

**Money Velocity and the Natural
Rate of Interest**

Luca Benati

17-06

October 2017

DISCUSSION PAPERS

Money Velocity and the Natural Rate of Interest^{*}

Luca Benati
University of Bern[†]

Abstract

Since World War I, M1 velocity has been, to a close approximation, the permanent component of the short-term nominal rate. This logically implies that, under monetary regimes which cause inflation to be $I(0)$, permanent fluctuations in M1 velocity uniquely reflect, to a close approximation, permanent shifts in the natural rate of interest. Evidence from the Euro area and several inflation-targeting countries is compatible with this notion, with velocity fluctuations being systematically strongly correlated with a Stock and Watson (1996, 1998) estimate of trend real GDP growth. I exploit this insight to estimate the natural rate of interest for the United Kingdom and Canada under inflation targeting: In either country, the natural rate has been consistently declining since the early 1990s.

Keywords: Money demand; Lucas critique; structural VARs; unit roots; cointegration; long-run restrictions: natural rate of interest.

^{*}I wish to thank Peter Ireland, Harald Uhlig, and participants at seminars at Oxford, Bonn, Bern, Banca d'Italia, ECB, Norges Bank, and Bundesbank for useful comments, and Giuseppe Cavaliere for useful suggestions. Thanks also to Ruth Judson for help in adjusting U.S. currency for the fraction held by foreigners; and to Juan-Pablo Nicolini for kindly providing data on U.S. Money Market Deposits Accounts. Usual disclaimers apply.

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

1 Introduction

Over the last several decades, the role of monetary aggregates in the conduct of monetary policy has been progressively downgraded, to the point that, at major central banks, they are today regarded as largely irrelevant.

There are two main reasons for this:

(I) the breakdown, around the early 1980s, of previously identified stable relationships between monetary aggregates, GDP, interest rates, and prices in several countries (for the United States see, e.g., Friedman and Kuttner (1992)).

(II) As documented, e.g., by Estrella and Mishkin (1997), the disappearance, at low inflation rates, of any reliable informational content of monetary aggregates for key variables such as inflation and nominal GDP.¹

As for (I), Benati, Lucas, Nicolini, and Weber (2017) have re-established stability of the U.S. *long-run* demand for M1 since World War I based on three of the alternative adjustments to the standard M1 aggregate originally suggested by Goldfeld and Sichel (1990, pp. 314-315). Further, they have produced very similar evidence, since the mid-XIX century, for most of the other 31 countries they have analyzed, with evidence of breaks or time-variation in the cointegration relationships being instead consistently weak to non-existent.² As for the *short-run* money demand, although this paper is exclusively concerned with (II), my results naturally point towards an alternative interpretation of the existing evidence of instability, which I discuss below in Section 2.4.3.

In this paper I reconsider the informational content of monetary aggregates from a perspective which is different from that of Estrella and Mishkin (1997), and of several other conceptually related papers. My evidence suggests that one specific variable—*M1 velocity*—does indeed contain information which is crucial for monetary policy purposes. The information, however, does *not* pertain to variables such as inflation, or nominal GDP: Rather, it pertains to the *natural rate of interest*—which, conceptually in line with Laubach and Williams (2003), I define as the permanent component of the real rate—and it becomes starkly apparent under monetary regimes which cause inflation to be I(0), such as inflation-targeting regimes (see Benati (2008)).

My main result is that M1 velocity is, to a close approximation, the permanent component of the short-term nominal rate, so that the time-series relationship between the two series is the same as that between consumption and GDP. In particular, *all* of the results obtained by John Cochrane (1994) in his classic investigation of the

¹In a nutshell, Estrella and Mishkin's (1997, pp. 300-301) argument is that the informational content of monetary aggregates—in particular, for inflation—appears only when money growth exhibits large and highly persistent fluctuations (such as, in the second half of the XX century, in Argentina, Chile, and Brazil). When both money growth and inflation are low and stable, on the other hand, money's information gets largely lost amid the 'noise' produced by non-monetary shocks.

²In the United States, in particular, there is no evidence whatsoever of breaks or time-variation in the cointegration relationship.

relationship between consumption and GNP, and dividends and stock prices, map, one-for-one, to the bivariate system for M1 velocity and the short rate.

Since the unit root in the short-term nominal rate originates from either permanent inflation shocks, or permanent shocks to the real rate,³ these results logically imply that, under monetary regimes which cause inflation to be $I(0)$, permanent fluctuations in M1 velocity *uniquely* reflect, to a close approximation, permanent shifts in the natural rate of interest. To put it differently, under these regimes M1 velocity is essentially a function of the natural rate of interest, e.g., $V_t = \alpha + \beta R_t^N + \epsilon_t$ (where the notation is obvious, and ϵ_t is a ‘small’ noise component), so that the natural rate is, to a first approximation, and up to a scale factor, *observed*. Evidence from the Euro area and several inflation-targeting countries is compatible with this notion, with velocity fluctuations being systematically strongly correlated with a Stock and Watson (1996, 1998; henceforth, SW) time-varying parameters median-unbiased (TVP-MUB) estimate of trend real GDP growth.

This means that the information contained in M1 velocity can be exploited in order to estimate the natural rate. Further, a consistent decrease in M1 velocity under a monetary regime causing inflation to be $I(0)$ —such as the protracted fall in velocity which has been going on in several inflation-targeting countries since the early 1990s—provides *direct evidence* of a fall in the natural rate of interest.

Once again, the parallel between these results, and those reported by Cochrane (1994) for GNP and consumption, is immediate and obvious. By disentangling permanent and transitory idiosyncratic shocks to their own income, consumers are providing policymakers crucial, and otherwise unavailable information about the permanent component of GDP. By the same token, agents are here disentangling permanent and transitory interest rate shocks, thus providing policymakers information about the permanent component of interest rates. I exploit this insight to estimate the natural rate of interest for the United Kingdom and Canada under inflation targeting: In either country, the natural rate has been consistently declining since the early 1990s.

The paper is organized as follows. The next section provides a simple illustration of this paper’s main findings and argument for the United Kingdom, for which evidence is so stark that it can be seen essentially *via* the naked eye. Section 3 describes the dataset, whereas Section 4 explores the unit root and cointegration properties of the data. Section 5 explores how the bivariate cointegrated system featuring M1 velocity and the short rate adjust towards equilibrium, whereas Section 6 disentangles permanent and transitory shocks to the system along the lines of Cochrane (1994). Section 7 provides evidence for monetary regimes which cause inflation to be $I(0)$. Section 8 concludes.

³This assumption is discussed at length in Section 2.4.1. (see footnote 11). In a nutshell, as I argue there, the assumption is meaningful *especially* from a monetary policy perspective.




Figure 1a Evidence for the United Kingdom, 1955Q1-2016Q4

2 A Simple Illustration

Although my main result is qualitatively the same for the vast majority of the countries I consider,⁴ for some of them it is especially stark, in the sense that it can be seen essentially with the naked eye. This is the case, in particular, for the United Kingdom over the post-World War II period. In this section I therefore illustrate the main results of this paper by drawing on the post-WWII U.K. experience. In Sections 4 to 7 I will present the corresponding evidence for all other countries.

2.1 The time-series relationship between M1 velocity and the short rate

The first panel of Figure 1a shows M1 velocity and the short rate for the post-WWII U.K.. Visual impression clearly suggests the following three facts, which, as I will discuss in Sections 4.1-4.2, are strongly confirmed by proper econometric techniques:

- (i) M1 velocity and the short rate are both I(1);
- (ii) the two series are cointegrated; and, crucially,
- (iii) up to a linear transformation, M1 velocity is, essentially, the stochastic trend of the short rate.

The implication is that when the cointegrated system is out of equilibrium, adjustment takes place *via* movements in the short rate towards its stochastic trend—i.e., (rescaled) velocity—rather than *via* movements in velocity. To put it differently, velocity is *always* approximately in equilibrium: It is rather the short rate which, featuring a transitory component which closely co-moves with the transitory component of GDP,⁵ is typically out of equilibrium.

The two panels in the second column of Figure 1a provide clear evidence on this, by showing the bootstrapped distributions of the two series' loading parameters on the cointegration residual in the estimated VECM: Whereas the estimate of the loading parameter for M1 velocity, at 0.002, is negligible, and is not significantly different from zero, the corresponding estimate for the short rate, at -0.187, is strongly statistically significant.⁶ In particular, the bootstrapped *p*-values for testing the null hypothesis that the two coefficients are equal to zero are equal to 0.166 and 0.007, respectively.⁷ These results are qualitatively the same as those obtained by Cochrane (1994): For the VECM featuring consumption and GNP, his estimates of the loading parameters on the two series reported in Table I were (*t*-statistics in parentheses) -0.02 (-1.23) and 0.08 (3.45), respectively, whereas for the VECM featuring dividends and stock

⁴The main exception is Japan.

⁵Reflecting the central bank's 'leaning against the wind' of the future inflationary/deflationary pressures signalled by a positive/negative output gap.

⁶The econometric methodology, which is off-the-shelf, is the same used by Benati, Lucas, Nicolini, and Weber (2017). Details are provided in Section 5.

⁷The *p*-values are reported in Tables A.3 in the Appendix.

prices the corresponding figures from Table II were 0.038 (0.47) and 0.225 (2.11), respectively.

2.2 Interpretation

A simple way of interpreting these results is the following. Assume that the nominal short-term interest rate, R_t , is equal to the sum of two orthogonal components, a random walk, R_t^P , and a stationary AR(1) process, R_t^T :

$$R_t = R_t^P + R_t^T \quad (1)$$

$$R_t^P = R_{t-1}^P + u_t \quad (2)$$

$$R_t^T = \rho R_{t-1}^T + v_t \quad (3)$$

with $0 \leq \rho < 1$, and u_t and v_t white noise. (Shortly, I will provide evidence that in the United Kingdom the short rate is indeed not a pure unit root process, and it rather features a sizeable transitory component. In Section 6 I will provide analogous evidence for the other countries.) Then, consider the following two linear specifications for money velocity, corresponding to what Benati *et al.* (2017) label as the ‘Selden-Latané’ money-demand specification, from Richard Selden (1956) and Henry Allen Latané (1960).⁸:

$$V_t = \alpha + \beta R_t + \epsilon_t \quad (4)$$

$$V_t = \alpha + \beta R_t^P + \epsilon_t \quad (5)$$

(As I show in Appendix B, the Selden-Latané specification is a special case of the ‘money in the utility function’ framework pioneered by Miguel Sidrauski (1967, 1968).) The key difference between (4) and (5) is that whereas in the former specification—in line with standard money-demand literature—velocity (and therefore its inverse, money balances as a fraction of GDP) depends on the nominal interest rate, in the latter specification it depends on its *permanent component*. It can be trivially shown that whereas (4) implies the following VECM representation for ΔV_t and ΔR_t :

$$\begin{aligned} \begin{bmatrix} \Delta V_t \\ \Delta R_t \end{bmatrix} &= \text{Constants} + \begin{bmatrix} 0 & \beta\rho \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \Delta V_{t-1} \\ \Delta R_{t-1} \end{bmatrix} - \\ &\quad - \underbrace{\begin{bmatrix} 1 \\ 0 \end{bmatrix}}_{\text{Loadings}} \underbrace{\begin{bmatrix} 1 & -\beta \end{bmatrix}}_{\text{Cointegration vector}} \begin{bmatrix} V_{t-1} \\ R_{t-1} \end{bmatrix} + \text{Shocks} \end{aligned} \quad (6)$$

⁸As discussed by Benati *et al.* (2017), the key reason for considering this long-forgotten specification is that for several low-inflation countries—first and foremost, the United States—the data seem to quite clearly prefer it over the traditional log-log and semi-log ones. This evidence will be discussed in Section 4 below.

(5) implies the following one:

$$\begin{bmatrix} \Delta V_t \\ \Delta R_t \end{bmatrix} = \text{Constants} + \underbrace{\begin{bmatrix} 0 \\ \frac{1-\rho}{\beta} \end{bmatrix}}_{\text{Loadings}} \underbrace{\begin{bmatrix} 1 & -\beta \end{bmatrix}}_{\text{Cointegration vector}} \begin{bmatrix} V_{t-1} \\ R_{t-1} \end{bmatrix} + \text{Shocks} \quad (7)$$

In plain English, the ‘traditional’ specification⁹ (4) implies that the VECM’s adjustment towards its long-run equilibrium takes place *via* movements in velocity, with *no* reaction of the short rate to disequilibria. Specification (5), on the other hand, implies that—in line with the evidence in the second column of Figure 1a—the adjustment takes place *via* movements in the short rate, with *no* reaction of velocity. This feature is a direct consequence of the fact that, according to (5), velocity is (up to a linear transformation) the stochastic trend of the short rate.

Expressions (6) and (7) provide a straightforward, and natural interpretation for the evidence reported in the two panels in the second column of Figure 1a: The dynamics of M1 velocity in the post-WWII U.K. is well described by (5), rather than by the traditional specification (4), thus implying that velocity has been systematically reacting to the *permanent* component of the short rate, rather than to the short rate itself. As we will see, this has been a robust feature of macroeconomic fluctuations in nearly all of the countries in my dataset—first and foremost, in the United States and the United Kingdom since World War I.

2.3 Impulse-response functions and variance decompositions

Expression (5) implies that

- (i) Assuming that ϵ_t is small, shocks to the permanent component of the short rate explain the bulk of the (forecast error) variance of velocity; and
- (ii) velocity only reacts to permanent shocks to the short rate, whereas it does not react to transitory shocks.

The first two panels in the first row of Figure 1b provide evidence on (i), whereas the corresponding panels in the second row report evidence on (ii). The fractions of forecast error variance (FEV) and the impulse-response functions (IRFs) have been computed based on a cointegrated structural VAR (SVAR) for the two series shown in Figure 1a. Conceptually in line with one of the identification schemes used by Cochrane (1994), the permanent shock driving the common trend in the system has been identified as the only shock impacting upon the *short rate*—rather than velocity—in the infinite long run. The bootstrapped confidence bands (the figure reports the 16th, 84th, 5th, and 95th percentiles of the bootstrapped distributions

⁹I label (4) as a traditional specification—in spite of the fact that Selden and Latané’s work had been essentially forgotten for six decades—because, according to (4), velocity is a function of the nominal rate, rather than of its permanent component.




Figure 1b Evidence for the United Kingdom, 1955Q1-2016Q4

of the relevant objects) have been computed based on the procedure proposed by Cavaliere *et al.* (2012; henceforth, CRT), which is briefly described in Section 4.2.¹⁰

Two features stand out:

First, in line with (5) and (7), the permanent shock to the short rate explains nearly all of the FEV of velocity at all horizons, whereas it explains between about 25 and 30 per cent of the FEV of the short rate itself at horizons up to five years ahead, and slightly more than 60 per cent ten years ahead. It is important to stress that this result has been obtained in spite of the fact that the shock has been identified as the one driving the permanent component of the *short rate*, rather than of velocity. The parallel with consumption and GDP is obvious: In his Table I, Cochrane (1994) reports that the permanent consumption shock explains 97 per cent of the variance of consumption growth, and only 30 per cent of the variance of GNP growth.

Second, V_t does not react to transitory shocks at any horizon, whereas the response of R_t is strongly statistically significant.

Both features stand in sharp contrast to the corresponding predictions of specification (4), which implies that velocity is also driven by, and reacts to, the transitory component of the short rate.

These results have several implications, which I discuss in turn.

2.4 Implications

2.4.1 The informational content of M1 velocity for the natural rate of interest

The fact that M1 velocity is, essentially, the permanent component of the short rate has the following implication. Basic economic logic suggests that R_t^P should be driven by

- permanent inflation shocks (*via* the Fisher effect) and
- permanent shocks to the real rate, i.e., shocks to the natural rate of interest,

that is, $R_t^P = \pi_t^P + r_t^N$, where π_t^P is the permanent component of inflation, and r_t^N is the natural rate of interest.¹¹ This implies that, under monetary regimes which cause

¹⁰Specifically, within the present context the model which is being bootstrapped is the VECM estimated conditional on one cointegration vector.

¹¹This assumption requires some discussion. The logical implication is that *all* other disturbances impacting upon nominal rates are transitory, including (e.g.) shocks to the risk premium. Now suppose—just for the sake of the argument—that shocks to the risk premium were partly transitory, and partly permanent (the argument applies to any other shock impacting upon nominal rates). Under these circumstances, my assumption would interpret permanent risk premium shocks as natural rate shocks. The key issue here, however, is that—*especially* for monetary policy purposes—this is *exactly* what we would want. The fact that, for a given equilibrium inflation rate, the equilibrium nominal Federal Funds rate increases by x per cent because of (say) a permanent risk premium

inflation to be $I(0)$ —so that $\pi_t^P=0$ —permanent shifts in M1 velocity should *uniquely* reflect permanent fluctuations in the natural rate of interest, so that $V_t = \alpha + \beta r_t^N + \epsilon_t$.

The two panels in the third column of Figure 1b provide simple evidence on this for the U.K. inflation-targeting regime.¹² The upper panel shows GDP deflator inflation: Visual evidence suggests that—in line with the evidence reported in Benati (2008)—under inflation-targeting U.K. inflation has been very strongly mean-reverting. In fact, as I discuss in Section 7, the null of a unit root is very strongly rejected, with p -values from Elliot *et al.*'s (1996) tests equal to or close to zero. By the same token, Hansen's (1999) bias-corrected estimate of the sum of the autoregressive coefficients in an $AR(p)$ representation for inflation is equal to -0.32, with the 90 per cent-coverage confidence interval equal to $[-0.75; 0.10]$. In plain English, under inflation-targeting U.K. inflation has been essentially white noise, thus implying that shifts in M1 velocity should have uniquely reflected fluctuations in the natural rate of interest. In turn, this implies that the protracted fall in M1 velocity experienced by the United Kingdom under inflation-targeting should have been driven by a corresponding decline in the natural rate of interest.

The lower panel presents simple evidence compatible with this notion. As discussed by Laubach and Williams (2003, p. 1063), within a vast class of models (i.e., Solow's growth model, and standard optimal growth models) the natural rate of interest is a linear function of the economy's trend growth rate. This implies that we should see a strong correlation between velocity and the trend growth rate of GDP in the United Kingdom under inflation-targeting. In Section 7 I estimate a time-varying trend for real GDP growth for the United Kingdom and several other countries based on SW's (1996, 1998) TVP-MUB methodology. Here I report a much simpler estimate—a linear time trend for GDP growth estimated *via* OLS—which is however in line with the results produced by SW's methodology (this can be seen by comparing the linear trend in Figure 1 with the TVP-MUB trend in Figure 8). The correlation between velocity and trend GDP growth, although not perfect, is very strong, with the former falling from 1.28 in 1992Q4 to 0.60 in 2015Q4, and the latter decreasing from about 2.3 to about 1.8 per cent over the same period. Although by no means does this evidence represent a hard proof that my argument is correct, it is, at the very least, compatible with such position. This implies that, in principle, it should be possible to estimate the natural rate of interest by exploiting the informational content of M1 velocity. In Section 7.2 I will provide a simple illustration of this for the United Kingdom and Canada under inflation-targeting.

shock—as opposed to a supposedly 'authentic' shock to the natural rate (due, e.g., to an acceleration of productivity growth)—is for monetary policy purposes, completely irrelevant. What matters is that the permanent component of the Funds rate has increased by x per cent. So, although in principle it is possible to further decompose permanent shocks to the real rate, such distinctions are, for monetary policy purposes, irrelevant.

¹²In the United Kingdom, inflation targeting was introduced in October 1992.

I now turn to additional implications of the finding that M1 velocity is the permanent component of the short rate.

2.4.2 What does a disequilibrium in the cointegrated system signal?

The money demand literature has routinely interpreted deviations from the long-run equilibrium between the short rate and velocity (or money balances as a fraction of GDP) as signalling possible future inflationary pressures. The *implicit assumption* behind such interpretation is that the presence of a disequilibrium in the cointegrated system implies that money balances are out of equilibrium. As they adjust towards equilibrium, pent-up inflationary pressures are released, and inflation increases.

Although this interpretation is intuitively appealing, my results show that it is incorrect (at least, for M1). The reason is that, as previously discussed, M1 velocity (and therefore real M1 balances) are always approximately in equilibrium: It is rather the short rate which is typically out of equilibrium. This implies that a disequilibrium in the relationship between velocity and the short rate (i.e., the cointegration residual being different from zero) does *not* signal future inflationary pressures: Rather, it signals future movements of the short rate towards equilibrium.

2.4.3 Meaninglessness of the notion of ‘instability of money demand’

A further implication is the following. The fact that M1 velocity is (to a close approximation, and up to a linear transformation) the permanent component of the short rate, logically implies that speaking of ‘money demand instability’, quite simply, *makes no sense*. Once again, the crucial point here is that velocity, and therefore real money balances (expressed as a fraction of GDP), are always approximately in equilibrium: It is rather the short rate which is typically out of equilibrium. The easiest way to grasp this point is by recalling the parallelism with Cochrane’s (1994) results for consumption and GNP. Surely, nobody would argue that the unit root component of GNP (i.e., consumption) exhibits an ‘unstable relationship’ with GNP itself, because such a statement would be manifestly non-sensical. The same logic applies here: The fact that M1 velocity is the unit root component of the short rate implies that speaking of instability of the relationship between velocity (or money balances) and the short rate equally makes no sense.¹³

In turn, this implies that the vast literature of the instability of money demand which originated from the work of Stephen Goldfeld (1976) is equally logically incorrect. To be sure, the relationship between the short rate and velocity (or money balances) does indeed exhibit instability. The correct interpretation of such instability, however, is not that money demand is unstable but rather that the short rate has exhibited unstable fluctuations around its stochastic trend, or, to put it differently, that the stochastic properties of the transitory component of the short rate

¹³Another way of making the same point is that speaking of instability of money demand is as meaningful as stating that potential GDP exhibits an unstable relationship with GDP itself.

have been time-varying (likely, because of changes in the systematic component of monetary policy). This is in line with Marvin Goodfriend’s (1985, pp. 223-224) insightful discussion of the (alleged) instability of money demand estimated equations first pointed out by Goldfeld (1976) In his words,

‘[...] the upward forecast bias could be due to a shift in the income or interest rate generating processes instead of a shift in true money demand. [T]he interest rate generating process is highly influenced by monetary policy. For example, monetary policy can affect the level of the interest rate, interest rate autocorrelation, and the variance of interest rate innovations [...]. Since these parameters, in turn, affect money demand regression coefficients, these regression coefficients can be expected to depend on the monetary policy being followed during the sample period over which the regression is estimated. It follows that post-sample predictive performance of a money demand regression could be adversely affected if monetary policy alters the post-sample interest rate generating process relative to the sample period.’

Another way of putting this is to say that the alleged instability of money demand is nothing but a simple consequence of the Lucas critique.

2.5 The short-long spread and the cointegration residual between velocity and the short rate

The last panel in Figure 1b provides evidence on another remarkably robust stylized fact which has held for all countries and periods in my dataset.¹⁴ The panel shows the cointegration residual between M1 velocity and the short rate, together with the difference between the short rate and a long rate. A striking negative correlation between the two series is readily apparent. Interestingly, the period following the collapse of Lehman Brothers—which featured the most violent phase of the recent financial crisis—does not exhibit any obvious difference with the rest of the sample. This suggest that such strong correlation originates from some deep, structural feature of the economy, so that it is not thrown out of kilter even by the largest macroeconomic shock since the Great Depression.

The simple model outlined previously points towards the following natural interpretation for this stylized fact. Assume that the long-term nominal interest rate, r_t , is equal to the permanent component of the short rate:

$$r_t = R_t^P \tag{8}$$

¹⁴To be precise: For all countries for which I could find data on a long-term nominal interest rate. Evidence is reported in Figure 3, and it is discussed in Section 3.

This specification is designed to capture, in an extreme fashion, the robust stylized facts that (i) short- and long-term rates are cointegrated, and (ii) the long rate consistently behaves as a low-frequency trend for the short rate,¹⁵ with (e.g.) its first-difference systematically exhibiting a lower volatility than the first-difference of the short rate.¹⁶ Equations (1) and (8) imply that the short-long spread is equal to the transitory component of the short rate, $R_t - r_t = R_t^T$. In turn, (5) implies that the cointegration residual between V_t and R_t is equal to $[V_t - \beta R_t] = \alpha - \beta R_t^T + \epsilon_t$, so that

$$[V_t - \beta R_t] = \alpha - \beta[R_t - r_t] + \epsilon_t \quad (9)$$

In plain English, the cointegration residual between velocity and the short rate is perfectly negatively correlated with the short-long spread, as documented in Figure 1b. On the other hand, under the ‘traditional’ specification (4) the cointegration residual would be equal to $[V_t - \beta R_t] = \alpha + \epsilon_t$.

I now turn to discussing the dataset.

3 The Data

Appendix A describes the data and their sources in detail. All of the data are from official sources, that is, either central banks or national statistical agencies. Almost all of the annual data are from the dataset assembled by Benati *et al.* (2017),¹⁷ which I have updated to the most recent available observation whenever possible (typically, I have added either one or two years).

All of the series are standard, with the single exception that, for the United States, I consider three of the alternative adjustments to the Federal Reserve’s standard M1 aggregate which 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, which had vanished around the mid-1980s. Specifically, I augment the standard M1 aggregate with either Money Market Deposits Accounts (MMDAs), as in Lucas and Nicolini (2015);¹⁸ Money Market Mutual Funds (MMDFs); or both MMDAs and MMDFs. Benati *et al.* (2017) show that whereas—in line with, e.g., Friedman and Kuttner (1992)—based on the standard aggregate there is no evidence of a stable long-run

¹⁵This fact was especially apparent during the metallic standards era (i.e., before World War I), when long-term rates typically exhibited a very small extent of low-frequency variation, and short-term rates systematically fluctuated around long rates, following the ups and downs of the business cycle.

¹⁶E.g., for the post-WWII U.K. the standard deviations of the first-differences of the short and long rates used to compute the spread shown in Figure 1b have been equal to 0.906 and 0.567 per cent.

¹⁷In several cases (South Africa, Taiwan, South Korea and Hong Kong, and Canada since 1967) I was able to find quarterly data for the same sample periods analyzed by BLNW (2017).

¹⁸As discussed by Lucas and Nicolini (2015), the rationale for including MMDAs in M1 is that they perform an economic function similar to the more traditional ‘checkable deposit’ component of the Federal Reserve’s official M1 series.




Figure 2a The annual raw series




Figure 2b The quarterly raw series

demand for M1, evidence of cointegration between velocity and the short rate is very strong based on Lucas and Nicolini's (2015) aggregate. Benati *et al.* (2017), on the other hand, do not analyze the other two aggregates I consider herein. Finally, for reasons of robustness, for either of the three 'expanded' U.S. M1 aggregates I also consider an alternative version, in which currency has been adjusted along the lines of Judson (2017), in order to take into account of the fact that, since the early 1990s, there has been a sizeable expansion in the fraction of U.S. currency held by foreigners. So, in the end, for the United States I consider *six* alternative M1 aggregates. As I discuss below, adjusting, or not adjusting for the fraction of U.S. currency held by foreigners does not make a material difference to the results, which originates from the fact that the currency component of M1 is ultimately quite small compared to the deposits component. For reasons of space, in what follows I only report results for the aggregate including MMDAs, and for the one including both MMDAs and MMFAs. Results for the aggregate just including MMFAs are qualitatively the same, and they are available upon request.

Appendix A discusses in detail a few countries in Benati *et al.*'s dataset which I have chosen not to analyze herein because, e.g., the data exhibit puzzling features (this is the case in particular for Italy and Norway). Further, in the present work I have chosen not to analyze the high- and very high-inflation countries in Benati *et al.*'s dataset, and to exclusively focus on low-to-medium inflation countries.¹⁹ This choice is motivated by the following considerations. Although high-inflation countries' extreme experiences are very useful for the purpose of identifying *cointegration* between velocity and the short rate, their macroeconomic dynamics is typically affected, to a non-negligible extent, by highly idiosyncratic shocks and events, which can be expected to distort the subtler features (IRFs and variance decompositions) investigated herein. Chile provides a stark illustration of this problem. Although Benati *et al.* (2017) detect cointegration between velocity and the short rate for Chile, as they discuss in Section G.2.1 of the online appendix, the two series exhibit dramatic fluctuations, and a strong *negative* correlation, in the first half of the 1970s, around the time of the economic and political turmoil which culminated with Augusto Pinochet's military *coup d'état* of September 1973.²⁰ The fact that, out of a sample of 56 years (1940-1995), about a decade of data has thus been significantly distorted suggests that the informational content of these data for features subtler than cointegration is likely limited. By the same token, in both Argentina and Brazil, the sharp disinflations of the 1990s have been followed by slow and belated falls in velocity, likely reflecting, at least in part, the public's gradual learning about the seriousness of the government's newfound commitment to low inflation. Finally, although for South Africa I have quarterly data starting in 1965Q1, I have decided to restrict my analysis to the sample starting in 1985Q1, because the relationship between M1

¹⁹A very partial exception to this is Mexico, for which, for one year and a half at the very beginning of the sample, inflation exceeded 100 per cent.

²⁰See Figure 2 in BLNW's (2017) online appendix.




Figure 3 The cointegration residual between M1 velocity and the short rate, and the short-long spread

velocity and the short rate during the previous two decades appears as manifestly different just based on a simple visual inspection of the data.²¹

Figures 2a and 2b show the raw data for M1 velocity and the short-term nominal interest rate. In line with the evidence for the United Kingdom, in several cases visual evidence quite clearly suggests that velocity and the short rate are cointegrated, and that the former is, essentially, the permanent component of the latter. This is the case, e.g., for Canada, Australia, Taiwan, and South Korea based on quarterly data. As we will see in the next three sections, econometric evidence does indeed confirm such visual impression.

Figure 3 shows the cointegration residual between velocity and the short rate (i.e., between the series shown in Figures 2a-2b), together with the difference between the short rate and a long-term nominal rate. Due to data limitations for the long rate, evidence for Switzerland starts in 1960, rather than in 1914 as in Figure 2a; and, more generally, the figure only shows evidence for a few countries. In line with the evidence for the United Kingdom in the last panel of Figure 1b, in nearly all cases the cointegration residual exhibits a strong, negative correlation with the short-long spread. The single exception is South Korea since the beginning of the new millennium (on the other hand, the correlation had been strong over the previous period). It is to be noticed, however, that the breakdown of the correlation for Korea over the last 15-20 years has been due to the anomalous behaviour of the spread, which has significantly increased compared to previous years, rather than to any obvious change in the behaviour of the cointegration residual. This means that for the purpose of this paper, whose focus is the relationship between velocity and the short rate, such a breakdown is immaterial.

Interestingly, in the United States the correlation had been thrown temporarily out of kilter by the introduction of MMDAs in 1982, but it reasserted itself in the second half of the 1980s, and it has consistently held since then (see the fourth panel in the second row). Further, in *all* cases²² the period following the collapse of Lehman Brothers—which featured the most violent phase of the recent financial crisis—does not exhibit any obvious difference with the rest of the sample. This provides additional support to the conjecture (see Section 2.4) that such a strong correlation reflects a deep structural feature of the economy. In particular, the fact that such a relationship has been holding steady at least since World War I, in spite of dramatic shifts in the monetary regime (the partial reintroduction, and then the disintegration of the Gold Standard in the interwar period; the Bretton Woods regime and its collapse; the introduction, in several instances, of inflation-targeting regimes in the 1990s; and the adoption of quantitative easing (QE) policies during the financial crisis) suggests that such a relationship might well be structural in the sense of the

²¹For the specific purpose of this paper, the results based on the full sample 1965Q1-2017Q1 are qualitatively the same as those presented herein. The *only* difference is that the IRF of the short rate to a permanent shock exhibits an implausible pattern. These results are available upon request.

²²With the just-mentioned exception of Korea.

Table 1 Bootstrapped p -values ^a for Johansen's maximum eigenvalue ^b tests for (log) M1 velocity and (the log of) a short-term rate				
Country	Period	Money demand specification:		
		Selden-Latané	Semi-log	Log-log
	I: Long-run annual data			
United States				
<i>standard</i> M ₁ + MMDAs	1915-2016	0.041	0.066	0.184
<i>standard</i> M ₁ + MMDAs + MMMFs	1915-2016	0.003	0.009	0.466
<i>Adjusting for currency held by foreigners:</i>				
<i>standard</i> M ₁ + MMDAs	1926-2016	0.067	0.130	0.150
<i>standard</i> M ₁ + MMDAs + MMMFs	1926-2016	0.151	0.027	0.580
United Kingdom	1922-2016	0.022	0.025	0.345
Switzerland	1914-2015	0.007	0.015	0.383
New Zealand	1934-2016	0.120	0.136	0.057
Canada	1935-2006	0.023	0.247	0.485
Japan	1955-2015	0.626	0.345	0.220
Australia	1941-1989	0.642	0.973	0.709
Belgium	1946-1990	0.361	0.016	0.010
Netherlands	1950-1992	0.349	0.286	0.401
Finland	1914-1985	0.622	0.659	0.839
	II: Post-WWII quarterly data			
United States				
<i>standard</i> M ₁ + MMDAs	1959Q1-2016Q4	0.013	0.023	0.332
<i>standard</i> M ₁ + MMDAs + MMMFs	1959Q1-2016Q4	0.002	0.001	0.027
<i>Adjusting for currency held by foreigners:</i>				
<i>standard</i> M ₁ + MMDAs	1959Q1-2016Q4	0.097	0.058	0.331
<i>standard</i> M ₁ + MMDAs + MMMFs	1959Q1-2016Q4	0.002	0.002	0.030
United Kingdom	1955Q1-2016Q4	0.058	0.104	0.588
Canada	1967Q1-2016Q4	0.020	0.150	0.000
Australia	1975Q1-2016Q3	0.078	0.089	0.512
Taiwan	1961Q3-2016Q4	0.001	0.240	0.293
South Korea	1964Q1-2017Q1	0.000	0.388	0.229
South Africa	1985Q1-2017Q1	0.382	0.378	0.419
Hong Kong	1985Q1-2017Q1	0.209	0.084	0.036
Mexico	1985Q4-2017Q1	0.051	0.041	0.444
^a Based on 10,000 bootstrap replications. ^b Null of 0 <i>versus</i> 1 cointegration vectors.				
^c Results from Benati <i>et al.</i> (2017).				

Lucas (1976) critique.

4 Integration and Cointegration Properties of the Data

4.1 Unit root tests

Tables A.1*a* and A.1*b* in the online appendix²³ report bootstrapped p -values²⁴ for Elliot, Rothenberg, and Stock (1996) unit root tests for (the logarithms of) M1 velocity and the short rate. All tests are with an intercept, but no time trend. In line with Benati *et al.* (2017), for the short rate, R_t , I also report results for $\ln(R_t+1)$, in which the simple series has been corrected along the lines of Alvarez and Lippi (2009), by adding to it a 1 per cent cost of either losing cash, or having it stolen.²⁵ In nearly all cases, evidence of a unit root in either series is very strong, with the p -values being almost uniformly greater than the 10 per cent significance level I take as the benchmark throughout the entire paper, and often significantly so. For Switzerland a unit root is rejected for $\ln(R_t)$, but not for $\ln(R_t+1)$. Because of the reason mentioned in the previous footnote, in what follows the analysis for the ‘log-log’ specification will be performed based on $\ln(R_t+1)$, rather than $\ln(R_t)$, and these results are therefore ultimately irrelevant.²⁶ For Korea the alternative lag orders produce contrasting evidence for velocity. In this cases I regard the null of a unit root as not having been convincingly rejected, and in what follows I therefore proceed under the assumption that the series is $I(1)$. Finally, for Taiwan a unit root is rejected for velocity based on either lag order. In the light of the evidence in Figure 2*b*—in which velocity has been consistently declining since 1961—I regard this result as a statistical fluke.²⁷

Tables A.2*a* and A.2*b* in the Appendix report bootstrapped p -values for Elliot *et al.* (1996) unit root tests for either the first differences, or the log-differences, of velocity and the short rate. In all cases the null of a unit root is strongly rejected, thus suggesting that the series’ order of integration is not greater than one.

²³The online appendix is available at: <https://sites.google.com/site/lucabenatiswebpage>.

²⁴ p -values have been computed by bootstrapping 10,000 times estimated ARIMA($p,1,0$) processes. In all cases, the bootstrapped processes are of length equal to the series under investigation. As for the lag order, p , since, as it is well known, results from unit root tests may be sensitive to the specific lag order which is being used, for reasons of robustness I consider two alternative lag orders based on annual data (either 1 or 2), and four based on quarterly data (either 1, 2, 3, or 4).

²⁵A key rationale for doing this is that this correction delivers a finite satiation level of real money balances at $R_t = 0$.

²⁶On the other hand, there is no point in implementing Alvarez and Lippi’s (2009) correction for either the Selden-Latané or the semi-log specification, since in both cases the short rate enters in levels.

²⁷When performing a large number of statistical tests, such as it the case here, a certain number of flukes should be expected. To be sure, the series I am analyzing here are not independent stochastic processes generated (e.g.) in MATLAB, but the same logic should approximately apply.

Table 2 Monte Carlo-based empirical rejection frequencies for the bootstrapped maximum eigenvalue^b tests for (log) M1 velocity and (the log of) a short-term rate, under the null of cointegration, based on annual data

<i>Country</i>	<i>Period</i>	<i>Selden- Latané</i>	<i>Semi- log</i>	<i>Log-log</i>
<i>I: Long-run annual data</i>				
United States				
<i>standard</i> M1 + MMDAs	1915-2016	–	–	0.628
<i>standard</i> M1 + MMDAs + MMMFs	1915-2016	–	–	0.430
<i>Adjusting for currency held by foreigners:</i>				
<i>standard</i> M1 + MMDAs	1926-2016	–	0.734	0.723
<i>standard</i> M1 + MMDAs + MMMFs	1926-2016	0.725	–	0.270
United Kingdom	1922-2016	–	–	0.383
Switzerland	1914-2015	–	–	0.700
New Zealand	1934-2016	0.685	0.690	–
Canada	1935-2006	–	0.846	0.621
Japan	1955-2015	0.363	0.596	0.605
Australia	1941-1989	0.168	0.079	0.200
Belgium	1946-1990	0.699	0.635	0.744
Netherlands	1950-1992	0.463	0.427	0.324
Finland	1914-1985	0.231	0.218	0.209
<i>II: Post-WWII quarterly data</i>				
United States				
<i>standard</i> M1 + MMDAs	1959Q1-2016Q4	–	–	0.544
<i>standard</i> M1 + MMDAs + MMMFs	1959Q1-2016Q4	–	–	–
<i>Adjusting for currency held by foreigners:</i>				
<i>standard</i> M1 + MMDAs	1959Q1-2016Q4	–	–	0.553
<i>standard</i> M1 + MMDAs + MMMFs	1959Q1-2016Q4	–	–	–
United Kingdom	1955Q1-2016Q4	–	0.576	0.232
Canada	1967Q1-2016Q4	–	0.773	–
Australia	1975Q1-2016Q3	–	–	0.382
Taiwan	1961Q3-2016Q4	–	0.993	0.755
South Korea	1964Q1-2017Q1	–	0.701	0.624
South Africa	1985Q1-2017Q1	0.494	0.494	0.299
Mexico	1985Q4-2017Q1	–	–	0.428
Hong Kong	1985Q1-2017Q1	0.820	–	–
^a Based on 10,000 bootstrap replications. ^b Null of 0 <i>versus</i> 1 cointegration vectors.				

4.2 Cointegration tests

Table 1 reports results from Johansen’s maximum eigenvalue tests²⁸ between velocity and the short rate based on either of three specifications considered by Benati *et al.* (2017): (i) the Selden-Latané specification, in which both series enter the system in levels, i.e., $Y_t = [V_t \ R_t]'$; (ii) the semi-log specification, with $Y_t = [\ln(V_t) \ R_t]'$; and (iii) the log-log specification, with $Y_t = [\ln(V_t) \ \ln(R_t+1)]'$. The cointegrated VECMs feature no deterministic time trend (so, to be clear, the VECM estimator I use is the one described in pages 643-645 of Hamilton (1994)), reflecting my judgement that, for strictly conceptual reasons, neither series should be expected to exhibit a deterministic trend.²⁹

As in Benati *et al.* (2017), I bootstrap the tests *via* the procedure proposed by CRT (2012). In a nutshell, CRT’s procedure is based on the notion of computing critical and p -values by bootstrapping the model which is relevant under the null hypothesis. This means that, within the present context, the model which is being bootstrapped is a simple, non-cointegrated VAR in differences. All of the technical details can be found in CRT, which the reader is referred to. I select the VAR lag order as the maximum³⁰ between the lag orders chosen by the Schwartz and the Hannan-Quinn criteria³¹ for the VAR in levels.

Monte Carlo evidence on the performance of CRT’s procedure can be found in CRT (2012), Benati (2015), and especially Benati *et al.* (2017). Either paper documents the excellent performance of the procedure conditional on Data-Generation Processes (DGPs) featuring *no cointegration*, with the null incorrectly rejected at close the nominal size irrespective of the sample length. Benati *et al.* (2017), however, also show that, if the DGP features *cointegration*, the tests have a harder and harder time detecting it (i) the shorter the sample length, and (ii) the more persistent the cointegration residual. This is in line with some of the evidence reported by

²⁸Results from the trace tests are in line with those from the maximum eigenvalue tests, and they are available upon request.

²⁹For the short rate, the rationale for not including a time trend is obvious: The notion that nominal interest rates may follow an upward path (the possibility of a downward path is ruled out by the zero lower bound), in which they grow over time, is manifestly absurd. For M1 velocity, on the other hand, things are, at first sight, less obvious. The reason for not including a time trend originates from the fact that what I am here focusing on is a demand for money *for transaction purposes* (so this argument holds for M1, but it would not hold for broader aggregates). The resulting natural assumption of unitary income elasticity logically implies that, if the demand for M1 is stable, M1 velocity should inherit the stochastic properties of the opportunity cost of money. In turn, this implies that the type of unit root tests we run for M1 velocity should be the same as those we run for the nominal rate.

³⁰We consider the maximum between the lag orders chosen by the SIC and HQ criteria because the risk associated with selecting a lag order smaller than the true one (model mis-specification) is more serious than the one resulting from choosing a lag order greater than the true one (over-fitting).

³¹On the other hand, we do not consider the Akaike Information Criterion since, as discussed (e.g.) by Luetkepohl (1991), for systems featuring I(1) series the AIC is an inconsistent lag selection criterion, in the sense of not choosing the correct lag order asymptotically.

Engle and Granger (1987) based on the Augmented Dickey-Fuller test, and it implies that if cointegration is not detected, (i) and/or (ii) are possible explanations.

4.2.1 Exploring the tests' ability to detect cointegration *via* Monte Carlo

As I discuss in the next sub-section, in several cases Johansen's tests fail to reject the null of no cointegration. Assuming that cointegration truly is there in all samples—which, e.g., appears as a reasonable conjecture based on the evidence shown in Figures 2a and 2b—there are (at least) two possible interpretations of these results. First, they might simply be due to the 'luck of the draw': Whatever the truth about the underlying DGP is, no statistical test will ever get it right 100 per cent of the times. Second, in line with Engle and Granger's (1987) just-mentioned point, lack of rejection might simply be the figment of a short sample length and/or a highly persistent cointegration residual. In order to gauge an idea about how plausible this explanation is, Table 2 reports evidence from the following Monte Carlo experiment. For all those cases for which, in Table 1, Johansen's tests do not reject the null at the 10 per cent level, I estimate the VECM imposing one cointegration vector. Then, I stochastically simulate the VECM 2,000 times, for random samples of length equal to the actual sample length, and based on each simulated sample I perform the same trace and maximum eigenvalue tests I have previously performed based on the actual data, once again bootstrapping the p -values as in CRT (2012). Table 2 reports the empirical rejections frequencies (ERFs) at the 10 per cent level, i.e. the fractions of times, out of the 2,000 Monte Carlo simulations, for which maximum eigenvalues tests³² reject the null of no cointegration.

4.2.2 Evidence

The evidence in Table 1 is in line with that reported by Benati *et al.* (2017) in their investigation of the long-run demand for M1 since the mid-XIX century.

Starting from the Selden-Latané specification—which, according to Benati *et al.*'s (2017) findings, appears to be the one preferred by the data in the case of low-to-medium-inflation countries such as those analyzed herein—evidence of cointegration is uniformly strong based on quarterly data. For the two countries for which cointegration is not detected at the 10 per cent level, evidence from the ERFs in Table 2 is mixed. The ERF for South Africa shows that if cointegration truly were there, Johansen's tests would have a slightly less than even chance of detecting it. The one for Hong Kong, on the other hand, suggests that there is less than one chance out of five that the lack of rejection in Table 1 might be due to the problem discussed by Engle and Granger (1987). Based on annual data, cointegration is detected for the United Kingdom, Switzerland, and Canada, and for the United States in all instances except based on the aggregate also including MMFAs, and adjusted for the

³²Results for the trace tests are near-numerically identical.

share of currency held by foreigners. In most of the cases in which cointegration is not detected, however, the ERFs are quite low, or very low. For Japan, Australia, the Netherlands, and Finland they range between 0.168 and 0.463, thus pointing towards a small chance of detecting cointegration if it truly were in the data. For the remaining three cases, the ERFs range between about two-thirds and three-fourths, thus pointing towards a higher chance. In order to better interpret the results in Tables 1 and 2, it is useful to get back to the raw data shown in Figures 2a-2b. Consider for example Australia and Belgium: In both cases, visual evidence clearly suggests that the two series share a common trend. Combined with the ERFs reported in Table 2—equal to 0.168 and 0.699, respectively—this suggests that cointegration is highly likely in the case of Australia, and still pretty much likely for Belgium. A similar argument holds for Finland, Japan, and the Netherlands, although the visual evidence is weaker. As for the single case for which cointegration is not detected for the United States, visual evidence (not reported) is as strong as that shown for the other aggregate in Figure 2a. Combined with an ERF equal to 0.725, this suggests that there is a good chance that cointegration is indeed there.

In line with Benati *et al.* (2017), the evidence of cointegration in Table 1 is somehow weaker based on the semi-log specification, and it is uniformly very weak based on the log-log. It is to be noticed, however, that for either specification the ERFs in Table 2 are, in the vast majority of cases, either low or very low, so that, strictly speaking, most of these results are still compatible with the presence of cointegration.

Based on these results, in what follows I will therefore mostly focus on the Selden-Latané specification and, to a lesser extent, on the semi-log one, and I will instead eschew the log-log. Further, I will work under the assumption that, based on either specification, cointegration is there in *all* samples. The rationale for this is that, even in those cases in which cointegration is not detected, the evidence in Table 2 is most of the times, compatible with the presence of cointegration.

I next turn to the issue of stability of the cointegration relationship.

4.3 Testing for stability in the cointegration relationship

Tables A.4a and A.4b in the Appendix report results from Hansen and Johansen's (1999) Nyblom-type tests for stability in either the cointegration vector, or the vector of loading coefficients, in the estimated VECMs. The p -values reported in the two tables have been computed by bootstrapping, as in CRT (2012), the VECMs estimated conditional on one cointegration vector and no break of any kind, and then performing Hansen and Johansen's (1999) tests on the bootstrapped series. Before delving into the results, however, it is worth briefly discussing evidence on the performance of the tests.

4.3.1 Monte Carlo evidence on the performance of the tests

Table A.3 in the Appendix reports Monte Carlo evidence on the performance of the two tests conditional on bivariate cointegrated DGPs, for alternative sample lengths, and alternative extents of persistence of the cointegration residual, which is modelled as an AR(1). The table also reports results for a third test discussed by Hansen and Johansen (1999), the ‘fluctuation test’, which is essentially a joint test for time-variation in the cointegration vector and the loadings. The main results in the table can be summarized as follows:

(i) the two Nyblom-type tests exhibit an overall reasonable performance, incorrectly rejecting the null of no time-variation, most of the time, at roughly the nominal size. Crucially, this is the case irrespective of the sample length, and of the persistence of the cointegration residual.

(ii) The fluctuation test, on the other hand, exhibits a good performance (irrespective of the sample length) only if the persistence of the cointegration residual is low. The higher the persistence of the residual, however, the worse the performance, so that, e.g., when the AR root of the residual is equal to 0.95, for a sample length $T = 50$, the test rejects at twice the nominal size. The crucial point here is that, as extensively documented by Benati *et al.* (2017) high persistence of the cointegration residual is the empirically relevant case, as far as long-run money demand is concerned. Because of this, Tables A.4a and A.4b only report results from the two Nyblom-type tests.

4.3.2 Evidence

The key finding in the two tables is that evidence of breaks in either the cointegration vector, or the loading coefficients, is weak to non-existent. Specifically, focusing on the Selden-Latané specification (evidence for the semi-log is near-identical), I detect a break in the cointegration vector only for Mexico and Japan, and a break in the vector of loading coefficients only for New Zealand (1934-2016) and Canada (1967Q1-2016Q4). In all other cases, no break is detected.³³

In the next three sections I discuss the empirical evidence, starting from the issue of how the system adjusts towards its long-run equilibrium.

³³When a break is detected, I also estimate the break date as in (e.g.) Andrews and Ploberger (1994), as that particular date which maximizes the log-likelihood of the data, with the likelihood being computed based on the VECM with the break. The results are reported in Table A.5 in the Appendix, but they are not especially notable, and I will therefore not discuss them in any detail.

5 How Does the Cointegrated System Adjust Towards Equilibrium?

The evidence in the second column of Figure 2a showed that, in the post-WWII U.K., velocity's loading coefficient on the cointegration residual has been close to zero, and statistically insignificant, thus implying that the system's adjustment towards equilibrium has taken place *via* movements in the short rate, rather than movements in velocity. This evidence is powerful because it is *reduced-form*, and it therefore does not hinge on imposing any identifying restriction upon the data.

Tables A.6, and A.7a-A.7b in the Appendix show evidence for the remaining countries in the dataset. Table A.6 reports bootstrapped p -values for testing the null hypothesis that the loading coefficients on the cointegration residual in the VECM are equal to zero. Tables A.7a-A.7b report, based on the Selden-Latané specification, the estimated loading coefficients on the cointegration residual, with 90 per cent bootstrapped confidence intervals.

Overall, evidence is mixed, and it does not point towards a robust, clear-cut pattern across countries and sample periods. Based on the Selden-Latané specification, in particular, the previously discussed pattern for the United Kingdom also holds for New Zealand, the Netherlands, Korea, South Africa, and, in one case, for the United States. It is clear, however, that this pattern is not a general one, and that reduced-form evidence does not robustly suggest that the system's adjustment towards equilibrium consistently takes place *via* movements in the short rate, with no reaction of velocity to disequilibria.

These results, however, reflect the limitations of reduced-form evidence, which can only take you that far. As I show in the next section, a permanent-transitory decomposition along the lines of Cochrane (1994) produces in most instances a consistent pattern. The key point here is that velocity's loading coefficient in the VECM being equal to zero is *not* a necessary condition for velocity to be, to a close approximation, the permanent component of the short rate.

6 Evidence from a Permanent-Transitory Decomposition

Figures 4 to 7 show, for all countries, results from the same bivariate structural VECM I previously estimated for the United Kingdom in Section 2.3, in which permanent shocks are identified as the only shocks impacting upon the short rate in the infinite long run. Specifically, Figures 4 and 6 show, based on quarterly and annual data, respectively, the fractions of FEV explained by the permanent shock, with the 16th, 84th, 5th, and 95th percentiles of the bootstrapped distribution. Figures 5 and 7 show, based on quarterly and annual data, respectively, the IRFs to the transitory shock, together with the same percentiles of the bootstrapped distribution. As in Section




Figure 4 Results from bivariate structural VECMs for M1 velocity and the short rate: Fractions of forecast error variance explained by the permanent shock (based on quarterly data)




Figure 5 Results from bivariate structural VECMs for M1 velocity and the short rate: Impulse-response functions to the transitory shock (based on quarterly data)

2.3., following CRT (2012), bootstrapping has been implemented based on the VECM estimated conditional on one cointegration vector. Finally, Figures A.1a-A.1b in the appendix report the IRFs to the permanent shock, whereas Figures A.2a-A.2b show scatterplots of the permanent and transitory components of the two series.

6.1 Evidence from post-WWII quarterly data

With the single exception of Taiwan, evidence based on quarterly data is in line with that for the United Kingdom. Specifically,

(i) the fractions of FEV of velocity explained by the permanent shock are consistently very high, and most of the time close to one at nearly all horizons. This is especially clear for Canada, Australia, Korea, South Africa, Hong Kong, and Mexico, whereas evidence for the United States³⁴ is slightly weaker. By contrast, the fractions of FEV of the short rate are systematically lower than those of velocity at all horizons, and in several cases they are quite remarkably low, especially at the short horizons. This is the case, in particular, for the United States, Korea, South Africa, and Hong Kong. As in Section 2.3, it is important to stress that this result has been obtained in spite of the fact that the permanent shock has been identified as the one driving the unit root in the *short rate*, rather than in velocity.

(ii) Turning to the IRFs in Figure 5, the response of the short rate to transitory shocks is strongly statistically significant for all countries except Taiwan. As for velocity, the response is statistically insignificant at (nearly) all horizons for Canada, Australia, Korea, South Africa, Hong Kong and Mexico. As for the United States, it is insignificant (and, in fact, close to zero) on impact, whereas it is strongly significant further out.

Overall, with the exception of Taiwan, the evidence in Figures 4 and 5 is in line with that for the United Kingdom in Section 2.3, and it suggests that M1 velocity is, to a close approximation, the permanent component of the short rate. The evidence for Taiwan, however, should not be taken at face value—and in fact it appears as puzzling—for the following reason. Simple visual evidence based on the raw series shown in Figure 2b suggests that velocity is, in fact, smoother than the short rate. Indeed, once the two series have been rescaled so that they have the same sample standard deviation, the variance of the first difference of the short rate is 2.85 times the variance of the first difference of velocity. Since the two series are cointegrated (the p -value for the maximum eigenvalue test in Table 1 is equal to 0.001), and they are therefore driven by the same permanent shock, this simple evidence is hard to square with the variance decomposition in Figure 4, suggesting that the *short rate* is, essentially, the stochastic trend in the system. Because of this, I would argue that the evidence for Taiwan should be discounted. Finally, I do not discuss the IRFs to the permanent shock in Figure A.1a in the appendix because they are as expected

³⁴As mentioned, this evidence is based on the aggregate including MMDAs. Evidence on the alternative aggregate also including MMFAs is very close, and it is available upon request.

(i.e., both variables increase permanently) and they are not especially interesting.

On the other hand, it is worth spending a few words on the scatterplots of the permanent and transitory components in Figures A.2a. The main finding emerging from the figure is that whereas the correlation between the permanent components of the two series is, as expected, uniformly strong and positive, that between the transitory components is, in the vast majority of cases, weak to non-existent (this is especially clear for the United Kingdom, Canada, and Australia). The only exception to this pattern is the United States, for which the correlation between the transitory components is positive, but weaker than that between the permanent components. The obvious interpretation of this result is the fact that, as documented in Figure 5, whereas the response of the short rate to transitory shocks is uniformly strongly statistically significant (with the exception of Taiwan), the response of velocity is most of the time insignificant at (nearly) all horizons.³⁵ Once again, it is instructive to recall the parallel with the relationship between GDP and consumption: By the permanent income hypothesis, under rational expectations and no constraint on their ability to borrow, consumers should only react to permanent income shocks. As a result, a transitory GDP shock, by leaving consumption unaffected, would produce a zero conditional correlation between GDP and consumption.

6.2 Evidence from long-run annual data

Turning to the evidence based on annual data, support for this paper's main thesis is provided by the results for the United States, the United Kingdom, Switzerland, New Zealand, Australia, the Netherlands, and Finland. In all of these cases, permanent shocks explain very high fractions of the FEV of velocity (in several cases, very close to one) at all horizons, whereas they consistently explain lower fractions of the FEV of the short rate. The contrast is especially stark for the United States, with the fraction of explained FEV of velocity being (based on point estimates) consistently beyond 80 at all horizons, whereas the corresponding fraction for the short rate is between 5 and 10 per cent at horizons up to two years ahead, and even ten years ahead it only rises to less than 70 per cent. By the same token, for either of these seven countries the reaction of the short rate to transitory shocks is strongly statistically significant, whereas the corresponding IRF for velocity is insignificant at all horizons for the United Kingdom, Switzerland, New Zealand, the Netherlands, and Finland;

³⁵In his review of Friedman and Schwartz's *Monetary History of the United States*, James Tobin (1965, p. 478) conjectured that the relationship between velocity and the short rate might be the same for both their trend and cycle components:

‘A second interpretation [...] is that velocity follows the pro-cyclical movement of interest rates. This has the scientific virtue of providing a unified theoretical and statistical [...] explanation of both trend and cycle in velocity.’

The evidence in Figure I.2a (and the analogous evidence based on annual data in Figure I.2b) shows that this intuitively plausible conjecture is, in fact, incorrect.




Figure 6 Results from bivariate structural VECMs for M1 velocity and the short rate: Fractions of forecast error variance explained by the permanent shock (based on annual data)




Figure 7 Results from bivariate structural VECMs for M1 velocity and the short rate:
Impulse-response functions to the transitory shock (based on annual data)

it is insignificant on impact, and at short horizons, for Australia; and it is instead mostly strongly significant for the United States. For Canada, Japan, and Belgium the fraction of FEV of the short rate explained by permanent shocks is (based on point estimates) consistently greater than the corresponding fraction for velocity (although for Belgium the difference is quite small). As for the IRFs to transitory shocks, they are uniformly insignificant at all horizons for either variable, and either country. Overall, the evidence based on annual data appears somehow weaker than that based on quarterly data, with three countries out of ten failing to support this paper's main thesis. As in the case of Taiwan, however, I would argue that the results for Canada and Belgium should be discounted. The reasons are the same I gave there: First, the visual evidence in Figure 2a quite clearly suggests that for both countries velocity is appreciably smoother than the short rate; Second, once the series are rescaled so that they have the same sample standard deviation, for both countries the first difference of the short rate is markedly more volatile than the first difference of velocity.³⁶ Again, since the two series share the same stochastic trend, this is hard to square with the notion that the short rate might be closer to such trend than velocity. If we accept this argument, this leaves us with Japan as the single country which truly seems to contradict this paper's argument. Finally, I do not discuss the IRFs to permanent shocks shown in Figure A.1b because, again, they are as expected, and not especially interesting. As for the scatterplots of the permanent and transitory components shown in Figures A.2b, the only point worth mentioning is that the correlation between the transitory components is weakly positive only for the United States and Australia; it is flat for the Netherlands; and it is weakly *negative* for all other countries.³⁷ As we will see in Section 8.3, this result arises naturally within the Sidrausky framework.

6.3 Summing up

Overall, the evidence in this section provides substantial—although by no means perfect—support to my thesis that M1 velocity is, to a close approximation, the permanent component of the short rate. Evidence is strong for eight countries out of nine based on quarterly data, and for seven countries out of ten based on annual data. On the other hand, evidence is negative—but, for the reasons I gave, it should arguably be discounted—for Taiwan based on quarterly data, and for Canada and Belgium based on annual data. This leaves us with only one country, Japan, for which evidence appears to quite clearly contradict my argument.

Finally, there are a few countries for which evidence also supports my thesis, but

³⁶For Belgium and Canada, respectively, the variance of the first difference of the short rate is 11.01 times, and 8.30 times greater than the variance of the first difference of rescaled velocity.

³⁷This provides additional perspective on James Tobin's (1965) conjecture (which I mentioned in footnote 29) that the relationship between velocity and the short rate might be the same for the trend and cyclical components of the two series.

whose results I have chosen not to report because their sample periods are quite short.³⁸ This is the case for Denmark and Sweden (for either country the sample period is 1993Q1-2017Q1): In particular, in both countries permanent shocks induce an insignificant response of velocity at all horizons, and a statistically significant response in the short rate (the fractions of explained FEV, on the other hand, are very high for both variables). Results for New Zealand based on quarterly data for the period 1988Q2-2016Q4 are qualitatively the same as those based on annual data discussed herein (in fact, they are significantly stronger).³⁹ The same holds for Switzerland for the period 1985Q1-2017Q1: Since these results are qualitatively the same as those reported in Figures 6 and 7, I have preferred not to report them. Finally, results for the Euro area since 1999Q1⁴⁰ exhibit exactly the same pattern as Denmark and Sweden.

I now turn to discuss evidence for monetary regimes which, historically, have caused inflation to be $I(0)$, such as inflation-targeting regimes. As previously mentioned, the most interesting feature of these regimes is that, by eliminating permanent inflation shocks, they cause M1 velocity—if my argument is correct—to be essentially a linear transformation of the natural rate of interest.

7 Evidence from Monetary Regimes Causing Inflation to Be $I(0)$

I start by discussing simple, *prima facie* evidence that under these regimes velocity might be driven, to a dominant extent, by the natural rate of interest. I then turn to estimating the natural rate for two inflation-targeting countries within a cointegrated SVAR framework.

7.1 Does velocity closely co-move with trend GDP growth?

As discussed in Section 2.4.3, within a vast class of models the natural rate of interest is a linear function of trend output growth. This implies that, if my argument is correct, under regimes causing inflation to be stationary we should see a strong correlation between velocity and trend GDP growth. Figure 8 provides evidence for eight such regimes, specifically: four inflation-targeting countries (United King-

³⁸These results are available upon request.

³⁹E.g., the fraction of FEV of velocity explained by permanent shocks is near-identical to one at all horizons, whereas the corresponding fraction for the short rate is below 10 per cent up to four years ahead. And the IRFs are very close to those for the United Kingdom in Figure 1*b*.

⁴⁰Data for the pre-EMU period are synthetic (i.e., reconstructed *ex post*), and so I decided to eschew them, to keep in line with my exclusive use, throughout the entire paper, of authentic (i.e., non reconstructed *ex post*) data from central banks and national statistical agencies.

dom, Canada, Australia, and New Zealand);⁴¹ European Monetary Union (EMU); Switzerland under the post-1999 ‘new monetary policy concept’, which is conceptually akin to EMU; West Germany/Germany up until the beginning of EMU (i.e., December 1998);⁴² and Denmark, which has consistently followed a policy of pegging the Krone first to the Deutsche Mark, and then to the Euro, thus importing the strong anti-inflationary stance of the Bundesbank, and then of the European Central Bank (ECB). In line with the evidence reported in Benati (2008) based on Hansen’s (1999) estimator of the sum of the autoregressive coefficients, for all countries—with the exception of Switzerland and the Euro area—Elliot *et al.*’s. (1996) tests strongly reject the null of a unit root for inflation.⁴³ In spite of the results from unit root tests, I have chosen to also consider Switzerland and the Euro area for the following reasons. As for Switzerland, results from unit root tests are most likely a figment of the short sample period: for the sample starting in 1980Q1 (when GDP deflator data start being available) a unit root is rejected very strongly. This is in line with Switzerland’s reputation as a hard-currency, low-inflation country.⁴⁴ As for the Euro area, visual evidence clearly suggests that the collapse of Lehman Brothers, which unleashed the most violent phase of the Great Recession, was associated with a downward shift in mean inflation, from 2.01 per cent over the period 1999Q1-2008Q3, to 0.99 per cent over the period 2008Q4-2016Q4. Once controlling for this break in the mean, a unit root is very strongly rejected, thus showing that the previous lack of rejection was a simple illustration of Pierre Perron’s (1989) well-known argument. My decision to also consider the Euro area reflects the fact that, in spite of such downward

⁴¹Inflation targeting was introduced in October 1992 in the United Kingdom; in February 1991 in Canada; and in February 1990 in New Zealand. As for Australia, there never was an explicit announcement of the introduction of the new regime. Here I follow Benati and Goodhart (2011) in taking 1994Q3 as the starting date of the inflation-targeting regime. The rationale is that, based on the central bank’s communication, during those months it became apparent that the bank was indeed following an inflation-targeting strategy. Finally, I do not report results for Sweden (the available sample is 1998Q1-2017Q2) because they are manifestly puzzling. Both a linear trend estimated *via* OLS, and simple averages computed for the first and second halves of the sample, clearly suggest that trend GDP growth has progressively decreased, which would be in line with the steady decrease in M1 velocity since 1998. SW’s estimate of trend growth, on the other hand, is essentially flat over the entire period.

⁴²I also consider the period after unification in order to have a longer sample period (quarterly data for West Germany’s nominal GDP start in 1970). Reunification caused a jump in both nominal GDP and M1, but, from a conceptual point of view, it does not cause any problem for the computation of velocity (i.e., their ratio). As for real GDP growth, I treat the very large observation for the quarter corresponding to reunification as an outlier, and, following Stock and Watson (2002), I replace it with the median value of the six adjacent quarters.

⁴³These results are reported in Table A.5 in the appendix. By the same token, Hansen’s (1999) bias-corrected estimate of the sum of the autoregressive coefficients in an $AR(p)$ representation for inflation clearly suggest that in all cases (again, with the exception of the Euro area and Switzerland) inflation is (close to) white noise.

⁴⁴Over the entire period since World War I, Swiss inflation has been equal, on average, to 1.9 per cent.




Figure 8 Evidence from monetary regimes causing inflation to be $I(0)$: M1 velocity and Stock and Watson (1996, 1998) TVP-MUB estimate of trend real GDP growth

shift in the mean of inflation, inflation expectations (as measured by the ECB’s Survey of Professional Forecasters) have remained well-anchored,⁴⁵ thus suggesting that agents have interpreted such shift as *temporary*.⁴⁶ Finally, I also show evidence for the United States for the period following the Volcker disinflation⁴⁷ for the following reason. Although Elliot *et al.*’s (1996) tests for sample periods following the end of the Volcker disinflation typically do not reject the null of a unit root,⁴⁸ this evidence does not square well with the fact that during this period inflation has been broadly stable. A likely explanation is that, after the Volcker stabilization, U.S. inflation has still exhibited a small unit root component. Results from Cochrane’s (1988) variance ratio estimator provide support for this conjecture: Over the period 1984Q1-2017Q2, the size of the unit root in U.S. GDP deflator inflation has been slightly below 10 per cent.⁴⁹ Such a small unit root component should introduce a small permanent ‘wedge’ between actual M1 velocity, and the value velocity would have taken in the absence of permanent variation in inflation. Such wedge should however be quite small, so that the same argument I made for monetary regimes under which inflation has been $I(0)$ should also approximately apply to the United States.

Figure 8 shows M1 velocity together with a SW (1996, 1998) TVP-MUB estimate of trend GDP growth, based on a time-varying parameters $AR(p)$. My implementation of SW’s methodology is exactly the same as in Benati (2007), which the reader is referred to for details. The correlation between the two series is strong for all countries, with the single exception of Australia. To be sure, by no means does this represent a hard proof that my thesis is correct. At the very least, however, it is compatible with my argument. For the United States the correlation with trend GDP growth is stronger for the velocity series based on the expanded M1 aggregate also including MMFAs, whereas it is somehow weaker based on the aggregate only including MMDAs.

I now proceed to exploit the insight that under such monetary regimes velocity is driven, to a dominant extent, by the natural rate of interest, in order to estimate it within a cointegrated SVAR framework.

⁴⁵Figure A.3 in the appendix shows the inflation forecasts from the ECB’s Survey of Professional Forecasters at three alternative horizons, 1-, 2-, and 5-years ahead. Over the entire period since 1999Q1, the 5-years ahead forecast has fluctuated between 1.8 and 2.0 per cent.

⁴⁶Indeed, as I write (August 2017) Euro area inflation has been moving, in recent months, back towards 2 per cent.

⁴⁷Specifically, for the period 1984Q1-2016Q4: Following Clarida, Gali, and Gertler (2000), I take 1983Q4 as marking the end of the disinflation.

⁴⁸For sample periods starting in the first quarter of each year from 1984Q1 to 1999Q1, a unit root in inflation can be rejected, at the 10 per cent level, only for the samples starting in 1988Q1, 1990Q1, and 1992Q1.

⁴⁹See Figure A.4 in the appendix. Bootstrapped confidence bands have been computed *via* spectral bootstrapping of the first-difference of inflation as in Benati (2007).




Figure 9 Estimated natural rate of interest for the United Kingdom and Canada under inflation targeting

7.2 Estimating the natural rate for two inflation-targeting countries

Figure 9 shows, for the United Kingdom and Canada under inflation targeting, estimates of the natural rate together with the 16th, 84th, 5th, and 95th percentiles of the bootstrapped distribution.⁵⁰ For either country, the VECM features M1 velocity, the central bank’s monetary policy rate, and three additional nominal interest rates.⁵¹ Since in both countries inflation is statistically indistinguishable from white noise, I do not enter it in the VECM. Rather, I subtract the average inflation rate over the inflation-targeting regime from either of the interest rates entering the VECM, thus converting them, to a close approximation, into *ex ante* real rates.⁵² For both countries, Elliot *et al.*’s (1996) tests strongly reject the null of a unit root for either of the spreads between rates at longer maturities and the central bank’s monetary policy rate.⁵³ This implies that, beyond the cointegration vector between velocity and the short rate, either country’s VECM features three additional cointegration vectors.⁵⁴ In turn, this implies that the system is driven by a single permanent shock, i.e., the shock to the natural rate, which I identify as the only shock affecting velocity in the infinite long run. I will not comment upon the results because they speak for themselves: In either country, the natural rate is estimated to have been declining since the early 1990s, slipping below zero following the collapse of Lehman Brothers.

7.2.1 A key advantage of using M1 velocity in order to estimate the natural rate of interest

A key advantage of using M1 velocity in order to estimate the natural rate of interest is that, although it is subject to the Zero Lower Bound (ZLB) exactly as nominal interest rates, as I now document such constraint is, in practice, much less binding. As I show, unless the ‘notional’ short rate—i.e., the short rate which *would* prevail in the absence of the ZLB—is ‘very’ negative, M1 velocity remains positive, and is therefore not affected by the ZLB constraint. This implies that whereas estimates of

⁵⁰Both estimation and bootstrapping details are the same as before, that is, I use Johansen’s estimator of the cointegrated VECM, and I bootstrap as in CRT (2012).

⁵¹For the United Kingdom, they are the 3-month Treasury bill rate, the 10-year government bond yield, and the yield on long-term consols. For Canada, they are the 10-year government bond yield, and the average yields on government bonds with maturities between 3 and 5, and 5 and 10 years.

⁵²The rationale, quite obviously, is that the conditional and unconditional forecasts of a white noise process are equal to its unconditional mean.

⁵³For Canada, this is the case for either the inflation-targeting regime, or the full sample. For the United Kingdom, this is the case only for the full sample. This is likely due to the very high persistence of the spreads, so that in small samples it is hard to reject the null of a unit root. Cochrane (1994) has a discussion of how the analogous lack of rejection of a unit root in the consumption/GDP ratio in small samples should be discounted.

⁵⁴For Canada, Johansen’s tests identify indeed four cointegration vectors. For the United Kingdom they only identify two. Since Elliot *et al.*’s (1996) tests strongly reject a unit root in all of the three spreads, I regard these results as a statistical fluke.

the natural rate exploiting the informational content of inflation and nominal interest rates (and possibly of other variables *except* M1 velocity) may be distorted by the fact that the ZLB is constraining nominal rates, this is much less likely to happen when the estimates also exploit the informational content of velocity. Intuitively, a series of negative shocks to the natural rate of interest which, within a low-inflation environment, drag the short rate down to the ZLB will leave velocity strictly above zero, unless they are very large. This means that, by focusing on velocity, an econometrician will be able to identify such shocks, which would not be the case if (s)he only focused on nominal rates and inflation.

Table 3 reports, based on the Selden-Latané specification,⁵⁵ the estimated nominal short-term interest rate corresponding to the ZLB on M1 velocity. For the sake of brevity I only report estimates based on post-WWII quarterly data, but the full set of estimates is available upon request. The estimates have been computed by setting $V_t = \epsilon_t = 0$ in expression (4), so that $R_{ZLB} = -\alpha/\beta$. I estimate (4) based on SW's (1993) estimator⁵⁶ of the cointegration regression, thus obtaining estimates of α and β , and therefore of the cointegration residual. Estimation of the VECM conditional on SW's (1993) estimate of the cointegration residual is then implemented based on the two-step procedure discussed in Luetkepohl (1991, pp. 370-371). Finally, in order to construct bootstrapped confidence intervals for R_{ZLB} , I bootstrap the estimated VECM as in CRT (2012), and I apply the same, just-described procedure to each bootstrapped replication.

⁵⁵For either the semi-log or the log-log specification, the ZLB on log velocity is not binding by definition. Since, as documented in Section 4, these specifications are not the ones the data seem to prefer (at least, for the low-to-medium inflation countries analyzed herein), this is however largely irrelevant.

⁵⁶The reason for using SW's estimator, rather than Johansen's, is that whereas both estimators produce estimates of the cointegration vector, only SW's produce an estimate of α in (4), that is, of the intercept of the cointegration regression *in levels*. Johansen's procedure, on the other hand, produces an estimate of the intercept for the *VECM representation*, which for the present purposes is useless.

Table 3 Estimated nominal short-term interest rate corresponding to the Zero Lower Bound on M1 velocity, with 90 per cent-coverage bootstrapped confidence interval^a		
United States		
<i>standard</i> M ₁ + MMDAs	1959Q1-2016Q4	-8.47 [-14.46; -6.21]
<i>standard</i> M ₁ + MMDAs + MMMFs	1959Q1-2016Q4	-5.93 [-24.52; -0.17]
<i>Adjusting for currency held by foreigners:</i>		
<i>standard</i> M ₁ + MMDAs	1959Q1-2016Q4	-9.42 [-15.49; -6.86]
<i>standard</i> M ₁ + MMDAs + MMMFs	1959Q1-2016Q4	-6.22 [-23.51; -0.64]
United Kingdom	1955Q1-2016Q4	-4.39 [-12.47; -2.16]
Canada	1967Q1-2016Q4	-7.65 [-17.95; -3.98]
Australia	1975Q1-2016Q3	-6.45 [-14.58; -4.56]
Taiwan	1961Q3-2016Q4	0.34 [-1.62; 1.29]
South Korea	1964Q1-2017Q1	-5.13 [-8.05; -4.01]
South Africa	1985Q1-2017Q1	-6.87 [-36.04; -1.50]
Hong Kong	1985Q1-2017Q1	-2.36 [-4.39; -0.17]
Mexico	1985Q4-2017Q1	-38.22 [-52.40; -27.34]
^a Based on 10,000 bootstrap replications.		

With the single exception (once again) of Taiwan, the point estimates in Table 3 are uniformly negative, and sometimes quite significantly so. This implies that unless the ‘notional’ short rate becomes quite significantly negative, M1 velocity will remain unconstrained by the ZLB, and—different from nominal interest rates—will keep correctly reflecting shocks to the natural rate.

8 Conclusions

In this paper I have shown that, since World War I, M1 velocity has been, to a close approximation, the permanent component of the short rate, so that the time-series relationship between the two series has been the same as that between consumption and GDP. The logical implication is that, under monetary regimes which cause inflation to be $I(0)$, permanent fluctuations in M1 velocity uniquely reflect, to a close approximation, permanent shifts in the natural rate of interest. Evidence from West Germany and several inflation-targeting countries is compatible with this notion, with velocity fluctuations being, most of the time, strongly correlated with a Stock and Watson (1996, 1998) estimate of trend real GDP growth. I exploit this insight to estimate the natural rate of interest for the United Kingdom and Canada under inflation targeting: In either country, the natural rate has been consistently declining since the early 1990s. A second implication of these results is that talking of ‘money demand instability’, or real money balances as ‘being out of equilibrium’, makes no

logical sense, as it is conceptually akin to talking of ‘instability’ of the relationship between GDP and potential GDP.

References

- ALVAREZ, F., AND F. LIPPI (2009): “Financial Innovation and the Transactions Demand for Cash,” *Econometrica*, 77(2), 363–402.
- ANDREWS, D. K., AND W. PLOBERGER (1994): “Optimal Tests When a Nuisance Parameter is Present Only Under the Alternative,” *Econometrica*, 62(6), 1383–1414.
- BENATI, L. (2007): “Drift and Breaks in Labor Productivity,” *Journal of Economic Dynamics and Control*, 31(8), 2847–2877.
- (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.
- BENATI, L., AND C. GOODHART (2011): “Monetary Policy Regimes and Economic Performance: The Historical Record, 1979-2008,” in *B. Friedman, B., and Woodford, M. (eds.), Handbook of Monetary Economics, Volume 3, North Holland*.
- BENATI, L., R. E. LUCASJR., J.-P. NICOLINI, AND W. WEBER (2017): “Long-Run Money Demand Redux,” *University of Bern, University of Chicago, Minneapolis FED, and University of South Carolina, mimeo*.
- 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.
- CLARIDA, R., J. GALI, AND M. GERTLER (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, CXV(1), 147–180.
- COCHRANE, J. H. (1988): “How Big Is the Random Walk in GNP?,” *Journal of Political Economy*, 96(5), 893–920.
- (1994): “Permanent and Transitory Components of GNP and Stock Prices,” *Quarterly Journal of Economics*, 109(1), 241–265.
- 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.
- ERCEG, C., AND A. LEVIN (2003): “Imperfect Credibility and Inflation Persistence,” *Journal of Monetary Economics*, 50(4), 915–944.
- ESTRELLA, A., AND F. MISHKIN (1997): “Is There a Role for Monetary Aggregates in the Conduct of Monetary Policy?,” *Journal of Monetary Economics*, 40, 279–304.
- FAMA, E. (1975): “Short-Term Interest Rates as Predictors of Inflation,” *American Economic Review*, 65(3), 269–282.
- 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*, 3:1973, 577–646.
- (1976): “The Case of the Missing Money,” *Brookings Papers on Economic Activity*, 3:1976, 683–739.
- HAMILTON, J. (1994): *Time Series Analysis*. Princeton, NJ, Princeton University Press.
- HANSEN, B. E. (1999): “The Grid Bootstrap and the Autoregressive Model,” *Review of Economics and Statistics*, 81(4), 594–607.
- HANSEN, H., AND S. JOHANSEN (1999): “Some Tests for Parameter Constancy in Cointegrated VAR Models,” *Econometrics Journal*, 2, 306–333.
- JUDSON, R. (2017): “The Death of Cash? Not So Fast: Demand for U.S. Currency at Home and Abroad, 1990-2016,” Federal Reserve Board, Division of International Finance Working Paper, forthcoming.
- LATANÉ, H. A. (1960): “Income Velocity and Interest Rates: A Pragmatic Approach,” *Review of Economics and Statistics*, 42(4), 445–449.
- LAUBACH, T., AND J. WILLIAMS (2003): “Measuring the Natural Rate of Interest,” *The Review of Economics and Statistics*, 85(4), 1063–1070.
- LUCASJR., R. E. (1976): “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.
- (1988): “Money Demand in the United States: A Quantitative Review,” *Carnegie-Rochester Conference Series on Public Policy*, 29, 137–168.

- LUCASJR., 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.
- PERRON, P. (1989): “The Great Crash, the Oil Price Shock and the Unit Root Hypothesis,” *Econometrica*, 57(6), 1361–1401.
- 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.
- SHIN, Y. (1994): “A Residual-Based Test of the Null of Cointegration against the Alternative of No Cointegration,” *Econometric Theory*, 10(1), 91–115.
- SIDRAUSKI, M. (1967a): “Inflation and Economic Growth,” *Journal of Political Economy*, 75(6), 796–810.
- (1967b): “Rational Choice and Patterns of Growth in a Monetary Economy,” *American Economic Review*, 57(2), 534–544.
- STOCK, J., AND M. WATSON (2002): “Has the Business Cycle Changed and Why?,” in *B. Bernanke and K. Rogoff, eds. (2003), NBER Macroeconomics Annuals 2002*.