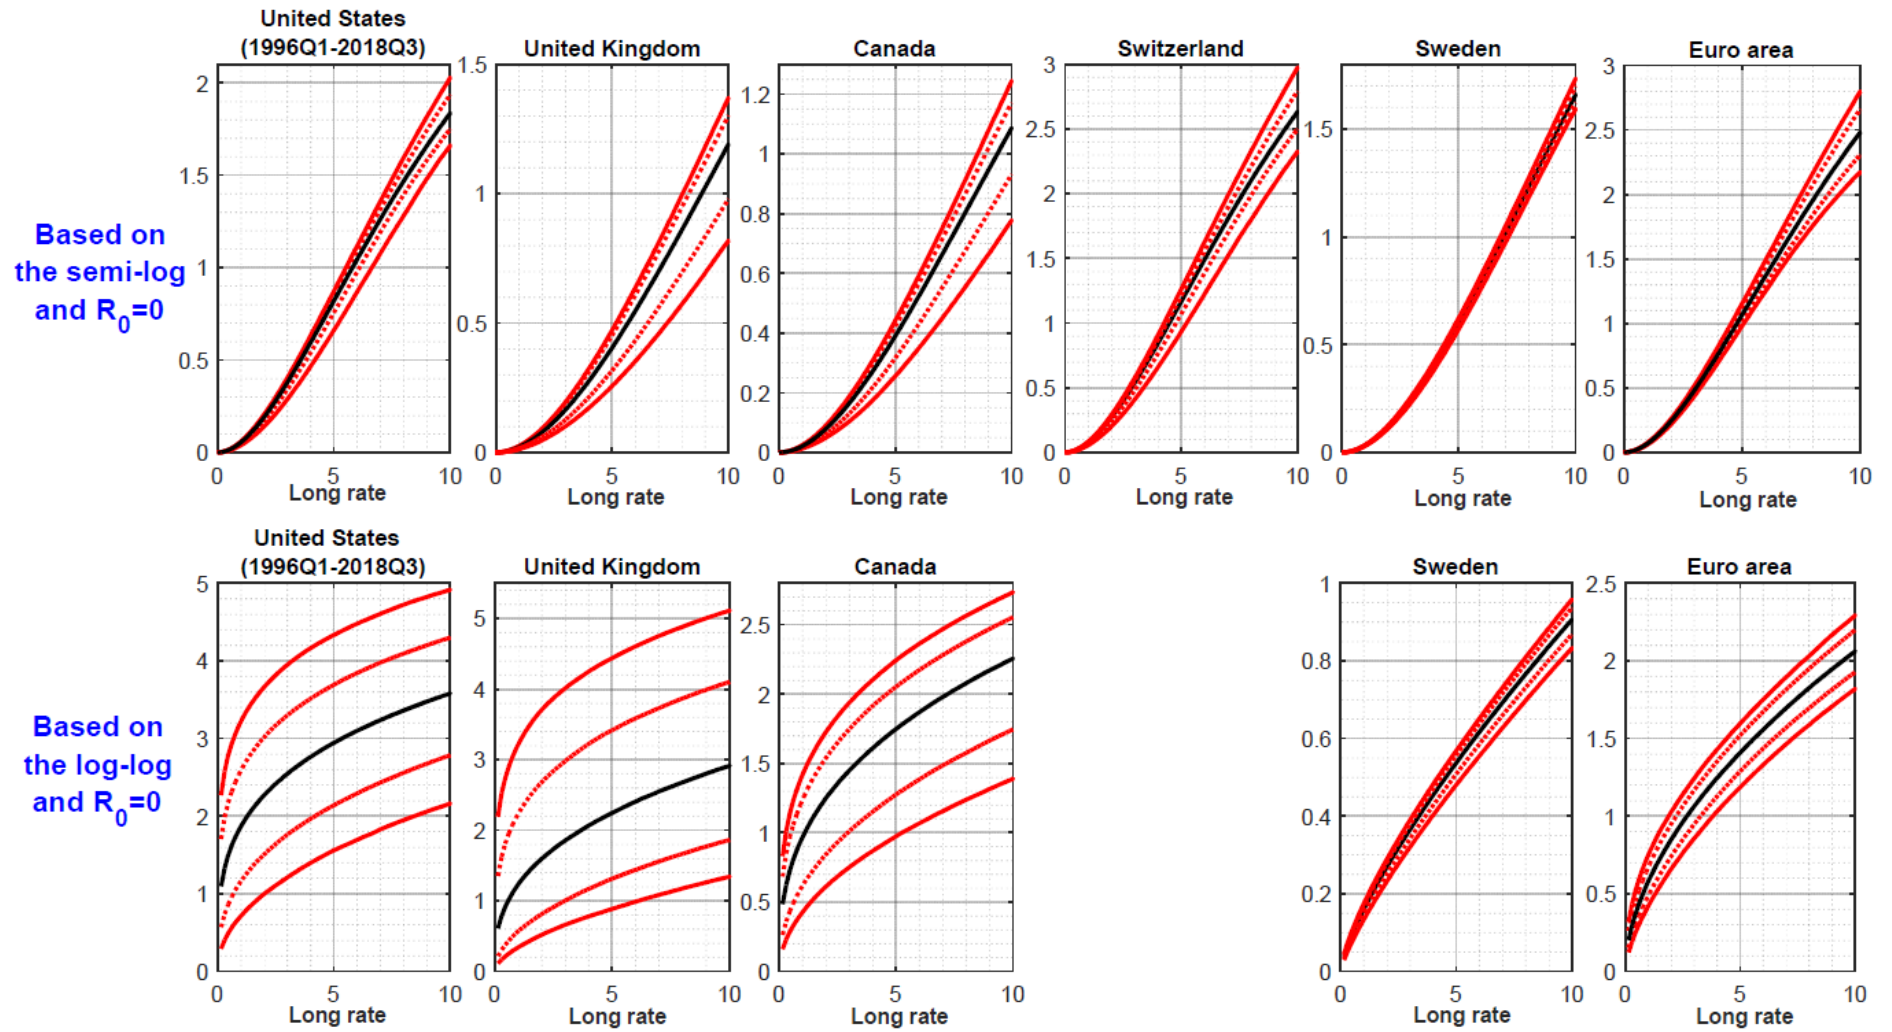


Figure 5 Estimated welfare cost functions based on the semi-log and the log-log, for $R_0=0$.

5 The Estimated Welfare Cost Functions

Figure 4 shows, based on the Selden-Latané specification, and for two alternative values of the ELB—either $R_0=0$ or $R_0=-1\%$ —the estimated welfare cost functions, with one- and two-standard deviations bootstrapped confidence bands. Figure 5 shows the corresponding welfare cost functions based on either the semi-log or the log-log for $R_0 = 0$, so that a comparison between the results in this figure and those in the first row of Figure 4 allows to gauge an idea about how alternative money demand specifications produce different estimates of the welfare costs. Finally, Table 6 reports, for the Selden-Latané, which is the specification preferred by the data, the estimated welfare costs assuming a plausible ELB of $R_0=-1\%$.

The methodology we use is standard. Following Luetkepohl (1991, pp. 370-371) we start by estimating *via* OLS the cointegrating regression corresponding to any of the three specifications, i.e. to either (1), (2), or (3). This gives us the point estimates we need in order to compute, based on either (6), (7), or (8) the point estimates of the welfare cost functions. Finally, we 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).

As Figure 4 shows, the welfare costs of inflation in the standard case in which the ELB is set to zero are quite substantial. For a long term interest rate of 5% they are between 1% and 2% of permanent consumption for the U.S., Switzerland and the Euro area, depending on the value assigned to the ELB. The costs are somewhat smaller for Canada, Sweden and the U.K.. The bottom panel of the Figure reports the results when considering a value for the lower bound of -1%, which we view as empirically plausible, given the recent experiences. In this case, the welfare costs are about twice as large. Table 6 shows the cost of increasing the target by 2% points, as a function of the initial level. Those range from 0.35 to 0.60% of permanent consumption.

In Figure 5 we reproduce the computations for the two most popular functional forms, even though one of them, the log-log, is clearly not the one preferred by the data. The reason is that this allows for a comparison with the results of Lucas (2000) and Ireland (2009). In most cases, the computations using the log-log deliver substantially larger costs than the ones we obtain based on either of the other two functional forms. However, as we argued in sub-section 4.5, the log-log is not the functional form preferred by the data. The welfare costs produced by the semi-log are slightly smaller than those based on the Selden-Latané specification. Thus, the two functional forms preferred by the data imply smaller welfare costs than those computed by Lucas (2000), although not as low as those estimated by Ireland (2009). Notice, however that the computations presented in Figure 6 are not directly comparable with Ireland's (2009), since he used a different monetary aggregate, and a different sample period.

Results for Switzerland and the Euro area show how 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 estimates of A , B , ξ and η are such as to satisfy the inequality

$$WC_{SL}(R_t) = \frac{A}{\xi} [1 - (1 + \xi R_t)e^{-\xi R_t}] > \frac{B\eta}{1 - \eta} R_t^{1-\eta} = WC_{LL}(R_t)$$

which is what we see in Figures 4 and 5 for Switzerland and the Euro area for $R_0=0$.

Table 6 Estimated welfare costs of inflation, in percentage points of GDP, for $R_0=-1\%$, for selected values of the long rate (median and 90% bootstrapped confidence bands)		
<i>Country</i>	<i>Money demand specification:</i>	
	Selden-Latané	Semi-log
<i>Long rate equal to 3%</i>		
United States	2.179 [1.187; 3.796]	0.724 [0.540; 0.779]
United Kingdom	0.987 [0.322; 2.287]	0.304 [0.184; 0.362]
Canada	0.877 [0.463; 1.386]	0.301 [0.188; 0.347]
Switzerland	1.399 [0.844; 1.952]	1.028 [0.753; 1.107]
Sweden	0.555 [0.509; 0.621]	0.473 [0.436; 0.491]
Euro area	1.603 [1.106; 2.871]	0.924 [0.834; 0.992]
<i>Long rate equal to 4%</i>		
United States	2.454 [1.421; 4.143]	1.001 [0.769; 1.071]
United Kingdom	1.220 [0.440; 2.647]	0.437 [0.278; 0.501]
Canada	1.070 [0.599; 1.643]	0.446 [0.275; 0.527]
Switzerland	1.713 [1.090; 2.345]	1.423 [1.067; 1.532]
Sweden	0.742 [0.683; 0.828]	0.683 [0.637; 0.708]
Euro area	1.932 [1.359; 3.310]	1.285 [1.173; 1.381]
<i>Long rate equal to 5%</i>		
United States	2.685 [1.643; 4.407]	1.279 [1.008; 1.365]
United Kingdom	1.432 [0.557; 2.950]	0.586 [0.379; 0.668]
Canada	1.240 [0.727; 1.861]	0.604 [0.379; 0.708]
Switzerland	1.996 [1.326; 2.692]	1.817 [1.398; 1.962]
Sweden	0.927 [0.856; 1.034]	0.910 [0.856; 0.941]
Euro area	2.220 [1.591; 3.723]	1.650 [1.514; 1.779]
^a Based on 10,000 bootstrap replications.		

How should a reader or a policymaker use the information contained in the figures and the table? The average long-short spread has been equal to 1.76% for the U.S.

since 1996Q1; and to 1.04% for the U.K., to 1.13% for Canada, to 1.44% for Switzerland, to 1.31% for Sweden, and to 1.61% for the Euro area. By assuming that such average values of the spread will also hold going forward, and by making assumptions about the level of the natural rate of interest, and the numerical target for inflation, it is then possible to obtain a corresponding reference value for the long rate. For example, assuming a natural rate of interest of 1% (probably high by current levels), a 2% inflation target would map into an average value of the long rate just slightly higher than 4% for the U.K. and Canada, and equal to about 4.75% for the U.S.. Based on the Selden-Latané specification the corresponding welfare costs in Table 6 are equal to 1.22% of GDP, with a 90% confidence interval of [0.440; 2.647], for the U.K., and to 1.07%, with a confidence interval of [0.599; 1.643], for Canada.

6 Conclusions

In this paper we have estimated the welfare costs of inflation originating from lack of liquidity satiation—as in Bailey (1956), Friedman (1969), and Lucas (2000)—for the U.S., U.K., Canada, and two countries/areas (Euro area, Switzerland and Sweden) in which interest rates have recently plunged below zero. We have paid special attention to the fact that, as shown by recent experience, zero cannot be taken as the effective lower bound (ELB); the possibility that, as discussed by Mulligan and Sala-i-Martin (2000), the money demand curve may become flatter at low interest rates; and the functional form for money demand.

We have reported three main findings. First, allowing for an empirically plausible ELB (e.g., -1%) materially increases the welfare costs compared to the standard benchmark of zero. Second, there is nearly no evidence that at low interest rates money demand curves may become flatter: rather, evidence for the U.S. (the country studied by Mulligan and Sala-i-Martin (2000)) clearly points towards a steeper curve at low rates. Third, welfare costs are, in general, non-negligible: this is especially the case for the Euro area, Switzerland, and Sweden, which, for any level of interest rates, demand larger amounts of M1 as a fraction of GDP. For policy purposes the implication is that, *ceteris paribus*, inflation targets for these countries should be set at a comparatively lower level.

In terms of directions for future research, one interesting extension of the present analysis is to estimate the welfare costs of inflation for very high inflation countries, and especially for hyperinflation episodes. For example, for the Weimar Republic's hyperinflation, Bresciani-Turroni (1937, Table XXII, p. 168) presents monthly estimates of the income velocity of circulation, from which the welfare costs can be easily extracted based on the methods used in the present work.

7 References

- Alvarez, F., and F. Lippi (2009): “Financial Innovation and the Transactions Demand for Cash”, *Econometrica*, 77(2), 363-402.
- Attanasio, O. P., L. Guiso, and T. Jappelli (2002): “The Demand for Money, Financial Innovation, and the Welfare Cost of Inflation: An Analysis with Household Data”, *Journal of Political Economy*, 110(2(April)), 317-351.
- Bai, J. and P. Perron (1998): “Estimating and Testing Linear Models with Multiple Structural Changes”, *Econometrica*, 66(1), 47-78
- Bai, J. and P. Perron (2003): “Computation and Analysis of Multiple Structural Change Models”, *Journal of Applied Econometrics*, 18(1), 1-22
- Bailey, M. J. (1956): “The Welfare Cost of Inflationary Finance”, *Journal of Political Economy*, 64(2), 93-110.
- Baumol, W. J. (1952): “The Transactions Demand for Cash: An Inventory Theoretic Model”, *Quarterly Journal of Economics*, 66, 545-372.
- Belongia, M.T., and P.N. Ireland (2019): “The Demand for Divisia Money: Theory and Evidence”, *mimeo*, May 2019
- Benati, L. (2008): “Investigating Inflation Persistence Across Monetary Regimes”, *Quarterly Journal of Economics*, 123(3), 1005-1060.
- (2015): “The Long-Run Phillips Curve: A Structural VAR Investigation”, *Journal of Monetary Economics*, 76(November), 15-28.
- (2020): “Money Velocity and the Natural Rate of Interest”, *Journal of Monetary Economics*, forthcoming
- Benati, L., R. E. Lucas Jr., J.P. Nicolini, and W. Weber (2019): “International Evidence on Long-Run Money Demand”, Minneapolis FED Working Paper 587, June 2019
- 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.
- 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.
- Friedman, B.M., and K.N. Kuttner (1992): “Money, Income, Prices, and Interest Rates”, *American Economic Review*, 82(3), 472-492.
- Goldfeld, S.M. (1976): “The Case of the Missing Money”, *Brookings Papers on Economic Activity*, 3, 683-730.
- Hamburger, M.J. (1977): “Behavior of the Money Stock: Is there a puzzle?”, *Journal of Monetary Economics*, 3, 265-288.
- Hansen, B. E. (1999): “The Grid Bootstrap and the Autoregressive Model”, *Review of Economics and Statistics*, 81(4), 594-607.
- Ireland, P. (2009): “On the Welfare Cost of Inflation and the Recent Behavior of Money Demand”, *American Economic Review*, 99(3), 1040-1052.
- 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. (2000): “Inflation and Welfare”, *Econometrica*, 68(2), 247-274.
- Lucas Jr., R. E., and J.P. Nicolini (2015): “On the Stability of Money 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.
- 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. (1967a): “Inflation and Economic Growth”, *Journal of Political Economy*, 75(6), 796-810.
- Sidrauski, M. (1967b): “Rational Choice and Patterns of Growth in a Monetary Economy”, *American Economic Review*, 57(2), 534-544.
- Stock, J. H., and M. W. Watson (1993): “A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems”, *Econometrica*, 61(4), 783-820.
- Tobin, J. (1956): “The Interest Elasticity of Transactions Demand for Money”, *Review of Economics and Statistics*, 38, 241-247.

A The Data

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

A.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.

A.2 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’).

A.3 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.

A.4 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).

A.5 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).

A.6 Euro area

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

B Constructing an Own Rate of Return for Lucas and Nicolini's (2015) M1 Aggregate

We construct an own rate of return for Lucas and Nicolini's (2015; henceforth, LS) M1 aggregate as follows. LS' aggregate is equal to the standard M1 aggregate until 1982Q3, and it is equal to the standard aggregate plus MMDAs starting from 1982Q4. The standard M1 aggregate, in turn, is defined as the sum of currency, which pays no interest, and checking accounts, which pay instead some small interest. We compute the own rate of return for LS' aggregate as the weighted average of the rate on checking accounts and, since 1984, MMDAs, where the weights are computed as the fractions of checking accounts and MMDAs in the overall LS' M1 aggregate. The rate on checking accounts is available since 1987, whereas the rate on MMDAs is available for the period 1987-2000. Since both rates are available, for these two periods, at both the annual and the quarterly frequency, for the missing periods we proceed as follows.

Working at the quarterly frequency for the period 1987Q1-2018Q3 (for the rate on checking accounts) and for the period 1987Q1-2000Q4 (for the rate on MMDAs), we estimate *via* OLS simple linear regression models linking the dynamics of the first-difference of either of the two rates to the *present* and *past* dynamics of a number of series which are available for the entire post-WWII period. The regressors we use are the vacancy rate (from Regis Barnichon’s web page); the unemployment rate (the St. Louis FED’s FRED II acronym is UNRATE); the rate of capacity utilization in manufacturing (CUMFNS); the first two principal components extracted from the panel of the first-differences of the 3- and 6-month Treasury bill rates, and of the 1-, 3-, 5-, and 10-year Treasury constant maturity rates (TB3MS, DTB6, GS1, GS3, GS5, and GS10, respectively);²⁴ and the first three principal components extracted from the panel of the successive spreads²⁵ among the same six interest rate series.²⁶ The regressions (with two lags for the MMDAS’ rate, and four for the rate on checking accounts, for which the available period is longer) produce R-squared equal to 0.919 and 0.612, respectively. Then, based on the estimated model for the first-difference of either of the two rates of interest, we compute predicted values for the missing quarters, and based on them we reconstruct predicted values for their levels. Based on the predicted quarterly rates for checking accounts and MMDAs, we then compute the corresponding predicted annual rates by taking annual averages.

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 for the U.S. (from the *Center for Financial Stability*, henceforth *CFS*) and for the U.K. (from the Bank of England). A second problem is that, for the U.S., the Divisia M1 series constructed by the *CFS* does not feature MMDAs (which are instead included in Divisia M2). This means that although the resulting index of monetary services has been constructed by optimally weighting the underlying individual assets, it suffers from the crucial shortcoming that it is not including a key component of the transaction technology. As a result, although Divisia M1 is in principle superior to the standard simple-sum M1 aggregate, it ultimately suffers from the same shortcoming of not including MMDAs.

So the key question is: What is more important? Including MMDAs, or optimally weighting the underlying assets? Figure C.1 provides evidence on this, by showing the same evidence shown in Figure 1*b* in the paper, but this time with velocity being computed based on Divisia aggregates. The figure speaks for itself, and provides

²⁴The first two principal components explain almost 99 per cent of the variance of the panel.

²⁵That is, 6-month minus 3-month, 1-year minus 6-month, ..., 10-year minus 5-year.

²⁶The first three principal components explain about 97 per cent of the variance of the panel.

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 1*b*, clearly shows that, for the purpose of detecting a stable long-run demand for M1, 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 it 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