

The Demand for Money at the Zero Interest Rate Bound

Tsutomu Watanabe
Tomoyoshi Yabu

Research Project on Central Bank Communication
702 Faculty of Economics, The University of Tokyo,
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
Tel: +81-3-5841-5595 E-mail: watlab@e.u-tokyo.ac.jp
<http://www.centralbank.e.u-tokyo.ac.jp/en/>

The Demand for Money at the Zero Interest Rate Bound*

Tsutomu Watanabe[†] Tomoyoshi Yabu[‡]

September 4, 2018

Abstract

This paper estimates a money demand function using US data from 1980 onward, including the period of near-zero interest rates following the global financial crisis. We conduct cointegration tests to show that the substantial increase in the money-income ratio during the period of near-zero interest rates is captured well by the money demand function in log-log form, but not by that in semi-log form. Our result is the opposite of the result obtained by Ireland (2009), who, using data up until 2006, found that the semi-log specification performs better. The difference in the result from Ireland (2009) mainly stems from the difference in the observation period employed: our observation period contains 24 quarters with interest rates below 1 percent, while Ireland's (2009) observation period contains only three quarters. We also compute the welfare cost of inflation based on the estimated money demand function to find that it is very small: the welfare cost of 2 percent inflation is only 0.04 percent of national income, which is of a similar magnitude as the estimate obtained by Ireland (2009) but much smaller than the estimate by Lucas (2000).

JEL Classification Numbers: C22; C52; E31; E41; E43; E52

Keywords: money demand function; cointegration; zero lower bound; near-zero interest rates; welfare cost of inflation; log-log form; semi-log form; interest elasticity of money demand

*We would like to thank Kosuke Aoki, Yoichi Arai, Hiroshi Fujiki, Andy Levin, and Keisuke Otsu for helpful discussions. This research forms part of the project on “Central Bank Communication Design” funded by a JSPS Grant-in-Aid for Scientific Research (No. 18H05217). Tomoyoshi Yabu gratefully acknowledges financial support from the JSPS through the Grant-in-Aid for Scientific Research No. 17K03663.

[†]Graduate School of Economics, University of Tokyo. E-mail: watanabe@e.u-tokyo.ac.jp, Website: <https://sites.google.com/site/twatanabelab/>.

[‡]Faculty of Business and Commerce, Keio University. E-mail: tyabu@fbc.keio.ac.jp.

1 Introduction

In regression analyses of money demand functions, there is no consensus on whether the nominal interest rate as an independent variable should be used in linear or log form. For example, Meltzer (1963), Hoffman and Rasche (1991), and Lucas (2000) employ a log-log specification (i.e., regressing real money balances (or real money balances relative to nominal GDP) in log on nominal interest rates in log), while Cagan (1956), Lucas (1988), Stock and Watson (1993), and Ball (2001) employ a semi-log specification (i.e., nominal interest rates are *not* in log).

Identifying empirically the proper specification of the money demand function is important for the following reasons. First, different specifications of the money demand function come from different economic models of money demand. For example, the inventory theoretic approach to money demand by Baumol (1952), Tobin (1956), and Miller and Orr (1966) implies a money demand function in log-log form. Similarly, the transactions time approach to money demand by McCallum and Goodfriend (1989) implies a money demand function in log-log form. On the other hand, the money in the utility function approach implies a money demand function in log-log form if the utility function is of constant relative risk aversion (Lucas (2000)) but a money demand function in semi-log form if the utility function is quasi-linear (Cysne (2009)). By empirically specifying the functional form of the money demand function, we can learn which economic model is more appropriate.

Second, different specifications of the money demand function have different implications for the welfare cost of inflation. Lucas (2000) shows that, with a log-log money demand function, a slight increase in the nominal interest rate from the zero lower bound generates a substantial decline in real money balances. Therefore, there is a significant welfare cost of deviating just slightly from the Friedman rule (Friedman (1969)). On the other hand, with a semi-log specification, a slight increase in the nominal interest rate from zero does not change money demand that much, so that the costs associated with it are smaller. Therefore, with a log-log specification, most of the benefits of reducing inflation come when moving from zero inflation to the optimal deflation rate needed to achieve the Friedman rule, but this is not the case with a semi-log specification.

Third, different specifications of the money demand function have different implications for the conduct of monetary policy near the zero lower bound. In the case of a log-log specification, money demand becomes infinitely large at the zero lower bound, so that the opportunity cost of holding money cannot go below zero. On the other hand, with a semi-log money demand function, the marginal utility of money becomes zero at a finite level of real money balances and becomes negative beyond that level.¹ Rognlie (2016), in the context of a negative interest rate policy, shows that, with a semi-log money demand function, the negative opportunity cost of holding money created by the negative interest rate policy may violate the Friedman rule in the opposite of the traditional direction, so that such a policy could deteriorate economic welfare.

Using annual data for the United States for the period 1900-1994, Lucas (2000) shows that a log-log specification fits the data more closely than a semi-log specification, and that the welfare cost of inflation is substantially large. For example, when the inflation rate is 10 percent, the welfare cost reaches 1.8 percent of national income. On the other hand, extending the observation period to 2006 so that observations in the period of near-zero interest rates are included, Ireland (2009) shows that a semi-log specification performs better than a log-log specification, and that the welfare cost of inflation is much smaller than estimated by Lucas (2000).

Key to identifying which specification is more appropriate – a semi-log or a log-log specification – is, as pointed out by Ireland (2009), the inclusion of a period of near-zero interest rates when estimating the money demand function, since the difference in money demand between the two specifications is extremely large near the zero lower bound. However, the period of interest rates below 1 percent in Ireland’s (2009) observation period is very short, comprising only three quarters (2003:Q3-2004:Q1), so that his results may

¹The marginal net benefits of holding money can become negative if the storage cost of holding money (especially the storage cost of holding cash) is non-negligible. Specifically, in the context of a negative interest rate policy, Rognlie (2016) and Eggertsson et al. (2017) consider the case in which the storage cost of holding money is positive and the marginal cost associated with holding money is positive as well. In this case, people hold money even if financial assets other than money yield a negative interest rate, so that the rate of interest can be negative in equilibrium. Note that if the marginal storage cost is an increasing function of real money balances (i.e., the storage cost is a convex function of real money balances), there is no lower bound on the nominal interest rate. This is in sharp contrast with the case discussed by Hicks (1937), who considers a situation in which the cost of holding money is negligible, so that the rate of interest must always be non-negative.

not be very robust. In this paper, we extend the observation period up until 2013, so that we have more observations with near-zero interest rates due to monetary easing after the global financial crisis. Specifically, we update sweep-adjusted M1 up to 2013 and repeat the exercises conducted by Ireland (2009). We show that inclusion of the more recent observations with near-zero interest rates changes the results substantially.²

The rest of the paper is organized as follows. In Section 2, we visually compare the log-log and semi-log specifications using annual data. In Section 3, we conduct cointegration tests, using quarterly data, for the money-income ratio and the nominal interest rate to examine which of the two specifications performs better. In Section 4, we check the robustness of our baseline result in three respects: (1) we check how the result changes if we use alternative measures for the opportunity cost of holding money; (2) we allow structural breaks in the cointegration relationships; and (3) we relax the constraint that the income elasticity of money demand is equal to unity. In Section 5, we discuss the welfare cost of inflation based on the estimation result of the money demand function. Section 6 concludes the paper.

2 Data overview

In this section, we conduct a visual examination of the relationship between money demand and the nominal interest rate. We use the same annual data from 1980 onward as Ireland (2009). Specifically, the nominal interest rate is the six-month commercial paper rate for 1980 to 1997 and the three-month AA nonfinancial commercial paper rate from 1998 onward. The nominal money balances are M1 for 1980 to 1993, and the constituent elements of M1 are currency held by the public, non-interest-bearing demand deposits, and interest-bearing negotiable order of withdrawal (NOW) accounts. For 1994 and later, we use retail sweep-adjusted M1 to avoid the influence of the introduction of the retail deposit sweep programs (see Dutkowsky and Cynamon (2003)). We extend the observation period

²A study closely related to ours is that by Mogliani and Urga (2018), who examine the stability of the money demand function using an annual data set that includes the recent near-zero interest rate period (1900-2013). They find structural breaks in 1945 and 1976 and show that, with these two structural breaks taken into consideration, there exists a cointegration relationship between the money-income ratio and the nominal interest rate. However, they focus only on the log-log specification and do not conduct a log-log versus semi-log comparison.

up to 2013, the latest year for which retail sweep-adjusted M1 is available.

Figure 1 shows a scatter plot of the US money-income ratio (sweep-adjusted M1 relative to nominal GDP) and the nominal interest rate. The years examined by Ireland (2009), 1980 to 2006, are represented by the circles. The newly added years, 2007-2013, are shown by x-marks. As argued by Ireland (2009), as far as the data up to 2006 are concerned, the money-income ratio seems to increase linearly as the interest rate declines and take a finite positive value at the zero interest rate bound, so that a semi-log specification looks like the appropriate specification. However, the more recent observations show that when the interest rate is closer to the zero lower bound, the money-income ratio increases more elastically as the interest rate approaches zero from above, suggesting that a log-log specification provides a better fit.

To conduct a more detailed comparison between the log-log and semi-log specifications, we plot the data into the semi-log graph in Figure 2(a) and into the log-log graph in Figure 2(b). Figure 2(a) shows that there exists a linear relationship between the log of the money-income ratio and the interest rate until 2006, but once the more recent observations are added, this linear relationship disappears. In this sense, Ireland's (2009) finding that the semi-log specification fits the data well is not robust to the addition of the more recent data. On the other hand, Figure 2(b) shows that, for the period until 2006, there appears to exist a linear relationship between the two variables – although it does not look as straight as that in Figure 2(a) for the period until 2006 – and that the linear relationship appears to survive even when adding the more recent observations. More specifically, the nominal interest rate is very close to zero from the start of monetary easing in 2008 onward, so that the dots for this period line up horizontally in Figure 2(a). In contrast, in Figure 2(b), the interest rate in log continues to decline even from 2008 onward, so that the linear relationship observed until 2006 remains unchanged.

3 Cointegration tests

In this section, we conduct cointegration tests to more rigorously examine the findings in the previous section based on casual examination of the data. Ireland (2009) conducts residual-based cointegration tests using quarterly data (1980:Q1-2006:Q4) to see which of the two

specifications is supported by the data. Specifically, if the residual from a regression of the log of the money-income ratio on the interest rate is stationary, this means that the two variables are cointegrated. In this case, the semi-log specification is supported by the data. On the other hand, if the residual from a regression of the log of the money-income ratio on the *log* of the nominal interest rate is stationary, the log-log specification is accepted. We follow this approach and examine in which specification a cointegration relationship exists using the data up until 2013:Q4.

Let m denote the ratio of retail sweep-adjusted M1 divided by nominal GDP and r the three-month US Treasury bill (TB) rate. For nominal GDP, we use the figures with base year 2009 rather than those with base year 2000. We start by conducting unit root tests – i.e., the augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests – for $\ln(m)$, $\ln(r)$, and r with only a constant term included to find that the null hypothesis of a unit root is not rejected for all of the three variables. Given this result, we examine in the rest of this section whether there exists a cointegration relationship between the variables.

We regress $\ln(m)$ on a constant and $\ln(r)$ or r to see whether the residual obtained is stationary or not.³ We employ three different cointegration tests: the ADF test proposed by Engle and Granger (1987); the PP test proposed by Phillips and Ouliaris (1990) with test statistics given by Z_t ; and the same PP test but with test statistics given by Z_α . Some simulation studies show that the ADF test has the least size distortions and is more reliable than the PP tests, but that the Z_α test has more power than the ADF and Z_t tests (see Haug (1996)). Based on this, Haug (1996) proposes to employ more than one cointegration test in applied research.⁴ Given that the different tests each have their relative advantages and disadvantages, it is useful to present the results from the three tests and compare

³Note that we assume here that the income elasticity of money demand is unity. Some previous studies using US annual data, such as Meltzer (1963), Lucas (1988, 2000), and Stock and Watson (1993), show that the income elasticity of money demand is close to unity. While we assume in this section that income elasticity is unity, we will relax this assumption in the next section to see how much the result depends on it.

⁴The finite sample properties of residual-based cointegration tests are similar to those of unit root tests. Several studies show that the PP tests have more size distortions than the ADF test when MA errors have a negative autocorrelation (see Phillips and Perron (1988)) and that the ADF test has better size and power when the errors have an AR structure (see DeJong et al. (1992)). This suggests that the PP tests are less reliable than the ADF test. On the other hand, some simulation studies suggest that the size-adjusted power of the PP tests is higher than the power of the ADF test (see Campbell and Perron (1991))

them. Note that Ireland (2009) reports the result of the PP test with the test statistics given by Z_t but does not discuss the other two tests.

Let us start by reproducing Ireland's (2009) result using the same observation period (1980:Q1-2006Q4). Table 1 shows the test statistics associated with the ADF, Z_t , and Z_α tests together with the static OLS estimates of the cointegrating vector, namely (α, β) .⁵ In the table, q represents the lag length in the case of the ADF test, and the optimal value of q is chosen based on the Akaike information criterion (AIC). In the case of the PP tests, q represents the truncation parameter to compute the long-run variance, and the optimal value of q is chosen based on Andrews' (1991) plug-in method. Superscript a indicates the optimal value of q in the ADF test and superscript b is the optimal value of q in the PP tests.

The estimation results for the semi-log specification show that the null hypothesis of no cointegration is rejected by the ADF and Z_t tests but not by the Z_α test at a significance level of 10 percent. In contrast, the results for the log-log specification show that the null hypothesis is rejected by the ADF test at the 5 percent significance level but not by the PP tests. These results suggest that the semi-log specification is better than the log-log specification but the difference between the two is not that substantial, which is consistent with what we saw in Figure 2.

As pointed out by Ireland (2009), the difference in $\ln(m)$ between the semi-log and the log-log specification is larger when the nominal interest rate comes closer to the zero lower bound. This suggests that it may be difficult to distinguish which of the two specifications provides a better fit unless the observation period contains a sufficiently long period of near-zero interest rates. While the data used by Ireland (2009) contain a period in which interest rates are below 1 percent, the period is very brief, consisting of only three quarters. This may explain why we do not obtain clear-cut results in Table 1. To see if this is the case, we conduct the same exercise but now use an observation period with a much longer period of near-zero interest rates.

Table 2 shows the results with the observation period extended until 2013:Q4. The

⁵Note that the base year for the nominal GDP data we use in this paper is 2005, while the base year for the nominal GDP data used by Ireland (2009) is 2000. Mainly due to this difference, our results differ slightly from those presented in Table 2 in Ireland (2009).

results for the semi-log specification show that the null hypothesis is no longer rejected by any of the three tests, indicating that there does not exist a cointegration relationship between $\ln(m)$ and r . The test statistics are all close to zero in absolute value regardless of the value of q . In contrast, the results for the log-log specification indicate that the test statistics are much larger in absolute value than in the case of the semi-log specification. In fact, the ADF test rejects the null hypothesis at a significance level of 1 percent. Taken together, the results in Table 2 indicate that the log-log specification performs better than the semi-log specification in the sense that it better captures the relationship between the two variables, especially during the period of near-zero interest rates. This again is consistent with what we saw in Figure 2.

Table 2 also presents the static OLS estimates of the cointegrating vector (α, β) . For the log-log specification, the constant is -2.0893 , and the interest rate elasticity is estimated to be 0.0551 . However, static OLS estimators may not be very efficient, so that we estimate cointegrating vector (α, β) based on dynamic OLS rather than static OLS. Specifically, we run a similar OLS regression but now add $\Delta \ln(r_{t-p}), \dots, \Delta \ln(r_{t-1}), \Delta \ln(r_t), \Delta \ln(r_{t+1}), \dots, \Delta \ln(r_{t+p})$ as explanatory variables. Note that, as long as the variables are cointegrated, dynamic OLS estimates are asymptotically efficient and conventional t-statistics obtained from dynamic OLS estimation have conventional normal asymptotic distributions (see Stock and Watson (1993) and Saikkonen (1991)). The estimation results are shown in Table 3, where $\text{std}(\hat{\beta})$ represents the HAC standard error of $\hat{\beta}$ and is obtained for various values of q . The estimates of the cointegrating vector (α, β) are almost the same as in Table 2: the constant α ranges from -2.0897 to -2.0961 , and interest rate elasticity β takes a value in the range from 0.0585 to 0.0571 . The estimated value of the interest rate elasticity is small, but the standard error associated with it is also small, so that the estimate is significantly different from zero. Note that the interest rate elasticity estimated here is much smaller than 0.5 , which is the estimate reported by Lucas (2000).

Finally, using a graph, let us check how well the log-log specification captures the relationship between the two variables. Figure 3 presents a scatter plot using quarterly data. The thick and thin lines respectively represent the fitted values for the log-log and the semi-log specification, which are calculated using the estimates of α and β in Table

1. It can be clearly seen that money demand increases substantially as the interest rate approaches the zero lower bound, but the semi-log specification fails to capture this. In contrast, the log-log specification performs well both in high and in near-zero interest rate periods.

4 Robustness checks

In this section, we check the robustness of the results obtained in the previous section. We check, first, how the results change when we use alternative measures for the opportunity cost of holding money. Second, we allow structural breaks in the cointegration relationships. Third, we relax the constraint that the income elasticity of money demand is unity.

4.1 Alternative measure for the opportunity cost of holding money

The opportunity cost of holding money is measured by the interest rates associated with financial assets that are outside the definition of money but are close to money. In our exercise, we use M1 as money, so in the previous section we used the three-month TB rate as a measure of the opportunity cost of holding money. However, there are many other candidates that would be appropriate as a measure of the opportunity cost of holding money. Also, it should be noted that when short-term interest rates are sufficiently close to the zero lower bound, people may make decisions on how much money to hold by comparing the marginal utility of money and medium- or long-term interest rates, which are still a long from the zero lower bound. If this is the case, we should use medium- or long- term interest rates as a measure of the opportunity cost of holding money. In this subsection, we replace the three-month TB rate with the three-year Treasury constant maturity rate and repeat the exercise.

Figure 4 shows a scatterplot of the annual data for 1980-2013 with the money-income ratio on the horizontal axis and the three-year rate on the vertical axis. The money-income ratio on the horizontal axis is in log in both Figures 4(a) and 4(b), while the three-year rate on the vertical axis is not in log in Figure 4(a) but in log in Figure 4(b). It can be clearly seen that there is no linear relationship between the two variables in Figure 4(a), indicating that the semi-log specification does not work well. On the other hand, there

appears to exist a linear relationship between the two variables in Figure 4(b), suggesting that the log-log specification performs well.

To confirm this visual observation, we conduct cointegration tests using the quarterly data for 1980:Q1-2013:Q4. Table 4 shows that, for the semi-log specification, the three test statistics are all close to zero in absolute value, and the null hypothesis of no cointegration cannot be rejected. On the other hand, for the log-log specification, the ADF and Z_α statistics are much larger in absolute value than in the case of the semi-log specification, and the null hypothesis is rejected. As for the Z_t statistic, this is larger in absolute value than that obtained in the previous section, but the null still cannot be rejected. Overall, we conclude that the results in the previous section are robust to switching the interest rate variable from the three-month TB rate to the three-year Treasury constant maturity rate.

4.2 Allowing for the possibility of structural breaks

The analysis in the last section shows that, as long as we employ the observation period used by Ireland (2009), 1980:Q1-2006:Q4, the money-income ratio and the nominal interest rate are cointegrated when the semi-log specification is used, but this cointegration relationship disappears once we switch to the longer observation period, 1980:Q1-2013:Q4. This raises the question why this happens. One possibility is the presence of structural breaks associated with the cointegrating vector. Specifically, the values of α and β in the recent period with near-zero interest rates may differ from those in the earlier period.

Gregory, Nason, and Watt (1996) and Gregory and Hansen (1996) point out that if one applies a simple cointegration test when some variables are actually cointegrated with structural breaks, then one obtains a result biased toward accepting the null of no cointegration. To investigate this issue in more detail, we employ the tests proposed by Gregory and Hansen (1996), which can detect a cointegrating relationship allowing for a structural break in the intercept and the slope. Specifically, following the procedure proposed by Gregory and Hansen (1996), we (1) consider the middle 70 percent of observations as potential candidates for the time that a structural break may have occurred, which we denote by T_B ; (2) obtain the residual from a regression of $\ln m_t$ on a constant and D_t , r_t (or $\ln r_t$),

and $D_t \times r_t$ (or $D_t \times \ln r_t$), where D_t is 1 for $t > T_B$ and zero otherwise; (3) conduct unit root tests for the residual and obtain the test statistics $ADF(T_B)$, $Z_t(T_B)$, and $Z_\alpha(T_B)$; and (4) search for the minimal values of these test statistics over all possible break points. The test statistics obtained in this way are denoted by Inf-ADF, Inf- Z_t , and Inf- Z_α .

Table 5 shows the values of Inf-ADF, Inf- Z_t , and Inf- Z_α for the semi-log and log-log specifications. Starting with the results for the semi-log specification, the test statistics are much larger in absolute value than in Table 2, suggesting structural breaks may be present. However, we still cannot reject the null hypothesis, implying that we fail to detect a cointegration relationship even when allowing for the possibility of a structural break.⁶ Turning to the log-log specification, we find that the test statistics do not change very much from those in Table 2, implying that the possibility of a structural break is negligible in the case of the log-log specification. The result from the ADF test indicates that the null hypothesis of no cointegration is rejected at the 10 percent significance level.⁷

4.3 No restriction on income elasticity

The cointegration tests we conducted in the last section are based on the assumption that income elasticity is unity. In this subsection, we relax this assumption and examine the cointegration relationship among three variables, i.e., real money balances, real income, and nominal interest rates.

Let M denote retail sweep-adjusted M1, Y the nominal GDP, P the GDP deflator, and r the three-month TB rate. We conduct cointegration tests for M/P , Y/P , and r . Specifically, we regress $\ln(M/P)$ on a constant, $\ln(Y/P)$, and $\ln(r)$ or r , and check whether the residual is stationary or not. Table 6 shows the three test statistics as well as the estimates for α , β_y , and β_r , where α is a constant, β_y is the coefficient on $\ln(Y/P)$, and β_r is the coefficient on $\ln(r)$ or r .

For the semi-log specification, the income elasticity turns out to be 1.0489, which is close to unity. All of the three test statistics indicate that the null hypothesis of no cointegration

⁶The date associated with the minimal value of the test statistics is 2008:Q4, when the three month TB rate fell below 1 percent (from 1.5% in 2008:Q3 to 0.3% in 2008:Q4).

⁷Note that the significance level associated with the ADF test is lower than that presented in Table 2. This may be due to a loss of the statistical power of cointegration tests when applying the Gregory-Hansen test to the case in which there exists a cointegration relationship with no breaks.

is not rejected. For the log-log specification, the estimate of the income elasticity is 1.0666, which is again quite close to unity. The three test statistics are all larger in absolute value than in the case of the semi-log specification, and the test statistics associated with the ADF test indicates that the null hypothesis is rejected at the 5 percent significance level. Overall, the results are essentially the same as those in the last section, indicating that the findings are robust to changes in the specification of the estimation equation.

5 Welfare cost of inflation

In this section, we calculate the welfare cost of inflation using the parameter estimates obtained in Section 3 and compare our results with those reported in previous studies such as Lucas (2000) and Ireland (2009). Table 7 shows the estimation results.⁸ In this calculation, we assume that the real interest rate at the steady state is 3 percent. For example, $r = 0.05$ means that the inflation rate at the steady state is 2 percent, which is the current target level set by the Fed. We start by reproducing the results by Lucas (2000) and Ireland (2009). Lucas (2000) uses annual data for 1900-1994 to obtain $\alpha = -1.036$ and $\beta = 7$ for the semi-log specification and $\alpha = -3.020$ and $\beta = 0.500$ for the log-log specification. The welfare cost of 2 percent inflation is 0.25 percent of national income in the case of the semi-log specification and 1.09 percent in the case of the log-log specification. These results indicate that the welfare cost of inflation is not negligible even if the inflation rate is only 2 percent, especially in the case of the log-log specification.

Turning to the result by Ireland (2009), we use the values for α and β estimated using the quarterly data for the period 1980:Q1-2006:Q4, which are shown in Table 1, to reproduce his result. As shown in the middle two rows of Table 7, the welfare cost of 2 percent inflation is now reduced to 0.03 percent of national income in the case of the semi-log specification and 0.06 percent of national income in the case of the log-log specification.

Finally, the welfare cost calculated based on our estimates for α and β is shown in

⁸Lucas (2000) shows, based on Bailey (1956), that the welfare cost of inflation, $w(r)$, is given by $w(r) = e^\alpha (\beta/(1 - \beta)) r^{1-\beta}$ for the log-log specification and $w(r) = e^\alpha [1 - (1 + \beta r)e^{-\beta r}] / \beta$ for the semi-log specification. Our calculation is based on these results.

the bottom two rows of Table 7. The results indicate that the welfare cost of 2 percent inflation is 0.04 percent in the case of the semi-log specification and 0.04 percent in the case of the log-log specification. Our welfare cost estimates are of almost the same size as those obtained for the shorter observation period used by Ireland (2009), suggesting that, as long as we use the data from 1980 onward, the estimated welfare cost of inflation is very small, irrespective of whether the recent period with near-zero interest rates is included or not.⁹ This is in sharp contrast with the result obtained by Lucas (2000), whose observations on the money-interest rate relationship include the period before 1980.

Figure 5(a) compares our estimates on the welfare cost of inflation with those obtained by Lucas (2000) and Ireland (2009). In the figure, the horizontal axis represents the nominal interest rate, while the vertical axis shows the estimated welfare cost of inflation. Once again, our estimate based on the log-log specification is much smaller than the estimate by Lucas (2000). Next, to compare the results we obtain based on the log-log and the semi-log specification more closely, Figure 5(b) focuses on the results for interest rates below 2 percent. The figure indicates that the estimate based on the semi-log specification is a convex function of r , while the estimate based on the log-log specification is a concave function of r . More interestingly, when r declines from 2 to 1 percent, $w(r)$ falls by 0.008 percentage points in the case of the log-log specification and by 0.006 percentage points in the case of the semi-log specification, so that the changes in $w(r)$ associated with a decline in r from 2 to 1 percent are of almost the same size in the two cases. However, when r declines from 1 to 0 percent, $w(r)$ in the case of the log-log specification falls by 0.009 percentage points, which is even greater than the welfare improvement from 2 to 1 percent, but $w(r)$ falls only by 0.002 percentage points in the case of the semi-log specification. This difference between the log-log and the semi-log specification – that is, that the welfare improvement associated with a decline in interest rates close to zero is larger in the case of the log-log than the semi-log specification – has been highlighted in previous studies such as Lucas (2000) and Wolman (1997). This difference between the two specifications arises because in the log-log specification money demand increases substantially as the interest rate falls

⁹This is consistent with the result obtained by Mogliani and Urga (2018), who used annual data for 1976-2013 to find that the welfare cost of inflation is very small at 0.06-0.14 percent of national income in the case of 2 percent inflation.

from 1 to zero percent, but such an increase in money demand does not occur in the case of the semi-log specification. An important implication of this difference is that it would make more sense for the central bank to reduce inflation and thus the nominal interest rate from 1 to zero percent, which corresponds to the optimal rate of deflation under the Friedman rule, if money demand follows a log-log functional form than a semi-log one.

6 Conclusion

Identifying the proper specification of the money demand function has important implications for the welfare cost of inflation and the conduct of monetary policy. However, in regression analyses of money demand functions, there is no consensus on whether the nominal interest rate as an independent variable should be used in linear or log form. Specifically, Ireland (2009) showed that the semi-log specification performs better than the log-log specification, which stands in sharp contrast with the result obtained by Lucas (2000).

In this paper, we examined the robustness of Ireland's (2009) results by extending the observation period so that it includes the recent period of near-zero interest rates. We showed through simple data plotting and formal cointegration tests that the log-log specification performs better than the semi-log specification. Specifically, we showed that the semi-log specification cannot account for the substantial increase in the money-income ratio during the period of near-zero interest rates since 2008, while the log-log specification can.

Our result on the shape of the money demand function has important implications for the conduct of monetary policy at the zero lower bound. A money demand function that takes a semi-log form implies that the marginal utility of money reaches zero at a finite value of real money balances and becomes negative beyond that level. In this case, the opportunity cost of holding money can go below zero, and in this sense there is no lower bound on nominal interest rates, as shown by Rognlie (2016). However, our result indicates that the marginal utility of money approaches zero as the opportunity cost of holding money falls, but it never reaches zero. Therefore, the opportunity cost of holding money cannot go below zero, which constrains the conduct of monetary policy.

We also computed the welfare cost of inflation based on our estimates of the key parameter values for the money demand function to find that the welfare cost of 2 percent inflation is very small, suggesting that the current inflation target of 2 percent set by the Federal Reserve does not create substantial distortions in the economy.

In this paper, we were able to extend the observation period only up to 2013:Q4 due to data limitations. However, by further extending the observation period, we may be able to learn more about the shape of the money demand function near the zero lower bound. In addition, it would be instructive to conduct similar analyses using data from other countries that experienced near-zero interest rates over a long period of time, such as Japan. We leave these tasks for future research.

References

- [1] Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59(3): 817-858.
- [2] Andrews, Donald W. K., and Christopher J. Monahan. 1992. "An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator." *Econometrica*, 60(4): 953-966.
- [3] Bailey, Martin J. 1956. "The Welfare Cost of Inflationary Finance." *Journal of Political Economy*, 64(2): 93-110.
- [4] Ball, Laurence. 2001. "Another Look at Long-Run Money Demand." *Journal of Monetary Economics*, 47(1): 31-44.
- [5] Baumol, William J. 1952. "The Transactions Demand for Cash: An Inventory Theoretic Approach." *Quarterly Journal of Economics*, 66(4): 545-556.
- [6] Cagan, Phillip. 1956. "The Monetary Dynamics of Hyperinflation." In Milton Friedman, ed., *Studies in the Quantity Theory of Money*, 25-117. Chicago: University of Chicago Press.

- [7] Campbell, John Y., and Pierre Perron. 1991. "Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots." *NBER Macroeconomics Annual* 6: 141-201.
- [8] Cysne, Rubens Penha. 2009. "Bailey's Measure of the Welfare Costs of Inflation as a General-Equilibrium Measure." *Journal of Money, Credit and Banking*, 41(2-3): 451-459.
- [9] DeJong, David N., John C. Nankervis, N. E. Savin, and Charles H. Whiteman. 1992. "The Power Problems of Unit Root Test in Time Series with Autoregressive Errors." *Journal of Econometrics*, 53(1-3): 323-343.
- [10] Dutkowsky, Donald H., and Barry Z. Cynamon. 2003. "Sweep Programs: The Fall of M1 and Rebirth of the Medium of Exchange." *Journal of Money, Credit and Banking*, 35(2): 263-79.
- [11] Eggertsson, Gauti B., Ragnar E. Juelsrud, and Ella Getz Wold. 2017. "Are Negative Nominal Interest Rates Expansionary?" *NBER Working Paper* No. 24039. National Bureau of Economic Research.
- [12] Engle, Robert F., and Clive W. J. Granger. 1987. "Co-Integration and Error Correction: Representation, Estimation, and Testing." *Econometrica*, 55(2): 251-76.
- [13] Friedman, Milton. 1969. "The Optimum Quantity of Money." In Milton Friedman, *The Optimum Quantity of Money and Other Essays*, 1-50. Chicago: Aldine Publishing Company.
- [14] Gregory, Allan W., and Bruce E. Hansen. 1996. "Residual-Based Tests for Cointegration in Models with Regime Shifts." *Journal of Econometrics*, 70(1): 99-126.
- [15] Gregory Allan W., James M. Nason, and David G. Watt. 1996. "Testing for Structural Breaks in Cointegrated Relationships." *Journal of Econometrics*, 71(1-2): 321-341.
- [16] Haug, Alfred A. 1996. "Tests for Cointegration: A Monte Carlo Comparison." *Journal of Econometrics*, 71(12): 89-115.

- [17] Hicks, John R. 1937. "Mr. Keynes and the "Classics": A Suggested Interpretation." *Econometrica*, 5(2): 147-159.
- [18] Hoffman, Dennis L., and Robert H. Rasche. 1991. "Long-Run Income and Interest Elasticities of Money Demand in the United States." *Review of Economics and Statistics*, 73(4): 665-674.
- [19] Ireland, Peter N. 2009. "On the Welfare Cost of Inflation and the Recent Behavior of Money Demand." *American Economic Review*, 99(3): 1040-52.
- [20] Jobst, Clemens, and Helmut Stix. 2017. "Doomed to Disappear? The Surprising Return of Cash Across Time and Across Countries," *CEPR Discussion Papers* No. 12327. Centre for Economic Policy Research.
- [21] Lucas, Robert E., Jr. 1988. "Money Demand in the United States: A Quantitative Review." *Carnegie-Rochester Conference Series on Public Policy*, 29: 137-167.
- [22] Lucas, Robert E., Jr. 2000. "Inflation and Welfare." *Econometrica*, 68(2): 247-74.
- [23] McCallum, Bennett T., and Marvin S. Goodfriend. 1989. "Demand for Money: Theoretical Studies." In John Eatwell, Murray Milgate, and Peter Newman, eds., *Money*, 117-130. London: Palgrave Macmillan.
- [24] Meltzer, Allan H. 1963. "The Demand for Money: The Evidence from the Time Series." *Journal of Political Economy*, 71(3): 219-46.
- [25] Miller, Merton H., and Daniel Orr. 1966. "A Model of the Demand for Money by Firms." *Quarterly Journal of Economics*, 80(3): 413-435.
- [26] Mogliani, Matteo, and Giovanni Urga. 2018. "On the Instability of Long-run Money Demand and the Welfare Cost of Inflation in the United States. " *Journal of Money, Credit and Banking*. doi:10.1111/jmcb.12480 (accessed September 2, 2018).
- [27] Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703-08.

- [28] Phillips, Peter C. B., and S. Ouliaris. 1990. "Asymptotic Properties of Residual Based Tests for Cointegration." *Econometrica*, 58(1): 165-93.
- [29] Phillips, Peter C. B., and Pierre Perron. 1988. "Testing for a Unit Root in Time Series Regression." *Biometrika*, 75(2): 335-346.
- [30] Rognlie, Matthew. 2016. "What Lower Bound? Monetary Policy with Negative Interest Rates." http://mattrognlie.com/negative_rates.pdf (accessed September 2, 2018).
- [31] Saikkonen, Pentti. 1991. "Asymptotically Efficient Estimation of Cointegration Regressions." *Econometric Theory*, 7(1): 1-21.
- [32] Stock, James H., and Mark W. Watson. 1993. "A Simple Estimator of Cointegrating Vectors Order Integrated Systems." *Econometrica*, 61(4): 783-820.
- [33] Tobin, James. 1956. "The Interest-Elasticity of Transactions Demand for Cash." *Review of Economics and Statistics*, 38(3): 241-247.
- [34] Wolman, Alexander L. 1997. "Zero Inflation and the Friedman Rule: A Welfare Comparison." *Federal Reserve Bank of Richmond Economic Quarterly*, 83(4): 1-20.

Table 1: Cointegration tests for 1980:Q1-2006:Q4

Semi-log function	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
$\ln(m) = \alpha - \beta r$	-1.8259	1.6114	0	-3.032	-3.032	-14.551
			1	-2.521	-3.133*	-15.881
			2	-2.305	-3.130* ^b	-15.840 ^b
			3	-3.801** ^a	-3.238*	-17.288*
			4	-3.136*	-3.268*	-17.700*
			5	-4.935***	-3.327*	-18.516*
			6	-4.623***	-3.410**	-19.672*
			7	-3.210*	-3.418**	-19.787*
			8	-3.931**	-3.413**	-19.716*
Log-log function	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
$\ln(m) = \alpha - \beta \ln(r)$	-2.1536	0.0777	0	-1.890	-1.890	-6.871
			1	-2.277	-2.078	-8.363
			2	-2.646	-2.218	-9.566
			3	-3.804** ^a	-2.394 ^b	-11.189 ^b
			4	-3.294*	-2.494	-12.169
			5	-3.728**	-2.574	-12.985
			6	-4.224***	-2.653	-13.805
			7	-3.356*	-2.683	-14.131
			8	-3.192*	-2.686	-14.163

Note: The test statistics Z_t and Z_α are computed using the Newey-West estimate of the long-run variance (Newey and West (1987)). ***, ** and * indicate that the null hypothesis of no cointegration can be rejected at the 1%, 5%, and 10% level, respectively. “a” indicates the lag chosen by the AIC, while “b” indicates the closest integer to the value chosen by the plug-in method of Andrews (1991).

Table 2: Cointegration tests for 1980:Q1-2013:Q4

Semi-log function	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
$\ln(m) = \alpha - \beta r$	-1.7777	2.2763	0	-1.197	-1.197	-4.948
			1	-0.889	-1.421	-6.333
			2	-0.622	-1.433	-6.414
			3	-2.279	-1.554 ^b	-7.234 ^b
			4	-1.253	-1.566	-7.315
			5	-2.510	-1.615	-7.665
			6	-1.955	-1.699	-8.277
			7	-0.602	-1.699	-8.278
			8	-1.977 ^a	-1.705	-8.322
Log-log function	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
$\ln(m) = \alpha - \beta \ln(r)$	-2.0893	0.0551	0	-1.949	-1.949	-9.153
			1	-2.141	-2.058	-10.049
			2	-2.702	-2.240 ^b	-11.645 ^b
			3	-2.710	-2.339	-12.563
			4	-3.095*	-2.44	-13.551
			5	-2.730	-2.475	-13.895
			6	-3.369*	-2.537	-14.523
			7	-4.195*** ^a	-2.605	-15.236
			8	-4.290***	-2.663	-15.850

Note: The test statistics Z_t and Z_α are computed using the Newey-West estimate of the long-run variance (Newey and West (1987)). ***, ** and * indicate that the null hypothesis of no cointegration can be rejected at the 1%, 5%, and 10% level, respectively. “a” indicates the lag chosen by the AIC, while “b” indicates the closest integer to the value chosen by the plug-in method of Andrews (1991).

Table 3: Dynamic OLS

$\hat{\alpha}$	$\hat{\beta}$	std($\hat{\beta}$)	p	q
-2.0959	0.0578	0.0035	1	2
		0.0042		4
		0.0047		6
		0.0051		8
-2.0961	0.0583	0.0034	2	2
		0.0041		4
		0.0046		6
		0.0049		8
-2.0948	0.0585	0.0033	3	2
		0.0041		4
		0.0045		6
		0.0048		8
-2.0897	0.0571	0.0038	4	2
		0.0046		4
		0.0052		6
		0.0056		8

Note: The regression equation is $\ln(m) = \alpha - \beta \ln(r)$. The observation period is 1980:Q1-2013:Q4. The dynamic OLS estimates for α and β are shown for various values of the lead/lag parameter p . The standard errors of $\hat{\beta}$, denoted by $\text{std}(\hat{\beta})$, are computed using the Newey-West estimate of the long-run variance for various values of the truncation parameter q (Newey and West (1987)).

Table 4: Alternative measure for the opportunity cost of holding money

Semi-log function	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
$\ln(m) = \alpha - \beta r$	-1.7555	2.2396	0	-0.657	-0.657	-2.507
			1	-1.093 ^a	-0.992	-4.199
			2	-1.422	-1.160	-5.166
			3	-2.501	-1.354	-6.389
			4	-1.194	-1.394 ^b	-6.657 ^b
			5	-1.652	-1.399	-6.690
			6	-1.798	-1.417	-6.815
			7	-1.268	-1.396	-6.674
			8	-2.095	-1.388	-6.620
Log-log function	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
$\ln(m) = \alpha - \beta \ln(r)$	-2.1995	0.0999	0	-2.089	-2.089	-11.106
			1	-2.979	-2.448	-14.495
			2	-2.766	-2.546	-15.492
			3	-4.137 ^{***a}	-2.748	-17.681*
			4	-3.308*	-2.838 ^b	-18.713 ^b
			5	-3.400**	-2.844	-18.778*
			6	-3.319*	-2.846	-18.801*
			7	-3.491**	-2.836	-18.687*
			8	-3.710**	-2.815	-18.447*

Note: The three-year Treasury constant maturity rate is used as a measure for the opportunity cost of holding money. The observation period is 1980:Q1-2013:Q4. The test statistics Z_t and Z_α are computed using the Newey-West estimate of the long-run variance (Newey and West (1987)). ***, ** and * indicate that the null hypothesis of no cointegration can be rejected at the 1%, 5%, and 10% level, respectively. “a” indicates the lag chosen by the AIC, while “b” indicates the closest integer to the value chosen by the plug-in method of Andrews (1991).

Table 5: Structural breaks

	Inf-ADF	Inf- Z_t	Inf- Z_α
Semi-log specification	-4.521	-3.295	-21.222
Log-log specification	-4.875*	-2.678	-16.109

Note: The three test statistics (Inf-ADF, Inf- Z_t , and Inf- Z_α) are computed using the procedure proposed by Gregory and Hansen (2006). ***, ** and * indicate that the null hypothesis of no cointegration is rejected at the 1%, 5%, and 10% level, respectively. For the ADF test, the lag length is chosen based on a sequential t-test with the maximum lag length set to 8. For the PP tests, the long-run variance is estimated using a prewhitened quadratic spectral kernel based on Andrews (1991) and Andrews and Monahan (1992).

Table 6: No restriction on income elasticity

	$\hat{\alpha}$	$\hat{\beta}_y$	$\hat{\beta}_r$	q	ADF	Z_t	Z_α
Semi-log specification							
$\ln(M/P) = \alpha + \beta_y \ln(Y/P) - \beta_r r$	-2.0221	1.0489	1.9354	0	-0.919	-0.919	-3.479
				1	-0.740	-1.165	-4.792
				2	-0.684	-1.224	-5.142
				3	-2.287 ^a	-1.372 ^b	-6.049 ^b
				4	-1.263	-1.409	-6.296
				5	-2.399	-1.471	-6.699
				6	-1.877	-1.556	-7.279
				7	-0.781	-1.569	-7.378
				8	-2.112	-1.585	-7.489
Log-log specification							
$\ln(M/P) = \alpha + \beta_y \ln(Y/P) - \beta_r \ln(r)$	-2.3691	1.0666	0.0466	0	-1.785	-1.785	-8.174
				1	-2.038	-1.949	-9.458
				2	-2.839	-2.182 ^b	-11.439 ^b
				3	-2.980	-2.324	-12.759
				4	-3.219	-2.446	-13.940
				5	-2.964	-2.502	-14.514
				6	-3.623*	-2.576	-15.275
				7	-4.254 ^{**a}	-2.641	-15.964
				8	-4.259 ^{**}	-2.684	-16.456

Note: The observation period is 1980:Q1-2013:Q4. The test statistics Z_t and Z_α are computed using the Newey-West estimate of the long-run variance (Newey and West (1987)). ***, ** and * indicate that the null hypothesis of no cointegration can be rejected at the 1%, 5%, and 10% level, respectively. “a” indicates the lag chosen by the AIC, while “b” indicates the closest integer to the value chosen by the plug-in method of Andrews (1991).

Table 7: Welfare cost of inflation

Observation period	Functional form	Parameter values		Welfare cost		
		α	β	$r = 0.13$	$r = 0.05$	$r = 0.03$
1900-1994	semi-log	-1.036	7.000	1.17%	0.25%	0.10%
	log-log	-3.020	0.500	1.76%	1.09%	0.85%
1980:Q1-2006:Q4	semi-log	-1.826	1.611	0.19%	0.03%	0.01%
	log-log	-2.154	0.078	0.15%	0.06%	0.04%
1980:Q1-2013:Q4	semi-log	-1.778	2.276	0.27%	0.04%	0.02%
	log-log	-2.089	0.055	0.10%	0.04%	0.03%

Note: It is assumed that the real interest rate at the steady state is 3 percent, so that $r = 0.03$ corresponds to price stability, $r = 0.05$ to 2 percent inflation, and $r = 0.13$ to 10 percent inflation. The welfare costs are computed using equations (3) and (4) in Ireland (2009) for the log-log and the semi-log specification, respectively. The parameter values of α and β are taken from Lucas (2000) for the period 1900-1994, from Table 1 for the period 1980:Q1-2006:Q4, and from Table 2 for the period 1980:Q1-2013:Q4.

Figure 1: Money Demand, 1980-2013





Figure 2: Semi-log vs. Log-log Plots

(a) semi-log plot



(b) log-log plot




Figure 3: Estimated Money Demand Functions





Figure 4: 3-Year Treasury Constant Maturity Rate

(a) semi-log plot



(b) log-log plot


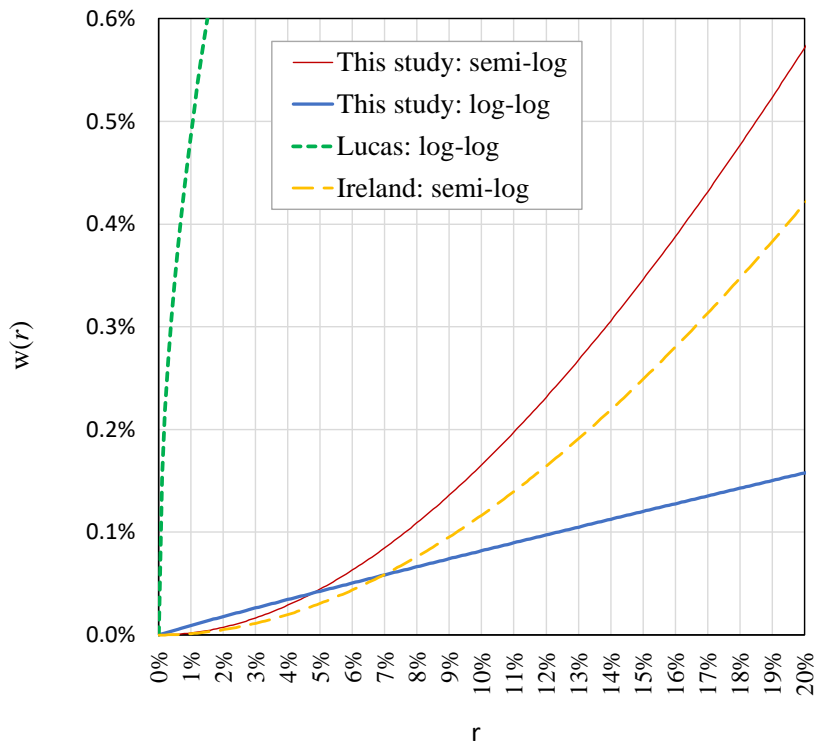


Figure 5: Welfare Cost of Inflation

(a)



(b)

