JSPS Grant-in-Aid for Scientific Research (S)

Central Bank Communication Design

working paper series

No.44 (December 2022)

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‡

December 2, 2022

Abstract

This paper undertakes both a narrow and wide replication of the estimation of a money demand function conducted by Ireland (*American Economic Review*, 2009). Using US data from 1980 to 2013, we show that the substantial increase in the money-income ratio during the period of near-zero interest rates is captured well by the log-log specification but not by the semi-log specification, contrary to the result obtained by Ireland (2009). Our estimate of the interest elasticity of money demand over the 1980-2013 period is about one-tenth that of Lucas (2000), who used a log-log specification. Finally, neither specification satisfactorily fits post-2015 US data.

JEL Classification Numbers: C22; C52; E31; E41; E43; E52
Keywords: money demand function; cointegration; zero lower bound; welfare cost of inflation; log-log form; semi-log form

*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).

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 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. We show that inclusion of the more recent observations with near-zero interest rates changes the results substantially. The money stock measure employed by Ireland (2009) and this paper, "M1 adjusted for retail sweeps," is available only until 2013:Q4. In the online appendix, we extend the sample period to 2022:Q2 by connecting the series to the original M1 series. We show that neither specification satisfactorily fits post-2015 US data.

The rest of the paper is organized as follows. In Section 2, we visually compare the log-log

Figure 1: Semi-log vs. Log-log Plots

(a) Semi-log plot (b) Log-log plot

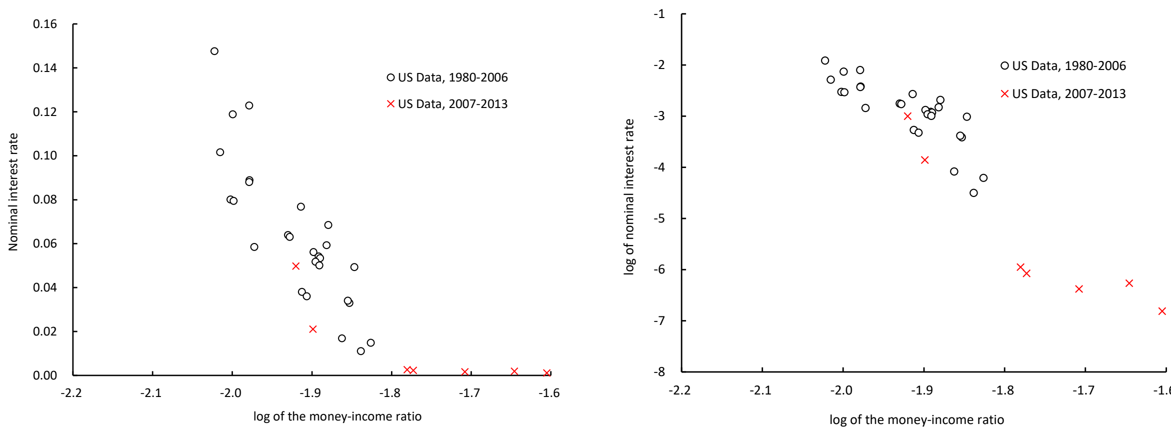

and semi-log specifications using annual data. In Section 3, we conduct cointegration tests, using quarterly data, for the money-income ratio and the nominal interest rate to examine which of the two specifications performs better. In Section 4, we discuss the welfare cost of inflation based on the estimation result of the money demand function. Section 5 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 up to 2013, the latest year for which retail sweep-adjusted M1 is available.

We plot the data into the semi-log graph in Figure 1(a) and into the log-log graph in Figure 1(b). Figure 1(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 1(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 1(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 1(a). In contrast, in Figure 1(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. 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. 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 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).[1] 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 are identical to the ones reported by Ireland (2009), suggesting that the semi-log specification is better than the log-log specification. However, the difference between the two is not that substantial, which is consistent with what we saw in Figure 1.

[1] Note that the base year for the nominal GDP data we use in this paper is 2009, 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).

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	-2.509	-3.133*	-15.881
			2	-2.294	-3.130*^{b}	-15.840^b
			3	-3.783**^{a}	-3.238*	-17.288*
			4	-3.121*	-3.268*	-17.700*
			5	-4.911***	-3.327*	-18.516*
			6	-4.600***	-3.410**	-19.672*
			7	-3.194*	-3.418**	-19.787*
			8	-3.911**	-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	-2.266	-2.078	-8.363
			2	-2.633	-2.218	-9.566
			3	-3.786**^{a}	-2.394^b	-11.189^b
			4	-3.278*	-2.494	-12.169
			5	-3.710**	-2.574	-12.985
			6	-4.203***	-2.653	-13.805
			7	-3.339*	-2.683	-14.131
			8	-3.175*	-2.686	-14.163

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

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 results for

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

	$\hat{\alpha}$	$\hat{\beta}$	q	ADF	Z_t	Z_α
semi-log function $\ln(m) = \alpha - \beta r$	-1.7777	2.2763	0	-1.197	-1.197	-4.948
			1	-0.886	-1.421	-6.333
			2	-0.620	-1.433	-6.414
			3	-2.271	-1.554[b]	-7.234[b]
			4	-1.248	-1.566	-7.315
			5	-2.500	-1.615	-7.665
			6	-1.947	-1.699	-8.277
			7	-0.600	-1.699	-8.278
			8	-1.969[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	-2.133	-2.058	-10.049
			2	-2.692	-2.240[b]	-11.645[b]
			3	-2.700	-2.339	-12.563
			4	-3.083*	-2.44	-13.551
			5	-2.719	-2.475	-13.895
			6	-3.356*	-2.537	-14.523
			7	-4.178***[a]	-2.605	-15.236
			8	-4.273***	-2.663	-15.850

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

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 much smaller in absolute value than in Table 1 and 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. First, the ADF test rejects the null hypothesis at a significance level of 1 percent. Second, the PP tests do not reject the null but the test statistics, Z_t and Z_α, are almost the same as in Table 1 and much larger in absolute value than the corresponding statistics for the semi-log specification.[2]

[2]Regarding the ADF and PP tests, previous studies have pointed out that (1) the PP tests have larger size distortions than the ADF test in the presence of negative moving average errors (see, for example, Phillips and Perron 1988), and (2) the ADF test has better size and power when the errors have an AR structure (see, for

Figure 2: Estimated Money Demand Functions

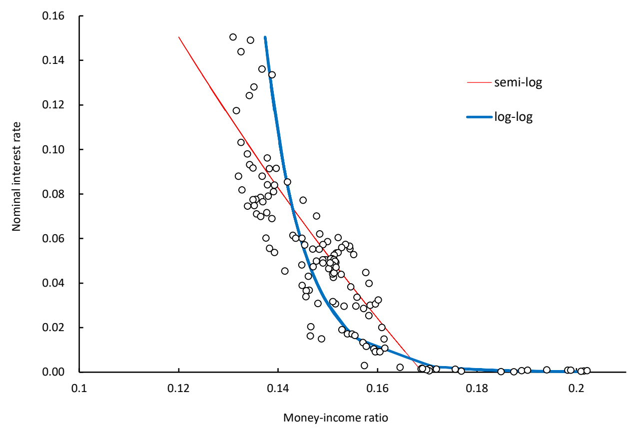

Taken together, the results in Table 2 indicate that the log-log specification performs better than the semi-log specification in the sense that it better captures the relationship between the two variables, especially during the period of near-zero interest rates. This again is consistent with what we saw in Figure 1. Table 2 also presents the static OLS estimates of the cointegrating vector $(\alpha,\ \beta)$. For the log-log specification, the constant is -2.0893, and the interest rate elasticity is estimated to be 0.0551.[3]

Finally, using a graph, let us check how well the log-log specification captures the relationship between the two variables. Figure 2 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 2. It can be clearly seen that money demand increases substantially as the interest rate approaches the zero lower bound, but the

example, DeJong et al. 1992). The correlogram of the log-log model residuals shows that the autocorrelation function gradually decays, while the partial autocorrelation function is close to 0 for all but the first order. This indicates that the residuals in the log-log model have an AR structure, and that this may be the reason why the ADF test has better size and power than the PP test in our analysis.

[3]The online appendix to this paper examines the possibility of structural breaks. Specifically, we conduct the test proposed by Gregory and Hansen (1996) to detect a cointegrating relationship even when there is a structural break. We show that the null of no cointegration cannot be rejected for the semi-log specification, which means that our test fails to detect a cointegration relationship even when allowing for the possibility of a structural break. On the other hand, the null of no cointegration is rejected for the log-log specification. We also conduct the test proposed by Kejriwal and Perron (2010) to examine whether there actually are structural breaks. We show that the null of no breaks cannot be rejected for the log-log specification but is rejected for the semi-log specification. For the semi-log specification, we detect a structural break in 2007:Q4.

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]

4 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 3 shows the estimation results. 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 associated with $r = 0.05$ 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.[5] These results indicate that the welfare cost of inflation is not negligible even if the interest rate is only 5 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 3, the welfare cost associated with $r = 0.05$ 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 the bottom two rows of Table 3. The results indicate that the welfare cost associated with $r = 0.05$ 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 is very small, irrespective of whether the recent

[4]The sample period of this paper ends in 2013:Q4 as "M1 adjusted for retail sweeps" is available only until 2013:Q4. In the online appendix, we extend the sample period to 2022:Q2 by connecting the series to the original M1 series. We show that, for the extended period, neither the semi-log form nor the log-log form fits the data well. Specifically, the Fed began raising the federal funds target rate in December 2015 and subsequently raised it nine more times through December 2018. Nonetheless, the money-income ratio *rose* during this period, albeit slightly, rather than falling. Neither the semi-log nor the log-log specification can account for this positive correlation. We also show that, for 2021:Q2 and onward, the money-income ratio has hardly decreased despite the upward trend in interest rates since the second quarter of 2021.

[5]Note that $r = 0.05$ means that the inflation rate is 2 percent if the equilibrium real interest rate is 3 percent.

Table 3: Welfare cost of inflation

Sample period	Functional form	Parameter values α	β	Welfare cost $r = 0.1$	$r = 0.05$	$r = 0.02$	$r = 0.01$
1900-1994	semi-log	-1.036	7.000	0.79%	0.25%	0.05%	0.01%
	log-log	-3.020	0.500	1.54%	1.09%	0.69%	0.49%
1980:Q1-2006:Q4	semi-log	-1.826	1.611	0.12%	0.03%	0.01%	0.00%
	log-log	-2.154	0.078	0.12%	0.06%	0.03%	0.01%
1980:Q1-2013:Q4	semi-log	-1.778	2.276	0.17%	0.04%	0.01%	0.00%
	log-log	-2.089	0.055	0.08%	0.04%	0.02%	0.01%

Note: The welfare costs are computed using the formulas given in Lucas (2000); $w(r) = \exp(\alpha)[\beta/(1-\beta)]r^{1-\beta}$ for the log-log specification and $w(r) = \exp(\alpha)[1-(1+\beta r)\exp(-\beta r)]/\beta$ for the semi-log specification. The parameter values of α and β are taken from Lucas (2000) for the period of 1900-1994, from Table 1 for the period of 1980:Q1-2006:Q4, and from Table 2 for the period of 1980:Q1-2013:Q4.

period with near-zero interest rates is included or not.[6] 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 3(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 3(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

[6] This 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.

Figure 3: Welfare Cost of Inflation

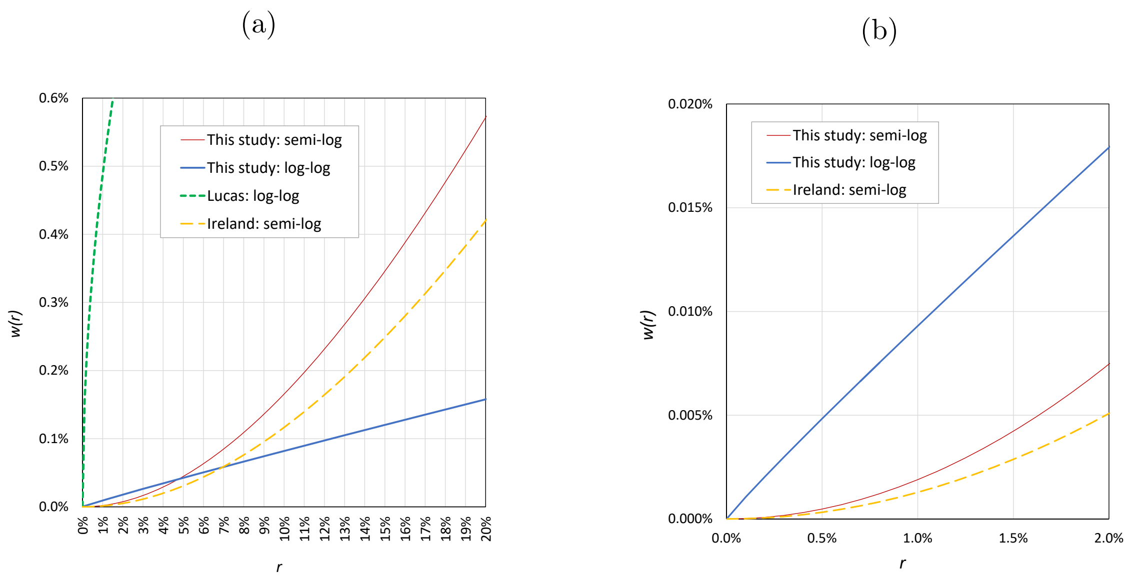

specification falls by 0.009 percentage points, which is even greater than the welfare improvement from 2 to 1 percent, but $w(r)$ falls only by 0.002 percentage points in the case of the semi-log specification.

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 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.

5 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.[7]

We also computed the welfare cost of inflation based on our estimates of the key parameter

[7] The log-log specification we employ in this paper, i.e., $\ln(m) = \alpha - \beta \ln(r)$, is based on the assumption that the cost of storing money is negligible, as in Hicks (1937). However, in the context of a negative interest rate policy, Eggertsson et al. (2019) consider the case in which the storage cost of holding money is non-negligible and the marginal cost associated with storing money is positive. In this case, the log-log specification changes to $\ln(m) = \alpha - \beta \ln(r + \theta)$, where θ is a positive parameter representing the marginal cost of storing money (see the online appendix for the derivation of this equation). This specification implies that the demand for money remains finite even at a negative interest rate as long as $|r|$ does not exceed θ. The fact that the demand for money remained finite even in countries with negative interest rates, including Japan, Sweden, and Switzerland, seems to suggest that $|r|$ is still smaller than θ. A task for the future is to estimate θ by running a log-log form regression relaxing the assumption that the storage cost is negligible.

values for the money demand function to find that the welfare cost of inflation is very small as long as the nominal interest rate is below 5 percent, suggesting that the current inflation target of 2 percent set by the Federal Reserve does not create substantial distortions in the economy.

The money stock measure employed by Ireland (2009) and this paper, "M1 adjusted for retail sweeps," is available only until 2013:Q4. In the online appendix, we extend the sample period to 2022:Q2 by connecting the series to the original M1 series to show that neither specification satisfactorily fits post-2015 US data. The Fed began raising the federal funds target rate in December 2015 and subsequently raised it nine more times through December 2018. Nonetheless, the money-income ratio *rose* during this period, albeit slightly, rather than falling. This is more or less similar to what Watanabe and Yabu (2019) found in Japanese data, in which the demand for money did *not* decline in 2006 when the Bank of Japan terminated quantitative easing and started to raise the policy rate. 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.

References

[1] Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59(3): 817-858.

[2] Ball, Laurence. 2001. "Another Look at Long-Run Money Demand." *Journal of Monetary Economics*, 47(1): 31-44.

[3] 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.

[4] 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.

[5] 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.

[6] Eggertsson, Gauti B., Ragnar E. Juelsrud, Laurance H. Summers, and Ella Getz Wold. 2019. "Negative Nominal Interest Rates and the Bank Lending Channel." *NBER Working Paper* No. 25416. National Bureau of Economic Research.

[7] Engle, Robert F., and Clive W. J. Granger. 1987. "Co-Integration and Error Correction: Representation, Estimation, and Testing." *Econometrica*, 55(2): 251-76.

[8] 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.

[9] Haug, Alfred A. 1996. "Tests for Cointegration: A Monte Carlo Comparison." *Journal of Econometrics*, 71(12): 89-115.

[10] Hicks, John R. 1937. "Mr. Keynes and the "Classics": A Suggested Interpretation." *Econometrica*, 5(2): 147-159.

[11] 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.

[12] Ireland, Peter N. 2009. "On the Welfare Cost of Inflation and the Recent Behavior of Money Demand." *American Economic Review*, 99(3): 1040-52.

[13] Kejriwal, Mohitosh, and Pierre Perron. 2010. "Testing for Multiple Structural Changes in Cointegrated Regression Models." *Journal of Business and Economic Statistics*, 28(4): 503-22.

[14] Lucas, Robert E., Jr. 1988. "Money Demand in the United States: A Quantitative Review." *Carnegie-Rochester Conference Series on Public Policy*, 29: 137-167.

[15] Lucas, Robert E., Jr. 2000. "Inflation and Welfare." *Econometrica*, 68(2): 247-74.

[16] Meltzer, Allan H. 1963. "The Demand for Money: The Evidence from the Time Series." *Journal of Political Economy*, 71(3): 219-46.

[17] 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*, 50(7): 1645-1660.

[18] Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703-08.

[19] Phillips, Peter C. B., and S. Ouliaris. 1990. "Asymptotic Properties of Residual Based Tests for Cointegration." *Econometrica*, 58(1): 165-93.

[20] Phillips, Peter C. B., and Pierre Perron. 1988. "Testing for a Unit Root in Time Series Regression." *Biometrika*, 75(2): 335-346.

[21] Rognlie, Matthew. 2016. "What Lower Bound? Monetary Policy with Negative Interest Rates." `http://mattrognlie.com/negative_rates.pdf` (accessed September 2, 2018).

[22] Stock, James H., and Mark W. Watson. 1993. "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems." *Econometrica*, 61(4): 783-820.

[23] Watanabe, Tsutomu, and Tomoyoshi Yabu. 2019. "How Large is the Demand for Money at the ZLB? Evidence from Japan." CARF Working Paper Series, CARF-F-465, September 2019.

[24] Wolman, Alexander L. 1997. "Zero Inflation and the Friedman Rule: A Welfare Comparison." *Federal Reserve Bank of Richmond Economic Quarterly*, 83(4): 1-20.